



From NASA's Visible Earth website, <http://visibleearth.nasa.gov>. Data courtesy of Marc Imhoff, NASA/GSFC, and Christopher Elvidge, NOAA/NGDC. Image by Craig Mayhew and Robert Simmon, NASA/GSFC.

Powering the Planet

By Nathan S. Lewis



This talk was the opening keynote speech at the first annual California Clean Innovation Conference, held at Caltech on May 11, 2007. The event, a partnership with UCLA and UC San Diego, included discussions on the futures of assorted energy technologies and how to finance them. In other sessions, clean-energy startup companies were given the opportunity to “fast pitch” their business plans, in three to five minutes each, to a panel of venture capitalists.

Nathan S. Lewis (BS '77, MS '77) is Caltech's Argyros Professor and professor of chemistry. Much more on global energy issues and on his own research in solar power can be found at <http://nsl.caltech.edu>.

This article was edited by Douglas L. Smith.


THE SCALE OF ENERGY

Energy is *the* single most important technological challenge facing humanity today. Nothing else in science or technology comes close in comparison. If we don't invent the next nano-widget, if we don't cure cancer in 20 years, like it or not the world will stay the same. But with energy, we are in the middle of doing the biggest experiment that humans will have ever done, and we get to do that experiment exactly once. And there is no tomorrow, because in 20 years that experiment will be cast in stone. If we don't get this right, we can say as students of physics and chemistry that we know that the world will, on a timescale comparable to modern human history, never be the same.

The currency of the world is not the dollar, it's the joule. Consider the image at left, for example. (I always have to explain to a lot of audiences, although I'm sure not this one, that this picture wasn't taken all at once.) You can see exactly where

Earth's city lights as seen from space. The brightest areas are not necessarily the most populous—compare China and India to the U.S. and Western Europe.

The world's energy diet is about four-fifths fossil fuels. "Biomass" means unsustainable burning of plant material; that is, burning it faster than it can be grown back. "Hydro" stands for hydroelectric power; "renew" means renewables—chiefly sustainable burning of biomass, but this category also includes solar, wind, and geothermal.



the consumption of electricity is. You can also see that there's an inordinate number of people who only have one candle to burn at night. They can't get out of poverty, they can't cure disease, they can't boil water, they can't do much of anything without energy. And they certainly can't save much energy.

Humanity's current energy consumption rate is 13 trillion thermal watts, or 13 terawatts. (My energy data all comes from peer-reviewed sources, primarily the World Energy Assessment report published by the United Nations Development Program, the latest version of which is available online at www.undp.org/energy/weaover2004.

The Stone Age did not end because we ran out of stones, and the fossil-energy age is not going to end any time soon because we've run out of cheap fossil energy. *Don't wait for that to happen.*

htm.) If you took the heat content of all the energy we consume in whatever form—kilowatt-hours of electricity, barrels of oil, cubic feet of natural gas—in a year, and divide it by the number of seconds in a year, you get thermal watts, which I will use as my standard unit, for ease of comparison. And, to refresh your memory, a watt is a joule per second. Politicians talk about changing a few light bulbs in Fresno to compact fluorescents. That's nothing compared to the 13 terawatts that the whole globe consumes, on average. This is the scale of energy.

The United States consumes a quarter of the world's energy, at a rate of about 3.3 terawatts, but I won't say anything more about the United States. To physicists, it's not important. I care more about the 13 terawatts. Of the global consumption,

about 85 percent comes from fossil fuel—coal, natural gas, and oil. These are primary fuels, that is, direct energy sources. And about 4.5 terawatts of that is used to make electricity—a form of secondary energy—resulting in the generation of about 1.5 terawatts of electricity.

I need to dissuade you up front from one important notion, that some low-cost process is magically going to take us away from fossil energy within the next 20 or 30 years. That's simply false. The Stone Age did not end because we ran out of stones, and the fossil-energy age is not going to end any time soon because we've run out of cheap fossil energy. *Don't wait for that to happen.* Any new energy-creating process is going to be a substitution product. It's not like the cell phone that's ringing in this audience as I speak, where people will pay a lot of money for the privilege of being the first person on the block to be able to annoy everyone else. Whether electricity comes from clean or green or mean does not matter to the end user. They only care that it comes out for a nickel a kilowatt, or less, because that's what electricity from coal and natural gas costs.

Selling a substitution product requires fostering a marketplace where the technology can come to scale and compete. You can't wait for the cost of a mature, competing technology that is already at scale to rise fast enough, soon enough, to make the new technology affordable. There is no way to compete with technology that consists of just taking concentrated energy sources, like coal and oil, pulling them out of the ground, and burning them. We can discuss the true costs of putting carbon into the atmosphere, but on the current economic basis, if we wait for price signals to drive us away from fossil energy, we'll be waiting a very long time.

Dividing our proven reserves by 1998 consumption rates shows that we have 40 years' worth of proven reserves of oil. This is what's in the ground that we can actually book with 90 percent confi-

dence. People look at this and say, "We're going to run out of oil in 40 years!" That's wrong. The ratio of proven reserves to consumption rates has been that same 40 years since the day after oil was discovered. If it costs a million dollars a day to drill a well, and three out of four wells turn up dry, it's not a good use of a corporation's capital to prove out more than 40 years of reserves. So you do that, and then you do something else with the money, like return it to your stockholders. On a net-present-value basis, it doesn't pay to prove out 100 years worth of reserves, so you *always* have about 40 years worth of proven reserves.


