ANALYTICAL PROOF OF CONCEPT / REVISION 1

DriftWing Technical Evidence Note

What the software demonstrates, what it does not, and the next physical proof

EVIDENCE CLASS / Modeled + implementation consistency
SYSTEM MATURITY / TRL 2 / no physical prototype
HISTORICAL STUDY / 21 causal model cases
NEXT GATE / Instrumented bench + tank demonstration

Technical Evidence Note

Version: 1.0 **Prepared:** July 25, 2026 **Project status:** Analytical proof of concept (TRL 2); pre-incorporation and pre-revenue **Evidence status:** Software-based. No DriftWing hardware, tank test, pressure test, controlled-water trial, or ocean deployment has been completed.

Purpose and evidence boundary

OceanStride is developing DriftWing as a recoverable, depth-selective profiling system. The operating hypothesis is that a vehicle can use vertical current structure for most horizontal transport, apply wing-assisted motion during selected depth transitions, surface for communication and re-planning, and deliver repeated water-column observations with less dependence on continuous propulsion.

The work completed to date tests that hypothesis at the mission-model and software-implementation levels. It does **not** establish vehicle hydrodynamics, endurance, sensor accuracy, reliability, recoverability, cost advantage, or field performance. The technically accurate conclusion is:

> The current software demonstrates that depth-selective, forecast-aware routing can be represented, exercised against NOAA model-data inputs, and reproduced consistently across independent implementations. A limited historical study contains successful modeled cases, but also shows a material robustness gap. Physical validation is the next proof.

System and model architecture

The present simulator is a mission-level kinematic model, not a six-degree-of-freedom hydrodynamic model. At each step, horizontal ground velocity is the sum of modeled ocean current, an estimated wave-drift sensitivity term, and any wing-assisted glide velocity. Depth is positive downward.

In the default month-scale `current_riding` mode, the planner may command one finite wing-assisted transition toward a selected current-bearing depth. When the target depth is reached, the model changes the action to a depth hold and removes the wing command. The model therefore assigns no wing-generated horizontal velocity while holding constant depth. A separate `sawtooth_experimental` mode exists for an intentionally propulsion-heavy comparison and is not the basis of the current-riding headline.

The reference software stack comprises:

Layer	Current implementation	Evidence status
Environmental data	MATLAB ingestion and normalization of NOAA current and wave products	Implemented in software
Mission model	Fixed-step kinematics, geodetic propagation, depth control, simplified energy accounting	Implemented; not physically calibrated
Planner	Passive, greedy-depth, causal look-ahead, oracle look-ahead, and issued-forecast modes	Implemented in software
Native implementation	Dependency-free C11 library and command-line interface	Built and deterministically tested
Batch implementation	C++ host path, CUDA kernel source, and CPU fallback	CPU fallback built and tested; GPU path not compiled locally
Physical system	Pressure body, buoyancy engine, wings, sensors, communications, recovery hardware	Proposed; not yet built

2013 causal historical-model study

Study design

The main month-scale experiment uses date-matched historical 2013 NOAA products:

- NOAA CoastWatch `hycom_gom310D` daily, depth-resolved Gulf of Mexico currents; and
- NOAA AOML `WaveWatch_2013` three-hourly bulk wind-sea and swell fields.

The cached environmental window is May 29 through August 4, 2013. The baseline matrix contains three offshore Gulf routes, seven release dates at five-day intervals from June 1 through July 1, and four policies: passive, greedy depth, causal look-ahead, and oracle look-ahead. This produces 84 baseline runs, including **21 causal cases**.

The causal policy does not see the archived future during planning. At each re-plan, it persists the issue-time environmental field through its look-ahead horizon, then executes the first 12 hours and re-plans. This is a causal persistence baseline, not a reconstruction of saved operational forecast cycles. The oracle policy sees the archived future and is an explicitly labeled upper bound.

Core modeled settings are:

Setting	Modeled value
Maximum mission duration	30 days
State integration step	15 minutes
Re-planning interval	12 hours
Planning horizon	36 hours
Rollout step	3 hours
Beam width	8
Planner depth choices	10, 50, 100, and 125 m
Arrival condition	Entry into a 5 km target circle
Nominal battery	2,500 Wh
Modeled reserve	250 Wh (10% of nominal capacity)
Modeled hotel load	2.0 W
Modeled wing-control load	0.5 W while active
Modeled pump cost	0.003 Wh per vertical metre
Modeled maximum vertical speed	0.05 m/s
Assumed glide ratio	3.0

These are simulation assumptions, not measured DriftWing specifications.

