

Capture of carbon dioxide from ambient air

K.S. Lackner¹

Columbia University, New York, and NY GRT LLC, Tucson, AZ, USA

Abstract. Carbon dioxide capture from ambient air could compensate for all carbon dioxide emissions to the atmosphere. Such capture would, for example, make it possible to use liquid, carbon-based fuels in cars or airplanes without negatively impacting the climate. We present a specific approach based on a solid sorbent in the form of an anionic exchange resin, that absorbs carbon dioxide when dry and releases it when exposed to moisture. We outline a particular implementation of such a moisture swing and discuss the scale of the collectors, the energy consumption, and the indirect carbon dioxide emissions related to the operation of carbon dioxide capture devices.

1 Introduction

Concerns over climate change are driving innovation in technologies for stabilizing the carbon dioxide (CO₂) concentration in the atmosphere. Here we describe a new technology for capturing CO₂ directly from ambient air at collection rates that far exceed those of trees or other photosynthesizing organisms, and that has a cost that would allow its widespread use in managing the anthropogenic carbon cycle. Air capture technology provides an important new tool for carbon management, making it possible to consider carbonaceous energy carriers in situations where their use would otherwise have to be phased out.² Air capture can compensate for any emitted CO₂ by capturing an equal amount of CO₂ at a different location and time. Air capture is independent of the source of emission and so can be applied to any source. Air capture applied at a large scale can reduce the CO₂ concentration in the atmosphere, thereby making the current excursion in greenhouse gas concentrations temporary. Finally, capture of CO₂ enables the closure of the carbon cycle by recapturing CO₂, so that it can again serve as the chemical feedstock that provides carbon for fuel synthesis. The other inputs are water, which provides hydrogen, and energy from a source that is carbon-free.

2 Feasibility and figures of merit

Capturing carbon dioxide from ambient air involves the separation of a fairly inert and dilute component from a large volume of gas. Yet the challenge is not as difficult as it is generally

¹ The science and technology described in this paper are based in part on early work performed at Los Alamos and Columbia University. In large part the development work occurred at GRT LLC, a private company founded by Gary Comer, Allen Wright and the author in 2004.

² Early discussions of air capture focused on synthetic fuel production, see for example [3, 16, 18, 21]. [9] explicitly aimed for air capture as part of a carbon capture and storage concept. An earlier paper [11] had also suggested a geo-engineering approach to CO₂ capture, which was motivated by carbon management. Air capture without subsequent production of synthetic fuels requires much more emphasis on energy efficiency and thus leads to substantial design changes. The strategic importance of air capture for carbon management was also highlighted by [7].

perceived. Even though CO_2 is far less reactive than, for example, SO_2 , it is still a reactive sour gas. There are highly selective sorbents which can easily reduce the partial pressure of CO_2 in a gas stream to a small fraction of a Pascal (Pa). The CO_2 partial pressure in air is approximately 40 Pa.

Removing CO_2 from air is not new; it has been practiced for decades in the context of producing CO_2 -free air [1, 17]. However, air capture is different from CO_2 scrubbing, because one does not need to extract all the CO_2 out of the air. Instead, the purpose of air capture is to collect CO_2 as efficiently as possible [9]. Collection becomes progressively more difficult and requires stronger sorbents as the CO_2 concentration in the gas stream decreases. Therefore, an economic optimization for air capture of CO_2 will typically result in a device that collects less than half of the CO_2 present in the air.

The CO_2 content of air is 400 parts per million (ppm) by volume. Such a high degree of dilution limits practical capture options to sorbent-based approaches. For sorbent-based capture, the energy cost of collection scales with the amount of CO_2 captured, rather than with the volume of air processed [9]. By contrast, energy investments in heating, cooling, compressing, or expanding air scale with the volume of air. Unless the amount of energy used per unit of air is exceedingly small, such energy investments are simply not affordable.

One mole of CO_2 emitted in the combustion of gasoline or diesel is associated with a heat release of 650 to 700 kJ. This becomes a figure of merit for the energy impact of capturing CO_2 and also provides a scale for gauging the energy consumption in the capture process. By capturing a mole of CO_2 from the air, one enables the release of another mole of CO_2 , which makes it possible to harness approximately 700 kJ of heat from fossil fuels in a carbon neutral manner [10].

The comparison with the heat of combustion of gasoline³ is inspired by the use of air capture for making the transportation sector carbon neutral. Other forms of energy will have different CO_2 emissions associated with them. Because there are many different ways of producing electricity, a comparison to electricity will depend on local circumstances. In the United States a mole of CO_2 is released for every 230 kJ of electricity generated; in Germany it is 290 kJ. In Brazil, which uses predominantly hydroelectricity, a mole of CO_2 is released for every 1700 kJ of electric energy. In France, due to a high reliance on nuclear electricity, a mole of CO_2 is emitted for every 1900 kJ of electric energy. In China, where power generation is very carbon intensive, only 190 kJ are generated for every mole of CO_2 .⁴

In terms of concentration, atmospheric CO_2 is very dilute. However, the CO_2 content of air is large, if measured in terms of its energy equivalent. A cubic meter of air contains 0.016 mol of CO_2 . Hence the combustion energy equivalent of CO_2 in air is 10,000 J/m³. This compares very favorably to the kinetic energy content of air. At a speed of 6 m/s, which is typical for a windmill location, the kinetic energy density in air is approximately 20 J/m³. This comparison suggests that windmills successfully extract from the air a value which in energy terms is five hundred times more dilute than the CO_2 captured in an air collector [10]. Therefore, a well-designed CO_2 collector that keeps up with the emissions of a fossil energy source could be two orders of magnitude smaller in its wind-facing cross section than a windmill that replaces the same energy source.