It's certainly true that most of the cheapest oil has been discovered, we believe. On the other hand, \$30 a barrel was thought to be prohibitively expensive three years ago, when the U.S. Energy Information Agency was forecasting \$24-a-barrel oil through 2025. Crude oil futures are now in the \$60-per-barrel range. And the higher the price goes, the more reserves you can access economically. The entire resource base—the best estimate of what's waiting to be discovered—gives us between 50 and 150 years at 1998 consumption rates. And if we should run out of oil, we have between 200 and 600 years of natural gas, and something like 2,000 years of coal. We know how to convert coal into oil—the Germans did it during World War II, and South Africa does it right now. In the United States, we could liquefy coal for \$40 a barrel, but investors don't even want to do that because they're not sure that even at that price it would be profitable in the long term.

IN THE YEAR 2050

"It's hard to make predictions, especially about the future." But that's never stopped us anyway. The graphs I'm about to show you come from a paper that Martin Hoffert, a physicist at New York University, and colleagues published in *Nature* in 1998, which in turn draws on data from the 1992 United Nations Intergovernmental Panel on Climate Change, or IPCC, report 15 years ago. The IPCC report was recently updated, but the findings remain essentially the same. So this is not new news.

I'm going to focus on the year 2050, which is not 43 years from now, it's five to 10 years from now. Our energy infrastructure has a capital-investment sunk cost that lasts for 40 years, so when you think about 2050, you think about that *now*. In addition, most of us—either our kids or ourselves—are going to be alive in 2050, so it's a good year to look at.

Obviously, people use energy. The world population is projected to be nine to 10 billion people by 2050 (we're at about six billion now), so I'll pick 10 as a round number. And I'll assume a gross domestic product, or GDP, growth of 1.6 percent per year per capita, which the IPCC calls the "business as usual" scenario, based on the average global GDP growth over the last century. The IPCC did not foresee, 15 years ago, 10 percent growth annually in China, and 7 to 10 percent in India. And the developed countries now believe that 4 to 5 percent growth is sustainable. But this doesn't matter, as the numbers just get worse as it gets higher. And

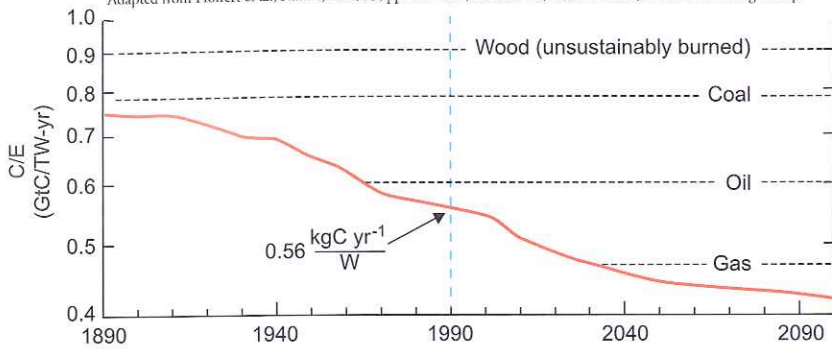


Top: This global population projection, taken from historical data and the Intergovernmental Climate Change Panel's "business as usual" scenario, hits 11 billion people by 2090 and keeps climbing—don't be fooled by the logarithmic vertical axis.

Bottom: The ratio of annual energy consumption per capita to gross domestic product per capita has been falling off in recent years as technology gets more efficient. The "business as usual" scenario assumes this will continue.

Adapted from Hoffert et al., *Nature*, vol. 395, pp. 881–884, October 29, 1998. © 1998, Nature Publishing Group.

Adapted from Hoffert et al., *Nature*, vol. 395, pp. 881–884, October 29, 1998. © 1998, Nature Publishing Group.



The IPCC's "business as usual" projection tracks how the carbon-to-energy ratio of our global energy mix has declined over time. But this trend cannot continue below the carbon-to-energy ratio of the cleanest carbon component without a substantial influx of carbon-free power. "GtC/TW-yr" stands for gigatons of carbon per terawatt-years.

no country that I'm aware of has a policy *against* economic growth.


With population and GDP growth conspiring together, we would then obtain a tripling of energy demand by 2050. This is partly mitigated, however, by the fact that we're using energy more efficiently per unit of GDP. The ratio of energy consumption to GDP has been declining at about 1 percent, globally averaged, per year. The United States actually saves energy at a faster rate, about 2 percent per year. Because we have such a high per-capita energy baseline consumption, it is easier for us to save off that base, whereas the developing countries save less. The "business as usual" scenario assumes that this will continue, and if we project that down, we will achieve an average energy consumption of two kilowatts per person within our lifetimes. (The United States now uses 10 kilowatts per person.) But factor in population growth and conservative economic growth, and we'll still need twice as much energy as we need now.

