THE XYLOMORPHIC PREMISE

Waste Heat, Physical Co-Computation, and the Architecture of Post-Carbon Infrastructure

Flyxion

Independent Researcher

June 2026

Abstract

Contemporary infrastructure design separates computation, heating, storage, and communication into distinct physical systems, each optimized for a single output and each discarding the byproducts of the others as waste. This essay argues that the dominant assumption underlying this architecture—that specialization is always efficient—is not a physical law but a design choice, and that the energy transition now underway creates the conditions under which a different choice becomes both technically viable and economically compelling. Drawing on the concept of xylomorphic architecture, which takes the multi-functional substrate of forest systems as its organizing analogy, the essay develops the proposal in two stages. The first is thermal integration: a distributed compute network in which residential GPU heater nodes co-locate computation with heat demand, using the existing fiber topology as the addressing layer and treating climate as a scheduling parameter. The second, more radical stage is physical co-computation: the recognition that pressurized fluid loops, phase-change materials, optical waveguides, and thermal diffusion layers are not merely cooling systems but computational substrates in their own right, capable of performing fluid optimization, matrix operations, smoothing, state storage, and spectral decomposition as intrinsic features of the same physical processes that generate and transport heat. In the physical co-computation architecture, the heater does not merely consume computation and produce heat; computation, heat transport, fluid flow, optical propagation, and thermal storage are manifestations of the same underlying physical substrate engineered together. The essay connects this proposal to the historical tradition of analog and physical computing—differential analyzers, hydraulic computers, optical processors, neuromorphic chips—and argues that the xylomorphic node represents a synthesis of that tradition with building infrastructure: a thermodynamic organism whose heating, sensing, communication, and computation are coupled processes occurring within the same physical object.

> *The forest does not heat itself.*
> *It computes the distribution of light, water, mineral,*
> *and structural load across a million simultaneous processes,*
> *discarding nothing that can be used downstream.*
> *We call this ecology.*
> *We could call it engineering.*
>
> *—field notes on xylomorphic substrates*

1. The Separated Substrate Problem

Modern infrastructure is organized around the principle of functional separation. A power plant produces electricity. A data center consumes electricity and produces computation and heat. A residential heating system consumes electricity or gas and produces heat. A refrigerator consumes electricity and produces cold, rejecting heat to the surrounding air. These systems are designed, sited, financed, regulated, and operated independently of one another, and their interactions are treated as externalities rather than as engineering parameters. The heat that a data center exhausts to the atmosphere is not connected, in the design of either system, to the heat that a nearby home requires to remain habitable through a northern winter. They are coincident in space and complementary in function, and yet the infrastructure that serves them treats their relationship as irrelevant.

This separation is not a consequence of physical law. It is a consequence of institutional history: the sequence in which electrical grids, telecommunications networks, computing infrastructure, and residential energy systems were developed and regulated as distinct sectors, each with its own capital structures, ownership models, regulatory jurisdictions, and engineering cultures. The separation made sense in a world where each function was modest enough in scale, and cheap enough in marginal energy cost, that the inefficiency of discarding complementary outputs was not worth the organizational complexity of recovering them. That world is ending. The global energy demand of computation is now large enough that its waste heat profile constitutes a significant thermal resource, and the energy cost of heating residential and commercial space in cold climates is large enough that recovering that resource would materially reduce total system energy consumption. The question is whether the institutional and engineering architectures that would allow that recovery can be built, and at what speed.

This essay argues that they can, and that the form they are likely to take diverges substantially from the intuitive picture of adding heat exchangers to existing data centers. The more radical and more defensible proposal is a reorganization of where computation happens—a shift from the concentrated, purpose-built data center toward a distributed network of thermally integrated compute nodes embedded in the residential and light-industrial fabric of dense settlements. The organizing principle of this network is not proximity to cheap power or favorable tax treatment, which is what currently determines data center siting, but proximity to heat demand. In climates where space heating constitutes a major energy expenditure for half the year or more, the co-location of computation and heat demand is not merely a marginal efficiency improvement. It is a structural redesign of how energy flows through the built environment.

2. XYLOMORPHIC ARCHITECTURE AND THE MULTI-FUNCTIONAL SUBSTRATE

The term xylomorphic designates a design philosophy derived from the structural and functional organization of arboreal systems. A tree does not optimize for a single output. Its physical substrate simultaneously provides mechanical structure, hydraulic transport, thermal regulation, chemical storage, atmospheric exchange, and biological habitat. Each of these functions is performed by the same material organized at different scales: cell walls, vascular bundles, bark, root networks, canopy geometry. The efficiency of the system derives not from the specialization of each component for a single function but from the degree to which a single physical process generates multiple useful outputs. Metabolic heat produced by cellular respiration is not wasted; it participates in the thermal gradient that drives water transport. The mechanical stress of wind loading is not merely a threat to be resisted; it is a signal that regulates growth allocation. Nothing that can be used is discarded.

