

# Choreographing the Way of Water: A Computational Framework for Aquatic Robotic Art

Aswin Ramachandran  
aramachandra@ethz.ch  
ETH Zurich  
Zurich, Switzerland

Noa Sendlhofer  
nsendlhofer@student.ethz.ch  
ETH Zurich  
Zurich, Switzerland

Christopher Golling  
cgolling@student.ethz.ch  
ETH Zurich  
Zurich, Switzerland

Jan Kamm  
jakamm@student.ethz.ch  
ETH Zurich  
Zurich, Switzerland

Sebastian Burmester  
sburmester@student.ethz.ch  
ETH Zurich  
Zurich, Switzerland

Ruiheng Jiang  
rjiang@student.ethz.ch  
ETH Zurich  
Zurich, Switzerland

Raffaello D'Andrea  
rdandrea@ethz.ch  
ETH Zurich  
Zurich, Switzerland




Figure 1: The Way of Water installation: A fleet of illuminated vessels performing synchronized motion on Lake Zurich.

## Abstract

Robotic choreography in open water is governed by nonlinear fluid dynamics, which impose significant challenges due to environmental disturbances and nonlinear system dynamics. This paper presents the cyber-physical architecture of Way of Water, a vertically integrated framework that orchestrates a fleet of autonomous surface vessels as a distributed choreographic platform. Moving beyond the surface-pixel paradigm, these vessels use laminar nozzles and multi-zone lighting to extend their expressive range from the 2D water plane into the 3D volumetric domain. Our primary contribution is the Way of Water Studio, a browser-based, timeline-compositing authoring paradigm that treats the fleet as a DAW-like instrument for music-responsive choreography. The Studio encapsulates Sequential Convex Programming for trajectory generation and Model Predictive Control for disturbance rejection presented through a visual timeline, broadening access to high-performance aquatic robotics for non-programmer artists. Grounding the Studio is the full cyber-physical stack: a custom holonomic chassis, a state-estimation and control stack

tuned for the aquatic domain, and an LTE/MQTT fleet link with RTK-GPS time synchronization. We report on the system's validation across two distinct deployments: an 18-vessel Swan Lake interpretation at Lake Zurich and an 8-vessel Time Space Existence 2025 Venice Biennale demonstration at Forte Marghera, establishing a foundational reference for the design and deployment of fluidic robotic swarms.

## CCS Concepts

• Computer systems organization → Robotics; • Applied computing → Performing arts.

## Keywords

robotic art, swarm choreography, musical expression, aquatic performance, GUI-based creative tools

## ACM Reference Format:

Aswin Ramachandran, Christopher Golling, Sebastian Burmester, Noa Sendlhofer, Jan Kamm, Ruiheng Jiang, and Raffaello D'Andrea. 2026. Choreographing the Way of Water: A Computational Framework for Aquatic Robotic Art. In *Proceedings of International Conference on New Interfaces for Musical Expression (NIME '26)*. ACM, New York, NY, USA, 8 pages.



This work is licensed under a Creative Commons Attribution 4.0 International License.

NIME '26, June 23–26, 2026, London, UK

© 2026 Copyright held by the owner/author(s).

## 1 Introduction: Robotic Agency in Fluid Environments

The domain of robotic art sits at the intersection of electromechanical precision and aesthetic expression. Specifically, the challenge of music-responsive choreography, where physical agents articulate the structural and emotional content of sound, requires a level of synchronization that is difficult to achieve in stochastic environments. While aerial swarms have successfully colonized the sky as "pixels in the sky" [25], surface-based aquatic systems struggle significantly due to the rigorous engineering required to operate in a boundary layer. The water surface imposes nonlinear disturbances such as wave dynamics, currents, and drift, which actively resist precise choreography. We present a framework that resolves these constraints, transforming a fleet of autonomous surface vessels (ASVs) into a distributed choreographic platform of *dynamical pixels*. The vessels are treated as compositional primitives, with their position, heading, light state, and water-jet output all addressable from a single timeline track.

This paper establishes the cyber-physical architecture of Way of Water and, centrally, the authoring paradigm through which artists compose for it. To achieve the fidelity required for musically coherent motion, we analyze the mechanical engineering required to sustain precise station keeping, velocity tracking, and synchronized motion. We detail the adaptation of control strategies, Model Predictive Control (MPC) and Sequential Convex Programming (SCP), which are essential not only for navigation but also for the autonomous, collision-free choreography required to align motion with musical structure.

### 1.1 The Narrative Arc of Aquatic Robotics

The history of performance robotics has evolved from individual automatons to coordinated swarms. Aquatic environments offer a complementary medium to aerial shows, characterized by high endurance, intrinsic safety, and the evolution of collective behaviors in real-world settings [5].

The Way of Water addresses the unique challenges of this medium through a vertically integrated approach. The vessels are not repurposed survey drones but purpose-built aesthetic platforms. Every design decision, from the holonomic X-drive propulsion to the custom refractive optics of the nozzle, serves the dual purpose of robustness and expression, using the fluid medium as a visual element to give physical form to music.

