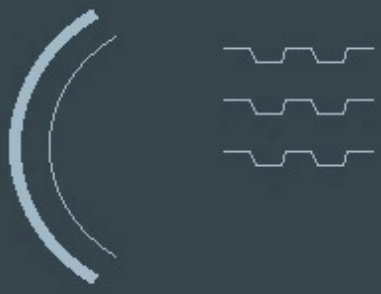
INDUSTRIES OF INDIA

THE HYDROELECTRIC DAM

Water, Concrete, Power



KETINENI RAMARAO

T H E
H Y D R O E L E C T R I C
D A M

Water, Concrete, Power

Ketineni Ramarao

Reference numbers in square brackets resolve in the companion volume, The Hydroelectric Dam: Source Companion.

Contents

PART I - THE RESOURCE 4
Chapter 1 - The mathematics of falling water 5
Chapter 2 - Three ways to catch a river 8
PART II - THE LONG HISTORY 11
Chapter 3 - The water wheel 12
Chapter 4 - The turbine men 15
PART III - THE ELECTRIC MARRIAGE 18
Chapter 5 - First lights 19
Chapter 6 - The long wire 22
PART IV - THE DAM 25
Chapter 7 - Holding back the river 26
Chapter 8 - Concrete, earth and the building of the wall 28
Chapter 9 - Water past the wall 31
PART V - THE MACHINE HALL 34
Chapter 10 - The modern turbine 35
Chapter 11 - From torque to current 38
Chapter 12 - The anatomy of a powerhouse 41
PART VI - THE GIANTS 44
Chapter 13 - Hoover and the federal river 45
Chapter 14 - The New Deal valleys 48
Chapter 15 - The planned giants 51
Chapter 16 - The binational giants 54
Chapter 17 - The Chinese century 56
PART VII - THE GRID’S WORKHORSE 59
Chapter 18 - The flexible generator 60
Chapter 19 - The water battery 62
PART VIII - THE BUSINESS 64
Chapter 20 - The economics of patience 65
Chapter 21 - The financiers’ retreat and return 67
PART IX - THE WORLD MAP 69
Chapter 22 - Where the water falls 70
Chapter 23 - The unbuilt rivers 72
PART X - THE COSTS DOWNSTREAM 74
Chapter 24 - The drowned valleys 75
Chapter 25 - The river interrupted 77
Chapter 26 - The carbon question 79

PART XI - WHEN DAMS FAIL 81
Chapter 27 - Banqiao 82
Chapter 28 - The aging fleet 84
PART XII - THE RECKONING 86
Chapter 29 - The rivals 87
Chapter 30 - Taking dams down 89
Chapter 31 - Conclusion: the long arc 91
Chapter 32 - A hydroelectric glossary 93

PART I - THE RESOURCE

Chapter 1 - The mathematics of falling water

Every hydroelectric station on earth is an answer to the same short equation. Power equals the density of water, times the pull of gravity, times the flow, times the height through which the water falls. Write it with the usual symbols and it fits on one line: P equals rho times g times Q times H. The density of water is about a thousand kilograms per cubic metre, gravity adds 9.81 metres per second squared, and the engineer supplies the last two terms - the flow of water through the machines, in cubic metres per second, and the head, the vertical distance the water drops to reach them [C35]. That product is the theoretical hydraulic power; what reaches the grid is less, discounted by friction in the waterway and the efficiency of turbine and generator - around nine-tenths in a well-designed modern station [C35].

The equation is simple enough for a school exercise, yet it organises a global industry that generated about 4,578 terawatt-hours of electricity in 2024, more than any other renewable source [C8]. Everything about a hydroelectric project - where it can be built, what it costs, how big its machines are, what kind of dam it needs - follows from those two site-given numbers, flow and head. A river with an enormous flow and a modest drop, like the Parana at Itaipu, demands one kind of machine and one kind of dam. A mountain torrent falling through hundreds of metres of pipe, as in the Alps or Norway, demands another. The water itself decides, and the entire craft of hydroelectric engineering consists of listening to what the water’s arithmetic allows.

Head is the quieter of the two numbers and often the more valuable. A cubic metre of water falling one metre releases about 9.81 kilojoules of energy, roughly the energy in a quarter of a millilitre of petrol - a seventeenth of a teaspoon - trivial in itself. But water arrives continuously, and the head multiplies it. Raise the drop to a hundred metres and the same cubic metre yields nearly a megajoule, enough to run a kettle for five minutes. Raise it to the three hundred metres and more of a high Alpine scheme and a steady stream that would fill a garden pond in a minute can drive a power station. This is why the map of hydropower is a map of relief. Mountains and plateaus give water its energy; flat rivers, however mighty their flow, give up their power grudgingly, a few metres of head at a time, and must be dammed across their whole width to make even that usable.

Flow is the louder number, the one that fills photographs of spillways in flood. It varies in ways engineers cannot negotiate with. A tropical river may carry ten times as much water in its wet season as in its dry; a snowmelt river peaks in late spring and shrinks by autumn. Because power is flow times head, a station’s output swings with the river unless something intervenes to smooth the flow. That

something is the reservoir, and the choice between storing water and taking it as it comes divides the industry into its two great families of plant [C1].

The first family takes the river as it finds it. A run-of-river station diverts part of the flow through a canal or a closed pipe called a penstock, lets the natural slope of the riverbed provide the head, and returns the water downstream [C1]. It has little or no storage, so it produces what the river offers, hour by hour and season by season. The second family builds a dam high enough to pond a reservoir behind it. Now the flow can be held back when it is not needed and released when it is, and the reservoir surface itself sets the head. Storage converts a weather report into a warehouse.

Between these two numbers and this one choice sits a third term that decides whether any of it is worth building: time. A power station’s rated capacity, measured in megawatts or gigawatts, says how much power it can produce at full stretch. What buyers and grids actually receive is energy - power multiplied by hours - and the ratio between what a plant actually generates in a year and what it could have generated running flat out is its capacity factor. Hydropower’s global generation of 4,578 terawatt-hours in 2024 came from about 1,443 gigawatts of installed capacity, an implied worldwide capacity factor of roughly 36 per cent [C8]. That figure conceals enormous range: a run-of-river plant on a steady glacier-fed river may exceed 60 per cent, while a storage plant kept in reserve for peak hours and drought years may run at a small fraction of its rating by deliberate design.

This is the first thing to grasp about the economics of falling water: capacity and energy are different goods, and a hydroelectric dam can be built to sell either. Some stations are built for volume, churning out cheap energy around the clock. Others are built for availability, sitting ready to deliver full power in the minutes when a grid most desperately needs it. The same equation governs both; the dam, the reservoir and the machines are simply tuned to different terms of it.

The equation also explains the industry’s obsession with efficiency. Unlike a steam engine, which throws away most of the heat it is fed, a water turbine converts nearly all the energy it receives. Modern hydroelectric machines routinely turn more than nine-tenths of the water’s energy into electricity - a conversion efficiency no combustion technology approaches [C35]. The losses that remain are friction in the waterways, turbulence in the runner, and resistance in the generator and transformer. When the resource itself is free and falls out of the sky as rain, efficiency is the whole game: every point recovered is pure gain, sold for decades.

And the resource is, within human timescales, endless. The sun lifts moisture from the seas, the winds carry it over the continents, and it falls on the highlands

and runs back down to the ocean. A dam taps one reach of that loop. It consumes no fuel and exhausts nothing but the height the water had anyway. The water that spins the turbines at Itaipu or Three Gorges flows on to the next barrage, the next irrigation canal, the estuary, the sea. What the industry captures is not the water but its fall - a distinction that sounds like a quibble and is in fact the whole business.

There is one more consequence of the simple equation, and it is the one that gives this book its subject. Flow is given by the river, but head is not. Head must usually be made. Nature provides a few ready-made drops - Niagara being the most famous - but most of the world’s hydropower exists because engineers raised the water themselves, building a wall across a valley and letting the river climb it. The dam is the device that manufactures head. Everything else in the story - the turbines, the generators, the transmission lines, the financiers, the displaced villages, the dried deltas, the resurrected salmon - follows from that single act of making water stand higher than it wants to.

To make electricity from a river you need only a wheel and a generator; men lit a country house that way in 1878 [C3]. To make power for cities and industries you need head measured in tens or hundreds of metres and flow measured in thousands of cubic metres per second, and there is no way to hold such water back except with a structure massive enough to resist it. Hoover Dam, the archetype, weighs more than six and a half million tons and the water at its base presses with a force of about 45,000 pounds per square foot [C5]. The hydroelectric dam is that structure wedded to a power station, and its double nature - part wall, part machine - runs through everything that follows.

Chapter 2 - Three ways to catch a river

Ask how many dams produce electricity and the answer startles most people: almost none of them. The United States has more than 90,000 dams, and fewer than 3 per cent of them generate power; the rest exist to store drinking water, irrigate fields, hold back floods, or simply to make lakes for boats and stock ponds [C1]. The hydroelectric dam is a specialised member of a much larger family of river walls, and the first thing the industry's own taxonomy does is separate the plants that store from the plants that borrow and the plants that recycle.

The impoundment plant is the classic image: a dam across a valley, a reservoir spreading behind it, a powerhouse at the foot of the wall. The reservoir is the point. It accumulates the river's uneven gifts - the spring flood, the wet season, the good year - and pays them out on the operator's schedule. Water released from storage drops through the turbines, spins them, and spins the generators coupled to them; the same releases can be shaped for flood control, irrigation, navigation and the fish, all from one ponded body of water [C1]. Impoundment is the most common form of large hydroelectric development, and it is the form that carries the industry's heaviest consequences, because a reservoir big enough to matter is a reservoir that drowned something.

The diversion plant, usually called run-of-river, builds little or no storage. It taps the river at one point, carries part of the flow along a canal or through a penstock to a powerhouse set lower down, and returns the water to the river below [C1]. Some diversions use a small weir rather than a true dam; a few need no barrier at all. Their virtue is lightness: a modest flooded area, a smaller civil works bill, a river that still flows. Their price is obedience. With no reservoir, a run-of-river station generates what the river happens to be carrying. In the wet season it may spill water it cannot use; in the dry season it shrinks toward a trickle. It sells energy, not reliability.

The third type abandons the river's energy budget almost entirely and builds a battery out of water. A pumped storage plant has two reservoirs at different elevations. When electricity is cheap and plentiful, the plant runs its machines backward, pumping water uphill; when electricity is scarce and dear, it lets the water fall back through the turbines [C1]. Round-trip, it consumes more energy than it produces - the laws of thermodynamics see to that - and yet it is one of the most valuable machines on a modern grid, because energy delivered at eight in the evening is worth far more than the same energy at three in the morning. Pumped storage is so effective at this arbitrage that it accounts for more than 90 per cent of all the energy storage installed in the world [C8].

These three are not rival designs so much as three answers to the question of time. Impoundment stores the river's own energy across seasons. Diversion declines to store and accepts the river's mood. Pumped storage stores other generators' energy across hours. The world's hydro fleets mix all three, and the mix is shifting: in 2024 the world held about 1,254 gigawatts of conventional hydropower - impoundment and run-of-river together - and 189 gigawatts of pumped storage, with annual pumped-storage additions nearly doubling in two years [C8].

Scale cuts across all three types, and the industry's size classes are worth fixing precisely because the word "hydropower" spans five orders of magnitude. In the definitions used by the United States Department of Energy, a micro plant makes up to 100 kilowatts - enough for a farm or a village; a small plant runs from 100 kilowatts to 10 megawatts; and large hydropower begins above 30 megawatts [C1]. At the top of the range the numbers stop making domestic sense. The largest station on earth, Three Gorges on the Yangtze, carries 32 main generating units behind a dam 2,309 metres long and 185 metres high [C23]. The smallest commercial schemes power a single lodge. The same equation governs both; only the zeros move.

It is tempting to arrange these types in a moral order - run-of-river gentle, impoundment brutal, pumped storage clever - and the industry's critics and defenders both reach for that framing. The record resists it. A run-of-river plant with a long diversion can leave a dried-out stretch of river as thoroughly as a dam can flood one. A reservoir can be managed as a flood shield, an irrigation bank and a fisheries refuge as well as a power source; the Tennessee Valley Authority was created in 1933 to do exactly that combination of jobs across an entire watershed [C31]. And a pumped storage plant is only as clean as the electricity it stores: pumping on coal-fired midnight power and calling the result green storage is an accounting choice, not a physical fact. What the physical fact does guarantee is flexibility, and flexibility has become the scarcest commodity on grids rich in wind and sunshine - a theme the book returns to when it follows the dam into the twenty-first-century grid.

One more member of the family deserves mention here, because it contradicts the tidy premise that dams make power. Most dams make nothing but water services, and many hydroelectric dams were not born as power projects at all. Hoover Dam was authorised for three purposes - irrigation water, flood control and, third, power [C2]. The reservoir behind Aswan was built to end the Nile's ancient alternation of flood and drought; the turbines in its base came along with the storage [C47]. Power often turned out to be the revenue that paid for the rest, the cash register bolted onto a water project. That inversion - the dam as a water

structure financed by its electricity - explains a great deal of the industry’s politics, and much of the twentieth century’s dam-building geography.

The geography itself is the subject of a later chapter, but its outline belongs to the resource. Hydropower flows where water and relief coincide: the Himalayas and their piedmonts, the Andes, the Congo basin, the great shield rivers of Canada and Russia, the monsoon-drained ranges of Southeast Asia, the fjord-cut plateaus of Norway. Some of the richest endowments remain almost untouched - Africa has harnessed only about a tenth of its technical potential on the industry association’s estimate, and the association put the continent’s unharnessed share at the centre of its 2025 outlook [C8]. Other regions built out their best sites generations ago and now spend their money maintaining, enlarging and rethinking what they have. Between the untouched river and the aging dam lies the whole arc of this book: how the falling water was first caught, how the catching was industrialised, what it built, what it cost, and what happens to it now.

PART II - THE LONG HISTORY

Chapter 3 - The water wheel

Long before anyone made electricity, civilisations were already taxing rivers for their fall. The Greeks were grinding grain with water wheels more than two thousand years ago, and the Egyptians had been lifting water with the Archimedes screw two centuries before that [C2]. The principle those machines exploited is the one in the first chapter of this book: moving water carries energy, and a wheel dipped into its path can take some of it. For roughly two millennia the water wheel was the most powerful prime mover humanity possessed, and the geography of early industry - the mill at the weir, the forge by the stream - is written in its terms.

The wheel’s evolution was a slow argument with efficiency. The earliest wheels were undershot: paddles dipping into a flowing stream, turned by the current’s push alone. They wasted almost everything, capturing only the water’s motion and none of its weight. The overshot wheel, fed at the top so that water filled its buckets and pulled it around by gravity, was far stronger, because it used the fall as well as the flow - head, in the language of the previous chapter, harvested one bucket at a time. Between those two designs lay most of the mechanical ingenuity of the pre-industrial world, and the distinction between impulse and weight, push and fall, would return centuries later to divide the water turbine into its own two families.

The ancient world’s most startling hydraulic complex was Barbegal, in Roman Gaul. Sometime in the second century AD, the aqueduct supplying the city of Arelate, modern Arles, was tapped to feed two parallel cascades of eight overshot wheels each - sixteen wheels in all, stepped down a steep hillside, the output of each feeding the fall of the next [C75]. It was a flour factory on an industrial scale, and archaeologists who have studied its carbonate-encrusted ruins call it one of the first industrial complexes in human history [C75]. Nothing on that scale was built again for well over a thousand years. What survived everywhere was the local mill. By the time of Domesday Book, England alone recorded thousands of watermills, and across medieval Europe, the Islamic world and China the wheel ground grain, fulled cloth, sawed wood, blew furnace bellows and lifted hammers. China had gone furthest in some respects: during the Han dynasty, between 202 BC and 9 AD, vertical water wheels already drove trip hammers that pounded and hulled grain, broke ore and worked in early papermaking [C3].

The mill shaped more than the economy; it shaped property and law. A mill site was a monopoly of position - the lord who owned the fall of water owned the power, and medieval peasants were often bound to grind at their lord’s mill and pay for the privilege. Wherever the wheel went it created a small geography of

rights: rights to the water, rights to the fall, rights against anyone upstream who might pond or divert it. Centuries later, when mills became power stations and millers became utilities, those riparian questions - whose water, whose fall, whose flood - returned as the legal architecture of the hydroelectric age. The arguments over dam licenses and water rights that fill modern regulatory dockets are recognisably the same arguments, conducted at continental scale.

Two limits held the wheel back, and both matter for what came next. The first was power density. Even a large wheel turned slowly and yielded only a few horsepower; industry grew by adding wheels, not by scaling them, and a great mill was a row of wheels rather than one big machine. The second was the wheel's appetite for site. It wanted a dependable stream with a modest fall, and it wanted the mill built directly over the water, because power could leave the wheel only as turning shafts, belts and gears. Mechanical power could not travel. The factory had to go to the river, and the river's geography became industry's geography.

That second limit is easy to underestimate, because the steam engine that broke it is remembered as a fuel story rather than a location story. Coal freed the factory from the riverbank; the mill town moved to the coalfield and the railway. What is less often noticed is that water power was never defeated on cost or efficiency - it was defeated on portability. The wheel remained cheaper to run than the engine wherever a good site existed. What it could not do was send its power anywhere. The whole future of hydropower depended on finding a way to carry the river's work to the customer, and that way would not be mechanical.

The eighteenth century brought the wheel to its perfection just as it was becoming obsolete. Engineers like John Smeaton in Britain ran systematic experiments on wheel design, measuring what the millwrights had known by rule of thumb and proving the overshot wheel's superiority with numbers. In France, the hydraulic engineer Bernard Forest de Belidor published his Architecture Hydraulique in the mid-1700s, a systematic treatise on harnessing water that the US Department of Energy's official history still credits as the founding document of modern hydropower thinking [C2]. Belidor's generation turned millwrighting into engineering: flows measured, heads calculated, blades shaped by theory. The intellectual machinery for the turbine was assembled a century before the turbine.

Richard Arkwright's Cromford Mill, opened in 1771 in the Derwent valley, shows the wheel at its industrial peak. Arkwright's water frame spun cotton by machinery, and he powered it with water - one of the world's first true factory systems ran on a river [C3]. Cromford is usually told as the beginning of the factory age. It is also, read from the other direction, nearly the end of an older story: the largest statement of the proposition that industry must live where the water falls. Within two generations the steam engine had refuted the proposition.

