

What the Sea Pushes Against

Choosing a coastal energy system when the site, not the technology, makes the decision

Matthew Schulz

Independent researcher · mattschulz97@protonmail.com

ORCID 0009-0000-0956-5743

CC BY 4.0 · Engineering specification, v21 · August 2026 · Not peer reviewed

Keywords: wave energy converter, oscillating water column, coastal defence, storm surge, submerged compressed-air storage, wave-driven desalination, DC-coupled microgrid, ice production, levelised cost sensitivity, appropriate technology, flywheel rotor materials, selection tool

A coastal machine should be selected by what it pushes against, and paid for by what it produces — and electricity is rarely the best product.

1. The problem, and why it is not a box

Coastal and island communities with weak grids or none pay some of the highest electricity prices on earth, and they sit next to one of the largest untapped energy resources on earth. The obvious response is to build a machine that turns waves into electricity and sell it to them.

That framing has produced four decades of prototypes and very few installations, and the reason is that there is no box to buy. A 2023 survey of one substantial sub-class of converters found that **most technologies have not passed technology readiness level 4, exactly one had reached level 8, and none was market-ready**, and named the cause as a failure to converge — a proliferation of concepts, slow development, device failures and investment losses [16]. **A document that helps somebody choose is therefore aimed at the problem the field says it actually has, which is not a shortage of inventions.** A wave machine is not one device. It is at least three, and each of the three is chosen independently, from different constraints, by different people:

- a **structure** that stands in the water and takes the load;
- a **conversion and transmission chain** that gets the energy from the water to somewhere useful;
- a **store** that holds what arrives when nobody wants it, until somebody does.

Change the site and any one of the three can flip while the other two stay put. A structure that works on a rock shelf is impossible on a mud delta. A conversion chain that suits an industrial supply chain is unmaintainable where there is no electrician. A store that makes sense on a stiff grid is worthless on an island. **So the useful output of this work is not a machine. It is a selection guide.**

We evaluated a number of approaches and deprecated several of them. Those are not walked through in the body — they are summarised with their reasons in **Appendix A**, so the pages ahead carry only what survived.

Section 2 gives the compatibility table: which structure works with which conversion chain and which store. Section 3 gives the site factors that force the choice — environmental, seasonal, geographic and financial, several of which are absolute and cannot be engineered around. Sections 4 to 6 review the survivors in generation, transmission and storage.

And Section 6 makes the argument this work exists to make. For many of these communities the most valuable thing a coastal machine can store is not electricity. It is **ice and fresh water** — because those are already the community's binding constraints, because storing them is an order of magnitude cheaper than storing the equivalent energy in batteries, and because making them locally takes load off infrastructure that is already inadequate.

2. The three components, and what can work together

2.1 The structure — what does it push against?

A wave carries momentum. An absorber takes some of it, and by Newton's third law that momentum must go somewhere. **Every wave machine must push against something**, and what it pushes against fixes the foundation, the siting envelope, the failure modes and whether the machine can also be a sea wall.

There are only three answers, and they are the families this document is organised around.

The categories are not ours. The EquiMar project proposed classifying wave energy converters by **referencing configuration — self-referencing, or seabed/shoreline referencing** — alongside the mode of motion of the primary conversion element, and "self-reacting" is long-standing vocabulary in the field with its own experimental literature [7]. The European Marine Energy Centre maintains a parallel and more widely used scheme that sorts converters into eight primary types by operating principle [18]; the two axes are complementary, and this document uses the referencing axis because it is the one that predicts where a machine can physically stand. **BEDROCK, BULWARK and DRIFTER are memorable names for an existing taxonomy, not a new one.** They are used because a reader carries a name further than a category label, and because the paper refers to them on every page.

Family	Pushes against	Archetype	Earns the defence credit?
BEDROCK	a fixed structure on the seabed	hinged blade	Yes — its foundation can be the barrier
BULWARK	the shore, or a wall built on it	oscillating water column	Yes, most of all — the structure is the barrier
DRIFTER	itself — float against float	articulated chain	No — it protects nothing

The consequences follow directly. BEDROCK's reference is fixed in space, so the machine lives in a depth band and *the tide moves the water relative to the machine* — it needs a shelf, a stable bed and a modest tidal range. BULWARK's moving element is a column of water inside a fixed chamber, and its working point is the *free surface*, which rises and falls with the tide by itself, so tide stops mattering. DRIFTER has no fixed reference at all — the reference travels with the device — so it works where no shelf exists, where the bed moves and where the tide is enormous, and pays for that by protecting nothing.

2.2 The compatibility table

This is the table the rest of the paper elaborates. Read a row to see what a family can be fitted with; read the notes to see why the exclusions exist.

	BEDROCK	BULWARK	DRIFTER
Capture	hinged blade; overtopping	oscillating water column, single or phased bank; overtopping	articulated chain; orbital rotor
Conversion	hydraulic to shore; sealed-tube generator; seawater pump	air turbine in the throat — the column is the piston	sealed-tube generator; seawater pump; hydraulic to a local hub
Transmission	oil pipe or DC cable to shore	direct — the machine is on the shore	DC cable to shore, with a subsea connector
Store — fast	accumulator, then flywheel, both onshore	accumulator, then flywheel, both onshore	accumulator at the hub; flywheel onshore
Store — slow	ice room, water tank, submerged air vessel	ice room, water tank, submerged air vessel	ice room, water tank, submerged air vessel
Excluded	—	a hydraulic circuit is redundant: the air path already does the work	anything requiring a fixed seabed reference

Two things in that table need saying out loud, because both confused an early reader.

The conversion chain is not interchangeable across families. BULWARK's air turbine sits in the chamber throat and has no equivalent in the other two; BEDROCK and DRIFTER need a hydraulic circuit, a sealed-tube generator or a pump because they have no air column to work with. Which conversion chain you can fit is decided by the structure, not chosen freely.

The slow stores are universal, and they are not attached to the machine. The ice room, the water tank and the submerged air vessel are separate objects sharing an electrical bus with the generator. The submerged vessel in particular sits in deep water some distance offshore and is connected to shore by its own air hose — **it is not bolted to a wave wall and does not store "wave energy" as such.** It stores whatever the bus has spare, from any source. That is why Section 6 treats storage as a site-level decision rather than a machine feature.

2.3 The one physical limit that governs everything upstream

For a regular deep-water wave of height H and period T , the power crossing each metre of wave crest is

$$P = \rho g^2 H^2 T / (32 \pi)$$

In words: the power arriving at each metre of structure grows with the **square** of wave height and only in direct proportion to period. Double the swell and you quadruple the power; double the period and you merely double it, because a longer wave delivers the same shape less often. Density and gravity set the scale and do not move — roughly one kilowatt per metre for a one-metre, one-second wave.

A caution worth carrying into any wave calculation. For an *irregular* sea described by a significant wave height and an energy period, the same expression carries 64π instead of 32π — **exactly half**. Mixing the two is the commonest error in this field, and it is a factor-of-two error in the direction that flatters the project. Every figure here uses the regular-wave form at $H = 2$ m, $T = 10$ s, which gives **39 kW per metre of crest**.

Two numbers follow, and they are different numbers:

Quantity	Value per kilometre of frontage	What it means
Incident flux	39 MW/km	The entire energy crossing that frontage. A hard ceiling. Nothing can exceed it.
Deliverable at 30% capture	11.8 MW/km	What a real machine takes out. This is the capacity of record.

In a 30 kW/m annual-average climate the annual-average delivery is roughly 9 MW/km.

2.4 The veto is seasonal, not annual

There is one condition that excludes every family, and stating it as an annual mean — as earlier drafts of this work did — gets the wrong answer in both directions.

Wave machines in these settings are not standalone plants. As §3.2 sets out, they are usually the wet-season half of a pair with photovoltaics, which are anti-correlated with them: the sun is strongest when the sea is calmest. **So the question is not what the resource averages over a year. It is what the resource delivers during the months when the other source is weakest.**

The two errors an annual test makes:

Site	Annual mean	Wet-season mean	Dry-season mean	Annual test says	Seasonal test says
Flat trade-wind coast	12.0 kW/m	12.0	12.0	passes	fails — no season where it beats solar
Mild monsoon, 60% in 4 months	10.0 kW/m	18.0	6.0	marginal	passes — carries the solar gap
Strong monsoon, 65% in 4 months	8.0 kW/m	15.6	4.2	fails	passes — the pairing case exactly
Very strong monsoon, 75% in 4 months	6.0 kW/m	13.5	2.2	fails	marginal — check the store

A site at 8 kW/m annual mean is told by the annual test to build solar instead — but 65% of its wave energy arrives in the four monsoon months at 15.6 kW/m, which are precisely the months when photovoltaic output is worst. That is the anti-correlated pairing this document recommends, and the annual test throws it away. **And a flat 12 kW/m coast passes the annual test** while having no season in which it beats solar at all.

The veto, stated correctly: below roughly 10 kW/m during the season when the complementary source is weakest, nothing in this document is worth building. A kilometre of any structure buys less power than a small solar farm, and the correct answer is more solar and more storage.

Applied that way it still removes the Arabian Gulf and Red Sea, where the resource is low in every season. It no longer removes monsoon coasts on the strength of an average that hides the case for building there — those must be screened on the wet-season figure and on what §6 can store through the dry months.

Which means the annual mean is the wrong number to collect. A site assessment for this work needs a **monthly** resource curve and the monthly curve of whatever it will be paired with. An annual figure cannot answer the question.

3. What the site decides

Each site brings factors that force the choice. Several are absolute — they cannot be engineered around, only designed away from by selecting a different family.

