# Thermal contradictions of electronic capital

## The second Jevons paradox and the prehistories of the Data Centre Cooling Crisis

Grégoire Chamayou

### Introduction

On 27 November 2025, temperatures inside a windowless building in Aurora, Illinois, began climbing toward 120°F (48.8°C). The facility, a data centre for CME Group, the world's largest derivatives exchange operator, was experiencing a major incident. Within hours, futures trading was halted across the board: West Texas Intermediate crude, ten-year US Treasuries, the S&P 500, the Nasdaq 100, the Nikkei, gold. A senior markets executive told Reuters he hadn't seen such a widespread outage in twenty years. The culprit was neither a cyberattack nor a power cut, but a failure of the air conditioning.

Electronic devices, in order to function effectively – or to function at all – must be kept below certain temperature thresholds. Beyond these, performance degrades, thermal throttling kicks in, and failure rates climb. Should cooling fail entirely, temperatures spike rapidly and everything must be shut down to prevent irreversible damage. The quiet art of 'thermal management' proves crucial to the life and survival of electronic capital. 'You can have the fastest processors', one expert sums up, 'but, if you cannot effectively remove the byproduct of all that work – heat – your entire operation will grind to a halt'.[1]

Data centres are not merely cooled: they are, as we will see, constitutively air-conditioned machines. Yet as AI-driven computing scales up, the chips stacked on their racks grow denser, more powerful and more energy-hungry with each generation, and the thermal loads they must dissipate mount accordingly, to levels that conventional forced air-cooling systems can no longer handle efficiently. The notion of the 'thermal wall' designates this breaking point.

A 'cooling crisis'[2] is raging in data centres. From this point of departure, I would like to sketch a brief genealogy of this kind of problem – not a linear account of cooling techniques,[3] but a placing of this contemporary crisis in series with kindred cases drawn from a longer past, in the hope that these sometimes unexpected connections shed new light on what is at stake.

I argue that such crises recur whenever the drive to expand technical capacity generates, at some threshold, the thermal conditions of its own arrest. Following this thread will lead, among other unexpected sites, to old mines and a forgotten corner of economic thought, where a neglected paradox lies buried.

What this genealogy adds to our understanding of industrial warming, I suggest, is a shift in scale: in this case, excessive heat does not only appear as a negative externality for the planet's climate, but also as a recurring internal contradiction within industrial micro-climates of production. This internal saturation, however, eventually joins an external one, as the final pages will show.

To this end, the argument proceeds not chronologically but morphologically: each section unfolds, through historical materials, a stratum of thermal contradiction, moving from the most immediate to the most expansive, progressively complicating the schema.

## Overheating machines and warming loops

Let us establish the basics: this is a problem of machines that heat up and, in doing so, risk overheating. Overheating represents an internal threat, a danger the machine generates from within and that turns back on itself. Conceptually, it is a case of self-contradiction that can lead to self-destruction. Devices that begin to smell of burning may catch fire entirely.

### *Simondon's cooling fins*

Hence the question: how to prevent overheating? Gilbert Simondon saw in the imperative of non-self-destructiveness a governing principle of the evolution of technical objects. He took the example of the engine. A first strategy, consisting in adding an autonomous cooling circuit with pump and pipes, had a weak point: should the water run out, the circuit would stop cooling while the engine continued to heat up. For Simondon, this vulnerability betrayed a design flaw, that of a still 'abstract' technical object, conceived as an assembly of monofunctional elements rather than a structural unity.

To this he opposed a movement of 'concretisation'. Around 1900, fins were added to a motorcycle engine. By extending the surface of the machine, they improved heat dissipation. Initially they behaved as simple 'defense structures'[4] with cooling as their sole function; but this changed: in a 1956 model, they had become ribbing that simultaneously reinforced the cylinder's structure and cooled it – a polyfunctionality characteristic of the 'concrete' technical object. The engine could now cool itself simply by running, air circulating naturally along its fins, in a configuration where 'the functioning itself implies cooling'[5]. This formula is crucial: it is precisely what defines a thermally concrete technical object.

Overheating problems are not, however, the exclusive preserve of heat engines. They have affected electrical and then electronic devices as well. Whatever their differences in operation, all such devices mobilise energy that is partly dissipated as heat. In the 1940s, 'vacuum tubes' – sealed glass cylinders in which a heated cathode emitted electrons – generated intense heat, to the point that 'heat engineering problems' took precedence over 'those of pure electronic character'.[6]

As one engineer explained, 'adequate cooling of a high-power vacuum tube is as critical an operating requirement as the cooling of the cylinders of an automobile engine'.[7] And the same solution was applied: tubes too were fitted with cooling fins.

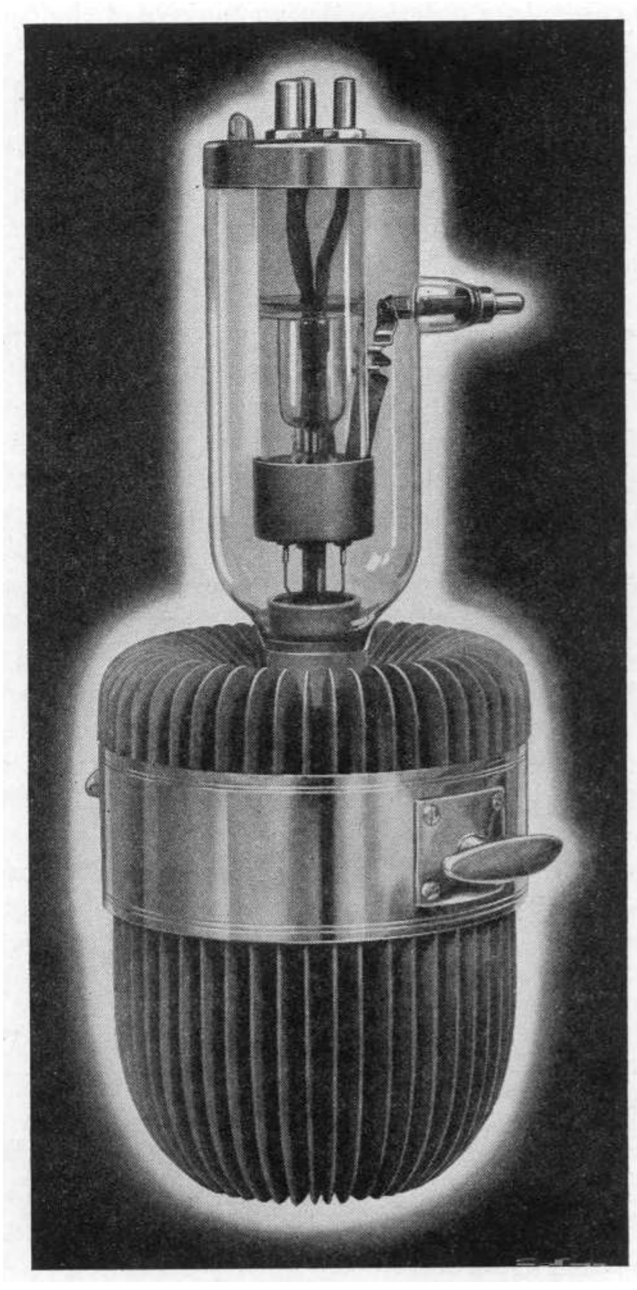

**Fig. 1** – Vacuum tube soldered in a copper air cooler with vertical fins.[8]

The transistors that subsequently replaced tubes faced the same thermal challenge: fins could be machined directly into their casing,[9] and many were mounted on a finned metal chassis called 'heat sink', a feature still present in computers today.

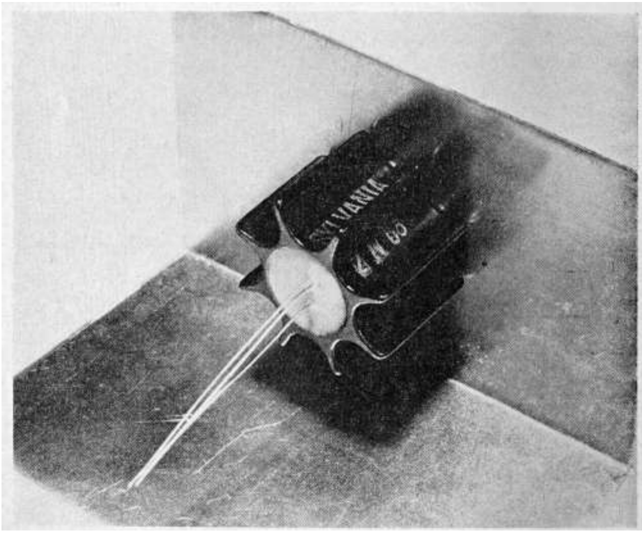

**Fig. 2** – High–power transistor with cooling fins.[10]

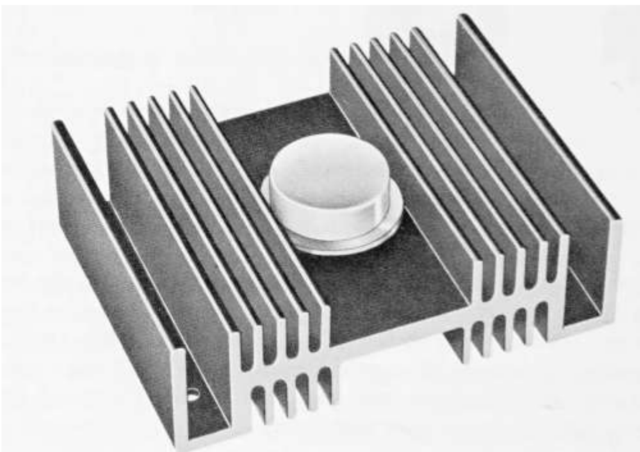

**Fig. 3** – Transistor heat sink.[11]

***'Thermal problem in electronic equipment'***

The 'thermal problem in electronic equipment' can be schematised as follows:

