Smart Energy

A Key Role for Computers

The miracle of the digital computer is giving us the opportunity to employ new levels of care and skill in solving our energy shortage problems

by ROBERT H. CANNON Jr.

IN WORKING out our energy shortage problems we may make more progress with finesse than we do with brute force — finesse in managing the energy we use in machines and processes, and finesse in understanding ahead of time what all the consequences are likely to be of choosing one or another course of action. Finesse in managing, finesse in understanding.

The new technology that is giving us the opportunity to employ this care and skill — this finesse — is the miracle of the digital computer.

Computers can help us get more from our energy at three levels. Level I is machines and processes, like automobile engines and glass factories. Level II is major operating systems, like regional power networks or the national air traffic control system. Level III is understanding whole segments of the economy, like transportation, or the complex of major industries involved in producing electric power from coal.

The opportunities are pervasive. For every example that space permits there are a dozen or two it does not.

LEVEL I: MACHINES AND PROCESSES

Example 1 is about smarter automobile engines. There are many ways to improve automobile fuel economy; one where finesse can really count is computer control of ignition.

The amount of power that gasoline delivers when you burn it in your engine depends on the particular crank angle (A_p) at which *peak pressure* is reached — at which maximum force is exerted on the piston by the exploding fuel mixture (just as pushing a child high in his swing depends on pushing at just the right time in each cycle). Angle A_p depends in turn on the *earlier* crank angle (A_s) at which you fire the spark plug to ignite the fuel. For good performance, peak pressure occurs after dead center (as in the child's swing), and ignition before dead center. But it is angle A_p that determines power; angle A_s merely controls A_p.

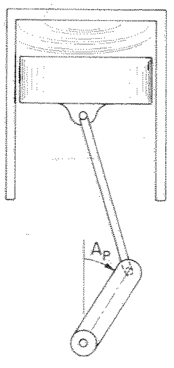

It turns out that (as with the child's swing) the angle A_p for best power is nearly the same for all operating conditions. The spark angle to produce it, however, varies greatly with temperature, altitude, fuel content, and humidity, as well as with engine speed. In automobiles, from the beginning, the right spark angle has had to be guessed at — first manually, using the little lever under the steering wheel on the Model T, and since then by using an automatic adjustment with engine speed, which guesses wrong a lot of the time. That's the problem.

The solution is to quit guessing at spark angle and to control A_p itself, using a simple sensor to measure angle A_p directly and a computer to continually adjust spark angle A_s until you get the A_p you want. This is what Professor David Powell and his colleagues at Stanford University have done. By measuring and controlling A_p directly, Powell's system always fires at exactly the right time, hot or cold, rain or shine, mountain or seashore, fast or slow, man or boy.

What would such a system cost? If you already have electronic ignition, you just add a peak-pressure sensor and a tiny computer chip that costs a dollar or so. The sensor Powell has designed fits inside a spark plug.

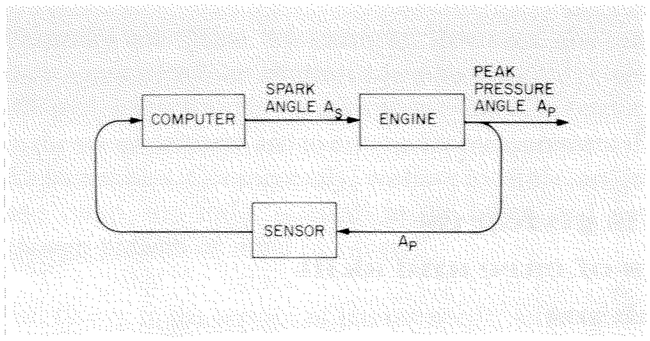

An inexpensive solution to the problem of getting maximum power from exploding fuel in an automobile engine — a sensor to measure crank angle A_P so a computer can continually adjust spark angle A_S.

(Notice that you don't need to know the pressure itself, only the crank angle at which pressure is maximum.) So, replacing one of your spark plugs with one of Powell's would perhaps cost less than five dollars. The whole system can then be very cheap.