### 1.2 From Script to Studio

The primary contribution of this work is the *Way of Water Studio*: a browser-based, timeline-compositing authoring paradigm for aquatic robotic swarms. It describes the evolution of the authoring workflow from a code-centric Python scripting framework toward a DAW-like interface that encapsulates the sophisticated mathematics of swarm coordination physics constraints, optimization solvers, and audio-structural mapping behind a visual timeline familiar to artists and musicians. This progression addresses the primary bottleneck in robotic performance: the friction between engineering rigidity and artistic fluidity. The Studio serves as an interface between the control parameters and the choreographer's musical intent, abstracting path-planning algorithms, collision constraints, and beat mapping behind a visual timeline. By documenting the complete system, from the hardware chassis to the web-based planning architecture, this work provides the necessary grounding for future research into

distributed software architectures and real-time, musically responsive swarm coordination.

## 2 Related Work: The Genealogy of Performative Engineering

The intersection of robotics and the arts has evolved from a niche curiosity into a rigorous discipline of performative engineering. We categorize the Way of Water within this lineage, tracing its roots from early cybernetics to modern distributed swarms.

### 2.1 Foundational Theory: From Feedback to Swarms

The genealogy of robotic art is rooted in the feedback loops of Cybernetic Serendipity (1968) and Gordon Pask's Conversation Theory, which posits that intelligence emerges from interaction [14]. The shift to collective behaviors was formalized by Reynolds' Boids [15], and carried into the artistic domain by practitioners such as Bill Vorn, whose *Hysterical Machines* and *Grace State Machines* treat collective robotic agitation as performative dramaturgy [24], and Ken Rinaldo, whose *Autopoiesis* (2000) demonstrated a gallery-scale ecology of robotic sculptures governed by local rules [16]. An aquatic lineage runs in parallel: Gilberto Esparza's *Plantas Nómadas* [6] couples autonomous robotic-plant hybrids to polluted waterways; John McCormack's *Holon* [11] stages an artificial aquatic ecosystem of reactive agents; and Audrey, Kelly, and St-Aubin's *Vessels* [1] deploys autonomous craft as performative sculpture. Waibel et al. popularized swarm behaviors as pixels in the sky for aerial shows [25], and Duarte et al. demonstrated real-world aquatic swarms under evolutionary constraints [5]. Way of Water extends this aquatic-art genealogy into high-performance, timeline-composed swarm choreography, while deliberately deferring Pask/Reynolds-style local-rule emergence as a future primitive class within the Studio.

### 2.2 Control Theoretic Frameworks

To translate musical intent into physical motion, the system must enforce strict temporal and spatial compliance against a chaotic fluid medium. Control theory serves as the bridge between the choreographer's idealized score and the physical reality of hydrodynamics. Without rigorous state estimation and trajectory optimization, the swarm's musicality and its ability to synchronize with a transient or maintain a formation would be lost to nonlinear dynamics and environmental disturbances.

Precise aquatic coordination begins with robust **state estimation**; we employ an Extended Kalman Filter (EKF) fused with vessel-specific parameter identification to capture the system's hydrodynamic reality. Building on this foundation, high-performance choreography requires advanced strategies to handle the high-dimensional state space of a swarm.

**Collision Avoidance:** Augugliaro et al. pioneered SCP for aerial choreography [2], utilizing it to resolve collision-free trajectories for agile quadrotors in 3D space. We adapt this formulation to the planar aquatic domain, reducing the problem dimensionality and removing the jerk constraints on acceleration necessitated by aerial dynamics. **Fleet Transition:** To transition the fleet, the Linear Sum Assignment Problem (LSAP) minimizes total travel distance [4], serving as a first layer of conflict resolution. **Tracking and Disturbance Rejection:** To proactively track, reject disturbance and drift, we employ MPC, optimizing control inputs over a finite horizon [13].



Figure 2: The Swan Lake performance, where autonomous vessels enact a balletic choreography, mirroring the elegance of the original score.

### 2.3 Authoring Ecosystems and Robotic Expression

The paradigm has shifted from choreography by code to choreography by design, evidenced by LLM-driven tools like SwarmGPT [21] and physical interfaces like Zooids [10]. Recent work expands this scope through distributed sensor networks [28, 29] and AI-augmented instruments [3].

This lineage traces back to Weinberg’s Robotic Musicianship, where robots listen and improvise [17, 20, 27]. Way of Water extends this agency from musical to spatial interaction: the vessels function as dynamical pixels [26] and interactive paintbrushes [19], strictly governed by physics to visualize Newton’s laws as aesthetic output. Recent MIR research using deep models such as Demucs [18] and CREPE [8] offers a forward path for driving procedural behavior from high-level musical features; the current Studio relies on the classical, offline-deterministic pipeline described in 5.5 and treats deeper MIR integration as future work.

Ultimately, Way of Water synthesizes robotics, physics-based modeling, and control theory into a unified framework for timeline-authored music-responsive choreography. While mathematical rigor ensures safety and coordination, it is the Way of Water Studio that serves as the critical abstraction layer, empowering artists to compose for aquatic swarms through a familiar visual timeline rather than code.

## 3 The Physical Platform: Mechanical and Hydrodynamic Engineering

The Way of Water vessel design balances several competing constraints: buoyancy, stability, agility, endurance, and visual impact. This section deconstructs the hardware specifications to provide a reference for the system’s capabilities.



Figure 3: Overview of the hardware design, showing the integration of the lifebuoy chassis, propulsion system, and multi-layered lighting.