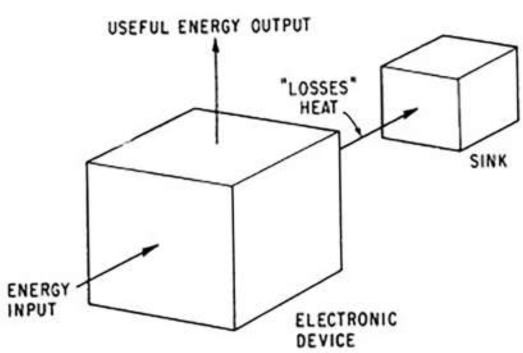


**Fig. 4** – Functional representation of thermal problem in electronic equipment.[12]

An electronic device receives an energy input that produces, beyond useful energy output, heat losses, heat waste that must be evacuated to prevent components from exceeding their peak internal temperature. Thermal management of electronic systems presents itself as an 'art of heat transfer' aimed at 'removing and transporting heat from individual components to the system thermal sink'.[13]

And yet 'the term heat sink is a misnomer':[14] contrary to what the metaphor suggests, this plate does not absorb the heat passing through it but, far from providing an ultimate receptacle, immediately diffuses it into the surrounding environment.

Heat sinks cool the apparatus, but they heat the room, warming the ambient air, the very air on which cooling depends. A feedback loop thus takes hold: a vicious circle in which electronic self-heating raises air temperature, which loses its cooling capacity, which in turn heats the components, and so on. This is another pattern of thermal contradiction: no longer simple overheating through self-heating, but an enlarged overheating via micro-climatic warming.

This problem was particularly acute for early computers that were vast cabinets housing thousands of heat-generating tubes.

***'Death Valley' in Philadelphia***

Long before Silicon Valley, there was 'Death Valley': the nickname employees had given to the Eckert Mauchly Corporation's Philadelphia site, a red brick building wedged at the foot of a hill in an alley with no airflow, where in 1949-50 the UNIVAC I, considered one of the world's first commercial digital computers, underwent testing. The building was not air-conditioned – few were at the time – and intense heat from the machine's myriad tubes compounded already elevated outdoor temperatures. Together they pushed temperatures to between 100°F and 120°F (38 and 48°C) in the 'test area'. Faced with this unbearable heat, some poured water over their heads; most stripped down until 'shorts and undershirts became the uniform of the day',[15] to the point that management circulated a memo asking engineers so attired to refrain from entering the Purchasing Department during business hours, so as not to shock visitors.

The tubes radiated heat so intense that they had 'made necessary a ventilation system to keep the components from disintegrating':[16] a 'hurricane of air' was blown through the computer to dissipate the 120 kilowatts it generated. But the air emerged warmer and was returned to the very room from which it had been drawn. Temperatures kept rising. Moving the air was not enough: it had to be truly chilled; hence the development of air conditioning systems specially designed for computer rooms.

### *Computers as air-conditioned machines*

In the late 1950s, a Univac computer was to be installed in the offices of General Electric. The building was already air-conditioned, but to no avail, since 'the heat liberated by the electronic units required cooling and air handling beyond the capacity of the building system'.[17] Hence the installation of a bespoke system, with ceiling ducts and extractor hoods, rather like those above a kitchen stove.

**Fig. 5** – Cooling unit for a computer at General Electric, Rome, Georgia, ca. 1957.[18]

At that time, a computer purchase came with an obligation: acquiring a 'specially engineered air conditioning system' capable of providing 'the climatic conditions required by the electronic brain'.[19] Beyond temperature, humidity and air purity also had to be controlled, since cardboard punched cards had to be kept dry and magnetic tapes free of dust.[20]

The historical thesis I wish to advance here is that the computer originally appeared as an air-conditioned machine. This was a *sine qua non*, such that, as was held at the time, 'computer operation without cooling is impossible'.[21]

But that vacuum tubes or transistors were fitted with cooling fins does not thereby make them concrete technical objects in Simondon's sense. At least from a thermal standpoint, one could argue that a motorcycle engine was more advanced: incapable of adapting to an associated milieu, and lacking internal autonomous regulation, such a computer was entirely dependent on a conditioned environment whose temperature and humidity had to be maintained within narrow bounds by an external regulator.

### *Cooling the brain*

If computers were born as air-conditioned machines, data centres were likewise born as air-conditioned infrastructures.

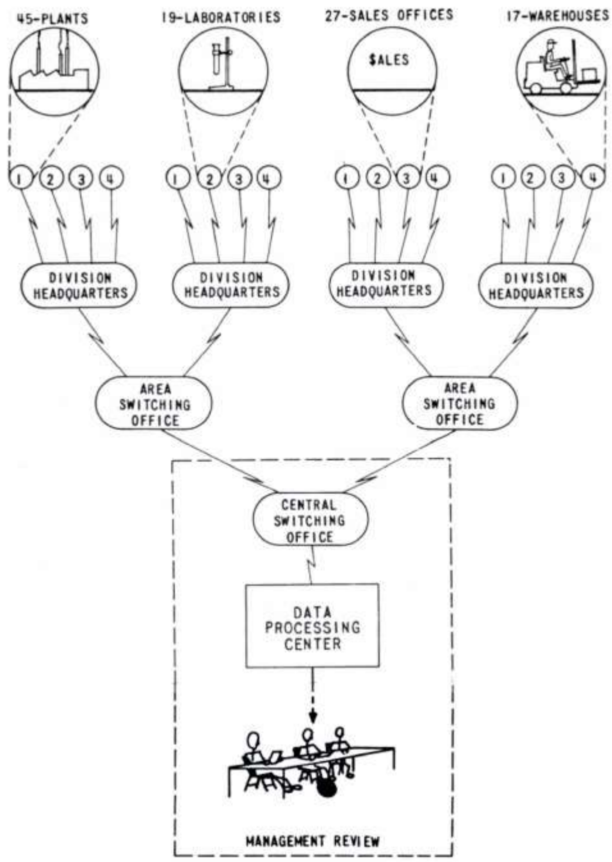


**Fig. 6** – 'Sylvania's integrated data processing system'[22]

One of the first 'electronic data processing centers' was inaugurated in March 1956 in Camillus, New York, by Sylvania Electric Products. This new kind of infrastructure, at once a communications hub and a firm-wide data processing centre, was designed to 'centralize the data processing functions of gathering, recording, computing and classifying' information from dozens of sites across several states.[23] Tasks previously performed laboriously by hand were delegated to the machine – a first step toward an avowed dream of 'administrative automation'.[24] It would henceforth suffice for a manager to press a button to get, say, an up-to-date payroll statement for the entire enterprise.

The 'brain' of the organisation was more prosaically a Sperry Rand Univac housing more than 5,000 vacuum tubes giving off immense heat. The task of cooling it was entrusted to the Carrier Corporation, which installed a 200-ton conditioner combining chilled-water circulation and forced ventilation.[25]

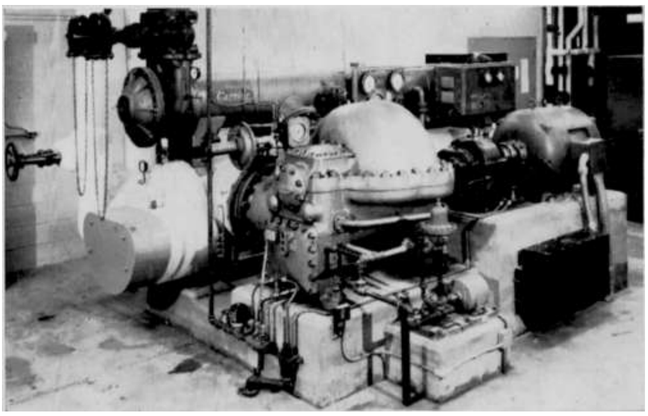

**Fig. 7** – Carrier conditioner at Sylvania's Data Processing Center, Camillus, New York, 1957.[26]

Carrier's involvement was no coincidence. Willis Carrier is credited with designing the first air conditioner in 1902 for a Brooklyn printing house[27] – industrial processes, not homes, were the first objects of this technology. Less well known is that the firm he founded was also, in the 1950s, at the forefront of computer cooling,[28] installing 'climate control systems for some of the largest electronic units in the country'.[29] These 'man-made weather' machines 'ended the tyranny of the macroclimate on industry',[30] yet at the cost of another dependence: the inaugural reliance of data centres on air-conditioning systems themselves.