Another part of the auto-engine efficiency story is, of course, fuel/air ratio. How lean a mixture can you burn with smooth performance? Here again a tiny computer can provide the precision control needed, and many laboratories are working on one. (In this case it is oxygen in the exhaust that is sensed, and the zirconia sensor needed will cost a little more than Powell's pressure sensor.) The payoff is not only the obvious reduction in gasoline needed when you use a lean mixture, but a major reduction in emissions as well. Precision control lets you operate safely very close to the lean limit, where emissions are low. Without such precise control you must use a much richer mixture (making emissions several times higher than necessary) in order to be sure that excursions around the nominal operating point never get you below the lean limit, where you will encounter nasty misfiring.

So, with two little computers you can provide yourself with an engine that is not only operating very efficiently when you get it back from a garage tune-up, but operating efficiently all the time, under all conditions. I expect to see Detroit including these computers as standard equipment before long.

Now you don't produce successful systems, even simple ones like these, by just thinking about them. First you make many measurements — many experiments — until you understand the nature of the system you're trying to improve. Next you try to evolve a quantitative description of how things seem to behave — a *model* of the system. Then you test the model with more experiments to be sure it is valid. That's the experimental method.

In addition, before you build your whole new system, you can do an inexpensive intermediate step, called *simulation*. This consists of putting your mathematical model on a general-purpose computer so that it behaves like the real system. You can then do preliminary experiments on this computer — much faster and much cheaper than you can do real experiments. And from them you can predict what will happen later with the real physical system. This is finesse in *understanding*. (Powell did this extensively for the automobile engine.) Then finally, with the detailed understanding you have thus developed, you can build the actual control system with much more confidence.

These two concepts — (1) the importance of experiment, and (2) the use of one kind of computer to manage energy (in machines and processes and operations) and another to model and *simulate* that performance (to gain understanding about it) — are common to all the examples we'll be discussing here, particularly as the systems grow more complex.

Example 2 is glass-making, a typical energy-intensive industry where, in fact, energy is more expensive than the raw materials.

Glass is made in a large vat. Across the top of the vat heated air flows from many orifices to melt, in just the right way, the incoming batch of constituents that will form the glass. At the bottom, molten glass is drawn off and sent to molds to make drinking glasses, or windshields, or building glass, or the myriad other glass things in our daily life.

The efficiency of the process depends critically on the distribution of temperature *all along* the top of the vat. And this whole temperature *distribution* is what must be controlled. To do this, many air flows must be controlled separately in the right proportions, based on many temperature measurements throughout the vat.

What has led to improved efficiency here is successful three-dimensional simulation, on a general purpose computer, of the heat and liquid flows and currents in the cauldronous glass-making vat. From this simulation has come the design of another, operational, computer to control air and fuel flows to optimize those distributed temperatures. This analysis was done by Professor Eugene Goodson's team at Purdue University. The computer uses finite element techniques to solve simultaneously a set of partial differential equations — a new and important achievement.

The glass industry produces about 13 million tons of glass per year. The 10 to 15 percent improvement in overall efficiency expected from computer control could save about half a million gallons of oil per day.

Example 3 is deep ocean mining. As we reach farther and farther for our energy resources, we're going to have to operate with machines where men cannot go — searching, probing, extracting, harvesting — all by remote control. A current example is the *Glomar Challenger,* a ship designed to drill holes in the bottom of the sea, to learn how the sea floor was formed, to learn more about the plate tectonics puzzle, and to learn where minerals and resources may be.

This means controlling, from the surface, a submerged drill string three or four miles long, threading a needle underwater by pushing on the other end of the thread about four miles away. The really interesting problem here is to reenter the same hole after you've broken a drill bit — or after you've been away for a week. Without this reentry capability, you can drill only as deep as one bit will take you. With reentry, there's no such limit.

The *Glomar Challenger* team solved this central problem with a sonar sensor on the end of the drill string, and a managing computer to operate the ship so it would move in just the right way to drop the drill string into the hole. Around the hole there are three sonar transponders on a 15-foot cone. The computer compares the three signals to determine where the drill bit is in relation to the hole. The computer on the ship must have built into it a full understanding of the dynamic response of the ship and its four-mile-long piece of spaghetti, so it can figure out how to have the ship move around on the surface to just drop the drill bit into the hole.

