A remote operations center framework for heterogeneous robotic zooplankton sampling: System design and field case study

Karoline Barstein[1*], Rabea Rogge[1], André Olaisen[2], Ahmed Abdelgayed[1], Muriel Dunn[3], Ralph Stevenson-Jones[3], Emlyn Davies[3], Leif Grimsmo[3], Martin Ludvigsen[1]

1 Department of Marine Technology, Norwegian University of Science and Technology, Trondheim, Norway
2 Department of Mathematical Sciences, Norwegian University of Science and Technology, Trondheim, Norway
3 SINTEF Ocean, Trondheim, Norway

* karoline.barstein@ntnu.no

This manuscript is a non-peer-reviewed preprint. A version has been submitted for journal peer review.

Abstract

Autonomous and heterogeneous marine robotic networks offer new possibilities for observing dynamic ocean phenomena beyond the limits of ship-centered and single-platform operations. Their scientific and operational value depends on coordinated system use: remote supervision, communication across constrained links, and onboard or near-real-time processing that converts raw sensor streams into compact information for live decision-making. This paper presents an evidence-based system framework for heterogeneous robotic operations in marine research coordinated through a Remote Operations Center (ROC). Zooplankton sampling is used as the motivating case because its patchy, dynamic biological structure is difficult to characterize efficiently with conventional ship-based approaches alone. The framework covers remote operation, edge processing and data quality, communication design, and targeted sampling with heterogeneous fleets. It is tested through a field case study conducted in Mausund, Norway, in 2024, where an autonomous underwater vehicle, two autonomous surface vehicles, and a support vessel were coordinated from a shore-based ROC approximately 100 km away. The results show that the ROC redistributed supervision, interpretation, and coordination from the field site to a shared shore-based environment, allowing these activities to continue in parallel with local launch, recovery, and support tasks. Compact edge-derived products made payload data more usable during operations, while layered communication design helped preserve command, coordination, and science-relevant information across links with varying capacity, latency, and reliability. The case study also showed how complementary platform roles supported broad-scale reconnaissance, local refinement, and targeted sampling. The findings indicate that the value of ROC-centered heterogeneous marine robotics lies in the combined observational capability of the system, linking platform roles, sensing, communication, coordination, and mission logic into one operational workflow.

Introduction

Marine research is increasingly dependent on observation strategies that can respond to dynamic, spatially distributed, and short-lived ocean phenomena. The scientific features of the ocean are often shaped by spatial patchiness, vertical structure, and temporal variability that are difficult to efficiently characterize with ship-centered sampling strategies or single robotic platforms alone [1–4]. This motivates multi-asset operational concepts that combine distributed sensing, adaptive coordination, and persistent situational awareness [5–7].

Heterogeneous marine robotic fleets offer one response to the multi-scale observational challenge. Autonomous underwater vehicles (AUVs), gliders, and autonomous or uncrewed surface vehicles (ASVs/USVs) can provide complementary sensing, coverage, and endurance, and can be assigned different roles within a shared mission [8–11]. Compared to ship-centered operations, robotic platforms can also reduce dependence on vessel time and improve operational efficiency with lower cost and carbon burden [10]. However, their scientific and operational value requires supervision, coordination, retasking, and interpretation across field and shore, especially when operations extend beyond line-of-sight or rely on bandwidth-limited links [12, 13].

In over-the-horizon (OTH) operations, communication systems and compact, context-relevant information transmission become central to the operational concept. Growing availability of 4G/5G backhauls and low-Earth-orbit satellite services has made sustained remote marine operations more practical, but high-volume payload streams are still rarely transferable [13]. Remote coordination depends on processed information products that retain the scientific content. The Remote Operations Center (ROC) serves as the point where telemetry, payload summaries, communication status, and operational context are combined to support coordination across a heterogeneous network. Recent work on autonomous marine robotic organizations has emphasized interoperability, distributed roles, and the need to design sensing, communications, and coordination as parts of a shared operational structure rather than as isolated subsystems [12, 14].

Zooplankton sampling provides a relevant use case for ROC-centered heterogeneous operations. Zooplankton aggregations are patchy and dynamic, often varying across spatial scales from meters to kilometers and temporal scales from hours to days [15, 16]. In the North Atlantic, Norwegian, and Barents seas, *Calanus finmarchicus* is a keystone species central to ecosystem functioning and fish production [17, 18], while its abundance and distribution are sensitive to environmental variability and climate-driven change [19–21]. Conventional methods such as net tows and ship-mounted acoustics remain essential, but they are constrained by vessel time, labor demands, and sparse coverage [22, 23]. Distributed robotic systems are relevant for improved detection and mapping of dynamic biomass distributions. This can support the broader goal of sustainable assessment and future utilization of low-trophic marine resources.

Early work on autonomous ocean observing proposed distributed sampling networks in which multiple autonomous platforms function as parts of a larger observing system [24]. Later studies demonstrated adaptive robotic sampling in practice, including multi-AUV adaptive sampling [25] and onboard model-based tracking of patchy subsurface features in the coastal ocean [26]. Other work has shown coordinated heterogeneous robotic exploration in dynamic marine environments [27, 28]. However, fewer studies present a field-grounded operational framework that combines sensing, edge-derived information products, communication design, remote supervision, and retasking within one coordinated system.

This paper addresses that gap by presenting an evidence-based system framework for ROC-centered heterogeneous marine robotic operations, motivated by the challenge of observing and sampling dynamic ocean phenomena and evaluated through a field case study on zooplankton sampling. The paper examines remote operation, edge processing, communication, and targeted sampling as connected parts of one operational design for heterogeneous robotic systems. The framework is organized around four connected components: remote operation, edge processing and data quality, communication design for over-the-horizon coordination, and targeted sampling with heterogeneous fleets.

The paper first presents a generic framework for ROC-centered heterogeneous operations and then evaluates its implementation through a field campaign for zooplankton sampling conducted in Mausund, Norway, in 2024.

System design and methods

The research method is design- and operations-driven: it starts from practical constraints and scientific requirements for over-the-horizon zooplankton surveys and treats the ROC, fleet, communication system, and sampling target as one coupled operational framework. Fig 1 summarizes this framework. The ROC and heterogeneous fleet form the operational backbone, while the target phenomenon defines the external environmental feature to be observed and sampled. Remote operation, communication design, edge processing and data quality, and targeted sampling describe the main functions that make the system operationally useful.

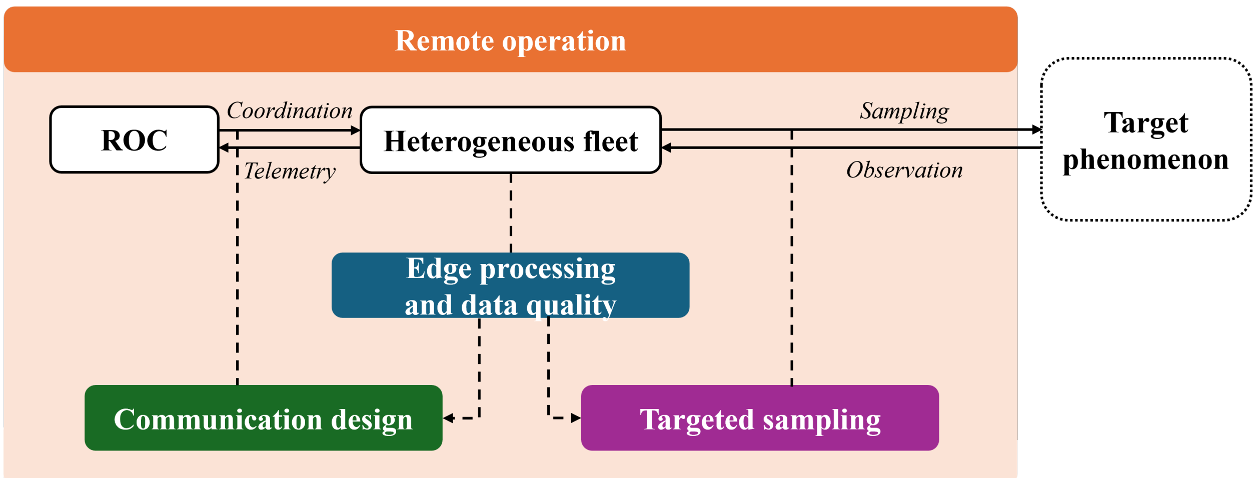

Fig 1. ROC-centered operational framework for heterogeneous marine robotic operations. Remote operation frames the ROC-centered coordination of the heterogeneous fleet, while the target phenomenon represents the external environmental feature to be observed and sampled. Communication design supports coordination and telemetry exchange between the ROC and fleet, edge processing and data quality convert observations into compact and trustworthy information products, and targeted sampling links fleet actions to the target phenomenon.

Remote operation concept and setup

Operational architecture and platform roles

The ROC design builds on systems-engineering work for the OceanLab subsea facility and on a later remote operation framework for networked marine robotic systems [13, 29]. In this architecture, the ROC acts as the shore-side coordination node linking vehicles, field personnel, and remote scientific users.