## The smaller, the hotter

The warming loop just described, in which the device heats the environment on which it depends for cooling, constitutes one type of thermal contradiction. Another one, of a different form, arises from the very dynamic of miniaturisation: the driving force of computing growth comes, through its own momentum, to generate the very effects that obstruct it.

### *Miniaturisation drive*

In the aftermath of the Second World War, a 'search for miniaturized electronics'[31] took hold: reducing component size made it possible to pack more of them into the same area, and very soon, in the domain of semiconductor electronics, 'achievement was calculated in packing densities'.[32]

The associated advantages were of particular interest to the military, the first clients of this nascent industry: the long-range missile guidance systems sought in the early 1950s had to be lightweight, compact and reliable – everything that vacuum tubes were not.[33]

It was also clear that radical miniaturisation was the precondition for the emergence of the powerful computers already being dreamed of. Given the components then available, to emulate certain capacities of the human brain, physicist Richard Feynman remarked, 'we would have to make it, perhaps, the size of the Pentagon'.[34] This would take up too much space, but above all, such a computer 'would be limited to a certain speed'.[35] Since execution time depends ultimately on the distance electrons must travel, the key was to shorten these micro-distances. Making computers 'faster and faster' meant making them 'smaller and smaller'.

### *Thermo-informational antinomy*

The 'tendency toward size reduction', Simondon added, is explained by an imperative peculiar to 'information technics'.[36] Whereas the 'technician of thermodynamics' tends toward gigantism, because the *energy output* he is after increases with the dimensions of the machine (a turbine, typically), the information technician, conversely, miniaturises in order to improve *information output* (146).

Here a fundamental tension appears. In electronics, electricity is mobilised in a specific way: through 'low current signals', that is, 'electric currents considered not as energy carriers, but as vehicles for information' (144). Yet this informational medium remains energetic: any electronic component through which current flows undergoes power dissipation in the form of heat. Electricity-as-information carries thermodynamic properties with it, inescapably.

From this duality arises an antinomy: even as 'the information technician is brought in to find the smallest possible dimensions', those dimensions must remain 'compatible with the residual thermodynamic requirements of the devices' (146). Yet as component dimensions shrink, their capacity to dissipate heat diminishes, until they risk 'reaching a temperature that would compromise their functioning' (147). The informational imperative, which demands ever-smaller dimensions, is thus held in check by a thermal constraint. The problem, in short, is how to shrink without overheating.

#### ***From tubes to transistors***

A distinction must be drawn here between two modes of size reduction with different thermal implications: *relative* reduction, in which the same component is produced in a smaller version, and *absolute* reduction, in which one family of components is replaced by another of an entirely different order of magnitude.

The transition from tubes to transistors illustrates this pattern exactly: from the moment of their invention, transistors were presented as the solution to the thermal burden of their predecessors: 'no heat is developed as in a vacuum tube', the *The New York Times* reported;[37] 'since there is no cathode to shoot out electrons when hot, there is nothing to heat up', *Popular Science* confirmed.[38] For a time, the problem seemed solved.

It was not to last. While the transistor allowed a 'dramatic reduction in power dissipation',[39] it did not truly resolve the problem. On the one hand, these new components proved far more sensitive to heat than tubes.[40] On the other, their density increased rapidly: as the space between them shrank, heat accumulated. This 'rapidly growing transistor packaging density' put thermal management back on the agenda.[41] The transistor had offered a thermal respite; densification closed it off. By 1965, the verdict was inescapable: 'The origin of the thermal problem is the continuous effort to use components that are lighter and smaller'.[42] The very force driving this industry was turning against it.

Two patterns of thermal contradiction thus emerge. The first, the 'warming loop', is static: devices heat their own microclimate beyond their tolerance threshold. The second, just examined, is dynamic: the very drive toward miniaturisation engenders, as a corollary of its own momentum, a growing thermal obstacle. What had hitherto appeared only as limits to functioning become limits to growth, thermal limits to growth. This second pattern is not unique to the history of computing. A remarkably early formulation of it can be found in an entirely different context.

## The Second Jevons Paradox: Thermal limits to growth

In Great Britain in the second half of the nineteenth century the spread of 'steam power' was driving a steep rise in coal consumption. Concerns were mounting over the exhaustion of coal mines. The question was not so much how many tons of coal remained buried underground in total, as how much remained exploitable, that is, at accessible depths, in seams above 'the limit of possible depth of working'.[43]

#### ***Thermo-physiological limitations***

This limit did not stem primarily from technical incapacity. Progress in the steam engine, itself coal-powered, had furnished increasingly powerful extraction machinery, making it possible to dig ever deeper.

The limiting factor lay elsewhere: in the climate of production, in the microclimate of the mine. It can be very hot underground, and the deeper you go, the hotter it gets. At 3,000 feet (915 metres), it was estimated, air temperature would reach 98°F (37.2°C), body temperature. And therein lay the problem: in the bodies of the miners. For there are limits to what the organism can endure: beyond certain thresholds, heat becomes physically 'incompatible with human labour'. This set the limit on extraction depth: a 'limit of practicable labour' defined by the 'limit of human endurance of high temperature'.[44]

There was, therefore, a paradox. On the one hand, increasing mechanical power enabled ever deeper penetration into the earth, but on the other, this very deepening, as descent brought rising heat, could degrade the work capacity to the point of making it impossible altogether. There was, in other words, a contradiction between the climate of production and the labour force understood as a living reality, as a thermosensitive entity: the deeper you dig, the less you will be able to dig.

This physiological threshold, set by some experts at 98°F, was poorly defined and widely overestimated, as physiologist John Scott Haldane's subsequent work would show.[45] But the essential point remains: there are thermophysiological limits to productive warming – thresholds of thermo-tolerance against which the extractivist dynamic of deepening came to bear, as it were, like a 'glass floor'.

This explains why the mining sector was a pioneering site for the scientific study of thermal stress. Ironically, one of the key indices currently used to assess the direct physiological effects of global warming, known as wet-bulb temperature, has its roots in the study of mine environments.[46]

### *Commercial impossibilities*

Geologists and physiologists were not the only interested parties. Economists too weighed in. In 1865 William Stanley Jevons published 'The Coal Question', in which he formulated the famous paradox that would come to bear his name.[47] The idea – now regarded as one of the 'pioneering insights into ecological economics'[48] – holds that making machines more energy-efficient does not necessarily reduce overall fuel consumption, since less wasteful machines are sold and used more widely, so that, as Jevons writes, 'new modes of economy will lead to an increase of consumption'.[49] If this first paradox is well known, the same reflections contain a second one, long forgotten, which I would like to reactivate here.

Jevons first established a distinction between two registers of impossibility: physical and commercial. A project may be technically feasible – it could be done – without being economically viable – it would not pay. In the case of mines, it becomes decisive: depths will certainly be reached at which the heat makes work impossible, but economic difficulties will be encountered well before these physical limits. Beyond a certain depth, operating costs – what Jevons calls the 'cost of sinking and working deep' – will rise unsustainably.

The greater the heat, the more it will 'enervate the labourers', eroding their work capacity to the point of extinguishing it. Jevons cites the example of a 'Cornish mine' where men do manage to work in an atmosphere varying from 110° to 120°F (43 to 49°C), but nearly naked, doused with cold water, and for no more than twenty minutes at a stretch (43).

Countermeasures could be taken, certainly:[50] but 'much increased ventilation will be a matter of expense and difficulty' (43). The question is not how deep it is physically possible to send miners, but 'how far can we follow them with profit, considering the growing costs and difficulty of deep mining' (30). This is less a question of technical feasibility than of profitable exploitation.

These costs, moreover, tend to spiral, since 'the whole capital and current expenditure of the mine grows in a higher proportion than the depth' (44). The coal problem is therefore not primarily one of scarcity but of mounting extraction costs, 'gradually augmenting from present rates to the limit of commercial impossibility'.[51]

### *The thermal tendency of the rate of profit to fall*

This contradiction did not escape Marx. In his manuscript notes on Jevons, lately published in the MEGA edition, he wrote, in his characteristic Anglo-German mishmash:

> Going deeper enhances the price 'und mit dieser Preiserhöhung gehn unser rate of progress' 'und special advantages' 'zum Teufel' 'Mit tiefrem Arbeiten': the heat grows more insupportable, … the ventilation more costly'.[52]

'*Zum Teufel*' (to the devil) evokes both definitive loss and images of an infernal underworld: in its descent into hell, profitability evaporates in the heat. Here lies the 'second Jevons paradox': the deeper one digs, the less profitable it becomes. In this specific case, the cause is rising warming costs, but the paradox applies more generally to configurations in which an increase in technical capacity produces an environmental counter-effect, a *negative internality*, whose mounting operating costs end up compromising the economic viability of the enterprise. Here, these internalities are thermal, and the counter-effect takes the form of unsustainable warming.