Modern infrastructure does not work this way. It is organized around what might be called the single-output principle: each system is designed to maximize performance on one metric while treating all other outputs as constraints to be minimized or externalities to be ignored. A data center is evaluated by computational throughput per unit of power consumed. The heat it produces is a constraint on cooling system design, not an output to be optimized. A residential heating system is evaluated by thermal output per unit of fuel consumed. The computation it does not perform is simply not on its evaluation scorecard because no one has asked that question.

The xylomorphic premise inverts this framing. It asks: given that a physical process must produce heat as a consequence of performing computation, and given that heat is a resource with real economic value in cold climates for a substantial fraction of the year, what would infrastructure look like if it were designed to maximize the sum of useful outputs rather than to maximize one output while discarding the others? The essay develops this question in two stages. The first stage is thermal integration: the co-location of digital computation and heat demand so that waste heat becomes useful heat, treating the existing fiber topology as the addressing layer and climate as a scheduling parameter. This is already a significant departure from current practice. The second stage is more radical, and it is where the xylomorphic analogy becomes most precise: physical co-computation. A tree does not have a central processor that runs thermal regulation as a separate subroutine. The hydraulic dynamics of water transport through the xylem are themselves a computation—the pressure distribution across the vascular network solves an optimization problem at every moment, and the solution is physically instantiated in the flow. The thermal diffusion of metabolic heat through the cambium performs a smoothing operation. The mechanical resonance of the trunk under wind performs a spectral decomposition. The chemistry of the root tip performs local gradient descent in a nutrient field. None of these are computations running on a separate substrate that happens to produce heat as a byproduct. They are computations occurring in the same physical substrate that performs heating, transport, and structure. The xylomorphic node, in its full development, aspires to the same integration.

3. THE WASTE-HEAT FIELD

To formalize the proposal, define the waste-heat field $Q_C(x,t)$ as the heat generated by computation at location x and time t, and the heat demand field $Q_H(x,t)$ as the space-heating requirement at location x and time t. In conventional infrastructure, computation is sited without reference to Q_H: the relationship between the two fields is approximately:

$$Q_C \perp Q_H,$$

meaning that computational heat generation and residential heat demand are statistically independent by design. Computation happens where data center economics are favorable. Heat demand happens where people live. The two fields overlap by coincidence rather than by intent, and the overlap that does exist is not recovered because the institutional boundary between the compute sector and the heating sector makes recovery organizationally expensive even when it is physically straightforward.

It is useful to distinguish three stages of infrastructure evolution to clarify precisely where the xylomorphic proposal sits. **Stage 1** is the current architecture of separation: Computer $\perp$ Heater. Computation and heat are independently optimized within distinct institutional domains, and the waste heat produced by computation is discarded. **Stage 2** is waste-heat recovery: computation happens where economics dictate, and a secondary system—a district heating loop, a heat exchanger, a hot-water network—captures the exhaust and transports it to where it has thermal value. Most existing proposals for using server heat in buildings operate at Stage 2, and while they represent genuine improvements over Stage 1, they preserve the fundamental separation between the compute and heating optimization domains. The data center operator still optimizes for throughput; the building operator still optimizes for thermal comfort; the recovered heat is a secondary output grafted onto a primary system that was designed without it. **Stage 3** is integration: Computer $+$ Heater $=$ Single Infrastructure Object. The compute node and the thermal device are the same physical object, co-located with heat demand, governed by a shared scheduler that treats computational throughput and thermal output as simultaneous first-class outputs of the same process. This is the xylomorphic proposal. Many readers will initially interpret it as a Stage 2 proposal because Stage 2 is the familiar category; making the distinction explicit is essential to prevent the idea from being collapsed into conventional waste-heat recovery.

The xylomorphic proposal is to replace the Stage 1 independence condition with a Stage 3 co-optimization objective. Define the ideal recoverable overlap of the two fields as:

$$\Omega(Q_C, Q_H) \;=\; \int_{\mathcal{X}} \min\bigl(Q_C(x,t),\; Q_H(x,t)\bigr)\, dx.$$

The quantity Ω measures the computational heat that would satisfy real heating demand if every node were perfectly thermally integrated with its surroundings. In practice, thermal integration is imperfect: an apartment tower can distribute heat through a water loop with high efficiency, while a detached house with poor insulation loses heat through the envelope before it reaches the living space. To capture these differences, introduce a spatially varying thermal transport coefficient $\eta(x) \in [0, 1]$ representing how effectively location x converts computational heat into useful thermal work:

$$\Omega_{\text{eff}}(Q_C, Q_H) \;=\; \int_{\mathcal{X}} \eta(x)\, \min\bigl(Q_C(x,t),\; Q_H(x,t)\bigr)\, dx.$$

