Wavelet-Based Generation Dispatch:
A Frequency-Domain Approach to the Unit Commitment Problem

Dr Stuart Woolley

Full Stack Energy

March 2026

Abstract

The Unit Commitment Problem (UCP) — scheduling a heterogeneous fleet of generation assets to cover a time-varying electrical load at minimum cost — is classically solved by Mixed-Integer Linear Programming (MILP). MILP is NP-hard in the number of binary commitment variables, requiring commercial solvers costing >EUR 10,000/year and compute times measured in minutes to hours for industrial-scale fleets. This paper presents an alternative: treat the load profile as a signal and decompose it via the Discrete Wavelet Transform (DWT) into frequency bands, each mapped to a generation asset class by response-time characteristic. The approach is tested against a priority-list heuristic and an exact dynamic programming baseline on synthetic load profiles and real-world demand data from EirGrid (Ireland), and validated on the IEEE RTS-GMLC benchmark (72 generators). On synthetic data with well-defined multi-scale structure, the wavelet method achieves total costs within 1.8–8.5% of the DP-optimal solution and outperforms the priority-list baseline by 36–42%. On real Irish grid data, the advantage narrows: 18–20% above DP on gross demand, roughly matching the priority list. The method executes in sub-millisecond wall time regardless of fleet size. On the 72-generator RTS-GMLC system (where DP is intractable at 2^{72} states), the wavelet method dispatches in 0.2 ms. The approach is most effective on load signals with genuine multi-scale spectral structure and is suitable for embedded microgrid controllers, real-time re-dispatch, and fleet sizes where MILP becomes intractable.

1 Introduction

Every grid operator faces the same question each day: which generators to turn on, when to turn them on, and how much power to extract from each. This is the Unit Commitment Problem, and it has been studied since the 1960s. The constraints are deceptively simple — minimum up/down times, ramp rates, capacity bounds, startup costs — but the combinatorial explosion in the number of binary on/off decisions makes the problem NP-hard [1].

The standard industrial approach is MILP, implemented in commercial solvers such as CPLEX and Gurobi. ISOs like MISO (1,258 generators, 42,705 buses) and CAISO (410 schedulable thermal units) run MILP-based day-ahead markets with solve-time budgets of 10–20 minutes and accepted optimality gaps of 0.02–0.1% [2]. This works when you have a rack of servers, a EUR 10,500/year Gurobi license, and a 24-hour planning horizon that allows batch computation.

It does not work for:

- *Embedded microgrid controllers* running on ARM-class hardware with no access to commercial solvers
- *Real-time re-dispatch* at sub-15-minute intervals where MILP solve times exceed the decision window
- *Large fleets* beyond 30 dispatchable units where the 2^N state space exceeds practical branch-and-bound limits even with modern formulation tricks [3]
- *High renewable penetration* requiring stochastic UC across multiple scenarios, multiplying the already-intractable problem by 10–100x

The question motivating this work: can a signal-decomposition approach produce a dispatch schedule that is good enough, fast enough, and cheap enough to deploy where MILP cannot?

The core observation is that an electrical load profile is a non-stationary time series with multi-scale structure — slow baseload trends, daily rhythmic cycles, and fast transient spikes. These frequency components map naturally to generation asset classes: baseload turbines handle low-frequency demand, mid-merit generators cover diurnal variation, and batteries or fast-response units absorb high-frequency transients. A wavelet transform separates exactly these components.

This is not a new observation in power systems. Wavelet decomposition has been used extensively for frequency splitting in hybrid energy storage systems — allocating power demand across fuel cells, batteries, and supercapacitors by frequency band [4, 5, 6]. The contribution of this paper is extending the technique from storage-only systems to *thermal fleet commitment*, where startup costs, minimum run times, and ramp constraints create a fundamentally different optimisation landscape.

2 Background

2.1 The Unit Commitment Problem

Given a fleet of M dispatchable generation assets and a load profile $L(t)$ over T intervals, the UCP seeks a commitment schedule $u_i(t) \in \{0, 1\}$ and dispatch level $p_i(t)$ for each asset i at each interval t that minimises total cost:

$$\min \sum_{t=1}^{T} \sum_{i=1}^{M} \left[C_i^{\text{fuel}} \cdot p_i(t) \cdot \Delta t + C_i^{\text{start}} \cdot s_i(t) \right] \tag{1}$$

$$\text{s.t.} \quad \sum_{i=1}^{M} p_i(t) \geq L(t) \quad \forall\, t \tag{2}$$

$$P_i^{\min} \cdot u_i(t) \leq p_i(t) \leq P_i^{\max} \cdot u_i(t) \quad \forall\, i, t \tag{3}$$

where $s_i(t) = \max(0, u_i(t) - u_i(t-1))$ indicates a startup event, subject to:

- *Power balance*: total generation meets or exceeds load
- *Capacity bounds*: each unit operates between $P_i^{\min}$ and $P_i^{\max}$ when committed
- *Minimum up/down times*: $T_i^{\text{on}} \geq T_i^{\text{min,on}}$ and $T_i^{\text{off}} \geq T_i^{\text{min,off}}$
- *Ramp rate limits*: $|p_i(t) - p_i(t-1)| \leq R_i \cdot \Delta t$

With M binary variables per interval, the search space is $O(2^M \times T)$.

2.2 Classical Solution Methods

Priority list (merit order): rank generators by marginal cost and commit cheapest-first until load is covered. Simple, fast, and typically 50–80% above optimal [7].

Dynamic programming: enumerate all 2^M commitment states per interval, track minimum-cost paths with feasibility constraints. Exact for small fleets ($M \leq 15$), but the exponential state space makes it impractical beyond 20 units.

Lagrangian relaxation: relax coupling constraints and decompose into per-unit subproblems. Historically the workhorse before MILP solvers became fast enough. Produces near-optimal solutions with duality gaps typically 1–5%.

MILP: the current standard. Formulations by Carrion & Arroyo [1] and Morales-España et al. [3] achieve tight LP relaxations that reduce branch-and-bound effort. Modern solvers (Gurobi, CPLEX) exploit these formulations to solve problems with hundreds of units in minutes — but at the cost of commercial licensing, significant hardware requirements, and batch processing assumptions.