### 3.1 Holonomic Propulsion Architecture

Unlike traditional rudder-steered boats, the Way of Water fleet utilizes a holonomic X-configuration with four brushless DC motors. This arrangement allows the vessel to translate in any direction ( $x, y$ ) while simultaneously rotating ( $r$ ) or holding a fixed heading to align the laminar fountain with artistic intent. Each thruster delivers 20N of thrust, allowing an operational speed of up to 0.8m/s. For safety, the thrusters are fully shrouded to enable operation in public waters.

### 3.2 Optical Payloads: Laminar and Peripheral

The vessel’s visual signature relies on a custom laminar nozzle that generates a turbulence-free, internally illuminated water column acting as a volumetric pixel up to 6 meters in height. Complementing this vertical element, the chassis integrates a three-zone peripheral lighting architecture (Figure 3) using individually addressable RGB LED strips an *Internal Ring* that highlights the chassis structure, a *Waterline Ring* that disperses streaks

across the surface, and a *Submersible Ring* that provides a diffuse underwater glow allowing independent control over the vessel's presence and its interaction with the fluid medium.

### 3.3 Electronics and Power Architecture

Operational endurance is a critical metric for public art. The fleet is powered by a custom Li-ion battery pack, enabling over 10 hours of continuous operation. The 25.2V architecture enables powerful thrusters, while a BMS provides active balancing and telemetry. The central electronics enclosure (E-Box), machined from aluminum, is sealed to IP67 and acts as a unified heatsink for the ESCs and battery. It houses a Raspberry Pi 5, providing significant onboard compute power for the decentralized execution of sophisticated solvers, such as OSQP for MPC, directly on the agent. The propulsion thrusters are fully submersible (IP68), and the external connectors are IP67-rated push-pull bayonet types; the platform has been operated for tens of hours per show across both freshwater (Lake Zurich) and brackish (Venice lagoon) deployments without sealing failures.

### 3.4 Fleet Communication and Time Synchronization

Each vessel carries an onboard LTE modem that streams telemetry and receives RTK-GPS corrections through an MQTT broker; an additional WiFi channel is used at the shore for high-bandwidth, low-cost communication during setup and mission upload. Fleet-wide show-clock synchronization is derived from the RTK-GPS pulse-per-second, providing sub-millisecond temporal alignment between vessels without requiring a dedicated master clock. The MQTT topic structure separates concerns for low-rate operator commands (start/abort/return-to-shore). Typical end-to-end planner-to-vessel latency over LTE is in the 80–250 ms range; show choreography is therefore not closed-loop over the radio link. Show data is pre-uploaded over WiFi at the shore, then executed under the shared GPS timebase: a vessel can tolerate a complete LTE blackout for several minutes while still hitting its choreography to within sub-RTK precision, and the planner detects stale-clock or off-pattern vessels by monitoring the deviation between expected and reported state in telemetry.

## 4 The Control Theoretic Framework: Taming the Stochastic Environment

The operational stability of the Way of Water relies on the suppression of environmental disturbances. To achieve smooth trajectories in the presence of wind and current, the control stack utilizes established state estimation and predictive control techniques adapted for the aquatic domain.

### 4.1 Sensor Fusion and Heading Estimation

To enable precise choreography, the system must robustly estimate each vessel's full 3-DOF state. We employ a Multi-Rate EKF with a state vector encompassing position, velocity, and orientation:  $\mathbf{x} = [x, y, \psi, u, v, r]^T$ . [7] **RTK-GPS:** Provides position  $(x, y)$  and velocity  $(v_x, v_y)$  at 10Hz with centimeter-level precision. **IMU:** Provides angular rates  $(\omega)$  and linear accelerations  $(a)$  at 100Hz, filling the gaps between GPS updates. Crucially, the system determines heading  $(\psi)$  through an observability-based estimator. By correlating the known control inputs with the observed inertial displacement from RTK-GPS, the EKF converges

on the true orientation. This sensor fusion strategy ensures accurate, drift-free tracking of the vessel's pose in an open-water environment.

### 4.2 Model Predictive Control

While reactive PID controllers are sufficient for station-keeping, they lack the anticipatory behavior required for fluid, organic motion in a high-inertia medium. Pure feedforward control, conversely, is prone to drift from wind and current.

We employ MPC to solve a finite-horizon optimization problem at 10Hz. This allows the vessel to anticipate turns, lean into maneuvers to counter inertia, and exhibit inherent recovery behaviors if a wave knocks it off course.

**Objective Function:** The cost function  $J$  minimizes deviation from the choreographed trajectory. The state-tracking weight  $\mathbf{Q}$  is tuned to prioritize heading precision (keeping laminar nozzles strictly parallel for visual coherence) while allowing mild positional compliance to absorb wave energy; the input-tracking weight  $\mathbf{R}$  penalizes deviation from the feedforward command to suppress jitter.

$$J = \sum_{k=0}^{N-1} [(\mathbf{x}_k - \mathbf{x}_{r,k})^T \mathbf{Q} (\mathbf{x}_k - \mathbf{x}_{r,k}) + (\mathbf{u}_k - \mathbf{u}_{r,k})^T \mathbf{R} (\mathbf{u}_k - \mathbf{u}_{r,k})] \quad (1)$$

Table 1 summarises the symbols. OSQP solves the resulting QP in real time on the Raspberry Pi. In the field, this controller delivers per-vessel RMSE of **0.038 m** in cross-track position  $(x, y)$  and **1.4°** in heading  $\psi$  (Figure 4). These figures are stable across fleets of 8 (Venice Biennale) and 18 (Lake Zurich) vessels under sustained wind and wave conditions, the operational envelope tested during the deployments. Equivalent post-hoc telemetry was not retained for Forte Marghera, so we report cross-deployment robustness only at the qualitative level. Beyond the tested envelope, performance degrades primarily through GPS-correction integrity rather than control bandwidth.