The coefficient $\eta(x)$ is high for well-insulated buildings with hydronic systems, intermediate for forced-air systems with duct losses, and low for locations with poor thermal coupling between hardware and conditioned space. This formulation allows apartment towers (high η, high heat density), detached houses with thermal storage tanks (moderate η with buffering), and industrial facilities requiring process heat to appear naturally in the optimization. The effective co-optimization objective is then:

$$\max_{Q_C}\; \Omega_{\text{eff}}(Q_C, Q_H) \quad \text{subject to} \quad \int_{\mathcal{X}} Q_C(x,t)\, dx \;\geq\; D(t).$$

This is a spatial scheduling problem whose solution is a distributed network of compute nodes whose locations and load allocations are determined jointly by computational demand, thermal demand, and the local capacity to convert computational heat into useful work. During a northern hemisphere winter, when $Q_H(x,t)$ is large in cold-climate residential districts and η is favorable for well-sealed buildings, the optimizer directs computational load toward those districts. During summer, when Q_H is small or negative, computational load migrates toward industrial facilities where η remains high for process heat, or is shed to latency-tolerant workloads that can wait for the next thermally favorable period.

4. THE FIBER TOPOLOGY AND THE EXISTING DISTRIBUTION NETWORK

The proposal just described may appear to require the construction of an entirely new communications infrastructure to connect residential compute nodes to the broader network. In fact, much of that infrastructure already exists. The fiber optic cables that carry internet traffic into residential neighborhoods, through street-level routers, and in many jurisdictions directly into homes already constitute a high-bandwidth, low-latency distribution network whose topology is not centralized but diffuse. Fiber runs through cities, subdivisions, and apartment buildings in patterns that roughly track the distribution of people rather than the distribution of industrial land or cheap power. The networking layer, in other words, already has the topology that a distributed thermal compute network requires.

This observation reframes the proposal considerably. The question is no longer whether a distribution network for residential compute nodes can be built; it is whether the existing distribution network can carry a different kind of traffic. A packet of data and a job scheduled to a compute node are structurally similar objects from the perspective of the network: both are addressed, both are routed, both have latency requirements, and both consume bandwidth proportional to their size. The fiber that currently carries a video stream to a household could, without modification to its physical layer, carry a compute job to a processor in that same household and return the result. The router that currently sits in a street cabinet or a building basement already performs packet scheduling. Extending that scheduler to include thermal state as a routing parameter is a software problem, not a hardware construction problem.

More precisely: the topology mismatch identified in the previous section is between centralized computation and distributed heat demand:

$$D(x,t) \;\rightarrow\; \text{few centralized nodes}, \qquad H(x,t) \;\rightarrow\; \text{millions of local nodes}.$$

The existing fiber infrastructure demonstrates that distributed addressing of individual homes is already solved at the communication layer. The xylomorphic proposal is to add a computational layer that follows the same topology:

$$D(x,t) \;\approx\; H(x,t),$$

routing computation to where heat has value just as the network currently routes data to where attention has value. Each residence becomes an addressable compute node in the same sense that it is already an addressable network node. The household IP address becomes, in this architecture, also a thermal address: a location with a known heating profile, a known fiber connection, and a compute node whose load can be adjusted by the scheduler in response to both computational demand and thermal demand.

Let $C_i(t)$ denote the computation allocated to household i at time t, $D_i(t)$ the computational demand available to be dispatched to that household (constrained by its connection bandwidth and processor capacity), and $H_i(t)$ the household's heat demand in compatible units. The scheduler assigns:

$$C_i(t) \;=\; \min\bigl(D_i(t),\; H_i(t)\bigr),$$

meaning that computation is dispatched to a household up to the lesser of what the network can send and what the household's thermal state can absorb. When heat demand is high, the node runs at full capacity and the household receives free heat. When heat demand is low—a warm spring afternoon, a well-insulated house with residual solar gain—the node enters an idle state: maintaining network synchronization, preserving memory state, and remaining available for dispatch, but generating only enough heat to remain warm in the technical sense of five degrees above its minimum operating temperature. The household is not running a server farm; it is running a standby node that occasionally becomes a heater when the scheduler determines that the combination of local heat need and available latency-tolerant workload makes dispatch worthwhile.

The idle state is important because it determines the baseline energy cost of network participation. A node that must run at full load continuously to remain in the network imposes a heating load whether or not the household needs heat, which is exactly the problem the proposal is trying to solve. A node that can run at minimal power in idle state and ramp to full load on scheduler instruction has a very different energy profile: its average consumption tracks heat demand rather than overriding it. The technical requirement is a processor architecture that supports wide dynamic power range—deep idle states drawing single-digit watts, full compute states drawing hundreds—with fast transition times so that a scheduler instruction can bring a node from thermal standby to full computational output within the time window of a latency-tolerant job. This requirement is consistent with existing trends in processor design, where power management has become a primary engineering constraint, but it implies specific design priorities for xylomorphic nodes that differ from the priorities of conventional server hardware.