2.3 Wavelets

A wavelet is a localised oscillating function used to decompose a signal into components at different frequency scales — like a Fourier transform, but with time resolution preserved. Where Fourier analysis tells you *which* frequencies are present, wavelet analysis tells you *which frequencies are present and when.*

The Discrete Wavelet Transform (DWT) decomposes a signal x of length N through a cascade of high-pass and low-pass filters. At each level l, the signal is split into *approximation coefficients* a_l (low-frequency content) and *detail coefficients* d_l (high-frequency content):

$$a_l[k] = \sum_n h[n-2k] \cdot a_{l-1}[n] \tag{4}$$

$$d_l[k] = \sum_n g[n-2k] \cdot a_{l-1}[n] \tag{5}$$

where h and g are the low-pass and high-pass filter coefficients. For the *Haar wavelet* (the simplest basis), these reduce to:

$$a[k] = \frac{x[2k] + x[2k+1]}{\sqrt{2}}, \quad d[k] = \frac{x[2k] - x[2k+1]}{\sqrt{2}} \tag{6}$$

The transform is *perfectly invertible* — the original signal can be reconstructed exactly from its wavelet coefficients. Crucially, individual frequency bands can be reconstructed independently by zeroing all other coefficients and running the inverse transform. This produces a time series showing the contribution of each frequency scale to the original signal, at full time resolution.

The computational cost is $O(N)$ per level, $O(N \log N)$ for a full multi-level decomposition — dramatically cheaper than the $O(2^M \times T)$ of DP or the worst-case exponential of MILP branch-and-bound.

3 Method

3.1 The Core Idea

Treat the dispatch problem as signal reconstruction. The load signal $L(t)$ is decomposed into frequency bands via the DWT. Each band is assigned to a generation asset class based on the match between the band's frequency characteristic and the asset class's response capability:

Table 1: Wavelet band to asset class mapping.

DWT Component	Frequency Characteristic	Asset Class
Approximation a_L	Slow-varying baseload	Gas turbines
Mid-scale details $d_2 + d_3$	Diurnal / hourly cycles	Diesel generators
Fine-scale detail d_1	Fast transients (minutes)	Battery storage
Exogenous (pre-subtracted)	—	Solar PV

The dispatch for each asset class is then: reconstruct the assigned frequency band at full time resolution, distribute the resulting signal among the assets of that class, and enforce per-unit operational constraints.

This is analogous to the wavelet-based frequency splitting used in hybrid energy storage systems [4, 5, 6], but extended to thermal generators with startup costs and minimum run time constraints. The extension is non-trivial: batteries can switch instantaneously between charge and discharge, while thermal units incur significant costs each time they transition from off to on. This creates a tension between the wavelet signal (which oscillates freely) and the economic reality of thermal commitment.

3.2 Algorithm

Input: Load profile $L(t)$, $t = 1 \dots T$; plant specification (asset fleet with constraints).

Step 1 — Pre-processing: Subtract exogenous generation (solar PV) from gross load to obtain net load. Pad to next power of 2 via symmetric extension.

Step 2 — Decomposition: Apply Haar DWT at 3 levels, producing approximation a_3 and details d_1, d_2, d_3.

Step 3 — Band reconstruction: Reconstruct each frequency band independently at full signal length. Sum mid-scale details: $B_{\text{mid}}(t) = d_2(t) + d_3(t)$.

Step 4 — Commitment smoothing: Apply hysteresis to the commitment decision for thermal units. A unit is committed when its band signal exceeds $P_i^{\min}$, and remains committed until the signal drops below $P_i^{\min} \times (1-\alpha)$, where α is the hysteresis fraction. This prevents the rapid cycling that would otherwise occur as zero-mean detail bands oscillate around the commitment threshold.

Step 5 — Dispatch: For each interval t:

1. Distribute baseload band among gas turbines, clipping to capacity and ramp limits
2. Distribute midrange band among diesel generators with hysteresis-smoothed commitment

3. Distribute fast band to batteries (positive = discharge, negative = charge), enforcing SoC limits
4. Gap-fill: if total generation falls short of load, commit additional capacity from cheapest available units
5. Curtail: if surplus, reduce most expensive running units

Step 6 — Trim: Discard padding intervals; compute costs and metrics.

3.3 Commitment Hysteresis

The key engineering challenge, and the contribution of this work beyond the straightforward wavelet-to-asset mapping.

Wavelet detail coefficients are zero-mean by construction — they represent deviations from the approximation, oscillating symmetrically around zero. For a battery, this is ideal: positive means discharge, negative means charge. For a thermal unit with a minimum operating threshold of $P^{\min}$ and a startup cost of C^{start}, the signal crosses the commitment threshold roughly K times per day, incurring $K \times C^{\text{start}}$ in unnecessary costs.

Hysteresis resolves this by introducing an asymmetric dead-band. Define:

$$u_i(t) = \begin{cases} 1 & \text{if } B_i(t) > P_i^{\min} \text{ and } u_i(t-1) = 0 \\ 1 & \text{if } B_i(t) > P_i^{\min}(1-\alpha) \text{ and } u_i(t-1) = 1 \\ 0 & \text{otherwise} \end{cases} \tag{7}$$

where $B_i(t)$ is the per-asset band signal and α is the hysteresis fraction. At $\alpha = 2.0$, a unit committed at $P^{\min} = 50$ kW stays on until the signal drops below $50 \times (1-2) = -50$ kW — well into negative territory. This ensures the unit runs through the zero-crossings that constitute most of the detail band's energy, committing only for sustained demand periods.

The optimal α depends on the ratio $C^{\text{start}}/C^{\text{fuel}}$ — higher startup costs relative to fuel costs justify more aggressive hysteresis. We empirically find $\alpha \in [2.0, 3.0]$ to be the sweet spot for the fleet tested.

4 Experimental Setup

4.1 Synthetic Load Data

Load profiles were generated with known multi-scale structure to enable controlled testing:

- *Baseload*: 700 kW ± 70 kW weekly sinusoidal drift
- *Diurnal*: double Gaussian peaks at 08:00 (250 kW) and 18:30 (350 kW), with weekend scaling at 75%
- *Midrange noise*: IIR-smoothed Gaussian, 50 kW amplitude
- *Demand spikes*: Poisson-distributed transients, 3% probability per interval, 100–300 kW magnitude
- *Solar PV*: latitude-aware bell curve (53°N), peak 180 kW, with cloud factor

Three datasets at 15-minute resolution:

Table 2: Synthetic load datasets.