**Table 1: MPC formulation symbols.**

Symbol	Meaning
$N$	prediction horizon (steps)
$\mathbf{x}_k, \mathbf{u}_k$	predicted state and control input at step $k$
$\mathbf{x}_{r,k}, \mathbf{u}_{r,k}$	reference trajectory and feedforward
$\mathbf{Q}, \mathbf{R}$	state and input weighting (heading-prioritised)

### 4.3 Collision Avoidance via SCP

Generating safe trajectories for a fleet of 24 robots is a non-convex optimization problem. We utilize SCP to solve this offline.

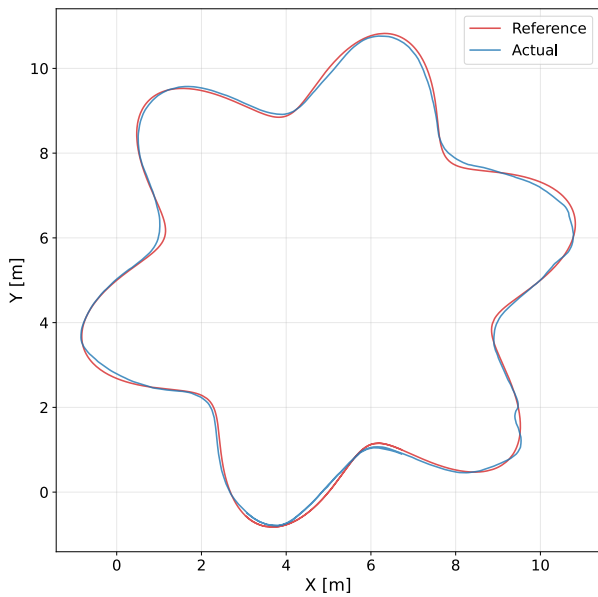
The pairwise collision avoidance constraint between robot  $i$  and robot  $j$  is defined as:

$$\|\mathbf{p}_i[k] - \mathbf{p}_j[k]\|_2 \geq R_{safe} \quad (2)$$

where  $\mathbf{p}[k]$  is the position at time step  $k$  and  $R_{safe}$  is the minimum safety radius (hull + buffer). Since this constraint is non-convex, SCP linearizes it around a candidate trajectory  $\bar{\mathbf{p}}$  from the previous iteration. The resulting convex constraint is:

$$\|\bar{\mathbf{d}}_{ij}[k]\|_2 + \frac{\bar{\mathbf{d}}_{ij}[k]^T}{\|\bar{\mathbf{d}}_{ij}[k]\|_2} (\mathbf{d}_{ij}[k] - \bar{\mathbf{d}}_{ij}[k]) \geq R_{safe} \quad (3)$$

where  $\mathbf{d}_{ij} = \mathbf{p}_i - \mathbf{p}_j$  is the relative position vector. This formulation defines a separating hyperplane between agents. The solver iterates this linear approximation until the trajectory converges



**Figure 4: MPC tracking performance under wave and current disturbance. The controller maintains high-precision station keeping, achieving a RMSE of 0.038 m in position  $(x, y)$  and  $1.4^\circ$  in heading  $(\psi)$ .**

to a solution that satisfies the original non-convex constraints, guaranteeing mathematically collision-free paths prior to deployment.

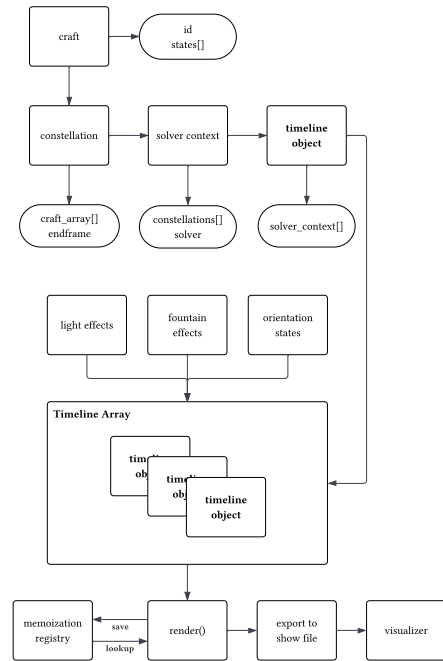
## 5 The Way of Water Studio

The Way of Water Studio is the primary contribution of this work: a browser-based, DAW-like authoring environment that treats a fleet of ASVs as a single timeline-composed instrument for music-responsive choreography (Figure 6). It transitions the workflow from code-based scripting to visual, timeline-based compositing, abstracting the rigorous control theory of Section 4 behind an interface designed for rapid artistic iteration. Instead of manipulating individual waypoints, choreographers define trajectories parametrically; these are pre-generated and validated against the fleet’s dynamic constraints to ensure physical feasibility before upload.

