Applied Vacuum Theory, Part Three

The Self-Tuning Mesh: A Machine Independent of Place, Time, Person and Material

J. Konstapel, Leiden, 25 September 2026 — Third edition

Abstract

Part One read six builders (Moray, Schauberger, Tesla, the Correas, Reich, and Moddel–Thibado–White) through the Vacuum.Net theory. Part Two named the control variables and computed where the leverage lies. This part designs the machine. Every device of the builders was tuned by one person to one place. Moved, it stopped. The machine described here contains the tuning as a component. It shifts its own setting until output appears, then holds. This makes it independent of place, time and person. It is not independent of material: the material sets the yield. Every variable in this article is tied to a quantity a physicist can measure or count. The one reported yield, 70 W/m² per layer, is recomputed here from published data; it has not been independently replicated. The actuator that makes the mesh adjustable is specified: an electrostatic MEMS mirror, since thermal tuning is shown to be too weak by two orders of magnitude. The missing step, the map from mesh geometry to local density deviation, is written as a counting rule that predicts the shape of the output curve, not its magnitude, and is checked against one published trend.

1. What the machine is

The image throughout this work is a fishnet in water. Its strands are everywhere and always under tension. Where strands wind around each other, knots appear; we call that matter. Between strands lie meshes: openings with a shape and a size. A winding persists only if it returns in phase inside its mesh. What does not fit does not stay.

A builder's device is a mesh imposed on the local net. Moray's box, Tesla's coil, Reich's layered chamber, Moddel's optical cavity: each is an opening of a chosen size and asymmetry. Where the mesh fits the local state, windings pass through it, and the passage appears as light, motion, warmth or current.

The builders fitted the mesh by hand. That is why each device was bound to its maker and its site. The machine of this article separates two things:

- **The mesh.** A material opening with one adjustable dimension.
- **The loop.** A rule that adjusts that dimension until the output is maximal, then holds.

The loop is the generic part. It is the same everywhere. The mesh is the specific part. It is chosen per application for its yield.

2. The variables and where to find them

Each variable below has a physical location. A physicist can go there and measure it, or count it.

χ — the local state of the net. Dimensionless. Two readings are in use.

- $\chi\_\Phi = |\Phi|/c^2$, with Φ the gravitational potential at the site. At the surface of the Earth: $GM = 3.986 \times 10^{14}$ m^3/s^2, $R = 6.371 \times 10^6$ m, $c^2 = 8.988 \times 10^{16}$ m^2/s^2. So $\chi\_\Phi = 3.986 \times 10^{14} / (6.371 \times 10^6 \times 8.988 \times 10^{16}) = 6.96 \times 10^{-10}$. This is the same quantity that sets gravitational redshift and clock rate. It is measured with a gravimeter or by comparing two atomic clocks at different heights.
- $\chi\_a = a/(cH_0)$, with a the local acceleration and $cH_0 = 2.998 \times 10^8$ m/s $\times 2.27 \times 10^{-18}$ s^{-1} = 6.80×10^{-10} m/s^2. At the Earth's surface $a = 9.81$ m/s^2, so $\chi\_a = 1.44 \times 10^{10}$. In the outskirts of a galaxy $a \approx 10^{-10}$ m/s^2 and $\chi\_a \approx 0.15$.

Whether the two readings are one χ is open point O4 of Part Two. This article uses χ only as "the state the mesh must fit", and the loop never needs its value.

δ — the local density deviation. Dimensionless. $\delta = (n - n_0)/n_0$, where n is the number of windings that persist inside the mesh and n_0 the number that persist in the same volume without the mesh. In an optical cavity, n and n_0 are mode counts. They are computed in Section 5. In a discharge tube, n and n_0 are set by gas pressure. In a layered chamber, by the sequence of materials.

γ — the lever. Dimensionless. The constitutive law of the net is

$g\_eff = \bar{g}\ (1 + \gamma\ \delta)$.

