OSH GLOBAL SYSTEMS, INC.

OSH: Physical Validation of an Autonomous Homeostatic Cognitive Architecture

From CPU-resident cognition to causal parity, external timing, and closed-loop execution on ESP32-S3

José Fabián Vallejos

Founder & Researcher, OSH Global Systems, Inc.

fabian.vallejos@oshsystems.com

Technical Research Paper

21 September 2026

Abstract

This paper reports a staged experimental validation of OSH, a proposed autonomous homeostatic cognitive architecture designed to execute as a resident dynamical mechanism rather than as a token-predictive language model. The validated implementation maintains an endogenous state, evolves trajectories, regulates attention and selection, binds selected transitions to compact semantic structure, and exposes language/phonation as downstream expression rather than as computational authority. Proprietary equations, coefficients, internal geometry, source code and raw knowledge-store contents are intentionally withheld.

Evidence spans CPU-only host benchmarks, cross-target compilation portability across unrelated embedded ISA families, native execution on two physical ESP32-S3 devices, five-level PC–MCU causal parity, endurance, independent host rebuilds, oscilloscope timing and a final physical closed loop. Step19 produced 150/150 parity checks across state, trajectory/vortex, homeostatic attention, selected edge identity and meaning. The final closed-loop experiment used ADC GPIO3 as physical reference, PWM GPIO14, an external $1\,\mathrm{k}\Omega/100\,\mu\mathrm{F}$ network and ADC GPIO2 feedback. It completed 100,000 cycles with $\mathrm{INC} = 5160$, $\mathrm{HOLD} = 93358$, $\mathrm{DEC} = 1482$, $\mathrm{ADC\_FAIL} = 0$ and $\mathrm{STATE\_FAIL} = 0$; the physical gate and final latch passed. Internal route timing was 24–30 µs, mean 29.525 µs. Independent oscilloscope captures preserved with this paper show a representative route pulse of 23.0 µs and a separately instrumented action/selection interval of 660 ns.

The supported conclusion is deliberately narrower than a claim of general intelligence: a native, repeatable, instrumentable, locally executable computational mechanism exists on physical microcontroller silicon and preserves the validated causal structure without requiring cloud inference. These data do not establish consciousness, biological life, or superiority in reasoning or language quality over large language models.

Keywords: homeostatic cognition; embedded cognition; dynamical systems; autonomous systems; ESP32-S3; causal parity; closed-loop control; non-Transformer architecture; reproducibility.

1 Introduction

Large language models demonstrate the utility of large-scale statistical sequence modeling. OSH investigates a different engineering question: whether a persistent cognitive mechanism can be organized around endogenous regulation, state evolution, trajectory, selective pressure, continuity and feedback while remaining small enough to execute natively on a microcontroller.

This is not an argument that language models do not work. It is an experimental report about a different computational architecture. The questions are whether the same frozen mechanism can be observed on a host CPU, reproduced on physical silicon, causally matched across implementations, sustained over long runs, instrumented externally and coupled to a physical feedback loop. The contribution is an evidence ladder rather than a single benchmark number.

2 What OSH Is

2.1 Operational definition

OSH denotes an *Organismo SemÃ¡ntico HomeostÃ¡tico* (Homeostatic Semantic Organism): a proposed cognitive architecture in which the active computational object is a continuously regulated endogenous state and its trajectory through structured state space. "Organism" is an architectural term for persistent internal regulation and coordinated functional subsystems; it is not a biological-life claim.

At the validated kernel boundary the system exposes a compact endogenous state including variables denoted C, P, X, T, H. Their proprietary update equations and coefficients are not disclosed. A public abstraction is

$$\mathbf{s}_{t+1} = \mathcal{F}(\mathbf{s}_t, \mathbf{u}_t, \mathcal{M}_t), \tag{1}$$

where $\mathbf{s}_t$ is endogenous state, $\mathbf{u}_t$ is perturbation/input and $\mathcal{M}_t$ is persistent contextual structure. A downstream abstraction is

$$a_t = \mathcal{S}(\mathcal{H}(\mathcal{V}(\mathbf{s}_{t+1}), \mathcal{G}_t)), \tag{2}$$

where $\mathcal{V}$ denotes trajectory/vortex, $\mathcal{H}$ homeostatic attention, $\mathcal{G}_t$ compact semantic structure and $\mathcal{S}$ selection. These equations expose tested functional boundaries only.

2.2 What OSH is not

The validated kernel is not presented as a large language model, Transformer, prompt wrapper, remote cognitive API, or agent whose cognitive authority is delegated to a foundation model. Nor is this paper a claim of superior intelligence or language quality. The physical results matter because the validated route executes locally on ESP32-S3 silicon.