5. THE GPU HEATER NODE

The physical instantiation of the distributed thermal compute network is what this essay calls the GPU heater node: a self-contained compute unit designed for residential or light-commercial installation, whose primary heat output is directed into the building's thermal envelope rather than exhausted to the exterior. The node performs computation for the network, receives payment for that computation in the form of energy cost offsets or direct compensation, and provides space heating as a co-product.

The proposal to replace conventional heaters—resistance heaters, heat pumps, gas boilers—with computation-performing heat sources is controversial because it immediately invites comparison on narrow efficiency grounds. A modern heat pump achieves a coefficient of performance between two and five in moderate climates, meaning it delivers two to five units of thermal energy per unit of electrical energy consumed. A GPU performing computation at full load achieves a coefficient of performance of exactly one from a thermal standpoint: one unit of electrical energy becomes one unit of heat. On this narrow comparison, the heat pump appears superior.

The comparison is incomplete in two respects. First, the heat pump performs no computation. Its coefficient of performance from a total-output perspective is:

$$\eta_{\text{total}}^{\text{pump}} = \frac{Q_{\text{thermal}}}{\text{energy input}} \in [2, 5], \qquad \eta_{\text{compute}}^{\text{pump}} = 0.$$

The GPU heater node performs computation whose economic value is separable from its thermal output:

$$\eta_{\text{total}}^{\text{node}} = \frac{V_{\text{compute}} + Q_{\text{thermal}}}{\text{energy input}},$$

where V_{compute} is the economic value of the computation performed per unit of energy consumed. If V_{compute} is nonzero—if the computation has real market value—then the total-output efficiency of the GPU heater node exceeds that of the heat pump by the magnitude of V_{compute}. The heat pump is efficient at heating. The GPU heater node is efficient at the joint production of computation and heating. These are different problems, and the second is the more relevant one for an economy in which computational demand is growing at a rate that makes large, purpose-built data centers an increasingly visible fraction of total energy consumption.

Second, the heat pump's efficiency advantage over a resistance heater depends on ambient temperature in ways that erode significantly in cold climates. At outdoor temperatures below $-15°$C, which are common in northern Canada for extended periods, air-source heat pump efficiency approaches that of resistance heating, and some units require supplemental resistance heating to maintain output. The GPU heater node's thermal coefficient of performance is invariant to outdoor temperature: one unit of electricity always produces one unit of heat regardless of ambient conditions, and that heat is produced inside the thermal envelope where it is needed. In the coldest climates and coldest periods—precisely those where heating demand is highest and heat pump efficiency is lowest—the GPU heater node's thermal performance is competitive with or superior to the heat pump on a strictly thermal basis, before the value of computation is counted at all.

A third and more important point concerns the framing of the comparison itself. The strongest version of the proposal is not:

$$\text{Heat Pump} \rightarrow \text{GPU Heater Node}$$

but rather the hybrid system:

$$\text{Heat Pump} + \text{GPU Node} + \text{Thermal Storage} \rightarrow \text{Integrated Thermal-Compute Plant}.$$

The GPU node provides baseline heat whenever the scheduler dispatches computation to it; the heat pump amplifies that heat when outdoor temperatures permit efficient operation; the thermal storage tank in the basement buffers the diurnal mismatch between computational dispatch windows (which track network demand) and heating demand peaks (which track occupancy and outdoor temperature). None of these components displaces the others; they operate as a coupled system whose collective thermal coefficient of performance, weighted by the economic value of computation, substantially exceeds any of its components in isolation.

5.1. *Workload Partitioning and Thermal Schedulability*

Not all computation is equally relocatable. Large AI training jobs have data locality constraints, memory synchronization requirements, and security considerations that limit where they can run. The proposal does not require that all computation migrate into homes—it requires only that enough computation be thermally schedulable to satisfy a meaningful fraction of residential heating demand.

A useful partitioning of computational demand distinguishes four classes:

$$D(t) \;=\; D_{\text{latency}}(t) \;+\; D_{\text{batch}}(t) \;+\; D_{\text{secure}}(t) \;+\; D_{\text{thermal}}(t),$$

where D_{latency} denotes latency-sensitive workloads (interactive services, real-time inference, financial transactions) that must run near their users regardless of thermal geography; D_{batch} denotes latency-tolerant batch workloads (model training, rendering, simulation, genomic processing, archival analytics) that can be deferred by hours or days; D_{secure} denotes workloads with data-sovereignty or security constraints that prohibit residential-node execution; and D_{thermal} denotes workloads explicitly designed for distributed thermal scheduling, including lightweight inference, cryptographic operations, and scientific computing with high per-watt thermal yield. The thermally schedulable fraction of total demand is $D_{\text{batch}} + D_{\text{thermal}}$, which by current estimates constitutes a substantial and growing share of global compute load as AI training and batch analytics workloads expand. The proposal does not need D_{latency} or D_{secure} to be thermally relocatable. It needs only the batch and thermal classes to be large enough to absorb a significant fraction of cold-climate residential heating demand—a constraint that the quantitative analysis in the following section shows is plausibly satisfied.