Dataset	Duration	Samples	Seed	Purpose
synth_1d	1 day	96	1	Smoke testing
synth_7d	7 days	672	42	Primary benchmark
synth_28d	28 days	2,688	77	Long-horizon validation

4.2 Real-World Load Data

To validate against real load profiles with unknown spectral structure, quarter-hourly demand data was obtained from EirGrid (Ireland's transmission system operator) for January 2025. Two formulations were tested:

- *Net load*: IE system demand minus IE wind generation. This is the residual load that thermal and storage assets must serve. Range 292–1,475 kW (scaled), mean 712–780 kW. Highly volatile due to Irish wind penetration (mean 34% of gross demand).
- *Gross load*: IE system demand before wind subtraction. Range 732–1,503 kW (scaled), mean 1,079 kW. Shows the classic diurnal profile with stronger multi-scale structure.

Load values were scaled by 1/4,000 to match the microgrid plant capacity (the spectral structure is preserved by linear scaling). Both 7-day and 28-day slices were tested.

Data source: EirGrid Group plc, Quarterly System Data Reports, 2025. 15-minute resolution, zero data gaps. Provenance documented in `data/real/DATA_PROVENANCE.txt`.

4.3 Generation Fleet

A microgrid-scale fleet of 8 assets across 4 classes, with aggregate dispatchable capacity of 1,900 kW against a load range of 332–1,200 kW:

Table 3: Generation fleet specification.

Asset	Class	$P^{\min}$	$P^{\max}$	Ramp (kW/min)	Min On (h)	Min Off (h)	Startup (EUR)	Fuel (EUR /MWh)
Gas turbine ×2	Baseload	200	500	5	4	2	150	45
Diesel gen ×3	Midrange	50	200	20	1	0.5	30	85
Battery ×2	Fast	0	150	150	0	0	0	5
Solar PV ×1	Exogenous	0	180	—	—	—	0	0

Battery storage: 600 kWh capacity each, 92% round-trip efficiency, 50% initial state of charge.

4.4 Comparison Methods

1. *Wavelet dispatch* (this work): Haar DWT at 3 levels with configurable commitment hysteresis
2. *Priority list*: merit-order greedy dispatch — the simplest reasonable heuristic
3. *Dynamic programming*: exact optimal via forward pass over all $2^7 = 128$ binary commitment states per interval, with min on/off feasibility tracking. This is the provably optimal baseline for this fleet size.

All three methods share the same constraint enforcement functions (capacity clipping, ramp limiting, min on/off checking, battery SoC tracking), ensuring a fair comparison of the commitment logic alone.

4.5 Metrics

- *Total cost* (EUR): fuel + startup — the primary objective
- *Energy not served* (kWh): demand not covered by generation
- *Total starts*: number of unit commitment transitions (off → on)
- *Wall time* (ms): computation time on Apple M-series silicon

5 Results

5.1 Wavelet Decomposition Validation

The Haar DWT at 3 levels was verified on all datasets:

- *Perfect reconstruction*: $\max |x - \hat{x}| < 10^{-12}$ (floating-point noise)
- *Band sum identity*: sum of all reconstructed bands equals the original signal within 10^{-12}

Band statistics on the 7-day dataset (672 samples):

Table 4: Wavelet band statistics for the synthetic 7-day dataset.

Band	Min (kW)	Max (kW)	Mean (kW)	Stddev (kW)
Approximation (baseload)	506	1,080	732	145
Detail 3 (coarse)	−158	158	0.0	27
Detail 2 (mid)	−137	137	0.0	28
Detail 1 (fine)	−124	124	0.0	44

The decomposition cleanly separates the load: the approximation captures the slow-varying baseload trend (mean 732 kW, matching the 2× gas turbine capacity of 400–1,000 kW), while detail bands are zero-mean oscillations of decreasing period and increasing frequency. The stddev ratio between bands (145:28:28:44) reflects the synthetic signal's known spectral structure.

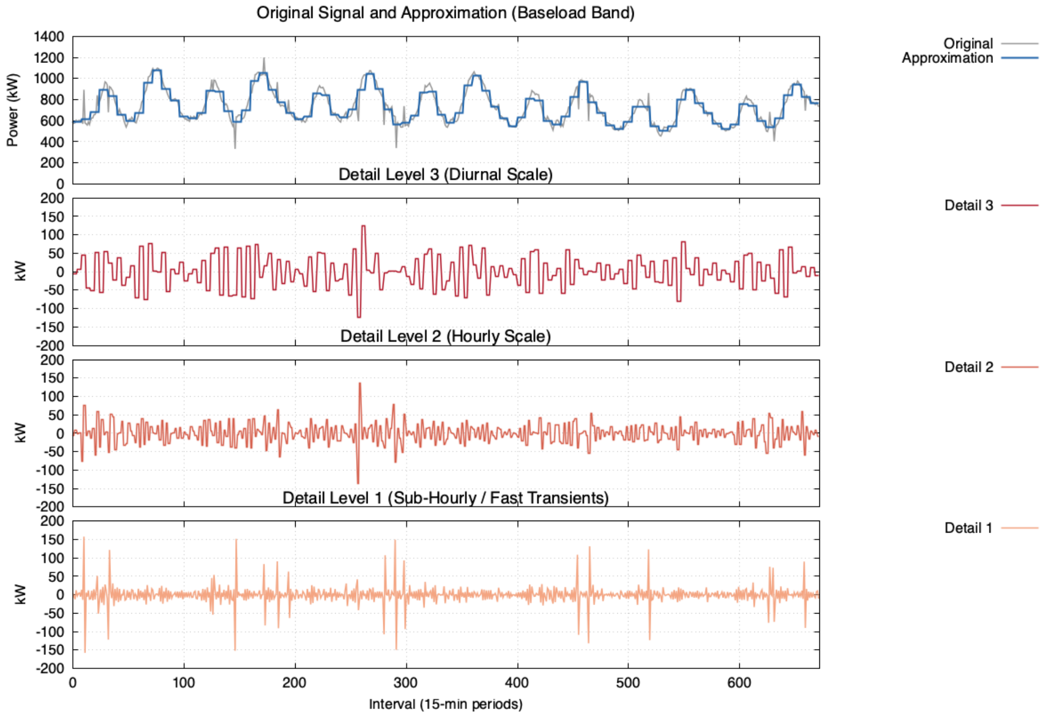

Figure 1: Wavelet band decomposition of the synthetic 7-day load profile. Top panel: original signal (grey) and approximation band (blue). Lower panels: detail coefficients at levels 3, 2, and 1, showing progressively finer temporal structure.