In terms of average thermal load, a person on a 2,000-calorie-per-day diet is basically a hundred-

watt lightbulb. And in our highly mechanized western agricultural system, the energy embedded in food—to run the farm and grow the food and transport it to the supermarket and put it in the refrigerator—is 10 to 20 times the energy content of the food itself. And the farther you live from the food source, the more embedded energy you consume. If we are 100-watt lightbulbs, this means that just keeping us fed requires one to two kilowatts.

The other thing we need to consider is the amount of carbon emitted per unit of energy produced, or the so-called carbon intensity of our energy mix on average. Back in the Stone Age, the carbon-to-energy, or C/E, ratio was quite high, as we were burning wood in caves. That's very inefficient. Most of the energy escapes into the air. We then moved to coal, and coal is not bad engineering, it's bad chemistry. We know how to burn coal efficiently, and when we burn all the carbon we get all carbon dioxide. When we burn natural gas, that's CH₄, we get one molecule of CO₂ but two H₂O's. So relatively more of the heat content in joules is delivered by making H₂O rather than forming CO₂. Natural gas is thus more energy-efficient on a carbon-emitted basis. And oil is in between, having a chemical formula of CH₂, on average. These figures are constants you can do nothing about. They are simply the products of the chemical formulas and the heats of combustion of coal, oil, and natural gas.


If we follow the "business as usual" C/E projection, which is hardly business as usual except for drawing straight lines into the future, it predicts by 2050 an average carbon intensity of 0.45, which is lower than that of the least-carbon-intensive fossil fuel, natural gas. And the only way you can do that is with a significant infusion of carbon-free or carbon-neutral power, to bring the overall average lower than the least of its carbon-based components. Furthermore, if you accept that we continue to burn oil and coal, because they are cheap, we'll need even more carbon-neutral energy to bring us



Adapted from Hoffert et al., *Nature*, vol. 395, pp. 881–884, October 29, 1998. © 1998, Nature Publishing Group.

The heavy black line shows humanity's primary-power consumption in the "business-as-usual" scenario. The red lines show the carbon-based power consumption reductions needed to stabilize atmospheric CO₂ at various levels.

The last 425,000 years' worth of data from ice cores drilled at Vostok, Antarctica, show that the levels of atmospheric carbon dioxide and methane, both greenhouse gases, go hand in hand with average temperatures.



down there. But I'll assume we will do that, too. So we've magically, somehow, added enough carbon-free power that we can stay on this decarbonization curve. And I'll further assume that we've implemented highly aggressive energy efficiency to reduce our total demand per person down to two kilowatts. This assumes that we can get the energy embedded in our food down to one kilowatt as part of that aggressive conservation program, and that leaves us with one kilowatt per person to heat our houses, get to work, play video games, and do everything else we do. And under those assumptions, if we relate the amount of carbon emitted to the amount of energy consumed, it is simple arithmetic to calculate the amount of carbon that we will release into our atmosphere. That set of

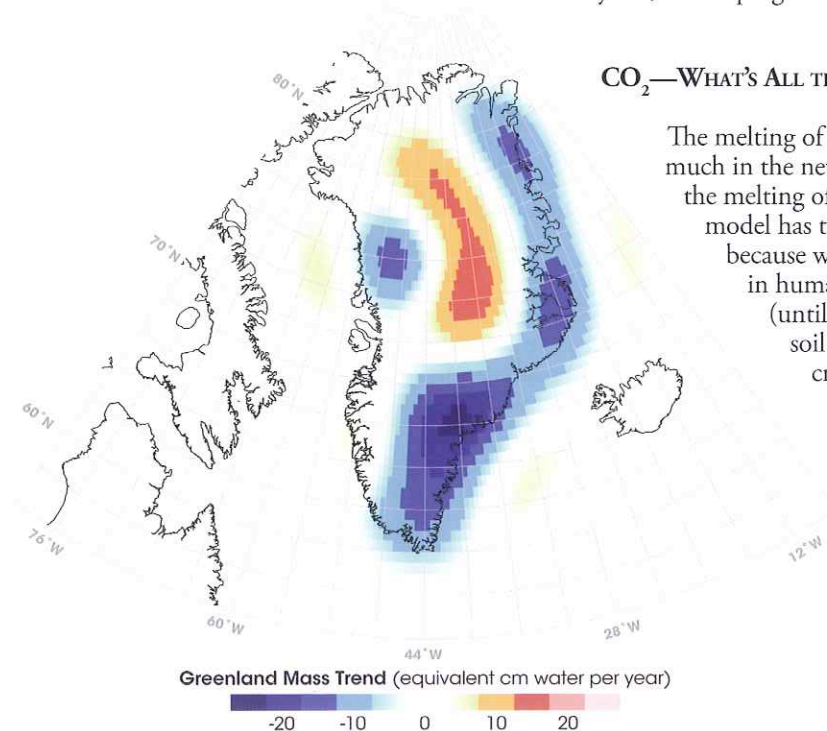
calculations brings us to the heavy black line labeled IS92a, which is the IPCC's shorthand name for this particular "business as usual" scenario.