Within four, the turbine and the dynamo would overturn it again from the other side, letting the river stay in its valley and sending only the power to town.

The wheel left one more legacy, less tangible than millraces and factory sites. It taught engineers to think of a river as a machine - an input of water and fall, an output of work, with losses to be hunted down. That framing, born at the millpond, is still the framing of the modern power station, whose control rooms display the same two numbers the miller watched: how much water is coming, and how far it falls. The turbine that replaced the wheel was not a rejection of two thousand years of practice. It was those two thousand years, redesigned for a world that had learned to move power down a wire.

Chapter 4 - The turbine men

The water turbine was invented to escape the water wheel’s ceiling. A wheel is large, slow and heavy; it must sit in the river it serves; and it wastes much of the head above its axle. The engineers who replaced it in the nineteenth century were not trying to build power stations - electricity had not yet arrived to ask them - they were trying to build a better mill engine. In doing so they designed the machine that still spins at the heart of every hydroelectric dam.

The first of them was Benoit Fourneyron, a French engineer trained at the famous school of mines at Saint-Etienne. In 1827 he produced the first practical version of what became known as the Fourneyron reaction turbine, a machine of about six horsepower in which water pressed outward through curved guide passages and pushed against a wheel of curved blades [C3]. Two ideas made it revolutionary. The water did not merely strike the wheel; it flowed through it under pressure, and the blades turned the pressure itself into work. And because pressure does the pushing, the whole machine could sit submerged, compact and fast-spinning, using every metre of head the site offered rather than the wheel’s diameter. By 1832 Fourneyron’s improved turbines were driving mills directly, and within a few years they were being exported across Europe and to America.

The second figure was James Bicheno Francis, a British-born engineer who emigrated to the United States and became chief engineer of the great mill city of Lowell, Massachusetts. Lowell ran on water: a network of canals fed hundreds of wheels, and Francis spent years measuring them, ruthlessly quantifying how much of the river’s energy each design actually captured. In 1849 he produced the machine that bears his name [C3]. In a Francis turbine the water enters around the full circumference of the runner, flows inward between its blades and leaves along the axle - a mixed-flow design that handles the middling heads and big flows of most of the world’s dam sites. It remains the most widely used water turbine in the world, a hundred and seventy-five years later [C3][C4].

The third man worked at the opposite extreme of the resource. When California’s gold country industrialised after 1849, the miners’ water arrived in small, fierce streams falling through hundreds of feet. Fourneyron and Francis machines drowned in such conditions; what the Sierra Nevada needed was a wheel that could take a jet. Lester Allan Pelton, a carpenter and millwright who had followed the gold rush west, spent the 1870s perfecting one. His wheel carried a ring of paired buckets, each split by a central ridge: the jet struck the ridge, divided cleanly, curled through nearly a hundred and eighty degrees in the polished cups, and gave up almost all its velocity before falling away. Pelton patented the design in 1880 [C3]. It was an impulse machine - pure jet against bucket, the water

wheel’s oldest idea refined to its limit - and it turned the high mountain valleys of the world, useless to reaction turbines, into prime power country.

The fourth name belongs to the lowlands. By the early twentieth century engineers wanted to harvest rivers where the head was only a few metres but the flow immense - the great flat rivers of Europe and America. A propeller in the stream would do it, but a fixed propeller is efficient only at one flow. Viktor Kaplan, an Austrian professor, solved the problem in 1913 with a propeller turbine whose blades pivot as the wicket gates open and close, so the machine reshapes itself to the river’s flow [C3]. The Kaplan turbine made the lowland river, and the run-of-river plant, fully economic. The dates themselves carry a small lesson in attribution: Kaplan’s biographers note that he arrived at his machine while trying to improve the Francis turbine, and even the year of invention is cited variously - 1913 by the industry association, 1919 in the US Department of Energy’s account, reflecting the gap between patent, prototype and proof [C4][C34]. Firsts in engineering are rarely single moments; they are arguments that take years to settle.

Between the four machines lies a complete map of the resource, and engineers eventually gave the map a coordinate: specific speed, a single number computed from head, flow and rotation speed that sorts every site toward its proper machine. High head and modest flow point to Pelton’s buckets. Middling head and generous flow point to Francis. Low head and great flow point to Kaplan’s adjustable propeller. A modern design office still begins exactly there, choosing the machine to fit the site’s arithmetic rather than the reverse.

The precision of that sorting is easy to underrate. A turbine is not a general-purpose engine detuned for a site; it is a machine optimised around one point in the head-flow space, and straying far from that point costs efficiency the operator will feel for decades. This is why the giants of the industry run to type: Itaipu and Three Gorges, with their enormous flows and middling heads, are Francis stations through and through, while the station-service units at Hoover - small machines fed from the full head of the reservoir - are Pelton wheels [C50][C63][C61]. The nineteenth century’s three answers have never needed a fourth; they have only needed to be made larger, smoother and more precisely shaped.

The family had one more branch worth noting. For the very lowest heads - tidal ranges, river mouths, canals - engineers eventually turned the propeller turbine horizontal and wrapped it in the flow itself: the bulb turbine, with the generator sealed inside a streamlined pod in the middle of the waterway. It is a long way from Fourneyron’s six-horsepower wheel, and yet the lineage is unbroken.

What the turbine men collectively achieved was bigger than any of their machines. They raised efficiency from the wheel’s wasteful half-or-less to figures

above ninety per cent; they shrank the prime mover from a building to a casting; and, decisively, they made the machine fast. A turbine spins at hundreds of revolutions per minute - precisely the speed at which a dynamo wants to be driven. When the electrical industry arrived in the 1880s looking for prime movers, it found a mature machine already waiting: compact, efficient, fast-spinning, and happy anywhere water falls. The marriage took less than a decade to consummate, and it is the subject of the next part.

PART III - THE ELECTRIC MARRIAGE

Chapter 5 - First lights

Hydroelectricity was born in a country house. In 1878, at Cragside in Northumberland, the industrialist William Armstrong - arms manufacturer, engineer, and one of the great Victorian collectors of gadgets - piped water from a lake on his estate down to a small turbine and lit a single lamp. It was the first time falling water had been turned into electric light [C3]. The demonstration was domestic, almost quaint, but the components of an entire industry were present in it: a head of water, a machine to catch the fall, a dynamo, a wire, a customer.

The customer multiplied quickly. In 1880 a dynamo driven by a water wheel lit brush arc lamps in a furniture factory at Grand Rapids, Michigan. The following year brought plants at Ottawa in Ontario, at Dolgeville in New York, and at Niagara Falls itself, where a small station supplied the local mills and lit a few buildings [C3]. These were direct-current machines serving customers within a stone’s throw of the water, because direct current could not travel. The pattern was set early and would take a decade to break: the power was where the water was, and if you wanted it you went to it.

The plant the industry conventionally calls its first is the Vulcan Street station at Appleton, Wisconsin, which opened in September 1882 and supplied a paper mill and a handful of homes along the Fox River - the first hydroelectric plant selling electricity to multiple customers [C3]. The date is worth a moment of precision, because the Appleton plant is often called the first hydroelectric station in the world, full stop, which it was not: Cragside was earlier, and Niagara’s small 1881 installation preceded it too. What Appleton began was not hydroelectricity but the hydroelectric utility - power as a business, sold to whoever would pay. Within ten years, by the industry’s own historical count, hundreds of hydroelectric stations were operating around the world [C3].

The early stations were improvised, nervous affairs. There were no standards for voltage or frequency, no textbooks, no spare parts; every plant was a prototype and every operator an experimenter. Records from the era are full of machines built for rivers that flooded them, dynamos wound for voltages nobody else used, promoters who sold power they had not yet learned to generate. The Appleton plant itself began in an Edisonian spirit of hopeful improvisation, and its output was famously unsteady - the voltage rose and fell with the river, and the lamps brightened and dimmed with it. What carried the industry through its experimental decade was the steadiness of the resource. A steam plant needed coal delivered forever; a water plant needed only the river to keep coming, and the river kept coming. Operating costs were so low that even clumsy plants made money, and the margins funded the learning.

The economics were so favourable that the industry grew even where the resource was modest. Most early stations were small - a wheel at a mill dam, a dynamo the size of a barrel - and they clustered where waterpower had already been industrialised: New England's mill towns, the glaciated lake country of the upper Midwest, the mill races of Ontario. By the turn of the century, hundreds of small hydroelectric plants were in operation worldwide [C3]. The pattern reached the far edges of the industrialising world with surprising speed: in 1905, a hydroelectric station with 500 kilowatts of capacity was built on the Xindian creek near Taipei, bringing electric light to Taiwan under Japanese rule [C3]. Wherever there was falling water and a customer within reach of a wire, the combination found investors.

It is worth dwelling on how odd the first product was. The early stations sold light - a luxury good - to a handful of customers, because electric motors were still novelties and there was almost nothing else to sell. Arc lamps were too bright for houses and were strung in streets and factories; incandescent lamps, newer and gentler, went indoors. A hydroelectric plant of the 1880s was, in effect, a way of converting a waterfall into street lighting, and its fortunes rose and fell with the price of the alternative: gaslight. Every improvement in either technology shifted the contest. What steadied the young industry was that its fuel cost nothing, so it could wait out any price war the gasworks started.

They also funded an appetite. The customers of the 1880s wanted light, but industry wanted power, and the biggest single market of the era turned out to be electrochemistry. The new processes for aluminium, abrasives, alkalis and carbides were gluttons for cheap electricity, and they went looking for it where it was cheapest - at the water's edge. Niagara Falls, which had been a scenery business with a mill stream attached, found itself courted by industrialists who wanted not the view but the fall [C37]. The relationship was reciprocal and lasting: energy-hungry industry anchored hydroelectric development, and hydroelectricity anchored the industrial geography of the early electrical age, a pattern that would repeat from the fjords of Norway to the dams of the New Deal.

The locations of those first plants tell the story in miniature. Grand Rapids powered a furniture factory; Ottawa, Dolgeville and Niagara lit mills and a scattering of buildings [C3]. These were manufacturing towns that already knew water power intimately - they had run their mills on it for decades - and electricity arrived as a better way to share a resource they already owned. The wheel under the factory floor became a dynamo beside the wheel, and the line shafting gave way to copper wire. Industrial historians sometimes describe electrification as a revolution that swept aside water power. On the ground, at the mills, it looked more like water power changing clothes.

Underneath the commercial scramble ran a technical problem that the early industry could not dodge. Every station in the 1880s was an island, because the electricity of the 1880s could not travel. Direct current at usable voltages faded within a mile or two of the dynamo; a city could be lit only by generators inside it, and a waterfall could power only what stood beside it. The water was often in the wilderness and the customers in the cities - Niagara’s immense fall sat twenty miles from Buffalo, and the most promising sites of the American West sat in mountains hundreds of miles from anywhere - and between them stood the whole unsolved problem of transmission. Solving it would make hydropower a continental industry rather than a local craft, and the solving is a story in itself.

Chapter 6 - The long wire

The argument that decided hydropower's future was fought over current. Direct current, the kind Edison's stations sold, flows one way and can be generated at whatever voltage the customer needs - but it cannot easily be changed from one voltage to another, and power pushed down a wire at low voltage bleeds away as heat within a few kilometres. Alternating current surges back and forth and is useless for some of the era's machines, but it has one decisive property: a transformer can step it up to very high voltage for the journey and step it back down at the far end. High voltage means low current for the same power, and low current means thin wires and modest losses over long distances. The battle between the two systems, the so-called war of the currents, was really a battle over whether electricity would remain an urban utility generated in the city or become a bulk commodity shipped in from wherever it was cheapest to make. Hydropower had everything riding on the answer.

The proof came in stages, and the first great demonstration was in Germany. In 1891, for the International Electrotechnical Exhibition at Frankfurt, engineers drove a hydroelectric plant at Lauffen on the Neckar river and transmitted its power about 175 kilometres to the exhibition grounds, where it lit lamps and ran a small artificial waterfall - a flourish the exhibitors could not resist [C3][C36]. The Institution of Mechanical Engineers' history of electrification treats the 1891 Frankfurt exhibition as the event that ended the direct-current era [C72]. It is worth pausing on why a trade fair mattered so much: the Lauffen-Frankfurt line did not merely show that high-voltage alternating current worked, it showed the assembled engineers and financiers of Europe that it worked reliably enough to build an industry on. Capital followed conviction, and conviction followed the cable. The United States got its equivalent two years later: in 1893 the Redlands Power Plant in California came online, using the natural flow of Mill Creek to supply communities by alternating current, and the Department of Energy credits it as the first commercial AC hydropower facility in the country [C2]. The technology spread with extraordinary speed once the principle was proven; Australia had its first publicly owned plant in the southern hemisphere by 1895 [C3].

The war of the currents was never only an engineering argument, and it was not settled by elegance. Direct current had the incumbent's advantages: Edison's name, installed systems, real customers. Alternating current had physics, and physics won where distance mattered. It is telling that the decisive contests were fought not in the cities, where either system could serve, but at the great remote power sites, where direct current simply could not play. Hydropower did not merely benefit from the outcome; hydropower was the outcome's strongest

argument. Every kilometre between a waterfall and a city was a vote for the transformer.

The deciding project, the one every textbook still names, was Niagara. The falls were the most famous power site in the world - an enormous, steady flow dropping about fifty metres - and the Niagara Falls Power Company spent the early 1890s deciding how to use it. The commission famously canvassed proposals from across the world before choosing Westinghouse's alternating-current system over the direct-current schemes. In 1895 the Edward Dean Adams Power Plant, the largest hydroelectric development of its time, produced its first power at the falls [C3][C37]. Within a year power was flowing the twenty-odd miles to Buffalo, lighting its streets and turning its factories. It was, as the Institution of Mechanical Engineers puts it, the first large-scale use of alternating current to transmit electricity from a hydroelectric scheme [C36].

The system's virtues compounded once installed. Polyphase alternating current could be generated directly, transformed to whatever voltage the distance required, and stepped down again for customers; and the same current drove the induction motor, the simple rugged machine that would electrify industry itself. The electrochemical companies that clustered at Niagara were only the most visible customers. Behind them came the ordinary load of a modernising economy - factories, streetcars, streetlights, homes - all of it reachable now from a single great powerhouse by the falls. Within a decade of the Adams Plant's first power, long-distance hydroelectric transmission was an established business on both sides of the Atlantic [C36][C37].

What Niagara proved was not merely technical but economic: power could be made in bulk at a great site and sold across a region. The electrochemical industries arrived first, as they always did - the falls soon supported aluminium, carbide and alkali works drawing power a decade into the new century [C37] - but the deeper lesson was about geography. Transmission collapsed the distance between the resource and the market. A dam no longer served a mill beside it; it served cities it would never see, and its economics could be planned at the scale of a state or a nation rather than a township. Every large hydroelectric project since, from Hoover to Three Gorges, has been built on that premise: make the power where the water is, move it to where the people are.

None of this arrived as a single plan. The utilities that strung the early long lines were private ventures chasing industrial customers; the grids they built were patched together town by town, frequency by frequency, decades before anyone spoke of national systems. But the direction was set the moment the Lauffen cable lit the Frankfurt exhibition: electricity would henceforth be made wherever it was cheapest and consumed wherever it was wanted, and the map of power stations

would follow the map of resources rather than the map of cities. For water power, confined by geography to its valleys, that was liberation.

The long wire also quietly changed what a dam was for. A mill dam served its mill; a power dam served a grid, and a grid has appetites of its own - morning peaks, evening peaks, seasonal swings - that no unassisted river matches. Meeting them pushed the industry toward bigger reservoirs and more controllable machines, and it pushed governments toward treating rivers as infrastructure to be planned rather than scenery to be admired or hazards to be endured. There was a further consequence, slower to arrive and larger when it did. Once power could travel, the power business could consolidate. The small station beside its mill gave way to the utility serving a region, and the utility serving a region needed capital on a scale mill owners had never imagined. The industry’s centre of gravity shifted from the machine shop to the bond market, and then - because rivers cross boundaries and electricity became a matter of national strategy - to the state. The stage was set for the era of the giants. But the giants needed more than turbines, wires and capital: they needed the wall itself, raised higher than anyone had raised one before. That is the dam’s own story, and it comes next.

PART IV - THE DAM

Chapter 7 - Holding back the river

A dam is a wager with a simple stake: the structure must weigh, wedge or anchor against the full push of the water it ponds, every hour of every day, for a century or more. The water at the base of Hoover Dam presses with about 45,000 pounds per square foot [C5], and it never stops pressing. There are only a few ways to meet that push, and they define the classic family of dam types.

The gravity dam answers water with weight. Its cross-section is a great triangle of concrete or masonry, broad at the base and narrow at the crest, so heavy that the reservoir cannot overturn or slide it. Grand Coulee on the Columbia, a concrete gravity dam 550 feet high, and the Three Gorges dam on the Yangtze, a gravity structure with a 2,309-metre axis, are the type carried to its extremes [C62][C50]. The gravity dam is honest engineering: no tricks, just mass. Its demands are equally blunt - enormous volumes of material and a foundation that will not yield.

The arch dam answers with geometry. Curved upstream like a bow, it hands the water’s load sideways into the rock of the valley walls. It can be far thinner than a gravity dam of the same height, but it is a snob about sites: it needs a narrow canyon with abutments of sound, strong rock to lean against. Hoover Dam is the great hybrid, an arch-gravity structure that carries the load both by its curve and by its weight [C5]. Kariba, on the Zambezi, is a true double-curvature arch, 128 metres high and 617 metres along the crest, holding back a reservoir of 185 cubic kilometres [C66]. Where the valley is narrow and the rock is good, the arch is the elegant solution; where it is not, no amount of elegance will help.

And where the valley is broad and the local ground is the cheapest building material available, engineers reach for the oldest answer of all: the embankment dam, a heap of earth or rock, armoured against the weather and sealed at its heart with a core of clay or a membrane of concrete. Aswan is the type at world scale - an embankment 111 metres high, 3,830 metres along the crest, containing some 43 million cubic metres of material [C65].

Between the great types sits a fourth, the buttress dam, which answers the water with structure rather than mass or curve: a sloping upstream deck of concrete held up by a row of buttresses, like a leaning wall propped by struts. It uses far less concrete than a gravity dam of the same height, and its openwork construction lets uplift pressure escape. It was a favourite of the early twentieth century, when cement was dear and labour cheap, and many survive in the older hydro districts of North America and Europe.