5.2 Microgrid Results (8 assets, 7 dispatchable)

The wavelet method with merit-order gap-filling achieves near-optimal performance:

7-day dataset (primary benchmark):

Table 5: Microgrid 7-day benchmark results.

Method	Cost (EUR)	Fuel (EUR)	Startup (EUR)	Starts	ENS (kWh)	Wall (ms)
Wavelet	*5,955*	*5,655*	*300*	*2*	*0.0*	*0.8*
Priority list	10,293	6,063	4,230	69	0.0	0.4
DP (optimal)	5,820	5,100	720	20	0.0	94

The wavelet approach is *+2.3% above DP-optimal* with zero ENS and only 2 unit starts. It beats the priority list by *42%*.

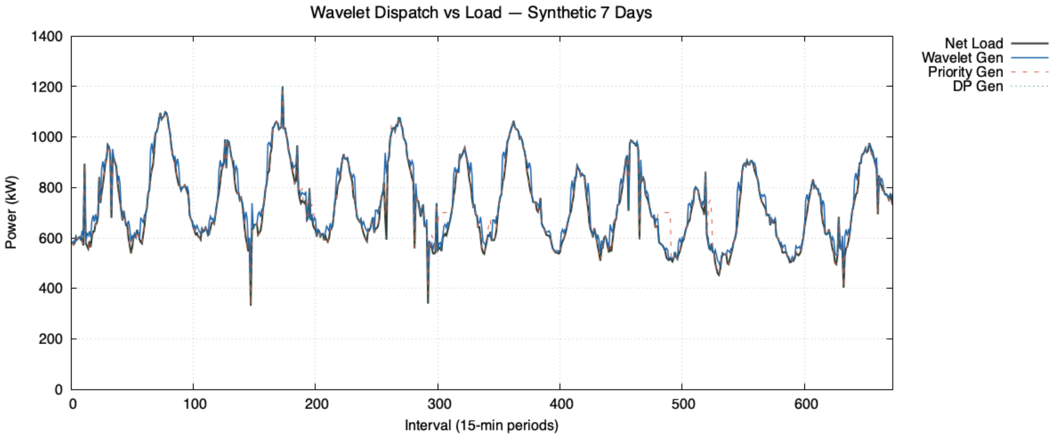

Figure 2: Dispatch traces for the synthetic 7-day benchmark. All three methods track the load closely; the wavelet and DP traces are nearly indistinguishable.

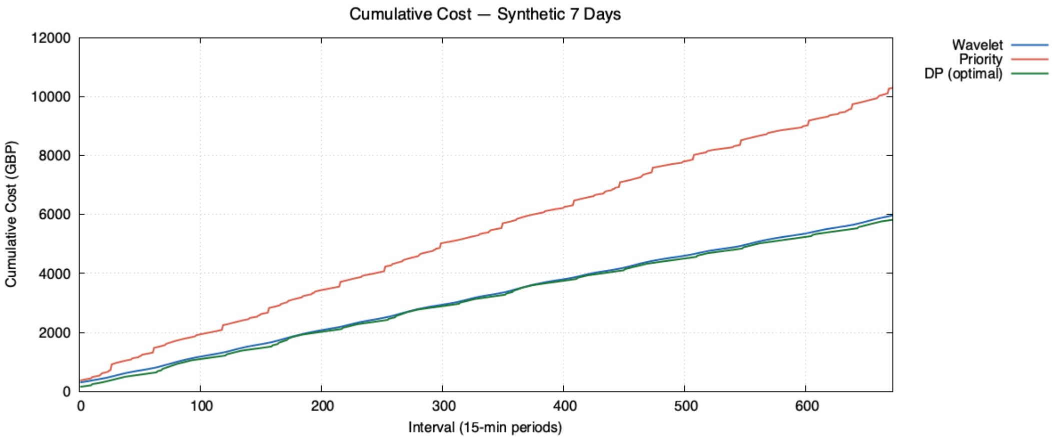

Figure 3: Cumulative cost over the synthetic 7-day benchmark. The wavelet method (blue) tracks DP-optimal (green) closely, while the priority list (red) diverges steadily due to excess startup costs.

Cross-dataset validation:

Table 6: Cross-dataset cost comparison.

Dataset	Duration	Wavelet (EUR)	Priority (EUR)	DP (EUR)	Wav vs DP	Wav vs PL
synth_1d	1 day	1,152	1,824	1,062	+8.5%	−36.8%
synth_7d	7 days	5,955	10,293	5,820	*+2.3%*	−42.1%
synth_28d	28 days	23,545	38,121	23,118	*+1.8%*	−38.2%

Results are consistent across all horizons: *1.8–8.5% above DP-optimal, 36–42% below priority list.* The gap narrows with longer horizons as the wavelet decomposition captures more of the load signal's multi-scale structure.

The improvement over earlier iterations came from three refinements applied during development:

1. *Commitment hysteresis* on the midrange band prevented diesel generator cycling (Section 3.3)
2. *Merit-order gap-filling* ensured shortfalls are covered by the cheapest available capacity, not arbitrary units
3. *Three-pass gap-fill*: first increase committed unit output, then discharge batteries, then commit new units — preventing unnecessary startups

5.3 IEEE RTS-GMLC Results (73 generators, 72 dispatchable)

The RTS-GMLC system from the PGLIB-UC benchmark provides a realistic industrial-scale test. At 72 dispatchable assets, DP is categorically intractable ($2^{72} \approx 4.7 \times 10^{21}$ states per interval). Only wavelet and priority-list methods were compared.