	BEDROCK pushes on the seabed <i>hinged blade</i>	BULWARK pushes on the shore <i>water column</i>	DRIFTER pushes on itself <i>articulated chain</i>
Wave resource in the season of need below ~10 kW/m	excluded	excluded	excluded
Spring tidal range above ~4 m	excluded	works	works
Deltaic or mobile seabed	excluded	marginal	works
Narrow steep-to shelf, no bench	excluded	works	works
Water deeper than ~30 m at the line	excluded	works	works
Reef, seagrass or protected area	excluded	marginal	marginal
Tropical cyclone exposure	marginal	works	marginal
Seasonal sea ice	excluded	marginal	excluded
Strong longshore sediment drift	marginal	excluded	works
Storm-surge defence also required	works	works	excluded

Read ACROSS a row to find which family survives that condition — not DOWN a column.

The first row is the only universal veto, and it is a SEASONAL test — see §2.3.

BEDROCK carries the most exclusions and is still right wherever a shelf exists, because BEDROCK and BULWARK are the only families that earn the coastal-defence credit.

Figure 1. Site conditions against family. **Read across a row** to find which family survives that condition. Only the first row excludes all three.

3.1 Environmental and geographic

Wave resource is the universal veto of §2.4 — and note that it is a *seasonal* test, not an annual one. It is the only condition that excludes everything. Everything else is a family selection. A six-metre tidal range that rules out BEDROCK leaves BULWARK untouched, because the chamber's working point floats. A deltaic seabed that buries a hinge is irrelevant to a moored DRIFTER. A steep-to volcanic shelf with nowhere to stand a blade array is a fine site for a chain.

Two conditions bite across the board. **Reef, seagrass and protected areas** remain marginal for all three — the nearshore habitat this design promises to create already exists at such a site, and altering sediment transport across it is the harm rather than the benefit. And **where surge defence is itself a requirement**, DRIFTER is excluded by definition.

Cyclone exposure excludes nothing; it inverts the economics. Survival load rather than production sets the capital cost, so the machine becomes a barrier that generates rather than a generator that protects — and the defence credit should rise to match.

3.2 Seasonal

A monsoon coast and a trade-wind coast pose different problems with the same annual mean, which is why §2.4's veto is seasonal. Where the resource is concentrated in one season, the machine is sized for that season and the store bridges the rest.

Photovoltaics and waves are often anti-correlated: the sun is strongest when the sea is calmest. That is not a complication to be managed — it is the reason the pair works, and it should be the default assumption rather than an afterthought. A microgrid running solar in the dry season, waves in the wet, and a generator for backup beats any single source, and the wave machine's job is to be **good enough during the months when the sun is not**, which is a much easier target than being good all year.

The corollary matters for procurement: **a wave machine on such a coast should be sized against the solar deficit, not against annual demand.** Sizing it as though it had to carry the year produces a machine several times larger than the system needs, and that oversizing lands directly on levers 1 and 2 in §7.

3.3 Supply chain

This is the constraint that decides the conversion chain, and it is frequently the binding one. A community can build a hull from bamboo, sealed drums or ferrocement and rope it with coir. It cannot build a shaft seal, a permanent-magnet generator, a power inverter or a subsea connector. Most published wave architectures quietly assume an industrial supply chain at exactly the point where the target site has none.

3.4 Financial

Two site facts change the answer by more than any engineering decision, and both are in §7: **the tariff the machine displaces**, and **whether the project is financed as a power plant or as coastal defence.** A site with island diesel at \$350/MWh and a site on a national grid at \$60/MWh are not the same project with different revenue — they are different machines making different products.

4. Generation — the survivors

Each entry gives what it is, **the cycle** — how energy goes in and comes back out — its numbers, and what remains unknown. Deprecated approaches and their reasons are in Appendix A.

4.1 The hinged blade · BEDROCK

What it is. A composite blade hinged to a seabed foundation in roughly 8–20 m of water, standing in the path of the incoming wave.

The cycle. The wave's orbital motion swings the blade through an arc of some $\pm 25^\circ$, driving a sealed hydraulic cylinder that pushes oil to shore through a fused pipe. Ashore, the oil is smoothed by an accumulator and spins a compact turbine at around 1,500 rpm. A check valve makes the flow one-way, so the push stroke and the return stroke both deliver.

Pairs with. Hydraulic circuit (§5.1) or sealed-tube generator (§5.2) or pump (§5.3). Its foundation carries the barrier, so it earns the defence credit.

Why the nearshore band, and it is a better reason than depth convenience. Wave force on a bottom-hinged flap *rises* in shallow water because of shoaling: for a ten-second wave in about 12 m of water the horizontal force is roughly **50% greater** than the same wave would exert in deep water [6]. The device is not tolerating the shallows — it is exploiting them.

Numbers. 30% wire-to-wave assumed. Modular — one blade replaceable without touching the others.

Unknown. Capture efficiency has never been independently estimated, and every capacity figure in this document rests on that 30%. Blade fatigue at 15.8 million cycles is uncomputed; hinge burial by sediment is unmodelled.

4.2 The oscillating water column · BULWARK

Origin. The oscillating water column is the oldest documented wave energy converter: the first recorded research was by the Japanese marine captain **Yoshio Masuda**, who developed a floating oscillating water column for navigation buoys from the 1940s [17].

What it is, physically. A concrete caisson standing at the shoreline, with a chamber inside it that is open to the sea below the waterline and closed above, trapping a volume of air. **There are no moving parts in the water at all.**

It is shore-connected, and it comes in two forms. It can be a **section of a longer wall** — one caisson among many in a breakwater or sea defence, which is how the phased bank of \$4.3 is built and how the defence credit of \$7.2 is earned. Or it can be a **single standalone caisson** on a stretch of shoreline, which is the village-scale form. **It is never an offshore island.** Everything about it — the reaction, the air turbine, the electrical connection — depends on it being founded on or against the shore, which is precisely what distinguishes BULWARK from DRIFTER.

The cycle. The passing wave drives the water level inside the chamber up and down like a piston. The trapped air is alternately compressed and expanded, and that oscillating air stream drives a bidirectional turbine in the throat.




Figure 2. The oscillating air-column chamber in section. Water column as mass, trapped air as spring, tuning valve changing the spring in service, perforated seaward face doing dissipation and storm relief as geometry rather than mechanism.

Pairs with. Its own air turbine and nothing else — the column is the piston, so a hydraulic circuit would be redundant. Sits on the shore, so no subsea cable and no connector.

Why it is the most forgiving capture element. The water column has mass and the trapped air has stiffness, so together they form an oscillator. In words: the natural period gets **longer** when there is more water to accelerate and **shorter** when the air above is stiffer, which happens when the trapped volume is smaller. A blade's tuning is fixed in the moulding shop and is wrong forever afterwards if the sea is not what was assumed. **A chamber's stiffness lives in a volume of air, and a volume of air has a valve on it — so the chamber can be re-tuned in service to the sea it actually gets.** For a poorly characterised coastline, which describes nearly every site here, that beats a higher peak, because peak efficiency at the wrong frequency is worth nothing.

The chamber has a classical resonant length too: a column open at one end and closed at the other resonates when its length is a quarter of the wavelength, so $L \approx \lambda/4$ is the geometric starting point, with the air spring providing adjustment around it.

Build the seaward face perforated. A perforated wall with a void behind it — a Jarlan-type caisson — lets the wave pass through the holes and lose energy to turbulence at every orifice, reflecting far less than a solid face. **This matters structurally more than it matters energetically.** A solid vertical face returns the wave almost intact; the reflected wave meets the next incoming one, and the standing pattern scours the seabed at the toe until the wall undermines its own foundation. Permeability is how a coral reef survives, and it does three jobs with one geometry here: cuts reflection, dissipates in normal operation, and in a storm passes the excess through rather than accumulating it as wall load. **That is storm relief with nothing to actuate and nothing to jam.**

This is not our reasoning and it is sixty-five years old. The perforated caisson was introduced by Gerard Jarlan at the National Research Council of Canada in 1961 [8], with early built examples at Comeau Bay in 1966 and at Kobe and Takamatsu in 1969–70. The breakwater literature independently establishes that such structures reduce reflection **and horizontal wave force** and **reduce scour of the seabed in front of the structure** — precisely the argument above, already demonstrated.

Numbers. 25–40% assumed. Cost dominated by concrete, which makes it the element that benefits most from the shore-side construction economics of \$7.4.

And the literature challenges this claim directly, which must be said here rather than buried. A 2023 review of hull-encapsulated converters surveys the frequency-tuning approaches proposed across the field — tunable springs, mass rearrangement, latching and declutching, gyro-speed control — and concludes that **none has been proven implementable in real applications**, citing complexity, high cost and very limited tuning ranges. Its recommended direction is the opposite of ours: **pursue non-resonant devices** with simple designs that sit tolerably across a wide frequency band at moderate capture, on the argument that robustness and frequency tolerance may reach a competitive cost of energy where sharp tuning has not [16].

That is a serious, sourced objection to the central advantage claimed for this family, and it is not answered here. Two things distinguish the chamber from the devices reviewed — its spring is a volume of air rather than a mechanism, and the adjustment is a valve rather than a moving mass — but **whether that difference survives contact with a real sea is exactly the untested claim**. If it does not, §4.3's phased bank becomes the only reason to prefer BULWARK, since a bank achieves breadth without tuning anything.

Unknown. The tuning range is uncomputed — that a valve can re-tune across a useful band is central to this entry and is currently an assertion, now contradicted by a published review. The perforation ratio is uncomputed and trades dissipation against the pressure differential that drives the column.

4.3 The phased chamber bank · BULWARK

What it is. Several chambers in a row with different water-column lengths, so each peaks at a different wave period.

