

Tunneling Through the Cost Barrier

Improving mindware — Optimize without compromise — More costs less — Seeing the obvious sooner — Big pipes, small pumps — Optimizing the system — Like eating a lobster — Thinking backward — Doing things in the right order — Solving for pattern

THE EXAMPLES OF THE HYPERCAR, ADVANCED INDUSTRIAL AND MATERIALS techniques, and green buildings all demonstrate that design is really just applied foresight. It's what you do now carefully and responsibly to achieve what you want later.

By the time the design for most human artifacts is completed but before they have actually been built, about 80–90 percent of their life-cycle economic and ecological costs have already been made inevitable.¹ In a typical building, efficiency expert Joseph Romm explains, “Although up-front building and design costs may represent only a fraction of the building’s life-cycle costs, when just 1 percent of a project’s up-front costs are spent, up to 70 percent of its life-cycle costs may already be committed. When 7 percent of project costs are spent, up to 85 percent of life-cycle costs have been committed.”² That first one percent is critical because, as the design adage has it, “All the really important mistakes are made on the first day.” This chapter presents ways to think differently — to use a different design mentality — on that first day.

We can make no better higher-leverage investments for the future than improving the quality of designers’ “mindware” — assets that, unlike physical ones, don’t depreciate but, rather, ripen with age and experience. Senior mechanical engineer Eng Lock Lee offers the following example. A typical colleague may specify nearly \$3 million worth of heating, ventilating, and air-conditioning (HVAC) equipment every year — enough to raise a utility’s summer peak load by a megawatt. Producing and delivering that extra megawatt conventionally requires the utility to invest several million dollars in infrastructure. If better

engineering education were ultimately responsible for the equipment's being made 20–50 percent more efficient (a reasonably attainable and usually conservative goal), then over a 30-year engineering career, the utility would avoid about \$6–15 million in present-valued investments *per brain*, without taking into account any of the savings in operating energy or pollution. This returns at least a hundred to a thousand times the extra cost of that better education. The savings would cost even less if good practitioners disseminated their improved practices through professional discourse, mentoring, or competition, so that educating just one engineer could influence many more. In addition, a good engineer's lifetime designs can improve comfort for perhaps 65,000 office workers, whose 30-year present-valued salary totals about \$36 billion. If increasing their comfort will increase their productivity on the lines suggested by the evidence mentioned in chapter 5,³ then society can gain perhaps a million times more benefit than the additional cost of the better engineering education.


Many architects, engineers, and other designers, however, are not being well taught. J. Baldwin, long the technology editor of *Whole Earth Review*, was told on his first day in design school that “design is the art of compromise.” Design, he was instructed, means choosing the least unsatisfactory trade-offs between many desirable but incompatible goals. He believed that this formulation described “a political technique masquerading as a design process,” and he realized it was wrong. His inspiration came as he gazed out the classroom window and saw a pelican catching a fish. For the past 3.8 billion years or so, nature has been running a successful design laboratory in which everything is continually improved and rigorously retested. The result, life, is what works. Whatever doesn't work gets recalled by the Manufacturer. Every naturalist knows from observation that nature does not compromise; nature optimizes. A pelican, nearing perfection (for now) after some 90 million years of development, is not a compromise between a seagull and a crow. It is the best possible pelican.

A pelican, however, is not optimized within a vacuum. It exists in an ecosystem, and each part of that ecosystem, in turn, is optimized in coevolution with the pelican. A change in the pelican or in any aspect of its ecosystem could have widespread ramifications throughout the system, because all its elements are coevolving to work optimally together. For the same reason, an engineer can't design an optimal fan except as

an integral part of its surrounding cooling system, nor an optimal cooling system without integration into the building around it, nor an optimal building without integration into its site, neighborhood, climate, and culture. The greater the degree to which the components of a system are optimized together, the more the trade-offs and compromises that seem inevitable at the individual component level becomes unnecessary. These processes create synergies and felicities for the entire system. And this in turn exposes a core economic assumption as a myth.