The Studio is a *timeline-authored choreography platform with audio-reactive payload modulation*, validated through two real-world case studies. It is not a real-time improvisatory instrument, nor an autonomous music-listening agent, nor a system in which motion itself is generated directly from audio. Vessel motion is composed by the choreographer against a pre-analyzed audio reference; the audio pipeline (Section 5.5) drives lights and water-jet payloads. Closing the gap between music-driven *payload* and music-driven *motion*, together with a more formal evaluation program, is the natural next research direction; the contribution here is the authoring paradigm and the cyber-physical stack that make such future work tractable.

### 5.1 Design Requirements: The Swan Lake Prototype

Prior to the development of the graphical Studio, the choreography for the Swan Lake performance was authored using a code-first approach. This framework, illustrated in Figure 5, served



**Figure 5: Architecture of the code-based planner used for Swan Lake. The system decouples the artistic Script from the Solver Context, validating physical feasibility against the dynamics model.**

as the foundational prototype for the current system. It relied on a scripted pipeline where choreographers defined trajectories using Python primitives. While rigorous and flexible, this approach lacked real-time visual feedback; it required specific programming knowledge and a slow compilation-simulation loop to validate even minor adjustments.

Insights from the Swan Lake deployment defined the architectural requirements for the new Studio, resulting in a system composed of three primary interaction layers: a timeline-based compositing interface, parametric primitives for behavior definition, and a real-time 3D viewport for rapid artistic iteration.

### 5.2 The Timeline Interface: Timeline-Based Compositing

Mirroring the layout of Digital Audio Workstations (DAWs), the Studio adopts a timeline-based compositing paradigm. This approach moves away from rigid waypoint lists, treating fleet behaviors as high-level parametric clips that can be dragged, resized, and layered non-destructively. As illustrated in Figure 6, the interface organizes these behaviors into synchronized sequencing layers:

**Motion Layer:** Defines the geometric topology of the fleet. As seen in the Motion Controls panel (Figure 6, right), users manipulate meta-parameters such as shape topology (e.g., circle, line), scale, and Cartesian offsets rather than individual robot coordinates. Transition effects transform the shapes from one to another, which can be created manually or using the LSAP-based solver.

**Payload Layer:** A track-agnostic sequencing layer that controls the vessel’s peripherals. This includes the Light Track for LED arrays (color, intensity, effects like Rainbow Chase) and the Jet Track for actuating the laminar nozzles.



**Figure 6: Parametric Motion Editor.** Double-clicking the timeline (bottom) creates an effect, which can be dragged and resized. There are five effect types are Motion, Transition, Jet, Light and Intensity, which control movements and payloads. Upon selecting an effect, its parameters can be tweaked (floating panel). The 3D viewport (center) previews the choreography.

Audio Reference: A visual waveform serves as the temporal anchor (Figure 6, bottom), allowing choreographers to align fleet transitions and payload effects with musical transients with visual precision.

### 5.3 Parametric Primitives

To bridge the gap between artistic intent and control constraints, the system relies on parametric primitives. Instead of manually positioning individual agents on a geometric shape, users define the meta-state of a formation.

For example, a circle primitive is not defined by 24 coordinate sets but by its global position ( $x, y$ ), rotation ( $\theta$ ), scale, and phase offset. The global position corresponds to the center of each primitive’s bounding box. The phase offset shifts all crafts along the parameterized vector graphic path that defines the shape. By choosing this approach, where primitives are vector graphics that are normalized based on their bounding boxes, it allows for quick changes to shapes at a later stage, while the crafts retain their approximate location within the choreography.

### 5.4 Transitions

Transitions are performed between the primitive shapes. The choreographer can manually match start and end positions or use the LSAP solver.

The LSAP solver uses squared Euclidean distance as weights. For this approach, it has been shown that the minimum distance between trajectories is guaranteed to be  $d_{min} = \sqrt{2}\delta/2$  where  $\delta$  is the minimum distance between any two crafts across the start

and end positions [9, 22]. For our crafts with a diameter of 50cm, this implies that a gap of  $2 * 50\text{cm} / \sqrt{2} - 50\text{cm} \approx 21\text{cm}$  is enough to guarantee collision-free paths. Note that this approach assumes that the crafts are identical and that each craft can assume any position in the target state. Manual adjustments of the LSAP results allow for artistic expression and desired sub-optimality.

Adjusting the transition smoothness parameters directly alters the resulting velocity profiles, enabling the choreographer to align movement dynamics with expressive qualities such as calmness or intensity. Note that the SCP solver from the Swan Lake planner could be integrated to allow for more intricate transitions that are more complex than straight lines, and allow for fixed end position assignments of crafts.

### 5.5 Audio Pipeline

To bridge the gap between musical structure and robotic action, the framework integrates an *offline* audio analysis pipeline powered by Librosa [12] and Scipy [23]. In the current Studio, choreographies are authored against a pre-analyzed audio track and rendered into a show file before deployment; live audio-driven operation is explicitly out of scope for this paper. We are also explicit about the division of labor between manual and audio-driven authoring: motion is composed in the timeline by the choreographer (with the audio waveform as the temporal reference), while the audio features described below primarily drive payload modulation of light intensity, jet height, and event triggers. Closing this gap so that motion itself becomes audio-reactive is a deliberate next step, not a hidden capability of the

current system. The frontend Studio interface focuses on timeline visualization, while the backend generates the complete choreography state file by combining the processed audio signal with the effect data from the timeline.