The choice among them is made by the ground more than by the designer. A narrow granite gorge invites an arch. A wide river valley with deep overburden rules one out. An embankment can be built almost anywhere there is fill to move,

which is why embankments outnumber every other kind of dam on earth. Geology sets the menu; economics picks from it; and the engineer’s reputation rests on reading the ground more carefully than the ground expects.

The types blur at the edges, deliberately. Hoover is called arch-gravity because it is both: thick enough to stand by weight alone, curved enough to lean on its canyon [C5]. Many embankments carry a concrete spillway section through which the floods and the penstocks pass, a hybrid of earth and concrete in one line. Guri, one of the largest power stations on earth, is a composite of gravity sections and earth and rock embankments stretching, all told, 7,426 metres across the Caroni valley [C64]. The designer’s real task is not to pick a type from a catalogue but to read a particular valley - its rock, its width, its materials, its river - and compose the structure that valley will accept.

Beneath the visible structure lies the part the visitor never sees and the designer never stops worrying about: the foundation and the water moving through it. Every dam leaks a little; the question is how much, and with what consequences. Water seeping under a dam lifts it - uplift pressure - and water moving through weak ground can carry the ground away grain by grain, the internal erosion called piping that has undermined more dams than any flood. The countermeasures are drilled and grouted into the rock: a grout curtain injected under the upstream edge to choke the seepage paths, drains to relieve the pressure that gets through. Hoover’s foundation and abutments were sound volcanic andesite breccia, and even there the builders drilled hundreds of thousands of feet of grout holes [C5]. The dam may be the monument, but the unseen work below it is what keeps the monument standing.

The hydroelectric dam adds one more demand to all of this. It is not enough for the wall to hold water; the wall must also be a plumbing system. Intakes must admit the flow to the turbines, control it, and protect the machines from ice, silt and driftwood; outlets must pass irrigation and environmental flows; and somewhere the design must provide for the one event no reservoir can swallow whole, the flood that exceeds the pond. That provision is the spillway, and it deserves its own chapter, together with the stranger problem of building the wall in a river that does not want to stop flowing while you work.

Chapter 8 - Concrete, earth and the building of the wall

The dam-building problem that defeated the nineteenth century was not strength but heat. Cement gives off heat as it cures, and a block of concrete the size of a dam is a furnace that cannot cool: the outside sets while the inside is still cooking, and as the interior finally cools it shrinks and cracks. A cracked gravity dam is a contradiction in terms - the cracks become seepage paths, the seepage becomes uplift, and the mass the design depended on is no longer a mass. For a structure the size of Hoover Dam, engineers calculated, natural cooling would have taken well over a century. The dam could not wait a century, so its builders invented industrial-scale cooling: they placed the concrete in a grid of 215 separate blocks, each pour limited to five feet in seventy-two hours, threaded the blocks with more than 582 miles of one-inch steel pipe, and circulated ice water from a refrigeration plant that made a thousand tons of ice a day. The cooling was finished in March 1935; the contracted grout was then forced into the joints between the shrunken blocks to lock the whole into a single monolith [C5]. Every large concrete dam since has been built on some version of that principle: tame the heat, or the heat will crack the dam.

Half a century later the industry found a second answer, and it came from the roads. Roller-compacted concrete is a very dry mix, stiff enough to be hauled by dump trucks, spread by bulldozers and compacted by vibrating rollers, placed in continuous layers instead of discrete blocks. With less water and less cement it generates less heat; with road-building machinery it goes down at a pace block placement cannot match. The world’s first roller-compacted concrete dam was built in Japan in 1980, and the method spread so fast that within a decade dozens stood worldwide, including the US Bureau of Reclamation’s first, at Upper Stillwater in Utah [C68][C69]. RCC did not make the classic block dam obsolete - the technique suits gravity sections better than thin arches - but it moved the economics of concrete dams closer to those of earthwork, and much of the late twentieth-century dam boom was built with it.

The embankment dam solves the heat problem by ignoring concrete altogether, and solves the economics by building the dam out of the valley itself. Its engineering is the engineering of soil: a watertight core of compacted clay, flanked by shoulders of sand, gravel and rock that give the heap its strength, armoured on the water face against waves and drained on the land face so seepage cannot turn the downstream slope to soup. Aswan’s 43 million cubic metres of earth and rock made it one of the largest structures humanity has ever piled up [C65]. The embankment’s vulnerability is the mirror of its virtue: it is a heap, and a heap overtopped is a heap destroyed. Earth cannot have water flowing over it.

An embankment dam lives or dies by its spillway - which is the next chapter's subject.

Materials were only half the building problem; the other half was that the river objected. To build a dam across a flowing river, the river must first be persuaded to go elsewhere. The standard persuasion is the diversion: tunnels blasted through the abutment rock or a channel cut around the site, into which the whole river is turned while the original bed is fenced off with temporary dams of earth and rock - cofferdams - and pumped dry [C71]. At Hoover, the diversion alone was an engineering campaign: the Bureau of Reclamation's account records a million and a half cubic yards excavated for the diversion tunnels, and behind the tunnels the main excavation removed another 1.76 million cubic yards for the dam, powerplant and cofferdams, with 750,000 more for the spillways and inclined tunnels [C5]. Before any of that could begin, the government had to build the infrastructure to build the dam: Boulder City to house the workers, seven miles of highway, twenty-two and two-tenths miles of standard-gauge railroad from the Union Pacific main line at Las Vegas, and a 222-mile transmission line from San Bernardino simply to supply construction power [C5]. A great dam does not begin with concrete. It begins with a town, a railway and a power line, and only then with a river persuaded to step aside.

At the peak of the Hoover campaign, 5,218 men were at work in a single month, June 1934, placing concrete at a rate that peaked above 275,000 cubic yards a month [C5]. The diversion is one of the most dangerous phases of any project, because the temporary works must survive whatever floods arrive before the permanent ones exist, and dam-safety literature treats it as a discipline of its own, with failures that have drowned worksites and set projects back years [C71]. When the diversion closes and the river begins, at last, to climb the new wall, the project crosses its point of no return. Impoundment is the moment the drawings become a lake.

The two material traditions - concrete and earth - also aged differently. Mass concrete, once cured, is nearly immortal: Hoover has stood for ninety years and its engineers speak of its life in geological terms. Embankments are more mortal, vulnerable to the slow mischief of seepage, burrowing animals and roots, and to the single fast mischief of overtopping; they are maintained rather than merely endured. Both traditions learned, expensively, that the river keeps its own accounts. Silt settles against every dam, subtracting storage year by year; floods arrive from outside the design's imagination; earthquakes test assumptions made by men long retired. The later chapter on failure takes up those reckonings. The point here is simpler: building the wall is the beginning of the dam's life, not the end of its cost.

What unites these methods is that a dam is not so much designed as negotiated - with the rock, the river, the weather, the available materials and the calendar. The great dams of the twentieth century were less objects than campaigns: the 21,000 men who worked on Hoover, the railroad and the 222-mile transmission line built just to supply its construction, the entire town of Boulder City raised from the desert to house them [C5]. The industry learned to build big because it learned to organise big, and the scale of the organisation, as much as any advance in concrete, is what made the age of the giants possible.

Chapter 9 - Water past the wall

Every reservoir eventually meets a flood it cannot store. Rain falls on the catchment, the river rises, and the pond fills toward the crest - and a dam with water flowing uncontrolled over its top is a dam in the process of failing, because the overflow scours the downstream face and foundation until the structure goes. The spillway exists to prevent exactly that: it is the dam’s safety valve, sized and shaped to pass the floods the reservoir cannot hold, and the engineering manuals state its purpose without romance - to pass moderate floods and to keep the dam alive through the very largest ones [C70].

The mechanics of the spillway vary with the dam. Concrete gravity and arch dams can let water pass over the structure itself, down a shaped ogee crest that keeps the flow pressed smoothly against the concrete; embankments cannot, so their spillways run around the dam, in channels cut into the abutment or in vertical morning-glory shafts that swallow the overflow like a drain and carry it through a tunnel. Whatever the form, the water arrives at the bottom carrying the energy of the reservoir’s full head, and that energy must be destroyed safely - in a stilling basin of churning water, a ski-jump bucket that throws the jet into mid-air to shatter, or a plunge pool the river itself is allowed to dig. Kariba’s spillway discharges have carved a plunge pool so deep that repairing its undercutting became, decades on, a rehabilitation project in its own right [C29]. A spillway is a controlled waterfall, and control is the whole art.

Designing that valve means deciding which flood to design for, and the decision is probabilistic, expensive and occasionally fatal. Hydrologists reconstruct the catchment’s flood history and extrapolate toward the worst event the river can physically produce - the probable maximum flood. A dam that cannot pass its design flood must be raised, enlarged or given more spillway; a spillway built for an exaggerated flood is concrete poured for nothing. The calculation is made harder by a trap the twentieth century learned the hard way: the reservoir itself changes the flood, smoothing and delaying it, so the spillway must be designed not for the river as it was but for the reservoir-river system as it will behave. At Banqiao in 1975, the storm that arrived exceeded anything the design had imagined - 1,631 millimetres of rain in five days - and the consequences of that shortfall belong to the book’s chapter on failure [C39].

Spillways take two broad forms. The service spillway is a gated channel - a weir with moveable gates the operators open as the flood demands. The emergency spillway is the last resort, an ungated overflow that runs only when the reservoir rises beyond everything else. The 2017 crisis at Oroville Dam in California showed what happens when both are tested at once. Heavy rains damaged the main spillway’s concrete chute so badly that operators had to throttle it back; the

reservoir rose until water went over the emergency weir for the first time in the dam's history; the hillside beneath it began to erode toward the weir's foundations; and the sheriff ordered the evacuation of about 188,000 people downstream [C74]. No one died and the dam held, but the independent forensic team's verdict was sobering: there was no single root cause. Physical weaknesses in the spillway combined with institutional failures - inspections that looked and did not see, an emergency spillway long assumed adequate on paper rather than proven [C73]. Oroville joined the short list of incidents every dam engineer studies, precisely because nothing in it was exotic.

The spillway is only the most dramatic of the openings through the wall. A hydroelectric dam is perforated by plumbing: intakes that draw water for the turbines, bottom outlets that can draw the reservoir down, sluices for silt, and, on multipurpose rivers, releases for irrigation, navigation and fish. Each opening is a controlled leak, and each control is a machine in itself - the great gates and valves that are among the heaviest moving parts in any power station. Hoover Dam carried some 21,670,000 pounds of gates and valves [C5]. They must work after years of immersion, under full reservoir pressure, in the dark, on demand - because the day a gate fails to open is the day the spillway design gets tested.

The intake side of the plumbing has its own subtleties. Water for the turbines should be drawn clean and cool, without the silt that grinds runners or the debris that jams them; intake towers stand in the reservoir like chess rooks, drawing from a chosen depth. And between the intake and the turbine runs the penstock, the pipe or tunnel that carries the water down the full height of the head. The penstock converts the reservoir's still pressure into directed speed, and it must survive the reverse as well: when the turbine gates slam shut, the moving column of water has nowhere to go, and the resulting hammer blow of pressure - water hammer - can split a penstock that was not designed for it. Surge tanks and air valves, the shock absorbers of the system, stand guard against that blow.

There is, finally, the question of what the water carries with it. Rivers move sediment, and a reservoir is a sediment trap: the flow slows at the back of the pond and drops its load. Managing that load - passing silt through low-level sluices, flushing reservoirs, drawing intakes from chosen depths - is one of the least visible arts of dam operation, and one of the most consequential, because every tonne of sediment that settles in the reservoir is a tonne of storage lost and a tonne denied to the river downstream. The Aswan High Dam stopped the Nile's silt almost completely, and the downstream consequences - starving the delta that the silt had built - belong to the chapter on costs. The wall holds back water; it also holds back everything the water was carrying.

The last openings in the system are the ones that serve purposes other than power, and they have grown steadily more important. Navigation locks, where the river

carries traffic, lift and lower ships around the dam; the Three Gorges project moves ten-thousand-tonne fleets through a double-way five-step ship lock and raises smaller vessels in a vertical ship lift [C50]. Fish ladders and lifts offer migratory fish a way around the wall - imperfectly, as a later chapter admits. And environmental releases attempt to give the downstream river something resembling a pulse of life. The modern dam is judged by all of its openings, not just the ones with turbines behind them: it is a gatekeeper for the whole river, and every purpose the river serves must somehow be passed through, over, or around the wall.

PART V - THE MACHINE HALL

Chapter 10 - The modern turbine

Walk into the powerhouse of a large hydroelectric dam and the first surprise is how little there is to see. The dam is spectacle; the machine hall is a long, clean, nearly empty room with a row of painted cylinders set into the floor, each humming slightly. Somewhere beneath each cylinder, invisible, a shaft of water falling from the reservoir is being turned into rotating shaft power with an efficiency no other prime mover matches. The machines that do it are the descendants of the four nineteenth-century turbines, brought to a pitch of refinement their inventors would recognise at once and marvel at for weeks.

The refinement begins with selection, because the first law of turbine design is that the site chooses the machine. Where the head is high and the flow modest - a mountain reservoir, a long penstock - the Pelton wheel still rules: an impulse machine, its runner spun by jets fired at its buckets from nozzles, operating in air at atmospheric pressure. Where the head is middling and the flow generous - the classic dam site - the Francis turbine still dominates, water spiralling inward through its guide vanes and out along the shaft. And where the head is low and the flow enormous - the great flat rivers - the Kaplan’s adjustable propeller still answers, pitching its blades to stay efficient as the river swells and shrinks [C4]. Engineers formalise the choice with specific speed, the single number that sorts any site toward its proper family, but the result is always the same three answers. The turbine families have proved as durable as the dam types, and for the same reason: each is matched to a permanent feature of the physical world.

Inside the families, a century of computation and model testing has chased the last points of efficiency. Cavitation is the turbine’s chronic disease: where pressure falls too low in the fast water, the water itself boils into vapour bubbles that collapse with microscopic violence against the steel, pitting it like rust. Designers shape blades to avoid it, operators learn to keep machines away from the zones where it bites, and maintenance crews weld and re-polish the scars it leaves. It is the tax the water levies for being hurried.

Every large runner begins life as a scale model in a hydraulic laboratory, where its curves are tested and tuned at small scale before a single tonne of steel is cast. The model tests predict not just peak efficiency but the shape of the whole efficiency hill - how the machine behaves at part load, where cavitation begins, how it pulls at start-up. A power station’s operators will live on that hill for half a century, choosing where to run; the laboratory’s job is to make the hill broad and its summit high.

Scale is the other axis of refinement, and it has run in both directions at once. The largest turbines ever built - Baihetan’s sixteen one-thousand-megawatt units, and

the 700-megawatt-class machines of Three Gorges and Itaipu - are buildings in themselves, each runner heavier than a loaded freight train, machined to fractions of a millimetre [C50][C63][C62c]. At the other extreme, packaged turbines of a few kilowatts bring the same principles to a village stream. Between them lies a manufacturing industry of a handful of great firms and many specialists, and a continuous trade in rehabilitation: turbines outlast their runners, generators outlast their windings, and one of the quietest businesses in hydropower is taking a machine built by grandfathers and returning it to the grid more powerful than it left.

The numbers behind that dominance are worth stating plainly. A Pelton wheel takes the whole head of the site in a single jet and converts it to velocity; a Francis runner takes pressure and flow together; a Kaplan trades pressure for volume. But all three, in their proper ranges, deliver the same result: conversion efficiencies that reach well above ninety per cent, and at the best points of the efficiency hill, higher still [C35]. The losses that remain - friction in the waterways, turbulence off the blades, resistance in the generator windings and the transformer - are measured in single percentage points, each one fought for, because over a fifty-year life each point is a river of energy sold or spilled.

The control of these machines is a discipline of its own. Each turbine is governed - its guide vanes opening and closing second by second to hold the shaft at exactly synchronous speed while the grid’s demand fluctuates around it. The governor is the turbine’s reflex system, and on an interconnected grid thousands of such reflexes act together to hold frequency steady. Hydropower’s machines excel at this work: they can sweep from near standstill to full output in a time measured in tens of seconds, a responsiveness no thermal plant can match. Dinorwig, the Welsh pumped storage station built inside a hollowed-out slate mountain, can bring 1,728 megawatts to full output in about twelve seconds [C60]. That agility, once a convenience, has become one of the grid’s most valued commodities, as the chapter on the grid’s workhorse explains.

The machine hall has one more quiet inhabitant worth meeting: the reversible machine. A pump-turbine is a Francis-type runner designed to work in both directions - spinning one way to generate as water falls, the other to pump water back uphill when the grid has energy to spare. It is the heart of pumped storage, the water battery that a later chapter follows, and it completes the turbine family’s logic: the same machine that answered every question of head and flow was finally taught to answer the question of time.

For all the refinement, the modern turbine’s deepest quality is the one Fourneyron demonstrated in 1827: it turns falling water into rotating shaft power almost without waste. Modern hydroelectric machines convert more than nine-tenths of the energy they receive [C35]. No steam turbine, no gas turbine, no engine of any

kind approaches that figure, because every combustion machine throws away most of its fuel as heat and the water turbine is given its energy in mechanical form to begin with. The consequence runs through the whole economics of the industry: once the dam is built and paid for, the machine at its foot wastes almost nothing and asks for almost nothing. The expensive part of hydropower is never the machine. It is the wall - and the waiting.

Chapter 11 - From torque to current

The turbine’s output is a spinning shaft, and the generator’s job is to make that rotation into electricity. The principle is Faraday’s: move a conductor through a magnetic field and current flows. The scale is industrial. In a large hydroelectric generator, a rotor carrying electromagnets spins inside a stator ringed with copper windings, and the excitation system - the small, precise supply that feeds the rotor’s field - decides how much voltage the machine makes and how it behaves under load. It is a beautifully direct chain: water turns the runner, the runner turns the shaft, the shaft turns the rotor, and the field sweeps current out of the stator. There is no fire anywhere in it.

