The Grid That Cannot Hold

What Is Wrong with the Electricity System, the Thinking Errors Behind It, and the Simple .Net Alternative

J. Konstapel, Leiden 31-8-2026.

Introduction

In August 2026 five major banks financed a 200-megawatt battery in Winschoten, in the north of the Netherlands. The battery does not generate anything. It absorbs electricity when sun and wind produce more than the grid can swallow, holds that charge for at most four hours, and releases it when demand peaks. Hundreds of millions of euros are being spent, in other words, on a machine whose only function is *to wait*.

That fact deserves more attention than it gets. A system that must buy the capacity to wait, at enormous cost and as an afterthought, is a system that was designed without that capacity — and a system designed without the capacity to wait is a system designed around a false picture of how flowing systems work. This essay reconstructs that false picture, names the thinking errors it produces, and then derives the alternative from first principles. The framework used for the derivation is the Vacuum.Net theory, whose primitive is a single self-crossing strand from which closures, layers and cycles are derived. No prior knowledge of that theory is assumed; every concept used here is introduced where it is needed.

The argument, in one paragraph: a healthy flowing system has four states — loading, holding, discharging, and rest — and moves through them as a cycle with a characteristic retention time. The electricity grid was built with only two of those states, fused into one: what is generated must be consumed *in the same instant*. Everything that is now going wrong — negative prices, curtailed wind parks, congested cables, emergency batteries, heat pumps that overload the grid on the coldest evening of the year — follows from that single missing capability and from a series of thinking errors that all share one root: confusing the *amplitude* of a flow with the *retention* of a state, and confusing the historically contingent *address* of the system with a law of nature.

Part I — The System as Built

A machine with two fused states

Alternating-current electricity cannot be stored as electricity. At every moment, generation must equal consumption plus losses; the balance is verified fifty times per second, because the grid frequency — 50 Hz in Europe — rises when there is surplus and falls when there is deficit. The entire operational apparatus of the grid, from balancing markets to spinning reserve, exists to keep that one number on its reference value.

Notice what this means structurally. In a complete cycle, *loading* (tension builds), *holding* (tension is retained), *discharging* (tension is released) and *rest* (the system relaxes) are distinct phases separated in time. The grid collapses the first three into a single instant: the loading of the system by a power plant **is** its discharge into a kettle somewhere, simultaneously, with the wires as nothing but geometry in between. There is no holding phase. There is no rest phase. The system is a machine for *instantaneous, obligatory discharge*.

For a century this did not matter, because the sources were controllable. Coal, gas, and hydro plants hold their energy upstream of the grid — in the coal pile, the gas field, the reservoir. The holding state existed, but it was hidden inside the fuel. The grid could afford to be a pure discharge machine because the *molecules feeding it* did the waiting.

The pulse changed sides

The energy transition replaces controllable sources with sun and wind. These are not merely "variable"; they are structurally different in kind. A gas plant discharges when the operator decides. The sun discharges when the sun decides. The pulse — the moment and magnitude of energy arrival — has moved from inside the system's control to outside it.

This is the precise point at which the missing states become fatal. A discharge-only machine fed by controllable sources works, because control substitutes for retention. A discharge-only machine fed by an external pulse cannot work, because an external pulse cannot be commanded — it can only be *received*, and receiving requires exactly the capability the grid lacks: holding. Hence Winschoten. Hence the negative electricity prices on sunny Sundays, which are the market screaming that discharge is being forced on a system with nowhere to put it. Hence curtailment, the deliberate destruction of free energy, which is what a system does when it can neither use nor hold what arrives.

The battery is therefore not an innovation. It is a prosthesis: the retrofitted hold-state of a system that was born without one.

The hidden second network

Meanwhile, roughly half of final energy demand — in the Netherlands and in most temperate countries — is not electricity at all but heat: space heating, hot water, industrial process heat. The current policy consensus answers this with a single word: electrify. Heat pumps in every home, electric boilers in industry, everything onto the one grid.

We will examine below why this is not a solution but a multiplication of the original error. For now, note only the shape of the move: a system that already cannot hold its own flows proposes to absorb the *other half of national energy demand* — the half with the strongest seasonal swing and the sharpest evening peaks — onto the same discharge-only machine.

Part II — The Thinking Errors

The dysfunction described above is not an engineering accident. It follows from identifiable errors of thought, each of which can be stated precisely. There are seven.