The robotic system is structured as a heterogeneous fleet for subsurface sampling, surface observation, and operational support. An AUV provides depth-resolved in situ sampling, while USVs support wider-area observation, communication, and mission coordination. The fleet combines acoustic, optical, and environmental sensing. Broadband echosounders are used for water-column backscatter observations, while optical sensors provide particle- or organism-level observations for interpretation and validation. Field deployment still requires a minimum level of crewed operational assistance for launch and recovery, maintenance, and recharging.

Remote operation workflow, roles, and operator stations

Remote operation is organized around three main phases: pre-mission planning, live execution, and post-mission review. In the planning phase, the coordinator and operations manager prepare a shared operational schedule specifying time windows, asset allocations, and key decision points. This schedule is used to align ROC staff, field personnel, and remote scientific participants.

During the execution, the *coordinator* maintains the operational overview and updates the shared schedule. The *operations manager* oversees mission priorities and higher-level decisions based on the schedule, available resources, field conditions, and operational constraints. *Asset operators* focus on specific vehicles or instruments, executing plans, monitoring vehicle and payload state, and proposing adjustments when conditions deviate from expectations.

Shared ROC displays support situational awareness across the operation. They provide a common operational picture by combining schedule information, event and log history, traffic and vehicle positions, map-based context, environmental conditions, vehicle-state information, and selected field video feeds.

Operator workstations use a common software stack, allowing operators to transition between vehicle and payload roles. While the shared displays provide the common operational picture, the workstations support task-specific interaction through mission-planning and control interfaces, detailed telemetry and payload views, communication tools, and the shared operational log.

Logging and reporting are integrated in an online workflow. Operator actions, plan uploads, parameter changes, vehicle events, and schedule updates are recorded on a shared timeline across the ROC and field team. This common record supports real-time coordination, later scientific interpretation of the operation, and iterative refinement of procedures and ROC configuration in subsequent campaigns.

Operational scenarios and cost–carbon framework

To quantify the operational effect of remote operation, a reusable comparison framework was defined for two application contexts: (i) *research operations* and (ii) *commercial surveys*. For each context, two scenarios are evaluated: (I) *ship-based operations* and (II) *ROC-supervised robotic operations*.

For each context and scenario, the framework reports two metrics: (a) financial cost, expressed in EUR, and (b) a fuel-based CO_2e proxy for operational emissions. Both can be defined across campaign types, while the exact line items, system boundaries, and reporting horizon remain case-dependent.

The financial cost is defined as the sum of all included line items converted to the chosen reporting horizon:

$$C = \sum_i r_i \, m_i, \tag{1}$$

where r_i is the rate assigned to line item i and m_i is the corresponding activity multiplier, such as duty hours, number of days, campaign count, or distance covered. Line items are specified as per-hour, per-day, per-campaign, or one-off values.

In the commercial context, results are additionally normalized as cost per trackline-kilometer, calculated as total cost divided by effective on-task trackline distance. Campaign totals are used for research operations, where mobilization and other fixed one-off costs are material. EUR/km is used in the commercial context to reflect delivered coverage and operating efficiency, excluding investment costs and other fixed one-off costs that depend strongly on ownership and utilization.

Operational emissions are represented using a simplified CO_2e proxy:

$$E = \sum_j q_j \, f_j, \tag{2}$$

where q_j is the quantity of relevant energy carrier or activity term j within the defined system boundary, and f_j is the corresponding CO_2e emission factor. This proxy supports structured scenario comparison but does not represent a full life-cycle assessment.

For ships, the proxy is limited to main-engine propulsion fuel. Hotel loads and scientific gear are excluded unless they can be represented consistently across scenarios without double counting. For ROC-based operations, building overhead is excluded, while mission-specific electric loads are included when metered or reasonably approximated. Fuel use is converted using emission factors from the Norwegian Environment Agency, and electricity use is converted using Norwegian Water Resources and Energy Directorate (NVE) electricity disclosure factors [30, 31].

The comparison is bounded by the defined scope and reporting horizon and should not be interpreted as a full life-cycle cost analysis unless such terms are explicitly included. Table 1 summarizes the two contexts and their assumptions, while S1 Table provides the detailed input schema used to structure the cost and CO_2e comparison.

Table 1. Scenario matrix for the cost–carbon comparison. Contexts and assumptions used for the research and commercial scenarios.

Context	Scenario	Short description	Key assumptions
Research	**I. Ship-based**	Coastal and/or offshore research vessel is the primary platform for daily operations.	Scientists and field engineers embarked; vessel crewed and operated per standard practice; duty hours as planned.
	II. ROC + robotic	ROC-supervised robotic fleet for research; zero ship days.	Scientists fully remote; field engineers on site as needed; ROC access available; field logistics as required.
Commercial	**I. Ship-based**	Traditional trawler-based surveys for zooplankton assessment.	Vessel crewed and operated per commercial practice; duty hours as planned.
	II. ROC + robotic	ROC supervises a robotic fleet for survey/search; zero ship days.	Vehicles operated under applicable usage rates; operator/technician coverage per concept of operations.

Edge processing, data quality, and detection capabilities

Edge processing refers to processing performed on or near the sensing platform, in near real time, to convert raw payload data into compact information that can support coordination, interpretation, and retasking during the mission.

The purpose of these outputs is to represent biologically relevant features in a compact and interpretable form while making data quality and confidence visible during the mission. In heterogeneous operations, this supports consistent interpretation and follow-up across platforms. Edge processing therefore links sensing to operational action by converting raw observations into compact information products for coordinated use.

Sensor configuration and calibration

For edge processing to support coordinated operations, sensor configuration must be selected with respect to the intended information products and documented together with calibration, timing, and relevant platform metadata. Acoustic configurations must preserve sensitivity to the target scattering structure over the required depth range, while optical configurations must maintain usable image quality and particle statistics under platform-specific flow and motion conditions.

Consistent interpretation of edge-derived outputs across the system requires configuration baselines, preserved calibration and timing metadata, and documentation of known platform effects. During system development and validation, biological reference samples are needed to establish how the acoustic and optical summaries relate to the target organisms.

Edge-derived information products

Edge-derived information products should capture mission-relevant features in a compact form that supports interpretation and follow-up during operations. For acoustic payloads, this means summaries that describe the presence, vertical extent, and confidence of scattering structure. For optical payloads, it means summaries that describe the observed particle field through time-resolved counts within selected size bands. Quality indicators should be reported with these estimates to distinguish weak or absent target features from degraded acquisition conditions.

For optical sensing, image frames are background-corrected and segmented to isolate candidate particles, after which regions of interest and per-particle summary statistics are retained for later interpretation, aggregation, and quality assessment.

The edge-derived outputs are reviewed together with vehicle position, depth, time, and environmental context. When platforms sample the same area or depth range within comparable time windows, acoustic and optical summaries can be checked against each other. Operators can then decide whether a feature should be accepted, re-examined, or followed up through additional or adapted sampling.

Communication design

The communication design distinguishes three main content categories: H = human communication between field personnel, ROC operators, and remote scientists; C2 = command-and-control and housekeeping telemetry for vehicles and instruments; and P = payload data. Payload data are further divided into compact edge products (E) and raw or near-raw data (R). These categories differ in bandwidth, latency, and operational criticality. Human communication (H) requires enough throughput for voice, video, and shared operational tools, but can tolerate moderate latency. Command-and-control traffic (C2) is low-volume, latency-sensitive, and safety-critical. Payload traffic is separated so that compact edge products (E) can support in-mission decisions over constrained links, while raw data and bulk logs (R) are transferred when higher-capacity links are available.

Layered network design

The communication architecture is organized as a layered design spanning the ROC, field gateways, support platforms, and autonomous vehicles. As shown in Fig 2, the design can be understood in four functional layers: a physical and link layer, a network and overlay layer, a messaging and middleware layer, and an operational application layer. Together, these layers provide a common communication structure across heterogeneous assets while allowing traffic classes to be handled according to their operational role and link constraints.