6. A NEIGHBORHOOD-SCALE THOUGHT EXPERIMENT

To move from conceptual proposal to preliminary engineering proposition, consider a concrete case: a hundred-home residential neighborhood in New Brunswick, Canada, during a typical January. New Brunswick is a useful reference because it combines cold winters (average January temperatures between -10°C and -15°C, with extended periods below -20°C), existing fiber infrastructure in most communities, a provincial grid with significant existing generation capacity, and a population that has historically faced above-average home heating costs relative to Canadian median incomes.

A well-insulated detached house in this climate with an average floor area of 150 m^2 requires between 5 kW and 8 kW of continuous thermal power to maintain 20 °C interior temperature at the design outdoor temperature. For a hundred homes, aggregate thermal demand is therefore approximately 500–800 kW continuous, or 360–580 MWh over a thirty-day January. A leaky older house—common in the New Brunswick housing stock, which skews toward pre-1980 construction—may require 10 kW or more, so 800 kW continuous is a reasonable upper bound for an unimproved neighborhood.

Now consider deploying 40 GPU heater nodes across those hundred homes, roughly two in every five houses: the houses with the highest heating loads, the best thermal coupling between compute hardware and living space, and the most thermally favorable $\eta(x)$ values. Each node runs at 400 W average when dispatched (a conservative figure for a modern mid-range GPU cluster under batch workloads), producing 400 W of heat with perfect thermal efficiency at the node. Forty nodes at 400 W each produce 16 kW of continuous thermal output—approximately 2% to 3% of the neighborhood's total heating demand. That sounds modest, but the economics are not primarily about the fraction of heat provided; they are about the value of the computation performed. If the computation has a market value of \$0.05 per kWh of electricity consumed (a conservative figure for batch cloud compute pricing), then 16 kW of continuous dispatch produces revenue of \$0.80/hour, or approximately \$575/month for the network across those forty households. Distributed equally, each participating household earns roughly \$14/month in compute revenue while receiving free supplemental heat.

To raise the thermal contribution to a significant fraction of neighborhood heating demand, a denser deployment is required. If every home hosts a node capable of 2 kW output when dispatched—achievable with a small rack of purpose-built compute modules rather than a single consumer GPU—and if the scheduler can dispatch 60% of nodes simultaneously during peak heating demand, total thermal output reaches:

$$100 \times 2\,\text{kW} \times 0.6 = 120\,\text{kW},$$

which covers 15% to 24% of the neighborhood's thermal demand without any other heating source. Combined with a heat pump in each home operating at a COP of 2.5 during moderate-cold periods and supplemented by the compute node during extreme cold, the integrated system covers the full heating load while providing 120 kW of thermally scheduled computation. The batch workload required to sustain that dispatch level—120 kW of compute capacity across a hundred nodes, drawing roughly 120 kW of electrical power—represents a nontrivial but not implausible share of the batch compute market for a mid-sized city's residential catchment area.

The case study does not prove that the proposal is cost-optimal under current market conditions. Current GPU hardware is not designed for residential installation; current regulatory frameworks do not recognize the residential compute node as a utility participant; current electricity tariff structures do not reward distributed thermal dispatch. What the case study shows is that the thermodynamic quantities are in the right order of magnitude for the proposal to be economically interesting rather than physically negligible. A 15%–25% contribution to neighborhood heating demand from distributed computation is not a rounding error; it is a number that would register meaningfully in any serious energy audit. Whether it is achievable at cost is a question that depends on hardware development, tariff design, and regulatory adaptation—institutional questions whose answers are not fixed by physics but by decisions that have not yet been made.

7. PHYSICAL CO-COMPUTATION: BEYOND WASTE-HEAT RECOVERY

The proposal described so far treats the GPU heater node as a digital computer whose waste heat is intentionally directed into the building's thermal envelope. This is Stage 3 in the schema introduced earlier: integration rather than mere recovery. But there is a further stage that the current framing does not capture, and that the xylomorphic analogy most strongly implies. Call it Stage 4: physical co-computation, in which the thermal, fluidic, and optical components of the node are not merely cooling or heating systems attached to a digital processor but are themselves computational substrates performing useful operations as intrinsic features of the same physical processes that generate and transport heat.

The distinction matters because it changes the philosophical ground of the proposal entirely. The Stage 3 claim is: computation creates heat; therefore heat should be reused. This is a resource-efficiency argument. The Stage 4 claim is: computation, heat transport, fluid flow, optical propagation, and thermal diffusion are manifestations of the same underlying physical dynamics and should therefore be engineered as a coupled system in which each process contributes to computation, not merely to thermal management.

Consider what the components of a xylomorphic node might compute if they were engineered to do so rather than merely to cool:

A *pressurized fluid loop* whose geometry is tuned to produce pressure distributions that solve a routing optimization problem. The physical solution is read off by pressure sensors at output ports. The computation is performed by fluid dynamics, not by a CPU stepping through an algorithm. The heat generated by fluid friction and pump operation is directed into the building envelope.