Matrix result

Policy	Modeled arrivals	Interpretation
Passive	0 / 21	Fixed-depth baseline
Greedy depth	1 / 21	Selects the locally helpful depth without look-ahead
Look-ahead causal	**3 / 21**	Persists the issue-time field; the principal defensible feasibility result
Look-ahead oracle	3 / 21	Uses archived future and must not be described as forecast performance

The causal result is a 14.3% modeled arrival rate within this defined matrix. All three causal arrivals occurred on the `central_to_east` route; none occurred on the other two routes. This is promising as a feasibility signal and equally important as evidence that the present planner is not robust across routes and release dates.

Exact successful causal cases

All three successful cases began 369.815061609696 km from the modeled target.

Scenario	Start (UTC)	Elapsed days	Final distance (km)	Energy used (Wh)	Final energy (Wh)	Wing-active time (h)	Wing duty (%)
`central_to_east_20130601`	2013-06-01	26.09375	4.53257272376007	1,255.11	1,244.89	1.5	0.239520958083832
`central_to_east_20130611`	2013-06-11	29.3020833333333	4.88951982388774	1,408.85	1,091.15	0.5	0.0710984713828653
`central_to_east_20130701`	2013-07-01	8.54166666666667	4.79277318050254	410.719999999999	2,089.28	0	0

The simulator's displacement accounting for those cases is:

Scenario	Vertical travel (m)	Current displacement (km)	Wave displacement (km)	Modeled glide displacement (km)
`central_to_east_20130601`	620	370.424788267049	1.33490078612591	0.323074765558732
`central_to_east_20130611`	700	361.666762424538	1.09510057599569	0.27
`central_to_east_20130701`	240	365.010354099926	0.62691619501046	0

The June 1 case is the recommended showcase because it visibly demonstrates the intended modeled operating logic: the trajectory entered the 5 km target region after 26.09375 modeled days, with 1.5 modeled wing-active hours and 0.3231 km of modeled glide displacement. It must always be accompanied by the broader **3 / 21** result and the label **“historical NOAA model-data backtest; not field data.”**