The cycle. Identical to §4.2, run in parallel across a band instead of at a line.




Figure 3. One chamber against a bank of five, irrationally spaced. The single chamber is sharp and narrow; the bank is broad and nearly flat. No individual chamber improved — the coverage did.

How to space them. Two chambers at a small whole-number ratio reinforce at the same frequencies instead of covering different ones. Spacing by a ratio that is *not* a simple fraction spreads the peaks as evenly as possible.

Pairs with. As §4.2. Belongs to BULWARK alone — staggering a blade array means a dozen different mouldings; staggering a chamber bank means one formwork poured to a dozen depths.

Numbers. Marginal cost over a single chamber is essentially the extra depth of pour.

Unknown. Depends entirely on §4.2's tuning range being real.

4.4 The articulated chain · DRIFTER

What it is. A line of floats connected by hinged joints, each joint carrying a spring and a piston. Nothing is anchored for reaction — only a slack mooring so it does not drift away.

The cycle. Adjacent floats sit at different points on the wave at the same moment, so they move differently from one another, and that *relative* motion drives the pistons.




Figure 4. The chain along a wave. The harvest is the articulation at the joints, and it exists only because the chain spans phases.

The one hard requirement. The chain must be long compared with the wave or nothing happens — if it is shorter than one wavelength, every float rides up and down together and no piston moves. Deep-water wavelength follows from the period alone:

$$\lambda = g T^2 / (2 \pi)$$

In words: the longer the wave takes to pass, the further apart its crests are, and the relationship is quadratic — a wave of twice the period is four times as long. At ten seconds that is **about 156 metres**, so a working chain wants one to two wavelengths. **A short chain is not a small chain; it is a raft.**

Two design consequences. A joint harvests only motion in its own plane, so capturing more than one axis needs **two hinge axes per joint**, doubling the joint count. And **most segments can be passive spacers** carrying load and cable and containing nothing, with generating joints at intervals — the length bought cheaply and the machinery count decided separately. One correction to the intuition: over a full wave period every joint sees the same total articulation, merely phase-shifted, so there is no clever placement that extracts more per generator. The saving is entirely in the spacers.

Pairs with. Sealed-tube generator (§5.2) or pump (§5.3), then a DC cable to shore through a subsea connector — which §8 identifies as the item most likely to kill this configuration on cost.

Numbers. Not estimated. Needs a full wavelength before the first watt, the worst first-increment threshold in the family.

Unknown. Chain response has not been modelled; the wavelength requirement is stated from geometry alone. Moorings and the dynamic cable become the principal wear items and are uncoded.

4.5 The orbital rotor · DRIFTER · unproven, kept on the board

What it is. A rotor with its axis horizontal and parallel to the wave crest, blades in the plane of the wave's orbital motion.

The cycle. Water particles under a deep-water wave do not move back and forth — **they move in circles.** Water above the axis moves one way while water below moves the other, and that difference is torque in a single consistent direction, renewed by every wave. **The sea is already rotating, so nothing has to be rectified** — no check valves, no ratchet, no reciprocating stroke.

The number that decides how to use it. A rotor driven by the orbit completes one revolution per wave, because that is what an orbit is. Its shaft speed is therefore the *wave* frequency — about **6 rpm** at a ten-second period. A conventional generator wants 250 times that, and direct drive at 6 rpm needs roughly **32,000 N·m** for a mere 20 kW. **It buys direction, not speed,** and must not be sold as a solution to the rate problem.

Unknown. Capture width uncomputed; the velocity differential across a rotor of practical diameter is under 0.1 m/s. **A model rotor and a torque meter in a flume settles it, and that is the cheapest experiment in this document.**

5. Conversion and transmission — the survivors

A wave arrives at about 0.1 Hz; a grid runs at 50 or 60. That ratio of five to six hundred is the central problem of every wave energy converter, because induced voltage depends on the *rate* of flux change and there is almost none at that speed. **Every conversion chain is a rate transformer.**

No gearbox anywhere in this family. A geartrain in seawater at fifteen million cycles is the failure mode that retired the original design twice over — see Appendix A. Distributed small generators are welcome; a gearbox is the rejected component in a new coat.

5.1 Hydraulic circuit — BEDROCK, DRIFTER

What it is. A sealed cylinder at the absorber, a fused oil line, and a turbine in a dry shed ashore.

The cycle. A gear ratio without gears. The absorber sweeps slowly through a large arc and displaces a small volume of oil at high pressure; that oil, smoothed by an accumulator, spins a compact turbine at 1,500 rpm. Slow and forceful in, fast and light out.

Why it survives. Only oil crosses to shore, and a pipe penetration is static. Seawater and oil never meet, there is no electrical connection in the water, and every short-lived component sits on dry land.

Numbers. \$1,800/kW assumed. Standard industrial components ashore.

Unknown. Line-breach volumes and biodegradable fluid selection unspecified.

5.2 The sealed-tube generator — BEDROCK, DRIFTER

What it is. A free magnetic slug — a *pig*, after the pipeline tool — inside a sealed tube, with coils wound on the outside.

The cycle. The wave drives the slug back and forth; its motion sweeps magnetic flux past the coils and generates directly. **Nothing crosses the pressure boundary at all** — no shaft, no rod, no gland — and the electrical machine lives on the dry outside where it can be replaced without opening anything.




Figure 5. The sealed-tube generator. The slug is the only moving part and never leaves the tube; the coils sit outside, dry and replaceable.

Numbers. Two thousand turns, 0.25 T in the gap, half a metre per second of stroke: **94 volts, about 700 newtons of reaction, 175 watts.** Twenty-nine tubes make five kilowatts using about 193 kg of copper — **roughly \$462/kW in wire**, against \$1,800/kW for the hydraulic route. Use **ferrite** for the magnet — iron oxide with barium or strontium, the material in every loudspeaker: cheapest available, abundant feedstock, and it does not corrode. Neodymium is three times stronger and rots in salt water.

The ceiling: about half the power heats the coil. At a matched load that is not a defect, it is what maximum power transfer means. Efficiency is bought back only by surrendering power — 175 W at 50%, or 112 W at 80%. Only speed buys both: at eight metres per second the identical hardware gives **44.9 kW instead of 175 W**, a factor of 256. Unobtainable from a wave, ordinary at a nozzle — **which is the quantitative case for \$5.3.**

Unknown. Force density unestimated, so the required magnet mass is unknown and may be prohibitive.

5.3 The seawater pump — all three families

What it is. A piston and two check valves. No windings, no magnets, no electronics, no synchronisation.

The cycle. The absorber drives the piston; pressurised seawater travels ashore; on land it either turns a Pelton wheel for electricity or goes straight through a **reverse-osmosis membrane**, because RO wants *pressure*, not volts. In a conventional plant the wave would become electricity, electricity would become pressure, and each conversion would take its cut. **The pump makes the pressure directly.**

Why it survives, and it is the \$3.3 answer. It can be maintained by anyone who has maintained a borehole pump. It is the only conversion chain that does not assume an industrial supply chain in the water.

Unknown. Check valves are wear items and uncoded. Membranes are consumables and are not locally made, so the supply dependency moves rather than vanishes; fouling and pretreatment are the real operating burden of every small RO plant; brine disposal must be designed. **Where a site has geothermal heat at 70-120 °C, multi-effect distillation makes fresh water with no membranes at all** — see \$6.5.

5.4 Transmission: rectify locally, combine on a DC bus

What it is. A rectifier at each unit, one DC busbar, and fault-managed feeders to the houses.

The cycle. Every unit rectifies at itself. The direct current combines at a busbar. Flywheel, battery, RO pump, ice plant and air compressor all hang off that same bar with no conversion between them. Feeders leave the compound at a voltage that is current-limited and monitored.




Figure 6. The single line, with the fence between skilled and unskilled territory drawn where it belongs.

Why the inverter stage disappears. A variable-speed marine generator produces variable-frequency alternating current that **must be rectified in any case**. In an AC system it is then inverted back to a fixed frequency, costing money, efficiency and a large field-unreparable component. On a DC bus the chain stops at the rectifier.

Why series AC would give nothing. Absorbers along a wall and segments along a chain sit at **different phases of the same wave** — that is the entire reason they move. Alternating sources at different phases add as *vectors*: five units spaced seventy-two degrees apart have a scalar sum of five units and a **vector sum of exactly zero**. Wiring them in series to build voltage yields nothing. Rectify locally, then combine — precisely a photovoltaic string, and a second, independent reason the bus must be DC.

Why voltage is the whole trick. Resistive loss is current squared times resistance, so delivering fixed power at ten times the voltage needs a tenth the current and loses a hundredth. Low voltage is used in these settings not for engineering reasons but because of safety rules and the licensed trades they require.

What fault-managed power buys. Class 4 is a circuit class introduced in the 2023 edition of the US National Electrical Code (NFPA 70) under a new Article 726, where such circuits are called **Fault Managed Power Systems** — the first new power class in over four decades. They are explicitly *not* power-limited, delivering hundreds or thousands of watts at up to **450 V AC or DC**, and they achieve safety by **intelligently limiting the energy that can enter a fault**, which mitigates shock and fire risk and permits installation using methods associated with power-limited circuits [5]. Power travels as a **pulsed waveform** from a transmitter to a receiver that converts it back to usable DC, and a fault of any kind — short circuit, line-to-line, ground fault, overcurrent, control malfunction — interrupts transmission within milliseconds, with the system resetting automatically once cleared [9]. The governing safety standards are **UL 1400-1** for systems and **UL 1400-2** for cable; only a system meeting UL 1400-1 is a Class 4 system. Published transmission figures for the class include more than 100 W beyond 100 m [5] and **520 W carried 500 m on a pair of 16 AWG conductors** [9]. At this time the permitted voltage is 450 V; review of the NFPA First and Second Draft Reports for the 2026 cycle indicates Class 4 systems **will be permitted in dwelling units**, a change from 2023 which restricted them to non-dwelling occupancies [10], and that is the regulatory basis for any domestic-retrofit case.