This nineteenth-century paradox finds unexpected echoes in the present. With the data centre 'cooling crisis', a similar pattern recurs: an ascending curve that collides, under its own momentum, with a thermal limit. The economic dynamics at play seem, however, to be reversed. Where Jevons shows coal mining to be 'a historically increasing cost industry',[53] the microprocessor industry long embodied the opposite: a spectacular instance of a 'historically decreasing cost industry', as was trumpeted for decades. For a long time, that held. Until the moment it seized up, running, precisely, into a thermal obstacle.

## The meltdown of Moore's law

'Moore's law' – in reality a forecast, famously put forward by engineer Gordon Moore in the mid-1960s – stipulated that the number of transistors on a chip would increase 'at a rate of roughly a factor of two per year' without a corresponding rise in cost.[54]

### *A forgotten heat problem*

This is well known. Less known is that Moore himself, in his 1965 article, had anticipated a potential obstacle to this progression:

> 'Heat problem / Will it be possible to remove the heat generated by tens of thousands of components in a single silicon chip? If we could shrink the volume of a standard high-speed digital computer to that required for the components themselves, we would expect it to glow brightly with present power dissipation'[55].

If one reasoned on the basis of older components, such densification seemed thermally untenable: chips, it appeared, would run white-hot. 'But it won't happen with integrated circuits' (84), he immediately countered. In raising the objection only to dismiss it, Moore had put his finger on the prediction's most serious vulnerability. But he set it aside, and the reason, according to him, was that 'shrinking dimensions on an integrated structure makes it possible to operate the structure at higher speed for the same power per unit area' (84).

### *Dennard scaling: Toward the end of free lunch*

Moore's intuition proved correct, for a time. What would later be called 'Dennard scaling' provided the mechanism that made it so: the principle that power density stays constant with transistor scaling. In a 1974 article, engineer Robert Dennard and his team showed how to place more transistors on a chip without increasing power consumption.[56] The key was to reduce supply voltage in proportion to component size, thereby keeping power density and heat emission constant.

Together, Moore's law and the Dennard principle enabled decades of exponential performance growth 'while staying within a similar power budget'.[57] This could be summarised in a formula: there was, for a time, something like a genuine period of 'free lunch' in computer engineering – performance gains came at no thermal cost.[58] Except that Dennard scaling broke down in the mid-2000s, disrupting the equilibrium on which the sector's growth rested. Before this became fully apparent, warning signs had been issued.

### *Hot As sun: A 2001 warning*

In February 2001, at the International Solid-State Circuits Conference, Patrick Gelsinger, Intel's senior technology executive, announced that an era was drawing to a close. The prodigious increases in transistors per chip had come at a price, namely 'dramatic increases in power consumption and power density'.[59] Despite efforts to reduce supply voltages, clock frequency, which governs the pace of circuit operations, was being driven relentlessly upward, demanding ever more power and shedding ever more heat.

Gelsinger offered his own prediction, a pessimistic counterpart to Moore's. If this continued, heat would climb exponentially: within a decade, a chip's surface would emit as much heat as a rocket nozzle, and eventually as much as the sun.

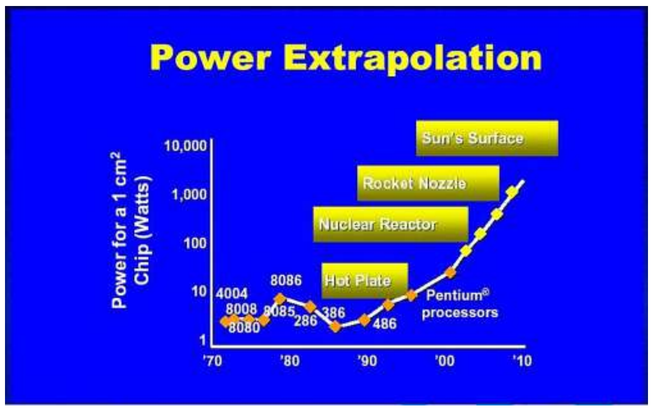


**Fig. 8** – Gelsinger's slide from ISSCC 2001.

In the past, new heat sinks and fans had contained thermal surges, but those fixes, Gelsinger warned, 'won't work in the future': nobody knows how to cool 'a 5,000-watt processor'.[60] To persist in this headlong flight, hoping that cooling would keep pace, was at the very least hazardous.

### *2004: The clocks hit the wall*

Three years later, Intel acknowledged having hit a 'thermal wall' on its microprocessor line. The term itself originates here: a limit against which an industrial trajectory comes to crash. 'Our strategy', Craig R. Barrett, Intel's CEO, had confided, 'is to go 200 miles per hour until we hit a brick wall'.[61] They had.

The approach had consisted in boosting performance by shrinking chips while 'relentlessly raising their clock speed'.[62] For a long time this worked: smaller sizes concentrated computing power while shorter signal paths enabled faster clock rates.[63] In the mid-2000s, however, both levers started to show signs of jamming.

Miniaturisation combined with lower supply voltages had reduced *dynamic power* – the consumption tied to active processor operations – but another source of dissipation remained: *static power*, the current that leaks through transistors even when switched off. Long negligible, this off-state leakage became problematic as transistors shrank: beyond a certain threshold, they became 'increasingly porous, permitting electric current – and heat – to cascade out even when switched off'.[64] At the nanometric scale, this leakage became the 'major contributor to power consumption':[65] the more one miniaturised, the more transistors leaked, and the more heat accumulated.

This uncontrolled thermal dissipation proved incompatible with continued clock acceleration, which was itself a source of warming. Previously, Dennard scaling had partially offset this, but leakage now added its effects without counterpart, making the combination unworkable. By the late 2000s, as thermal packaging specialist Wataru Nakayama would later note, further increases in clock frequency had become 'unsustainable, due partly to high cost of cooling'.[66]

Moore's law was not dead: miniaturisation and densification remained possible, but performance gains now came with a rising energy bill and increasing heat output. As one trade magazine put it in 2005: 'the more you scale, the hotter you get'.[67]

### *Deconcretisation of a technical object*

The effects manifested at several levels, starting with the chips themselves. Tracing the evolution of graphics processing units (GPUs) over twenty-five years, engineer Steven Woo shows that the 'breakdown of Dennard scaling around 2005 prompted designers to place a major emphasis on thermal management', leaving them caught in a 'tug of war between increasing computing performance and dealing with the resulting heat'. A typical mid-1990s graphic card had no noticeable thermal equipment. Then, progressively, cooling structures appeared, grew, multiplied – heatsinks, fans, heat pipes, copper plates, radiator structures – until they occupied most of an expanding volume. 'The amount of work and space required to deal with the heat is really growing over time, and is something that's really unsustainable', he concluded.[68]

Like Simondon before him, Woo was tracing the lineage of a technical object. But this was the world upside down. Where Simondon had identified a progressive movement of concretisation, here increasingly unwieldy defensive structures accumulated on the technical object. It was rather the spectacle of a trajectory of 'deconcretisation' against a backdrop of unresolved thermal contradiction.[69]

### *Unsustainable costs*

This warming tendency appeared again, magnified, at the scale of data centres. Chips, denser and more numerous with each generation, grouped on racks and massed in warehouses, drove thermal loads to extreme levels. Conventional cooling devices were approaching their limits. But as in Jevons's analysis, the problem was above all one of economic viability.

In the second half of the 2000s, experts, foremost among them data centre executive Ken Brill, sounded the alarm. Everyone in the sector could see the symptoms: rising electricity consumption, temperatures 'rising dramatically', cooling problems across installations. But, Brill warned, 'very few understand that the meltdown is fundamentally a high-stakes economic problem, rather than an engineering problem'.[70]

Power and cooling costs, traditionally negligible, had begun to weigh heavily on operating expenditure, to the point of potentially exceeding the cost of the servers themselves.[71] The reason lay in the end of Dennard scaling: previously, miniaturising transistors had simultaneously increased computing capacity and reduced power per operation, so that performance and energy efficiency advanced in concert. That alignment was now broken: transistor counts continued their exponential climb while energy efficiency stagnated. The gap had to be compensated. Hence increased electricity consumption, driving up cooling costs accordingly.[72]

Projecting the curves was enough to reveal the bottleneck: this was what Brill called the 'meltdown of Moore's law', not an overheating of components but, metaphorically, the melting away of profitability itself. The sector's legendary profitability was itself going '*zum Teufel*'. And here they were, from coal to chips, from mining to data mining, the returning limits, thermal limits to growth.

In this analogy – partial, heuristic –, the dynamic of miniaturisation, densification and acceleration stands in relation to data centres as the dynamic of extractive deepening stood in relation to mining: in both cases, the microclimate of production warms until it can no longer be efficiently contained by existing mitigation devices, and becomes economically unsustainable, before becoming physically so by crossing thermal tolerance thresholds.

## The cyclicality of thermal walls

These alerts dated from the mid-2000s. They were taken seriously, prompting sustained efforts to improve energy efficiency,[73] but this amounted to damage control more than any genuine resolution. For the contradiction endures, and today takes a still more spectacular turn. 'Note that projected power dissipation in future is expected to surpass air cooling capability', a 2014 *Encyclopedia of Thermal Packaging* already warned.[74] A decade later, we are there.