2.3 Autonomy and cloud independence

"Autonomous" has an operational meaning: the validated kernel can evolve state, execute trajectory, perform homeostatic attention and selection, and drive a physical action/feedback path without cloud inference. Network access may supply information in applications, but it is not required to instantiate the validated cognitive route. Internet is an acquisition boundary, not cognitive authority.

3 Cognitive Organization

3.1 Endogenous state and homeostasis

External or internal perturbations displace a persistent internal state; the architecture evolves a bounded response rather than treating each input as an isolated stateless query. Homeostasis is the organizing constraint under which state evolution, attention and action are evaluated. In the final physical experiment the observable policy exercised INC, HOLD and DEC under real analog perturbation and feedback.

3.2 Trajectory, vortex and HAL

A state endpoint alone does not specify a dynamical process. OSH preserves trajectory information. The "vortex" label denotes a dynamical stage; Step19 separately tests settled-state and vortex-path parity. The Homeostatic Attention Loop (HAL) evaluates candidate transitions in relation to endogenous condition, and Step19 treats HAL candidates as an independent causal level.

3.3 Selector, compact graph and continuity

Candidate evaluation feeds a selector over compact semantic structure. Step19 checks selected edge identities before meaning, so causal selection must be preserved rather than merely similar surface text. The broader architecture also includes continuity/scar and persistence components intended to preserve consequences of prior state and use; their implementation remains black-box protected but artifact hashes are published.

4 How OSH Produces Cognition and Speech

The validated route can be summarized as

$$\text{perturbation} \rightarrow \text{state} \rightarrow \text{trajectory} \rightarrow \text{HAL} \rightarrow \text{selection} \rightarrow \text{meaning/action}. \tag{3}$$

Surface language is downstream. Text is not used as proof that cognition occurred; the internal route is tested independently.

The language path converts selected semantic structure into surface form; phonation renders that form. Step21 tested authority boundaries: VOICE OFF passed 30/30, VOICE ON passed 30/30, 10,030 phonation events were submitted/rendered, zero were dropped, and meaning/selection digests remained invariant. This is a host boundary qualification, not physical I2S speaker validation.

Step22 similarly qualifies acquisition: valid ingest/quarantine, activation, pre-activation invariance, post-activation benchmark invariance, 30/30 offline local parity and corrupt-payload rejection were observed, with no direct response authority and no selector/surface network

dependency in the tested boundary. The evidence does not include a physical live Wi-Fi/HTTP run.

5 Evidence Ladder

Table 1: Evidence progression. Metrics from different workloads are deliberately kept separate.

Stage	Measurement	Result
PC runtime	CPU-only resident-runtime protocol on Intel Atom D2500	measured
Cross-target portability	Freestanding compact C substrate across ARM Cortex-M, 32-bit RISC-V, AVR and MIPS32 profiles	compile PASS
Tier-A silicon	Native float32 trajectories on physical ESP32-S3	**PASS**
Cross-device	Two physical ESP32-S3 units reproduce Tier-A	**PASS**
Step19	Five causal levels, 30 cases per level	150/150 **PASS**
Step20	Physical endurance	100,000 cycles; 0 cognitive failures
Step21	Voice/phonation host boundary	30/30 OFF + 30/30 ON; 10,030 events
Step22	Acquisition host boundary	30/30 offline parity; corrupt payload rejected
External timing	Oscilloscope GPIO instrumentation	route + action captured
Closed loop	Analog reference $\rightarrow$ OSH $\rightarrow$ PWM/RC feedback	100,000 cycles; **PASS**
Integrity	SHA-256 native modules, build, firmware and evidence set	locked

6 Benchmark 3: Host Runtime

The preserved CPU-only protocol compared OSH Resident Runtime V1 with TinyLlama 1.1B Q4_K_M. Median latency was 6.466 ms versus 18,186.463 ms, a 2,812.69× gap under that protocol. Working sets were 24.883 MB versus 662.770 MB, a 26.64× ratio. Both used GPU=0; TinyLlama was constrained to exactly one generated token per input. A separate recorded run reported 5.537 ms and a 3,349.99× ratio. These are runtime/resident-memory measurements, not an intelligence or language-quality comparison.

7 Cross-Architecture Portability Before Physical Silicon Validation

Before the physical ESP32-S3 validation, the compact embedded OSH C substrate was subjected to a separate cross-target portability assay. The source unit was compiled with Clang 17 in freestanding mode using `-Oz`; the tested unit required no network, operating system, Python runtime, dynamic allocation or application framework. This assay addresses a different question from physical execution: whether the disclosed compact substrate is structurally tied to a single host or instruction-set family.