Error 1: Amplitude thinking — mistaking capacity for retention

When the grid congests, the reflex is to lay thicker cables: raise the amplitude the system can carry. But congestion in a renewable grid is not primarily an amplitude problem; it is a *timing* problem. The energy arrives at the wrong moment relative to the demand. The relevant quantity is not megawatts but the **retention time τ** — how long the system can hold a state between loading and discharge — and the grid's τ is zero.

A thicker cable with $\tau = 0$ simply delivers the mismatch faster. The distinction matters because the two quantities scale differently: amplitude problems are solved with copper, retention problems are solved with structure. Confusing them means spending copper on a structural deficiency, which is exactly what grid-reinforcement programs across Europe are doing. The correct diagnostic question for any proposed investment is not "how much can it carry?" but "how long can it hold, and how asymmetric are its loading and discharge times?" A battery that loads slowly across a sunny afternoon and discharges fast into an evening peak has an asymmetric retention structure — and that asymmetry, not its megawatt rating, is what the system actually buys.

Error 2: The single time scale

The grid is operated as if it had one characteristic frequency: the real-time balance. But the mismatches it must absorb live on at least four distinct time scales: seconds (frequency stability), hours (the daily solar cycle against the evening peak), days (wind fronts passing), and months (winter heat demand against summer solar surplus).

A system with one instrument on one time scale cannot see, let alone treat, the others. The Netherlands is currently building hour-scale batteries in quantity — the Winschoten class — while the largest mismatch of all, the *seasonal* one, has no instrument whatsoever: there is no month-scale storage in the plan at any meaningful size. This is not because month-scale retention is exotic; it is because a two-state mental model has no slot in which a "seasonal hold" could even be represented. What the model cannot represent, the planner cannot miss.

The requirement that follows is not "more storage" but a **ladder of retention times**: flywheels at seconds, batteries at hours, pumped and thermal storage at days, molecules at months. Each rung is a different instrument for a different frequency. A single rung, however large, leaves every other frequency untreated.

Error 3: Averaging a context-dependent system with context-free instruments

Grid planning, market design and policy all operate on averages: average load profiles, average household consumption, capacity factors, standard scenarios. An average is a context-free instrument — it parses every situation with the same formula, blind to local structure.

But the emerging grid is context-*dependent* in an exact sense. A neighbourhood with thirty solar roofs, two heat pumps and one fast charger is not a scaled average household; its behaviour is dominated by local, discrete, correlated events. There is a boundary — analogous to the mean free path in a gas — below which the statistical description stops being valid, because individual trajectories rather than collective averages carry the state. Recent work in classical statistical mechanics has demonstrated this sharply: in regions narrower than the particle mean free path, flux ratios are provably non-Boltzmannian — the averaging law does not bend there, it *fails* there, and no choice of boundary conditions repairs it. The lesson generalizes: **a context-free instrument cannot parse a context-dependent system.** Not "parses it poorly" — cannot, as a matter of structure.

Applied to the grid: every planning tool built on averaged profiles is silently crossing that validity boundary as decentralized generation grows. The tools do not degrade gracefully; they produce confident numbers about a system whose actual state they can no longer represent. The negative prices and local congestion that "surprise" the models every year are not outliers. They are the residue of a morphological mismatch between instrument and system.

Error 4: Transporting the ground state

Every energy carrier has a height, its exergy: the fraction of it convertible into work. Electricity stands at the top — fully convertible into anything. Low-temperature heat stands at the bottom; it is where all conversion losses end up, the ground state of the energy system, the energetic equivalent of rest. This ladder has a transport rule that follows directly from it: **transport high, discharge low, and discharge at the destination.** High carriers travel well — electricity loses a few percent over hundreds of kilometres, molecules cross oceans by ship at almost no loss en route. The ground state travels catastrophically: heat leaks toward everywhere by its very nature, so every kilometre of district-heating pipe is a continuous discharge into the soil, on the order of a fifth to a third of the throughput, delivering lukewarm water that can do nothing except be lukewarm — a carrier with no convertibility, chained to one use and one supplier.

District heating at city scale therefore inverts the transport rule: it discharges energy to the bottom of the ladder first, and then tries to *move the bottom*. The only defensible heat flows are measured in metres, not kilometres — waste heat captured within the same building or plant where it is released, where the "transport" is a wall. Beyond that boundary, the physics of the ground state wins, always.