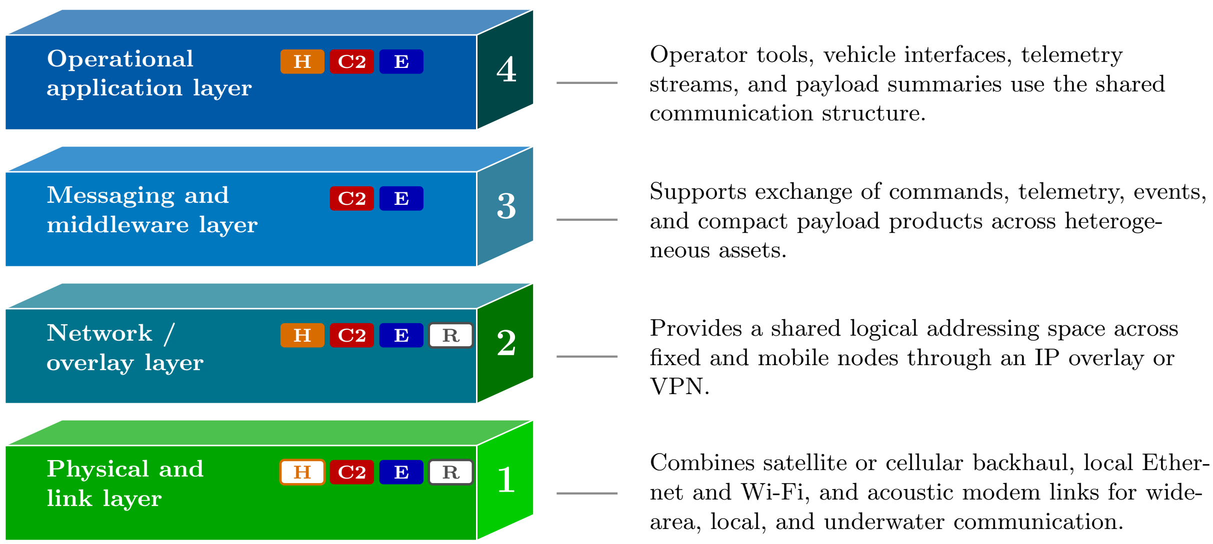

Fig 2. Layered communication design for ROC-centered heterogeneous operations. Filled tags indicate traffic classes that are routinely supported at a given layer; outlined tags indicate selective or capacity-dependent support. H = human communication, C2 = command-and-control and housekeeping telemetry, E = compact edge products, and R = raw or near-raw data.

At the physical and link level, the architecture combines satellite or cellular backhaul, local Ethernet and Wi-Fi, and acoustic modem links to cover OTH, near-field, and underwater communication needs. Above this, an IP-based overlay provides a shared logical addressing space across fixed and mobile nodes through an overlay network or virtual private network (VPN) spanning the underlying access technologies. On top of the network layer, a common messaging and middleware layer supports exchange of commands, telemetry, events, and compact payload products across heterogeneous assets. In practice, this layer is implemented through vehicle middleware or message buses that provide shared message formats, routing, and logging across the system. At the operational application level, operator tools, vehicle interfaces, telemetry streams, and payload summaries use this shared communication structure while preserving content roles and transfer priorities.

The figure also indicates how the main traffic classes relate to the layered design. Command-and-control traffic (C2) and compact edge products (E) are central throughout the communication stack. Human communication (H) is primarily associated with the operational application and network layers, but depends on available link capacity at the physical and link layer. Raw or near-raw data (R) is treated more selectively, reflecting that these products are generally transferred only when network and link capacity allow.

A key design objective is to decouple operational connectivity from any single access technology. Using overlay networking and gateway functions, the system maintains a common logical network even when the underlying links differ in bandwidth, latency, or intermittency. This supports flexible participation by ROC operators, local field personnel, and remote scientists, while allowing vehicles and field nodes to appear as part of the same operational environment despite being connected through separate physical channels.

For underwater and over-the-horizon systems, multiple communication modes coexist. High-bandwidth radio or wired links are used when assets are locally connected, whereas lower-bandwidth acoustic links are used for submerged vehicles and satellite links are used beyond local radio range. This creates a layered connectivity model in which critical C2 traffic is prioritized across available links, compact edge products are transmitted over constrained links when link budgets allow, and larger raw-data transfers are deferred to periods with greater capacity.

Content roles and transfer priorities

The framework does not require one specific protocol. It requires commands, telemetry, events, and compact payload summaries to use message formats that can be routed, logged, and interpreted consistently across asset types, gateways, and operator interfaces under intermittent connectivity.

Payload transfer is governed by operational role and available link capacity. Compact edge products (E) are prioritized for in-mission transfer because they support live interpretation and retasking over constrained links. Raw or near-raw data (R) are retained for scientific review, archiving, and higher-fidelity reprocessing, and are transferred when capacity allows.

The communication design is content-aware: link capacity, latency, and intermittency determine which content types are transferred, when, and at what fidelity. High-bandwidth links can support human communication (H), command-and-control (C2) traffic, compact edge products (E), and selected raw or near-raw data transfers (R), whereas low-bandwidth links are primarily allocated to C2 traffic and compact edge products.

Mission planning and sampling strategies

Targeted sampling is the mission layer that allocates sampling effort toward features of biological interest. It includes the search for candidate targets, their refinement, and their follow-up through additional observation or sampling.

In this framework, targeted sampling is a systems function built on remote operation, edge processing, and communication design. It translates those capabilities into sampling strategies and evaluates how they affect mission execution, sampling outcomes, and operational constraints in heterogeneous robotic operations.

Strategy classes and sampling logic

Targeted sampling approaches are grouped into three operational classes. *Reconnaissance strategies* establish broad spatial or temporal context and identify candidate areas or depth intervals of interest. *Adaptive single-platform strategies* modify the behavior of one vehicle on the basis of its own onboard or near-real-time observations. *Cooperative multi-platform strategies* use observations from one platform to cue, constrain, or retask another.

The main difference between these classes lies in how information is generated and acted upon. Reconnaissance strategies are used to reduce the search space and identify candidate targets for later follow-up. Adaptive single-platform strategies use observations from the same vehicle to adjust local sampling behavior during the mission. Cooperative multi-platform strategies distribute these roles across the fleet, allowing one platform to detect or refine a target and another to validate or sample it.

Across these classes, the operational logic is similar: a candidate feature is identified, the supporting evidence is assessed, and a sampling response is selected. The resulting sampling response takes the form of retasking, trajectory or depth adjustment, repeated coverage, or follow-up by a second platform for validation or in situ sampling.

The exact trigger rules, autonomy logic, and optimization criteria differ between strategies and implementations. The strategy classes are used to describe how sensing, edge processing, and communication capabilities are translated into targeted sampling actions in the overall system.

Evaluation basis for targeted sampling strategies

Targeted-sampling strategies are evaluated using three dimensions: operational implications, task success, and limitations. Operational implications describe how the strategy affects mission execution beyond the immediate sampling task. Task success describes whether the intended sampling response was achieved, including whether the target was detected, characterized, validated, or sampled as intended. Limitations describe factors that reduce confidence in the outcome or constrain practical use under the observed operating conditions.

This structure provides a basis for comparing strategy classes while allowing the specific indicators to be tailored to each mission and operating context.

Case-study implementation and results

The generic system framework was applied in a field campaign conducted in Mausund, Norway, in June 2024. The campaign examined how ROC-centered heterogeneous robotic operations could support field observation, detection, and targeted sampling of zooplankton.

Case-study implementation

The robotic system comprised LAUV Thor, the USVs Grethe and AutoNaut, and NTNU's workboat Flyer. The autonomous platforms carried EK80 broadband echosounders configured for zooplankton scattering-layer detection. The AUV additionally carried a SilCam for in situ particle imaging [8], together with CTD and chlorophyll-a sensing. The workboat served as a local support and communication node linking field operations and the ROC.

The campaign was organized into five experiment types. Broad-scale USV reconnaissance and USV-based layer detection provided reconnaissance context. Adaptive AUV missions represented adaptive single-platform sampling. Cooperative USV–AUV missions represented cooperative multi-platform sampling, with USV detections used to cue targeted subsurface follow-up by the AUV. AUV depth-resolved sampling provided subsurface observations across these activities.

Remote operation and coordination

Operational coordination outcomes

The operational implication considered for remote coordination was the ability to run parallel activities across assets with a small on-site team while supervision, interpretation, and coordination were handled in the ROC. This shifted planning, monitoring, and cross-platform interpretation away from the field site, allowing the local team to focus on launch, recovery, and support tasks.

During parallel operations, ROC operators monitored ongoing AUV or USV missions while the local field team handled launch, recovery, and support for other assets. A common schedule, shared displays, communication channels, and integrated logging maintained a shared operational picture across the distributed team. This required active status sharing between ROC operators, the local field team, and remote scientific participants, including vehicle positions, mission status, planned and ongoing activities, communication status, and relevant payload observations. The ROC setup used to support this workflow is shown in Fig 3.

Fig 3. Remote Operations Center during parallel USV and AUV operations. The photograph shows shared displays and two operator workstations, with the USV station to the left and the AUV station to the right.

The main operational limitations were weather windows and launch-recovery constraints. Within those constraints, the campaign was coordinated from shore with a small on-site team across multiple assets and mission types.

Cost and emissions in research operations

The scenario framework was applied to the research context using three research-operation scenarios. Research Scenario I-a and Research Scenario I-b are counterfactual ship-based reference scenarios, representing campaigns based on a smaller coastal research vessel, RV *Gunnerus*, and a larger ocean-going research vessel comparable to RV *Helmer Hanssen*, respectively. Research Scenario II represents the ROC-supervised robotic campaign conducted in Mausund. The selected metrics were total campaign cost and the corresponding operational CO_2e proxy, both evaluated over the defined five-day horizon and system boundary.

Table 2 summarizes the cost comparison. In this comparison, the ship-based reference scenarios are driven mainly by vessel access and embarked personnel, whereas the ROC-supervised robotic scenario is driven mainly by field logistics, local support, and ROC-related operations. The ROC-supervised robotic scenario had lower operational cost than both ship-based reference scenarios.