TUNNELING THROUGH THE COST BARRIER


Economic dogma holds that the more of a resource you save, the more you will have to pay for each increment of saving. That may be true if each increment is achieved in the same way as the last. However, if done well, saving a large amount of energy or resources often costs *less* than saving a small amount.⁴ This assertion sounds impossible, and indeed, most economic theorists can “prove” it won’t work. Blissfully unaware of economic theory, however, intelligent engineers put it into practice every working day as part of an approach called *whole-system engineering*.



If you build a house, you’ll be told that thicker insulation, better windows, and more efficient appliances all cost more than the normal, less efficient versions. If you build a car, you’ll be told that lighter materials and more efficient propulsion systems are more expensive options.

These statements are often true — but at the level of single components considered in isolation. On the cost-versus-savings graph shown on page 113, as you save more energy (that is, as you move from the lower left end of the curve toward the right), the cost of saving the next unit of energy initially rises more and more steeply. This is called “diminishing returns.” When you’ve struggled up to the limit of cost-effectiveness, you should stop additional outlays of money, because they’re no longer justified by their results. This part of the curve illustrates the common principle that better usually costs more, a principle that has taken a death grip on our consciousness.

Actual engineering practice, however, presents a different possibility. Only recently noticed is an additional part of the curve further to the right (see the graph below): There, saving even more energy can often “tunnel through the cost barrier,” making the cost come *down* and the return on investment go up. When intelligent engineering and design are brought into play, big savings often cost even less *up front* than small or zero savings. Thick enough insulation and good enough windows can eliminate the need for a furnace, which represents an investment of more capital than those efficiency measures cost. Better appliances help eliminate the cooling system, too, saving even more capital cost. Similarly, a lighter, more aerodynamic car and a more efficient drive system work together to launch a spiral of decreasing weight, complexity, and cost. The only moderately more efficient house and car do cost more to build, but when designed as whole systems, the *superefficient* house and car can often cost less than the original, unimproved versions.



There are two main ways to achieve this more-for-less result. The first is to integrate the design of an entire package of measures, so that each measure achieves multiple benefits, such as savings on both energy *and* equipment costs.⁵ The second method is to piggyback on improvements being made anyway for other reasons, such as renovation of aging equipment, renewal of deteriorating building facades, or removal of such hazards as CFCs, asbestos, and PCBs. These two practices, which can also be combined, rely not on some arcane new technology but on well-known engineering fundamentals rigorously applied. A well-trained engineer will be guided by the following three precepts:

- The whole system should be optimized.
- All measurable benefits should be counted.
- The right steps should be taken at the right time and in the right sequence.

Most engineers would agree with these principles in the abstract but have actually been trained to do something different. Perhaps the scheme is too *simple*. (As broadcaster Edward R. Murrow once remarked, “The obscure we always see sooner or later; the obvious always seems to take a little longer.”) Tunneling through the cost barrier requires not a change in what we know but a shift of what we already know into new patterns — patterns that can lead to innovations as rich and diverse as the Hypercar, the superefficient passive building, the New Urbanist neighborhood. That shift can ultimately reach the scale of an industry, city, or society, but it must start at a more immediate and fine-grained level: at the building or factory, and even earlier, at their constituent systems and subsystems. This chapter addresses design at the latter level, the realm of machinery and infrastructure, while the following chapter considers the broader implications of this approach for manufacturing and industrial development.

INTEGRATING DESIGN TO CAPTURE MULTIPLE BENEFITS

Motors use three-fifths of the world’s electricity. Their largest use, at least a fifth of their total output, is pumping. Almost every factory or major building is full of huge pumps, often running around the clock. In industrial pumping, most of the motor’s energy is actually spent in fighting against friction. But friction can be reduced — indeed, nearly eliminated — at a profit by looking beyond the individual pump to the whole pumping *system* of which it is a part.