Error 5: The laboratory efficiency fallacy — the heat pump

The heat pump is promoted on its coefficient of performance: three to four units of heat per unit of electricity, measured under test conditions. The system-level analysis reads differently, and the difference is instructive because it is a *timing* error, not a thermodynamic one.

The COP falls as the outdoor temperature falls, because the machine must pump heat over a larger gap — precisely when heat demand peaks. On the cold, dark, windless winter evening that sizes the entire system, an air-source heat pump delivers a COP near two, supplemented by resistive heating, fed by a grid whose marginal winter source is a gas plant running at roughly fifty percent conversion loss. The full chain — gas to electricity to transport to a degraded heat pump — arrives at the radiator with little more than the gas boiler achieved by burning the molecule directly at the destination, but now with grid reinforcement, conversion losses, and a new synchronized national peak added on top.

The deeper error: mass heat-pump deployment does not *shift* discharge in time, it *synchronizes* it — every unit responds to the same weather at the same hour. A discharge-only grid, already unable to hold, is being handed a demand layer engineered for maximum simultaneity at minimum source availability. The instrument is judged by its behaviour in isolation; the system is destroyed by its behaviour in chorus.

Error 6: Mistaking the address for a law of nature

Some properties of the grid are frozen decisions; others are physics. The system's builders and inheritors routinely confuse the two. Call the frozen layer the **address**: the choices fixed at the system's founding — alternating current, centralized generation, one national synchronous machine, radial distribution — which no amount of operation ever revisits. Call the live layer the **tension**: the actual flows, imbalances and prices of the moment. Address is topology, fixed at closure; tension is dynamics, treatable daily. The two must never be confused, because treatments aimed at the wrong layer always fail: you cannot solve an address problem by managing tension harder.

The clearest example is alternating current itself. AC defeated Edison's DC in the 1890s for exactly one reason: only AC could then be transformed to high voltage for long-distance transport. That reason has evaporated — modern power electronics convert DC voltages trivially, and the longest, most efficient transmission lines in the world today are DC. Meanwhile virtually every device in a modern home is internally DC: LEDs, electronics, chargers, solar panels, home batteries, cars. Solar power is generated as DC, inverted to AC to enter the house, rectified back to DC to enter the battery or the laptop — three conversions where zero were needed, each one a small discharge to the ground state, multiplied across every device in every building.

The AC address was correct for its century and is wrong for this one, but because it is mistaken for physics rather than recognized as a revisable founding choice, it is defended instead of examined. The same holds for centralization itself. History proves the address is contingent: before electricity won, factories distributed power mechanically through shafts and belts, Paris ran a compressed-air network, London a hydraulic one, and Manhattan runs district steam to this day. Electricity prevailed because it is the most convertible carrier — the highest rung — not because a single synchronous national machine is the only possible architecture. The steam engine that started it all was, ironically, structurally complete: boiler loading, pressure holding, piston discharge, condenser rest — the full four-state cycle in iron. The grid that replaced it kept the discharge and threw the other three states away.

Error 7: Believing the pulse can be imposed

The operational philosophy of the grid is central command: a control room commands plants, balances the machine, imposes the rhythm. With controllable sources this works, because imposition and control are the same act. With an external pulse it fails by definition — the sun takes no dispatch orders — and the reflex response, ever-finer central management of ever-more devices, fights the structure of the problem rather than accepting it.

There is a mathematical result lurking here that the engineering culture has not absorbed: in systems of coupled asynchronous sources, the corrective rhythm is *intrinsic*. It emerges from the structure of the ensemble and cannot be externally imposed — imposition produces oscillation and instability, not order. The practical translation: a grid of millions of small sources, holds and flexible loads must be designed as a **swarm** whose global smoothness emerges from local asynchrony, not as one machine with ten million actuators on a single leash. Breadth and desynchronization across a distributed ensemble carry more stability than any individual component's controllability — the sum of many uncorrelated pulses is smoother than any member. Central imposition, by synchronizing responses, destroys exactly the asynchrony that stabilizes.

Part III — The .Net Alternative

The alternative is not a technology package. It is a small set of structural rules from which the technology choices follow. Stated compactly:

Rule 1 — Complete the cycle. Every level of the system, from device to dwelling to district to nation, must possess all four states: load, hold, discharge, rest. A level without a hold state is incomplete and will export its incompleteness upward as instability. The battery, the buffer vessel, the insulated mass, the fuel store are not accessories; they are the third state, without which the cycle cannot close.