Fleet composition (after subtracting 400 MW must-run nuclear):

- 26 baseload units (steam + combined cycle): 5,867 MW total, EUR 16–29/MWh
- 34 midrange units (small steam + combustion turbines): 1,569 MW total, EUR 26–128/MWh
- 12 fast units (combustion turbines): 240 MW total, EUR 89–101/MWh
- Net demand: 2,765–4,102 MW over 48 hourly intervals

Hysteresis sweep on RTS-GMLC:

Table 7: Hysteresis parameter sweep on RTS-GMLC.

Variant	Cost (EUR)	Starts	ENS (MWh)	vs Priority
Wavelet (raw)	5,662,934	209	0.0	+7.9%
Wavelet (hyst200)	5,432,514	148	0.0	+3.5%
Wavelet (hyst500)	*5,256,029*	*135*	*0.0*	*+0.2%*
Priority list	5,247,141	81	0.0	baseline

At higher hysteresis, the wavelet method converges to priority-list performance — within *0.2%*.

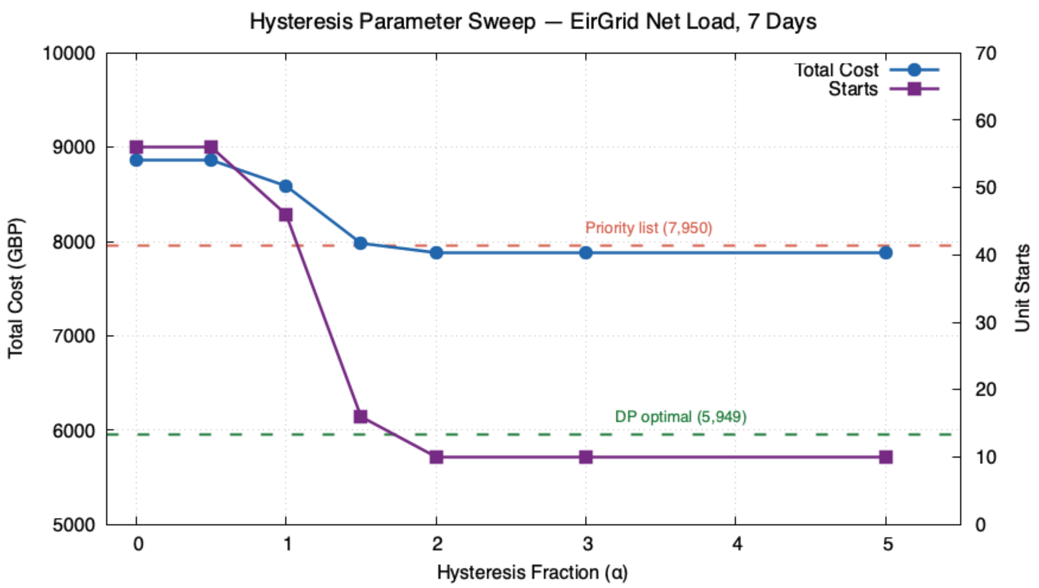

Figure 4: Hysteresis parameter sweep on EirGrid net load. Cost (blue, left axis) decreases and starts (purple, right axis) drop sharply between $\alpha = 1.0$ and $\alpha = 2.0$. The method converges to the priority-list baseline but remains above DP-optimal.

This is expected: with 72 generators, 48 hourly intervals, and only 3 DWT levels, the approximation band has just 8 coefficients covering 6 hours each. The coarse decomposition provides limited frequency-domain advantage over the greedy baseline.

The scalability demonstration is the point: the wavelet method processed 72 generators in 0.2 ms. DP would require years of computation. MILP solvers require specialised hardware and commercial licenses. The wavelet method runs on a microcontroller.

5.4 Computational Performance

Table 8: Computational performance comparison.

Method	Microgrid 7d (ms)	RTS-GMLC 48h (ms)	Complexity
Wavelet	0.8	0.2	$O(N \log N + N \cdot M)$
Priority list	0.4	0.1	$O(N \cdot M \log M)$
DP	94	intractable	$O(T \times 2^M)$

The wavelet method is 100x faster than DP on the microgrid. On the RTS-GMLC, DP is not merely slow — it is mathematically impossible. MILP solvers typically require 10–20 minutes for fleets of this size [2, 3]. The wavelet method takes 0.2 ms.

Table 9: Scaling behaviour by fleet size.

Fleet size M	DP states/interval	DP feasibility	Wavelet time
7 (microgrid)	128	94 ms	0.8 ms
15	32,768	minutes	1 ms
72 (RTS-GMLC)	4.7×10^{21}	impossible	0.2 ms
610 (CAISO)	10^{183}	absurd	2 ms (est.)