**Feature Extraction:** Upon upload, each audio file is loaded at its native sample rate to preserve temporal fidelity. A Short-Time Fourier Transform (STFT) with a Hann window of 2048 samples and a hop size of 512 samples serves as the basis for all spectral analysis. The pipeline extracts the following feature sets: **Rhythm features:** Global tempo estimation and beat tracking, onset detection using the onset strength envelope with subsequent peak picking, and onset strengths normalized to  $[0, 1]$ . **Energy features:** Root-mean-square (RMS) energy computed per frame, with global energy peaks. **Spectral features:** Spectral centroid (timbral brightness), spectral rolloff at the 85th percentile, and zero-crossing rate for texture characterization.

**Frequency Band Decomposition:** For audio-reactive effects, the system computes band-limited energy envelopes on demand. The STFT magnitude spectrogram is partitioned into seven perceptually motivated frequency bands Sub Bass (20–60 Hz), Bass (60–250 Hz), Low Mids (250–500 Hz), Midrange (500–2000 Hz), High Mids (2–4 kHz), Presence (4–6 kHz), and Air (6–20 kHz) by summing magnitude bins within each range and normalizing to unit peak. Additionally, a custom frequency band can be specified. Per-band peak detection with configurable height, distance, and prominence thresholds allows designers to trigger choreographic events from specific spectral regions. These per-band results are cached independently, so repeated queries incur no recomputation.

**Intensity Modulation:** The extracted audio features drive a multi-stage intensity modulation chain that transforms base color values or water jet intensities into musically responsive output.

**Jumping and Flash Patterns:** For spatially distributed effects, detected beats or frequency-band peaks serve as trigger events. A *jumping* pattern sequentially or randomly activates individual crafts at each trigger, with per-craft intensity decaying between activations. The inverse *flash* pattern maintains all crafts at peak intensity and introduces an anticipation dim before each trigger, creating a rhythmic breathing effect synchronized to the music’s temporal structure.

## 5.6 Authoring Workflow and Deployments

The workflow is designed to minimize cognitive load while maximizing control authority. Users begin by initializing a project and uploading the master audio track; the physical fleet configuration (number of active USVs, show area bounds) is then set so the planner can enforce the specific hardware and geometric constraints. By abstracting the underlying robotic control layers, the system lets artists focus on the expression of music through light, motion, and fluid dynamics.

**Comparative authoring assessment:** We conducted an informal comparative assessment across two production cycles. The Swan Lake performance (18 vessels across an ~60 m show area on Lake Zurich) was authored entirely through the code-first pipeline, which restricted authorship to two engineers with Python fluency and required 4–8 hours per scene with a 5-minute script-compile-simulate iteration loop. The Lake Zurich demonstration with 18 crafts and the Venice Biennale demonstration (8 vessels in the smaller lagoon footprint at Forte Marghera) were produced entirely within the Studio; authoring time dropped to 90–120 minutes per song, and the edit-to-preview cycle to under 10 seconds.

Most significantly, the Studio eliminated the requirement for programming expertise, expanding the authoring team. These observations stem from a single production team and are preliminary; a formal usability evaluation with external choreographers and musicians remains important future work.

**Failure modes and aesthetic resilience:** Across both deployments, no catastrophic failure has occurred to date, but two incidents are worth documenting. In one performance, a single vessel lost its show-clock synchronization and continued executing its full program out of phase; rather than reading as a fault, this manifested as a formation asymmetry that audience members interpreted as an intentional perturbation, an unintended but illustrative form of choreographic noise. In another case, a single vessel suffered a mechanical fault mid-show; its lights were remotely disabled, and it drifted out of the active formation unnoticed.

**Informal audience and collaborator reflections:** Collaborator responses during the Venice Biennale run most often framed the work in terms of the medium itself “*water dancing to music, with lights and motion.*” A recurring informal observation was a collective hush in the spectator field: children who had been playing loudly prior to a show would fall silent for the duration of the choreography and resume only afterward. We report these reactions as anecdotal rather than evaluative; systematic audience and artist studies, particularly with choreographers and composers external to the development team, are a future step.

## 6 Conclusion

We have presented the Way of Water: a vertically integrated cyber-physical framework for music-responsive choreography on open water with an authoring paradigm, the Way of Water Studio, that treats a fleet of ASVs as a timeline-composed instrument rather than as a programming problem. The Studio encapsulates SCP for collision avoidance, the LSAP for fleet transitions, and MPC for disturbance rejection behind a DAW-like timeline, allowing choreographers with musical rather than engineering backgrounds to compose for a distributed aquatic swarm.

Two deployments demonstrate versatility: the 18-vessel Swan Lake interpretation on Lake Zurich confirmed precision in station-keeping and laminar stability, while the 8-vessel Venice Biennale demonstration at Forte Marghera proved robustness in open water under agile, high-velocity maneuvers. The project sits at the intersection of several communities: engineers, choreographers, musicians, and public audiences, whose convergence is itself part of the contribution. The Biennale spectator field foregrounded a public reading of robotic choreography as an ecological-aesthetic spectacle rather than a technical demonstration. We see the Studio as a step toward making such interdisciplinary, audience-facing work reproducible by other teams.