Here $\bar{g}$ is the background coupling and g_eff the coupling inside the mesh. Part Two derives $\gamma = 2\pi$ from the sign structure of the winding energy $e(n) = g_2n^2 + g_3n^3$: if the n^3 term contributes 2π units where the sum is one unit, the n^2 term is $1 - 2\pi = -5.283$ units. Of the gross stiffness, $(2\pi - 1)/2\pi$ = 84.1% is cancelled and $1/2\pi$ = 15.9% remains. A medium with most of its stiffness cancelled stands near its critical point, and there small δ gives large response.

The empirical anchor for γ is the galactic acceleration threshold. The crossover between the two regimes of g_eff lies at $\gamma\delta = 1$. With $\delta = k\ \chi\_a$ this gives $a_0 = cH_0/(\gamma k)$. The measured threshold is $a_0 \approx 1.2 \times 10^{-10}$ m/s^2 (Milgrom 1983; McGaugh, Lelli and Schombert 2016). So $\gamma k = 6.80/1.2 = 5.7$, within 10% of $2\pi = 6.28$. A physicist who doubts $\gamma = 2\pi$ can take $\gamma k = 5.7$ from the rotation-curve data and carry it through every formula below.

α and β — why the fine-structure constant is not the knob. $\alpha = 1/137.036 = 0.0072974$. The net composes $g = g(\alpha(\chi))$ with $\Delta\alpha/\alpha = \beta\ \Delta\chi$ and $\beta = \alpha/4\pi = 0.0072974/12.566 = 5.81 \times 10^{-4}$. A 1% shift in α needs $\Delta\chi = 0.01/(5.81 \times 10^{-4}) = 17$. The Earth's χ_Φ is 7×10^{-10}. Nothing reachable moves α. All control goes through δ. This also matches the laboratory bound: α is constant to better than one part in 10^{17} per year in atomic-clock comparisons.

L — the mesh size. Metres. The one adjustable dimension of the mesh. In an optical cavity it is the dielectric thickness. In a resonant circuit it is the wavelength set by the tuning capacitor. In a discharge tube it is the electrode gap.

s — the closure signal. Watts. The electrical power the mesh delivers to a load. For a linear current–voltage curve, $s = |I\_sc\ V\_oc|/4$, with I_sc the short-circuit current and V_oc the open-circuit voltage. Both are read on a source meter. This is the only quantity the loop uses.

p — the setting. The current value of L, or of whatever the actuator controls. Metres, farads or hertz, depending on the mesh.

K, η — the loop state and its step. K is a running estimate of s. $\eta = 0.1$ is the learning rate. Both are numbers in a microcontroller.

N, A — count and area. N is the number of mesh elements. A is their total active area in square metres. Output scales with both; that scaling is measured, see Section 4.

Q — the address. The Hopf charge $Q \in \mathbb{Z}$ of a winding, and its trit address. Not continuously adjustable. It changes only when a strand breaks. The machine does not set Q; it sets δ and phase so that a chosen Q is the fixed point. Q plays no role in the loop and is listed for completeness.

3. The loop

The loop replaces the builder's hand. It has three states, named in Part Two as soften, write and harden.

Soften. Release the setting p. Set K to the current reading of s.

Write. Two modes.

Discovery. Large steps over the whole range of p. Record s at each step. This gives the full curve s(p) and locates its maximum. It runs once at start-up and again after a jump in the local state.

Tracking. Small dither around the current p, local slope only, continuous. This holds the maximum. It runs the rest of the time.

In tracking mode, repeat:

1. Move p by one step Δp in the direction of the last measured slope.
2. Read s with a lock-in: superimpose a small periodic dither on p, of amplitude about five times the actuator resolution, and read s phase-sensitively over an integration time of 1 to 10 s. At picowatt levels this is not a refinement but a necessity; the signal-to-noise ratio grows as the square root of the integration time. The lock-in also returns the local slope ds/dp directly.
3. Update $K \leftarrow K + \eta\,(s - K)$.
4. If the slope has changed sign, halve Δp.
5. Stop when $|ds/dp|$ falls below the noise floor, or when Δp is below the actuator's resolution.