The computer can do it. Making many reentries, the *Glomar Challenger* recently drilled a hole 1,000 feet into the sea bottom in 15,000 feet of water.

Harvesting oil in water that deep is another interesting challenge. In fact, there are many things to be done on the ocean floor. That's the new frontier — and we'd better learn to operate there, for energy, for food, and for resources of all kinds.

Sometimes, unmanned, untethered vehicles will have to do things for us at great depths, in contact with people on the surface only through underwater sound communication. This adds to the problem of manipulation at a distance the problem of having to wait as long as eight seconds after you send the command before you see that command obeyed — because sound signals take four seconds to travel from the surface to the bottom and from the bottom to the surface. It is like trying to drive your car with the arrangement that after you turn the steering wheel there will be an eight-second delay before the wheels turn. Experiments show this is truly impossible.

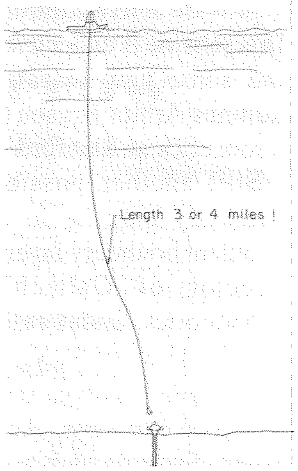

The solution to the time-delay problem has to be to place tactical control in a local computer. The commands that you send must be strategic ones only. You cannot say, "Turn left — quick!" but "Go around that rock," or maybe "Go from point A to point B and don't run into any rocks." The only thing that makes this conceivable is, again, the miracle of the minicomputer, which has such enormous amounts of computing power in a tiny space, and *in situ.*

Interestingly, the ocean-bottom remote control problem is apt to get solved by the space program. William Whitney and his team at JPL have the problem of controlling an exploratory vehicle on Mars. The communication-time delay (in a radio signal, in this case) from here to Mars and back is many minutes. As on the ocean floor, the task is to probe, manipulate, and examine. One difference is that on Mars the light is going to be better.

There are, of course, many energy-related requirements for fail-safe remote manipulation. Handling toxic materials such as plutonium is one. Mining coal is another.

Example 4 is aircraft efficiency. About 10 percent of all transportation oil is used by aircraft. That percentage will probably go up. So reducing aircraft drag can save important amounts of oil.

Part of airplane drag comes from the large vertical tail. Now you do need such a tail for feathering — so the plane will be stable and always point in the direction you want to go. One way to reduce drag is to replace the large vertical tail with a much smaller tail; but if you do only that, you get an unpleasant fish-motion ride. On the other hand, you can replace the large tail by a small tail, plus a sensor of yaw motion, plus a computer to control that motion by operating the rudder in just the right time sequence to give you a very smooth ride indeed. This system is known as a yaw damper, and there is now one on most commercial aircraft — to make your ride more comfortable and to save some fuel by not having to add that huge tail.

Another way to reduce aircraft weight, and hence fuel consumption, is by controlling the flutter of aircraft wings. The wings can be made amply "stiff" with considerably less structural weight if you use feedback to provide active control of their up-and-down motion. This scheme is under test for military aircraft at Boeing, and some new research is under way by Caltech alumnus Arthur Bryson (MS '49, PhD '51), who heads aeronautics at Stanford.

Still another possibility for improving aircraft efficiency may lie in controlling, by feedback, the turbulent boundary layer on the wing. Hans Liepmann, director of Caltech's Graduate Aeronautical Laboratories, is currently addressing this possibility.

Example 5 is about simulation in aeronautics and the energy it can save. Aeronautics is where simulation began. Systems such as the yaw damper and automatic pilot are always tested extensively, using computers to simulate the aircraft's dynamic response under the whole range of flight conditions, before the planes are ever flown.

Another example of simulation that saves energy is the program for familiarizing commercial airline pilots with airports they have not landed in recently. Before really good simulation was available, flight crews made large numbers of practice landings with full-sized, empty commercial aircraft. The economic penalty of tying up the capital investment in an empty airplane, compounded by the rapidly rising cost of jet fuel, motivated vigorous development of the art of simulating these landings so that pilots can practice without limit in a special room on the ground, with all the realism of actual flight.