The machines grew up with the systems they serve. The first Niagara generators of 1895 were two-phase machines, state of the art for a technology barely out of the laboratory; within a decade the falls’ powerhouse had multiplied its units and the industry had settled on the three-phase standard that rules transmission to this day [C37]. Hoover Dam’s two powerhouses carry seventeen main Francis units plus two small Pelton machines for station service, a nameplate of about 2,080 megawatts in all, and from 1939 to 1949 the complex stood as the largest hydroelectric installation in the world [C61]. Each generation of plants scaled the same design upward: bigger rotors, stronger windings, higher voltages, tighter tolerances.

That directness gives hydroelectric generation a character no thermal plant can imitate. A steam turbine must be nursed to temperature over hours; a hydroelectric machine can be brought from standstill to full load in minutes, and its output can be swung up and down almost as fast as the wicket gates can move. The governor, the machine’s reflex system, watches the shaft speed second by second and trims the gates to hold it exactly synchronous with the grid. Every generator on an interconnected system spins in lockstep with every other, and the system’s frequency - sixty cycles a second in North America, fifty across most of the rest of the world - is the heartbeat they all share. Hydropower’s speed and obedience made it the grid’s natural regulator long before anyone built a battery or a smart inverter.

The generators earn their keep in a second, less visible way as well. An alternating-current grid needs not just energy but voltage support - the reactive push that keeps the system stiff and the voltage steady - and a synchronous generator supplies it as naturally as it spins, by adjusting its excitation. Hydropower’s machines, big synchronous units that can be dispatched at will, have always been among the grid’s chief suppliers of that stiffness. Some hydro units are even run, from time to time, as synchronous condensers: spinning

without water, making no power at all, purely to steady the system. It is a strange service to sell - rotation without output - and grids pay for it gladly.

Between the generator terminals and the grid stands the transformer, and its role is the one the long wire requires: to step the machine's voltage up to transmission levels, hundreds of kilovolts on the largest systems, so the power can cross a landscape without bleeding away. Guri sends its power out at 765 kilovolts, among the highest transmission voltages ever used, because its energy travels hundreds of kilometres to Venezuela's cities and industries [C64]. The switchyard - the open-air geometry of busbars, breakers and lightning arresters beside every powerhouse - is where the plant meets the network, and where protective systems stand ready to disconnect a faulted machine in a fraction of a second before the fault can spread.

The electrical side of a hydroelectric plant is, in short, the obedient half of the partnership. The civil works are massive, slow and fixed; the electrical works are fast, precise and controllable. The dam decides how much energy there is to sell; the machines decide, instant by instant, how much of it to sell now. When a grid operator somewhere far away watches demand surge and dispatches power into the evening peak, it is this obedient half that answers - gates opening, excitation strengthening, megawatts rising in the time it takes a thermal plant to clear its throat.

The governor that controls all this began, like everything else here, as a nineteenth-century mechanism: whirling weights on a spindle, rising and falling with the shaft's speed, linked to the gates. The modern governor is a digital controller, but its task is unchanged and its reflexes have only sharpened: sense the grid's frequency, open or close the gates, hold the machine in step. On a large interconnected system, thousands of such reflexes act in concert within seconds of any disturbance - a tripped line, a lost generator, a sudden surge of demand. The grid's stability is, in large part, the sum of its machines' reflexes, and no fleet of machines has faster or more willing reflexes than the hydro fleet.

The result of all this - storage above, machines below, wires beyond - is a power plant with a split personality that grids have learned to treasure. It is at once an energy producer, churning out some of the cheapest kilowatt-hours ever sold, and a system servant, holding itself ready to regulate, to balance, to restart. A thermal plant sells heat converted to motion; a hydroelectric plant sells falling water converted to obedience. The distinction, made in the machine hall, becomes the central economic fact of the book's later chapters, when the grid fills with wind and sunshine and goes shopping for exactly the obedience hydro has always sold.

There is one more service the machine hall performs, so valuable that it has no substitute. When a grid collapses - blacked out region-wide by cascading failure -

most power stations cannot restart, because a thermal plant needs electricity to make electricity: pumps, fans, controls. A hydroelectric unit can start with almost nothing. Open the gates, spin the machine, energise a dead line, and begin nursing the grid back section by section. The capability is called black start, and grid planners treat hydroelectric stations that possess it as the emergency services of the power system. Itaipu, Three Gorges, the great storage dams of every continent - each is, among its other roles, the match with which a darkened grid can be relit. The wall in the valley turns out to be insurance as well as industry.

Chapter 12 - The anatomy of a powerhouse

Follow the water through a hydroelectric dam and you trace the whole anatomy of the business. It begins in the reservoir, still and heavy with stored height. The intake draws it in - clean, cool, screened against driftwood and ice - and commits it to the penstock, the pipe or tunnel that carries it down the full drop of the site. In the penstock the water trades altitude for pressure: every metre of fall adds its increment, and by the bottom of the drop the pipe holds a column straining to move.

At the powerhouse the penstock divides, and each branch delivers its stream to a machine. In a Francis unit the water enters a spiral casing - the scroll case, curled like a nautilus shell around the runner - which distributes the flow evenly to the ring of adjustable guide vanes, the wicket gates. The gates angle the water onto the runner blades; the runner spins; the water, its work done, exits downward through the draft tube, a flared passage that slows the outflow and recovers a last increment of energy from the leaving water. Then the tailrace returns it to the river, cooler and calmer, to continue downstream as if nothing had happened - except that the height it fell through now moves down a wire as electricity.

The intake works deserve a closer look, because they are where the dam becomes a machine. Hoover’s reservoir feeds four intake towers standing in the lake upstream - great cylinders of concrete, built in the dry of the diversion and now standing in deep water, drawing the reservoir’s flow through gates that can be closed against an emergency [C5]. From the towers the water enters tunnels, and from the tunnels the steel penstocks that descend to the powerhouse: at Hoover the system moves the reservoir’s entire working flow from lake to runner under continuous control. It is plumbing on the scale of architecture, and like the gates and valves that work it - more than twenty million pounds of them at Hoover [C5] - it must function for a century in conditions no plumber would tolerate.

The powerhouse itself comes in as many forms as the sites. At Hoover, it sits at the toe of the dam in a great gallery at the canyon’s base, the classic image of the type. At run-of-river plants it is a building beside a weir. At the most dramatic sites it disappears entirely: the powerhouse at Dinorwig in Wales was hollowed out inside the mountain Elidir Fawr, in the caverns of a worked-out slate quarry, with sixteen kilometres of tunnels threading the hill [C60]. Underground powerhouses leave the landscape above them almost untouched, and on rivers where the valley walls are steep or the scenery defended, the industry learned to hide its machine rooms in the rock.

The draft tube earns a paragraph of its own, because it is the least understood element of the chain. A runner’s exit water still carries speed, and speed is energy

the machine has failed to capture. The draft tube - a passage that widens like a trumpet as it descends to the tailrace - slows that exit flow and converts its leftover velocity back into suction beneath the runner, recovering energy that a naked discharge would waste. The detail is typical of the whole design tradition: a century of refinement has been spent hunting the last few per cent of the falling water’s gift, in scroll cases and blade profiles and draft tubes, because the water is free and every fraction of it recovered is pure profit.

Hidden or not, every powerhouse contends with the same physics, and the most dangerous of its problems is the water’s own momentum. A penstock holds a moving column of water that may be kilometres long; when the wicket gates close, that column cannot stop at once, and the pressure surge it hurls back up the pipe - water hammer - can rupture steel. The defences are the surge tank, an open standpipe that lets the surge rise harmlessly as a column of water, and the careful choreography of the governors, which close the gates no faster than the plumbing can bear. The same physics in reverse greets the gates’ reopening: a water column takes time to get moving, which is why even the nimblest hydro plant ramps over seconds and minutes rather than instantly.

The newest generation of powerhouses continues the retreat underground. The great pumped storage stations are nearly all cavern plants now - machine halls excavated in sound rock between an upper and a lower reservoir, reached by tunnels and shafts rather than roads along a riverbank. Fengning in China, the largest pumped storage station on earth at 3,600 megawatts, packs its twelve units into such a complex in the mountains of Hebei [C59]. The invisibility is partly environmental manners and partly economics: a cavern plant can be built where the head is best, regardless of what stands above it.

A powerhouse is also, finally, a workshop. The machines inside it will run for half a century or more, through floods and droughts, and they are built to be taken apart: overhead cranes the size of the building itself, erection bays where a runner can be laid down and rebuilt, galleries and tunnels for inspection. Much of the world’s hydro fleet is now old enough that this maintainability is being exercised at full scale - units rewound, runners replaced, governors modernised - and the renewal of an aging powerhouse has become one of the industry’s steadiest businesses. The dam is built once; the machine hall is rebuilt continuously.

It is worth standing back from the machinery to see the whole arrangement as one device, because that is what it is. The reservoir stores energy; the intake selects it; the penstock concentrates it; the scroll case aims it; the runner catches it; the draft tube salvages the last of it; the generator converts it; the transformer launches it. A hydroelectric station is a single continuous transformation from still water at height to alternating current on a wire, and every element exists to serve that one conversion. The dam gets the photographs and the controversies. The powerhouse

gets the megawatts. Between them they are the subject of this book - and the twentieth century, which built them by the hundred, is where the story goes next.

PART VI - THE GIANTS

Chapter 13 - Hoover and the federal river

The hydroelectric dam became an instrument of national policy on the Colorado. The river was everything the young industry could want and everything it feared: a big flow draining the Rocky Mountains through desert canyons, wild with silt, savage in flood, and coveted by seven states and two countries. Through the 1920s the Southwest argued over who owned its water while farmers in the Imperial Valley, irrigating at the river's whim, watched floods take their intakes and droughts take their crops. The Boulder Canyon Project was Washington's answer: one great dam that would store the river, end the floods, guarantee the irrigation supply and - third on the list of purposes - sell power to pay for it all [C2]. Congress authorised it in 1928, and the Bureau of Reclamation, created in 1902 to bring water to the arid West, got the largest job it had ever attempted [C2].

The site was Black Canyon, where the Colorado cuts between Nevada and Arizona through walls of hard volcanic rock - andesite breccia, sound enough to take the thrust of an arch [C5]. The design that rose there between 1931 and 1936 was an arch-gravity dam 726 feet from foundation to crest, the highest dam yet attempted, holding back a reservoir that stretched more than a hundred miles. The numbers still startle: three and a quarter million cubic yards of concrete in the dam itself, placed at a peak rate above 275,000 cubic yards a month; 21,000 men employed over the job, 5,218 of them at work in the peak month; a whole town, Boulder City, built to house them, and a railroad and a 222-mile transmission line built to supply the site before the dam could supply anyone [C5]. The contractors were allowed seven years; the concrete was finished in five [C5].

The construction itself was a study in logistics under pressure. Before the first concrete could be placed, the river had to be moved: four great diversion tunnels driven through the canyon walls, carrying the whole Colorado around the site while the bed was pumped dry behind cofferdams [C5]. The excavation that followed took out five and a half million cubic yards of material - for the tunnels, the foundations, the spillways, the intake towers - and the concrete came down at an average of 160,000 cubic yards a month for two years [C5]. The dam was built as a grid of 215 interlocking blocks, each pour limited to five feet in seventy-two hours so the heat of curing could be drawn off through 582 miles of embedded pipe by a refrigeration plant making a thousand tons of ice a day [C5]. Nothing on that scale had been done before, and the methods invented to do it - block construction, artificial cooling, high-speed placement - became the standard grammar of every large concrete dam built since.

The engineering has already had its chapters - the diversion tunnels, the 215 cooled blocks, the 582 miles of ice-water pipe - but the deeper significance of

Hoover was institutional. It proved that a government could plan an entire river as a single system: storage here, irrigation there, power everywhere, flood control over all. The power was the revenue engine that made the system payable, sold by long-term contracts to the cities of Southern California and the states of the basin. When the powerhouse was completed, its seventeen main units and 2,080 megawatts made it the largest hydroelectric installation on earth, a title it held from 1939 to 1949 [C61]. The dam that was authorised for irrigation and flood control had become, in effect, a power station with a water project attached - exactly the inversion the industry's economics kept producing.

The water politics mattered as much as the wattage. The Colorado Compact of 1922 had divided the river's flow between the upper and lower basin states before anyone knew what the flow really was; the dam made the division deliverable, storing the wet years against the dry. Hoover's reservoir, Lake Mead, became the bank account of the Southwest - and, a century later, as the basin's demands outgrew the river's real supply, the visible ledger of its overdraft, its white bathtub ring of mineral deposits marking the waterline's long retreat. The dam worked exactly as designed; the river turned out to be smaller than the promise. That tension, between the permanent concrete and the variable water, runs through the book's later chapters on climate and limits.

Hoover also fixed the image. The clean curve of the arch between the canyon walls, the white art-deco towers, the spillway in flood: it became the most photographed engineering work in America and the template for what a dam meant - modernity, mastery, abundance. Politicians across the world learned the lesson that a great dam is a monument that generates electricity, and the New Deal took the lesson to scale: Hoover was followed by Grand Coulee on the Columbia, by Bonneville, by the Tennessee Valley Authority's whole family of dams, until by 1940 hydroelectric plants were producing a full 40 per cent of America's electricity [C2]. The era of the giants had begun, and it was unmistakably a public enterprise. Private capital had built the industry's childhood; the state built its prime.

The powerhouse, though, kept quietly paying. Hoover's electricity was sold under fifty-year contracts to the cities and utilities of the basin, and the revenue retired the dam's cost with interest while the irrigation water it guaranteed turned desert valleys into some of the most productive farmland on the continent. It is the purest expression of the bargain at the heart of this book: the river's fall, captured once by a wall, pays for the river's management forever after. When later chapters ask why governments kept building these vast, slow, capital-hungry structures, Hoover is the standing answer - because, built right and sited well, the machine works.

That prime came at prices the public ledgers did not all record. The Colorado’s water was divided by a compact written in a wet decade, promising more water than the river usually carries - a miscalculation the basin still struggles with. Lake Mead, the largest reservoir of its time, drowned canyons and rapids and set the pattern for a West plumbed behind dams. And the model itself - the great federal dam as development’s engine - was exported wholesale after the war, to the Dnieper and the Nile, the Zambezi and the Parana, carrying both its achievements and its assumptions. The next chapters follow it.

Chapter 14 - The New Deal valleys

If Hoover proved the federal dam, the rest of the New Deal made it a system. The 1930s in the United States were the great convergence: a depression demanding public works, rivers demanding control, an electricity industry demanding customers, and a federal government newly willing to build all three at once. Work on Grand Coulee Dam began in the summer of 1933, the year Franklin Roosevelt took office promising a New Deal, and the Columbia Basin Project it anchored was conceived as nothing less than a planned landscape - power, irrigation and navigation engineered together for an entire region [C32].

Grand Coulee outgrew even Hoover. A concrete gravity dam 550 feet high on the Columbia River in Washington State, it became, with its later third powerplant, a 6,809-megawatt complex - the largest power station of any kind in the United States, producing more than 21 billion kilowatt-hours a year and serving eight western states and Canada [C62]. Its reservoir, Franklin D. Roosevelt Lake, covers 82,300 acres [C62]. But Grand Coulee's purpose was always double: the dam's pumps lift Columbia water into the Grand Coulee itself, the ancient dry channel that gives the dam its name, creating the storage for the Columbia Basin's irrigated farms. Power paid for water; water justified power. The formula Hoover invented, Grand Coulee perfected.

The human argument for these projects was as explicit as the engineering. The Columbia basin was semi-arid land waiting for water; the Tennessee Valley was a region of eroded farms, malaria and out-migration; the whole country was in depression, and a dam was a payroll as much as a structure. The New Dealers said so plainly: public power would be a yardstick to measure - and discipline - the private utilities' rates, and public works would put men back to work. The projects delivered both in measure: the Columbia Basin became one of the great irrigated agricultural districts of the continent, and the Pacific Northwest's cheap power underwrote its wartime and postwar industry.

The Tennessee Valley Authority took the idea furthest. Created by act of Congress on 18 May 1933, the TVA was charged with the whole life of a river basin: to control its floods, open it to navigation and generate cheap electric power [C31]. It was regional planning as an act of government - a federal corporation owning dams, power plants, fertiliser works and demonstration farms across a valley that had been one of America's poorest regions. Today the TVA operates twenty-nine hydroelectric dams and a pumped storage plant across the Tennessee River system [C2]. Whatever the debates over its economics and its politics, the TVA fixed a second template beside Hoover's: not the single monumental dam but the integrated river-basin authority, a model that would be copied from the Snowy Mountains to the Mekong.

The Bonneville project completed the Pacific Northwest's transformation. In 1937 Roosevelt signed the Bonneville Project Act, creating the Bonneville Power Administration to market the Columbia's power, and the Bonneville Lock and Dam began operation the same year [C2]. The BPA's cheap power went on to rewire the region's economy - and, during the Second World War, to power the aluminium plants and shipyards of the Pacific war effort. The pattern the industry association describes for the whole era - low-cost hydropower tied to energy-intensive industry like aluminium smelters - was nowhere truer than on the Columbia [C3].

The government's entry also settled, in America's case, a question the industry had dodged for decades: whose river is it? The Federal Power Commission's authority was extended in 1935 to all hydroelectric projects, public or private, engaged in interstate commerce [C2] - licensing, in effect, the nation's rivers as a regulated commons. The Bureau of Reclamation became the second-largest producer of hydropower in the nation, the Army Corps of Engineers, authorised to build hydroelectric plants in the 1920s, became the fifth-largest electric supplier, and the TVA ran its valley as a fief [C2]. The American system that emerged was a hybrid nobody had designed: private utilities under federal licence, public power under regional authorities, and the great rivers managed as national assets. Versions of that hybrid, with the roles differently assigned, spread everywhere the dam-building age reached.

The aggregate effect was a transformation of national scale. In 1920 hydropower had been a minority source; by 1940, the New Deal construction programmes had more than tripled hydroelectric capacity, and water power supplied 40 per cent of all electricity generated in the United States [C2]. The programmes were employment projects as much as power projects - the Fort Peck dam in Montana alone put 10,500 people to work [C2] - and they were built with an explicit social argument: that cheap public power would lift whole regions out of poverty. That argument, and the dams that embodied it, became one of the most influential exports of the American century. After 1945, the world went dam-building, and it built to the American plan.