The following are stated engineering input from the author's commercial practice rather than from the citations above: stored energy of order fifteen joules at the source capacitor; operation below the let-go threshold for muscular contraction; automatic shutdown on skin contact; installation without a licensed electrician as a general statement; and an anticipated future rise toward 850 V.

The efficiency belongs to the voltage, and any high-voltage scheme would get it. What fault management adds is permission — the right to run that voltage in a small cable, through a village, installed by people who are not electricians. Which is precisely the §3.3 constraint.

Put every difficult thing inside one fence and make everything downstream unskilled.

Collection is high-voltage DC on conventional gear — dangerous, properly protected, skilled, done **once**, inside a compound the size of a shed. Distribution is fault-managed DC: small cable, no trenching crew, no licensed installer, little to no arc threat at the point of use. This also splits the DC arc problem, which is real: direct current has no zero crossing at which an arc self-extinguishes, so the voltage that makes protection hard stays behind the fence with the switchgear designed for it.

Unknown. Every fault-managed figure above is stated input, not verified. For DRIFTER only, the **subsea connector** joining the floating unit to the shore cable is uncosted and is the single item most likely to kill that configuration — see §8.

6. Storage — and the argument this paper exists to make

Four timescales, four different answers. And the ones that matter most do not store electricity at all.

Store	Holds	Bridges	Attached to
Hydraulic accumulator, 250 L, 200–350 bar	2.80 MJ	2.3 s	the conversion chain
Flywheel, 2 t at 1.5 m, 2,000 rpm, usable to half speed	37.0 MJ	30.8 s	the DC bus, onshore
Submerged air vessel, 500 m³ at 100 m	369 kWh	hours	the DC bus, via its own air hose offshore
Ice room and water tank	a product, not energy	days	the DC bus, onshore

6.1 The fast pair: accumulator and flywheel

The cycle. The accumulator absorbs each stroke and releases it smoothly. The flywheel absorbs the difference between one wave and the next by speeding up, and gives it back by slowing down.

Numbers. An accumulator holds approximately $E \approx p_1 V_0 \ln(p_2/p_1)$ — in words, what it holds depends on how much gas you have and how hard you squeeze it, and because the relation is logarithmic, squeezing harder buys much less than it appears to. A flywheel holds $E = \frac{1}{2} I \omega^2$ with $I = \frac{1}{2} m R^2$ — in words, energy climbs with the **square** of speed and only linearly with mass, which is why flywheel design buys speed before weight, and why radius counts twice: once for the mass it carries and again for how fast the rim travels. Two tonnes at 1.5 m and 2,000 rpm stores **49.35 MJ, or 13.71 kWh**.




Figure 7. What each store bridges against what the sea asks for. The accumulator cannot reach from one crest to the next; the flywheel can.

The whole argument in one line. The accumulator is a *ripple filter* and the flywheel is a *period filter*, moving the horizon from about two seconds to about thirty — an order of magnitude, and precisely the span between crests. The flywheel costs shaft speed: because energy varies as the square of rotational rate, absorbing one full wave period of rated output drops the rotor **13%**, which is acceptable for a variable-speed machine with an inverter and fatal for a directly-coupled synchronous one.

Where the flywheel is mounted, and the free fix

A spinning rotor exerts no force on its mountings from spinning. A reaction torque appears only when something tries to **turn the spin axis**, and its magnitude is the angular momentum times the rate at which the axis is being swung:

$$\tau = I \omega \Omega$$

In words: how hard the flywheel fights back depends on how much angular momentum it holds and how fast you are rotating its axis. Sitting still on a fixed foundation, that is zero. **Bolted to a flap that pitches with every wave, it is not.**

Worked at a plausible scale. A 100 kW flap on an eight-second wave delivers about 800 kJ per stroke; buffering half of that across a 10% speed band — which releases only 19% of what is stored, because energy goes as the square of speed — needs a rotor holding about **2.1 MJ**. In steel at roughly 45 kJ/kg that is a **47 kg rim at 0.4 m radius turning 7,150 rpm**, with a rim speed of 299 m/s and an angular momentum of about **5,600 N·m·s**. Pitch that through $\pm 20^\circ$ at an eight-second period — a peak swing rate of 0.274 rad/s — and the gyroscopic couple is **1.5 kN·m**, perpendicular to both axes, reversing through **3.9 million full cycles a year, forever**. For a structure meant to last fifty years, that is precisely the load that writes the failure story.

And the fix costs nothing, because the torque is a cross product. Mount the flywheel with its spin axis parallel to the hinge axis, and pitching the flap rotates the spin axis about itself — which is no rotation at all. The gyroscopic reaction is exactly zero, by geometry.

Spin axis, relative to the hinge axis	Gyroscopic couple
0° — parallel	0 kN·m
15°	0.40 kN·m
30°	0.77 kN·m
45°	1.09 kN·m
90° — perpendicular	1.54 kN·m

Same hardware, same sea, entirely different mounting problem. **Better still, put the flywheel on the fixed frame or ashore rather than on the moving absorber and the question never arises** — which is what §9.1 already requires for unrelated reasons.

And an entire device family does the exact opposite, deliberately. Gyroscopic converters — the Inertial Sea Wave Energy Converter developed at Politecnico di Torino is the best known — encapsulate a spinning rotor inside a hull and **use the gyroscopic reaction as the power take-off**, tuning the machine by controlling gyro speed and gyro torque [16]. They maximise precisely the couple this section sets out to eliminate. That is not a contradiction: **it is the same physics read in both directions**, and it is worth knowing that the load you are avoiding is somebody else's product.

Arraying: what it shares and what it cannot

Splitting the store across several rotors is worth doing, but three separate things get run together under "sharing the load."

Shaft and hub torque does share. Power splits across N units, so each shaft sees $P/(N\omega)$. Four rotors, a quarter of the torque per hub.

Gyroscopic load can be cancelled outright. Run the rotors as **counter-rotating pairs** and the angular momenta sum to zero, so the platform may pitch as it likes with no net reaction on the foundation. Each rotor's bearings still carry their own couple, but the two are equal and opposite and close through the common frame. The cancellation is only as good as the speed matching, so a pair must be controlled as a matched set, with a stated plan for the failure of one.

Internal rotor stress does not share, and this is the one to be clear about. Hoop stress is $\sigma = \rho v^2$, which depends **only on rim speed** — in words, how hard the material is being pulled apart depends on how fast the rim is moving and on nothing else. Ten rotors at 299 m/s carry exactly the same 704 MPa as one rotor at 299 m/s. **Arraying ruggedises the system and cannot make the material's life easier.** The only lever on rim stress is slowing down, and that costs energy as the square.

What arraying does buy: N times the energy; the option to **stagger the speed bands**, cascading rotors one at a time through an efficient window instead of swinging one large rotor wide, which is a real round-trip gain on an irregular wave stream; and redundancy, so losing one rotor costs 1/N of capacity rather than the machine. Set against that, housings, bearings and containment do not scale down proportionally — expect roughly 20–40% more total mass for the same energy, and more bearings to maintain. For manufacturability the array argues the same way steel did: several modest rotors are commodity parts from any shop, where one large rotor is a long-lead forging from a short list of suppliers.

What the flywheel rotor should be made of

A flywheel stores energy in *motion* rather than in a chemical or electrical difference, which is why it does not degrade with cycling the way a battery does — there is nothing to wear out chemically, only bearings and windage to lose energy to. That makes the rotor material the interesting question, and it has a tidy answer. For a rotating rim the specific energy per unit **mass** is a shape factor times tensile strength divided by density, so the energy per unit **volume** is simply

$$e_v \approx K \sigma$$

In words: **volumetric energy density is the shape factor times the tensile strength, and density cancels out entirely.** A heavier rotor is not a denser store. A *stronger* one is.

Rotor material	Volumetric (thin rim, $K = 0.5$)	Per unit mass
Steel, 1.0 GPa	139 kWh/m ³	~18 Wh/kg
Glass fibre, 1.4 GPa	194 kWh/m ³	~90 Wh/kg
Carbon fibre, 2.4 GPa ultimate	333 kWh/m³	~208 Wh/kg
Carbon fibre at 50% design stress	167 kWh/m ³	~104 Wh/kg

Read those numbers carefully, because the headline figure is easy to over-claim. The 333 kWh/m³ is *rotor material volume at ultimate strength*, with no safety factor and excluding the vacuum housing, magnetic bearings and burst containment. Apply a realistic 50% design stress for cyclic duty and it halves; count the installed volume and it

falls several times again. Lithium-ion is roughly 200–400 kWh/m³ at **pack** level, so "comparable to lithium" is true rotor-to-cell and **not** true installed-to-installed.

The verdict for this design, and it is not the obvious one. Carbon fibre's real advantage is not volume — at a working design stress it is barely ahead of steel per cubic metre. Its advantage is **mass**, where it is roughly an order of magnitude better. And §9.1 puts this flywheel **onshore, in a dry shed**, where both mass and volume are cheap.

A third reason applies wherever the drivetrain is impulsive. A ratchet or a check valve delivers torque in impulses rather than smoothly, and impulsive loading is exactly where steel's toughness pays and where a composite's weak interlaminar shear hurts most. A steel rotor is also a single forging with no hub-to-rim joint, so the delamination failure mode does not exist, and it accepts ordinary non-destructive inspection — ultrasonic, magnetic particle, dye penetrant — which matters for an asset that can only be visited in good weather.