Table 2. Research operations cost comparison. Five-day campaign cost comparison across ship-based and ROC-supervised robotic scenarios.

Item	Factor used (cost)	Cost (EUR)
Research Scenario I-a		
Ship day rate (incl. fuel, crew, workboat)	6,535 EUR/day × 5 days	32,675
Personnel travel to embarkation port	20 EUR/person × 13 persons	260
Accommodation on ship	46 EUR/(person·day) × 13 persons × 5 days	2,990
Cruise allowance (scientists)	172 EUR/(person·day) × 10 persons × 5 days	8,600
Subtotal		**44,525**
Research Scenario I-b		
Ship day rate (incl. fuel, crew, workboat)	21,500 EUR/day × 5 days	107,500
Personnel travel to embarkation port	250 EUR/person × 13 persons	3,250
Accommodation on ship	46 EUR/(person·day) × 13 persons × 5 days	2,990
Cruise allowance (scientists)	172 EUR/(person·day) × 10 persons × 5 days	8,600
Subtotal		**122,340**
Research Scenario II		
ROC rent	27 EUR/h × 12 h/day × 5 days	1,620
Workboat	43 EUR/h × 12 h/day × 5 days	2,580
Workboat fuel	1.72 EUR/L × 11 L/h × 12 h/day × 5 days	1,135
Field-station lodging	370 EUR/day × 5 days	1,850
Meals (field engineers)	42 EUR/(person·day) × 3 persons × 5 days	630
Field allowance	77 EUR/(person·day) × 3 persons × 5 days	1,155
Field-team car rental (one-off)	1,462 EUR	1,462
Robotic system transport: crane truck (one-off)	4,385 EUR	4,385
Field-team transport fuel	1.72 EUR/L × 25 L	43
Subtotal		**14,860**
Δ (I-a − II)		**29,665**
Δ (I-b − II)		**107,480**

Price basis. Costs are expressed in EUR at nominal 2024–2025 prices and rounded to whole euros. NOK values were converted using an average 2024 exchange rate of 1 EUR = 11.63 NOK. Gasoline was estimated at 20 NOK/L. Costs use actual 2024 rates or claims where available; otherwise, public list rates or market-based estimates were used. Personnel travel differs between Research Scenario I-a and Research Scenario I-b, with local travel assumed for Research Scenario I-a and domestic return travel for Research Scenario I-b. The Research Scenario I-b ship day rate is a conservative estimate based on comparable published vessel rates.

Table 3 reports the corresponding emissions proxy. In this comparison, the ship-based reference scenarios are dominated by vessel fuel use. In Research Scenario II, the emissions proxy is dominated by local workboat fuel, with field transport and ROC electricity contributing smaller shares. The ROC-supervised robotic scenario had a lower operational fuel, electricity, and travel emissions proxy than both ship-based reference scenarios.

Table 3. Research operations emissions proxy. Five-day operational and travel CO_2e proxy for ship-based and ROC-supervised robotic scenarios.

Item	Factor used (emissions)	CO_2e (t)
Research Scenario I-a		
Ship fuel (diesel)	90 L/h × 12 h/day × 5 days × 2.66 kg CO_2e/L	14.364
Workboat fuel (gasoline)	11 L/h × 12 h/day × 5 days × 2.32 kg CO_2e/L	1.531
Personnel local travel	13 persons × 20 km/person × 0.151 kg CO_2e/km	0.039
Subtotal (Scenario I-a)		**15.934**
Research Scenario I-b		
Ship fuel (diesel)	225 L/h × 24 h/day × 5 days × 2.66 kg CO_2e/L	71.820
Workboat fuel (gasoline)	11 L/h × 12 h/day × 5 days × 2.32 kg CO_2e/L	1.531
Personnel domestic travel	13 persons × 1,600 passenger-km/person × 0.195 kg CO_2e/passenger-km	4.056
Subtotal (Scenario I-b)		**77.407**
Research Scenario II		
Workboat fuel (gasoline)	11 L/h × 12 h/day × 5 days × 2.32 kg CO_2e/L	1.531
Field-team transport fuel (gasoline)	25 L × 2.32 kg CO_2e/L	0.058
Robotic system transport: crane truck	356 km × 0.572 kg CO_2e/km	0.204
ROC electricity (grid)	63.888 kWh/day × 5 days × 0.01185 kg CO_2e/kWh	0.004
Subtotal (Scenario II)		**1.797**
Δ (I-a − II)		**14.137**
Δ (I-b − II)		**75.610**

Notes. Ship fuel rates are represented as 50/50 cruise–standby averages for both ship-based reference scenarios. The smaller coastal-vessel scenario Research Scenario I-a assumes 12 h/day ship operation, while the offshore-vessel reference scenario Research Scenario I-b assumes continuous 24 h/day ship operation. Workboat fuel is included in all research scenarios as in Table 2. Personnel travel is local for Research Scenario I-a and domestic return travel for Research Scenario I-b. Crane-truck distance is based on a 356 km Trondheim–Mausund–Trondheim round trip. ROC electricity use is based on the assumed equipment load over 12 h/day, detailed in S2 Table. Fuel and electricity conversion factors are described in the methods text.

Cost, coverage, and emissions in commercial survey scenarios

The same framework was also applied to a commercial-survey context, where cost was normalized by delivered coverage and reported as EUR/km. In this context, the ship-based case is represented by Commercial Scenario I, while the ROC-supervised robotic case is represented by Commercial Scenario II. Table 4 summarizes the cost and coverage comparison.

Table 4. Commercial survey cost and coverage comparison. Daily cost, delivered coverage, and cost per trackline-kilometer for ship-based and ROC-supervised robotic scenarios.

Item	Factor used (cost / coverage)	Cost (EUR/day)	Coverage (km/day)	EUR/km
Commercial Scenario I				
Trawler (incl. fuel and crew)	12,898 EUR/day; 9.26 km/h × 12 h	**12,898**	**111.12**	**116**
Commercial Scenario II				
LAUV Thor	440 EUR/day; 3.6 km/h × 12 h	440	43.20	
USV Grethe	258 EUR/day; 7.2 km/h × 12 h	258	86.40	
AutoNaut	258 EUR/day; 3.6 km/h × 12 h	258	43.20	
ROC operator coverage	589 EUR/day; no direct survey coverage	589	–	
Subtotal (Scenario II)		**1,545**	**172.80**	**9**
Δ (I − II)				**107 EUR/km**

Notes. EUR/km is computed as cost/day divided by km/day and rounded to whole euros. Operator coverage contributes to cost but not directly to delivered trackline coverage. Platform rates for Commercial Scenario II are treated as service rates and assumed to include ordinary deployment and support unless otherwise specified. Costs are expressed in EUR at nominal 2024–2025 prices. NOK values were converted using an average 2024 exchange rate of 1 EUR = 11.63 NOK. Actual 2024 rates or claims were used where available; otherwise, public list rates or market-based estimates were used. No CPI adjustment was applied.

Under the stated assumptions, the robotic scenario had lower cost per delivered trackline-kilometer than the ship-based scenario. The trawler scenario concentrated cost in a single high-capacity vessel, whereas the robotic scenario distributed cost across multiple autonomous assets and ROC operator coverage. The comparison therefore indicates lower cost per delivered trackline-kilometer for the robotic scenario under the stated service-rate and coverage assumptions.

Table 5 reports the corresponding operational CO_2e proxy.

Table 5. Commercial survey emissions proxy. Daily operational CO_2e proxy for ship-based and ROC-supervised robotic scenarios.

Item	Factor used (emissions)	CO_2e (t/day)	CO_2e (kg/km)
Commercial Scenario I			
Trawler fuel (diesel)	90 L/h × 12 h/day × 2.66 kg CO_2e/L	2.873	25.85
Commercial Scenario II			
ROC and platform electricity	100 kWh/day × 0.01185 kg CO_2e/kWh	0.001	0.01
Δ (I − II)		**2.872**	**25.84**

Notes. Emissions are reported as an operational fuel/electricity proxy, not a full life-cycle inventory. Trawler fuel use follows the lower-end diesel rate used for the smaller research-vessel scenario. The robotic scenario includes 100 kWh/day for ROC operation and autonomous-platform charging. AutoNaut propulsion was treated as wave-powered, and USV Grethe as electrically powered. CO_2e per km is computed from daily emissions and delivered coverage in Table 4.

The commercial emissions proxy was dominated by trawler fuel use in the ship-based scenario, while the robotic scenario was represented by ROC and platform electricity use.

Edge processing for data quality and zooplankton detection

Case-study implementation of edge processing

In the Mausund 2024 case study, acoustic edge processing was implemented on the EK80 computer using a near-real-time pipeline based on the Echopype Python package [32]. Incoming echograms were processed file by file to derive compact acoustic summaries for live operational use without relying on transfer or inspection of full acoustic records during the mission. The acoustic outputs consisted of estimates of the upper and lower bounds of the dominant scattering layer together with recent data-quality indicators, including ping dropout.