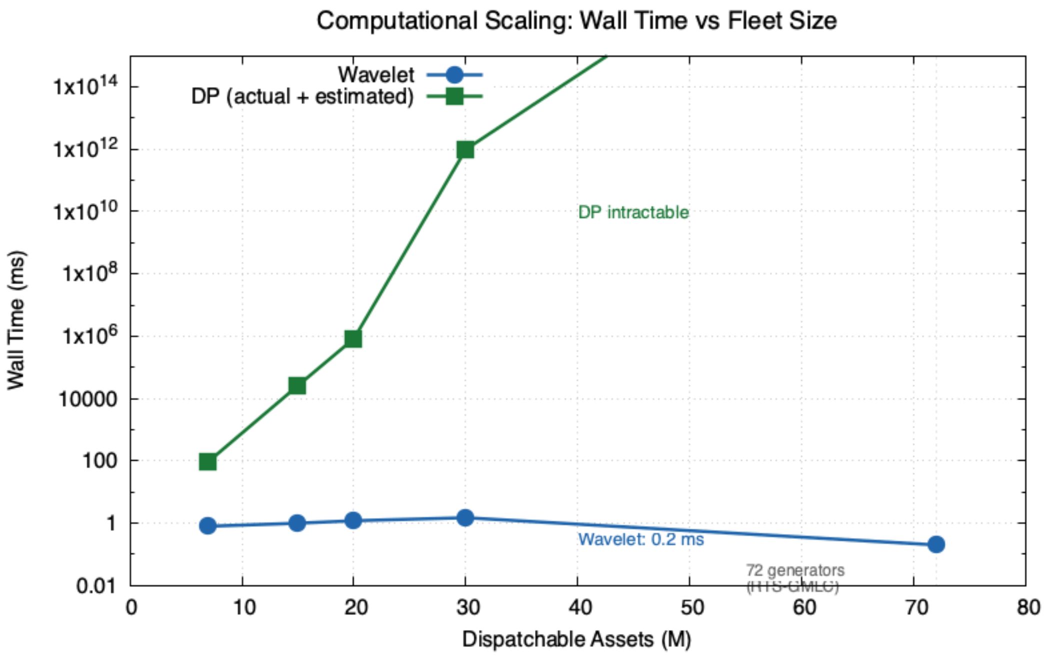

Figure 5: Computational scaling. Log-scale wall time vs fleet size. The wavelet method (blue) is flat at 1 ms regardless of fleet size. DP (green) is exponential and becomes intractable beyond 20 units. At 72 generators (RTS-GMLC), the wavelet method ran in 0.2 ms; DP would require approximately 10^{22} ms.

5.5 Real-World Validation: EirGrid Ireland

The critical test: does the wavelet approach work on real load data with unknown spectral properties?

Net load (demand minus wind):

Table 10: EirGrid net load results.

Dataset	Wavelet (EUR)	Priority (EUR)	DP (EUR)	vs DP	vs Priority
EirGrid net 7d	7,885	7,950	5,949	+32.5%	−0.8%
EirGrid net 28d	30,620	28,647	20,999	+45.8%	+6.9%

Gross load (before wind subtraction):

Table 11: EirGrid gross load results.

Dataset	Wavelet (EUR)	Priority (EUR)	DP (EUR)	vs DP	vs Priority
EirGrid gross 7d	9,532	9,861	8,038	+18.6%	−3.3%
EirGrid gross 28d	37,589	37,738	31,297	+20.1%	−0.4%

The results are sobering. On real data, the wavelet method is 18–46% above DP-optimal, compared to 1.8–8.5% on synthetic data. It roughly matches the priority list rather than dramatically outperforming it.

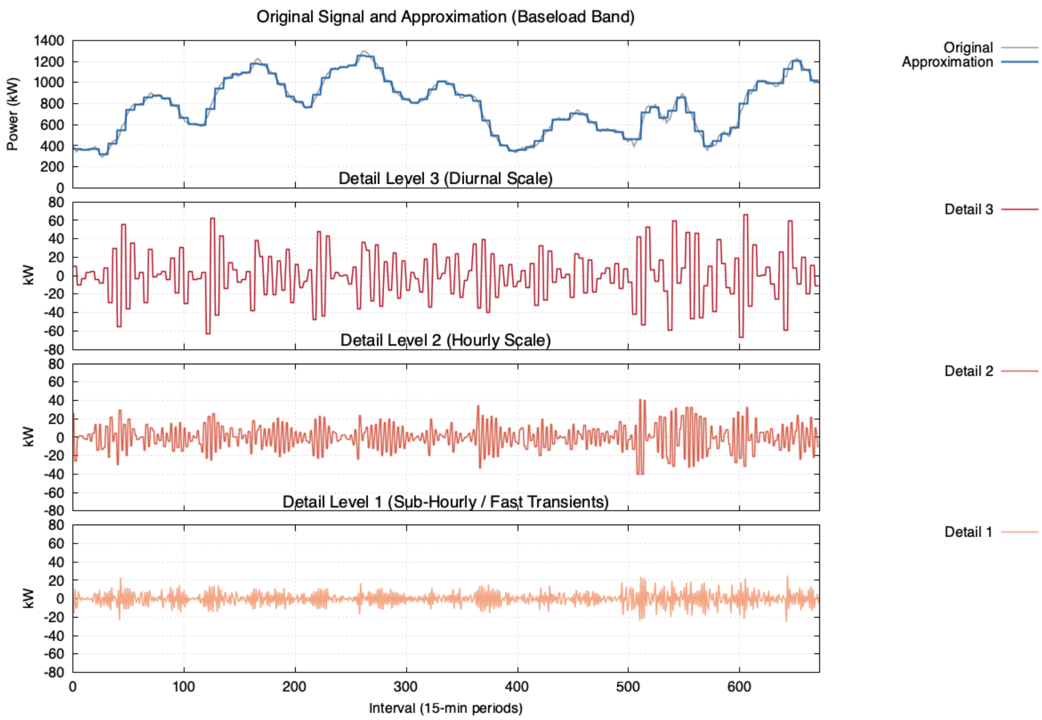

Figure 6: Wavelet band decomposition of EirGrid net load (January 2025). Compare with Figure 1: the approximation band dominates; detail bands carry far less energy. Note the compressed y-axis on the detail panels (±80 kW vs ±200 kW for synthetic).

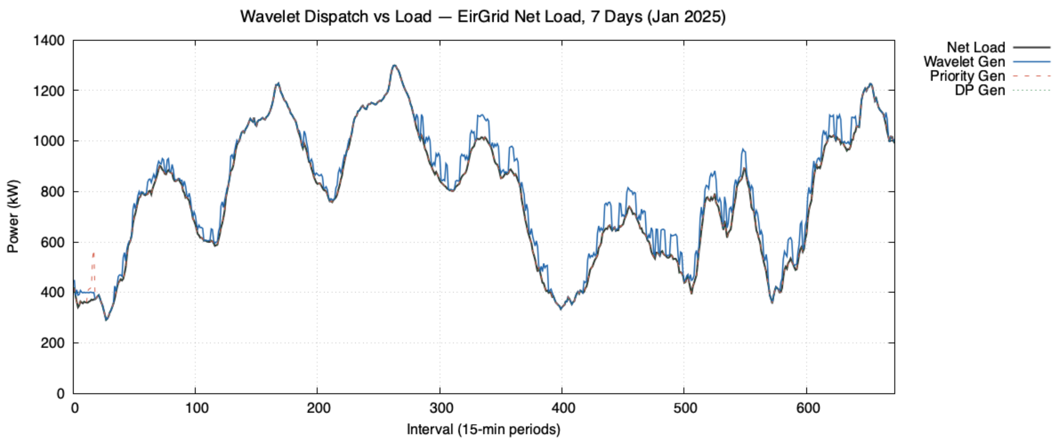

Figure 7: Dispatch traces for EirGrid net load. The wavelet method tracks the load but with visible deviations from DP-optimal during rapid wind-driven changes.