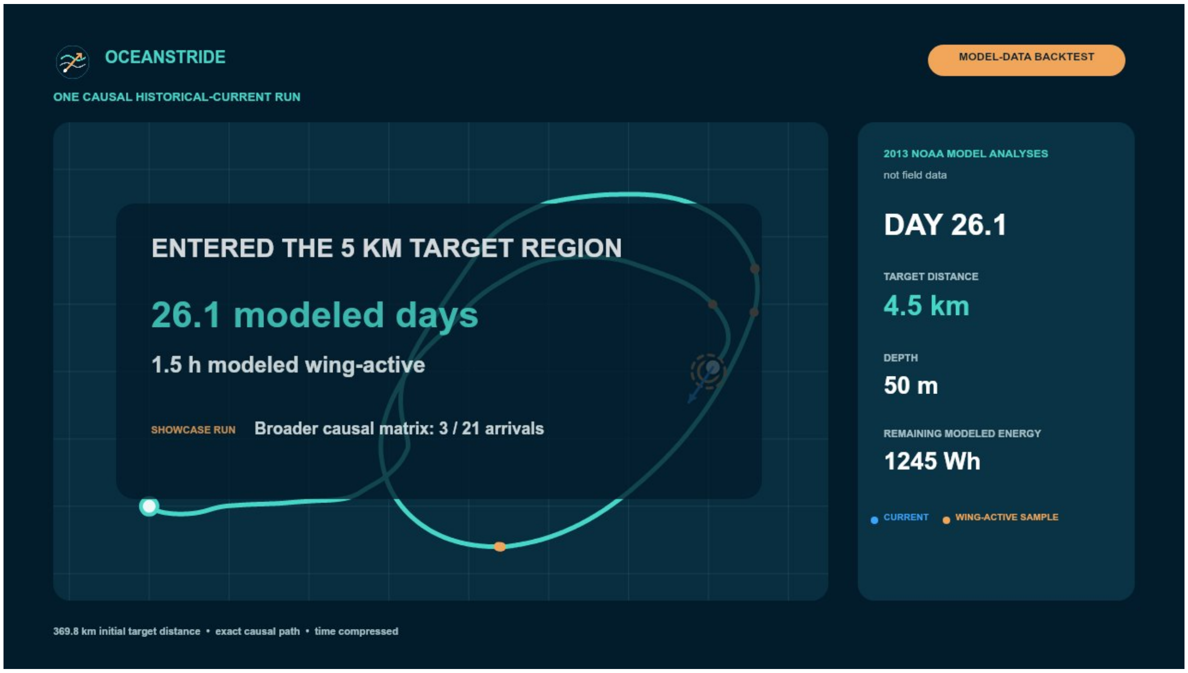

Recorded path for one successful causal historical-current model case. Modeled output—not a field track.

Short 42-hour numerical regression

The original short comparison uses the NOAA AEC Gulf of Mexico climatology for June 1-3, 2000. It begins near 25 degrees N, 90 degrees W, places a waypoint 20 km east and 5 km north, uses a 3 km arrival radius, and runs for no more than 42 hours. The arrival radius is comparable to the source-grid scale, and the case was intentionally configured to be reachable. It is useful for deterministic regression and cross-language verification, not as evidence of an operational deployment.

Policy	Modeled result	Final distance (km)	Elapsed hours	Energy used (Wh)	Wing duty (%)	Vertical travel (m)
Passive	No arrival	22.1627490409225	42	83.9999999999809	0	0
Depth only	No arrival	18.2252230599726	42	84.1199999999811	0	40
Hybrid	Modeled arrival	2.99024059924407	41.4166666666667	125.891666666665	99.7991967871486	7,450

The hybrid result used a 99.8% modeled wing duty cycle and repeated sawtooth motion. It must **not** be cited as evidence of intermittent wing use, low-energy current riding, or a validated DriftWing duty cycle. Its appropriate use is to show that a known scenario produces the same decisions and trajectory across implementations.

Cross-language numerical consistency

MATLAB is the reference implementation. A shared `WBGRID1` file supplies identical latitude-fastest `[latitude, longitude, depth, time]` field values to MATLAB, C, and the C++/CUDA codebase. The recorded NOAA regression grid contains 195 x 251 x 23 x 3 samples. Its provenance ledger records:

- source NetCDF SHA-256: `dcde6abbc0b6f5ff290219f856df91fbb8fa3a85ac33543c158a764ddd539f39`; and
- exported `WBGRID1` SHA-256: `1f13492bf8cafb4d206a59cdd7d66c361958f2afd983ee06b720cfa96ca3323b`.

For the passive, depth-only, and hybrid regression policies, the recorded comparison has exact agreement on policy, sample count, time, target-depth commands, wing commands, and 64-bit action signatures. The largest logged continuous-state or current difference between C and the MATLAB reference is:

3.637978807091713 x 10^-12, at the depth-only trajectory's east-position value on row 217.

This is below the documented 4 x 10^-12 comparison tolerance. The same checker also passed for the CUDA codebase's CPU fallback traces. The C deterministic test passed interpolation, domain, kinematics, planner, and repeatability checks. The CUDA CPU-fallback test log records **14 passed, 0 failed, and 2 skipped**; one skip is the platform-specific `/dev/full` write-failure test and the other is CPU/CUDA device parity.

The GPU boundary is important: CUDA kernel source is authored for a batch model that assigns one independent mission to each GPU thread, but this workspace did not have a CUDA toolkit or `nvcc`. The GPU translation unit has therefore not been compiled, executed, or benchmarked here. No GPU speedup claim is supported.

Numerical agreement verifies implementation consistency. It does not verify that the model accurately predicts a physical vehicle or the ocean.