Rule 2 — Build the τ-ladder, not the τ-point. Retention is needed at every native frequency of the mismatch: seconds, hours, days, months. Each rung is its own instrument with its own physics — flywheels, batteries, thermal mass, molecules — and no rung substitutes for another. The current build-out, clustered entirely on the hour rung, leaves the season rung empty; the ladder rule makes that gap visible as a gap.

Rule 3 — Two layers, never mixed. Distinguish rigorously between the address (the frozen founding choices: current type, topology, centralization) and the tension (the live flows). Treat tension daily; *revisit the address explicitly* whenever its founding reason has died, as AC's has. A system that cannot name its own address cannot know which of its problems are structural.

Rule 4 — Close locally; minimize chaotic transitions. Every conversion between carriers at a wide, thermalizing boundary discharges part of the flow to the ground state. The losses are a property of the *boundaries*, not of the distances: it is the number and morphology of transitions that leaks, and boundary morphology decides what can pass at all. Therefore: the shortest loop wins. Generate where you consume, hold where you generate, capture waste heat within the same closure that produces it (metres, not kilometres), run one DC spine through the building instead of thirty adapters. A house that closes its own cycle and exchanges only its residue with the grid removes an entire cascade of conversions — and removes its contribution to the synchronized peak at the same time.

Rule 5 — Match the carrier to the demand's height. Half of demand is heat; heat is the bottom of the ladder; answering bottom-rung demand with the top-rung carrier squanders the height difference twice — once in making the electricity, once in throwing its convertibility away in a radiator. Heat demand is answered first by not existing (insulation is the only source with zero conversions: tension never created never needs discharge), then by heat made *at the destination* from a carrier that transports well, then by seasonal heat held in the ground *under the building it serves*. Transport high, discharge low, discharge on site.

Rule 6 — The reference stands beside the measurement, never inside it. The grid already contains, in embryo, its most important design pattern: the 50 Hz reference. Every deviation is measured against a fixed reference that the deviations themselves never alter — the blueprint stands beside the measurement, so every instant is a scored prediction. Generalize that pattern to every level: each closure (house, district, region) carries its own fixed reference state against which its live tension is read and treated. And never let the frequency of an event determine the placement of the structure: how often a peak occurs must not decide where the buffer lives — the address is computed from the structure of the demand, not from its statistics.

Rule 7 — Let the pulse be intrinsic. Design the ensemble so that stability emerges from asynchrony instead of being imposed by command. Concretely: maximize the diversity of source rhythms (wind and sun across uncorrelated regions, connected high on the ladder by DC where it must travel at all), give every closure its own hold so that no closure is forced to discharge on the ensemble's worst moment, and let flexible loads follow local surplus instead of central dispatch. The control room's role shrinks to guarding the references — which is the one thing a centre *can* do that a swarm cannot.

What this looks like

None of these rules requires undiscovered physics. Together they describe a recognizable system: buildings as small closures with DC spines, local generation, hour-scale electrical and day-scale thermal holds, and ground-coupled seasonal storage beneath them; districts as swarms of such closures exchanging residues; a national layer that transports only high on the ladder — DC electricity and molecules — and holds the season in molecules and in the ground; heat that never travels; and a reference hierarchy in place of a command hierarchy. The Winschoten battery fits inside this picture, but demoted from headline to rung: one hold, on one rung of one ladder, of one closure among many.

The deepest point is the simplest. The current system is not failing because renewable energy is difficult. It is failing because a two-state machine is being asked to receive a four-state world. The pulse now comes from outside, and an external pulse cannot be commanded — only received, held, and released in its own time. Everything wrong with the grid is a variation on refusing that sentence; everything in the alternative is a variation on accepting it.

This essay was developed within het MAZE-project, the applied implementation of the Vacuum.Net theory. The four-state cycle (load, hold, discharge, rest), the two-layer law (address is topology, fixed at closure; tension is dynamic and treatable), the retention time τ, the morphology prohibition (context-free instruments cannot parse context-dependent systems) and the intrinsic pulse are derived constructs of that theory; their application to the electricity system is presented here as a self-contained argument requiring no external reference.