The assay covered the reported cross-target profile set spanning unrelated embedded ISAs; Table 2 reproduces the target/profile rows preserved in the benchmark. The result is deliberately labeled *compile portability*; it is not promoted to physical execution on boards that were not physically tested. This stage is compiler-level evidence. No emulator result is used here as a substitute for physical execution.

Table 2: Cross-target compilation portability of the compact embedded C substrate. “PASS” in this table means successful target compilation, not physical-board execution.

Target profile	Representative family / interpretation	ISA / float profile	.text bytes
Cortex-M0	STM32F0 / RP2040-class	ARMv6-M / software FP	952
Cortex-M3	STM32F1-class	ARMv7-M / software FP	1,018
Cortex-M4	STM32F4 / nRF52-class	ARMv7E-M	820
Cortex-M7	STM32H7 / NXP RT-class	ARMv7E-M hard-float	752
Cortex-M33	STM32U5 / modern secure MCU	ARMv8-M Mainline	764
RV32IM	generic 32-bit RISC-V	integer + multiply/divide	1,728
RV32IMF	RISC-V MCU with F extension	single-precision float ABI	960
RV32IMC/IMAC	ESP32-C3/C6-class architectural proxy	compressed 32-bit RISC-V	1,310
AVR	ATmega328P	AVR / software float	2,108
MIPS32R2	legacy embedded MIPS	MIPS32 little-endian	1,192

The portability assay supports the proposition that the tested compact substrate is not tied to x86 or the ESP32-S3 execution environment. It does *not* establish physical timing, peripheral behavior, power behavior, memory behavior or full production integration on ARM, RISC-V, AVR or MIPS boards. Those remain separate silicon-validation questions. This distinction is essential: cross-compilation, emulated/virtual execution and physical-silicon execution are different evidence levels and must not be collapsed into one PASS label.

The portability-assay source fingerprint is:

`408deada3466d53f51354a3ae2c4ece53e2bf6dc1f4eb0c1bd7c692d179aa999`

This result strengthens the architecture-level portability thesis while preserving the strict evidence boundary that, in the supplied evidence set, ESP32-S3 is the family carrying physical-silicon validation.

8 Native ESP32-S3 Evidence

The validated target was ESP32-S3 rev. v0.2; the preserved dossier observed two CPU cores, 32 MB Flash, 16 MB Octal PSRAM and ESP-IDF v5.4.4. Native Tier-A trajectories passed at candidate scales 10/64/256/1024. Maximum absolute error against the golden trajectory was $5.96046448 \times 10^{-8}$ for OSH_PRESENCE and $2.98023224 \times 10^{-8}$ for OSH_NO_ACTIVATION, both inside the 10^{-6} gate. Two physical units produced 29.4728 and 29.4721 µs/cycle, a recorded 0.0024% cross-device delta.

9 Step19: Five-Level Causal Parity

Thirty cases were checked at five boundaries: $C/P/X/T/H$ state parity 30/30; settle+vortex path 30/30; HAL candidates 30/30; selected edge IDs 30/30; meaning 30/30. The combined result is 150/150 with zero authority-boundary failures. This is stronger than surface-output agreement because it checks intermediate causal identity.

10 Step20 and Independent Revalidation

The preserved physical Step20 run completed 100,000 cycles with zero cognitive failures. Internal free memory remained 368,691 B and PSRAM free 16,774,624 B at recorded checkpoints. Stack HWM moved from 488 to 440 B without cumulative growth. A percentage-display anomaly was traced to 32-bit telemetry arithmetic; absolute memory did not decline.

Step19, Step21 and Step22 were also independently reconstructed and rerun under Linux x86_64/GCC 14.2.0 with UBSAN. All compiled, ran and passed UBSAN. Across 100 executions per artifact each produced one output SHA-256.

11 Physical Closed-Loop Experiment

The final experiment couples OSH to a real analog perturbation and feedback path:

$$\text{B25K} \rightarrow \text{GPIO3 ADC} \rightarrow \text{OSH} \rightarrow \text{GPIO14 PWM} \rightarrow (1\,\text{k}\Omega, 100\,\mu\text{F}) \rightarrow \text{GPIO2 ADC}. \quad (4)$$

Reference was GPIO3/ADC1_CH2, feedback GPIO2/ADC1_CH1 and PWM GPIO14. Native Core, HAL, Vortex, Selector and Compact Graph files were unchanged and hash-locked.