For optical sensing, SilCam data were processed with the PySilCam Python package [33] into segmented particle detections and per-particle summary tables, allowing counts and size-based summaries to be reviewed without full raw-image inspection during operations. These products formed the optical outputs available for ROC-side review and operational interpretation during the campaign.

Operational data assessment

ROC-side operational data assessment focused on whether EK80 layer-bound outputs were consistent with raw and processed echograms, vehicle motion, sea state, and mission context. One operational quality indicator was ping dropout, defined as the number of dropped pings within a time window of 120 s. A ping was classified as dropped when the received signal strength S_v fell more than 40 dB below the global maximum.

For SilCam, available AUV missions produced segmented particle detections and per-particle summary tables with count- and size-based information along the vehicle track. These summaries were reviewed without inspection of the full raw-image stream and were used to assess particle occurrence during AUV sampling.

Cross-platform interpretation

For overlapping segments, acoustic detections from surface platforms and optical observations from the AUV were reviewed together to interpret candidate scattering layers. The review assessed whether acoustic features occurred within the depth range sampled by the AUV and coincided with particle observations from SilCam. Fig 4 shows representative examples of the acoustic observations, optical particle regions of interest, and SilCam-derived particle summaries used during this interpretation.

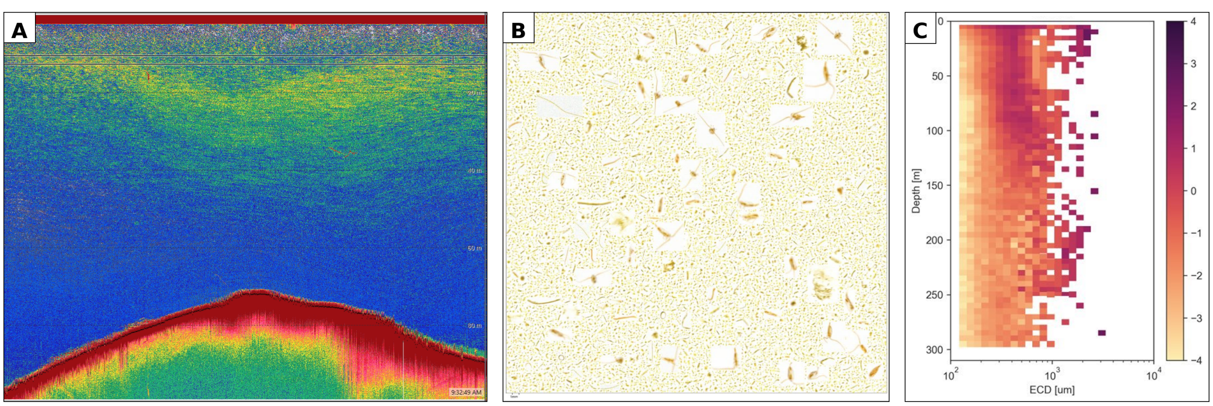

Fig 4. Representative acoustic and optical data products from the Mausund campaign. (A) Example EK80 echogram segment illustrating water-column backscatter structure used during acoustic interpretation. (B) Montage of representative SilCam particle regions of interest from AUV sampling, illustrating optical observations available for particle-level interpretation. (C) Example SilCam-derived particle size distribution by depth. The horizontal axis shows equivalent circular diameter (ECD), the vertical axis shows depth, and colors indicate log volume concentration within each size and depth bin.

For missions separated in time or location, interpretation was limited to mission-scale consistency across neighboring transects and successive missions. At this scale, acoustic, optical, and net-sampling observations supported the interpretation that zooplankton-relevant scattering and particle features were present in the study area. The campaign did not include enough co-located and synchronized acoustic-optical samples to estimate quantitative agreement between acoustic layer detections and optical particle observations for the targeted-sampling events.

Communication design outcomes: usefulness for science and operations

Case-study communication architecture

In the Mausund case study, the communication architecture connected the ROC, field site, support platforms, autonomous vehicles, and remote participants through multiple communication layers. The physical and link layer combined Starlink satellite backhaul, cellular connectivity, local Wi-Fi and Ethernet, acoustic modems, and Iridium low-bandwidth satellite links. A mesh VPN overlay connected ROC workstations, field nodes, support platforms, vehicles, and remote participants into a shared logical network across these access technologies.

On top of the network layer, vehicle and operator communication relied on the LSTS software and protocol stack, including IMC-based message exchange, DUNE onboard middleware, and Neptus for command, monitoring, and log review [34]. This provided a common operational environment across vehicles, gateways, and ROC workstations while allowing payload-specific acquisition systems to run on dedicated onboard computers.

At the field site, Starlink provided the primary backhaul between the local operational setup and the ROC, with 4G/5G available as an additional or backup path. AutoNaut primarily used 4G connectivity, with Iridium as a low-bandwidth fallback outside cellular coverage. The ROC was connected through a terrestrial internet connection, and remote participants joined through the same mesh VPN.

For the AUV, Wi-Fi was used for close-range operations and higher-throughput transfer near the workboat or during launch and recovery. Acoustic communication was used when the vehicle was submerged and beyond radio range, carrying essential C2 traffic and minimal edge telemetry. GSM or Iridium links were used when the vehicle was surfaced outside Wi-Fi range, supporting limited C2 traffic and compact payload summaries. These channels were aggregated through the local gateway setup at the field site and forwarded to the ROC when connectivity was available.

Link performance and continuity

Recorded uplink rates during active missions were 60–460 kB/s. Within this operating range, the connection supported simultaneous navigation monitoring, camera feed, and graphical remote-desktop access to the payload computer during operational testing. Periods of degraded performance occurred intermittently, particularly during heavy vessel motion and adverse weather, but these were brief and did not prevent continued mission supervision.

In the cooperative USV–AUV missions, compact acoustic layer products derived on the surface platform were transmitted and used for vehicle-to-vehicle coordination during execution. Communication continuity was also demonstrated during one recovery delay, when poor weather prevented immediate AUV recovery and the ROC continued to monitor vehicle state until local recovery could resume.

Compression performance and communication fit

Onboard data reduction made SilCam and EK80 products small enough for operational transfer and retained the information needed for targeting and review.

For SilCam, the compressed product consisted of summary tables and segmented ROI files aggregated over June 6–7, 2024. Compression is reported per image for two baselines: continuous acquisition at 4 Hz over the mission span (92,832 images total) and a detections-only baseline counting frames containing at least one particle (12,256 images). Relative to the raw-image baseline, the reductions were approximately 8,836× (99.99%) for continuous acquisition and 1,167× (99.91%) for detections-only output.

For EK80, the vehicle-to-vehicle (V2V) profile reduced each ping to a compact summary containing only the information needed for coordination and transmission over constrained links, while the vehicle-to-ROC (V2ROC) profile provided a richer export for scientific inspection and archiving. The corresponding reductions were approximately 1,000× (99.90%) for V2V and 2,052× (99.95%) for V2ROC.

Table 6 summarizes representative compression results for products selected for transmission or exchange under the layered communication setup.

Table 6. Compression overview for SilCam and EK80 data. SilCam values are reported per image using two baselines: continuous acquisition at 4 Hz over the mission span and detections-only acquisition, counting only raw images containing detected particles. EK80 values are reported per ping and per raw file, matching the export profiles used for vehicle-to-vehicle (V2V) and vehicle-to-ROC (V2ROC) transfer.

Data	Baseline	Raw (kB)	Compressed (kB)	Raw/Compressed
SilCam	Per image, continuous	14 703	1.664	8 836.1
SilCam	Per image, detections only	14 703	12.604	1 166.6
EK80	Per ping, V2V	150	0.15	1 000.0
EK80	Per file, V2ROC	30 784	15	2 052.3

These results show that both SilCam and EK80 products were reduced substantially while retaining compact summary formats suitable for transmission under the communication profiles used in the campaign.

Targeted sampling and heterogeneous fleet coordination

Adaptive AUV sampling

The adaptive AUV missions examined whether onboard decision-making based on chlorophyll-a could guide sampling toward candidate high-interest regions with limited human intervention. In these missions, the AUV used fluorometer measurements to update an onboard statistical model. A decision criterion then used this model to select the next sampling location, prioritizing regions associated with elevated chlorophyll-a. Within the tested missions, the AUV reached and explored depth intervals associated with the highest observed chlorophyll-a concentrations.

These missions also included onboard optical processing of SilCam data. Although the adaptive control logic was based on chlorophyll-a observations, the simultaneous SilCam processing showed how in situ optical information could support confirmation or refinement of candidate targets during decision making.

In addition to the adaptive objective, the missions produced concurrent optical and environmental observations that could be interpreted relative to local stratification and particle occurrence. In the tested cases, SilCam observations indicated elevated particle or copepod-sized particle occurrence in the same upper mixed-layer region where chlorophyll-a was highest. The evidence nevertheless remained limited in scope, since the number of experiments was small and the observations were restricted to a specific season and location.

Cooperative USV–AUV targeted sampling

A cooperative USV–AUV targeting strategy was tested in missions where a surface platform first searched for candidate acoustic layers and the AUV then performed subsurface follow-up at selected positions. The USV provided surface-based acoustic screening, and the AUV provided follow-up sampling at the selected sites.