The accelerated deployment of high-performance computing demanded by the AI boom involves specialised chips, graphics cards on steroids that send power densities skyrocketing. Where a standard rack dissipated between three and five kW, AI-dedicated installations routinely reach 20 to 50 kW, with some clusters already exceeding 100 kW.[75]

Faced with this, fans can no longer keep up, and once again, the saturation point is primarily economic. Since 'fan power consumption increases roughly with the cube of the fan speed',[76] doubling blade speed does not double the bill, it multiplies it by eight. Under this non-linearity, cooling costs mount steeply, to the point that 'the parasitic power consumption of the cooling infrastructure' ends up 'aggressively eroding profit margins'.[77]

### *Immersion cooling: Nothing new under the sun*

However, reaching the limits of air cooling does not close the story. When one family of solutions is exhausted, others may take over. Among the candidates to 'solve the thermal challenges that traditional air cooling cannot address', the industry today favours immersion: directly submerging components in a dielectric liquid whose boiling, through phase change, removes heat far more effectively.[78] Immersion is often presented as an innovative solution, a breakthrough technology. It has, however, a long history in electronics, with traces going back to the 1950s.

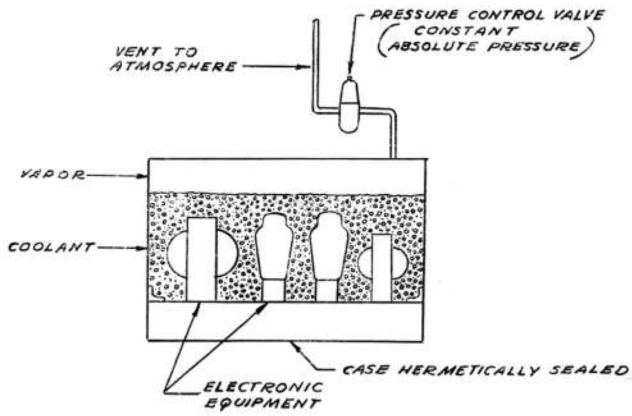


**Fig. 9** – Direct vaporisation cooling system for electronic equipment, 1955.[79]

More important than its age or its non-novelty is its historically recurrent character. As 'past experience indicates a cyclic demand for high capacity cooling',[80] the present turn to immersion is best understood as a revival, in similar circumstances of saturation, of a technical fix already mobilised during other thermal walls. For there have indeed been several, not just one.

### *Back to the future: The 1980s cooling crisis*

To show what is at stake in such moments, I would like to dwell on one of them: the thermal crisis of the mid-1980s.

Contrary to Moore's optimistic 1965 postulate, integrated circuits had not by themselves resolved the heat problem. His forecast was being fulfilled, but only through intense and largely invisible thermal management work. From the early 1970s onward, experts noted that the thermal problem 'has been aggravated in recent years because of developments in circuit and component technologies, which have caused large increases in volumetric heat generation, while concurrently reducing the available surface area for heat dissipation'.[81]

The trend continued into the 1980s, driven by the 'rising trend of heat flux on very large scale integration

(VLSI) chips',[82] until it became manifestly untenable: at the prevailing rate, chip thermal flux could within a decade reach the intensity of the sun's surface.

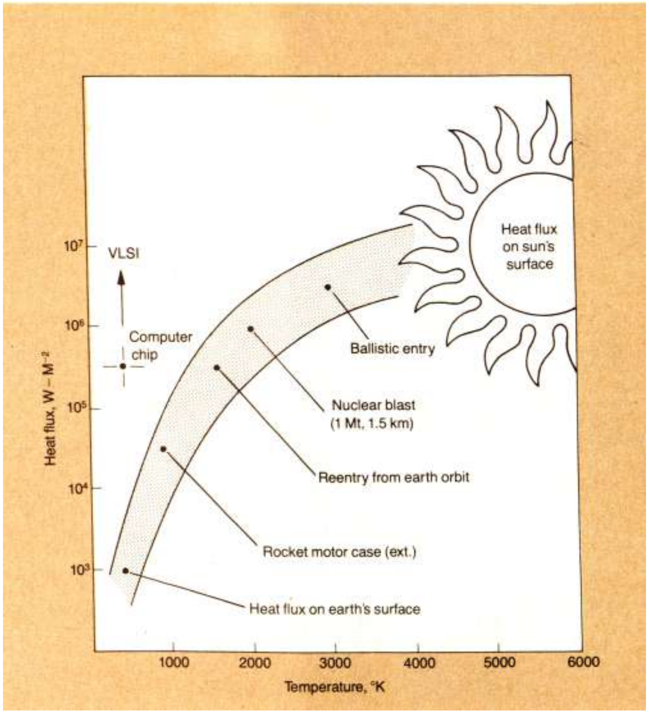


**Fig. 10** – 'The heat flux of today's microelectronic chip is only 20 times less than that on the surface of the sun'.[83]

This was a thermal wall 'avant la lettre', and it is striking to encounter, decades apart, the same discourses and the same visions. When it 'became obvious that air-cooling no longer could cope with the rising heat dissipation',[84] research turned to 'high-performance cooling schemes', including immersion: the photograph reproduced below shows a 'boiling scene from an experiment conducted at Hitachi Mechanical Engineering Laboratory in the mid-1980s'.[85] This technology was nonetheless implemented only on a handful of machines, foremost among them the CRAY-2, the first computer to have 'the central processing unit immersed in liquid'.[86]

### *From Bipolar to CMOS: A pattern*

This was a costly move, reserved for an elite of supercomputers. But another was taking shape: a pivot to a different component family.

In the 1980s, large computers ran on bipolar transistors, powerful but energy-hungry. Another type of component was available, however: the Complementary Metal Oxide Semiconductor (CMOS), used in low-power consumer devices such as battery-powered calculators. Low-heat but slow, CMOS microprocessors were soon run at higher frequencies as miniaturisation progressed, without a corresponding surge in power consumption. In the early 1990s, faced with rising cooling costs, computer manufacturers switched from bipolar to CMOS. This pivot 'set a brake on the rising heat dissipation from computers'.[87] It was the golden age of Dennard scaling.

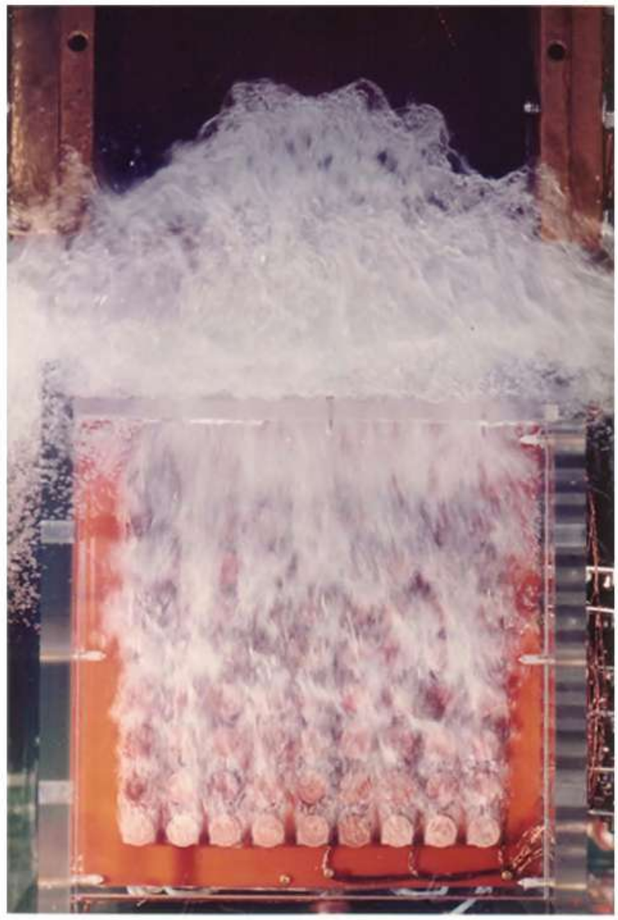

**Fig. 11** – Boiling cooling of simulated chips in the 1980s.[88]

Once again, the respite did not last. Through the 1990s, CMOS chips grew faster and denser, and thermal dissipation resumed its climb: although Dennard scaling allowed continued reductions in energy per operation, 'the thirst of the marketplace for greater IC speed and functionality fueled an acceleration in chip power dissipation'.[89] In the early 2000s, the same prediction resurfaced: chip heat was approaching the intensity of the sun's surface. The same graph reappeared too – the one Gelsinger displayed in 2001 was, more or less, that of 1986. The cycle had come full circle.

From these episodes a typical sequence emerges. A given component type is miniaturised and densified, but its thermal load increases until it exceeds the capacity of the prevailing cooling method. When the repertoire of fixes is exhausted, the only recourse is a new class of components with lower heat output. If found, the problem recedes for a time before the same dynamic re-emerges. The transistor seemed to solve the problem posed by tubes; it recurred. The same with printed circuits; the same with CMOS.