Harden. Hold p. Keep reading s at a slow rate. Two thresholds apply. A fast one: if s drops by more than 5% of K, return to soften. A slow one: a running mean over days, to separate ageing of the mesh (10–20% in six months, measured) from a passing fluctuation.

This is a one-dimensional search for the maximum of s(p). Its convergence is simple to bound. The plain version, stepping and halving, reaches the optimum in at most $|p_0 - p^*|/\Delta p$ steps plus $\log_2(\Delta p/\text{resolution})$ halvings. With $\Delta p = 1$ nm, a starting error of 80 nm and a resolution of 0.1 nm, that is 80 + 4 = 84 steps. Discovery mode with 15 points over the range, followed by tracking, needs roughly 20 lock-in readings to lock. Each reading costs seconds to minutes, so the factor of four matters.

The loop moves a voltage V, not a length. The conversion L(V) is part of the model and is specified in Section 4. Alongside s, the controller logs V, the mirror temperature T and, where an interferometer is present, L. A change in s that tracks T or V rather than L is then visible in the log.

Nothing in the loop refers to χ, to the site, to the operator or to the material. That is the sense in which the machine is generic. What it needs is a mesh whose one dimension can be moved by an actuator, and a load across which s can be read.

4. The actuator

The loop needs a mesh whose one dimension moves. The measured stack of Section 5 is fixed after fabrication. This section states what moves L, with the arithmetic that rules out the cheapest option.

Thermal tuning, ruled out. The fitting condition is $\lambda_{max} = 2nL$. Heating changes both n and L. For PMMA, $dn/dT \approx -1 \times 10^{-4}$ per kelvin and the linear expansion coefficient is about 7×10^{-5} per kelvin. At L = 33 nm:

$d\lambda_{max}/dT = 2L\cdot(dn/dT) + 2n\cdot L\cdot(7 \times 10^{-5}) = 66\ \text{nm} \times 10^{-4} + 100\ \text{nm} \times 7 \times 10^{-5} = 0.0066 + 0.007 \approx 0.014$ nm/K.

To shift λ_{max} by 1 nm, the equivalent of moving L by 0.33 nm, takes about 70 K. To cover a search range of 80 nm in L takes thousands of kelvin. Thermal tuning cannot reach the resolution the loop needs. It is listed here so that no one has to rediscover this.

Electrostatic MEMS mirror, chosen. Replace the fixed aluminium mirror by a mirror on a spring, driven by a voltage between mirror and base. This is the Fabry–Pérot tunable filter of optical telecommunications, made in volume since the 1990s. Resolution 0.1 to 1 nm; range hundreds of nanometres; response time microseconds. The cavity is then air or vacuum instead of PMMA, so n = 1 and $\lambda_{max} = 2L$. The upper electrode and insulator of the stack stay as they are. The mirror voltage is the setting p.

Electro-optic layer, kept as alternative. A layer whose index n changes with an applied field (a liquid-crystal film or a $BaTiO_3$ thin film) tunes $\lambda_{max} = 2nL$ without moving anything. Sub-nanometre equivalent resolution. Range smaller than the MEMS mirror. Suitable once the working L is known and only drift correction is needed.

L(V), the conversion the loop needs. The loop sets a voltage V. The mirror moves by electrostatic force against its spring:

$L(V) = L_0 - \Delta L(V),$

with L_0 the rest gap and ΔL growing roughly as V^2 until pull-in, the voltage at which the spring no longer holds and the mirror snaps down. The usable range is the part of L(V) below pull-in, in practice up to one third of L_0. For a rest gap of 300 nm that is 100 nm of travel, more than the 80 nm search range of Section 3. The characterisation of one prototype records: L_0, the pull-in voltage, $\Delta L(V)$ up and down (hysteresis), the settling time after a step, the smallest voltage step that gives a resolvable change in L, the repeatability of L at a fixed V, and the shift of L per kelvin. These are standard measurements for any MEMS mirror and take one afternoon with an interferometer.