In 1997, leading American carpet maker Interface was building a factory in Shanghai. One of its industrial processes required 14 pumps. In optimizing the design, the top Western specialist firm sized those pumps to total 95 horsepower. But a fresh look by Interface/Holland's engineer Jan Schilham, applying methods learned from Singaporean efficiency expert Eng Lock Lee,⁶ cut the design's pumping power to only 7 horsepower — a 92 percent or 12-fold energy saving — while *reducing* its capital cost and improving its performance in every respect.

The new specifications required two changes in design. First, Schilham chose to deploy big pipes and small pumps instead of the original design's small pipes and big pumps. Friction falls as nearly the fifth power of pipe diameter, so making the pipes 50 percent fatter reduces their friction by 86 percent. The system then needs less pumping energy — *and* smaller pumps and motors to push against the friction. If the solution is this easy, why weren't the pipes originally specified to be big enough? Because of a small but important blind spot: Traditional optimization compares the cost of fatter pipe with only the value of the saved *pumping energy*. This comparison ignores the size, and hence the capital cost, of the *equipment* — pump, motor, motor-drive circuits, and electrical supply components — needed to combat the pipe friction. Schilham found he needn't calculate how quickly the savings could repay the extra up-front cost of the fatter pipe, because capital cost would fall more for the pumping and drive equipment than it would rise for the pipe, making the efficient system as a whole cheaper to construct.

Second, Schilham laid out the pipes first and *then* installed the equipment, in reverse order from how pumping systems are conventionally installed. Normally, equipment is put in some convenient and arbitrary spot, and the pipe fitter is then instructed to connect point A to point B. The pipe often has to go through all sorts of twists and turns to hook up equipment that's too far apart, turned the wrong way, mounted at the wrong height, and separated by other devices installed in between. The extra bends and the extra length make friction in the system about three- to sixfold higher than it should be. The pipe fitters don't mind the extra work: They're paid by the hour, they mark up the pipe and fittings, and they won't have to pay the pumps' capital or operating costs.

By laying out the pipes before placing the equipment that the pipes connect, Schilham was able to make the pipes short and straight rather than long and crooked. That enabled him to exploit their lower friction

by making the pumps, motors, inverters, and electricals even smaller and cheaper.

The fatter pipes and cleaner layout yielded not only 92 percent lower pumping energy at a lower total capital cost but also simpler and faster construction, less use of floor space, more reliable operation, easier maintenance, and better performance. As an added bonus, easier thermal insulation of the straighter pipes saved an additional 70 kilowatts of heat loss, enough to avoid burning about a pound of coal every two minutes, with a three-month payback.

Schilham marveled at how he and his colleagues could have overlooked such simple opportunities for decades. His redesign required, as inventor Edwin Land used to say, “not so much having a new idea as stopping having an old idea.” The old idea was to “optimize” only part of the system — the pipes — against only one parameter — pumping energy. Schilham, in contrast, optimized the *whole* system for *multiple* benefits — pumping energy expended plus capital cost saved. (He didn’t bother to value explicitly the indirect benefits mentioned, but he could have.)

Such whole-system life-cycle costing, in which all benefits are properly taken into account over the long run, is widely accepted in principle but almost always ignored in practice. Instead, single components are usually considered in isolation. Designing a window without the building, a light without the room, or a motor without the machine it drives works as badly as designing a pelican without the fish. *Optimizing components in isolation tends to pessimize the whole system* — and hence the bottom line. You can actually make a system less efficient while making each of its parts more efficient, simply by not properly linking up those components. If they’re not designed to work with one another, they’ll tend to work against one another.

A charrette improving the design of a chemical plant noticed a big pump whose function was to send fluid up a pipe. Because it had such an important task, the pump required an adjacent, identical spare pump. The designer had drawn two identical rectangles, side by side, representing the two pumps. Up out of each rectangle came a line, representing a pipe. The two lines bent at right angles, came together and joined, bent upward again, and continued on together as the common exit pipe, with a valve on each of the three sides of the T-junction. As a drawing, it was a clear enough design intention. The trouble was that it had been built exactly as drawn.