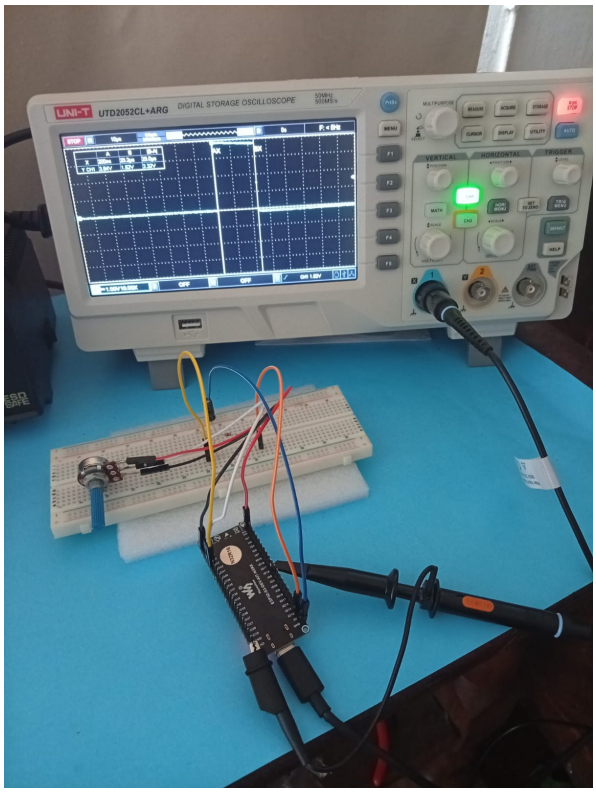

(a) Physical apparatus during route instrumentation.

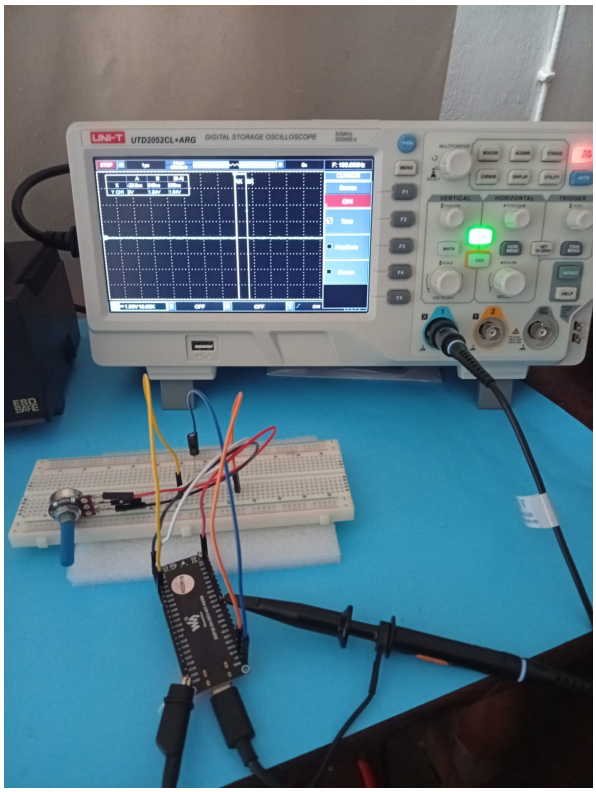

(b) Independent apparatus view during action instrumentation.

Figure 1: Physical ESP32-S3, B25K/RC network and oscilloscope. Photographs document bench context; numeric gates remain primary evidence.

11.1 Final 100,000-cycle result

The run first re-executed Step19 and returned all five 30/30 levels. The gates `STACK_RUNTIME_GATE`, `PC_ESP32_PARITY_GATE` and `OVERALL` all returned PASS. Physical hardware initialization then passed. At cycle 10, `REF_RAW=4095` and `FB_RAW=1295` produced INC. By cycle 10,000 all three regimes had been exercised. The terminal result was:

Table 3: Final physical closed-loop result.

Metric	Result
Cycles	100,000
INC	5,160
HOLD	93,358
DEC	1,482
ADC failures	0
State failures	0
IAE	10,862.551863
Max state delta	0.141184
Route min	24 µs
Route max	30 µs
Route mean	29.525 µs
Elapsed	1000.005270 s
DONE GPIO15	1
FAIL GPIO4	0
PASS GPIO5	1

The firmware emitted PHYSICAL_100K_GATE=PASS and FINAL_LATCH=PASS. Thus the run simultaneously satisfied endurance, ADC validity, state validity and non-zero occupancy of all three observable policy regimes.

12 External Oscilloscope Instrumentation

Software timers are useful but not sufficient as the only timing evidence. GPIO instrumentation was therefore measured with a UNI-T UTD2052CL+ARG digital storage oscilloscope using a 10X probe.