The result is pretty impressive. What the computer does is generate, for the pilot, a three-dimensional picture of the airport, complete with approach lights, runway and taxi markings, and even lighting from the airport terminal. The system is so realistic, including the motions of turns, banks, and runway touchdown, that the FAA has approved it for use in all pilot proficiency training. United Airlines, for example, has installed the system on its 737 and DC-8 flight simulators in Denver. This one company says it saves about 55 million gallons of jet fuel each year with its flight simulator pilot training.

Here again is computer finesse at work. One of the designing geniuses behind this particular computer simulation and display is Ivan Sutherland, who has just joined Caltech as professor of computer science.

LEVEL II: LARGE OPERATIONAL SYSTEMS

Example 6 is the National Air Traffic Control System. Air traffic control began with the Berlin Airlift where, in zero-visibility weather, controllers on the ground used radar to talk the pilots down to the runway. Since those days much operational development and much research have been done. Recently the FAA installed it in its Third Generation Air Traffic Control System, a major computer achievement now in use throughout the country.

Before the new computer was in, each controller had to keep track of all the blips of light on his radar screen and figure out which aircraft was which — and remember them all as he directed their movements. As the number of aircraft mounted, pressure on the controllers also mounted until the problem began to reach a crisis level around 1969. There was anxiety for safety; and droves of aircraft were kept orbiting airports for hours (burning fuel the while).

What the new computer does is tag each aircraft on the controller's radar screen. It provides a picture not only of the raw radar return but also of much other information, including each aircraft's identity, altitude, speed, and heading. With this information the controller is relieved from routine pressures and is free to lay out his strategy and manage the group of aircraft in his area with certainty and safety — and minimum expenditure of fuel.

In concert with surer control of aircraft in the vicinity of the airport, en route control is also being improved now to take advantage of computer-maintained, up-to-the-minute knowledge of where the winds aloft are, and of where thunderstorms are, so that safe, minimum-fuel routes can be laid out (with safe spacing between aircraft). That's why, when you fly across the country, you may take any of several different routes, depending on what the computer has determined is the most fuel-efficient one.

This capability will continue to improve. There will come a time fairly soon when commercial aircraft will be controlled gate-to-gate. An aircraft will not leave the passenger gate until its entire path all the way to the passenger gate at its other terminal is cleared. It will be controlled on the runway, on takeoff, throughout the flight, and all the way to the next gate, to minimize fuel consumption. This, finally, is the full *systems* approach, completely eliminating the fuel waste due to flying under adverse conditions, due to waiting in the landing pattern, and due to waiting on the ground with the engine running.

The ultimate air traffic control system will probably use satellites instead of radar. Three satellites (at 22,000 miles altitude) can pinpoint simultaneously the location of all the aircraft over the country to about 50 yards and can provide all the information in a single computer (with fail-safe redundancy added, of course,

continued on page 28

probably with a backup computer in another part of the country).

Example 7 concerns goods movement. Goods movement by fuel-efficient combinations, like rail plus truck, is currently inventory *in*efficient because of great uncertainty as to when a shipment will arrive. This occurs because shipments, and indeed whole freight cars, are lost track of for days, and often weeks. The problem will be remedied only when computers begin monitoring and providing efficient *continual* management of each shipment from its origin to its destination — managing loading and transferring and train makeup and train movement. The first step was to paint a code on every freight car. This has now been done: You can now see, on every freight car that passes, a series of colored horizontal lines painted on its side; and code readers are deployed throughout the rail network. These constitute the "sensor" input to the computer. Dozens of independent railroads are cooperating to take the next step. They have a long way to go.

Example 8 addresses electric power consumption. Power companies must provide the region they serve with electricity by drawing on many sources — steam turbines using oil or coal, gas turbines for peak loading, hydro (waterfalls), perhaps some day also solar collectors, windmills, and others. And they must make prudent use of energy storage if more than a few percent is to come from wind or solar power.

To see how the computer saves oil, consider a simple hydro-plus-oil system. With reasonable manual control, the part of the load provided by oil is found inevitably to show high overloads at peak hours. These must be met with fuel-expensive oil-burning auxiliary generators. But with a computer optimizing the mix, oil need provide only the *average* sector of the energy, because the computer, knowing with precision about all parts of the system, can draw on hydro at just the right time to avoid the inefficient oil-based peaking.