The MEMS mirror is the first prototype's actuator. Its voltage is what the loop moves, and its position is what a physicist reads off with an interferometer to confirm L.

5. The one reported yield

The only published mesh with an area-scaled power measurement is the optical-cavity device of Moddel, Weerakkody, Doroski and Bartusiak (2021). The stack, from bottom to top (fixed mirror; the actuator of Section 4 replaces it in the prototype):

Layer	Material	Thickness
Base electrode	nickel	38–50 nm
Insulator	native NiO_x + Al_2O_3	≈ 1 nm + 0.7–2.3 nm
Upper electrode	palladium	8–24 nm; output peaks near 12 nm
Cavity dielectric	PMMA (n = 1.52) or SiO_2 (n = 1.49)	11–1100 nm
Mirror	aluminium	150 nm

Total thickness of one element at a 33 nm cavity: 50 + 3 + 12 + 33 + 150 = 248 nm.

Reported: at 33 nm cavity thickness the maximum power is 1.4 pW from an active area of 0.02 μm^2. Recomputed here: 0.02 $\mu m^2 = 2 \times 10^{-14}$ m^2, so 1.4×10^{-12} W / 2×10^{-14} m^2 = 70 W/m^2. This is a published value, recomputed, not independently replicated.

Reported scaling: short-circuit current is linear in active area up to 10,000 μm^2. In 4 × 4 arrays, current in parallel and voltage in series are each four times the single-element value. Output appears only after the mirror is deposited; the same stack without mirror gives no current. Output decreases monotonically with increasing cavity thickness. After six months the output had degraded by less than 10% (PMMA) and less than 20% (SiO_2).

Three facts from this data feed the design:

1. The mesh dimension that controls output is the cavity thickness L. That is the dimension the loop should move.
2. The asymmetry (mirror on one side only) is what produces the passage. A symmetric structure gives nothing.
3. The material sets the yield. 70 W/m^2 is a property of nickel, palladium, PMMA and aluminium at these thicknesses, not of the net in general.

Sizing. A Dutch house draws on average about 340 W of electricity and about 1,100 W of heat, together 1,450 W. At 70 W/m^2 that is 1,450/70 = 20.7 m^2, one layer. Halving heat loss by insulation gives 890 W and 12.7 m^2. If layers of 248 nm can be stacked, 100 layers are 25 μm thick and give 7,000 W/m^2; the house then needs 0.21 m^2. This is extrapolation. The hypothesis is $P_N = N \cdot P_1$. The two-layer measurement of Section 8 has three possible outcomes: $P_2 < 2P_1$, the layers interact and the yield saturates; $P_2 \approx 2P_1$, the layers are independent and the extrapolation holds; $P_2 > 2P_1$, the layers couple cooperatively. Each outcome is a number, and each sets the design of the stack.

Voltage. One element delivers microvolts. A usable voltage needs elements in series. With 4 × 4 arrays measured, the series scaling is known; a 1,000-element series string is a fabrication choice, not extra area.

6. The missing step: from mesh to δ

Part Two names as open point O2 the rule that turns a mesh geometry into a δ. Without it, the design copies a measured mesh instead of deriving one. Here the rule is written as a counting procedure. It is the second axiom of the net applied literally: a winding persists only if it closes on itself inside the mesh.

Step 1. Which windings fit. Take a mesh with two reflecting walls a distance L apart, filled with a medium of refractive index n. A winding of wavelength λ returns in phase after one round trip if $2nL = m\lambda$, $m = 1, 2, 3, \ldots$ The longest wavelength that fits is $\lambda\_max = 2nL$. For Moddel's 33 nm PMMA cavity: $\lambda\_max = 2 \times 1.52 \times 33$ nm = 100 nm.

Step 2. Count, per octave. Let λ_min be the shortest wavelength the walls still reflect; for aluminium that is in the near ultraviolet, of order 100 nm. The number of fitting longitudinal windings is $m\_max = 2nL/\lambda\_min$. For the 33 nm PMMA cavity, $m\_max = 100/100 = 1$: exactly one winding fits. The counting rule therefore applies from $L = \lambda\_min/2n = 33$ nm upward; Moddel's best point sits on that boundary.