There was a darker register to the boom as well, though the era's ledgers mostly omitted it. The reservoirs that rose behind the New Deal dams drowned farms, towns and burial grounds, and the people displaced - poor farmers in the Tennessee Valley, Native nations along the Columbia whose fishing grounds the dams destroyed - paid costs the project's accounting treated as external. The salmon runs of the Columbia, among the richest on earth, began their long decline as the dams rose. These costs were not hidden so much as uncounted; the reckoning with them, decades later, became the environmental movement's first

great fight over rivers, and it reshapes the industry to this day. The book returns to it in the chapters on the costs downstream.

It is worth registering what the New Deal dams changed in the industry's soul. Before them, hydropower was a business with a technology; after them, it was infrastructure with a politics. The questions that dominate the rest of this book - who decides, who pays, who benefits, who bears the cost - all arrived with the federal dams, and they have never left.

Chapter 15 - The planned giants

The dam-building age spread outward from the United States on the current of the century's two great projects of state: communist industrialisation and post-colonial development. Both took the American template - the state-planned river, the multipurpose dam, the power-fed industrial boom - and built it larger than America ever had.

The Soviet Union went first and fastest. The Dnieper Hydroelectric Station, DneproGES, rose between 1927 and 1932 at the rapids below Zaporizhzhia, the flagship of the first Five-Year Plan and, at its opening in October 1932, the largest hydroelectric station in Europe [C30]. American engineers and equipment helped build it, a detail the Soviets both used and resented; the political point was the opposite of dependence. The dam was proof, poured in concrete, that socialism could out-build the capitalists at their own most modern game - and its power fed the metallurgical complex the plan demanded. DneproGES was destroyed in the Second World War, blown up to deny it to the advancing Germans, and rebuilt afterward; it remains the anchor of the Dnieper cascade [C30]. The pattern it set - the river planned as a staircase of dams, each step serving power, navigation and irrigation - became standard practice on every big river the twentieth century organised.

Egypt took the model to the Nile. The British had built a first Aswan dam in 1902 to tame the flood for cotton; after the revolution of 1952, Nasser's government conceived something far larger - a High Dam that would end the Nile's ancient alternation of flood and drought altogether, store a year's water and more, and electrify the delta [C47]. The geopolitics were as large as the engineering: the United States and Britain offered financing and withdrew it in 1956 as the Cold War's price lists shifted, and the Soviet Union stepped in with money, equipment and engineers. The dam rose between 1960 and 1970: an embankment 111 metres high and 3,830 metres long, holding back Lake Nasser, a reservoir of some 162 to 169 cubic kilometres that backs the Nile up 500 kilometres into Sudan [C65] [C43]. Its 2,100 megawatts supplied about half of Egypt's electricity when it opened [C65]. The Nile, the most managed river of antiquity, became the most managed river of the modern world - flood-free, irrigated year-round, and, as the chapters on consequences will show, transformed in ways its builders neither intended nor admitted.

The Soviet cascade-builders took their own model still further after the war, damming the Volga and the great Siberian rivers at a scale to match the ideology. The Sayano-Shushenskaya dam on the Yenisei, completed decades later at 6,400 megawatts, became Russia's largest power station - and the site, in 2009, of the industry's worst machine-hall disaster, when a turbine tore loose and flooded the

powerhouse, killing seventy-five people [C62b]. That accident belongs to the chapter on failure, but its lesson belongs here too: the giants the mid-century built are now aging fleets, and their maintenance is a burden their builders bequeathed.

Africa's great post-colonial dam was Kariba. Built between 1956 and 1959 across the Zambezi in the Kariba Gorge, on the border of what are now Zambia and Zimbabwe, it was a double-curvature concrete arch 128 metres high, and its reservoir - 185 cubic kilometres at maximum retention - is commonly cited as the largest man-made lake by volume on earth [C66][C29]. Kariba's power stations, on both banks of the river, total 2,130 megawatts and still supply much of both countries' electricity [C66]. It was built by a federal government that no longer exists, for a copper belt that boomed on its power, across a valley whose people - the Tonga - were removed by the tens of thousands to make way for the rising lake. The dam's story contains, in one project, the whole mixed legacy of the mid-century giants: genuine development, genuine power, and costs settled on people who were never consulted.

What made these projects irresistible to young states was exactly what had made them irresistible to Washington: they were development in visible form. A dam was electrification, irrigation, flood control, national prestige and a payroll, all rising from one valley. The financing, though, came with flags attached - Soviet engineers at Aswan and on the Volga, American and European contractors and World Bank loans across Latin America and Asia - and the dams of the third world became instruments of the Cold War's competition as much as of their rivers' management. The reckoning over who paid, and who profited, would come later.

The same template reached Ghana at Akosombo, Brazil's southeast, India's river valleys and dozens of smaller basins, and it carried the same structure everywhere: a state, a river, a dam, an industrial promise. The industry association's summary of the era is exact: from the 1940s to the 1970s, state-owned utilities built great hydroelectric developments across Western Europe, the Soviet Union, North America and Japan, and low-cost hydropower was tied deliberately to energy-intensive industries - aluminium above all [C3]. The dam had become the engine room of the development state. The builders, it should be said, were not blind to all of it. Egypt moved the temples of Abu Simbel in the most famous archaeological rescue ever mounted, an international UNESCO campaign that dismantled the monuments and re-erected them above the rising lake [C44]. Resettlement programmes, however grudging, were written into projects. But the era's accounting always ran the same way: the benefits were national and electrified, the costs local and rural, and the people who paid them rarely had a voice in the capital. DneproGES, Aswan, Kariba - each was, in its own register, a wager that the river's fall could be converted into a nation's future.

The wager paid, often spectacularly. The bill, when it arrived, was itemised in a different currency: drowned valleys, displaced communities, rivers whose floodplains and deltas had been built by the very floods the dams ended. The giants made the modern world. The modern world has spent fifty years arguing about the price.

What none of the builders yet reckoned with was that the model's costs were accumulating as steadily as its reservoirs, and that the reckoning, when it came, would arrive from the rivers themselves.

Chapter 16 - The binational giants

The largest power stations on earth sit on borders, and that is no accident. A great river that divides two countries offers its fall to whoever can agree to share it, and the agreement itself becomes a work of engineering as delicate as the dam. Two projects stand for the type, and between them they generated more electricity than any machines humanity has ever built.

Itaipu is the greater. The Parana River, second only to the Amazon among South America’s rivers, forms the border between Brazil and Paraguay, and in April 1973 the two countries signed the Treaty of Itaipu to develop its fall together [C27]. The arrangement was unique: a binational entity, owned fifty-fifty, to build and run one enormous plant and split its energy. The dam that rose over the following decade is a composite giant - concrete gravity sections, buttress wings, embankment dams stretching the barrage across the whole valley - and its powerhouse carries twenty 700-megawatt Francis units, 14,000 megawatts in all [C63][C26]. Itaipu came online through the early 1980s and immediately settled into a role no plant had ever played: it supplies roughly 8.6 per cent of all the electricity consumed in Brazil and the overwhelming majority of Paraguay’s [C26]. In 2016 it generated 103.098 terawatt-hours, a world record for annual production that stood until Three Gorges passed it in 2020 [C63].

The treaty’s economics were as engineered as the dam. Paraguay, a country of a few million people, could never absorb half of Itaipu’s output, so the treaty obliges it to sell its unused share to Brazil at prices fixed by the agreement - a clause that made the dam, for Paraguay, both a revenue engine and a standing grievance. The arrangement was written for fifty years, and the mandated revision of its financial annex, opened when the treaty turned fifty in 2023, became one of the most consequential energy negotiations in South America’s recent history: a tariff agreement in 2024 fixed the power price through 2026 and opened Brazil’s market to Paraguay’s surplus power, and the full annex revision was still being negotiated at the time of writing [C27b]. A dam is a permanent fact; a binational dam is a permanent negotiation.

Itaipu’s generation record deserves a word of precision, because the industry trades in “largest” claims that split three ways. Three Gorges is the largest power station by capacity at 22.5 gigawatts; but capacity is only potential. By actual energy produced in a single year, Itaipu’s 2016 record of 103.098 terawatt-hours stood until Three Gorges, helped by an exceptional wet year, recorded 111.8 terawatt-hours in 2020 [C63][C50]. And by cumulative lifetime output, Itaipu - running at full stretch for four decades on a reliable river - remains the benchmark against which every other station’s productivity is measured. The Parana’s gift is

steadiness: a huge, well-watered catchment and a reservoir that smooths it, feeding machines that run near capacity year after year.

Guri, on Venezuela's Caroni River, is the second type of giant: a national project of resource nationalism. The Caroni drops from the Guiana Highlands through a canyon country of extraordinary hydropower wealth, and Venezuela built there what became, through staged construction from the late 1960s to the mid-1980s, a 10,235-megawatt complex in two powerhouses behind a dam 7,426 metres long [C64]. Guri supplies more than 60 per cent of Venezuela's electricity and sends it out on 765-kilovolt lines among the highest-voltage in the world [C64]. It powered the aluminium industry Venezuela built to spend its oil wealth on something durable. And when Venezuela's wider economy and institutions decayed in the 2010s, Guri's output - and the blackouts when it faltered - became a barometer of the state itself. A nation that builds its grid on one great river builds its vulnerability there too.

The construction of Itaipu was itself a campaign in the New Deal tradition, scaled up. The Parana had to be diverted around the works through a channel two kilometres long; the concrete was placed in volumes that rivalled Hoover's whole dam every few months; and the workforce, at its peak, numbered in the tens of thousands. The project came with the era's costs, notably the drowning of the Guaira Falls - the greatest waterfall by volume on earth - behind the rising reservoir, a loss that Brazil mourned and then largely forgot as the power flowed. Itaipu is the mid-century dam idea at its maximum: the state, the river, the wall, the industrial current. That its power now underwrites a tenth of the largest economy in Latin America is the strongest argument the type ever produced [C26].

The binational and national giants of the Parana and the Caroni closed out the American century of hydropower. Between Hoover in 1936 and Itaipu's completion in 1983, the industry had executed the twentieth century's great proposition: that the fall of the world's great rivers could be captured by the state and converted into national development. The proposition built the modern electricity industry, powered two industrial booms, and left a chain of monuments from the Colorado to the Caroni. It also left the industry with a question it had never had to ask: once the great sites are taken, what next? The answer came from the east, and it came on a scale that dwarfed everything before it.

Chapter 17 - The Chinese century

Hydropower’s centre of gravity crossed the Pacific at the turn of the twenty-first century, and it landed on the Yangtze. China had built dams for decades, but the Three Gorges project, authorised in 1992 and built over the following two decades, was of a different order: the largest power station on earth, on the river that defines Chinese geography, undertaken as an explicit instrument of national development. The dam is a concrete gravity structure 181 metres high with an axis 2,309 metres long; its powerhouse carries thirty-two 700-megawatt main units plus two 50-megawatt station units, 22,500 megawatts in total; and its double-way five-step ship lock and vertical ship lift keep the Yangtze’s traffic moving past the wall [C50]. The official cost was 184.9 billion yuan, reported as within budget by the State Council’s project office in 2009; unofficial totals that count financing and transmission run considerably higher [C50]. The reservoir displaced about 1.27 million people by the project’s own accounting - the largest resettlement ever undertaken for a dam [C50].

Three Gorges was built for three reasons, and the order matters. First, flood control: the Yangtze’s floods had killed hundreds of thousands across the centuries, and the 1954 and 1998 floods were national traumas. Second, navigation: the ship locks turned the gorges’ dangerous reaches into a deep-water shipping channel into China’s interior. Third, power - and the power, when it came, was prodigious. In 2020 the station generated 111.8 terawatt-hours, the most electricity any single plant has ever produced in a year [C49]. The dam that was justified by floods became the anchor of the world’s largest electricity system.

What followed Three Gorges was a cascade boom without precedent. China moved up the Jinsha - the upper Yangtze - and its tributaries, building a staircase of giants. Baihetan, the second-largest power station on earth at 16 gigawatts with sixteen one-thousand-megawatt units, entered full operation in December 2022 [C62c]. Wudongde, Xiluodu, Xiangjiaba and others filled the cascade, and China’s dam-builders became the world’s: by one industry-federation estimate in 2019, Chinese companies held over 70 per cent of the international hydropower construction market and had participated in 320 projects totalling 81 gigawatts across 140 countries [C51]. The figure is an industry claim and should be read as one, but the phenomenon it describes is real: where the World Bank and the Western utilities had retreated from big dams, Chinese state firms and policy banks advanced, financing and building the dams of Africa and Asia.

The engineering establishment that did this was built deliberately. China’s dam agencies and construction conglomerates - Sinohydro, the China Three Gorges Corporation, PowerChina - learned on the domestic cascade and then exported the

capability, pairing construction contracts with policy-bank finance in a package no Western competitor could match. Boston University's Global Development Policy Center, tracking China's overseas development finance, counts tens of billions of dollars in loans from the China Development Bank and the Export-Import Bank of China for hydropower projects totalling more than 27 gigawatts across 38 countries [C51]. After Beijing's 2021 pledge to stop financing overseas coal, the same apparatus pivoted harder into renewables, with state contractors dominant and hydropower prioritised in the least-developed countries [C51b]. The dam, which the mid-century had made an instrument of American and Soviet influence, became an instrument of Chinese influence - with the same mixture of genuine development and hard leverage.

The cascade logic deserves emphasis, because it is how China thinks about its rivers. A single dam exploits a single fall; a cascade of dams, stepped down a river's whole descent, harvests nearly all of it, each reservoir feeding the next powerhouse's intake. On the Jinsha alone, four stations - Wudongde, Baihetan, Xiluodu, Xiangjiaba - form such a staircase above Three Gorges, and together they constitute the largest concentration of hydroelectric capacity ever built on one river system [C62c]. State media present the cascade as a "clean energy corridor," feeding the coastal cities through ultra-high-voltage lines. The corridor is real, and so is its corollary: the upper Yangtze is now one of the most completely harnessed great rivers on earth, its fall converted to current from the Tibetan foothills to the sea.

The domestic programme was matched, inevitably, by frictions. The Yangtze cascade changed a great river's pulse; Three Gorges' reservoir drowned cities and farmland and displaced more people than any dam before it; and China's dams on the Mekong's upper reaches placed the water of six downstream nations under upstream control, turning every dry season into a diplomatic event [C50][C24][C82]. The project's defenders point to the floods it has prevented and the shipping it has enabled; its critics to the resettlement and the reservoir's unstable slopes. Both are describing the same structure, which is precisely the point: Three Gorges is the dam age's largest single statement, and it contains the whole argument - power, water, people and politics - at its largest scale.

The boom made China the undisputed centre of the industry. In 2024 China added 14.4 gigawatts of hydropower capacity - more than the rest of the world combined - and held more than 91 gigawatts of pumped storage under construction, on track to blow past its 2030 target of 120 [C8][C48]. The Chinese century of hydropower is, like the American one before it, a story of the state treating the river as infrastructure for national development, executed at a scale and speed that earlier age never approached. It is also, increasingly, the setting for the industry's hardest questions - about displaced people, dammed international rivers, and the

climate the dams were built to serve - which the later chapters of this book take up in turn.

PART VII - THE GRID'S WORKHORSE

Chapter 18 - The flexible generator

A power grid is the largest machine ever built, and it has no storage tank. Every second, the electricity generated must equal the electricity consumed, plus losses; the frequency that every generator shares is the system's pulse, and it must be held within a narrow band or the machine begins to tear itself apart. Keeping that balance, second by second and year by year, is the hidden work of the power system, and hydropower does more of it, in more ways, than any other technology. The dam that the nineteenth century built for energy turned out to be the twenty-first century's indispensable servant.

The first service is the simplest: dispatchability. A reservoir hydroelectric plant can be told to generate, and it generates - gates open, water flows, megawatts rise in minutes. Within its water budget, the operator chooses when to run, and that choice is worth money, because electricity's value swings wildly through the day. A storage dam sells into the evening peak and rests through the cheap hours, earning not just energy prices but timing. Run-of-river plants, by contrast, sell the river's timing - cheap, steady energy that arrives whether or not anyone wants it. The two products, energy and flexibility, are both sold under the single name "hydropower," and much of the industry's economics turns on the difference.

The second service is regulation. The grid's frequency drifts with every imbalance, and something must correct it continuously. Hydropower's governors - the reflex systems met in the machine hall chapter - do this superbly: the machines can sweep from part load to full load in tens of seconds, absorbing or supplying power as the system requires. Grids pay explicitly for these services - frequency response, spinning reserve, voltage support - and hydro fleets, with their big synchronous machines and fast gates, supply a disproportionate share of them. In the United States, the Department of Energy's Hydropower Vision study and its HydroWIRES initiative, launched in 2019, were built around exactly this recognition: that hydropower's value to a changing grid lies as much in its flexibility as in its kilowatt-hours [C2].

The system value of these services is easiest to see where they are scarcest. On grids with little storage, every solar farm and wind park forces some flexible generator to shadow its output - ramping down when the sun floods the system, ramping up when a cloud front or the evening arrives. Where hydro exists, it does the shadowing naturally, and its reservoirs grow more valuable with every megawatt of variable capacity added. Where hydro does not exist, the shadowing falls to gas turbines and, increasingly, batteries - technologies that must be paid for the readiness hydro was born with. The International Hydropower Association's claim that no country has approached a fully renewable grid without

hydropower in the mix is an advocacy body's claim, but the physics behind it is plain [C7b].

The third service is the one no other generator can easily provide: recovery from catastrophe. When a grid blacks out, thermal stations cannot restart unaided - they need power to make power. A hydroelectric unit can start from nothing, energise a dead line, and begin nursing the network back to life. Black-start capability is insurance the whole system buys, and the great storage dams are its underwriters.

The fourth service arrived with the twenty-first century's new problem: wind and sunshine. Solar and wind generation are excellent energy sources and terrible servants - they produce when the weather allows, not when the grid asks. A system rich in them needs something that fills the gaps: the calm night, the still evening, the week of grey weather. Hydropower is the natural partner. Its reservoirs are, in effect, storage for the variable renewables' surpluses and shortfalls, and the International Hydropower Association's 2025 outlook makes the pairing the centre of its case: hydropower delivers the flexibility and resilience that systems dependent on variable sources require [C48]. Norway sells this service across borders, banking its reservoirs against Europe's wind; every grid with both resources learns the same complementarity.