12.1 Route pulse

A preserved route capture used a 10 µs/division timebase and time cursors. The displayed cursor values are A=200 ns and B=23.2 µs, yielding $|B - A| = 23.0$ µs. This is a representative externally observed route pulse. It should not be confused with the 29.525 µs mean over the final 100,000-cycle internal timing run; they are different observations under different captures.

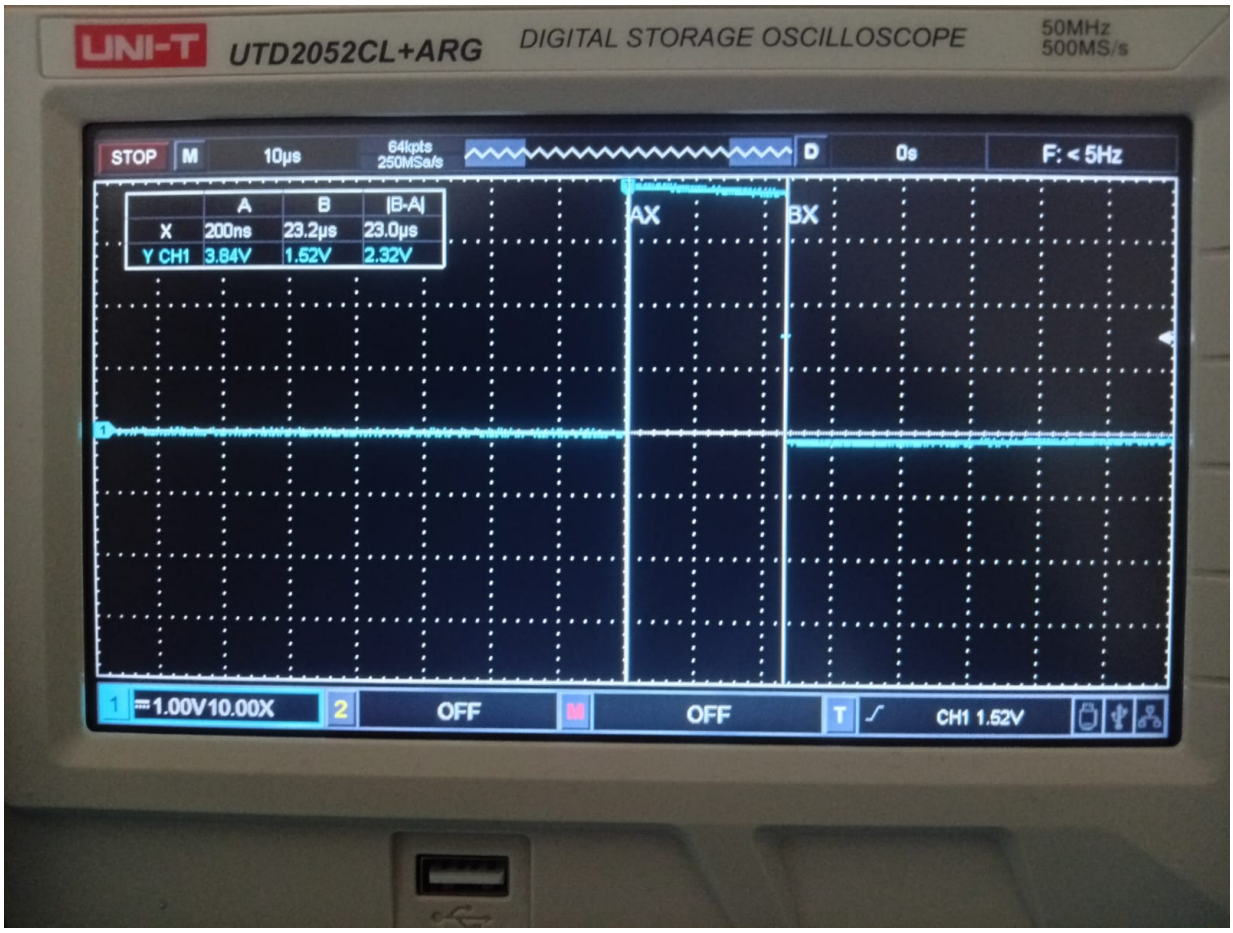

Figure 2: External route timing. The oscilloscope display reports $|B - A| = 23.0$ µs at 10 µs/division.

12.2 Action/selection interval

For action timing, the outer route electrical marker was disabled in a separate instrumentation build and the action marker was mapped to the known-good GPIO18 measurement path. The cognitive math and native modules were not changed. At 1 µs/division, the cursor display reports A=-20 ns, B=640 ns and $|B - A| = 660$ ns.