A second reason points the same way, and it is the §3.3 constraint again. A carbon-fibre rim is filament-wound on controlled tooling in a specialist facility and arrives as a finished import. A steel rotor can be forged, cast or machined in a great many more places, inspected by ordinary means, and rebalanced locally after a bearing change. For a design whose whole premise is that the community can maintain what it owns, **that difference outweighs the energy density.**

So steel is the right rotor here. Carbon fibre becomes the answer only where the powerhouse footprint is genuinely constrained, or where a future configuration puts the store on something that floats or must be craned — and in either case it should be recognised as an imported component with an imported repair path.

Storage selection rule. The accumulator is always fitted. The flywheel is fitted where the coupling is weak or absent — islanded operation, microgrid service, or where the installation is a significant fraction of local capacity. It is omitted on a stiff interconnection, where its foundation and shaft volume are nonetheless **reserved at build time**. Grid stiffness is a property of the decade, not the site: large continuously-drawing loads are being added faster than generation in several markets, and reserving concrete is cheap while coring it out later is not.

6.2 The submerged air vessel — a site-level store, not a machine part

What it is. An upside-down concrete cup on the seabed, open at the bottom, in deep water offshore, connected to the powerhouse by its own air hose. **It is not attached to any wave machine.** It stores whatever the bus has spare, from any source — waves, solar, wind or a generator running below rated.

This is underwater compressed air energy storage, and it is an established field. Variants have been proposed since at least the 1980s; **Seymour at UC San Diego proposed a rigid accumulator — a pipe or tank with ballast bins — in 1997**, the closest published antecedent to the vessel described here. Garvey and Pimm at the University of Nottingham, with Thin Red Line Aerospace, tested a five-metre energy bag in 25 m of seawater at EMEC, Orkney, in 2012 [11]; Hydrostor built a grid-connected 1 MW demonstration in Lake Ontario in 2015. **The constant-pressure argument below is Garvey's, not ours** — he states that at depth the pressure stays almost constant across all fill levels, giving roughly three times the storage per cubic metre of a fixed-volume vessel at the same upper pressure.

The cycle — read this before the properties, which only make sense after it. To **charge**, pump air down the hose. The air cannot escape upward through the concrete or sideways through the walls, so it pushes the water down and out of the open bottom. You now have a bubble held under the sea, and the sea is squeezing it — constantly, at whatever pressure that depth provides, because the sea does not run out. To **discharge**, open a valve: the sea pushes the water back up into the cup, and that pushes the air back out through the hose and through a turbine. **It is a battery whose charge is a bubble and whose pressure vessel is the ocean.**

The property that makes it worth building. A rigid tank's pressure *collapses* as you draw it down, so the turbine sweeps across its efficiency curve and spends most of the discharge off its best point. A vessel at depth does not — the sea holds it at hydrostatic pressure, **constant from full to empty**.

Numbers.

Depth	Sea pressure	Store per m ³ of vessel	Vessel for 360 kWh	As a hemisphere
30 m	4.0 bar	0.16 kWh	2,321 m ³	r = 10.3 m
50 m	6.0 bar	0.30 kWh	1,197 m ³	r = 8.3 m
100 m	11.1 bar	0.74 kWh	488 m ³	r = 6.2 m
200 m	21.1 bar	1.79 kWh	201 m ³	r = 4.6 m

The ocean's slowness helps here rather than hurting. Compressing air to 10 bar adiabatically costs about 326 kJ per cubic metre and delivers it at roughly 283 °C, all thrown away. Isothermal compression costs 230 kJ — **29% less** — and industry cannot reach it because compressors run fast. **A wave compresses over about five seconds into an infinite constant-temperature bath.** This is the one setting where isothermal compression is not a textbook fantasy.

It is pumped storage with the sign reversed. Pumped hydro lifts water against gravity; this pushes water *down* against buoyancy, and the energy expression is identical — density times gravity times height times volume. It is available to an island with deep water offshore but no mountain and no valley to dam. Per cubic metre of vessel it is 2.6× denser at 100 m than pumped hydro — but a reservoir is a pond and a vessel is concrete, so **cost per cubic metre decides, and where a site has real topography pumped hydro should win.**

How to build it and hold it down. Cast it as **two half domes from one mould** — formwork dominates one-off marine concrete, the joint falls at the equator where geometry is simplest, and the halves nest for tow. A flexible bag would need a gas-tight membrane with a fatigue life, and gas-tightness is what natural rubber is worst at; **the open-bottom dome needs no membrane at all**, and a water surface does not fatigue.

A 500 m³ vessel full of air displaces 512 tonnes of water; its shell holds down 83. **Net uplift is 429 tonnes**, needing about 229 m³ of caged local rock. That is not a flaw — it is what storing energy *as buoyancy* means. And the siting rule is counter-intuitive: buoyancy depends on displaced volume, not depth, so ballast at 50 m and at 100 m is identical while stored energy rises from 150 to 369 kWh. **Depth buys storage at no ballast cost.**

Depth must be set by a length, never by a force balance. An untethered vessel at neutral buoyancy is unstable in both directions: push it two metres deeper and the air compresses, buoyancy drops about nine tonnes, and it sinks further; lift it and it runs for the surface. Slung beneath a surface float or standing above seabed ballast, a length cannot run away.

That arrangement gives three things free. Ballast below the centre of buoyancy is a **pendulum** — five metres of sling gives about 1.8 MN·m of righting moment at five degrees, so it barely tilts. Cable tension becomes a **direct state-of-charge readout**, measurable with a load pin. And **notching the rim** makes overfill relief free: excess air burps out under the edge with no valve to jam.

Unknown, and the field's own warning is sharper than ours. Round-trip efficiency unmeasured here; air dissolution into seawater at pressure unquantified; ecological footprint of a seabed structure at depth unassessed. **And the UW-CAES literature reports that the onshore section works well while the underwater section remains the problem** — marine components are described as the principal barrier to return on investment, and planned projects at Lake Huron and Aruba were terminated on those grounds. Any proposal built on this section must answer that before it answers anything else.

6.3 Ice and water — the argument

The cheapest store is the one where the stored thing is already what somebody wanted.

What it is. A water tank and an ice room, on the DC bus alongside everything else.

The cycle. Surplus electricity runs the RO pump or the ice plant when it is available; the water and the ice sit there until wanted. **You never get the electricity back — and that is the point.** You did not need electricity. You needed cold and water.

Numbers. At a coefficient of performance around three, a tonne of ice costs roughly 27 kWh and holds about 93 kWh of cooling. A day's water for two hundred households is around 32 m³, about 112 kWh of freely deferrable load. Water in a tank costs a fraction of the battery holding the same energy.

Somebody has already built this, and it is the strongest thing in this section. A system developed by AIREF with GIZ Indonesia and PT ASTB uses **thermal energy storage instead of batteries** to produce up to **one tonne of ice per day entirely off-grid**, powered by a 25 kWp photovoltaic array, with a cooling capacity of roughly **100 kWh per tonne** — no grid electricity and no large battery, explicitly so that locally caught fish can be refrigerated, and described as scalable and replicable for coastal regions facing similar logistics [12]. Their 100 kWh per tonne and our 93 kWh per tonne are two independent routes to the same number.

And the durability argument is better than the cost argument. A microgrid case study running ice thermal storage alongside batteries reports that the **batteries were vulnerable to damage from cycling and could not be fully discharged without permanent damage, while the ice store could** — and that ice storage is proven, durable and has no moving parts, so maintenance is minimal [13]. A battery degrades by being used. A tank of ice does not.

Why this is the paper's central claim, and it is about infrastructure rather than thermodynamics. For a fishing community, ice is not a convenience — it is the difference between selling a catch and losing it, and in most such places the ice arrives by truck or boat from somewhere else. Fresh water frequently arrives the same way. **Both are already being transported into the community at cost, over infrastructure that is already inadequate.** Making them locally does not merely save the electricity that would have made them; it removes a delivery burden from roads, boats and cold chains that have no spare capacity.

And the second-order effect is larger than the first. In a microgrid sized to peak, halving a major load *releases capacity*. If ice is forty percent of demand and its energy per tonne halves, total load falls about twenty percent — so the same generator serves roughly a quarter more households, or the generator itself shrinks by a fifth. Since conversion cost is the second-ranked lever in §7, **a thermodynamic improvement in the ice plant propagates directly into the two levers that dominate the economics.**

6.4 Halving the cost of ice with deep water

The prerequisite is deep water close to shore — **the same prerequisite as §6.2**, so the two share a pipe route.

Condenser cooling	Condensing temperature	Realistic COP	Electricity per tonne of ice
Tropical air	35 °C	3.0	30.8 kWh
Surface seawater	30 °C	3.4	26.9 kWh
Deep cold water	12 °C	7.1	13.1 kWh

A factor of 2.4, from one pipe. Same compressor, no new physics. Deep water used directly for space cooling is better again, because no refrigeration cycle is involved at all.

6.5 Where a site has geothermal heat

Geothermal is rare on the coastlines this document addresses and, where it exists, is usually blocked by finance rather than physics — see Appendix A. But where a resource is already proven, two combinations change the picture and both need the deep water this document already requires.

Pair geothermal heat with deep cold water. A 150 °C resource against 5 °C deep water is a 145 K gradient where the Carnot ceiling is roughly 34% rather than the 6.7% available from the ocean thermal gradient alone. Few sites have been evaluated that way.

Better still, do not make electricity with it. Low-grade geothermal heat at 70-120 °C runs **multi-effect distillation**, which produces fresh water using heat rather than pressure and **requires no membranes**. That removes the one imported consumable \$5.3 admits it cannot source locally, and it is the same move as §6.3 — spend the resource on what it is good at.