In the collaborative run, the USV identified three sites of interest, and the AUV was directed to a preset depth of 15 m at each target location. Fig 5 documents the USV-side decision trace used during this run. The ping-drop indicator shows whether acoustic data quality was degraded, the layer-confidence output shows support for a detected scattering layer, and the mission-state transitions show when the logic moved from scouting to evaluation and AUV launch. Together, these outputs document how surface-based edge products supported target selection during the cooperative mission.

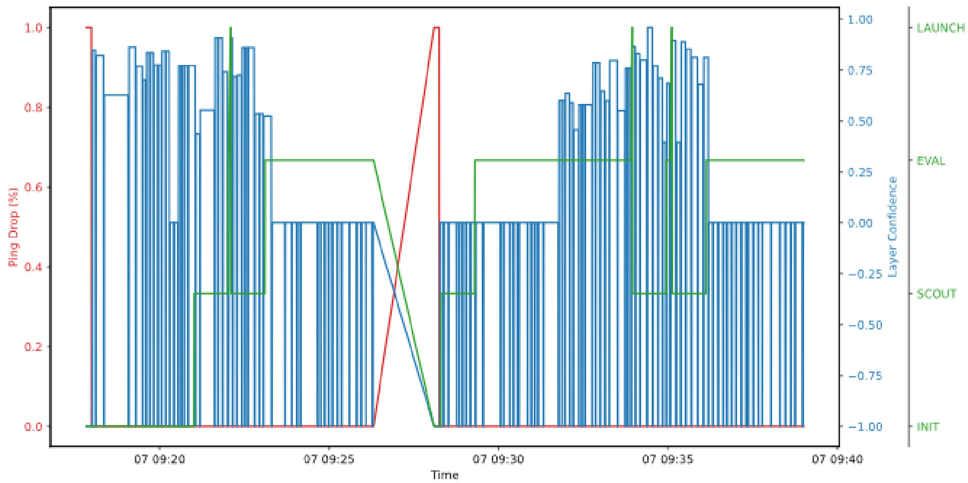

Fig 5. USV-side decision trace from the cooperative USV–AUV targeted-sampling mission. The plot shows the surface-platform metrics used during the scouting phase: ping-drop indicator (red), layer-confidence output (blue), and mission-state transitions (green). The ping-drop indicator represents acoustic data-quality status, the layer-confidence output represents support for a detected scattering layer, and the mission states show the progression from initialization (INIT) and scouting (SCOUT) to evaluation (EVAL) and AUV launch toward a selected target (LAUNCH).

The operational goal was to reduce AUV submerged time by using USV detections to select a limited number of follow-up locations. Task success was evaluated in terms of whether the cooperative strategy could guide the AUV to USV-selected targets with reduced dive effort. In the reported run, the three targeted dives corresponded to a total AUV dive time of 21 min, compared with 40 min for the full mapping survey used as the reference case.

Task-level evaluation of the USV component focused on acoustic data quality and the ability to identify candidate scattering layers from the surface-based acoustic data. In the reported USV-only runs, sea state was calm and no ping-dropout issue was observed during the relevant tests.

The main limitations were end-to-end validation and processing reliability. SilCam did not record usable data during the AUV verification phase because of an operational error, so the intended in situ confirmation step could not be completed. The run also experienced a lost connection to the WBT Mini processing unit, visible as a mid-run reinitialization of the algorithm.

Biological reference sampling

Net sampling provided an independent biological reference through two successful surface-towed hauls from MS *Frida* on 7 June 2024. The gear was a fine-mesh net with an opening diameter of approximately 3.2 m, towed with the center of the opening at approximately 2 m depth. A haul in Mausundfjorden covered approximately 950 m at an average tow speed of 1.34 kn and yielded approximately 0.5 kg of *Calanus*. A haul near the edge of Frohavet covered approximately 1200 m at an average tow speed of 1.21 kn and yielded approximately 2 kg, indicating higher abundance there. Preserved subsamples from both hauls were dominated by late-stage *Calanus finmarchicus*, mainly copepodite stages C5 and C6.

Discussion

Remote operation and coordination

The Mausund case study shows that remote operation redistributed operational responsibilities between shore and field. The ROC handled mission supervision, interpretation, replanning, and cross-platform coordination in parallel with field execution, while the field team remained focused on launch, recovery, and local safety-critical tasks. This reduced the need to colocate all scientific and operational expertise at the field site and allowed decision support to remain active when local conditions constrained the pace of operations.

A key benefit was the ability to combine multiple information sources in real time. Vehicle telemetry, payload summaries, environmental context, communication status, and operational logs were reviewed side by side and interpreted jointly. This improved the ability to detect emerging opportunities or constraints during missions and supported more timely adjustments than would be practical in a field-only setup. In that respect, the ROC contributed by integrating these information sources into a shared basis for interpretation and decision-making.

The results also show that remote operation changed staffing patterns rather than simply reducing staffing. In the campaign, effort shifted away from field presence and toward supervision, monitoring, coordination, and interpretation. This was operationally useful when expertise needed to be distributed and shared, but it also created a risk of diffuse roles or information overload. Remote operation therefore required explicit ownership of operational functions, clear handover procedures, and agreed decision rights.

The cost and emissions comparison supports the same general interpretation. Within the assumptions used in this study, the ROC-supervised robotic scenarios had lower recurring operational cost and lower operational CO_2e than the ship-based alternatives. The comparison also showed a change in cost structure: ship-based scenarios were dominated by vessel access and fuel, whereas the ROC-supervised robotic scenario concentrated costs in logistics, local support, and distributed operational infrastructure.

The campaign also highlights the limits of the ROC-centered concept. Its effectiveness still depended on field execution quality, weather windows, launch and recovery constraints, and reliable communication between field and shore. Remote operation did not remove these dependencies, but it made it easier to absorb delays and continue useful work elsewhere in the system when they occurred. The operational benefit lay in greater resilience and flexibility in how work was organized across shore and field.

Operational value of edge processing

The main contribution of edge processing in this campaign was to bring payload interpretation closer to the live operational loop. By converting raw acoustic and optical payload data into compact products during operations, including layer bounds, quality indicators, and particle summaries, edge processing reduced the dependence on post-mission interpretation before operational decisions could be made. This was particularly relevant for a dynamic target field, where the value of information depends on whether it is available while the vehicles are still deployed. In this respect, edge processing functioned as the bridge between sensing and operational decision-making.

A central outcome was that data quality became part of live operational assessment rather than only post-processing. Acoustic layer products and ping-dropout indicators were used to judge whether surface-based detections were reliable enough to support follow-up action, while optical particle summaries were reviewed to assess whether the resulting in situ observations were operationally usable. This was especially important in the cooperative missions, where one platform identified a candidate feature and another was tasked to validate or sample it. In that setting, weakly supported detections could lead to inefficient or misdirected mission actions.

The campaign also indicates that edge products were most useful when interpreted jointly across platforms and modalities. Acoustic products from the surface vehicles described the depth and structure of candidate scattering layers, while the AUV's optical observations provided in situ information on particle occurrence and size structure in the sampled water column. This combination supported interpretation of whether an acoustic layer was a plausible biological target and helped guide follow-up sampling decisions. Net sampling strengthened that interpretation by confirming the presence of *Calanus* in the area.

The main limitation was that the evidence was stronger on operational usefulness than on formal validation. The case study showed that edge products were generated, reviewed, communicated, and used in practice, but it contained limited quantitative analysis of detection performance, uncertainty, or agreement across products. The results therefore support improved operational usability, whereas stronger statements about detection accuracy or cross-platform validation require more detailed analysis. A practical implication was that interpretable processing outputs were useful because operators could inspect them during live operations.

Communication robustness in heterogeneous networks

The communication results support the design choice of separating traffic into human communication (H), command-and-control and housekeeping telemetry (C2), compact edge products (E), and raw or near-raw data (R), and matching these categories to the available link classes. In the Mausund setup, higher-capacity links supported operator interfaces, telemetry, and transfer of selected payload products, whereas constrained links were used mainly for C2 traffic and compact edge products. This was evident in the cooperative USV–AUV missions, where compact layer products were transmitted and used in ongoing coordination, and in periods of degraded connectivity, where mission supervision continued even when richer data access was limited.

The case study also indicates that the layered design was useful because it combined redundancy with prioritization. This was evident in the cooperative ASV-AUV missions and during intermittent link degradation, when shore-based supervision continued despite reduced access to richer data products.

The large compression reduction factors obtained for both SilCam and EK80 products show that the system did not depend on transferring raw payload streams in order to remain operationally useful during the mission. Compact edge products supported targeting and coordination, whereas richer exports were transferred when capacity allowed. This points to a close coupling between communication design and edge processing in the overall operational logic.

The results also show that communication robustness in the heterogeneous network did not depend on any single technology. Although Starlink was highly enabling in this campaign, continued operation still depended on multiple link types, transition points, and platform-specific constraints. A key property of the system was its ability to degrade gracefully while preserving C2 traffic, compact feature products, mission-state updates, and platform telemetry in forms that remained interpretable across platforms, gateways, and operator tools. This operational coherence made multi-vehicle coordination possible under variable link conditions.