What's wrong with this picture? The primary flow, coming from the first pump 99-plus percent of the time, must nevertheless always pass through two right-angle bends and two valves. To combat that added friction, the pump, motor, motor-drive controller, and electric supply must all be larger and hence costlier, and will use more energy forever after. Instead, the designer should have drawn (and the contractor installed) the pipe from the primary pump going directly to its destination *with no bends and* (usually) *no valves*. The pipe from the backup pump, in turn, should have come up and joined the main pipe at a shallow angle, using probably just one valve. This layout may look less orderly, but it works better, makes less noise, has fewer parts to fail, offers better maintenance access, and costs less both to build and to run. It also requires less space, one or two fewer valves to buy, install, and mend when they jam or leak, and less pipe fitting.

This novel pipe layout, like Schilham's rethinking of his pumping system, requires a *change of design mentality*. Once that change happens, it tends to be irreversible. An engineer exposed to so simple and adhesive an idea is unlikely ever again to use the traditional right-angle-bends layout and skinny, twisting pipes — at least not without squirming. And that transformation in design mentality opens the mental door to others: Layout is only the first step in reducing the friction in piping systems, and friction is only one of the forces that pumps must overcome.

Traditionally poor designs often persist for generations, even centuries, because they're known to work, are convenient, are easily copied, and are seldom questioned. One story traces the standard fifty-six-and-a-half-inch U.S. rail gauge back through British railways, trams, and wagons, back for two millennia to the spacing of ruts in ancient roads built by the Romans. So if, the story concludes, you look at some modern specification and wonder what horse's ass designed it, you may be exactly right in your assessment — because those ruts were made by chariots designed to fit the back ends of two Imperial Roman Army warhorses.

Saving a lot of energy, or any other resource, at low cost is like eating a lobster. To do it successfully requires both a grasp of system anatomy and attention to detail. There are big, obvious chunks of meat in the tail and the front claws. There's also a roughly equal quantity of tasty morsels hidden in crevices, requiring skill and persistence to extract but

worth the effort. It was this “whole-lobster” approach, as described in chapter 5, that eliminated heating and cooling systems both in the Davis house and in Rocky Mountain Institute’s headquarters. Both structures, in climatic extremes ranging over 160 Fahrenheit degrees, perform as well as or better than conventional houses but cost less to build.⁷ Their success resulted from combining the right details with an important underlying principle that ignored practically every textbook’s description of how to select the basic design elements for energy-efficient buildings. That description instructs you to add more insulation, buy more heat-tight windows, and purchase more efficient appliances only to the point justified by the value of how much energy *each of those individual components* will save over time. But this is an instruction for designing a wall or a window by itself, not a house that combines them. For the whole house, it gives the wrong answer. America has \$6 trillion worth of houses whose thermal efficiency rests on a methodological design error.

The fallacy is the same one that Schilham found in pipe-diameter selection: Counting saved energy costs as the only benefit ignores the additional savings available in *capital* equipment, such as heating and cooling systems, that can be reduced or eliminated if efficiency is sufficiently increased. This avoided capital expense made the far more efficient houses cheaper to build, by reducing construction cost more than the efficiency measures increased it. In the RMI building, this involved simply a substitution of superinsulation, superwindows, and ventilation heat recovery for a heating system, including associated fuel and power supplies, vent, ductwork, plumbing, wiring, and controls.⁸ The Davis house used a more complex series of substitutions, but the net effect was the same — tunneling through the cost barrier to achieve much larger savings at negative cost.⁹ In short, neither house had heating or cooling equipment for the simplest possible reason: Each cost less to build that way.