Several directions remain open. **Scalability:** the current centralized solver practically limits the fleet to ~24 vessels; larger swarms will require distributed middleware and hierarchical planning. **Site and environment:** the installation relies on fixed audio infrastructure and is sensitive to weather and seasonal visibility; distributed audio streaming and more robust perception are needed to expand the operational envelope. **Choreographic primitives:** current primitives are top-down parametric formations; a natural extension is a class of Reynolds/Pask-style local-rule primitives that would let the Studio host bottom-up emergent behaviors alongside scripted ones. This keeps the Studio’s compositional grammar uniform while opening an emergent design

space. *Live operation*: the current audio pipeline is offline; extending it to real-time listening and improvisation would bring the platform closer to the lineage of robotic musicianship. *Evaluation*: formal usability studies with external choreographers and audience-experience research remain the most important immediate next step.

## Media Links

- Skyfall: <https://youtu.be/gprHw5OZqsQ>
- The Most Beautiful Boy: <https://youtu.be/VHYqJ1hxct0>
- Forte Marghera: <https://youtu.be/G4cM6xbG7PA>

## Ethical Standards

This work was supported by ETH Zurich. The authors declare no competing financial interests or conflicts of interest regarding the publication of this paper. Motivated by the need for sustainable, reusable alternatives to pyrotechnics, this research prioritizes environmental stewardship. To minimize ecological impact, the Way of Water fleet utilizes zero-emission electric propulsion and fully shrouded impellers, preventing chemical contamination and physical injury to aquatic wildlife. Operational safety in public venues (Lake Zurich, Venice Lagoon) was guaranteed through a multi-layered control architecture, where all autonomous trajectories were validated via Software-in-the-Loop (SITL) simulation prior to physical execution. The evaluation of the authoring workflow relied solely on the internal production team's professional practice; as no external human subjects were recruited, informed consent procedures were not applicable, and formal Institutional Review Board (IRB) approval was not required.

## Acknowledgments

We extend our gratitude to our technical staff Daniel Wagner, and Matthias Müller, as well as the Focus Project 2025 participants Jonathan Nowack, Fjodor Lundgren, Janis Weyrich, and Daniel Gelfenbaum for their hard work and continued support during the development of the fleet.