An *optical interference network* consisting of waveguides, beam splitters, and phase modulators performs matrix-vector multiplication at the speed of light and at vanishingly small energy cost per operation. The optical components generate negligible heat; the photonic-electronic interface generates modest heat; all heat produced is delivered to the thermal envelope.

A *phase-change material layer* whose state—solid, liquid, or mixed—encodes information. The transition between states stores and releases latent heat, which serves both as thermal buffering for the building and as a read/write operation in a thermal memory. The state of the material is computational state and thermal state simultaneously.

A *thermal diffusion layer* in which a temperature field evolves over time according to the heat equation. The steady-state temperature distribution across the layer is the solution to a Laplace equation over the boundary conditions, which can encode the solution to a potential-field problem. The computation is performed by heat diffusion; the heat is delivered to the room.

A *resonant acoustic chamber* whose normal modes encode a spectral decomposition of an input signal. The chamber is simultaneously a sound insulator (reducing noise transmission through walls), a mechanical resonator, and a signal-processing element. Its thermal losses are directed into the wall assembly.

The functional density formulation. To formalize what distinguishes these systems from conventional infrastructure, define the functional density of a physical object as:

$$\mathcal{F} = \frac{\sum_i U_i}{M},$$

where U_i are the distinct useful services the object provides and M is its mass, volume, or embodied energy. A conventional resistance heater has $\mathcal{F} = U_{\text{thermal}}/M$: one service per unit of material. A conventional server has $\mathcal{F} = U_{\text{compute}}/M$: one service. A Stage 3 GPU heater node has $\mathcal{F} = (U_{\text{compute}} + U_{\text{thermal}})/M$: two services. A Stage 4 physical co-computation node has:

$$\mathcal{F} = \frac{U_{\text{digital}} + U_{\text{fluid}} + U_{\text{optical}} + U_{\text{thermal}} + U_{\text{acoustic}} + U_{\text{storage}}}{M},$$

where each term represents a distinct class of useful computation or useful physical output. The xylomorphic imperative is to maximize $\mathcal{F}$: to extract the maximum number of useful services from the minimum amount of physical substrate. Forest systems achieve high $\mathcal{F}$ not by packing more devices into the same space but by engineering materials in which each physical process contributes to multiple utility functions simultaneously.

Historical precedents. The idea that physical dynamics can perform computation is not new. Differential analyzers used gear-and-shaft mechanisms to integrate differential equations; the computation was performed by the mechanical dynamics of the device. Hydraulic computers, including the MONIAC built by Bill Phillips in 1949, used water flows through tanks and valves to model economic systems; the computation was performed by fluid pressure and flow. Analog optical computers used lenses and interference to perform Fourier transforms. Neuromorphic chips exploit the native dynamics of transistors operating near threshold—rather than driving them into hard digital states—to perform energy-efficient neural computation. What these systems share is the recognition that physical dynamics already perform mathematical operations, and that engineering systems to exploit those dynamics directly can produce computation that is faster, more energy-efficient, or qualitatively different from what is achievable through digital abstraction.

The xylomorphic node synthesizes this tradition with building infrastructure. The pressurized fluid loop is not merely a cooling system; it is a hydraulic computer whose problem domain is chosen to match problems the network needs solved. The phase-change layer is not merely thermal mass; it is a non-volatile memory whose read/write operations are performed by heating and cooling. The optical layer is not merely a communications interface; it is an optical co-processor whose matrix operations are performed at photonic speeds. And the heat produced by all of these processes—the friction of the fluid, the phase transitions of the PCM, the photon-electron interface losses, the acoustic damping of the resonant chamber—is delivered into the building envelope as useful heat, because the node is sited at a location where heat has economic value.

The deeper implication is that the boundary between computation and thermodynamics begins to dissolve at Stage 4. In a conventional digital computer, computation and heat are in fundamental tension: heat is the cost of computation, the enemy of performance and reliability, the quantity that must be removed as quickly as possible. In a physical co-computation substrate, heat is one of the dimensions in which computation occurs. A thermal diffusion field that solves a potential problem is performing computation through heat, not despite it. The heater that thinks is not a poetic formulation; it is a precise description of an architecture in which thinking and heating are the same physical process viewed from two different engineering perspectives.

8. THE THERMAL UTILITY FUNCTION AND CLIMATE AS SCHEDULER

To integrate the spatial and temporal dimensions of the proposal, define a thermal utility function $U(x, t)$ that combines the value of computation and the value of heat at location x and time t:

$$U(x,t) \;=\; C(x,t) \;+\; \lambda(x,t)\, H(x,t),$$

where $C(x,t)$ is the computational value of a unit of processing at (x,t), $H(x,t)$ is the heating value of a unit of thermal output at (x,t), and $\lambda(x,t)$ is a spatiotemporally varying weight that reflects the local price of heat relative to the local price of computation.