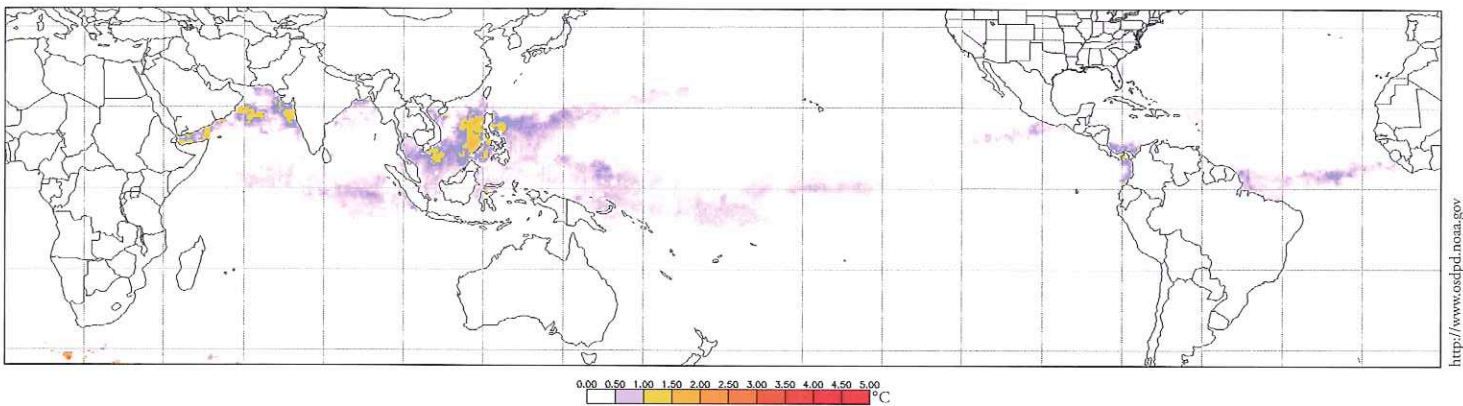
However, this is *still* insufficient to stabilize the atmospheric levels of CO₂ at any reasonably acceptable levels. Ice cores taken near Vostok Station, Antarctica, show that the CO₂ level has been in a narrow band between 200 and 300 parts per million by volume (ppmv) for the last 425,000 years; data from other cores have extended this back to 670,000 years. Current CO₂ levels are about 380 ppmv. "Business as usual" will require 10 trillion watts, 10 terawatts, of carbon-free power, and it never stabilizes CO₂ levels—they just keep going up. So even on that track, we are betting against data that goes back for almost a million straight years, and hoping that this time, we get lucky.

CO₂—WHAT'S ALL THE FUSS ABOUT?

The melting of Greenland's ice pack has been much in the news, but let's talk instead about the melting of the permafrost. No climate model has that nonlinear effect built in, because we have no experience of it in human history. Permafrost is the (until now) permanently frozen soil of the tundra, and as the ice crystals in it melt, it reflects less light and turns darker, absorbing more light, and that melts more permafrost. Helium dating of trapped bubbles in the permafrost shows that we're melting permafrost now that hasn't been melted in 40,000 years. And there's enough CO₂ and methane (another

Observations made by NASA's Gravity Recovery and Climate Experiment (GRACE) satellites show that between 2003 and 2005, Greenland's low coastal areas shed 155 gigatons (183 cubic kilometers) of ice per year, while snow accumulation in the interior was only 54 gigatons per year. This two-year ice loss is roughly equivalent to the amount of water that flows through the Colorado River in 12 years.





Oceanic hot spots on June 11, 2007, as compiled by the National Oceanic and Atmospheric Administration's satellites. A hot spot is defined as a region where the sea-surface temperature is at least one degree Centigrade greater than the maximum expected summer temperature. These warmer waters can lead to the bleaching and eventual death of coral reefs.

greenhouse gas) trapped in the permafrost to have the greenhouse gas levels not go up by a factor of two but by a factor of 10.

The world was there at least once before, most recently in the Permian era 250 million years ago. There was a massive release of isotopically light carbon from unknown causes, and CO₂ levels rose by a factor of 10. (The fast release rate and the isotope ratio suggest it was some sort of self-catalyzing event, such as permafrost melting, as opposed to, say, a volcanic release.) Temperatures spiked for on the order of tens of thousands of years, and the fossil record shows that about 90 percent of the species on the planet went extinct. We do not know if this will happen again. We do know that there is only one way to find out.

The CO₂ we produce over the next 40 years, and its associated effects, will last for a timescale comparable to modern human history. This is why, within the next 20 years, we either solve this problem or the world will never be the same.

How different that world will be, we won't know until we get there.

We also know that, unfortunately, there is no natural destruction mechanism for carbon dioxide in our atmosphere. Unlike ozone depletion, it will not heal by itself through chemical processes. In our highly oxidizing atmosphere, CO₂ is an end product. The lifetimes of CO₂ in the atmosphere are well known, and the time for 500 to 600 ppmv of CO₂ to decay back to 300 ppmv is between 500 and 5,000 years. Which means that the CO₂ we produce over the next 40 years, and its associated effects, will last for a timescale comparable to modern human history. This is why, within the next 20 years, we either solve this problem or the world will never be the same. How different that world will be, we won't know until we get there.