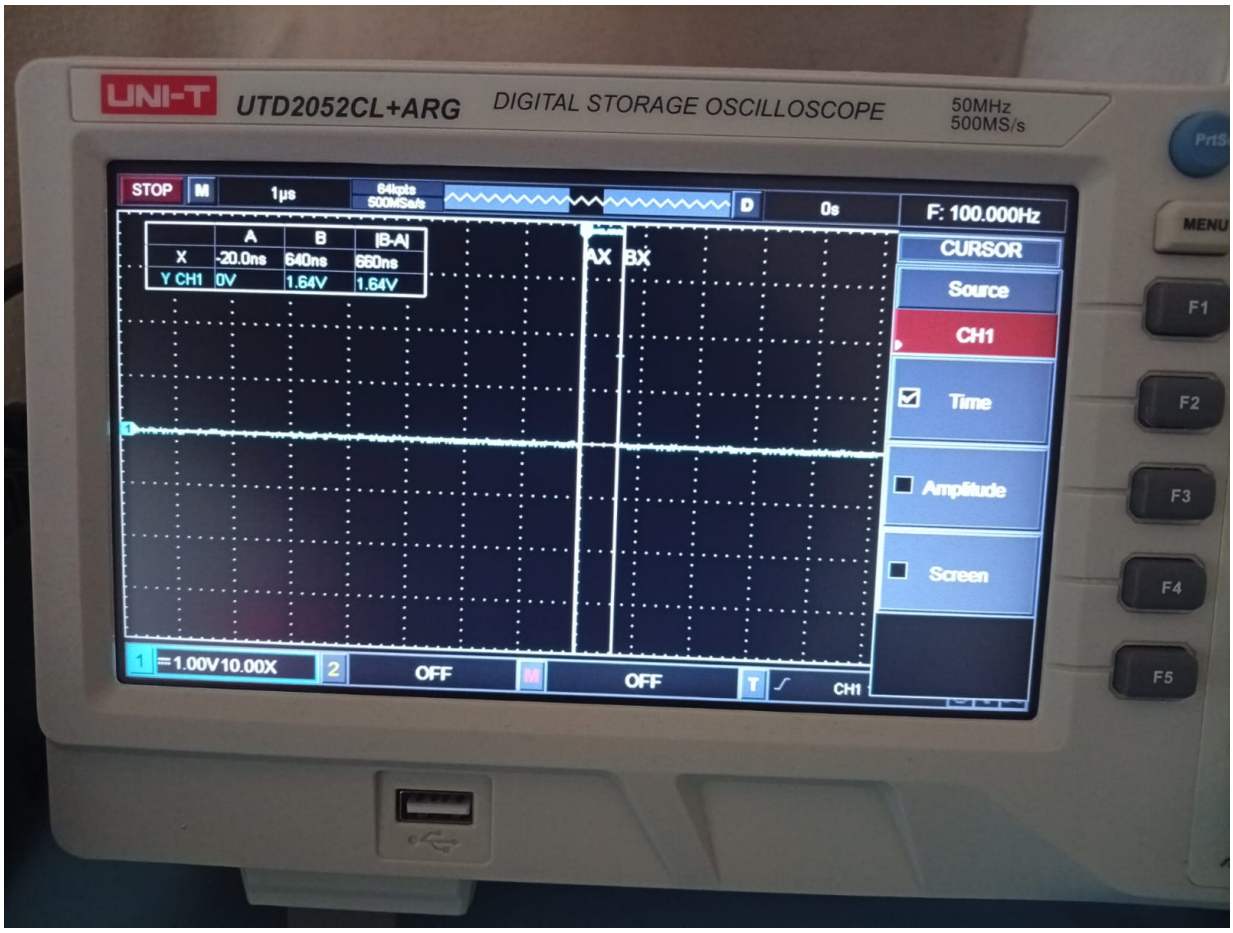

Figure 3: Externally measured action/selection interval: 660 ns.

Table 4: Timing evidence kept by workload and measurement method.

Measurement	Method/scope	Result
Tier-A unit 1	physical ESP32-S3 Tier-A workload	29.4728 µs/cycle
Tier-A unit 2	second physical ESP32-S3	29.4721 µs/cycle
Final closed-loop route	internal timer, 100,000-cycle physical run	mean 29.525 µs; 24–30 µs
Representative route pulse	external oscilloscope, GPIO marker	23.0 µs
Action/selection	external oscilloscope, dedicated marker	660 ns

13 Why the Physical Result Matters

The final closed loop closes several alternative explanations. The reference is a physical analog quantity, not a hard-coded test vector. The response leaves the MCU as PWM, passes through an external RC network and returns through a different ADC input. INC, HOLD and DEC are therefore observed while the system is coupled to an external physical process. At the same time,

Step19 parity is rerun before the physical loop, tying the experiment to the same validated causal route.