There is a fifth service, so basic it is usually invisible: inertia. The spinning mass of a large synchronous generator resists changes in the grid's frequency the way a flywheel resists changes in speed - it is the first line of defence in the seconds after a disturbance, before any governor can react. A grid rich in spinning machines is physically stiff; a grid rich in inverters, which feed power through electronics with no inherent inertia, is not. As solar and wind displace thermal and hydro machines, system operators have had to start buying inertia as a service in its own right - and hydroelectric generators, big heavy synchronous machines, provide it by their mere existence. It is a strange fate for a power plant to be valued for refusing to change speed, but that is what a high-renewables grid increasingly pays for.

The result is a quiet reversal in the industry's identity. For a century, hydropower sold bulk energy and was judged on kilowatt-hours. Increasingly it sells readiness - capacity, flexibility, inertia, insurance - and is judged on services a meter barely records. The oldest major power source on the grid has become the enabling technology of the newest ones. And its most direct expression, the machine that exists purely to shift energy through time, is the subject of the next chapter.

Chapter 19 - The water battery

Pumped storage is hydropower’s strangest and most valuable branch: a power plant that consumes more electricity than it produces, on purpose. The idea is disarmingly simple. Build two reservoirs at different heights; when electricity is cheap and abundant, pump water uphill; when it is scarce and dear, let the water fall back through the turbines. The plant loses perhaps a fifth to a quarter of the energy round-trip and profits anyway, because energy at eight in the evening is worth far more than the same energy at three in the morning. It is arbitrage performed with water and gravity, and it is the oldest, largest and most proven form of energy storage humanity possesses [C6].

The technology is older than its reputation. The first large-scale pumped storage development in the United States was the Rocky River plant in Connecticut, begun in 1929, built to raise the effective capacity of the Housatonic River’s hydro system by storing off-peak power in the newly flooded Candlewood Lake [C2]. For decades pumped storage served a steady, unglamorous role: soaking up the surplus of coal and nuclear plants that preferred to run flat out through the night, and repaying it into the daytime peak. The great stations of the late twentieth century - Dinorwig in Wales, opened in 1984 with six reversible units inside a hollowed slate mountain, able to reach full output in about twelve seconds; Bath County in Virginia, which began operation in 1985 as the world’s largest at what is now 3,003 megawatts with some 24,000 megawatt-hours of storage - were built in that mould [C60][C41]. They were the power system’s shock absorbers, famous among engineers and invisible to everyone else.

The twenty-first century made them famous. As wind and solar capacity exploded, the value of storage - of anything that could shift energy from when it floods in to when it is needed - exploded with it. Pumped storage, a mature technology sitting on more than a century of operating experience, suddenly found itself the anchor of every serious plan for a high-renewables grid. The numbers tell the story. Global pumped storage capacity reached 189 gigawatts in 2024, having grown 5 per cent in a single year; annual additions have nearly doubled in two years; and pumped storage still accounts for more than 90 per cent of all the energy storage installed on earth - more than every battery in the world combined [C8]. The industry’s development pipeline, as of 2024, held some 600 gigawatts of proposed pumped storage - nearly three times what the association’s previous pipeline survey, in 2023, had found [C8].

China has seized the role with the same scale it brought to conventional dams. It added 7.75 gigawatts of pumped storage in 2024 alone, had more than 91 gigawatts under construction - on course to exceed its 120-gigawatt 2030 target - and a further 136 gigawatts moving through development [C8]. Its Fengning

station in Hebei, at 3,600 megawatts with twelve units, became the world’s largest pumped storage plant when it reached full operation at the end of 2024 [C59]. The Global Energy Storage and Grids Pledge, signed by fifty-eight countries at COP29, commits to 1,500 gigawatts of storage by 2030 - a target that, on current technology, pumped storage must largely carry [C8].

The economics of the water battery are the economics of the dam itself, turned on their side. Like a conventional hydro plant, a pumped storage station is nearly all civil works - reservoirs, tunnels, powerhouse cavern - with a long life and low running costs once built. Bath County cost about 1.6 billion dollars to build between 1977 and 1985 and, with a turbine refurbishment completed in 2009, continues to serve as the largest battery on the North American grid nearly half a century later [C41][C2]. The pattern recurs: pumped storage is expensive to build, cheap to run, and almost indestructible, so the fleets of the 1970s and 1980s are being refurbished and uprated rather than replaced, while an entirely new generation of projects is drawn up beside them. It is the closest thing the energy business has to a perpetual-motion annuity - the same water, pumped up and let down for a century, selling the same hours over and over.

The machines themselves are a quiet marvel of adaptation. Early pumped storage plants used separate pumps and turbines; the modern plant runs a reversible pump-turbine, a Francis-type machine that spins one way to generate and the other to pump, switching roles in minutes. Dinorwig’s six reversible units can go from standstill to full output in about twelve seconds - a response speed that made the station, for decades, the first resort of the British grid when a million kettles went on at once [C60]. Bath County’s six units, each over 500 megawatts, make it the largest battery on the North American grid nearly forty years after it entered service [C41]. The technology is not exotic; it is the hydroelectric machine, taught to run backward - and that simplicity is precisely its durability.

The water battery’s virtue is also its limit. It needs two reservoirs and a hill between them; the best sites, like the best dam sites, are finite and often contested. The new generation of projects - closed-loop systems off rivers, seawater schemes, disused-mine proposals - is an attempt to loosen that constraint. But the central fact stands: the oldest storage technology in the electricity business is also, by an enormous margin, the largest, and the grid of the future is being planned around a machine the industry has been building for a hundred years. The dam, having stored the river’s water, was asked to store the grid’s time. It answered with the same wall.

PART VIII - THE BUSINESS

Chapter 20 - The economics of patience

A hydroelectric dam inverts the normal economics of power generation. A gas plant is cheap to build and expensive to feed; a dam is ferociously expensive to build and almost free to run. The consequence is an industry that lives on the longest timescales in the energy business: decades to plan and license, half a decade or more to build, a century to operate. Hydropower is the economics of patience, and almost everything distinctive about its business - who finances it, who owns it, what can go wrong - follows from that.

The cost structure is dominated by civil works. The dam, the tunnels, the powerhouse cavern, the spillways - concrete, earth and excavation - typically outweigh the machinery in the project's accounts, and they are bespoke: every site is different, every design a one-off, every geology a fresh negotiation. The World Bank's guide for developers and investors puts it plainly: each hydropower plant is custom-designed to the conditions of its site, and the civil works that dominate its cost are precisely the works that cannot be standardised [C35b]. This is why hydropower has no learning curve like solar panels or wind turbines, whose costs fall as factories scale. A dam is not manufactured; it is built, once, in a particular valley, and the second one is no easier.

Against that stands the operating account. Once built, the plant needs no fuel, few staff and modest maintenance, and it keeps working for a working life that dwarfs any other power technology. Hoover has been generating for ninety years; DneproGES, rebuilt after the war, for most of a century; the turbines at Niagara's successors run on foundations from the age of Edison. Levelised over such lifetimes, hydropower's cost of electricity is among the lowest of any source: the International Renewable Energy Agency's cost series consistently places hydro at the cheap end of the generation range, and notes that in 2024 renewables broadly undercut fossil alternatives for new capacity [C10]. The catch is the word "levelised" - the cheapness is real only if the capital is patient enough to wait for it.

Patience is precisely what private capital hates and states possess. This is why the industry's ownership history reads as a history of the state: the New Deal authorities, the Soviet cascades, the post-colonial development dams, the Chinese conglomerates. Private utilities built the industry's childhood, when sites were small and paybacks quick; the giants required balance sheets that could absorb a decade of spending before the first kilowatt-hour, and risk appetites for geology, hydrology and politics that no shareholder meeting enjoys. Even today, the International Hydropower Association's core complaint is that markets alone will not deliver the investment the technology needs - that hydropower's benefits are

system-wide and long-term, while its revenues are plant-level and near-term [C48].

The operating account has one more line that separates hydro from every rival: the machines can be renewed without rebuilding the plant. Turbine runners are replaced with more efficient designs, generators rewound, governors digitised, and a station's capacity can be uprated by ten or twenty per cent in the process - new megawatts from old concrete, the cheapest capacity in the industry. Bath County's 2009 refurbishment lifted its output and reclaimed its title as the world's largest pumped storage plant for another fifteen years [C2]. The United States, whose fleet is among the world's oldest, has made this its principal hydro strategy: the Department of Energy's analyses find the great majority of the country's hydroelectric potential now lies not in new dams but in powering existing non-powered dams and upgrading the plants already built [C2b]. The dam, it turns out, is the one power asset that appreciates with reinvestment rather than merely depreciating.

The revenue side has its own subtleties. A dam sells into markets that were rarely designed for it: energy markets that pay for kilowatt-hours and undervalue the flexibility, capacity and insurance the plant also provides. The industry's modern argument - pressed hardest in liberalised markets - is that it is paid for the wrong product. The solar and wind buildout has begun to change that, creating explicit markets for capacity and ancillary services that pay hydro for its readiness rather than only its output. There is, finally, the line that no account captures and every operator fears: hydrology. The dam's revenue is rain, deferred. A dry year cuts output across a whole system - the 2023-24 droughts in Latin America and southern Africa dented global hydropower generation measurably - and a sequence of dry years can break a utility's finances [C8]. Thermal plants buy fuel; dams inherit weather. The industry's longest-standing business risk is the one the first chapter stated as physics: the resource is free, but it is not promised. The financing structures, the reservoir rules and the regional compacts of the hydro world are all, at bottom, devices for sharing that risk across years and across borders.

The business of the twenty-first-century dam, where it works, is the business of being paid to wait: storing water, holding readiness, and selling the moments when the grid has no one else to call. It is an old trade, and the world's newest energy economy has made it newly valuable.

Chapter 21 - The financiers’ retreat and return

For half a century, the world’s dams were financed by the world’s development bank, and the story of the industry’s conscience can be told through the World Bank’s relationship with them. The Bank was, from its founding, the most important international funder of large dams: by one careful accounting, it supported more than 550 dams around the globe with over 90 billion dollars in loans and guarantees [C16]. The money built the post-colonial dam age - the Karibas and Tarbelas and Akosombos - and it built, alongside the reservoirs, a growing body of evidence about what the reservoirs cost.

The dams the Bank financed were rarely mistakes in engineering; they were frequently tragedies in accounting. The benefits were real - power, irrigation, flood control - and so were the costs, but they fell on different ledgers. The power went to cities and industries; the displaced went to resettlement colonies whose promised land and jobs often failed to materialise. The Bank’s own Operations Evaluation Department, reviewing its large-dam portfolio in the 1990s, found resettlement outcomes so consistently poor that the review itself became part of the case for change [C45]. The lesson the institution slowly absorbed was not that dams were wrong but that their economics had been computed with some of the costs left out - and that the people who bore those costs had no seat at the table where the computation was made.

The retreat, when it came, was driven by the evidence. Through the 1980s and 1990s, dam projects drew mounting criticism over displaced people, drowned ecosystems and resettlement promises broken, and the Bank’s own reviews conceded the pattern: its dam financing slowed by about half from the mid-1980s to the mid-1990s [C45]. The controversy culminated in an experiment without precedent: the World Commission on Dams, an independent body of twelve commissioners drawn from industry, governments, academia and activism, established in 1998 with the Bank’s blessing to review the development effectiveness of large dams. Its final report, Dams and Development, delivered in November 2000 before Nelson Mandela, reached a deliberately balanced verdict: dams had delivered significant benefits, but at a price too often borne by people who shared neither the power nor the profits, and decisions about them had systematically excluded the affected [C15]. The Commission proposed a framework of rights-and-risks assessment that satisfied almost no one entirely - industry thought it a recipe for paralysis, activists for legitimisation - but it fixed the terms on which every dam proposal since has been judged.

The vacuum the retreat created did not last. As the multilateral banks stepped back from big dams in the 1990s and 2000s, a new financier stepped forward: China. Its policy banks - the China Development Bank and the Export-Import Bank of

China - and its state construction conglomerates offered developing countries a package the West no longer would: finance, equipment and construction, delivered fast and with few environmental conditions. Boston University's Global Development Policy Center counts more than 44 billion dollars in Chinese policy-bank loans for hydropower projects totalling over 27 gigawatts across 38 countries, and a Chinese industry-federation official claimed in 2019 that Chinese companies held over 70 per cent of the international hydropower construction market [C51]. The figure is an industry claim and the lending data carry their own uncertainties, but the structural fact is plain: the centre of dam finance moved east, and with it much of the industry's momentum. After Beijing's 2021 pledge to stop financing overseas coal, the same machinery pivoted further into renewables, with hydropower prioritised in the least-developed countries [C51b].

The debate the Commission settled into has never really closed, because it is not resolvable in principle: it is the argument, endemic to all large infrastructure, about who may impose costs on whom for a collective benefit. Dams merely make the argument concrete - literally. What changed after 2000 is that no major project can now proceed as if the argument did not exist. Environmental and social impact assessment, resettlement action plans, sustainability certification: these are the institutional residue of the retreat, imperfect and often evaded, but real. A dam proposed today, anywhere, must answer questions that Aswan and Kariba never had to hear.

The result is the split personality of modern dam finance. In the rich world, the money goes to refurbishment, pumped storage and grid services - maintaining and re-valuing what was built in the last century. In the developing world, it still goes to new giants on the great remaining rivers, financed increasingly by China and built to the same proposition that animated the New Deal: that a river's fall can fund a nation's rise. And in the background, the multilateral institutions have edged back toward hydropower, chastened, under the banner of sustainability standards - the International Hydropower Association's Sustainability Standard, the certification schemes that now badge individual projects - trying to institutionalise the lessons the World Commission on Dams wrote down [C48]. The financiers' argument, like the river, never really stops. It only changes course.

PART IX - THE WORLD MAP

Chapter 22 - Where the water falls

Hydropower’s world map is drawn by two features of the planet: rain and relief. Where both coincide - mountains in the path of wet winds - the resource concentrates; where either is missing, no engineering can summon it. The result is a geography as fixed as geology, and the industry’s history is largely the story of nations discovering which piece of that geography they happened to inherit.

The map’s first fact is Asia’s dominance, and within Asia, China’s. The Himalayas and the Tibetan Plateau - the water tower of a continent - feed the Yangtze, the Mekong, the Brahmaputra, the Salween, the Irrawaddy and the Indus, rivers whose combined fall is the richest hydropower endowment on earth. China has built the upper Yangtze into the world’s greatest cascade; the Tibetan rivers that flow south and east remain, in large part, the industry’s last great unbuilt frontier, and its most geopolitically charged. In 2024 China added 14.4 gigawatts of hydroelectric capacity, more than the rest of the world combined [C8].

The second fact is the Americas’ long relationship with the resource. South America draws about 45 per cent of its electricity from hydropower - the highest dependence of any major region - built on the Parana, the Caroni and the Amazon’s tributaries, with Brazil at its centre [C8]. North America built its great fleets generations ago: the United States on the Colorado, the Columbia, the Tennessee and the Tennessee’s neighbours; Canada on the St Lawrence, the Columbia’s Canadian reaches, and above all the shield rivers of Quebec, whose James Bay project made Hydro-Quebec one of the largest hydroelectric producers on earth. Both countries now live off the inheritance: in 2024 North and Central America added a modest 541 megawatts of new capacity, and the work of the region is modernisation, not expansion [C8].

Europe’s hydro geography splits three ways. The Alps gave the industry its early giants and still anchor the fleets of Switzerland, Austria, France and Italy; the fjord-cut plateaus of Norway made that country the most hydro-dependent major economy on earth, drawing nearly all its electricity from its reservoirs and selling their flexibility across the North Sea; and the Iberian and Balkan ranges fill in the map. Europe’s rivers are nearly all spoken for - the industry’s future there is pumped storage and refurbishment, and 2024’s generation, at 680 terawatt-hours, was a decade high driven by wet weather rather than new dams [C8].

The map’s unfinished half is Africa. The continent holds more than 600 gigawatts of technical hydropower potential and has harnessed only about a tenth of it; hydro already supplies about a fifth of Africa’s electricity from that small fraction [C48]. The Congo River, second on earth by flow, carries at its Inga Falls a theoretical endowment of more than 40 gigawatts - the Grand Inga scheme, if ever

built in full, would be the largest power station project in the world by a factor of two. For decades it has remained a plan, defeated not by engineering but by finance, governance and the sheer risk of the site. Africa's newer projects tell the same story at smaller scale: Ethiopia's Renaissance Dam on the Blue Nile, the continent's largest, filling its reservoir against Egypt's anxiety; a scatter of Chinese-built dams from Uganda's Karuma to Cameroon's Nachtigal, fully commissioned in 2024 [C8]. The map of where hydropower will be built in the twenty-first century is, more than anywhere else, the map of Africa - if the finance can be found.

Russia holds a distinctive place on the map: enormous endowment, mid-century buildout, aging fleet. The Volga and the great Siberian rivers - the Yenisei, the Angara, the Lena's basin - carry cascades built in the Soviet decades, crowned by Sayano-Shushenskaya's 6,400 megawatts [C62b]. The Soviet planners paired their dams with aluminium smelters as deliberately as the Americans paired theirs with war industry, and the Siberian hydro-metallurgical complexes remain among the largest electricity consumers on earth. Russia's hydro story since has been maintenance and, occasionally, catastrophe - the 2009 Sayano-Shushenskaya accident killed seventy-five people and flooded one of the world's great machine halls [C62b].

India sits in a distinctive position on the map: a large, fast-growing electricity system with a substantial Himalayan and Western Ghats endowment, a long history of state dam-building from Bhakra Nangal onward, and a development pipeline that, at 132 projects in various stages in 2024, is among the world's largest [C8]. Its hydro story, like China's, is inseparable from the politics of shared rivers - the Indus waters treaty with Pakistan, the Brahmaputra with China - and from the long domestic argument over large dams that the Narmada campaigns made famous. It is also, increasingly, a pumped storage story, as the country's solar buildout creates the same demand for storage that is driving the technology worldwide.

The final fact of the map is that it does not move. Solar and wind resources are diffuse; hydropower is fixed to particular valleys, and the valleys that remain unbuilt are unbuilt for reasons - distance, poverty, politics, ecology - that no improvement in turbines will cure. The industry knows this. It is why its centre of gravity has shifted from building the map to working it: storing more in its reservoirs, squeezing more from its machines, selling its flexibility to the grids that now surround it. The geography was settled long ago. The argument now is over what the geography is worth.