Although major uncertainties remain, most climate-change researchers set 550 ppmv as the upper limit of what would lead to about a two-degree-Centigrade mean global temperature rise. This is projected to have significant, but possibly not catastrophic, impacts on the earth's climate. For example, the coral reefs would probably all die. But we, as humans, would probably be able to adapt, at some level, to such a change. On the other hand, most people in the modeling effort feel that 750 ppmv or higher would be quite serious.

If we want to hold CO₂ even to 550 ppmv, even with aggressive energy efficiency we will need as much clean, carbon-free energy within the next 40 years, online, as the entire oil, natural gas, coal, and nuclear industries today combined—10 to 15 terawatts. This is not changing a few lightbulbs in Fresno, this is building an industry comparable to 50 Exxon Mobils. Furthermore, if we wait 30 years, the amount of carbon-free energy we'll need will be even greater, and needed even faster, because in the meantime we will have put out 30 years of accumulated CO₂ emissions that will not go away for centuries to millennia. So stabilizing at 550 ppmv will then require about 15 to 20 terawatts of carbon-free power in 2050.

These results underscore the pitfalls of "wait and see." Because "wait and see" is "wait and do."

KICKING THE CARBON HABIT

We absolutely have to have universal, government-based policies to drive this transformation if we are going to make such a transition on this rapid a timescale. As I said, if a substitution product has to compete on a cost basis from Day One with our cheapest energy sources and their economies of scale, we won't get there. "If carbon dioxide is free, we'll take 10." And, contrary to assertions, we simply do not have the technology on the shelf to provide that much carbon-free pow-

er cost-effectively today. You will hear people say we have the technology, all we need is the political will. We have the technology to go to the moon, too, but just because we have the political will to give Southwest Airlines a few gates at La Guardia doesn't mean that they'll fly you to the moon and back on a \$49 Internet special. It's a question of scale, as well as cost, not solely technology.

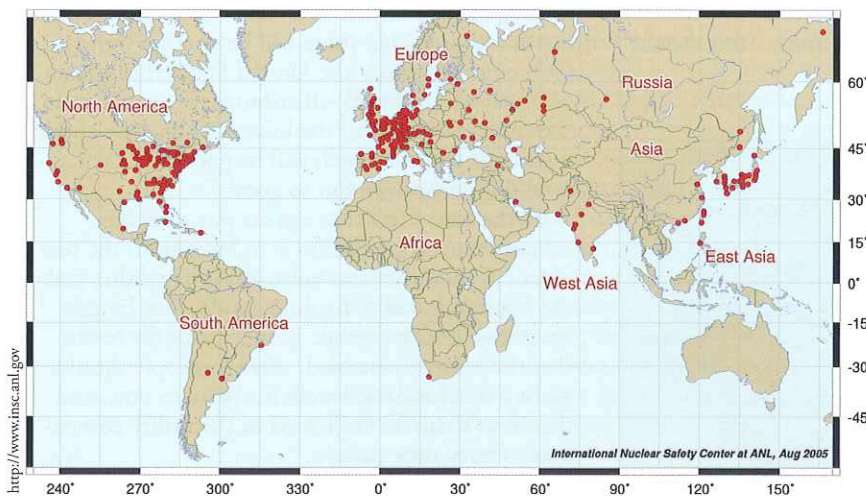
Let's talk first about energy efficiency. It's much cheaper to save a joule of energy than it is to make it, because the losses all along the supply chain are such that saving a joule at the end means you save making, say, five joules at the source. So lowering demand with energy-efficient LED lighting, fuel cells, "green" buildings, and so on is going to pay off much sooner than clean energy supplies. On the other hand, if we save as much energy as we

currently use, combined, we will still need to make at least as much carbon-neutral energy by 2050 as we currently use, combined, merely to hold CO₂ levels to double where they are now. That's the scale of the challenge.

So let's look at carbon-neutral energy sources. We could go nuclear, which is the only proven technology that we have that could scale to these numbers. We have about 400 nuclear power plants in the world today. To get the 10 terawatts we need to stay on the "business-as-usual" curve, we'd need 10,000 of our current one-gigawatt reactors, and that means we'd have to build one every other day somewhere in the world for the next 50 straight years. I've been giving this talk in one version or another for five years—we should have already built on the order of 1,000 new reactors, or double what's ever been built, just to stay on track. So we're really behind.

There isn't enough terrestrial uranium on the planet to build them as once-through reactors. We could get enough uranium from seawater, if we processed the equivalent of 3,000 Niagara Falls 24/7 to do the extraction. Which means that the only credible nuclear-energy source today involves plutonium. That's never talked about by the politicians, but it's a fact. Forgive my facetiousness, but on some level we should be thanking North Korea and Iran for doing their part to mitigate global warming. We'd need about 10,000 fast-breeder reactors and, by the way, their commissioned lifetime is only 50 years. That means that after we choose this route, we're building one of them every other day, or more rapidly, forever.