Targeted sampling with heterogeneous fleets

The results suggest that the main value of targeted sampling in this system lay in the way multiple strategy classes were combined into one coordinated sampling logic. Broad-scale reconnaissance, adaptive single-platform behavior, and cooperative inter-platform retasking served distinct roles in reducing the search space and concentrating effort where expected scientific return was highest. In this respect, targeted sampling functioned as a system-level capability rather than a vehicle-level autonomy function only.

The adaptive AUV missions show that single-platform adaptation was operationally useful when linked to a clear objective, here the exploration of chlorophyll-a structure with limited human intervention. These missions also show that adaptive mission logic can be embedded within a broader ROC-centered system and can generate candidate information for later biological follow-up. Their main limitation was interpretive scope: the number of experiments was small, and the observations were too context-specific to support broader ecological claims.

The cooperative ASV–AUV missions provided a clearer example of the heterogeneous systems concept. One platform acted as a scout and another as the higher-resolution follow-up sampler, with edge products and communication linking the two. In the reported comparison, the targeted dives reduced AUV dive time relative to the full mapping reference, indicating a practical benefit of cooperative targeted sampling. At the same time, the campaign showed that the full value of this strategy depended on reliable end-to-end execution from detection and retasking to in situ follow-up and validation.

Together, these results indicate the operational value of assigning distinct sampling roles to different assets. Broader-area platforms supported search and refinement, while in situ platforms supported follow-up and validation. The case study therefore supports the framework view of targeted sampling as a coordinated mission logic across the fleet, although the quantitative evidence remains limited for judging which specific strategy is best.

Limitations and transferability

These results point to a common pattern across the framework components: the main strength of the system lay in the integration of operational roles, compact information products, communication-aware coordination, and targeted mission logic into one coherent workflow. The ROC was central because it provided the environment in which these components were interpreted together and turned into decisions. In that respect, the case study supports the framework view that heterogeneous robotic operation is most effective when treated as a systems-integration problem rather than as a collection of isolated technical components.

The case study illustrates why zooplankton is a relevant use case. For dynamic low-trophic resources, useful assessment depends on the ability to detect, map, and characterize biomass distributions with sufficient spatial and temporal resolution. In this campaign, the robotic observations were complemented by net sampling, which confirmed the presence of *Calanus* in the area and provided independent biological context for interpreting the robotic results.

The study also has important limitations. First, the campaign was short and geographically specific, restricting the generality of ecological interpretations and operational performance claims. Second, several of the strongest concepts in the paper, especially edge-enabled validation and cooperative targeted sampling, were demonstrated more clearly in operational logic than in complete quantitative validation. Third, the evaluation frameworks were intentionally simplified, especially the cost–carbon comparison. Finally, the system was built around a particular fleet, communication stack, and operational context, all of which shaped what was feasible in practice.

The transferability of the work lies primarily in the design principles rather than in the exact configuration. Separating communication by content role, reducing raw data to operationally meaningful edge products, dividing labor across heterogeneous platforms, and using the ROC as a point of evidence integration are principles that can be applied in other fleets and environments. Transfer to harsher or more remote settings would require stronger fallback autonomy, more conservative communication assumptions, and more explicit operational minima.

A practical implication is that future improvements should focus on tighter coupling between sensing, communication, coordination, and validation. In this study, the clearest opportunities were clearer uncertainty reporting on edge products, more explicit operational trigger logic, and stronger links between surface detections and in situ confirmation.

Conclusion

This paper has presented a system framework for ROC-centered heterogeneous robotic operations in zooplankton-oriented marine research and examined it through a field campaign conducted in Mausund in 2024. The case study showed how remote operation, edge-derived information products, communication design, and targeted sampling can be combined within one operational workflow.

The results indicate that the main value of the ROC was shifting supervision, interpretation, and coordination away from the field site and into a shared shore-based environment. Edge processing contributed compact products that were usable during operations, and the layered communication design helped preserve command-and-control traffic and compact science-relevant products across links with varying capacity, latency, and reliability. The campaign also showed that heterogeneous platform roles can be combined in broad-area search, local refinement, and in situ follow-up, supported by independent biological context from net sampling.

The work demonstrates the value of treating heterogeneous robotic operations as an integrated system problem. Future work should focus on stronger validation of edge products, clearer uncertainty handling, more robust cooperative sampling chains, and broader testing across campaigns, environments, and fleet compositions.

Funding

This work was supported by the Research Council of Norway through SFI Harvest, project number 309661. André Olaisen was supported by the Research Council of Norway through MASCOT—Maritime Autonomous Sampling and Control, project number 305445.

Data availability

The data underlying the findings reported in this study are provided within the manuscript and its supporting information. The manuscript reports the numerical values used for the cost, emissions, compression, communication-performance, and mission-outcome comparisons, and the supporting information provides the input schema and equipment-load assumptions used for the cost–CO_2e framework.

Use of AI-assisted tools

AI-based tools were used to assist with language editing, formatting, and evaluation of text based on the authors' own drafts. The authors reviewed and edited all content and take full responsibility for the work. AI tools are not listed as authors.

Acknowledgments

The authors thank the 2024 SFI Harvest field campaign participants for their contributions to the Mausund field campaign. The authors also thank the AUR-Lab engineers for technical support before and during the campaign, the AutoNaut team led by Tor Arne Johansen for support with the AutoNaut operations, and Mausund Field Station for local facilities and field support. The authors also acknowledge the local crew and sampling support at Mausund.

References

1. Lévy M, Franks PJS, Smith KS. The role of submesoscale currents in structuring marine ecosystems. Nature Communications. 2018;9(1):4758. doi:10.1038/s41467-018-07059-3.

2. Wiebe PH, Benfield MC. From the Hensen net toward four-dimensional biological oceanography. Progress in Oceanography. 2003;56(1):7-136. doi:10.1016/S0079-6611(02)00140-4.

3. Chai F, Johnson KS, Claustre H, Xing X, Wang Y, Boss E, et al. Monitoring ocean biogeochemistry with autonomous platforms. Nature Reviews Earth & Environment. 2020;1(6):315-26. doi:10.1038/s43017-020-0053-y.

4. Fossum TO, Eidsvik J, Ellingsen I, Alver MO, Fragoso GM, Johnsen G, et al. Information-driven robotic sampling in the coastal ocean. Journal of Field Robotics. 2018;35(7):1101-21. Available from: https://onlinelibrary.wiley.com/doi/abs/10.1002/rob.21805. doi:10.1002/rob.21805.

5. Curtin TB, Bellingham JG. Progress toward autonomous ocean sampling networks. Deep-Sea Research Part II: Topical Studies in Oceanography. 2009;56(3):62-7. doi:10.1016/j.dsr2.2008.09.005.

6. Whitt C, Pearlman J, Polagye B, Caimi F, Muller-Karger F, Copping A, et al. Future Vision for Autonomous Ocean Observations. Frontiers in Marine Science. 2020;7:697. doi:10.3389/fmars.2020.00697.

7. Pinto J, Dias P, Ribeiro M, Costa M, de Sousa JB. Networked Vehicle Systems: From Vision to Reality. In: OCEANS 2017 - Aberdeen. Aberdeen, UK; 2017. p. 1-6. doi:10.1109/OCEANSE.2017.8084823.

8. Davies EJ, Brandvik PJ, Leirvik F, Nepstad R. The use of wide-band transmittance imaging to size and classify suspended particulate matter in seawater. Marine Pollution Bulletin. 2017;115(1):105-14. Available from: https://www.sciencedirect.com/science/article/pii/S0025326X16309833. doi:10.1016/j.marpolbul.2016.11.063.

9. De Robertis A, Lawrence-Slavas N, Jenkins R, Wangen I, Mordy CW, Meinig C, et al. Long-term measurements of fish backscatter from Saildrone unmanned surface vehicles and comparison with observations from a noise-reduced research vessel. ICES Journal of Marine Science. 2019

07;76(7):2459-70. Available from: https://doi.org/10.1093/icesjms/fsz124. doi:10.1093/icesjms/fsz124.

10. Handegard NO, De Robertis A, Holmin AJ, Johnsen E, Lawrence J, Le Bouffant N, et al. Uncrewed surface vehicles (USVs) as platforms for fisheries and plankton acoustics. ICES Journal of Marine Science. 2024 09;81(9):1712-23. Available from: https://doi.org/10.1093/icesjms/fsae130. doi:10.1093/icesjms/fsae130.

11. Ohman MD, Davis RE, Sherman JT, Grindley KR, Whitmore BM, Nickels CF, et al. Zooglider: An autonomous vehicle for optical and acoustic sensing of zooplankton. Limnology and Oceanography: Methods. 2019;17(1):69-86. Available from: https://aslopubs.onlinelibrary.wiley.com/doi/abs/10.1002/lom3.10301. doi:10.1002/lom3.10301.

12. Skaugset K, de Sousa JB, Sørensen AJ. Autonomous robotic organizations for marine operations. Science Robotics. 2025;10:eadl2976. doi:10.1126/scirobotics.adl2976.