This does not prove every proposed future capability of OSH. It does demonstrate that the tested mechanism is more than a web interface, API call, prerecorded output or cloud-dependent inference path.

14 Artifact Integrity and Black-Box Protection

Reproducibility and intellectual-property protection are treated as compatible goals. The public evidence discloses hashes and measured gates while withholding equations, coefficients, internal geometry and source.

The native module hashes are:

Table 5: Frozen native module SHA-256 fingerprints.

Artifact	SHA-256
Core	`e09f1d7928f291a64ab5c8a8d6f9af62cb2b293a1022e6897044fd04288fd363`
HAL	`da9fa7a321a6bf1df797054f9dcac132349930eef8aa2ba11e3c078a188f422f`
Vortex	`0c405d2e86ec9dfbb54b03363d984bb8c745ab84cfdafdab94205351a7e73d78`
Selector	`26c4971a12209460ab74e188c4c234f5bd80db17ac15167ca3802faecd2b630b`
Compact graph	`fe5dcaa69ec4cc1d9a111a4ad8cd60057ddd4a9bbf63984167d4c7b7beef80eb`

The final physical harness and binaries are additionally bound by:

- final main: `331c7c5298acfcc5a9d8020af6a036e7d77f253ddcebf0276f9d6afbd67c5000`;
- firmware BIN: `86457e990bd1b96a093f418c5ce8f71b60d41f55c4065757df4ca543c7b4863d`;
- firmware ELF: `30cbbf62952b442b41aa2bf1398bd9377bd1401a18a5da41a1ea74dd601f1ff3`;
- evidence-set SHA-256: `d38648587bc627409e770f5b3fd3b2972e9abe31094db8550257fc8c9a6a0106`.
- portability-assay source: `408deada3466d53f51354a3ae2c4ece53e2bf6dc1f4eb0c1bd7c692d179aa999`.

The evidence-set digest binds the ordered artifact manifest for the final physical build.

15 Interpretation Boundaries

A technically strong claim must be scoped. The present evidence supports:

- cross-target compilation portability of the tested compact C substrate across ARM Cortex-M, 32-bit RISC-V, AVR and MIPS32 profiles;
- native execution of the validated kernel on physical ESP32-S3 silicon;
- reproduction of Tier-A behavior on two physical units;
- 150/150 five-level causal parity against the frozen PC reference;
- 100,000-cycle physical endurance with zero cognitive failures in the preserved Step20 route;
- host qualification of voice and acquisition authority boundaries;
- externally instrumented route and action pulses;
- a 100,000-cycle analog closed loop with INC/HOLD/DEC all non-zero, zero ADC failures and zero invalid-state failures;
- local execution of the validated route without cloud inference;
- cryptographic identity of the disclosed evidence artifacts.

The present evidence does *not* establish:

- consciousness, sentience or biological life;
- general intelligence;
- superiority over LLMs in reasoning, knowledge or language quality;
- physical execution on ARM Cortex-M, RISC-V, AVR or MIPS32 boards solely from the portability assay;
- physical I2S voice validation;
- physical live Wi-Fi/HTTP acquisition;
- universal timing independent of workload, compiler, target or instrumentation.

16 Threats to Validity

Instrumentation. GPIO markers add small overhead and oscilloscope cursor placement has finite resolution. External route timing is therefore reported as a representative pulse, not substituted for the internal 100,000-cycle distribution.

Harness dependence. The final closed loop includes a specific B25K reference, GPIO mapping and RC network. The PASS result validates this configuration and architecture route; broader plant dynamics require separate testing.

Black-box reproducibility. Public readers can verify hashes and reported artifacts but cannot independently reconstruct proprietary internals from this paper alone. This is intentional IP protection and limits full open-source reproducibility.

Benchmark comparability. Benchmark 3, cross-target compilation, Tier-A, Step20, final closed-loop route and oscilloscope captures are distinct evidence types/workloads. Their latency values must not be merged as if they measured the same operation.

Semantic scope. Five-level parity strongly constrains implementation equivalence for the tested cases, but it is not an exhaustive proof over all possible inputs or future modules.