Another way to avoid peaking is by smoothing out the demand — by "load leveling." There is a particularly good idea called "load shedding," which began with individual industrial companies noting that the charge for electricity use is scheduled in steps based on the amount used in each half hour interval. The companies contrived to have a computer turn off non-essential equipment, by priority, just long enough each half hour to avoid the next step in cost. Two companies that are friends of Caltech pioneered this idea — Thompson-Ramo-Wooldridge and Lockheed Electronics — and others are probably now in the market.

Here's how it works at TRW's Tapco Plant near Cleveland: The computer has been told that air conditioning has lowest priority; it can be turned off for five minutes of every half hour and will hardly be missed. Next lowest priority is the air compressor, because there is an accumulator that can be turned off for a short while. Then come certain presses and furnaces in this particular factory. So the computer cuts things in and out by priority whenever it predicts that the maximum allowable demand would otherwise be exceeded.

How do the power companies like this idea? In the city of Burbank, where Lockheed employs the scheme, the utility company found two problems from the city's viewpoint. First, many little companies in Burbank that couldn't afford the computer were not getting the same advantages from load shedding as the big companies. Second, everybody was shutting off their equipment all at once at the end of each half hour. However, the savings inherent in the Lockheed scheme were obviously impressive. So the City of Burbank Utility Company is in the process of putting together a deal that it will make with Lockheed, as well as several film studios, NBC, and numerous manufacturers in Burbank. The city plans to propose a rate reduction for each plant, in exchange for which the city, from *its* central computer, may cut out certain of that plant's equipment already agreed to ahead of time — and for only the length of time that plant has agreed to. The signal for doing all this can move through the existing 60-cycle power line, so very little additional equipment needs to be installed.

With the resulting savings in peak power requirement, the City of Burbank expects to be able to delay purchase of costly new generator equipment by as much as seven years — a very important saving.

LEVEL III: UNDERSTANDING ECONOMIC RELATIONS

At Level III the computer has only one role — to provide better understanding of economic relations and of how energy can be most prudently used. Computer "management" of human affairs is antithetical to a free society.

Example 9 concerns strategies for energy-efficient transportation. This concerns a 1971 study done at the U.S. Department of Transportation's Transportation Systems Center in Cambridge, Massachusetts, to improve everyone's understanding of which strategies would significantly reduce energy consumption by transportation — and which would not. (Transportation uses about two-thirds of our petroleum.) The study quickly found, of course, that highway vehicles are the major culprits. After looking at over 50 possible strategies, some first estimates were made of what could and could not be done to save significant amounts of this energy.

When the OPEC crisis hit two years later, the study suddenly got a lot of attention in Washington, and the Cambridge Center was asked to select some short-term winners from the strategies that had been studied. Carpooling and the national 55-mile-per-hour speed limit were quickly singled out. These could produce savings of about 4 percent on the automobile's consumption of energy, and that is a big saving.

The Center had also made a first projection of what longer-term strategies could save how much energy. These included reducing air drag (on which Anatol Roshko's group has done important research here at Caltech), proper transmissions, lighter materials, radial tires, and, of course, combustion controls. Beyond this, there can be additional improvements from new engines and other new technology that will take longer — a decade — to find their way into the fleet. The auto companies provided part of the data for these studies, which addressed not only technological improvements but fleet-mix versus time, retooling times, and so on. We checked all these data at the Center against our own measurements and estimates and every other source of data and judgment that we could find.

Based on this computer simulation of projected national energy consumption by automobiles, President Ford called for a new-auto efficiency improvement of 40 percent.

The role of the Center — and of the Department of Transportation — was to provide the government with some understanding of what actions were likely to pay off. But DOT depended (with some regulatory teeth) on the private auto industry to create and initiate the new engineering to improve the fleet. That's how it ought to work.

Moreover, it is important to realize that simulation of an economic system is really different in kind from simulation of a power plant, or an aircraft, or of an air traffic control system. Simulation of physical systems has been developed to a highly reliable art. The stability of an aircraft can be predicted precisely before it is ever built. Landing at the Seattle airport can be experienced with rather high fidelity inside a building in Denver. But with simulation of economic systems, such confidence is far from possible. Indeed, done without extreme prudence, presented without humility, and accepted without the most sceptical judgment, economic simulation could be downright perilous.