Heat output varied across these sequences, rising as each technology approached its limits and falling as the next was adopted; yet the long-term trend was unmistakably upward, toward 'higher chip-power dissipations'.[90] Beneath the continuous curve lie broken waves: rising, hitting a ceiling, being superseded, replaced by a cooler technology that in turn heats up – but each crest higher than the last.

### *Lessons of a circular history*

What the smooth big-picture curve of Moore's law conceals, as economists Burg and Ausubel showed through a close analysis of Intel chip data, is a series of 'accelerations and decelerations':[91] typically, six years of growth, three of near-stagnation, repeated. These plateaus correspond to points of saturation before a technological switch: one component family reaches its limits, another is adopted, and densification resumes. What Moore's curve displays is therefore not a homogeneous progression but the ridge line connecting the crests of successive waves of technological substitution.

Thermal constraint is, this genealogy suggests, one important driver of this dynamic among others, and perhaps a more significant one than has generally been recognised. In any case, the two major slowdown phases, in the 1980s and in the mid-2000s, correspond precisely to two documented thermal walls. This, in turn, invites a strong qualification of the standard narrative of the Dennard scaling breakdown, the story of a smooth growth suddenly ruptured. There is no thunderclap in a clear sky, but a tenacious underlying tension expressing itself through recurrent cycles.

One can read this history in two ways. A techno-optimist will argue that since past thermal walls were surmounted, this one will be too. A historian can only note that these overcomings were never more than provisional, and that the unresolved contradiction has returned each time with greater intensity. Whether the current wall will prove different is not for this paper to say.

What can be said with confidence is that recourse to immersion usually marks a moment of end-of-cycle. It functions as a last-resort technology, without any known backstop, and experts already warn that it could prove insufficient for 'smaller-node technologies coming down the pipeline'.[92] Historically, its deployment signals that the corresponding component technology has run its course, and the current return of immersion to centre stage would be rather a signal that its days are numbered.

## Conclusion

The parallel between nineteenth-century coal mines and contemporary data centres may seem far-fetched: what do silicon chips have to do with miners' bodies? In fact, whether living or nonliving, they are both thermosensitive factors of production with hard upper tolerance limits. In each case, an endogenously warming microclimate of production becomes the obstacle, as much an economic limit as a physical one. Jevons establishes this for the mines, Brill for the data centres. The miners' bodies and the chips, however different their nature, are both heat-sensitive elements within the production process, the ones that overheat, and whose tolerance threshold sets the economic ceiling.

Cooling costs rise alike, and the expected outcome is confounded: the deeper you dig, or the more you scale, the less you are able to do it. But the dynamic does not so much encounter an external obstacle as trigger it, as a consequence of its own movement. Heat is the counter-effect here at work, a case of a broader configuration, in which the environmental costs of productive intensification return as growing negative internalities, the kind of costs that a discrete productive entity has to bear. And one form it takes is unbearable heat.

### *Thermal wall and power wall*

This thermal wall is neither the first nor, in all likelihood, the last. Yet it carries its own specificities.

Rack power densities are now soaring, and with them the heat these installations generate. This is partly due to familiar trends of miniaturisation and densification, but also to the specific demands of high-performance computing and AI. Manufacturers increasingly supplement transistor shrinking with what one industry observer calls 'the brute-force stacking of power density',[93] pushing the 'thermal design power' of chips, that is the maximum heat they are rated to generate, to extreme levels.

A second novelty is that the AI boom is fuelling a sharp rise in data centre power demand – up to 40% of which today is absorbed by cooling.[94] Not only does each installation consume more, but ever more are being built, multiplying the sector's aggregate energy demand, to the point where, if trends continue, this growth could 'exceed power utilities' capacity'.[95] The cooling crisis may thus tend to combine with a capacity problem at grid scale. To the internal thermal wall could be compounded an external 'power wall'.

### *The electro-economy in crisis*

This further dimension inscribes itself within a broader crisis of the 'electro-economy'.[96] Since the 1970s, the informatisation of production has made capital increasingly dependent on electricity, which has become 'the indispensable fuel of its most sophisticated technologies'.[97] Yet electricity production has not kept pace. The acceleration of technological innovation and the abbreviation of fixed capital's life-cycle that Mandel associated with the Third Technological Revolution unfolded against a background of stagnant energy performance:[98] no 'established energy production or conversion technique' has ever been capable of following the kind of exponential improvement curve that computing traced.[99] For a long time this mattered little: the computing boom remained energy-light, and the asymmetry posed no problem so long as it did not produce a surge in electricity demand.[100]

The AI boom has changed that. The massive concentration of computing power in AI data centres has initiated a dynamic of electricity consumption that itself takes on an exponential character, something that might be called a 'Moore's curse' in another sense: electricity consumption is growing at its own exponential rate, as if the same logic had expanded its terrain, spilling over from performance into consumption, leading to a cannibalisation of energy-electricity by information-electricity, with the paradox that, extrapolating further, the energy required for the cybernetic control of the economy might come to exceed that required for its primary functioning, putting the economy's digital intellect layer at odds with its corporeal substrate.

### *Thermopolitics: Struggles for the means of cooling*

This unsustainability is not only economic but social and ecological, the grounds on which thermo-economic problems convert into thermo-political conflict.

Such conflicts are rooted in rivalry over resources: given a finite good, appropriating it reduces what others can access. A slogan from a November 2025 protest against a New Mexico data centre, targeting the waste of drinking water for evaporative cooling, put it bluntly: 'We can't drink data'.[101] Such struggles over the means of cooling are erupting across the world, and what applies to water may apply to electricity.

Economic historians of technology have a concept that turns out to be strangely apt here: 'clustering of saturations'. According to Grübler, this notion designates a configuration in which several technological lifecycles reach their end simultaneously, producing cumulative stress effects:[102] taken individually, each could have been absorbed; falling together, they combine their weight and produce crisis. Thermal saturation and energy saturation seem today to be converging, to which a dimension of political saturation may well be added. One scenario would thus take the form of a clustering of saturations – financial risks of speculative collapse aside – the conjunction of a thermal wall, a power wall and a climate wall.

### *AI energy dystopia*

I would like to give another example of these antagonisms in formation, drawn from *Situational Awareness*, a widely circulated text published in 2024 by Leopold Aschenbrenner, a former researcher at OpenAI, one of the many 'organic intellectuals' of AI, whose gesture consists in appropriating the future, in pre-empting imaginations the better to make politics in the present from this projected world. He identifies a major problem for the hyperscalers: energy demand. To illustrate it, he plots two curves: one representing annual electricity generation in the United States, the other showing back-of-the-envelope estimates of near-future total AI electricity demand. The second shoots upward until it collides with the first shortly after 2030: within a few years, on his reckoning, this demand could absorb US production almost entirely.[103] Faced with this finding, the author does not conclude that the AI scaling narrative, so manifestly unsustainable, should be abandoned. Not at all. The tech giants, he suggests, should become their own energy suppliers: building their own nuclear plants, for instance. But this takes time, and international competition leaves little of it. The immediate solution is gas: rapidly constructing gigantic gas plants to power data centres. There is, however, a political precondition: dispensing with those 'well-intentioned but rigid climate commitments' that 'stand in the way of the obvious, fast solution',[104] and this means deregulating and dismantling the remaining climate policies seen as impediments to the surge in electricity production that AI's primitive accumulation requires. And here is the picture of the future these people have in mind:

**Fig. 12** – 'The trillion–dollar cluster. Credit: DALL–E'.[105]

This involuntarily dystopian vignette merely reveals an open secret: the blind commitment to AI scaling, by virtue of its own energy requirements, implies an active and strategic alignment of tech capital with the anti-climate agenda of a position critically labelled as 'fossil fascism'.[106]

Having reached this point, the thermal contradiction we have been tracing expands to its widest scale. It is no longer only a matter of cooling overheating chips, but of doing so in a world where the climate itself is overheating – and where global warming already threatens to exceed, in the coming decades, the limits of human thermal tolerance, that is, the limits of possible survival, in many parts of the world.[107] The politics of cooling becomes, in this light, a matter of life and death. And yet unprecedented resources are being invested in cooling data centres while billions of people will still lack access to potentially vital air conditioning as heat waves worsen.[108] It is a question of what is worth preserving, and the course the answer is taking is written into the infrastructure being built and the climate commitments being dismantled: cooling chips over people, tech capital over living beings.