17 Discussion

The evidence is notable because it converges across levels. Benchmark 3 shows a small resident CPU runtime under a narrow protocol. The cross-target assay then shows that the tested compact C substrate compiles across unrelated embedded ISA families without converting those compile results into physical-execution claims. Tier-A subsequently demonstrates native float32 execution on physical MCU silicon and cross-device reproduction. Step19 checks causal identity at five internal boundaries. Step20 demonstrates endurance. Step21 and Step22 isolate voice and acquisition from cognitive authority. Oscilloscope captures provide an external timing channel. The final closed loop connects physical perturbation, endogenous computation, action and physical feedback. SHA-256 fingerprints bind the implementation artifacts without disclosing the black box.

This combination is the basis for describing OSH as a proposed autonomous homeostatic cognitive architecture rather than as an API facade. The strongest evidence is not any single speed ratio or photograph; it is the agreement of independent evidence types.

The architecture also suggests a different deployment regime from large foundation models. A resident mechanism that fits and executes on microcontroller-class hardware can in principle be placed near sensors, actuators, robots, vehicles and embedded systems without making cloud inference a prerequisite for the validated cognitive loop. Those applications remain engineering directions, not results of the present paper.

18 Conclusion

This work presents a cumulative validation of OSH from host runtime to physical closed-loop execution. A frozen architecture was carried from host execution through cross-target compilation portability and then through native ESP32-S3 execution, cross-device reproduction, five-level causal parity, endurance, independent artifact reruns, external oscilloscope instrumentation and a final 100,000-cycle analog feedback experiment.

The final run completed with INC=5,160, HOLD=93,358, DEC=1,482, ADC_FAIL=0, STATE_FAIL=0, PHYSICAL_100K_GATE=PASS and FINAL_LATCH=PASS. The internally measured route mean was 29.525 µs; the external evidence set contains a representative 23.0 µs

route capture and a 660 ns action/selection capture. Native modules, build configuration, BIN and ELF are cryptographically fingerprinted.

The defensible conclusion is therefore concrete: OSH exists as a native, instrumentable computational mechanism on physical microcontroller silicon; the tested causal route is reproducible against its frozen PC reference; the mechanism can sustain long execution and participate in a real closed feedback loop without cloud inference. Further claims—general intelligence, consciousness, universal superiority, or untested physical modalities—remain outside the evidence presented here.

A Final Physical Gate Extract

```
MILESTONE=100000
REF_RAW=1610 FB_RAW=1295
INC=5160 HOLD=93358 DEC=1482
ADC_FAIL=0 STATE_FAIL=0
ROUTE_MIN_US=24
ROUTE_MAX_US=30
ROUTE_MEAN_US=29.525

PHYSICAL_100K_GATE=PASS
CYCLES=100000
ADC_FAIL=0 STATE_FAIL=0
INC=5160 HOLD=93358 DEC=1482
IAE=10862.551863
MAX_STATE_DELTA=0.141184
ELAPSED_S=1000.005270
DONE_GPIO15=1

FINAL_LATCH=PASS
DONE_GPIO15=1
FAIL_GPIO4=0
PASS_GPIO5=1
```

B Portability Evidence Fingerprint

The cross-target portability assay source SHA-256 is `408deada3466d53f51354a3ae2c4ece53e2bf6dc1f4eb0c1bd7c692d179aa999`. The assay used Clang 17.0.0, `-Oz`, freestanding compilation and no link-time vendor SDK. Physical-board claims remain restricted to separately validated silicon.

C Prior Validation Fingerprints

Additional preserved component fingerprints from the consolidated benchmark are:

- Continuity/scars: `376ac8491c712e9582431d2e4b88036203f9fc50fdef2cc704571ba0f00b80bd`
- Persistence: `85fbaf60682b10cdaa1cd6a0bc7ca1a0f04507358b877fd8281f4f7b20c0deb3`
- Phonation: `7cb3c0021792b749edbc4ab622b94a15bdf75afdf9bf42439f31591f4677498c`
- Acquisition: `23461bdd9eee9fd728c05c6e7ace7452953767c491ba1583953b17219bef739b`

D Overleaf Evidence Files

Upload this `.tex` file and the four JPEG files with exactly these names:

```
fig_physical_setup_route.jpg
fig_route_cursor.jpg
fig_physical_setup_action.jpg
fig_action_cursor.jpg
```

The separate file `OSH_PAPER_PHOTO_SHA256_2026-09-21.txt` contains SHA-256 fingerprints of the four originals/copies used by the paper.

E Reproducibility Note

The public reproducibility model is intentionally evidence-oriented rather than source-disclosure-oriented. A verifier can compare published hashes, hardware identity, numerical gates, photographs, oscilloscope cursor values and run logs. Reproduction of the proprietary internal equations requires controlled access and is outside this public artifact.