13. Vasilijevic A, Bremnes JE, Ludvigsen M. Remote Operation of Marine Robotic Systems and Next-Generation Multi-Purpose Control Rooms. Journal of Marine Science and Engineering. 2023;11(10). Available from: https://www.mdpi.com/2077-1312/11/10/1942. doi:10.3390/jmse11101942.

14. Dallolio A, Quintana-Diaz G, Honoré-Livermore E, Garrett JL, Birkeland R, Johansen TA. A Satellite-USV System for Persistent Observation of Mesoscale Oceanographic Phenomena. Remote Sensing. 2021;13(16). Available from: https://www.mdpi.com/2072-4292/13/16/3229. doi:10.3390/rs13163229.

15. Folt CL, Burns CW. Biological drivers of zooplankton patchiness. Trends in Ecology & Evolution. 1999;14(8):300-5. Available from: https://www.sciencedirect.com/science/article/pii/S016953479901616X. doi:10.1016/S0169-5347(99)01616-X.

16. Mackas DL, Denman KL, Abbott MR. Plankton Patchiness: Biology in the Physical Vernacular. Bulletin of Marine Science. 1985;37(2):652-74.

17. Melle W, Runge J, Head E, Plourde S, Castellani C, Licandro P, et al. The North Atlantic Ocean as habitat for Calanus finmarchicus: Environmental factors and life history traits. Progress in Oceanography. 2014;129:244-84. North Atlantic Ecosystems, the role of climate and anthropogenic forcing on their structure and function. Available from: https://www.sciencedirect.com/science/article/pii/S0079661114000743. doi:10.1016/j.pocean.2014.04.026.

18. Steinberg DK, Landry MR. Zooplankton and the Ocean Carbon Cycle. Annual Review of Marine Science. 2017;9:413-44. Available from: https://www.annualreviews.org/content/journals/10.1146/annurev-marine-010814-015924. doi:10.1146/annurev-marine-010814-015924.

19. Beaugrand G. Long-term changes in copepod abundance and diversity in the north-east Atlantic in relation to fluctuations in the hydroclimatic environment. Fisheries Oceanography. 2003;12(4-5):270-83. Available from: https://onlinelibrary.wiley.com/doi/abs/10.1046/j.1365-2419.2003.00248.x. doi:10.1046/j.1365-2419.2003.00248.x.

20. Helaouët P, Beaugrand G. Macroecology of Calanus finmarchicus and C. helgolandicus in the North Atlantic Ocean and adjacent seas. Marine Ecology Progress Series. 2007;345:147-65. doi:10.3354/meps06775.

21. Richardson AJ. In hot water: zooplankton and climate change. ICES Journal of Marine Science. 2008 03;65(3):279-95. Available from: https://doi.org/10.1093/icesjms/fsn028. doi:10.1093/icesjms/fsn028.

22. Cowles TJ, Desiderio RA, Carr ME. Small-Scale Planktonic Structure: Persistence and Trophic Consequences. Oceanography. 1998;11(1):4-9. Available from: https://doi.org/10.5670/oceanog.1998.08. doi:10.5670/oceanog.1998.08.

23. Harris R, Wiebe PH, Lenz J, Skjoldal HR, Huntley M, editors. ICES Zooplankton Methodology Manual. San Diego: Academic Press; 2000.

24. Curtin TB, Bellingham JG, Catipovic J, Webb D. Autonomous Oceanographic Sampling Networks. Oceanography. 1993;6(3):86-94.

25. Fiorelli E, Leonard NE, Bhatta P, Paley DA, Bachmayer R, Fratantoni DM. Multi-AUV Control and Adaptive Sampling in Monterey Bay. IEEE Journal of Oceanic Engineering. 2006;31(4):935-48. doi:10.1109/JOE.2006.880429.

26. Fossum TO, Fragoso GM, Davies EJ, Ullgren JE, Mendes R, Johnsen G, et al. Toward adaptive robotic sampling of phytoplankton in the coastal ocean. Science Robotics. 2019;4(27):eaav3041. Available from: https://www.science.org/doi/abs/10.1126/scirobotics.aav3041. doi:10.1126/scirobotics.aav3041.

27. Pinto J, Costa M, Mendes R, Lima K, Dias P, Pereira J, et al. Coordinated Robotic Exploration of Dynamic Open Ocean Phenomena. Field Robotics. 2022;2:843-71. doi:10.55417/fr.2022028.

28. Mo-Bjørkelund T, Majaneva S, Fragoso GM, Johnsen G, Ludvigsen M. Multi-vehicle adaptive 3D mapping for targeted ocean sampling. PLOS ONE. 2024;19(5):e0302514. Available from: https://doi.org/10.1371/journal.pone.0302514. doi:10.1371/journal.pone.0302514.

29. Berg CH, Røyrvik J, Fyhn H, Danielsen BE, Hauan T, Fossum K, et al. OceanLab – Subsea Facility ICT-Platform: Needs and Requirements. Trondheim, Norway: CIRiS – Centre for Interdisciplinary Research in Space, NTNU Samfunnsforskning AS; 2020. CIRiS-TN-034. Issue 1, 14 Aug 2020.

30. Norwegian Environment Agency. Utslippsfaktorer i klimagassregnskap for Norge; 2026. Accessed: 2026-06-22. Available from: https://www.miljodirektoratet.no/ansvarsomrader/klima/klimagasser-utslippstall-regnskap/utslippsfaktorer-klimagassregnskap/.

31. Norwegian Water Resources and Energy Directorate. Strømdeklarasjoner; 2026. Accessed: 2026-06-22. Available from: https://www.nve.no/energi/energisystem/energibruk/stroemdeklarasjoner/.

32. Lee WJ, Setiawan L, Tuguinay C, Mayorga E, Staneva V. Interoperable and scalable echosounder data processing with Echopype. ICES Journal of Marine Science. 2024 Dec;81(10):1941-51. Available from: https://doi.org/10.1093/icesjms/fsae133. doi:10.1093/icesjms/fsae133.

33. Davies E. PySilCam: SINTEF SilCam acquisition, processing and plotting; 2024. Software repository. Accessed 22 May 2026. https://github.com/SINTEF/PySilCam.

34. Pinto J, Marques E, Pinto J, Sousa J, et al. The LSTS Toolchain for Networked Vehicle Systems. In: OCEANS 2013 - Bergen; 2013. p. 1-9. Available from: https://www.dcc.fc.up.pt/~edrdo/publications/papers/oceans13.pdf. doi:10.1109/OCEANS-Bergen.2013.6608107.

Supporting information

S1 Table. Input schema for the cost–CO_2e comparison framework

This appendix provides a concrete input schema for the cost and CO_2e comparison framework used in the study. The schema lists generic activity items, their typical rate or quantity forms, and indicates which items are used in each context–scenario combination in this study: research ship-based (Research Scenario I), research ROC + robots (Research Scenario II), commercial ship-based (Commercial Scenario I), and commercial ROC + robots (Commercial Scenario II).

S1 Table. Input schema for cost and CO_2e comparison, showing which items apply in each context and scenario.

Item	Rate / quantity form	Use	RS I	RS II	CS I	CS II
Campaign length	days or fixed campaign	C	✓	✓		
Duty time	hours on task per day	C+E	✓	✓	✓	✓
Number of embarked scientists	persons	C	✓			
Number of field engineers	persons	C		✓		✓
Research vessel / ship access	rate per day or campaign	C+E†	✓			
ROC facility access	rate per hour/day/campaign	C		✓		✓
Accommodation / per-diem, ship	rate per person and day	C	✓			
Field-station lodging / meals	rate per day or per person and day	C		✓		
Logistics, mobilization/demob	one-off amount and/or travel distance	C+E	✓	✓		
Robotic asset use, AUV/ASVs	rate per hour/day/campaign	C				✓
Operator / technician coverage	rate per hour	C				✓
Trawler access	rate per day or campaign	C+E			✓	
Distance-based operating cost	EUR/km or distance derived from speed and time	C			✓	✓
Fuel use, CO_2e proxy	L/h or L/day	E	✓		✓	
Electric load, CO_2e proxy	kW or kWh/day	E		✓		✓

Use: C = used in cost comparison; E = used in CO_2e proxy; C+E = used in both. RS = research scenario; CS = commercial scenario. Scenario I represents ship-/trawler-based operation, while Scenario II represents ROC-supervised robotic operation.
†In this study, ship/trawler access is monetized using a day rate; the CO_2e proxy uses a separate fuel-use row.

S2 Table. ROC equipment list and assumed loads used for the electricity estimate

S2 Table. ROC equipment list and assumed loads used for electricity estimate.

Item	Unit load (W)	Quantity	Total (W)
Laptops	50	8	400
Desktop workstations	350	12	4,200
LED monitors	12	12	144
Large displays / TVs	100	5	500
Servers (light-duty)	40	2	80
Sum (connected load)			**5,324**
Derived daily energy at 12 h/day			5,324 W $\times$ 12 h = 63.888 kWh/day