The network scheduler maximizes total utility:

$$\max_{\{w(x,t)\}} \int_{\mathcal{X}} U(x,t)\, w(x,t)\, dx \quad \text{subject to} \quad \int_{\mathcal{X}} w(x,t)\, dx \;=\; D(t),$$

where $w(x,t)$ is the computational workload allocated to location x at time t, and $D(t)$ is total demand.

The key insight in this formulation is the role of $\lambda(x,t)$. In a Canadian residential district in January at -20°C, the local price of heat is high: the alternative is gas or electric heating at full cost. The local price of computation is set by the network's going rate for its workload class. If λ is large, the scheduler directs load toward that location even if its raw computational efficiency is not the highest available, because the joint utility of computation plus free heat exceeds the utility of more efficient computation that discards its heat. Climate becomes a scheduling parameter not metaphorically but literally: the optimizer reads outdoor temperature, building heat loss coefficients, current occupancy patterns, and grid electricity prices to determine where computation should run.

This produces a network with qualitatively different behavior from a conventional content delivery network or compute cluster. Latency-sensitive workloads are constrained to low-latency paths regardless of thermal profile. But a large and growing class of workloads—AI training, rendering, simulation, batch analytics, cryptographic operations, genomic processing—is latency-tolerant at the scale of hours or days. These workloads can be scheduled thermally: run in cold places when those places need heat, migrate when they do not. The infrastructure cost of this migration is essentially zero because the workload is software; only the direction of the packets changes.

9. Institutional Frictions and the Xylomorphic Transition

The physical and mathematical arguments developed above are not, in themselves, the difficult part of the proposal. The physics is well understood: resistive heating has a coefficient of performance of one, computation also has a coefficient of performance of one from a thermal perspective, and if the computation has economic value then the joint system is more efficient than either separately. The mathematics of the scheduling problem is a variant of standard network utility maximization with an additional thermal term. Neither of these represents a conceptual barrier.

The difficult part is institutional. The distributed thermal compute network requires the simultaneous reorganization of at least four sectors that currently operate with substantial independence: the compute infrastructure industry, the residential and commercial heating industry, the electrical grid and its balancing mechanisms, and the regulatory frameworks that govern each. A GPU heater node installed in a residential building is simultaneously a piece of consumer electronics, a heating appliance, a grid-connected distributed energy resource, and a node in a commercial compute network. No existing regulatory category comfortably covers all four of these descriptions, and the jurisdictions that govern each are distinct.

This is not an argument against the proposal. It is an identification of the primary constraint on its realization, which is institutional rather than technical. The history of infrastructure transitions suggests that this kind of regulatory fragmentation is a normal feature of the early period of a new architecture: the institutional categories that govern a new system are typically inherited from the systems it is displacing, and they fit imperfectly. Electricity regulation was inherited from gas lighting regulation. Internet regulation was inherited from telecommunications regulation. The regulatory framework for distributed thermal compute nodes will need to be constructed, and its construction will require the same kind of patient institutional work that preceded every previous infrastructure transition—which is to say, it will require time, conflict, and the gradual accumulation of precedents that eventually constitute a new regulatory logic.

The GPU heater node also faces a manufacturing challenge that is worth naming honestly. High-performance compute hardware is currently produced in highly concentrated supply chains optimized for the large, air-cooled server architectures that data centers require. A GPU heater node has different requirements: it must be physically compact, thermally integrated with residential HVAC rather than with industrial cooling systems, robust to the humidity and particulate environment of occupied buildings, repairable by building maintenance staff rather than specialized data center technicians, and manufacturable in sufficient volume to penetrate the residential heating market at price points competitive with conventional heaters. None of these requirements is technically impossible; some of them push against the current trajectory of high-performance chip design, which tends toward larger, hotter, and more power-dense configurations that are increasingly difficult to cool in residential environments. The xylomorphic node requires a different design trajectory: toward modularity, thermal legibility, and field serviceability rather than toward raw throughput maximization.

10. Infrastructure Transformation Versus Substitution

The distributed thermal compute network is not a prediction. It is an existence proof that the space of futures that current institutional categories regard as inaccessible contains architectures that are genuinely novel rather than merely incremental improvements on existing systems. From the perspective of a planner working within current categories, a residential heater that simultaneously participates in a planetary compute network and receives compensation for doing so is a category error: heaters are in one industry, compute is in another, and the regulatory and economic frameworks that govern each have no mechanism for their intersection. This is not because the intersection is impossible but because the categories that would make it legible have not yet been constructed.