δ cannot be counted over the whole band $[\lambda\_min, \infty)$, because that band has no finite reference and the removed fraction then depends entirely on where the band is cut. The count is made per octave instead. For each scale λ, count the windings in $[\lambda, 2\lambda]$ with the mesh, $n(L, \lambda)$, and without it, $n_0(\lambda)$. Then

$$\delta(L, \lambda) = (n(L, \lambda) - n_0(\lambda)) / n_0(\lambda).$$

For $\lambda < 2nL$ the octave is untouched and $\delta = 0$. For $\lambda > 2nL$ the octave is emptied and $\delta = -1$. The deviation is therefore a step at $\lambda = 2nL$. What the mesh does is move that step: a smaller L pushes it to shorter wavelengths. Weighted by winding energy, which rises as the inverse cube of wavelength, the total weight removed grows as L^{-3}. That is the shape prediction: output rises as L falls, steeper than linear.

Step 3. Asymmetry. With a mirror on one side only, δ is set on that side and not on the other. The difference $\delta\_mirror - \delta\_open$ is the passage. The device of Section 4 measures exactly that difference as a current.

Step 4. Response. The constitutive law gives the change in coupling in each octave: $\Delta g/g = \gamma\, \delta(L, \lambda) = 2\pi\, \delta(L, \lambda)$. In the emptied octaves the coupling shift is the full lever; in the untouched octaves it is zero. The net responds to a deviation with 2π times that deviation. This is the second half of O2: how χ moves when δ is imposed. No single number for δ is quoted here, because the counting rule fixes the shape of the response, not its magnitude; the magnitude passes through the material factors of Step 5.

Step 5. From δ to s. Between the counted δ and the measured s stands a function

$$s(L) = F[\delta(L, \lambda), T\_Pd(\lambda), T\_ins(\lambda), R\_mirror(\lambda), A, N],$$

with T_Pd the transmission of the upper electrode, T_ins the tunnelling through the insulator, R_mirror the reflection band of the mirror, A the area and N the count. The form of F is open. O2 therefore has three parts. O2a: the geometric count, mesh $\rightarrow \delta(L, \lambda)$, written above. O2b: $\delta \rightarrow$ coupling, $\Delta g/g = 2\pi\, \delta$, Step 4. O2c: coupling $\rightarrow$ electrical output, the function F. Only O2a and O2b are written. O2c is where a mode count becomes watts, and it is material, not net.

Step 6. Check. The rule predicts that output rises monotonically as L falls, with a curve steeper than linear. Moddel's Figure 4b shows a monotonic rise as L falls from 1100 nm to 33 nm. That is one confirmed trend, not a fit. Until F is written, the rule is a shape predictor. Measurement 0 of Section 8 records s(L) point by point on one element, open-loop, and that curve is what F is then fitted to.

The counting in Steps 1–2 is what Casimir (1948) did for two perfect mirrors. The net rule adds nothing to that arithmetic. What it adds is the reading: the removed windings are not "missing energy"; they are strands that no longer close inside the mesh, and the passage is the net restoring closure across the asymmetry.

What Claude Code runs. On the lattice of the net kernel: impose a two-wall mesh of size L, count the windings that close, count without the mesh, tabulate δ(L) for L from 10 to 1000 lattice units, then repeat with one wall removed to obtain the asymmetric difference. Output: a table δ(L) and the difference curve. The attestation is whether the table reproduces the monotonic trend of Section 4.

7. Independence, stated exactly

Of	The loop	The mesh
Place	Independent: never reads χ	Independent once tuned by the loop
Time	Independent: re-enters soften on drift	Ages; measured 10–20% in six months
Person	Independent: no hand tuning	Independent
Material	Independent: reads only s	Dependent: material sets yield

The design is material-independent. The yield is not. This is the split between invariant and state that runs through the whole Vacuum.Net work: the loop is the invariant, the mesh the state.