We don't have time for the physicists to figure out how to make nuclear fusion reactors—they've been saying it will be demonstrated (although not economical) in 35 years, and they've been saying that for the last 50. If we assume they're right this time, then ITER, a multinational demonstration fusion reactor being built in the south of France, will demonstrate break even—that is, it will put



The red dots show nuclear power reactor locations. (Map courtesy of the International Nuclear Safety Center at Argonne National Laboratory.)

Incomplete burning of coal and wood leads to a buildup of haze in eastern China, where mountains and weather patterns can trap it for days at a time.

Here the haze extends from the edge of the Gobi Desert (left) to the South China Sea (right)—a distance of well over 2,000 kilometers. (Image courtesy of the SeaWiFS Project, NASA/Goddard Space Flight Center, and ORBIMAGE.)



out as much energy as it takes to run it—in 35 years, and it will run for all of one week before the entire machine will, by design, disintegrate in the presence of that high-neutron radiation and temperature flux. And in the meantime we would have to build a commercial fission reactor every day for the next 30 years. It's not going to happen.

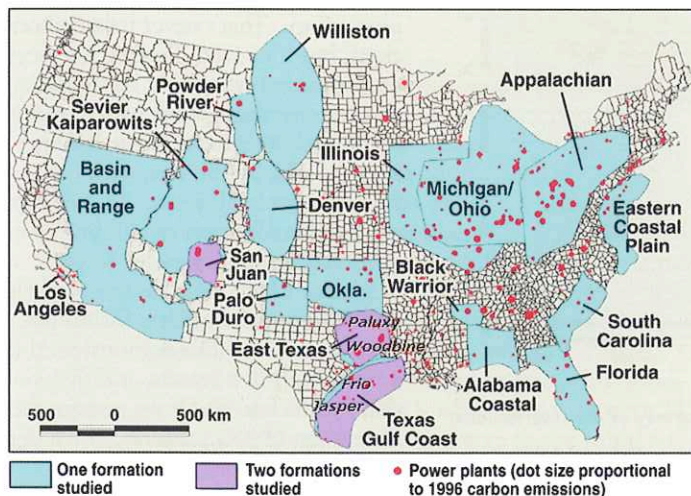
We could get there by sequestering the carbon. We have plenty of cheap coal, globally. China is building two gigawatts' worth of coal-fired electric power plants every week now. We could pipe the CO₂ out to the deep ocean, but CO₂ dissolved in water becomes carbonic acid, and estimates are that in some places the local pH change would be about 0.1 pH units. That's probably not a good idea. We could pump the CO₂ into deep oil and natural gas wells, but there aren't enough of them to hold all the CO₂ we will make during the next 50 years. We could put it in deep aquifers, where there's about 100 to 200 years' worth of total capacity, which would give us enough time to bridge to something else—if it works technically. You should

not assume that it works yet. The decay time of CO₂ in the atmosphere is, as I said before, between 500 and 5,000 years. That means that if one percent of the CO₂ in the reservoirs leaks, in 100 years the flux to the atmosphere would be identical to what you intended to mitigate in the first place. We know that CO₂ migrates underground. It bubbled up in Lake Nyos, Cameroon, on August 26, 1986, and killed some 1,700 people. So we're going to have to demonstrate within the next 10 years that it will leak less than 0.1 percent, globally averaged, for the next millennium in thousands of different aquifers around the world.

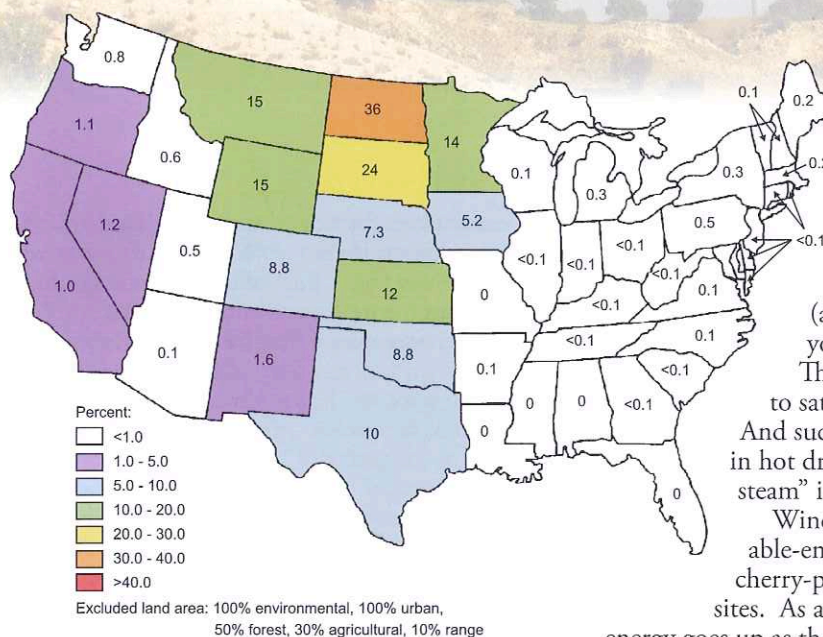
Every site is geologically different. So even if you validate sequestration at one site, that doesn't mean that it will work at the other thousands of sites we'll need. (Of which, by the way, nobody knows whether China has basically any.) And be careful what you wish for, because you might actually get it. If it works, a quick calculation based on the density of supercritical CO₂ at 1,000 meters' burial depth indicates that there will be enough buried CO₂ emissions from the United States that within 100 years, if uniformly distributed, it would cause a rise in the elevation of the lower 48 states by about five centimeters. Which will be good if the sea level rises; otherwise not so good.