Most projections of the energy transition are substitution narratives. The expected future looks like:

$$\text{Gas Furnace} \to \text{Electric Heat Pump}, \qquad \text{Coal Plant} \to \text{Solar Farm}.$$

In each case the function remains the same and the energy source changes. This is a coherent and important project—direct substitution of clean for dirty energy sources is responsible for a substantial share of feasible near-term emissions reduction. But substitution narratives systematically underestimate the degree to which large infrastructure transitions reorganize functions rather than merely replacing their energy substrate. Electrification did not merely substitute electric motors for steam engines; it dissolved the physical proximity constraint between power source and production floor, which reorganized factory layout, then factory location, then the geography of industrial cities. The railroad did not merely substitute steam traction for horse traction; it restructured the time-distance relationships of commerce, which reorganized supply chains, then retail, then urban form.

The xylomorphic proposal is a transformation narrative rather than a substitution narrative:

$$\text{Heater}+\text{Computer}+\text{Network Node} \rightarrow \text{Single Integrated Infrastructure Object.}$$

The successor to the heater is not a cleaner heater. It is a device that performs heating as one output of a multi-output process whose other outputs include computation and network participation, and that receives compensation for those other outputs in a way that partially or fully offsets the cost of the heating it provides. This object does not currently have a regulatory category, a product standard, a utility tariff structure, or a common name. Its absence from the planning horizon of current infrastructure policy is not evidence that it is technically impossible or economically unviable; it is evidence that the transition has not yet generated the institutional forms that would make it visible as a category.

The analogy to forest systems is not decorative. A forest does not have a centralized processor and a separate heating department. It distributes processing, transport, thermal regulation, chemical storage, and structural load across the same substrate at different organizational scales, and the efficiency of the system derives precisely from this integration. A leaf is not merely a photosynthetic device; it is simultaneously a gas exchanger, a water regulator, a structural element, and a thermal surface. Each fungal connection in the mycorrhizal network is simultaneously a nutrient transporter, a chemical signal relay, and a structural anchor. The intelligence of the forest emerges from coordination across semi-autonomous local nodes, not from the centralization of function in a single large organ. That is the topology the xylomorphic network proposes to translate into the built environment: not a replacement for the data center but a dissolution of the category, replaced by a wood-like distribution of semi-autonomous thermodynamic nodes whose collective behavior produces both computation and heat without designating either as waste.

11. CONCLUSION: THE HEATER THAT THINKS

The GPU heater node is a deliberately provocative formulation. It is designed to surface a question that the current organization of infrastructure makes difficult to ask: why do we build systems that produce heat as a byproduct and discard it, while simultaneously building other systems whose primary purpose is to produce heat and whose discarded byproduct is the absence of computation? The inefficiency this question reveals is not small. Global data center energy consumption is on the order of several hundred terawatt-hours per year, a substantial fraction of which is exhausted to the atmosphere as waste heat in climates that simultaneously spend comparable amounts of energy on space heating. The overlap of these two fields is not a rounding error. It is a resource of civilizational scale that is being systematically discarded because the institutional categories that govern each system have no mechanism for treating the other's waste as their own input.

The xylomorphic premise does not require that every home contain a server farm, or that every data center be converted into a district heating plant, or that the entire compute network be reorganized around thermal geography. It requires only that thermal output be treated as a design parameter rather than a design constraint, and that the spatial distribution of computational workload be treated as a scheduling decision that can be made jointly with the spatial distribution of heating demand. These are modest requirements at the level of physical principle. At the level of institutional design, they are substantial—which is where the actual work of the transition lies.

The path from the current architecture to the xylomorphic one is not a single invention but a sequence of institutional and engineering decisions whose cumulative effect is to dissolve the boundary between computation and heat, and to replace the single-output data center with a substrate that performs both functions simultaneously without sacrificing either. The civilization that emerges from this substitution will not have a word for "waste heat" in the way we currently use it, because it will have built systems in which heat is not wasted. What it will call those systems, and what other functions they will perform, are questions that belong to the future that is not yet describable from inside the present transition. The heater that thinks is, for now, the nearest approximation we have to what that future looks like.

BIBLIOGRAPHY

REFERENCES

[1] Smil, Vaclav. *Energy and Civilization: A History*. MIT Press, 2017.

[2] Hughes, Thomas P. *Networks of Power: Electrification in Western Society, 1880–1930*. Johns Hopkins University Press, 1983.

[3] Geels, Frank W. *Technological Transitions and System Innovations*. Edward Elgar, 2005.

[4] Kelly, Nathan, and Jonathan Cockroft. "Analysis of Retrofit Air Source Heat Pump Performance: Results from Detailed Simulations and Comparison to Field Trial Data." *Building Services Engineering Research and Technology* 30, no. 1 (2009): 43–62.

[5] Masanet, Eric, et al. "Recalibrating Global Data Center Energy-Use Estimates." *Science* 367, no. 6481 (2020): 984–986.

[6] Patterson, David, et al. "Carbon Emissions and Large Neural Network Training." *arXiv preprint* arXiv:2104.10350 (2022).

[7] Winner, Langdon. *The Whale and the Reactor: A Search for Limits in an Age of High Technology*. University of Chicago Press, 1986.