7. What sets the cost

The base case for every figure below is one kilometre of BEDROCK — hinged blades, hydraulic conversion to shore, accumulator and flywheel, grid-tied — at default assumptions, landing at \$185/MWh. Each parameter was then swept one at a time through a plausible range and the swing in levelised cost measured.

#	Lever	LCOE range	Swing
---	-------	------------	-------

1	Capacity factor, 20 → 45%	\$123 → \$278	\$154
2	Conversion chain, \$900 → \$3,600/kW	\$136 → \$283	\$147
3	Discount rate, 3 → 12%	\$131 → \$263	\$132
4	Capture efficiency, 15 → 45%	\$164 → \$248	\$84
5	O&M, 2 → 6%/yr	\$156 → \$234	\$77
6	Marine build premium	\$185 → \$236	\$51
7	Structure, \$6k → \$24k/m	\$168 → \$219	\$51
8	Absorbers, \$3k → \$12k/m	\$171 → \$214	\$43
9	Defence credit, 0 → 100%	\$170 → \$208	\$38
10	Contingency, 15 → 40%	\$170 → \$207	\$37
11	Flywheel cost	\$167 → \$203	\$36
12	Design life, 25 → 100 yr	\$181 → \$207	\$25
13	AC versus DC coupling	\$185 → \$203	\$18
14	HVDC switchgear	\$182 → \$192	\$10
15	Fault-managed distribution	\$185 → \$185	\$1

How to read the swing column. It is the difference the lever makes between its best and worst plausible value with everything else held fixed. A lever with a \$150 swing can move the answer by more than the entire electrical layer; a lever with a \$1 swing cannot move it at all.

The entire electrical layer sits at the bottom. Coupling, switchgear and fault-managed distribution together move under \$30/MWh. That is no argument against them — DC is free efficiency and fault management solves an installation-labour problem with no other answer — but the project is not won there.

7.1 Clip the generator — the best single move available

Levers 1 and 2 are the top two and one decision pulls both. **Rate the conversion chain below the peak sea state.** You lose a little energy at the top of the distribution and buy a far smaller power train, and conversion is roughly 38% of capital. Clipping about a fifth takes the cost from \$185 to **\$129/MWh**; clipping two fifths reaches **\$99/MWh**.

Solar and wind have done this for years and there it is a genuine energy loss. **Here it is not**, because the stores of \$6 exist precisely to catch what is clipped. **The store is not bolted onto the generator — it is what lets the generator be small.**

7.2 Procure it as defence, not as power

Levers 3 and 9 are one decision. As a merchant power project — eleven percent money, no defence credit — the same machine in the same sea lands at **\$282/MWh**. As adaptation infrastructure at three and a half percent concessional money, 85% of the barrier charged to defence, hundred-year life, it lands at **\$122/MWh**. A factor of 2.3, with nothing about the engineering changed.

And a coupling that only shows on a second look: design life is nearly worthless alone. Doubling from fifty to a hundred years saves 3.3% of the capital charge at seven percent money and 0.5% at merchant rates — discounting has erased everything past year forty. At three and a half percent it saves 15.2%. **Durability and cheap money multiply.**

7.3 What escalation does and does not do

Rising electricity prices never move levelised cost by a cent; it is a cost, not a value. They move the worth of what is displaced. Against island diesel at \$350/MWh, one megawatt-hour a year is worth about \$2,300 over fifty years flat, \$3,800 at two percent real, \$6,300 at four. Against a weak mainland grid at \$140/MWh the project is under water flat and at two percent and clears only at four. **Escalation rescues marginal sites and barely changes good ones.**

And the standing comparison. Forecast levelised costs for first commercial-scale wave projects run roughly \$120–470/MWh against utility solar near \$43/MWh. **Against grid solar this loses by a factor of four or five and no assumption set rescues it.** Against island diesel at \$200–500/MWh it is competitive. So the electricity case exists only where the grid is weak or absent — and everywhere else the machine must be sold on its other products.

7.4 Build it dry, and build it in pieces

Marine construction is expensive for reasons unrelated to the structure — barge work inside weather windows at marine labour rates, with vessel day-rates running whether or not anything is poured. The same cubic metre placed in a field costs a fraction. **So the biggest decision is not how it is designed but where it is assembled**, and this applies to every BEDROCK and BULWARK element, to the submerged vessel of \$6.2, and to the unit bases of \$9.2.

Pour upright in a graving dock or dug basin using ordinary slip-forming and land plant; flood the basin, float out on the structure's own buoyancy, tow to station, and set down by controlled flooding of internal ballast cells. Large offshore concrete gravity structures have been built this way for decades precisely because it moves the expensive work ashore. **And it makes the poorest-community case realistic for the first time** — a basin, formwork and a concrete plant can exist where no marine construction industry does.

Build dry, float out, ballast down — the expensive work happens ashore




Figure 8. The float-out sequence. Every expensive hour happens ashore.

Incremental build changes the SHAPE of the spending, not the total




Figure 9. The same total capital spent two ways. Incremental building moves the first revenue forward.

Segment-at-a-time construction is financial rather than logistical. These projects fail to get built not because the total capital is impossible but because all of it must be committed before any of it produces revenue. Incremental construction converts one large sunk commitment into a series of small ones, each paying before the next is made.

One hard limit applies to DRIFTER only — a chain does not produce until it spans a meaningful fraction of a wavelength, so its smallest useful increment is of order 156 m. BEDROCK and BULWARK have no such threshold: one blade or one chamber generates on the day it is commissioned.

7.5 Two structural moves not yet taken

One conversion chain shared by many absorbers — a common hydraulic manifold feeding a single turbine — attacks lever 2 where a machine per absorber does not. No version of this design has made that choice explicit.

The structure is a breakwater, so it makes a harbour. A sheltered basin has real value, and in many target communities the absence of a safe landing constrains the fishing economy more than the absence of electricity. Aquaculture on the structure is a second such line. Both are credits against lever 7.

7.6 Running this yourself — the selection tool

Everything in Sections 2, 3 and 7 is a set of gates and a set of sums, which means it can be handled better by a small interactive tool than by a reader flipping between tables. **One accompanies this paper.**

It is a single self-contained HTML file that runs in any browser with no installation, no network and no account. You set the site — wave resource, tidal range, water depth, wave obliquity, and the presence or absence of a mobile bed, a steep-to shelf, cyclone exposure, longshore drift, protected habitat or sea ice — and pick a family. The screen re-reads on every change, and when a family fails it **names which family to switch to** rather than telling you not to build. A second switch flips the whole system between AC and DC coupling and shows the conversion chain redrawing, with the inverter and transformer stages visibly struck through on DC.

The part that matters most is that every cost assumption is an editable field with a provenance badge. Nothing is hard-coded and nothing is hidden. Figures with no traced source are marked as such, figures supplied from commercial practice are marked separately, and the comparators carry their own labels. **The purpose is not to tell anyone what a wave wall costs.** It is to let somebody with better numbers than ours replace ours and get their own answer in a few minutes — and to make it obvious which inputs are doing the work, which is the same argument \$7's swing column makes.

It also carries a scale check we found necessary. Enter a village-sized load and it reports the ratio against a deployed 3–5 kW community system, because a kilometre of wall is roughly three thousand times that and the mismatch should be visible before anyone pitches the wrong machine to the wrong customer.

Its limits, stated as plainly as the paper's. It models a wave machine and not the mixed system §3.2 recommends, so it cannot yet show the seasonal pairing with photovoltaics that most of these sites will actually build. It contains no wave-resource database — you supply the number, and the regional presets are unverified working estimates, not hindcasts. And its arithmetic is the arithmetic in this paper, so every caveat in §10 applies to it unchanged.

8. Three worked selections

Each is a walk through the same gates — structure, conversion, transmission, store — with the site forcing each choice. Regional resource figures are working estimates and are **UNVERIFIED**.

8.1 The port machine

Site. An existing harbour, town or airport with a shelf, a working grid connection, and active erosion or surge risk.

Gates. *Structure:* BULWARK if the tide is large, BEDROCK if it is not — both earn the defence credit. *Conversion:* air turbine for BULWARK, hydraulic for BEDROCK. *Transmission:* direct to the town grid, DC bus at the powerhouse. *Store:* accumulator always; flywheel reserved but omitted on a stiff grid; submerged vessel if the water offshore is deep; ice at the fish dock.

Scale and money. A kilometre at 12–40 MW/km, built incrementally by float-out, **procured as adaptation infrastructure** per §7.2 — which is where the factor of 2.3 lives.

Products in order. Surge defence, sheltered water, electricity to the town grid, ice.

8.2 The village unit

Site. A fishing village of fifty households with no grid, on any coast that passes §2.3.

The arithmetic that defines it. At roughly 11.8 kW per metre of frontage, fifty households need **about 1.3 metres** of absorber, and a small community facility a third of a metre. **A village machine is not a small wall — it is a single unit.**

There is a precedent for what that means commercially. **Photovoltaics did not win on efficiency** — a panel converts a modest fraction of what lands on it and always has. It won by becoming **a repeatable unit you can buy one of**: same part, same mounting, same wiring, strung and combined at a busbar, extended next year when there is money. Nobody engineers a solar farm from first principles; they instance a unit. Marine energy has never had that, which is much of why it remains a sequence of prototypes.

Gates. *Structure*: BEDROCK where there is a shelf, DRIFTER where there is not. *Conversion*: sealed-tube generator (\$5.2) or pump (\$5.3) — chosen on \$3.3, the supply chain, not on efficiency. *Transmission*: one cable to a shore busbar, rectified at the unit. *Store*: water tank and ice room first; flywheel fitted because the grid is absent; submerged vessel if the water is deep.