## References

- [1] Sofian Audry, Stephen Kelly, and Samuel St-Aubin. 2019. Vessels. Robotic art installation. <https://sofianaudry.com/works/vessels/> A performative installation of autonomous aquatic vessels exhibited in artistic venues..
- [2] Federico Augugliaro, Angela P Schoellig, and Raffaello D'Andrea. 2012. Generation of collision-free trajectories for a quadcopter fleet: A sequential convex programming approach. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, IEEE, New York, NY, USA, 1917–1922.
- [3] Lancelot Blanchard, Perry Naseck, Eran Egozy, and Joseph A Paradiso. 2024. *Developing Symbiotic Virtuosity: AI-Augmented Musical Instruments and Their Use in Live Music Performances*. Technical Report. MIT Media Lab, Cambridge, MA, USA. <https://doi.org/10.21428/e4baedd9.69c11de7> An MIT Exploration of Generative AI.
- [4] Rainer Burkard, Mauro Dell'Amico, and Silvano Martello. 2009. *Assignment Problems*. SIAM, Philadelphia, PA.
- [5] Miguel Duarte, Vasco Costa, Jorge Gomes, Tiago Rodrigues, Fernando Silva, Sancho M Oliveira, and Anders Lyhne Christensen. 2016. Evolution of collective behaviors for a real swarm of aquatic surface robots. *PLoS One* 11, 3 (2016), e0151834.
- [6] Gilberto Esparza. 2010. Plantas Nómadas. Artwork installation. <https://gibertoesparza.net/portfolio/plantas-nomadas/> Autonomous hybrid robotic-plant organisms that migrate along polluted waterways. VIDA 13.2 Art and Artificial Life International Awards, Fundación Telefónica..
- [7] Simon J Julier and Jeffrey K Uhlmann. 1997. New extension of the Kalman filter to nonlinear systems. In *Signal Processing, Sensor Fusion, and Target Recognition VI*, Vol. 3068. SPIE, SPIE, Bellingham, WA, USA, 182–193.
- [8] Jong Wook Kim, Justin Salamon, Peter Li, and Juan Pablo Bello. 2018. Crepe: A convolutional representation for pitch estimation. In *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, IEEE, New York, NY, USA, 161–165.
- [9] Vit Kratky, Robert Penicka, Jiri Horyna, Petr Stibinger, Tomas Baca, Matej Petrlik, Petr Stepan, and Martin Saska. 2025. CAT-ORA: Collision-Aware Time-Optimal Formation Reshaping for Efficient Robot Coordination in 3D Environments.
- [10] Mathieu Le Goc, Lawrence H Kim, Ali Parsaei, Jean-Daniel Fekete, Pierre Dragicevic, and Sean Follmer. 2016. Zooids: Building blocks for swarm user interfaces. In *Proceedings of the 29th Annual Symposium on User Interface Software and Technology*. ACM, New York, NY, USA, 97–109.
- [11] Jon McCormack, Toby Gifford, Patrick Hutchings, Maria Teresa Llano, Matthew Yee-King, and Mark d'Inverno. 2019. Holon: A Framework for Simulation and Exhibition of Artificial Biota. In *Proceedings of the 10th International Conference on Computational Creativity (ICCC)*. Association for Computational Creativity, Charlotte, NC, USA, 8 pages. *Holon* (2018) is an interactive aquatic robotic ecosystem..
- [12] Brian McFee, Colin Raffel, Dawen Liang, Daniel PW Ellis, Matt McVicar, Eric Battenberg, and Oriol Nieto. 2015. librosa: Audio and music signal analysis in python. In *Proceedings of the 14th Python in Science Conference (SciPy 2015)*, Vol. 8. SciPy, Austin, TX, USA, 18–25.
- [13] Daniel Morgan, Soon-Jo Chung, and Fred Y Hadaegh. 2014. Model Predictive Control of Swarms of Spacecraft Using Sequential Convex Programming. *Journal of Guidance, Control, and Dynamics* 37, 6 (2014), 1725–1740.
- [14] Gordon Pask. 1969. The Architectural Relevance of Cybernetics. *Architectural Design* 39, 9 (1969), 494–496.
- [15] Craig W Reynolds. 1987. Flocks, Herds, and Schools: A Distributed Behavioral Model. In *Proceedings of the 14th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '87)*. ACM, New York, NY, USA, 25–34.
- [16] Ken Rinaldo. 2002. Autopoiesis. *Leonardo* 35, 4 (2002), 395–396. Artist statement for the robotic sculpture *Autopoiesis* (2000)..
- [17] Amit Rogel, Qiaoyu Yang, Jack Hayley, and Gil Weinberg. 2025. Do Re Mi Fa So Pass the Tool: Using Melodic Prediction to Improve Human-Robot Fluency. In *2025 34th IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*. IEEE, IEEE, New York, NY, USA, 199–206.
- [18] Simon Rouard, Francisco Massa, and Alexandre Défossez. 2023. Hybrid transformers for music source separation. In *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, IEEE, New York, NY, USA, 1–5.
- [19] Maria Santos and Magnus Egerstedt. 2020. Interactive Multi-Robot Painting Through Colored Motion Trails. *Frontiers in Robotics and AI* 7 (2020), 580415.
- [20] Richard Savery, Amit Rogel, Lisa Zahray, and Gil Weinberg. 2023. How Happy Should I be? Leveraging Neuroticism and Extraversion for Music-Driven Emotional Interaction in Robotics. In *Sound and Robotics*. Chapman and Hall/CRC, Boca Raton, FL, USA, 199–218.
- [21] Martin Schuck, Dinushka Orrin Dahanagammaarachchi, Ben Sprenger, Vedant Vyas, Siqi Zhou, and Angela P Schoellig. 2025. SwarmGPT: Combining Large Language Models with Safe Motion Planning for Drone Swarm Choreography. *IEEE Robotics and Automation Letters* 10, 11 (2025), 12237–12244.
- [22] Matthew Turpin, Nathan Michael, and Vijay Kumar. 2014. Capt: Concurrent assignment and planning of trajectories for multiple robots. *The International Journal of Robotics Research* 33, 1 (2014), 98–112. <https://doi.org/10.1177/0278364913515307>
- [23] Pauli Virtanen, Ralf Gommers, Travis E Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, Stéfan J van der Walt, Matthew Brett, Joshua Wilson, K Jarrod Millman, Nikolay Mayorov, Andrew R J Nelson, Eric Jones, Robert Kern, Eric Larson, C J Carey, İlhan Polat, Yu Feng, Eric W Moore, Jake VanderPlas, Denis Laxalde, Josef Perktold, Robert Cimrman, Ian Henriksen, E A Quintero, Charles R Harris, Anne M Archibald, António H Ribeiro, Fabian Pedregosa, Paul van Mulbregt, and SciPy 1.0 Contributors. 2020. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods* 17 (2020), 261–272. <https://doi.org/10.1038/s41592-019-0686-2>
- [24] Bill Vorn. 2016. Artistic Approaches in Robotics: A Personal Journey. In *Robotic Art: Exploring an Unlikely Symbiosis*, Damith Herath, Christian Kroos, and Stelarc (Eds.). Springer, Cham, Switzerland, 243–262. Discusses works including *Hysterical Machines* and *Grace State Machines*..
- [25] Markus Waibel, Bill Keays, and Federico Augugliaro. 2017. Drone shows: Creative potential and best practices. Verity Studios White Paper. <https://veritystudios.com>
- [26] Wei Wang, Luis A Mateos, Shuguang Park, Pietro Galdems, Dylan Kellogg, Fabio Duarte, Carlo Ratti, and Daniela Rus. 2020. Roboat II: A novel autonomous surface vessel for urban environments. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, IEEE, New York, NY, USA, 1740–1747.
- [27] Gil Weinberg and Scott Driscoll. 2006. Toward robotic musicianship. *Computer Music Journal* 30, 4 (2006), 28–45.
- [28] Irma Wicaksono, Lancelot Blanchard, Sam Chin, Cristian Colon, and Joseph A Paradiso. 2024. KnitworkVR: Dual-reality Experience through Distributed Sensor-Actuator Networks in the Living Knitwork Pavilion. In *SIGGRAPH Asia 2024 Emerging Technologies*. ACM, New York, NY, USA, 2 pages.
- [29] Irma Wicaksono, Don Derek Haddad, and Joseph A Paradiso. 2022. Tapis Magique: Machine-knitted Electronic Textile Carpet for Interactive Choreomusical Performance and Immersive Environments. In *Proceedings of the 2022 ACM Conference on Creativity and Cognition*. ACM, New York, NY, USA, 326–339. <https://doi.org/10.1145/3527927.3532812>