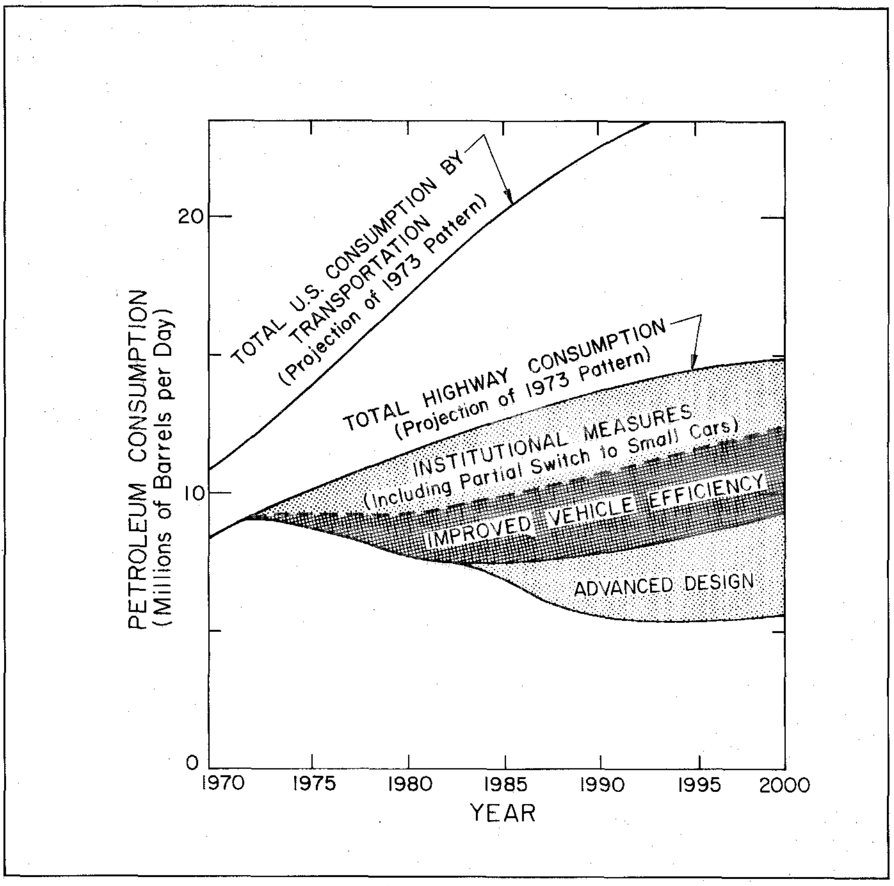

A computer simulation of projected national consumption of petroleum by automobiles indicates some of the strategies that might lead to improvements in our energy situation.

Recently some progress has been made in economic modeling. The progress has come from one thing: Computers can now handle massive amounts of data about the economy — about what is built, what is purchased, and what energy is used. These data can be used through a technique called "regressive analysis" to test various models on the computer in order to get some idea of what might happen. But the uncertainty is very great. There is an even more fundamental difference: There is no direct way to do an experiment, and experiments are vital for a reliable model. The economic system involves millions of people exercising free choices — people in all their splendid inscrutability. We have only the most approximate models for these phenomena. Moreover, and above all, individual privacy must be rigorously protected.

Can we, then, simulate in any useful way the behavior of the economy under various assumptions? I think the answer is: only in certain respects, only with large uncertainty, only where massive and reliable data are available, and only if the very best human judgment, caution, and humility are brought to bear.

What simulation can do is establish the boundaries of what is possible. And it can help predict, quantitatively, some pitfalls. It can, for example, combine a mass of estimates of earth's natural energy resources and other information about production and market, and can produce some understanding of the environmental and safety costs of pursuing them.

But, limited as they are, these are still important advances, important helps to our understanding. We must draw on this new tool to help us manage our data and — with care and scepticism — help us organize our understanding. For the stakes are high, and we need to bring to bear every tool we have.□