PIGGYBACK ONTO RENOVATIONS ALREADY PLANNED

A 200,000-square-foot all-glass-and-no-windows¹⁰ curtainwall office tower near Chicago needed its twenty-year-old windows replaced because they were starting to leak as the seals failed, and its large air-conditioning systems needed renovation to renew the moving parts and replace their ozone-eating CFC refrigerant. Analysis revealed that

changing the renovation design to a whole-systems approach could dramatically improve comfort, quadruple energy efficiency, and cost about the same as normal renovations. Superwindows, deep daylighting, and efficient lights and office equipment could reduce the cooling load (except that caused by the occupants) by 85 percent. This in turn could make the replacement cooling equipment three-fourths smaller than the original system, four times as efficient, *and* \$200,000 cheaper — a sum large enough to pay for the other improvements. The annual energy bill would then fall by 75 percent, or by \$1.10 per square foot per year — at least ten times the competitive rent difference in the local market. The fourfold energy efficiency improvement would cost essentially the same as the standard renovation that was about to be done anyway (its extra cost would pay for itself in between minus five and plus nine *months*), with far better amenity, aesthetics, and rentability.¹¹ By the time America's 100,000 or so glass office buildings now ripe for such renovation have been retrofitted, another generation of roughly as many similar structures will have reached the age of rehabilitation. If the building discussed above was typical in all respects (not too bad an approximation), then redesigning the routine renovation of all big U.S. office towers in similar fashion could save about \$45 billion a year.

Reducing this project's total capital cost depended on spending the renovation money in different places than a standard rehabilitation would have — more on windows and daylighting and efficient lights, less on the downsized air-conditioning system. This required *optimizing the entire building as a system*, not value-engineering its individual components. Normal "value engineering," which is about neither value nor engineering, would have cut out the costlier windows and any other component that wasn't the cheapest possible commodity, considered in isolation. But like a squeezed balloon, the costs would then have bulged out elsewhere — in this case, as fourfold bigger air-conditioning equipment.

The key is whole-system engineering with meticulous attention to detail. Close enough attention often reveals more than just two benefits per technology. Not surprisingly, superwindows have ten engineering-economic benefits. These include radiant comfort, no under-window radiators, smaller ducts, better blocking of noise and ultraviolet rays, no condensation, better daylighting, and simpler controls. Some common technologies have even more benefits: eighteen each for premium-efficiency motors and dimming electronic ballasts, for example. Those


multiple benefits have been demonstrated in a wide range of applications.¹² They are the key to extraordinary economic performance. They make superwindows often the most important single technology for highly efficient, comfortable, and cheaper-to-construct buildings. They also make possible comprehensive motor- and lighting-system retrofits that, applied nationwide, could inexpensively save upward of half of all U.S. electricity used.¹³

TO LEAP FORWARD, THINK BACKWARD

Much of the art of engineering for advanced resource efficiency involves harnessing helpful interactions between specific measures so that, like loaves and fishes, the savings keep on multiplying. The most basic way to do this is to “think backward,” from downstream to upstream in a system. A typical industrial pumping system, for example (as illustrated below), contains so many compounding losses that about a hundred units of fossil fuel at a typical power station will deliver enough electricity to the controls and motor to deliver enough torque to the pump to deliver only ten units of flow out of the pipe — a loss factor of about tenfold.

But turn those ten-to-one compounding losses around backward, as in the drivetrain of the Hypercar, and they generate a one-to-ten compounding *saving*. That is, saving one unit of energy furthest downstream (such as by reducing flow or friction in pipes) avoids enough compounding losses from power plant to end use to save about *ten* units of fuel, cost, and pollution back at the power plant.

A TYPICAL INDUSTRIAL PUMPING SYSTEM



From the *Drivepower Technology Atlas*. Courtesy of E SOURCE, www.esource.com.

Those compounding savings represent significant economic and environmental leverage — the same principle that a Hypercar uses to multiply its reduced air and rolling resistance into big fuel savings. This compounding effect also enables each successive component, as you go back upstream, to become smaller, simpler, and cheaper. This in turn means that *downstream savings merit the greatest emphasis*. The reason is simple. In a chain of successive improvements, all the savings will multiply, so they appear all to have equal *arithmetic* importance. However, the *economic* importance of an energy-saving measure will depend on its position in the chain. Savings furthest downstream will have the greatest leverage in making the upstream *equipment* smaller, and this saves not just energy but also capital cost. Downstream savings should therefore be done first in order to save the most money.