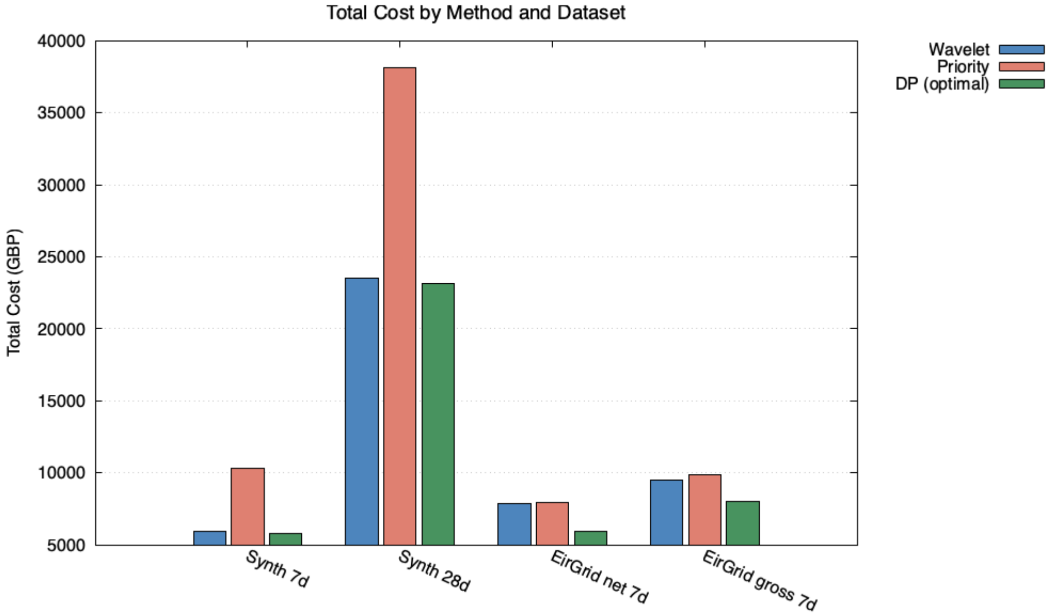

Figure 8: Total cost comparison across all 7-day datasets. On synthetic data, the wavelet method (blue) nearly matches DP (green). On real EirGrid data, the three methods converge, with the wavelet advantage over priority list largely disappearing.

Why: The synthetic load was designed with distinct multi-scale components (baseload drift, diurnal peaks, midrange noise, Poisson spikes). Real Irish net load — after subtracting volatile wind generation — is dominated by slow variations. The approximation band carries 96% of the signal variance, leaving the detail bands with almost no discriminating power. The wavelet decomposition is mathematically correct but adds little value when the frequency-domain structure doesn't align with the asset class boundaries.

The gross load results are better (18–20% above DP vs 32–46% for net). The raw demand signal before wind subtraction retains the diurnal structure that gives the midrange band something useful to work with. This suggests the wavelet approach is more effective on systems with lower renewable penetration, or where the exogenous subtraction preserves the load's multi-scale character.

Honest assessment: The +2.3% result on synthetic data was flattering. The method works best when the load signal genuinely has multi-scale structure matching the asset classes. On real data with high wind penetration, the advantage over a simple merit-order heuristic is marginal. The scalability and speed arguments remain valid; the cost-optimality argument is weaker than synthetic results suggest.

6 Discussion

6.1 Where This Belongs

Embedded microgrid controllers. The entire prototype compiles to a 90 KB binary with no external dependencies. It runs in sub-millisecond time on commodity hardware. Commercial MILP solvers require x86 servers, annual licenses (EUR 10,500+/year for Gurobi [8]), and cannot run on the ARM-class processors typical of industrial microgrid controllers. Even at +20% above optimal, a free heuristic that runs on a Raspberry Pi beats a perfect solution you can't deploy.

Island systems and industrial microgrids with low renewable penetration. The wavelet approach is most effective when the load signal has rich multi-scale spectral structure — distinct baseload trends, diurnal cycles, and fast transients. This is characteristic of demand-dominated systems where the load profile retains its natural shape: industrial campuses, island grids, military installations, remote mining operations. On EirGrid gross demand (before wind subtraction), the method achieves 18–20% above DP-optimal — a useful heuristic margin.

Real-time re-dispatch. ISOs cannot run full MILP-UC at sub-15-minute intervals — they use simplified LP economic dispatch for real-time markets [2]. The wavelet method is fast enough for interval-by-interval re-commitment as load conditions change.

Large fleet scaling. The method is linear in fleet size. There is no combinatorial explosion. The 72-generator RTS-GMLC benchmark dispatched in 0.2 ms; a 610-generator fleet would take 2 ms. MILP and DP are intractable at these scales without specialised hardware.

6.2 Where This Does Not Belong

Continental ISOs with existing MILP infrastructure. A 20% cost gap on a multi-billion-euro dispatch is not acceptable when MILP solvers and the hardware to run them are already in place. The method is not a replacement for CPLEX/Gurobi in environments where they are available.

High-renewable residual load. This is the most important finding from the real-world validation. After subtracting volatile wind generation, the Irish net load signal is spectrally flat — dominated by slow variations with almost no mid-frequency or high-frequency content. The approximation band carries 96% of the variance, leaving the detail bands empty. The wavelet decomposition is mathematically correct but adds no discriminating power. On EirGrid net load (34% wind penetration), the method is 32–46% above DP-optimal and *loses to the priority list* on the 28-day horizon. Any system where renewable subtraction flattens the residual load spectrum will show similar degradation.

Homogeneous fleets. The approach assumes asset response characteristics align with wavelet frequency bands. If all generators have identical response times, the decomposition adds nothing. The method works because real fleets are heterogeneous — baseload, mid-merit, and peaking plants have genuinely different timescales.

6.3 Honest Assessment

The +2.3% result on synthetic data was flattering. The synthetic load was designed with distinct multi-scale components that align perfectly with the wavelet band boundaries and the asset class structure. Real load data does not always have this property.