By the way, I feel that a great way to make money from sequestration is to learn from the past. Think of the American railroads—they didn't make the big money off of hauling goods, they bought up all the land and made money from the towns that the railroads enabled. And so people should go buy up abandoned wells for pennies now, and then rent them for millennia to the utility companies to bury their carbon.

Areas of the continental United States where deep saline aquifers may allow CO₂ sequestration. Many coal-fired power plants are not near such aquifers, which means that CO₂ would have to be piped to them. (Map courtesy of the U.S. Department of Energy.)



The electricity-generating potential of wind from Class 4 and higher sites, i.e., places where the average wind speed at a height of 50 meters above the ground is 28 kilometers per hour or better. The percentage figures compare this potential to 1990 electricity consumption.



our planet, if you captured all of the heat flux at 100 percent efficiency (a small second-law problem!), you might get 11 terawatts. The heat of the earth isn't close to satisfying our thirst for energy. And such deep geothermal wells in hot dry rock tend to "run out of steam" in about five years.

Wind is the cheapest renewable-energy source now, because we cherry-pick the high-wind-velocity sites. As a bonus, the wind's potential energy goes up as the cube of the wind speed— $1/2 mv^2$ times the mass of air per unit time, which introduces another factor of v . And wind energy is relatively economic, about five cents per kilowatt hour in *very* high-wind-speed areas, but, even adding in the lower-wind-speed areas, when you calculate the total kinetic energy that we can get at the surface of the earth, there is to be had in practical terms about two to four terawatts.

If we assume that the net energy return from biomass equals the gross energy production—that is, that it takes negligible energy input to run the farm and harvest the crop—generating 20 terawatts would require 31 percent of the total land area of the planet— 4×10^{13} square meters. The problem is that photosynthesis is fundamentally inefficient. Leaves should be black instead of green. They have the wrong band gap, and they convert less than 1 percent of the total energy they receive from sunlight into stored energy on an annual basis.

And, by the way, the fastest-growing plants known are a mere factor of two or so under their ultimate CO_2 fixation rate. CO_2 is dilute in the atmosphere, so unless there's a transport system sucking carbon dioxide down from above, the natural mass-transport rates limit plant growth to a factor of two or so over the fastest that we already have. So if someone shows you pictures of little

RENEWABLE ENERGY

Which brings us last to renewable resources—biomass, hydroelectric, geothermal, wind, and solar.

Hydroelectric power is a model renewable resource, but all the kinetic energy in all the rivers, lakes, and streams on our planet combined adds up to a rate of 4.6 terawatts. And we can't tap all of that, because we can't dam up the Okeefenokee Swamp and get much energy. So as a practical matter, there's 1.5 or so terawatts available, but that includes places like the Hudson River, and we only want to dam that if the Yankees fire Joe Torre. Similar economic considerations leave us 0.9 terawatts, and we've already built 0.6. So forget about hydroelectricity. It's cheap, it's abundant, and we've pretty much maxed it out.

You'll hear a lot about geothermal energy. The sustainable geothermal heat flux works out to 0.057 watts per square meter. That's from the temperature at the center of the earth, the thermal conductance of the earth, and the diameter of the earth. So from the entire continental surface of

I believe in the Willie Sutton school of energy management. . . . One hundred twenty thousand terawatts of solar energy hits the earth, so Willie Sutton would say go to the sun because that's where the energy is.

tomatoes and big tomatoes, and extrapolates from tall switchgrass to 20-times-taller switchgrass, that's defying the laws of physics.

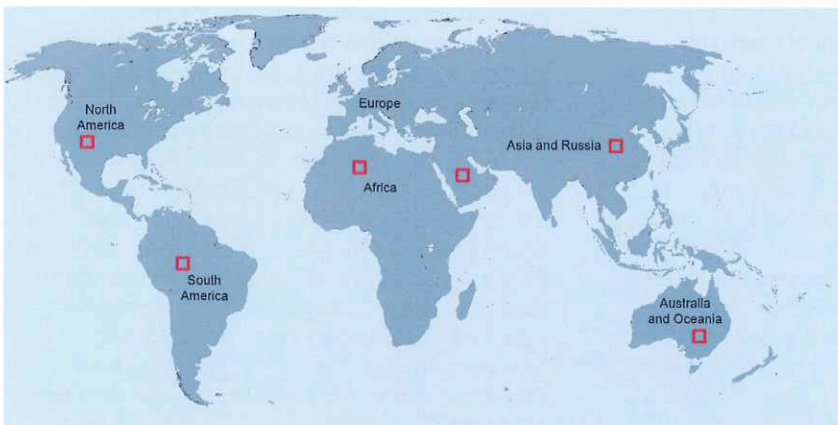

You hear a lot about schools of management. I believe in the Willie Sutton school of energy man-

agement. The Willie Sutton principle is simple. Willie Sutton was a famous bank robber, and when they finally caught him someone asked, "Why do you rob banks, Mr. Sutton?" He said, "Because that's where the money is." I believe in that, too.