8. Status ledger

- **Reported and recomputed, not independently replicated.** 70 W/m² at 33 nm cavity (Moddel et al. 2021). Linear area scaling to 10,000 μm². Series and parallel scaling in 4 × 4 arrays. Monotonic rise of output with decreasing L.
- **Measured elsewhere, used as calibration.** γk = 5.7 from a_0 (Milgrom; McGaugh et al.).
- **Derived within Vacuum.Net.** γ = 2π. The agreement between 6.28 and 5.7 is a consistency check between derivation and calibration, not an independent confirmation of the constitutive law.
- **Derived.** $\beta = \alpha/4\pi = 5.81 \times 10^{-4}$. α stiff, δ soft. δ(L, λ) as a counting rule (O2a) and Δg/g = 2π δ (O2b), Section 6.
- **Designed, not built.** The loop of Section 3, with lock-in reading and parabolic interpolation: 84 steps plain, 15–25 interpolated. The MEMS-mirror actuator of Section 4. Thermal tuning computed at 0.014 nm/K and ruled out.
- **Extrapolated, not measured.** Linear scaling beyond 0.01 mm². $P_N = N \cdot P_1$ for stacking in depth.
- **Open.** O2c: the function F from coupling to electrical output. L(V) of the prototype mirror, until characterised. O4: identity of χ_Φ and χ_a. Note that γk = 5.7 is calibrated on χ_a alone; if χ_Φ and χ_a are not one χ, that calibration hangs on χ_a only.

Measurement order. Claude Code runs the lattice count of Section 6; the physical measurements below are for whoever fabricates the prototype, and the author attests each result. 0. Characterise the mirror: L(V), pull-in, hysteresis, settling time, repeatability, temperature shift (Section 4).

1. One element, open-loop: set L by hand at 10, 15, 20, 25, 30, 33, 40, 50, 75, 100, 150, 250, 500, 750 and 1000 nm, and record s at each. This is the curve s(L). It is taken before the

loop is switched on, so that the loop later has a known maximum to find. It also fixes F of Section 6. An identical element with a fixed mirror is measured alongside as reference.

2. The same element, loop active: start from a random p_0 and record the settling of p, s and the time to lock. The loop has found the maximum of step 1, or it has not.
3. A 100-element series string: millivolt range. This settles interconnect resistance and capacitive coupling before the 1,000-element string.
4. A two-layer stack: P_2 against $2P_1$, with the three outcomes of Section 5.
5. The same prototype at a second site, in another orientation and temperature, started from a random p_0: record L*, s* and the time to lock at both sites. The spread between sites is the measured meaning of "independent of place" in Section 7.
6. Six months of drift with the loop active: this is the test of "independent of time" in Section 7.

9. Prototype specification

Build and measurement parameters only.

Parameter	Prototype
Cavity	air or vacuum, n = 1
Adjustable dimension	L, by electrostatic MEMS mirror
Rest gap L_0	≈ 300 nm
L range	≥ 80 nm, below pull-in
L resolution	≤ 0.1 nm
Actuator variable	V
Feedback variable	s = |I_sc V_oc|/4
Logged alongside s	V, T, L (interferometer)
Sensing	lock-in on a dither of p, 1–10 s integration
Modes	discovery (full curve), tracking (local slope)
First active area	one element, 0.02 μm^2 to 10,000 μm^2
Reference	identical element with fixed mirror
First scaling test	two layers
Second scaling test	100 elements in series
Site test	same prototype at two sites
Stability test	six months, loop active

10. What a physicist can do with this article

Recompute χ_Φ from GM, R and c. Recompute β from α. Recompute 70 W/m² from 1.4 pW and 0.02 μm^2. Recompute γk from cH_0 and a_0. Count the modes in a 33 nm PMMA cavity and find m_max = 1. Recompute the thermal tuning rate of 0.014 nm/K. Bound the loop's convergence from Δp and the starting error. None of these steps needs a reference to be opened. The references say where the measured numbers come from.