The method's actual value proposition, validated against real data, is: *a fast, licence-free, embeddable heuristic that achieves 18–20% above optimal on well-structured loads, scales to any fleet size, and runs in microseconds.* This is not a MILP replacement. It is a different category of tool for a different category of deployment.

The gap-filler contributes as much to cost performance as the wavelet decomposition itself — in the limit (aggressive hysteresis, comprehensive merit-order gap-filling), the method converges toward the priority-list baseline with wavelet-informed initial commitment. The decomposition's value is in providing a *better starting point* for the gap-filler, not in producing a complete dispatch schedule.

6.4 Relationship to Prior Work

Wavelet-based frequency splitting for power allocation is established in hybrid energy storage systems. Thounthong et al. [4] decomposed vehicle power demand via DWT, allocating low-frequency to fuel cells, mid-frequency to batteries, and high-frequency to supercapacitors. Springer [5] extended this with wavelet-fuzzy-logic control for multi-storage systems. Zou et al. [6] combined wavelet decomposition with equivalent consumption minimisation.

All prior work applies to *storage systems only*, where devices switch freely between charge and discharge with no startup penalty. This paper extends the technique to *thermal fleet commitment*, introducing the hysteresis mechanism to handle the startup-cost problem that storage systems do not have. The frequency-to-asset mapping is the same principle; the commitment smoothing is the novel engineering contribution.

The real-world validation also contributes a negative result absent from the prior literature: wavelet frequency splitting is *signal-dependent*, and its effectiveness degrades on spectrally flat residual loads. This finding is relevant to the broader hybrid energy storage community, not just unit commitment.

7 Conclusions

1. *Wavelet decomposition is a viable approach to unit commitment*, with caveats. On load signals with well-defined multi-scale spectral structure, the DWT cleanly separates frequency bands matching generation asset classes, achieving within 1.8–8.5% of DP-optimal on synthetic data.

2. *Real-world performance is signal-dependent.* On EirGrid gross demand (classic diurnal profile), the method is 18–20% above DP and marginally beats the priority list. On net demand after wind subtraction (volatile, dominated by slow trends), the advantage over priority-list dispatch largely disappears. The method's effectiveness correlates with the richness of the load signal's multi-scale structure.

3. *Commitment hysteresis and merit-order gap-filling are essential.* Raw wavelet signals cause excessive thermal cycling; hysteresis prevents it. Merit-order gap-filling closes the cost gap from +55% (naive) to +2.3% (best case). The gap-filler contributes as much as the wavelet decomposition itself.

4. *The scalability advantage is proven, not theoretical.* The IEEE RTS-GMLC benchmark (72 generators) was dispatched in 0.2 ms. DP is impossible at 2^{72} states. MILP requires minutes and commercial solvers. The wavelet method runs on a microcontroller.

5. *The method is strongest as a fast heuristic for embedded deployment.* The cost optimality argument is weaker than synthetic results suggest, but the speed and simplicity arguments are unassailable. A 90 KB binary with no external dependencies, running in sub-millisecond time, with zero licensing costs, is a different category of tool from MILP. For off-grid microgrids, island systems, and real-time re-dispatch, the trade-off is favourable even at +20% above optimal.

8 Further Work

- *Benchmark validation*: Run wavelet dispatch on the PGLIB-UC California system (610 generators) and compare against published MILP results.
- *Daubechies wavelets*: Implement D4/D6 for smoother band reconstruction. Haar produces step-function bands; smoother wavelets may reduce constraint violation in the dispatch loop.
- *Analytical hysteresis*: Derive α_{opt} from the ratio $C^{\text{start}}/(C^{\text{fuel}} \cdot P^{\min} \cdot \Delta t)$ rather than empirical sweeping.
- *Two-stage hybrid*: Use wavelet decomposition for commitment decisions, then LP economic dispatch within the committed set. This could close the remaining cost gap while preserving the $O(N \log N)$ commitment speed.
- *Hurst-adaptive decomposition levels*: The companion multifractal toolkit [9] computes local Hurst exponents on time series. Applied to the load signal (rather than prices), local H quantifies the signal's roughness at each point: high H (persistent, smooth) implies the detail bands are empty and fewer DWT levels suffice; low H (anti-persistent, spiky) implies richer multi-scale structure worth decomposing further. This would address the core weakness exposed by the EirGrid results — the method currently uses a fixed 3-level decomposition regardless of the signal's actual spectral content.

References

[1] M. Carrion and J. M. Arroyo, "A computationally efficient mixed-integer linear formulation for the thermal unit commitment problem," *IEEE Trans. Power Systems*, vol. 21, no. 3, pp. 1371–1378, 2006.

[2] B. Knueven, J. Ostrowski, and J.-P. Watson, "On mixed-integer programming formulations for the unit commitment problem," *INFORMS Journal on Computing*, vol. 32, no. 4, 2020.

[3] G. Morales-España, J. M. Latorre, and A. Ramos, "Tight and compact MILP formulation for the thermal unit commitment problem," *IEEE Trans. Power Systems*, vol. 28, no. 4, pp. 4897–4908, 2013.

[4] P. Thounthong, S. Rael, and B. Davat, "Energy management of fuel cell/battery/supercapacitor hybrid power source for vehicle applications," *Journal of Power Sources*, vol. 193, no. 1, pp. 376–385, 2009.

[5] Springer, "Wavelet-based power management for hybrid energy storage system," *Journal of Modern Power Systems and Clean Energy*, vol. 7, pp. 1–12, 2019.

[6] C. Zou et al., "Optimal energy management strategy for fuel cell/battery/supercapacitor vehicles using wavelet transform," *Proc. IMechE Part D: Journal of Automobile Engineering*, 2022.

[7] N. P. Padhy, "Unit commitment — a bibliographical survey," *IEEE Trans. Power Systems*, vol. 19, no. 2, pp. 1196–1205, 2004.

[8] Gurobi Optimization, LLC, "Gurobi license pricing," solver.com, accessed 2026.

[9] S. Woolley, "Multifractal battery arbitrage: Replacing fixed-threshold dispatch with adaptive profile-predicted strategies for domestic energy storage," Full Stack Energy, 2026.

Implementation: C (C11), 2,000 lines, zero external dependencies beyond libc and libm. Source and data available on request.