Principal limitations

1. **No physical article exists.** There is no measured buoyancy engine, wing mechanism, pressure housing, sensor payload, communications subsystem, or recovery system.
2. **Kinematics are simplified.** The model omits full rigid-body dynamics, stability margins, lift and drag polars, stall, turning dynamics, actuator transients, trim, and vehicle-environment coupling.
3. **Energy inputs are assumptions.** Hotel load, control load, pump cost, capacity, efficiency, temperature effects, battery aging, and reserve behavior have not been calibrated to hardware.
4. **The historical “causal” study is not an operational forecast hindcast.** It uses persistence of the issue-time historical analysis; archived future data are used only by the oracle upper bound.
5. **Wave forcing is a sensitivity model.** Bulk wind-sea and swell statistics are converted to a depth-decaying drift estimate; this is not a spectrum-resolved Stokes product.
6. **Geospatial constraints are incomplete.** Missing HYCOM cells indirectly represent land or invalid water, but there is no independent bathymetric-clearance, collision, entanglement, traffic, geofence, or recovery-risk model.
7. **Coverage is narrow.** The matrix covers one historical Gulf window, three routes, and seven starts on a deliberately strided grid. It is not a reliability estimate for another season, region, or operating envelope.
8. **Arrival is a modeled geometric event.** Entering a 5 km circle is not equivalent to maintaining station, completing a survey pattern, surfacing successfully, transmitting data, or being recovered.
9. **Sensor performance is outside the model.** No statement about dissolved oxygen, conductivity, temperature, fluorescence, turbidity, timing, calibration drift, fouling, or data quality has been physically demonstrated.

Physical validation gates

Acceptance thresholds should be fixed before each test and retained with raw data, calibration records, scripts, anomalies, and pass/fail decisions. The following gates convert the present analytical proof into progressively stronger evidence.

Gate	Go/no-go question	Required evidence before advancing
0. Requirements and safety review	Is there a coherent, testable architecture for the selected customer use case?	Customer-informed mission requirements; environmental and recovery envelope; mass, buoyancy, power, communications, and sensor budgets; interface drawings; preliminary FMEA/hazard log; reviewed test plan and acceptance criteria
1. Instrumented dry bench	Can the core subsystems actuate repeatably and be measured?	Buoyancy and wing actuation over a predefined repeated-cycle test; current, voltage, energy, position, timing, and temperature logs; control-electronics and fail-safe demonstration; anomaly and wear record
2. Pressure and ingress	Can the mechanical article tolerate its proposed depth envelope?	Independently reviewed pressure-test method and proof factor; calibrated pressure record; seal inspection; verified absence of ingress; post-test functional check; documented operating limit
3. Instrumented tank	Does the integrated article produce controllable, repeatable motion?	Repeated profiles; measured vertical speed, stability, trim, turning response, wing behavior, energy per cycle, and telemetry; video and synchronized sensor logs; comparison against pre-registered model predictions
4. Controlled-water learning mission	Can the system operate and be recovered safely outside the tank?	Site and safety approvals; supervised deployment; geofence and abort behavior; communications and recovery evidence; repeated trajectory and energy logs; issue register and corrective actions
5. Paired-data pilot	Can OceanStride deliver a useful monitoring outcome?	Multiple recoverable missions; co-located reference CTD/sonde comparisons; calibration and QA/QC record; data completeness and latency; customer-reviewed report; measured mission cost and operator effort

Only after the relevant gates are passed should OceanStride publish physical depth, endurance, sensor-accuracy, reliability, recovery-rate, or cost-performance claims.

Evidence provenance

The authoritative internal artifacts for this note are:

- `README.md` — model purpose, data sources, workflows, and limitations;
- `config/makeLongMissionParams.m` — month-scale simulation assumptions;
- `src/makeLongStudyScenarios.m` — route and release-date matrix;
- `src/planLongAction.m`, `src/makePlannerActions.m`, and `src/stepVehicleLong.m` — planner, finite glide-to-depth logic, kinematics, and energy accounting;
- `results/month_study_2013/month_summary.csv` — 84 baseline runs plus three additional battery-sweep rows;
- `results/month_study_2013/paths/central_to_east_20130601__lookahead_causal.csv` — recommended exact causal showcase path;
- `results/mission_summary.csv` — short 42-hour comparison;
- `validation/noaa_aec_golden.json` — reference hashes, action signatures, trajectories, and numerical tolerance;

- `c/build-cmake/Testing/Temporary/LastTest.log` — recorded C deterministic test;
- `cuda/build/cpu-release/Testing/Temporary/LastTest.log` — recorded CUDA CPU-fallback validation status; and
- `cuda/README.md` — intended batch architecture and explicit local GPU-compilation limitation.

Figures and rounded presentation copy should be treated as derived views. The CSV and JSON artifacts above remain authoritative for numerical claims.