The unit: one absorber, one base, one sealed generator, one cable, one connector — combined at a DC busbar exactly like a photovoltaic string.

What is local and what is not. Hull, ballast, mounting, concrete and labour are local. Copper and ferrite are imported regardless — but **coil winding is unskilled work with a jig**, so wire arrives on a spool and is wound on site. The concrete hemispheres of \$6.2 are cast by the same crew.

Products in order. Ice, water, electricity.

The falsification target is the connector, not the physics. Photovoltaic strings are cheap because air-insulated cable and a screw terminal are cheap. Subsea wet-mate connectors are neither and are a standard failure point. On a unit of a few tens of kilowatts a single connector could rival the generator in cost, and if it does the string architecture collapses into bespoke installation. **That is a supplier quote, not an experiment, and it should be answered before any other work on the unit.**

8.3 The mainland coastal retrofit

Site. A coastline with a real grid — a southern Indian fishing coast, a West African city and its satellites. Moderate resource, active erosion, ice arriving by road.

Gates. *Structure*: BEDROCK or BULWARK, whichever the tide allows. *Conversion*: hydraulic or air. *Transmission*: grid-tied, so the flywheel is omitted and its space reserved. *Store*: ice room and water, sized to displace what is currently trucked in.

Money. The electrification case largely evaporates here; the ice case does not, and neither does erosion. The economics rest on **defence plus cold** rather than on generation, and \$7.3's escalation analysis does its work, because these are the marginal sites.

Why one document covers all three. Real coastlines carry all of them at once: an island has a port and satellites; a national ministry finances the harbour machine and buys village units from the same supply chain; and a country addressing remote villages may buy the small configuration precisely because it can be built in-country with local labour and materials.

9. Design rules this work produced

9.1 The concrete outlives the copper. A marine concrete structure is a fifty-to-hundred-year asset; an electrical machine is a twenty-year asset. Any architecture welding them together throws away the long-lived thing when the short-lived one dies. **Put the interface at a pipe, a busbar or a plug you can reach, and keep every short-lived part on land.** The test for any new idea: *if servicing it in year twenty means opening something poured in year zero, it is the wrong architecture.*

9.2 Resonance is the objective for the absorber and the enemy for the structure. An absorber that does not ring at the wave frequency takes little energy; a structure that does, cracks. The dynamic-excitation check is therefore **scooped to structural members and does not govern conversion elements**, and the mount between a resonant

absorber and a non-resonant structure becomes a compliance item in its own right. Margins and the separate self-excited-instability check are in the companion standard [1]; **resonance avoidance does not protect against flutter, galloping or vortex lock-in.**

9.3 Depth is set by a length, never by a force balance. See §6.2.

9.4 No gearbox anywhere. See §5.

9.5 Geometry beats mechanism. The perforated wall does storm relief with nothing to actuate; the notched rim does overfill relief with no valve; the folded blade and the flooded chamber both fail into being a wall. **Every member of this family fails by ceasing to generate and continuing to do its other job.**

10. What would falsify this

Open analyses, required before construction and not yet performed:

1. **Absorber capture efficiency has never been independently estimated** for any family. The 30% used throughout is assumed and every capacity figure rests on it.
2. **Flutter, galloping and vortex-induced vibration have not been analysed** for the blade or the chain's joints. Most likely place for an unpleasant surprise.
3. **Chamber tuning range has not been computed** (§4.2) — **and a published review reports that no frequency-tuning approach in the field has proven implementable, recommending non-resonant designs instead [16].** This is the most directly contradicted claim in the document.
4. **Chain response has not been modelled** (§4.4).
5. **The perforation ratio has not been computed** (§4.2).
6. **The submerged vessel's round-trip efficiency is unmeasured**, air dissolution at pressure unquantified, ecological footprint unassessed (§6.2).
7. **The orbital rotor's capture width has not been computed** (§4.5).
8. **The sealed tube's force density has not been estimated** (§5.2).
9. **The subsea connector cost is unknown** (§8.2) — the cheapest item here to resolve and able to kill the village unit on its own.
0. **Fatigue life at 15.8 million cycles has not been computed** for any candidate material.
1. **Sediment transport and scour have not been modelled** for any family.
2. **Every fault-managed power figure in §5.4 is stated engineering input, not verified.**
3. **The gyroscopic mounting analysis assumes a rigid rotor on a rigidly pitching body.** Compliance in the mount, and misalignment tolerance in practice, are uncomputed — and the table above shows the penalty for getting the axis wrong is steep and immediate.
4. **Every regional resource figure in §8 is an unverified working estimate**, not a hindcast — and per §2.4 an annual figure is the wrong quantity anyway. **A monthly resource curve, paired with the monthly curve of the complementary source, is the minimum data a real site screen requires**, and no site in this document has one.
5. **No cost figure in §7 is traceable to a published source.**

Claims that would falsify the architecture if measured false:

- If extraction does *not* measurably reduce reflection at the toe, the structural argument of §4.2 fails and these are walls with generators bolted on.
- If a chamber's air-spring tuning range is too narrow to track a real sea state, §4.2's advantage over a blade evaporates.
- If the accumulator alone delivers grid-code-compliant quality on a weak coupling, the flywheel is unnecessary and §6.1's rule collapses to "never fit one."
- If fouling detunes absorbers faster than a five-year rotation can address, tuned resonance fails as a strategy across all three families.
- If the isolated mount cannot pass both the absorber's required compliance and the structure's required stiffness, §9.2 fails and every family must choose between capture and structural safety.

- If local demand for ice and water is smaller than §6.3 assumes, the paper's central argument fails and the machine must be sold on electricity after all.

11. Attribution

This section is deliberately incomplete and is the next work item. No systematic search has been performed, and it is stated as a gap rather than omitted.

Almost nothing here is novel as a mechanism, and that should be said first rather than conceded later. The seabed-hinged surge converter with hydraulic take-off, the shore-mounted oscillating water column, the overtopping device, the articulated floating attenuator, the perforated-face caisson, float-out construction of concrete gravity structures, the sealed-tube linear generator, wave-driven pumping and wave-driven reverse osmosis all have substantial existing literature and identified originators. **Each needs finding and citing by name.** The perforated caisson is the pattern for the whole exercise: recognising it as long-established breakwater practice saved inventing a relief port that would have been worse.

Two search passes have been completed and they moved four of five claims downward. That is recorded here rather than quietly absorbed.

Element	Status after two attribution passes
Reaction-class taxonomy	Withdrawn. It is the EquiMar referencing classification renamed [7]
Ice and water as the store	Not ours. Deployed off-grid in Indonesia for exactly this customer [12]
Submerged constant-pressure store	Not ours. Underwater compressed air energy storage, with a 1997 rigid-vessel antecedent. Only the two-half concrete casting method may remain
Absorber integrated with a surge barrier	Downgraded. Multi-functional structures combining coastal protection with generation are an existing review field
Two-stage store with a grid-strength selection rule	Untested. No antecedent found in two passes — the only element still standing

The honest position is therefore not novelty. Almost every element was done first by somebody who did it properly, and their results are available to cite. What this document offers is an **integration and a decision procedure** — the parts assembled against a site screen and a product screen, with the numbers attached, the failures recorded in Appendix A, and a tool in §7.6 that lets a reader run it on their own inputs.

That is offered as a contribution to convergence rather than to invention. The field's own diagnosis is a proliferation of concepts and a failure to converge on any of them [16]; more than a thousand designs exist and almost none has reached the market. A document whose purpose is to help somebody *choose* — and which publishes the things that did not work, with the numbers that killed them — is aimed at that problem rather than adding to it. **If a single coastal community uses the screen in Section 3 to avoid building the wrong machine, this will have been worth writing.**

One overlap is acknowledged rather than resolved. Bertram et al. (2020) describe a three-stage systematic approach for selecting the most suitable converter for a given site, whose first stage is a classification of the devices then in development [14]. **That is the same shape as Sections 2 and 3 of this document, and the reader should know it.**

The full text of that paper was not accessible to the authors of this one. It sits behind a publisher paywall, and no open-access or repository copy could be located. The abstract has been read; the methodology has not. **The extent of the overlap with Sections 2 and 3 is therefore unassessed,** and nothing here should be taken as a claim that the approach described in those sections is distinct from theirs. If a reader with access finds that it is not, Bertram et al. have priority and this document should be read as an application of their method rather than as an alternative to it.

This is stated plainly rather than omitted because a reader deserves to know where an author could not check. It is also a small illustration of the problem this work is aimed at: a selection method that communities and small operators cannot read is a selection method they cannot use.

Standing attribution notice. We have searched for prior art on the mechanisms described here and have not identified controlling patents or definitive originating publications for every element. That search was not exhaustive and should not be read as a claim of novelty. **If you know who got there first, tell us and we will credit them.**

Who operates it is the last unowned assumption. Every configuration presumes someone reviews telemetry, schedules the deferrable loads and notices a fouled absorber before it costs a season. Remote microgrid operations is a solved problem with existing providers running large fleets of village-scale systems across many countries, and adopting one beats inventing one.

A scale note, because the mismatch is easy to miss. Community systems in at least one substantial deployed programme — village microgrids in the Guatemalan highlands — are of the order of **3 to 5 kW each**. That figure is specific to that programme and should not be read as a universal, but it is the right order of magnitude for a village facility, and a kilometre of wall at 11.8 MW is roughly three thousand times it. **The natural meeting point between this work and village-scale operations is \$8.2's unit, not \$8.1's wall.**

Appendix A — Approaches evaluated and deprecated

Kept because an idea merely deleted comes back, and because the reason it lost is site-dependent — the next site may reverse several of these.