One hundred twenty thousand terawatts of solar power hits the earth, so Willie Sutton would say go to the sun because that's where the energy is. It is the *only* natural energy resource that can keep up with human consumption. Everything else will run up against the stops, soon. In fact, more solar energy hits the earth in one hour than all the energy the world consumes in a year.

For a 10-percent-efficient photovoltaic system, and the latest systems are 15 percent or better, we could supply all the United States' energy needs with a square of land some 400 kilometers on a side. As you can see in the map at left, this would cover the Texas and Oklahoma panhandles, part of Kansas, and a wee slice of Colorado. The good news is that this area is pretty lightly populated, and the residents of even a few counties there would make enough energy to become full-fledged members of OPEC. And six of these boxes would power the globe. Unfortunately, solar is also far and away the most expensive way we have of making electricity today, with costs ranging from 25 to 50 cents per kilowatt-hour for photovoltaic systems, that is to say solar panels. Solar thermal systems, which I'll talk more about in a moment, run 10 to 15 cents per kilowatt-hour, which is still too expensive. Nobody is going to pay that much for a substitution product, when they can get the original one for four cents a kilowatt-hour.

The only way that we can get this to happen is if we lower the cost of solar converters to something like \$10 a square meter. It has to be something you'd buy at Home Depot to paint your roof with. You can't use single-crystal silicon—at this cost, you have to think potato chips, not silicon chips. You have to use really cheap materials, so my lab is trying to make solar cells out of fool's gold and rust.



Top: The nation's entire energy needs could be met by tiling a 400 × 400 kilometer parcel of land in the sunny Midwest with solar panels.

Bottom: Six such squares, appropriately sited, could power the world.

Solar thermal systems, in which a parabolic dish of mirrors focuses the sun's energy on a collector, produce cheaper electricity than photovoltaic cells.

They can also be easily mated to electrolyzers (the building and cooling towers in the background) to transform that electricity into storable hydrogen fuel. Unfortunately, they don't scale up well.



And we're working on paintable materials based on nanorods of TiO_2 , which is the white pigment in paint. The folks at Behr Paint called yesterday to see if we had a bucketful that they could test, and we had to say no. We're still working on it.

And, by the way, if we succeed and make really cheap solar cells, that alone will not solve much in the big picture of energy. Because as Johnny Cochran might have said, "If it does not store / You'll have no power after four." Solar cells convert sunlight into electricity. And there's no good way to bottle up and store vast quantities of electricity. If you have one, go buy electricity off the grid at five cents a kilowatt-hour at night, outside of peak load hours, and then sell it back to the grid at 25 cents per kilowatt-hour in the daytime to balance the load, and laugh all the way to the bank.

I believe that the best way to store massive quantities of electricity is to convert it into chemical fuel. The best technology for that purpose that we have now uses a solar thermal system that collects and concentrates solar energy to electrolyze water. You get H_2 for fuel, which you can distribute through pipelines and store in tanks. And then you can pump it out of the tank whenever you like and run it through a fuel cell, which converts it back into electricity and water. The problem is, the existing technology is not scalable. The setup in the photo above makes about a kilogram of hydrogen—the energy equivalent of about a gallon of gasoline—every day. And we would have to build one of these every second, for 50 straight years, just to hold the CO_2 concentrations to 550 ppmv. We need to find a better way to make fuel from sunlight directly so that we can bring energy to whoever wants it whenever they want it—day or night, summer or winter. My lab and other labs at Caltech are working on that, too.

So, in summary, we're going to need more energy in order to lift people out of poverty and have economic growth. Even if we keep demand flat, it doesn't help us very much because CO_2 emissions

are cumulative. And the globe has *never* had a year in which it has used less energy in a year than it did the year before.

No rational energy program would start without promoting energy efficiency. We should do all we can there. But no amount of saving energy ever turned on a lightbulb. No amount of saving energy actually put food on somebody's table. Energy efficiency is simply not enough to bridge the demand gap. On the supply side, there are only three big cards to play, in some combination: coal sequestration, if we dare; nuclear fission involving plutonium, if we double dare; or finding a way to make cheap, storable energy from the other big card that we have, which is the sun. But solar has to be *really* cheap, and scalable, *and* we've got to find a way to store it.

I haven't talked much about economics, but I will say that it's easy to prove, thinking 100 years out, on a risk-adjusted net-present-value basis, that the earth is simply not worth saving. It's a fully depreciated, four-billion-year-old asset. Unless you have policy incentives that reflect the true cost of doing this experiment, the economically efficient thing to do is just what we are doing now. On the other hand, with the appropriate policy incentives, the financial opportunities are commensurate with 50 Exxon Mobils on the supply side, and, in devising ways to lower our energy consumption from triple to double by 2050, 50 more Exxon Mobils on the demand side. This is both the challenge and the opportunity.

I leave it to you to decide whether this is something that we cannot afford to do, or something at which we simply cannot afford to fail. Remember, we get to do this experiment exactly once. And that time, like it or not, is now. □