Chapter 23 - The unbuilt rivers

For every great river harnessed, others remain wild, and the reasons why are the industry’s frontier. The largest of the unbuilt prizes is the Congo. The river is second only to the Amazon in flow, and at Inga, near its mouth, it falls nearly a hundred metres through a series of rapids in a short reach - a concentration of power unmatched anywhere on earth. Engineers have dreamed of Grand Inga for half a century: a scheme of more than 40 gigawatts that would dwarf Three Gorges and could, in principle, supply much of a continent. Two smaller Inga dams, built under Mobutu, stand there now, decayed and under-producing. Grand Inga remains a drawing, not because it cannot be built but because it cannot yet be financed, governed or trusted - a standing lesson that the limit of hydropower is rarely the river.

The second frontier is the Himalayan arc. The rivers that drain the Tibetan Plateau - the Yarlung Tsangpo above all, but also the Salween, the Mekong’s upper reaches, the Brahmaputra’s tributaries - carry the largest unharnessed fall on earth, and they are held between countries that regard each other’s dams with strategic suspicion. China’s dams on the Mekong’s upper reaches have already put downstream drought and flood under upstream control, to the alarm of the six nations that share the river; every new project on the Brahmaputra is read in Delhi as a threat as much as a power scheme [C82][C83]. Water was always politics; on the Himalayan rivers it is geopolitics with the taps attached.

The third frontier is quieter: the small rivers. The industry’s giants monopolise its image, but the great majority of its plants are small, and much of its untapped potential lies in small streams, old mill sites and, above all, existing dams that have no turbines. The United States has more than 90,000 dams and generates power at fewer than 3 per cent of them; the Department of Energy’s analyses treat that stock of non-powered dams - built for flood control, water supply and recreation - as one of the cheapest sources of new hydroelectric capacity in the country [C1][C2b]. The logic extends worldwide: the world is already dammed, for other purposes, and bolting turbines onto structures that already exist avoids nearly all of the cost and most of the controversy of building anew.

The map of the unbuilt also includes a category the industry prefers not to advertise: the sites that will never be built because the world decided they were not worth the cost. The great environmental battles of the late twentieth century - over the Amazon’s tributaries, the Mekong’s mainstream, the wild rivers of Patagonia and the Balkans - ended with some dams cancelled and some rivers protected outright. Chile’s HidroAysen scheme, planned for the wild Baker and Pascua rivers of Patagonia, was abandoned after years of protest; the campaign against it became a national referendum on what development was for. Every

unbuilt dam of that kind is a data point in the argument the World Commission on Dams started: the question is no longer whether a river can be dammed, but whether it should be, and who gets to say.

And then there is the frontier that is not a place but a change: the climate. Hydropower's resource is rain and snow, and both are shifting. The droughts that cut global generation in 2023 and rebounded in 2024 are a taste of the variability the industry expects to deepen [C8]. Some regions will get wetter and some drier; snowpacks that stored water for summer melt are shrinking; the design floods that dams were built to survive are being recalculated upward. The industry's maps - of resource, of risk, of revenue - were drawn from the rivers of the twentieth century. The twenty-first is redrawing them underneath the concrete, and no dam can be unbuilt to match. There is a final irony in the geography. The regions with the largest unbuilt potential - Africa, the Himalayan arc - are also the regions where the need for power is most acute and the alternatives most expensive. The Congo could light a continent that remains the least electrified on earth; the Himalayan rivers could feed the fastest-growing electricity markets in history. The dams are unbuilt not because the power is unwanted but because the projects are hard - hard to finance, hard to govern, hard to justify against the rivers they would change. Whether the twenty-first century builds them, and on what terms, is one of the largest open questions in the world's energy future.

The unbuilt rivers and the changing ones meet at the same point: hydropower's future is bounded less by what engineers can build than by what the water, the politics and the climate will allow.

PART X - THE COSTS DOWNSTREAM

Chapter 24 - The drowned valleys

Every reservoir drowns something. That is not rhetoric but geometry: to store water you must cover land, and the land behind a great dam was, almost always, already in use - farmed, lived on, worshipped in, fished, remembered. The World Commission on Dams, reviewing the global record in 2000, put the worldwide number of people physically displaced by dams at between 40 and 80 million - a figure the Commission itself stressed was an estimate built on incomplete records, but whose scale no one disputes [C15]. The history of hydropower is inseparable from that number, and the industry's conscience is the long, unfinished argument about it.

The largest single displacement ever undertaken for a dam was Three Gorges. The official accounting records about 1.24 million people relocated by June 2008 and 1.27 million by June 2009, with resettlement spending of 69.3 billion yuan [C50]. The Yangtze's gorges had been farmed and settled for millennia; the reservoir rose over cities, towns, factories and farmland, and the people moved - some to new towns nearby, some to distant provinces, some to livelihoods that never recovered. Independent reporting through the project's completion found relocated families facing dwindling prospects and, in places, further forced moves as unstable new slopes failed [C24]. The Chinese state built the dam, in part, to end the floods that had killed far more people than the reservoir displaced; the arithmetic is real, and so is the fact that the people who paid it were not the people it protected.

The pattern recurs across the dam age with monotonous geography: the valley is rural, the power is urban, and the displaced are almost always the people with the least say. Kariba's rising water removed the Tonga people of the Zambezi valley by the tens of thousands in the 1950s, scattering communities whose lives had been organised around the river for centuries [C66][C29]. Aswan's Lake Nasser drowned Lower Nubia, and tens of thousands of Nubians were resettled from a homeland their culture had occupied since antiquity, most never to return [C43] [C44]. The monuments at least were saved: the UNESCO campaign that moved Abu Simbel was the first great international rescue of its kind, and the contrast it draws - the world mobilised for temples, not for villages - is one the displaced themselves have not failed to notice [C44].

The mechanics of displacement deserve plain statement, because they explain why it so often went badly. A reservoir does not merely cover houses; it removes the land that fed a community, the commons that grazed its animals, the river that fished and the neighbours that made its economy. Cash compensation, where it was paid at all, bought new houses but not new livelihoods; land-for-land promises, where they were made, ran against the simple shortage of equivalent land. The World Bank's own reviews found resettlement outcomes so consistently

poor across its portfolio that the failures became a central exhibit in the case for reform [C45]. The displaced were, as a rule, poor before the dam and poorer after it - and the benefits, flowing to distant cities, never reached them.

The record is not uniformly bleak, and honesty requires the qualifiers. Some resettlement, in some places, delivered on its promises; some communities gained from irrigation, flood protection and power. The World Commission’s verdict was not that dams should never be built but that the people who bear their costs must have a voice in building them - that rights and risks must be assessed together [C15]. What changed after the dam age’s peak is that the displaced found ways to be heard: the anti-dam movements that grew from the Narmada valley’s long resistance in India to the global campaigns of the 1990s made resettlement a central political question rather than an administrative footnote. No large dam is proposed today without an answer, however contested, to the question the old projects answered by force: who gets to say no.

The drowned valleys are the industry’s oldest debt, and they are still being paid down - in compensation schemes, in resettlement standards, in the simple new fact that displacement is now a project risk priced into every feasibility study. The water behind the dam stores the river’s fall. It also stores, for as long as anyone downstream remembers, the names of the places the river took in exchange.

Chapter 25 - The river interrupted

A dam is a wall across a working river, and the river notices. The water that leaves the reservoir is not the water that entered it: clearer, colder or warmer by depth, stripped of its sediment, released on the operator’s schedule rather than the seasons’. Downstream of every dam, the river becomes a managed artefact, and the consequences run from the riverbed to the sea.

The first consequence is the sediment. Rivers move the ground they cross, and a reservoir is a trap: the flow slows at the pond’s edge and drops its load. The Aswan High Dam stopped the Nile’s silt almost completely, and with it the annual gift of mud that had fertilised the valley and built the delta for five thousand years [C65][C47]. Egyptian farmers turned to artificial fertiliser to replace it; the delta, no longer fed, began to erode into the Mediterranean; and the riverbed below the dam, scoured by clear water hungry for its missing load, began to cut downward. Every large reservoir runs the same account: the storage that makes the dam useful is slowly filled by the very sediment it stops - a loss measured, over decades, in dead storage and, downstream, in starving banks and deltas.

The second consequence is the fish. Rivers are migration routes, and a dam is a full stop. The salmon rivers of the Columbia and Snake systems were among the richest on earth before the dams; the runs collapsed as the walls rose, despite fish ladders and lifts and, eventually, a vast apparatus of barges and hatcheries. The sturgeon of the Volga, the migratory fish of the Mekong - the pattern repeats on every dammed river that had migratory fish to lose. Fish passage technology has improved enormously - ladders, lifts, bypass channels, surface collectors - and on some rivers the runs are being coaxed back; but no ladder makes a dam transparent to a river, and the industry’s honest accounting is that fish passage mitigates, it does not restore.

The third consequence is the flow itself. A river’s life runs on its pulse: the floods that scour channels and trigger spawning, the low water that concentrates fish, the seasonal rhythm the whole ecosystem evolved to expect. Dams flatten the pulse. The flood that should arrive in spring is stored; the water released in winter for power arrives at the wrong season. The World Commission on Dams documented the pattern across its case studies - ecosystems simplified, floodplains dried, deltas starved [C15]. Modern operators increasingly release environmental flows, deliberately re-creating something of the natural pulse; it is a real improvement, and it is also an admission of how much the dam changed.

The fourth consequence is the one that reaches furthest: the estuary and the sea. A river delivers more than water to its coast - it delivers the sediment that builds beaches and deltas, the nutrients that feed fisheries, the fresh water that holds

back the salt. Dam the river, and the coast changes: deltas erode, salt wedges push upstream, near-shore fisheries decline. The Nile delta's retreat after Aswan is the textbook case [C47][C44b2], but the same signature appears at the mouths of dammed rivers worldwide. The wall in the valley, it turns out, is a coastal structure too.

The costs downstream are not the whole ledger, and the next chapters weigh the other side. But they are the ledger's most durable entries. A dam's benefits are largely monetised and largely captured; its ecological costs are largely neither. There is a countercurrent worth naming honestly. Reservoirs also create ecosystems - the lake fisheries of Kariba and Nasser became significant food sources; the regulated flows of managed rivers can protect downstream towns from floods that once killed; and the water stored behind dams is, in much of the world, irrigation and drinking water as much as it is electricity [C44]. The river interrupted is not simply a river destroyed. It is a river converted - from a wild system to a managed one - and the conversion trades one set of goods for another. The argument of the dam age, still unsettled, is over the exchange rate.

The river interrupted keeps its own account, and it settles it slowly - in deltas and fisheries and floodplains, downstream and downstream of downstream, long after the powerhouse's first megawatt has been sold and forgotten.

Chapter 26 - The carbon question

Hydropower’s claim to be clean energy is mostly true and was, for decades, simply assumed. A dam burns nothing; its electricity displaces coal and gas; the industry association estimates that hydropower’s output avoids billions of tonnes of greenhouse gas a year compared with fossil generation, and the IPCC’s lifecycle assessments place hydro’s median emissions alongside the lowest-carbon sources - only wind and nuclear lower [C56]. That is the claim, and it is largely justified. The question that complicated it came from the reservoirs themselves.

A reservoir floods land, and the flooded vegetation and soils begin to decay underwater. Decay without oxygen makes methane, a greenhouse gas far more potent than carbon dioxide, and the reservoir’s surface becomes a slow, bubbling source of it - along with carbon dioxide from ordinary decomposition. The landmark synthesis of the science, by Bridget Deemer and colleagues in 2016, pulled together measurements from reservoirs worldwide and estimated that reservoir water surfaces emit about 0.8 petagrams of carbon-dioxide-equivalent a year - roughly 1.3 per cent of all human-caused greenhouse gas emissions, comparable to the national emissions of Canada - with methane doing most of the warming [C17][C18]. The same research established the crucial qualifier: emissions vary enormously from reservoir to reservoir, driven less by age or latitude than by productivity - the amount of organic matter feeding the decay. A deep, cold, nutrient-poor reservoir in Scandinavia is a different proposition from a shallow tropical one flooded over rainforest.

The complication did not stop at the water’s surface. Subsequent research found that a large share of the emissions may escape not through the reservoir surface at all but downstream - degassing as the water passes through the turbines and spills into the tailrace. A 2023 modelling study estimated that turbine degassing could add substantially to the surface estimates, though with wide uncertainty, not least because the depths from which turbines draw are rarely reported [C55]. And the field measurements sometimes contradict the models: one 2025 review of the standard G-res estimation tool found sign errors against field data - models predicting emission where instruments measured uptake - and warned that the tool’s training data over-represent boreal reservoirs [C56b]. The honest summary is that reservoir emissions are real, material in aggregate, highly variable from site to site, and still imperfectly measured.

The picture that emerges is a spectrum, not a verdict. At one end, deep, cold, steep-sided reservoirs with little flooded biomass have lifecycle emissions comparable to wind; at the other, shallow tropical reservoirs drowned over dense vegetation can, for a time, emit greenhouse gases at rates that rival fossil generation per unit of electricity - with the heaviest emissions concentrated in the

first ten to twenty years after flooding [C56]. The emissions can even, in some cases, be negative: reservoirs that bury carbon in their sediments faster than they emit it act as net sinks [C56]. Hydropower is low-carbon. It is not, everywhere and always, zero-carbon, and the difference is not a footnote - it is a design parameter. Where a dam is built, what it floods, and how it is operated all change the answer.

The policy world has begun to catch up with the science. Sustainability standards now ask developers to estimate reservoir emissions before construction and, where they are high, to change the design - clearing biomass before flooding, drawing water from different depths, choosing different sites. The tools are improving even as their limits are acknowledged. And the comparison that matters, the one that never goes away, is the counterfactual: the relevant question is not whether a reservoir emits, but whether it emits less than the power station it replaces. On that measure, nearly all hydropower wins comfortably over coal and most wins over gas - which is why the IPCC’s median figures place hydro among the lowest-carbon sources despite the reservoir effect [C56]. The carbon question does not convict the dam. It merely insists that the dam, like everything else, be judged on evidence rather than reputation.

The carbon question completes the accounting this part of the book has been building. The dam’s costs are real: the drowned valley, the interrupted river, the warming reservoir. Its benefits are real too: the power, the flood control, the irrigation, the displaced fossil fuel. The case for any given dam is the balance of the two, struck at a particular site by particular people. What the science of the last three decades has done is make the balance harder to fake - which is, from the industry’s own long-run interest, the best thing that could have happened to it.

PART XI - WHEN DAMS FAIL

Chapter 27 - Banqiao

In the first days of August 1975, Typhoon Nina came ashore in China and stopped. Blocked by a cold front over Henan province, the storm parked itself above the headwaters of the Ru river and rained. Over five days it dropped a maximum of 1,631 millimetres of rain - a deluge beyond anything the region's engineers had designed for [C39]. In the valleys below stood dozens of dams built in the 1950s during the great dam campaigns of the early People's Republic, among them the Banqiao dam on the Ru and the Shimantan on a tributary. What happened in the night of 7-8 August remains, by any accounting, the worst dam disaster in history.

The sequence was a cascade in the worst sense. The reservoirs filled past their design limits; communications failed; warnings, where they were given at all, came too late. When the Banqiao dam failed, its breach outflow peaked at somewhere around 78,000 cubic metres per second against an inflow of 13,000 - the reservoir emptying in hours; 701 million cubic metres of water went through the gap in six hours, and the Shimantan dam, downstream, failed under it with a breach peak of 30,000 cubic metres per second [C38]. In all, 62 dams failed in the event - the two large dams, two medium ones, and fifty-eight small ones [C39]. A wave of water tens of metres high swept down onto the plain below, where millions of people lived. Some 12,000 square kilometres were inundated [C38].

The death toll is a lesson in the politics of catastrophe. The official figure, published years later, was about 26,000 dead - and it counted only those who died in the immediate flooding [C39]. The Henan provincial party committee's preliminary figure in August 1975 was 85,600 dead in the province, with the total under 100,000; members of the Chinese People's Political Consultative Conference who investigated in the 1980s put the figure at 230,000, counting those who died of disease and starvation in the aftermath, and Human Rights Watch cited a similar figure in 1995 [C39]. The widely quoted figure of 171,000 is a mid-range estimate that blends the immediate and aftermath deaths without a clear primary source [C39]. What is not in doubt is the scale: tens of thousands died in the water, and far more in what the water left behind. More than ten million people were affected [C39].

Banqiao stands, in the engineering record, with a small company of disasters that defined the limits of dam design. At Malpasset in France in 1959, a thin arch dam failed when the water pressure opened a fissure in its abutment rock, and the wave killed 421 people in the town of Frejus - the lesson that a dam is only as strong as the rock it leans on [C12b]. At Vajont in Italy in 1963, the dam itself did not fail at all: a mountainside of 250 million cubic metres slid into the reservoir during filling, drove a wave 250 metres high over the crest, and erased the town of

Longarone below, killing about 2,000 people - the lesson that the reservoir, not the dam, can be the hazard, and that the geology of the slopes matters as much as the concrete [C46][C14]. In each case the technical failure was compounded by warnings discounted: Malpasset's abutment, Vajont's creeping mountainside and Banqiao's flood forecasts all gave notice, and all were in some measure set aside.

The causes were not only meteorological. The dams had been designed for floods far smaller than the one Nina delivered, on hydrological records that were short and an understanding of extreme storms that was inadequate; and the design philosophy of the era had prioritised storing water for irrigation over passing floods safely. Warnings were suppressed and information controlled, and the disaster itself was concealed from the outside world for years - the government kept news of it from national broadcast, and the full account emerged only slowly [C58b]. The industry's response to its disasters has been, on the whole, to learn from them, because it cannot hide them. Dam-failure case studies - Banqiao, Malpasset, Vajont, and the near-misses like Oroville - are now the core curriculum of dam safety, maintained and studied by bodies like the Association of State Dam Safety Officials [C11][C74]. The result is a profession organised around a single sobering fact: a dam is a structure whose failure kills downstream, and whose safety therefore depends not only on how it is built but on how it is watched, maintained and, when necessary, evacuated around. Banqiao is the extreme case, the worst imaginable. The discipline it helped create exists so that it remains the only one of its kind.