#	Approach	Where it applied	Why it was deprecated	What replaced it
A1	Surface flap → rod → rack, pinion, ratchet	BEDROCK conversion	The rod crosses the wall through a dynamic seal in seawater at ~15.8 million cycles over fifty years ; the ratchet is a repeated-impact device and the first wear item	Hydraulic circuit (\$5.1) and sealed tube (\$5.2), neither of which crosses the pressure boundary
A2	Any geartrain in the water	all conversion	Same cycle count, same seawater	Distributed small generators on a DC bus (\$5.4)
A3	Dedicated storm relief port in the chamber	BULWARK structure	One more thing that can jam	Perforated seaward face (\$4.2) — same job, no moving part
A4	Proportioning a chamber by the golden ratio	BULWARK geometry	Resonance is set by a mass and a stiffness — a length in metres and a volume in cubic metres against a period in seconds. A dimensionless ratio connects to neither	An irrational ratio is the right rule for the spacing between chambers in a bank (\$4.3), where the requirement is to avoid small whole-number ratios between neighbours
A5	Surplus power → hydrogen → fuel cell → desalination	storage	Wave to electricity to electrolysis to fuel cell to electricity to pump to membrane loses well over half before a drop of water	Pump straight to pressure (\$5.3), and store <i>water</i> rather than energy (\$6.3)
A6	Harvesting electrical potential between ocean depths	generation	Motional voltage is speed × field × separation; ocean currents move at ~0.5 m/s, so a kilometre-long electrode pair develops about 23 millivolts — and seawater's 4.5 S/m conductivity short-circuits it anyway . Dissimilar metals give real power but consume the anode, which is a battery, not a harvest	Nothing — the resource is not there. Salinity-gradient power at a river mouth is genuine and deserves a look where a river exists
A7	Thermoelectric conversion of the ocean thermal gradient	generation	At 20 K the Carnot ceiling is 6.7%; a good real engine reaches ~2.7% and a thermoelectric about 0.5% . No material improvement rescues it — Carnot sits above it	Use the cold as cold (\$6.4). If a low-moving-part engine is wanted, the thermoacoustic engine and the Fluidyne are the real options
A8	Raising steam from a submarine hydrothermal vent	generation	Vents sit at 200–300 bar where 350 °C water does not boil; recovery needs a multi-kilometre insulated riser; vent fluid precipitates sulfides fast enough to build chimneys, so any pipe scales shut; and vent fields are ecologically singular	Conventional geothermal where a proven resource exists (\$6.5)
A9	Island geothermal as a default option	generation	The barrier is risk capital, not technology. A well costs millions and its productivity is unknown until drilled; three dry holes can consume a small utility's entire budget, and no commercial lender finances an unproven resource	Kept as a <i>conditional</i> option where a resource is already proven and public risk-mitigation funding exists (\$6.5)

Appendix B — Three lessons that transfer beyond this project

B1 — Check the dimensions before designing the geometry. A1 and A4 both failed the same test: a proposed rule connected quantities that do not share dimensions. Resonance is a mass against a stiffness. A dimensionless constant cannot set either.

B2 — Count the conversions in the chain before optimising any one of them. A5 lost not because any stage was badly designed but because there were six of them. The winning path had one.

B3 — Check Carnot before designing a heat engine. A7 was ruled out before any hardware existed. At a twenty-degree gradient the ceiling is 6.7%, and no device — however elegant, however solid-state — gets above the ceiling.

References

- [1] Schulz M. *Two Checks for Structures in a Moving Environment: will the environment shake it apart, and will the environment outgrow it?* Zenodo (2026). Concept DOI 10.5281/zenodo.22105828 (always resolves to the latest version); this version 10.5281/zenodo.22105829. — the companion standard, referred to throughout and specifically at §9.2.
- [2] Falnes J. *Ocean Waves and Oscillating Systems: Linear Interactions Including Wave-Energy Extraction*. Cambridge University Press (2002).
- [3] Salter SH. Wave power. *Nature* **249**, 720–724 (1974).
- [4] European Marine Energy Centre. Wave resource assessments.
- [5] Eaves S. *What is Class 4 Fault Managed Power?* NECA-BICSI Summit, June 2023. — Class 4 as defined in NFPA 70 (2023) Article 726; Fault Managed Power Systems; up to 450 V AC or DC; safety by limiting fault energy rather than by limiting power.
- [6] Whittaker T & Folley M. Nearshore oscillating wave surge converters and the development of Oyster. *Phil. Trans. R. Soc. A* **370**(1959), 345–364 (2012). doi:10.1098/rsta.2011.0152. See also Folley M, Whittaker T & Osterried M, *The Oscillating Wave Surge Converter*, 14th Int. Offshore & Polar Eng. Conf., Toulon (2004).
- [7] EquiMar project — classification of wave energy converters by referencing configuration (self-referencing, or seabed/shoreline referencing) and mode of motion. **Primary EquiMar deliverable to be located and cited directly; presently known through secondary review literature.**
- [8] Jarlan GE. A perforated vertical breakwater. *The Dock and Harbour Authority* **41**(486), 394–398 (1961). See also Jarlan GE, *The application of acoustic theory to the reflective properties of coastal engineering structures*, Quarterly Bulletin, NRC Canada (1965).
- [9] Ibrahim M. *The Future of Power Distribution is Here: Introducing Fault Managed Power Systems*. Panduit (July 2024). — UL 1400-1 and UL 1400-2; pulsed transmitter-receiver architecture; 520 W over 500 m on 16 AWG.
- [10] Miller M. *Progress on Class 4 Circuits*. Electrical Association. — NFPA 2026 draft reports indicating Class 4 permitted in dwelling units.
- [11] Pimm A, Garvey S et al. Design and testing of Energy Bags for underwater compressed air energy storage. *Energy* **66**, 496–508 (2014). See also Pimm A & Garvey S, *Underwater Compressed Air Energy Storage* (book chapter).
- [12] WIPO. *Green Technology Book — Energy transition in fisheries and aquaculture*, §3.4 (2025). — AIREF / GIZ Indonesia / PT ASTB off-grid solar ice plant using thermal storage in place of batteries.
- [13] CALMAC microgrid case study — ice thermal storage alongside batteries; battery cycling damage versus ice store durability. **Secondary source; a peer-reviewed equivalent should replace it.**
- [14] Bertram DV, Tarighaleslami AH, Walmsley MRW, Atkins MJ & Glasgow GDE. A systematic approach for selecting suitable wave energy converters for potential wave energy farm sites. *Renewable and Sustainable Energy Reviews* **132**, 110011 (2020). doi:10.1016/j.rser.2020.110011. **Full text not yet read — see §11.**
- [15] Falcão AF de O. Wave energy utilization: A review of the technologies. *Renewable and Sustainable Energy Reviews* **14**, 899–918 (2010). — the standard review of the field.

[16] Maheen MH & Yang Y. Wave energy converters with rigid hull encapsulation: A review. *Sustainable Energy Technologies and Assessments* **57**, 103273 (2023). doi:10.1016/j.seta.2023.103273. — technology readiness levels across the field; the finding that no frequency-tuning approach has proven implementable; the recommendation to pursue non-resonant designs; gyroscopic converters including ISWEC.

[17] Masuda Y. An experience of wave power generator through tests and improvement. In: *Hydrodynamics of Ocean Wave-Energy Utilization*, Springer (1986), 445–452. — the originating oscillating water column work from the 1940s.

[18] European Marine Energy Centre. Wave device classification — eight primary types.

References [2], [3] and [4] carry substantive claims and should be verified against the publisher record before deposit.

16. Tool and computational resource disclosure

What the Sea Pushes Against · Matthew Schulz · ORCID 0009-0000-0956-5743 · CC BY 4.0 · DOI: TBD

Tools used. Claude (Anthropic) was used for literature and prior-art search, for drafting and restructuring prose, for performing and checking the arithmetic reported throughout, for generating the figures, and for producing the typeset artifact. Figures were generated with Python and matplotlib. Typesetting used python-markdown and wkhtmltopdf, driven by a build script that performs a font-coverage preflight before rendering. The interactive selection tool referenced in §7.6 was written in the same collaboration and is deposited alongside this document.

Responsibility. Responsibility for the correctness and adequacy of the arguments and results presented here, and for the completeness and accuracy of the citations, rests entirely with the named human author. No automated system is an author of this work.

Verification performed. Every number reported in this document was computed from stated inputs by a script rather than transcribed, and the scripts are deposited with this record. Where a figure was inherited from an earlier draft it was reproduced independently before being kept — the flywheel energy of §6.1, the capacity figures of §2, the gyroscopic couple of §6.1 and the sensitivity table of §7 were each recomputed from first principles, and the gyroscopic analysis was verified against nine separately stated quantities before it was accepted. The built artifact was checked for orphaned headings, for split heading–formula pairs, and for internal cross-reference consistency: every §-pointer, every figure pointer, every bracketed citation and every appendix row reference in this document was verified against its target by an automated pass over the source.

Attribution limits. Attribution is materially incomplete and §11 states this in full rather than burying it. Specifically: the full text of Bertram et al. [14] could not be obtained, so the overlap between that work and Sections 2 and 3 of this document is **unassessed**; the primary EquiMar deliverable [7] is known here only through secondary review literature and has not been read directly; reference [13] is a secondary source that should be replaced with a peer-reviewed equivalent; and the twelve remaining search targets listed in the accompanying attribution register have not been searched at all. **Several claims initially believed to be ours were found to have prior art during preparation and have been withdrawn or downgraded in §11.** Readers who can identify further prior art are asked to say so.

Computational resources. None material. All computation described here runs in seconds on ordinary hardware.