*Grégoire Chamayou is a research scholar in philosophy at the Centre National de la Recherche Scientifique (CNRS) / Ecole Normale Supérieure (ENS-Ulm) in Paris. His publications include* The Ungovernable Society: A Genealogy of Authoritarian Liberalism *(2021).*

## Notes

**1.** Charles Nehme, 'How do data centers stay cool?' (Amazon Digital Services LLC, 2025).
**2.** Ajay Deewan, 'The cooling crisis: AI's biggest problem is not hallucinations, it is heat', '*Medium*', 30 August 2025, https://pub.towardsai.net/the-cooling-crisis-ais-biggest-problem-is-not-hallucinations-it-is-heat-dabe6fd17c18.
**3.** For this type of work, see R. R. Schmidt and H. Shaukatullah, '*Computer* and telecommunications equipment room cooling', *ITherm 2002*, 30 June 2002, 751–66; and Tyler Sapasap and Xicotencatl Rojas, 'Reviewing cooling strategies for post-Dennard era computing', *Journal of Student Research* 12:4 (2023).
**4.** Gilbert Simondon, *On the mode of existence of technical objects* (Minneapolis: Univocal Publishing, [1958] 2017), 27.
**5.** Simondon, *On the mode of existence of technical objects*, 30.
**6.** I. E. Mouromtseff, 'Water and forced-air cooling of vacuum tubes: Nonelectronic problems in electronic tubes', *Proceedings of the IRE* 30:4 (1942), 190.
**7.** A. L. London, 'Air–coolers for high power vacuum tubes', *Transactions of the IRE Professional Group on Electron Devices* 1:2 (1954), 9.
**8.** Mouromtseff, 'Water and forced-air cooling', 197.
**9.** Garner, *Transistor circuit handbook*, 28.
**10.** Louis E. Garner, *Transistor circuit handbook* (Educational Book Pub. Division, Coyne Electrical School, 1957), 28.
**11.** Allan D. Kraus, *Cooling electronic equipment* (Englewood Cliffs, NJ: Prentice-Hall, 1965), 232.
**12.** Joseph Kaye, 'Review of industrial applications of heat transfer to electronics', *Proceedings of the IRE* 44:8 (1956), 979.
**13.** Gerald L. Ginsberg, 'Thermal management', in *Electronic equipment packaging technology*, ed. Gerald L. Gins-

berg (New York: Springer US, 1992), 215.
**14.** Kraus, *Cooling electronic equipment*, 227.
**15.** Herman Lukoff, *From dits to bits: A personal history of the electronic computer* (Portland, OR: Robotics Press, 1979), 99.
**16.** Lukoff, 'From dits to bits', 99.
**17.** E. L. Watford, 'One engineering solution for temperature and humidity control when computers are added', *Air Conditioning, Heating and Ventilating* (1959), 89.
**18.** Watford, 'One engineering solution', 89.
**19.** 'Electronic computer requires large cooling air supply', *Air Conditioning, Heating and Ventilating* 54:3 (1957), 62.
**20.** N. E. Sheldon, C. A. Nolph and G. A. Fleischer, 'Cooling robot brains', *Electronic Industries & Tele–tech* 16:7 (1957), 46, 102.
**21.** B. V. Anisimov and V. N. Tchetverikov, *Fundamentals of the theory and design of digital computers* (Ohio: Wright-Patterson Air Force Base, 1968), 429.
**22.** USASC, 'Concepts and applications', 106.
**23.** USASC, *Concepts and applications of electronic data processing* (Washington, DC: United States Government Printing Office, 1957), 105.
**24.** 'Sylvania's unique processing center opens at Camillus, New York', *Office Management* 17:8 (1956), 40.
**25.** Leon Albertson, 'Brain "washed" by cold water and fans', *Electrical World*, 3 October 1955, 138.
**26.** Sheldon, Nolph, and Fleischer, 'Cooling robot brains', 46.
**27.** Gail Cooper, *Air-conditioning America: Engineers and the Controlled Environment, 1900-60* (Baltimore: Johns Hopkins University Press, 1998), 40.
**28.** At the same period, Carrier Corporation engineers were pursuing another major air conditioning project for electronic equipment: cooling the RCA's Bizmac, a huge military data centre in Detroit (see Sheldon, Nolph and Fleischer, 'Cooling robot brains').
**29.** 'Air conditioning essential to electronics growth', *DATAmation* 3:8 (1957), 43.
**30.** Steven Gonzalez, 'Cloud ecologies: An environmental ethnography of data centers', Phd thesis, Massachusetts Institute of Technology, 2023, 160.
**31.** Steven J. Dick, *Historical studies in the societal impact of spaceflight* (National Aeronautics and Space Administration, 2015), 168.
**32.** Ernest Braun and Stuart MacDonald, *Revolution in miniature: The history and impact of semiconductor electronics* (Cambridge: Cambridge University Press, 1978), 104.
**33.** Kaye, 'Review of industrial applications', 978.
**34.** Richard P. Feynman, 'There's plenty of room at the bottom', *Engineering and Science* 23:5 (1960), 25.
**35.** Feynman, 'There's plenty of room at the bottom', 26
**36.** Simondon, *On the mode of existence of technical objects*, 147. All subsequent page references are to this edition.
**37.** 'The news of radio', *The New York Times* (July 1.1948).
**38.** 'This capsule challenges vacuum tube', *Popular Science* (September 1948).
**39.** Avram Bar-Cohen, 'Computer-related thermal packaging at the millennial divide', *Electronics Cooling* 6:1 (2000), 33.
**40.** Their thermal tolerance range was considerably narrower (Kraus, *Cooling electronic equipment*, 226), and they introduced a new vulnerability: thermal runaway. 'The cooling problem thus became critical even though the power dissipation was relatively low' (A. E. Bergles, 'Evolution of cooling technology for electrical, electronic, and microelectronic equipment', *IEEE Transactions on Components and Packaging Technologies* 26:1 (2003), 10).
**41.** Bar-Cohen, 'Computer-related thermal packaging', 33.
**42.** Kraus, *Cooling electronic equipment*, 214.
**43.** William Armstrong, 'On the coal question', *Nature* 7:172 (1873), 293.
**44.** Armstrong, 'On the coal question', 293.
**45.** In 1905, the Scottish physiologist J. S. Haldane published the results of a physiological research begun in mine environments, establishing 'wet-bulb temperature' as the relevant index of heat stress, showing that above a wet-bulb temperature of 78°F (25.5°C), 'continuous hard work becomes impracticable' (John Scott Haldane, 'The influence of high air temperatures', *The Journal of Hygiene* 5:4 (1905), 513).
**46.** Grégoire Chamayou, 'The birth of thermopolitics: Wet-bulb temperatures, industrial microclimates, and class struggle in the early 20th century', *Social Studies of Science* 55:6 (2025), 862–90.
**47.** The expression 'Jevons paradox' is itself a retrospective coinage: it was introduced in 1990 by Fred Jevons 'in homage to my distinguished namesake', as he put it, William Stanley Jevons, 'who pointed it out more than 100 years ago' (Fred Jevons, 'Greenhouse – a paradox', *Search* 21:5 (1990), 172). The concept was subsequently taken up in ecological economics (Mario Giampietro, 'Energy efficiency and sustainability in human societies', in *Implikationen der ökologischen Ökonomie für die Regionalokonomie*, ed. Jörg Köhn and John M. Gowdy (Rostock, 1997); John M. Polimeni et al., *The Jevons' Paradox and the Myth of Resource Efficiency Improvements* (London: Routledge, 2008)) and in ecosocialist analysis (John Bellamy Foster, Brett Clark and Richard York, 'Capitalism and the curse of energy efficiency', *Monthly Review* 62:6 (2010)), and has since become current in broader debates on energy, technology and resources.
**48.** Brett Clark and John Bellamy Foster, 'William Stanley

Jevons and the coal question', *Organisation & Environment* 14:1 (2001), 95.
**49.** William Stanley Jevons, *The coal question: An enquiry concerning the progress of the nation, and the probable exhaustion of our coal-mines* (London: Macmillan, 1865), 103. All subsequent page references are to this edition.
**50.** As one industrialist replied: 'you can convey fresh air in such abundance, and at such a temperature to the miner at almost whatever depth he may be' – given that there are 'various ways in which an artificially low temperature could be maintained' (Henry Hussey Vivian, Speech on the coal question, delivered in the House of Commons, 12 June 1866 (London, 1866), 9).
**51.** 'Report of the Commissioners on Coal', *The British Quarterly Review* (July 1872), 44.
**52.** Karl Marx, 'The duration of our supply of coal', in *Exzerpte und Notizen*, Bd. IV/19, Marx-Engels-Gesamtausgabe digital (Berlin: Berlin-Brandenburgische Akademie der Wissenschaften), https://mega-digital.bbaw.de/M0004847.

