Gradient-Stabilized Recursive Filtering (GSRF Practical)

Technical note — log-space soft thermostat for pre-model signal hygiene

Author: Carl Boon · Organization: BoonMind · Contact: **info@boonmind.io**
Product: Zero Overshoot (GSRF Practical) · Site: https://www.zeroovershoot.com/
Public papers/eval: https://github.com/CBoon99/gsrf-papers · Document version: 2026-08-08 (replaces Jan 2026 marketing PDF)

Claim-safe notice. This note restates the public claim ladder only. It does **not** claim a fixed false-alarm residual detector, a universal non-overshoot guarantee on all plants, medical device performance, or that GSRF always beats EMA. Live locked numbers: https://www.zeroovershoot.com/evidence.html

1. Abstract

Zero Overshoot implements **GSRF Practical**: a deterministic nonlinear filter in log-space with a restoring spring toward a declared normal `x*`, a bounded observation pull (`tanh`), and optional memory. It is designed to reduce mid-band oscillation energy and compress peaks toward normal on strictly positive signals—not to maximize tracking MAD, not to tolerate multi-minute feedback delay as a command smoother, and not to win fixed false-alarm residual event detection versus EMA.

Evidence is produced by the GSRF Audit Workbench under pre-declared rules. Public marketing may only climb the published claim ladder. Evaluation materials (Try snippet / gsrf-bench) are non-production.

2. Identity (what it is)

For observation `E_t > 0`:

```
x_obs = log(max(E_t, ε))
G = −(x − x*)   // restoring gradient toward normal
R = tanh(x − x_prev)   // bounded memory
O = tanh(x_obs − x)   // bounded observation pull
x ← x + β (b + k_return·G + mem·R + w_obs·O) · Δt
output = exp(x)
```

Typical balanced-style defaults used in public characterization: β=0.3, k_return=1.0, mem=0.2, obs_weight=0.6, Δt=1.0. `x*` is a declared normal in log-space (e.g. median log of a calm prefix, or adaptive trailing median). **Reference** mode is an open-attractor twin (weak observation coupling) used for comparison—not the primary product tracker.

3. Supported public claims (locked-style)

- **Mid-band oscillation:** ~+35% osc31 vs best EMA on phase6.2 actuator-class packs (3 seeds, full series); sweet-band spectral power ~−70% vs EMA on the same packs.
- **Peak compression toward normal:** ~38% mean peak-deviation advantage; 49/52 dataset wins in excellence hunt v2.
- **Soft-setpoint identity:** pure-sine AC gain plateaus ~0.37 while EMA → ~1.0 at long periods; step 50→65 with frozen x* settles near ~55 (EMA fully tracks 65).
- **Adaptive x*:** permanent regime jumps can be followed without dropping the spring; caution on flat SpO2-like traces.
- **Industrial batteries (claim-safe):** synthetic Phase 5 SCADA-style packs under frozen rules; DAMADICS Lublin real-plant first look published with partial and loss—not “wins industrial actuators” without qualifiers.

4. Explicit non-claims (do not invert)

- Not a general MAD tracking champion (0/52 wins vs best EMA in excellence hunt).
- Not a multi-minute delay closed-loop command smoother (MAD-to-measured worse vs EMA at τ≈3/6/10 min on locked packs).
- Not a fixed false-alarm residual event detector (EMA residual recall higher at matched FA on locked probes).
- Not universal always-beats-EMA on every frequency band or plant.
- Not a medical device; no clinical false-alarm reduction claims.

- Not a quantum decoder scoreboard on this document (quantum is a separate NDA ladder).
- Demo/snippet/eval package grants **no** production rights.

5. Parameter envelopes (Practical)

Public characterization commonly stays near: $\beta \in (0, 1]$, mem $\in [0, 0.5]$, k_return > 0 for the product niche, obs_weight trading usability vs pure attractor behavior. Ablations show k_return is the soul of the oscillation/peak niche: k_return=0 loses those advantages and behaves more like a tracker path.

6. Deployment posture

- Prefer measurement hygiene and shadow mode before any command path.
- Audit plant/network delay; multi-minute delay command smoothing is a published red light.
- Keep residual/alarm detectors on a separate design path—do not force GSRF residual to be the FA engine.
- Validate on your traces with gsrf-bench; local results do not rewrite the public evidence ledger.
- Production requires a commercial license (Pro / Enterprise / Quantum tiers).

7. How to evaluate

Try: https://www.zeroovershoot.com/try.html
gsrf-bench: https://www.zeroovershoot.com/gsrf-bench.html
Eval package (LICENSE-EVAL): https://github.com/CBoon99/gsrf-papers/tree/main/eval
Evidence: https://www.zeroovershoot.com/evidence.html
Sources & DAMADICS citation: https://www.zeroovershoot.com/sources.html

8. Citation

Boon, C. / BoonMind (2026). Gradient-Stabilized Recursive Filtering (GSRF Practical) — technical note aligned to the Zero Overshoot claim ladder. https://www.zeroovershoot.com/evidence.html · Contact: info@boonmind.io