Revision note

Second edition, 25 September 2026. After a review that recomputed every quantity and found no arithmetic error, four changes. Section 3 adds lock-in reading of s and parabolic interpolation. Section 4 is new: the actuator, with thermal tuning computed and ruled out and an electrostatic MEMS mirror chosen. Section 6 replaces the whole-band count by a per-octave count, drops the example value $\delta = -0.01$, which did not follow from the rule, and states that the rule predicts shape, not magnitude. Section 8 adds the measurement order and the note that γk is calibrated on χ_a alone.

Third edition, 25 September 2026. After a second review. The 70 W/m^2 is labelled throughout as reported and recomputed, not replicated. Section 3 replaces parabolic interpolation by a discovery mode and a tracking mode, and logs V, T and L alongside s. Section 4 adds L(V) and the mirror characterisation. Section 5 states stacking as the hypothesis $P_N = N \cdot P_1$ with three outcomes. Section 6 writes the function F and splits O2 into O2a, O2b and O2c. Section 8 separates $\gamma = 2\pi$ (derivation) from $\gamma k = 5.7$ (calibration), adds the open-loop s(L) curve as the first measurement and the two-site measurement as the fifth. Section 9 is new: the prototype specification.

Annotated references

Casimir, H. B. G. (1948). "On the attraction between two perfectly conducting plates." Proceedings KNAW 51, 793. Why read: the original mode count between two mirrors. Section 5 uses the same arithmetic. Read it for the counting, not for the interpretation.

Konstapel, J. (2026). "Applied Vacuum Theory: Setting the State of the Net." Fourth edition, Constable Research, Leiden. Why read: Part Two. Defines χ, δ, γ, β and the open points O1–O4 that this article builds on. Section 4 there derives $\gamma = 2\pi$ step by step.

Konstapel, J. (2026). "Builders of the Net." Constable Research, Leiden. Why read: Part One. Reads Moray, Schauberger, Tesla, the Correas, Reich and Moddel–Thibado–White as meshes fitted to a place. This article's design answers the question that Part One raises: why each device stopped when moved.

McGaugh, S. S., Lelli, F. and Schombert, J. M. (2016). "Radial acceleration relation in rotationally supported galaxies." Physical Review Letters 117, 201101. Why read: the cleanest measurement of the acceleration threshold a_0 across 153 galaxies. Gives the number from which $\gamma k = 5.7$ follows.

Milgrom, M. (1983). "A modification of the Newtonian dynamics as a possible alternative to the hidden mass hypothesis." Astrophysical Journal 270, 365. Why read: first statement of $a_0 \approx 1.2 \times 10^{-10}$ m/s^2. Read only the first section; the rest is his own framework, not needed here.

Moddel, G., Weerakkody, A., Doroski, D. and Bartusiak, D. (2021). "Optical-Cavity-Induced Current." Symmetry 13, 517. Preprint arXiv:2101.03085. Why read: the only published mesh with power, area scaling, array scaling and eight artifact tests. Table 1 gives every layer thickness. Figure 4 gives the 1.4 pW and the thickness trend. Read in full; it is the empirical floor of this article.

Thibado, P. M., Kumar, P., Singh, S., Ruiz-Garcia, M., Lasanta, A. and Bonilla, L. L. (2020). "Fluctuation-induced current from freestanding graphene." Physical Review E 102, 042101.

Preprint arXiv:2002.09947. Why read: the second mesh type, a vibrating membrane with two diodes. Yield is far lower than Moddel's, but the circuit is the model for a loop that reads s across a load.

Markley, F. L., Cheng, Y., Crassidis, J. L. and Oshman, Y. (2007). "Averaging quaternions." Journal of Guidance, Control, and Dynamics 30, 1193. Why read: the eigenvalue definition of the quaternionic order parameter R_Q used in Part Two for phase coherence. Not needed for the loop of this article, which uses only s; listed so that the phase variables of Part Two remain traceable.