Downstream-to-upstream thinking is thus a special case of a more general rule: Do the right things *in the right order*. For example, if you're going to retrofit your lights and your air conditioner, do the lights first so you can make the air conditioner smaller. If you did the opposite, you'd pay for more cooling capacity than you'd need after the lighting retrofit, and you'd also make the air conditioner less efficient because it would either run at part-load or cycle on and off too much. There is a similarly logical sequence for such common efficiency improvements as improving office lighting¹⁴ or providing hot-weather comfort.¹⁵ Once you've done the right things in the right order, so as to maximize their favorable interactions, you'll have very little energy use left: Successive steps will have nibbled away at it a piece at a time, with each improvement saving part of what's left after the previous steps. The arithmetic of these multiplying terms is powerful.

Efficient distribution of ventilation air, generally from the floor toward the ceiling, is another system that captures this multiple benefit of “thinking backward.” Indeed, properly designed, such an air system offers many additional advantages:

- It enables people to stay happier and healthier by eliminating toxic materials, improving thermal comfort, providing the options of individual ventilation control and even of operable windows or vents, and helping air to flow without fans by means of gravity, breezes, and other natural forces.
- It distributes the delivered fresh air more effectively to the people's bodies, and particularly to their noses.

- It minimizes friction, from downstream (grilles) to upstream (ducts, filters, silencers, fans).
- It makes the resulting smaller fans, and their controls and power supplies, more efficient, and it reoptimizes them for their new operating conditions.

Air-handling, in turn, interacts with other systems in the building: superwindows, lighting, daylighting, cooling, and a whole range of design elements. For example, smaller fans heat the air less, requiring less cooling and hence smaller fans.

The world's master of the new design mentality in fluid-handling and air-conditioning systems — the Singaporean engineer Eng Lock Lee — was trained in the same engineering principles as everyone else. He buys hardware from the same companies and looks up data in the same handbooks. Yet his designs are typically about three to ten times more efficient, deliver better services, and cost less to build. The trick is all in how he *thinks*. He wrings out friction and waste of every kind, downstream to upstream, end to end. To save land (very costly in Singapore), he untangles and compacts plant layouts so they take up less space, yet are easier to maintain. Space, money, metal, energy, time, words — he uses just the right amount of every resource, in the right place and time and manner. Every input and result is measured, nothing is guessed. Energy is used frugally, then recaptured and reused until almost nothing is left. When he was once congratulated on devising an especially clever way to use a building's outgoing air to pre-dry its incoming fresh air, using no energy and no moving parts, Mr. Lee replied: "Like Chinese cooking. Use everything. Eat the feet."

Inevitably, great engineering like Lee's is elegantly simple. Simplicity and elegant frugality are natural partners. Using less material means there is less to go wrong, less work involved, less cost, and better performance. All are products of the same design mentality. All reflect what farmer-poet Wendell Berry calls "solving for pattern" — finding solutions that are "good in all respects," solutions that improve not just the part that seems to be the problem but all parts of the system that contains it.¹⁶ As Village Homes developer Michael Corbett put it, "You know you are on the right track when your solution for one problem accidentally solves several others. You decide to minimize automobile use to conserve fossil fuels, for example, and realize that this will reduce noise, conserve land by minimizing streets and parking, multiply

opportunities for social contact, beautify the neighborhood, and make it safer for children.” Corbett was solving for pattern as Christopher Alexander teaches in his famous design text, *A Pattern Language*:¹⁷ “When you build a thing, you cannot merely build that thing in isolation, but must also repair the world around it, and within it, so that the large world at that one place becomes more coherent, and more whole; and the thing which you make takes its place in the web of nature, as you make it.”