Rendered fully in English: 'Going deeper enhances the price and with this rise in price our rate of progress and special advantages go to the devil. With deeper working: the heat grows more insupportable, … the ventilation more costly.'
**53.** Glen Lawhon Parker, *The coal industry: A study in social control* (Washington, DC: American Council on Public Affairs, 1940), 4.
**54.** G. E. Moore, 'Cramming more components onto integrated circuits', *Proceedings of the IEEE* 86:1 (1965), 83.
**55.** Moore, 'Cramming more components', 84. All subsequent page references are to this edition.
**56.** R. H. Dennard, F. H. Gaensslen, V. L. Rideout, E. Bassous and A. R. LeBlanc, 'Design of ion-implanted MOSFET's with very small physical dimensions', *IEEE Journal of Solid-State Circuits* 9:5 (1974).
**57.** Jacob Barton Leverich, *Future scaling of datacenter power-efficiency* (Stanford University *Computer* Science Department, 2014), 7.
**58.** Vikram Sekar, 'How Dennard scaling allowed transistors to shrink', *Vik's Newsletter*, 12 January 2025, https://www.viksnewsletter.com/p/how-dennard-scaling-actually-works.
**59.** Stephan Ohr, 'Intel technologist cites power as biggest issue', *EE Times*, 5 February 2001, https://www.eetimes.com/intel-technologist-cites-power-as-biggest-issue/.
**60.** Tom Mainelli, 'Future chips: Headed for heat problems', *PC World*, 7 February 2001.
**61.** John Markoff, 'Intel's big shift after hitting technical wall', *New York Times*, 17 May 2004.
**62.** Markoff, 'Intel's big shift'.
**63.** Avram Bar-Cohen, *Encyclopedia of thermal packaging*, set 2: Thermal packaging tools (World Scientific, 2014), 1.
**64.** Markoff, 'Intel's big shift'.
**65.** N. S. Kim et al., 'Leakage current: Moore's law meets static power', *Computer* 36:12 (2003), 69.
**66.** Wataru Nakayama, 'Heat in computers: Applied heat transfer in information technology', *Journal of Heat Transfer* 136 (2013), 6.
**67.** 'Silicon on insulator's new momentum', *Electronic Business* (March 2005), 42.
**68.** Steven Woo, 'The thermal challenges of Moore's law: Part 1', video presentation, Rambus, 16 April 2020, https://www.rambus.com/blogs/the-thermal-challenges-of-moores-law-part-1/.
**69.** It should be noted that the industry's response was not limited to this inflation of cooling structures. Other strategies were explored, which more closely resemble what Simondon calls processes of concretisation. This was notably the case with the development of multicore architectures, based on parallel processing: distributing tasks across several lower-frequency cores rather than pushing a single one harder, thereby reducing both power consumption and thermal dissipation. Or again the so-called 'dark silicon' strategies, which consist in putting successive portions of the circuitry into standby to avoid overheating (David Atienza et al., 'A 20-year retrospective on power and thermal modeling and management', *IEEE Design & Test* 42:6 (2025), 1). These approaches illustrate other ways of arbitrating the trade-off between performance and heat. Hence, if in practice 'the heat produced does not scale with the number of transistors … This is because of judicious design, introduction of sleep modes where parts of the microprocessor are inactive when not needed, the use of multicores, etc.' (Krishna Seshan, 'Scaling – its effects on heat generation and cooling of devices', in *Handbook of Thin Film Deposition* (Norwich, NY: William Andrew Publishing, 2012), 41).
**70.** Kenneth G. Brill, *The invisible crisis in the data center: The economic meltdown of Moore's law* (Uptime Institute, 2007), 2.
**71.** Christian Belady, 'In the data center, power and cooling costs more than the IT equipment it supports', *Electronics Cooling*, 9 February 2007, https://www.electronics-cooling.com/2007/02/in-the-data-center-power-and-cooling-costs-more-than-the-it-equipment-it-supports/.
**72.** As Brill put it: 'The widening delta between the faster rate of increase of server computational performance and the slower rate of increase in energy efficiency growth results in a dramatic increase in … cost burdens of providing power and cooling' (Brill 2007, 2).
**73.** Richard Brown et al., 'Report to Congress on server

and data center energy efficiency: Public Law 109-431', 1 January 2007, https://doi.org/10.2172/929723.

**74.** Bar-Cohen, *Encyclopedia of thermal packaging*, 2.

**75.** Nehme, 'How do data centers stay cool?'

**76.** AmiNext, 'The thermodynamics of capital: When liquid cooling becomes an economic imperative', 1 June 2026, https://www.aminext.blog/en/post/liquid-cooling-economics-ai-infrastructure-cost-1.

**77.** AmiNext, 'The thermodynamics of capital'.

**78.** Immersion IQ, 'The cooling crisis: Why traditional data center cooling has hit the wall', 27 June 2025, https://immersioniq.io/the-cooling-crisis-why-traditional-data-center-cooling-has-hit-the-wall/.

**79.** Cornell Aeronautical Laboratory, *Heat transfer in miniaturized electronic equipment* (Bureau of Ships, Navy Dept., 1955), 34.

**80.** Wataru Nakayama, 'Exploring the limits of air cooling', *Electronics Cooling*, 9 August 2006, https://www.electronics-cooling.com/2006/08/exploring-the-limits-of-air-cooling/.

**81.** John H. Seely and R. C. Chu, *Heat transfer in microelectronic equipment: A practical guide* (New York: M. Dekker, 1972), 1.

**82.** Wataru Nakayama, 'Heat in computers', 2.

**83.** S. Oktay, R. Hannemann and A. Bar-Cohen, 'High heat from a small package', *Mechanical Engineering* 108:3 (1986), 37.

**84.** Wataru Nakayama, 'Heat sinks in computers' in *Thought-evoking approaches in engineering problems*, ed. Yoshimo Ito (Cham: Springer, 2014), 127.

**85.** Nakayama, 'Heat in computers', 5.

**86.** Nakayama and Bergles, 'Cooling electronic equipment', 25.

**87.** Nakayama, 'Heat sinks in computers', 126.

**88.** Nakayama, 'Heat sinks in computers', 5.

**89.** Bar-Cohen, '*Computer*-related thermal packaging', 32.

**90.** Oktay, Hannemann and Bar-Cohen, 'High heat from a small package', 36.

**91.** Burg and Ausubel, 'Moore's law revisited', 14.

**92.** James Myers, 'Future chips will be hotter than ever', *IEEE Spectrum*, 16 April 2025, https://spectrum.ieee.org/hot-chips.

**93.** AmiNext, 'The thermodynamics of capital'.

**94.** Karthik Ramachandran et al., 'As generative AI asks for more power, data centers seek more reliable, cleaner energy solutions', *Deloitte Insights*, 19 November 2024. Data centres in the U.S. already consume around 4% of the country's electricity, AI workloads currently accounting for 10% to 20% of this consumption (Martin C. Offutt and Ling Zhu, 'Data centers and their energy consumption: Frequently asked questions', *CRS Report* 48646 (2025)). But it is the trend that matters: AI training drove a tenfold increase in computing power usage between 2018 and 2022 (Alexandra Sasha Luccioni et al., 'The climate and air pollution impacts from AI's energy consumption'. *Nature Sustainability* (2025)). Globally, data centres account for approximately 1.5% of electricity consumption, a share projected to more than double by 2030 (IEA, *Energy and AI*, 10 April 2025, https://www.iea.org/reports/energy-and-ai).

**95.** Gartner, 'Gartner predicts power shortages will restrict 40% of AI data centers by 2027', 12 November 2024, https://www.gartner.com/en/newsroom/press-releases/2024-11-12-gartner-predicts-power-shortages-will-restrict-40-percent-of-ai-data-centers-by-20270.

**96.** David Thomas, 'Keeping the lights on: Oil shocks, coal strikes, and the rise of electroculture', *Mediations* 31:2 (2018), 99.

**97.** Thomas, 'Keeping the lights on', 103.

**98.** Ernest Mandel, *Late capitalism*, rev. ed. (London: Verso, 1975), 9.

**99.** Vaclav Smil, *Energy transitions: History, requirements, prospects* (Santa Barbara, CA: Praeger, 2010), 121.

**100.** Vaclav Smil, *Energy and civilization: A history* (Cambridge, MA: MIT Press, 2017), 200.

**101.** Leila Chapa and Paloma Chapa, 'Protesters pressure lawmakers during NM legislature opening day', *The Daily Lobo*, 26 January 2026, https://www.dailylobo.com/article/2026/01/protesters-pressure-lawmakers-during-nm-legislature-opening-day.

**102.** Arnulf Grübler, 'Technology diffusion in a long–wave context: The case of the steel and coal industries' in *Life Cycles and Long Waves*, eds. Tibor Vasko and Robert U. Ayres (Berlin: Springer, 1990), 143.

**103.** Leopold Aschenbrenner, *Situational awareness: The decade ahead* (2024), 83.

**104.** Aschenbrenner, 'Situational awareness', 85.

**105.** Aschenbrenner, 'Situational awareness', 76.

**106.** Cara Daggett, 'Petro–masculinity: Fossil fuels and authoritarian desire', *Millennium* 47:1 (2018), 25–44. See also Zetkin Collective and Andreas Malm, *White skin, black fuel: On the danger of fossil fascism* (London: Verso. 2021).

**107.** Daniel J. Vecellio et al., 'Greatly enhanced risk to humans as a consequence of empirically determined lower moist heat stress tolerance', *Proceedings of the National Academy of Sciences* 120:42 (2023), e2305427120.

**108.** Up to four billion people by 2050, it is estimated (Giacomo Falchetta et al., 'Inequalities in global residential cooling energy use to 2050', *Nature Communications* 15:1 (2024), 7874). At the same time, the world already faces a looming 'cold crunch': energy demand from air conditioners could more than triple by 2050, accounting for 30% of global energy consumption and putting enormous strain on electricity systems (IEA, *The Future of Cooling* (2018)).