Banqiao is therefore two failures at once: an engineering failure, of dams overwhelmed by a flood beyond their design; and an institutional failure, of a system that could neither anticipate the event nor respond honestly to it. Both lessons entered the world's dam-safety practice. The probable-maximum-flood standards, the emergency-action planning, the institutional candour that modern dam safety demands - all were written, in part, in Henan's flood.

Chapter 28 - The aging fleet

The world’s dams are, on average, old. The great buildout of the mid-twentieth century means that a large share of the global fleet is now past fifty years of age, and the design life written into many of them was not much longer. The concrete endures - mass concrete ages on geological timescales - but the dam is more than concrete. Gates and valves corrode; turbines wear; spillways erode; instruments drift out of calibration; and the institutional memory of how the thing was built and what it was designed to survive walks out the door with each retiring generation of engineers. The safety problem of the twenty-first-century dam is, increasingly, not a failure of design but a failure of upkeep.

The 2009 accident at Sayano-Shushenskaya made the point in the most violent way possible. Russia’s largest power station, a 6,400-megawatt giant on the Yenisei, suffered a turbine failure on 17 August 2009: a runner, weakened over years and held down by aging bolts, tore free under the reservoir’s pressure, and the machine hall flooded in moments. Seventy-five people died [C62b]. The official investigation found not a single cause but a chain - design weaknesses, deferred maintenance, operational decisions - the same chain the forensic teams at Oroville would later describe in California [C73][C62b]. The lesson both events teach is that the danger in an aging fleet is rarely a single dramatic flaw; it is the slow accumulation of small degradations that no one decision ever quite forces anyone to fix.

The problem is sharpened by a change in what the dams are asked to survive. The design floods of the mid-twentieth century were calculated from the rainfall records of the time; a warming atmosphere, which holds more water and can drop it in heavier bursts, is forcing hydrologists to revise those estimates upward. A spillway adequate for the floods of 1960 may be marginal for the floods of 2060, and re-rating a dam’s flood capacity - enlarging spillways, raising crests, adding fusegates - is among the most expensive works a dam can require. The Oroville crisis of 2017, in which both the service and emergency spillways of America’s tallest dam proved inadequate to a wet season, was read across the industry as a warning of exactly this: the fleet was designed for a climate that no longer exists [C74][C73].

Against that stands the profession’s own response, and it is substantial. Dam safety has become a formal discipline: regular inspection regimes, instrumentation that reads a dam’s behaviour continuously, emergency action plans that decide in advance who is warned and how, and review boards of independent experts for the largest structures. The failures that continue to occur are, overwhelmingly, not at the great engineered dams but at the tens of thousands of small, old, often privately owned structures - farm ponds, mine tailings, municipal reservoirs - that

were built to lower standards and inspected less often. The Banqiao cascade of 1975 failed sixty-two dams, but the killing was done by two large ones; the tailings dam failures of recent decades, from Brazil's Brumadinho to Bulgaria's Ajka, are a reminder that the most dangerous dams are often the ones nobody calls dams at all.

There is, finally, the question the aging fleet poses most sharply: what to do with the dams whose costs now exceed their value. Some are refurbished; some are relicensed with new conditions; and some, increasingly, are removed. The Klamath River removals of 2024 - four dams taken down in the largest dam-removal project in American history - were driven by exactly this calculus: the dams were old, their power small, their cost to the river's salmon immense, and removal cheaper than the upgrades a new licence would have demanded [C52] [C53]. The dam age built the fleet; the present age is deciding, dam by dam, what to keep. That decision, and the return of the rivers it frees, is where this book ends.

The economics of aging are as important as the engineering. Refurbishing a dam is expensive and unglamorous; the money competes with new construction that brings power and headlines. Yet the refurbished dam is almost always the cheapest capacity available - the structure already exists, the reservoir already full, the grid connection already built. The uprating and refurbishment of the mid-century fleet, and the safety work that comes with it, is the least visible and most consequential work in modern hydropower. The giants of the dam age are now old enough that their continued service is a decision, renewed with every inspection and every repair budget. The walls hold. Keeping them holding is the industry's quiet, permanent task.

PART XII - THE RECKONING

Chapter 29 - The rivals

For most of its history, hydropower's competition was fuel: coal first, then nuclear, then gas, each challenging the dam as the cheapest way to make bulk electricity. Hydropower held its ground on cost where the sites were good, but conceded the ground where they were not - you cannot build a dam wherever the customer is, and the twentieth century's grids were largely built, in the end, on fuel that could be shipped. The twenty-first century has changed the rivalry entirely. The new competitors are wind and solar, and they do not merely undercut the dam's energy; they change what the grid wants from every other machine on it.

The cost reversal has been extraordinary. Solar photovoltaics and onshore wind are now, on a levelised basis, among the cheapest sources of new electricity ever built; the International Renewable Energy Agency's 2024 cost report found that 91 per cent of newly commissioned utility-scale renewable projects delivered electricity below the cost of new fossil generation [C10]. Hydropower, with its bespoke civil works, cannot match the learning curves of manufactured technologies that halve in cost as their factories scale. On pure kilowatt-hour price for new capacity, the dam has largely lost the argument to the panel and the turbine.

And yet the same revolution that undercuts hydropower's energy has made its other product - flexibility - the scarcest commodity on the grid. Wind and solar generate when the weather says, and a grid built on them needs, more than anything, machines that generate when the weather does not: to fill the evening ramp, to cover the windless week, to hold the system steady second by second. The International Hydropower Association's central argument is precisely this: hydropower delivers the flexibility, storage and resilience that high-renewables grids cannot do without, and no country has approached a fully renewable electricity system without it [C48][C7b]. The rival and the complement are the same machine, seen from two ends.

The honest comparison, then, is not dam against panel but system against system. Solar plus batteries can do much of what hydro does, and the cost of batteries is falling as fast as the cost of panels did; but pumped storage still holds more than 90 per cent of the world's installed energy storage, and the storage a grid needs over days and seasons - not just hours - remains hydropower's natural ground [C8]. The realistic future, in most of the world, is not one technology displacing the other but the two interlocking: wind and sun for cheap energy, water for storage and stability. The dam that the fuel age outflanked has become the renewable age's necessary partner.

The geography of the rivalry matters as much as the economics. In the rich world, the question is rarely whether to build new dams but how to value the ones that exist - and the answer has been to pay them for what the new grid needs. Pumped storage is booming; old plants are being refurbished and uprated; grid operators pay for the capacity and flexibility hydro has always provided free. In the developing world, the calculus is different: the need for energy is vast, the unbuilt sites remain, and the alternative to a new dam is often a new coal plant. The Chinese-built dams of Africa and Asia are, in large part, the world’s answer to that calculus - a continuing expansion of the fleet where the resources and the need coincide, financed by the one country still building dams at scale [C51][C51b].

What this means for new dams is narrower but real. The age of building dams for bulk energy alone is largely over in the rich world, where the best sites are taken and the alternatives are cheap; it continues in the developing world, where the sites remain and the need is greatest. But the age of hydropower as a system resource - storage, flexibility, resilience - is only beginning. The industry’s future is written less in concrete than in the value of what its reservoirs and machines can do for a grid that increasingly runs on weather. The rivals turned out, in the end, to be the dam’s best customers.

Chapter 30 - Taking dams down

On 2 October 2024, the last of four dams on the Klamath River, on the border of Oregon and California, came down - the largest dam-removal project in the history of the United States [C19][C52]. The removals opened more than 400 miles of salmon habitat that had been blocked for a century, and they capped a campaign that had run for decades, led by the river's tribes, who had fought for the dams' removal through regulatory hearings, corporate negotiations and one of the great environmental campaigns of the American West [C52][C21]. The Klamath is the clearest statement yet of a proposition that would have been unthinkable in the dam age: that some rivers are worth more flowing free.

The removal itself was an engineering campaign to match any construction. The dams - Copco No. 2 first, in late 2023, then Iron Gate, Copco No. 1 and J.C. Boyle through 2024 - had to be taken down without releasing the decades of sediment accumulated behind them in a single choking pulse [C52][C19]. The reservoirs were drawn down in a controlled winter release, the river given back its channel, and the exposed lakebeds - some 2,200 acres of formerly submerged land - replanted in a restoration effort that will run for years [C52]. It is one thing to breach a dam; it is another to return a river. The Klamath did both, deliberately, and the watching salmon found their way home within months.

The case for removal is, at bottom, the same accounting that built the dams, run in reverse. The four Klamath dams were old, their combined power modest, their cost to the river's salmon immense. When their operator, PacifiCorp, faced relicensing, the upgrades required - fish passage, water quality, flow - would have cost more than the dams were worth. Removal was cheaper. A settlement agreement capped the cost at 450 million dollars, shared between a surcharge on PacifiCorp's customers and a California bond measure, and the dams came down on budget and ahead of schedule [C53][C19]. The lesson the industry absorbed is uncomfortable and liberating at once: a dam is not forever, and when its costs exceed its value, the river can be given back.

The Klamath is the largest removal, but not the only one, and the movement behind it is the mirror image of the movement that built the dams. Across the United States and Europe, hundreds of small, obsolete dams - mill dams, defunct power dams, structures whose purposes have vanished - have been removed over the past three decades, and the rivers have responded with a speed that has surprised even the ecologists. Salmon and steelhead return within seasons; sediment rebuilds estuaries; water temperatures fall. The removals are, in a sense, the final chapter of the river's interruption: the proof that the interruption was a choice, and that it can be unchosen.

The politics of removal are as instructive as the engineering. The Klamath dams came down not because the dams failed but because the coalition for the river outlasted the coalition for the dams - tribes, fishermen and conservationists pressing a case across decades, against a utility for whom removal had become, eventually, the rational choice. It is the dam age's central argument, settled by persistence: the question of who may impose costs on whom for a collective benefit, asked this time of the dams rather than by them. The answer the Klamath gave is that the accounting can be reopened, and the river can win.

It is important to be precise about what removal is not. It is not the end of the dam age - the world's great dams, the Hoovers and Itaipus and Three Gorges, will be maintained for as long as their power and their storage matter, which is to say indefinitely. Removal is the discipline of the fleet's margin: the sorting of structures into those worth keeping and those worth removing, on a case-by-case accounting that the dam age never performed. The Klamath matters less for the concrete removed than for the principle established - that a dam's permanence is a decision, revisited as the costs and values change, and not a fate. The river, given the chance, does the rest.

Chapter 31 - Conclusion: the long arc

This book began with a line of arithmetic: power equals the density of water times gravity times flow times fall. Everything that followed has been the working out of that equation into a world. The water's fall had to be made, so the dam was built; the fall had to be caught, so the turbine was invented; the power had to travel, so the wire was strung; and the whole apparatus, once assembled, proved so useful that the twentieth century built it on the world's great rivers. The hydroelectric dam is what happened when humanity learned to turn geography into electricity.

The story's first truth is that the dam is a double structure, wall and machine at once, and that its two natures pull in different directions. As a machine it is nearly perfect: it converts the water's energy at efficiencies no combustion technology approaches, burns nothing and asks for almost nothing once built, and answers the grid's needs with a speed few other machines can match. As a wall it is an imposition: it drowns valleys, interrupts rivers, displaces communities, and carries a small, irreducible risk of catastrophic failure. The dam age was, at bottom, a long argument about whether the machine justifies the wall - and the argument was never really about engineering. It was about who decides, who pays, and who benefits.

The second truth is that the dam's economics are the economics of patience, and they select their owners accordingly. The dam costs everything up front and pays back over a century; it cannot be manufactured, only built, once, in a particular valley; and its revenue is rain, deferred and unpromised. These are not the economics of private appetite but of the state, and so the history of hydropower is a history of public power - the New Deal valleys, the Soviet cascades, the post-colonial development dams, the Chinese century. Where the state has led, the dams have been built; where it has withdrawn, the industry has had to learn to sell not energy but readiness - the flexibility, storage and insurance that the new grids of wind and sun cannot do without.

The third truth is that the dam's costs were real and were, for a long time, hidden. The drowned valleys and displaced millions, the interrupted rivers and starving deltas, the warming reservoirs - none of it was secret, but all of it was, for decades, left out of the accounting, because the people who paid were not the people who profited. The reckoning, when it came, changed the industry: the World Commission on Dams, the retreat of the financiers, the rise of the environmental flows and the sustainability standards, and finally the removals. The Klamath's dams came down not because they failed but because the world learned to ask, of a dam as of anything else, whether it was worth what it cost - and to accept the answer.

The last truth is the one on which the machine depends: the water cycle continues. Rain still falls on the highlands and runs to the sea, and the fall is still free, and the machine that catches it is still among the most efficient prime movers humanity has. The dam age - the century of building - is largely over in the rich world and continuing in the poor one; but the dams that exist are among the most valuable energy assets on earth, and the work of the next century is to keep them, renew them, and use them well. The unbuilt rivers will be argued over; the aging fleet will be refurbished or, at the margins, removed; and the falling water will go on driving the grid as it has for a hundred and fifty years.

A dam is a bet that the future will want what the river offers. The twentieth century placed that bet at a scale never matched before or since, and it won - the lights are on, the floods are tamed, the power is among the cheapest ever sold. The twenty-first is learning the bet’s price and deciding, river by river and dam by dam, what it is willing to pay. That is the long arc of the hydroelectric dam: from the wheel in the stream to the wall in the canyon to the question, asked at last honestly, of what the water’s fall is worth - and to whom.

Chapter 32 - A hydroelectric glossary

ARCH DAM. A dam curved upstream like a bow, handing the water’s load to the rock of the valley walls. Thin but site-fussy: it needs a narrow canyon of sound rock. (Chapter 7)

BLACK START. The ability of a power station to restart with no outside power and begin re-energising a blacked-out grid. Hydroelectric units, needing almost nothing to start, are the grid’s natural emergency services. (Chapter 11)

BUTTRESS DAM. A dam of a sloping concrete deck propped by a row of buttresses - structure instead of mass. A favourite of the early twentieth century, when cement was dear and labour cheap. (Chapter 7)

CAPACITY FACTOR. The ratio between the energy a plant actually generates in a year and what it could have generated running flat out. Hydropower’s global figure is about 36 per cent - a number that conceals a wide range between steady run-of-river plants and storage dams held for peaks. (Chapter 1)

COFFERDAM. A temporary dam of earth or rock that fences off a construction site so the riverbed can be pumped dry and the real dam built. (Chapter 8)

DIVERSION. In plant types, a run-of-river station that channels part of a river through a canal or penstock without large storage. In construction, the temporary rerouting of a river around the worksite, usually through tunnels. (Chapters 2, 8)

DRAFT TUBE. The flared passage beneath a reaction turbine that slows the leaving water and recovers a last increment of its energy. (Chapter 12)

EMBANKMENT DAM. A dam built of the valley itself - earth or rock fill around a watertight core of clay. The most common type on earth, and the one that cannot be overtopped. (Chapter 7)

FRANCIS TURBINE. The most widely used water turbine in the world: water enters around the runner’s full circumference, flows inward and leaves along the shaft. Suited to middling heads and generous flows - the classic dam site. Developed by James Francis in 1849. (Chapters 4, 10)

GOVERNOR. The turbine’s reflex system: the mechanism - now digital - that opens and closes the wicket gates second by second to hold the machine at synchronous speed as the grid’s demand swings. (Chapter 11)

GRAVITY DAM. A dam that answers the water with weight: a great triangle of concrete or masonry too heavy to overturn or slide. Grand Coulee and Three Gorges are the type at its largest. (Chapter 7)

GROUT CURTAIN. A barrier of cement grout injected into the rock beneath a dam's upstream edge to choke the seepage paths under the foundation. (Chapter 7)

HEAD. The vertical distance the water falls to the turbines - the multiplier in the power equation. Head is usually made, not found: the dam is the device that manufactures it. (Chapter 1)

IMPOUNDMENT. The classic hydroelectric plant: a dam across a valley storing a reservoir, whose water is released through the turbines on the operator's schedule. The most common form of large hydropower. (Chapter 2)

KAPLAN TURBINE. A propeller turbine whose blades pivot as the flow changes, keeping it efficient on low-head, high-flow rivers. Developed by Viktor Kaplan in the early twentieth century. (Chapters 4, 10)

PELTON WHEEL. An impulse turbine: a runner of paired buckets struck by jets of water, operating in air. Suited to high heads and modest flows - the mountain site. Patented by Lester Pelton in 1880. (Chapters 4, 10)

PENSTOCK. The closed pipe or tunnel that carries water down the full head of the site to the turbines, trading altitude for pressure. (Chapter 12)

PUMPED STORAGE. The water battery: two reservoirs at different heights, water pumped uphill when power is cheap and released through the turbines when it is dear. More than 90 per cent of the world's installed energy storage. (Chapters 2, 19)

ROLLER-COMPACTED CONCRETE (RCC). A very dry concrete mix placed with road-building machinery - dump trucks, bulldozers, rollers - in continuous layers rather than blocks. Faster and cooler than conventional placement; first used on a dam in Japan in 1980. (Chapter 8)

RUN-OF-RIVER. A hydroelectric plant with little or no storage, generating what the river happens to carry. Cheap energy, sold on the river's timing rather than the grid's. (Chapter 2)

SCROLL CASE. The spiral casing that wraps a Francis runner and distributes the incoming water evenly to the wicket gates. (Chapter 12)

SPECIFIC SPEED. The single number, computed from head, flow and rotation speed, that sorts any site toward its proper turbine family - Pelton, Francis or Kaplan. (Chapters 4, 10)

SPILLWAY. The dam's safety valve: the channel that passes the floods the reservoir cannot store, keeping water from overtopping the wall. The service spillway is gated; the emergency spillway is the ungated last resort. (Chapter 9)

SURGE TANK. An open standpipe connected to a penstock that lets the pressure surge of a suddenly closed gate rise harmlessly as a column of water - the shock absorber of the waterway. (Chapter 12)

TAILRACE. The channel that returns water from the turbines to the river below the plant. (Chapter 12)

UPRATING. Raising an existing plant’s capacity by renewing its machines - new runners, rewound generators - without rebuilding the dam. The cheapest new megawatts in the industry. (Chapter 20)

WATER HAMMER. The pressure surge hurled back up a penstock when the turbine gates close suddenly and the moving column of water has nowhere to go. (Chapter 12)

WICKET GATES. The ring of adjustable guide vanes around a reaction turbine’s runner, opened and closed by the governor to control the flow onto the blades. (Chapter 12)