To estimate the size of a CO_2 collector, consider the following calculation: a square meter opening through which air flows at 6 m/s will pass 120 W of kinetic energy and 3.7 g/s of CO_2 . In the course of a day, such an opening would pass 300 kg of CO_2 , or about five times the daily per capita emission in the US. On the other hand, the primary energy consumption in the US is 10 kW per person, or eighty times the wind energy flowing through the same opening.

Flue gas scrubbing in conventional power plants provides a standard of comparison for air capture. For it to be of economic interest, air capture cannot be many times more expensive than flue gas capture. Because of the much higher CO_2 concentration in the flue gas, flue gas scrubbing is easier than air capture. Flue gas concentrations range from 3% to 5% CO_2 , for

³ Here we consider diesel fuel, kerosene, jet fuel and gasoline as virtually interchangeable. They all have nearly the same energy content per unit of CO_2 emitted.

⁴ All carbon intensity data are from Energy Information Administration (EIA): <http://www.eia.doe.gov/oiaf/1605/techassist.html> (last visited February 2, 2009).

natural gas fired power plants, to 10–15% for coal fired power plants. They are a hundred to three hundred times higher than CO₂ in air. Both technologies have in common that they are sorbent based. Because of the much higher CO₂ concentration, a flue gas sorbent does not have to be quite as strong as a sorbent for air capture, and the size of the flue gas scrubbers is much smaller than air capture devices with similar collection capacity. The advantage of air capture is that there is no need to scrub all the CO₂ out of the air. By contrast, flue gas scrubbing cannot achieve its goal of carbon neutrality unless it collects nearly all the CO₂ from the flue gas stream.⁵

Even though the air capture collector is large compared to a flue gas scrubber, it is small in absolute terms. A unit on the size scale of a meter can easily keep up with one person's CO₂ emissions [10]. Because air collectors are small, the cost of contacting the air is likely to be small as well. The contacting cost includes the cost of the scaffolding and of the materials necessary for exposing large sorbent covered surfaces to the wind. The contacting cost does not include the cost of separating the CO₂ from the sorbent and of processing it further so that it can put into a pipeline or directly into storage.

If one were to assume, overly simplistically, that at a wind speed of 6 m/s, the costs and collection efficiencies of contacting the air for CO₂ and for wind energy were the same for the same amount of air flow, then a cost of \$0.05/kWh for wind electricity would be equivalent to a collection cost of \$0.50/ton CO₂.⁶ By contrast, the cost of sorbent recycling is likely to be much larger. The cost of flue stack sorbent recycling is already measured in tens of dollars per ton [12].

The free energy required to separate one mole of CO₂ from a gas mixture is given by⁷

$$\Delta G = RT \ln \left(\frac{P}{P_0} \right). \quad (1)$$

Here, R is the universal gas constant, T is the temperature of the gas, P is the remaining partial pressure of the CO₂ at the exit of the scrubber, and P_0 is the pressure of the gas to be separated, which we take to be the pressure of ambient air. At $T = 300$ K and $P_0 = 100,000$ Pa, the free energy of the absorption must be at least $|\Delta G| = 20$ kJ/mol. In scrubbing the flue gas of a power plant, the CO₂ partial pressure at the beginning and the end of the separation step is very different. The sorbent must work even with the low partial pressure at the end of the extraction step, which we assume is less than 10,000 Pa. In most installations one can expect flue gas to be warmer than ambient air. Assuming $T = 350$ K, the minimum requirement on the free energy of binding is more than 13 kJ/mol. Practical sorbent binding energies tend to be several times larger than the free energy change and are typically well above 50 kJ/mol [12]. As a result, many sorbents used in flue gas scrubbers are strong enough to work with CO₂ from air.

Because the required binding energies are small in both cases and only scale logarithmically with the concentration at the exit of the collector, the effort in separating the CO₂ from the sorbent will be very similar for air capture and flue gas scrubbing. As long as it is possible to find a good sorbent that is capable of absorbing CO₂ without introducing environmental or other problems, the cost of sorbent recycling for air capture will not be much different from the cost estimates that are currently being discussed for CO₂ capture in retrofitted power plants.⁸ The energy demand in sorbent recycling alone suggests that the sorbent recycling cost will far exceed the cost of contacting the air stream. However, the sorbent recycle cost for air capture

⁵ Current state of technology typically captures 85–90% of the CO₂ in the exhaust stream (Metz et al. 2005). In a world that strives for carbon neutrality, this is not enough. Air capture could address the remaining balance [22].

⁶ Throughout this paper the unit “ton” refers to a metric ton which equals 1000 kg. The use of MT for metric ton is confusing and we avoid it here.

⁷ For a detailed discussion of the thermodynamics of the separation process see [8].

⁸ To convert energy costs into dollars per ton we note that 50 kJ/mol of CO₂ is equivalent to 1.1 GJ/ton of CO₂. The cost of energy varies depending on the source, and ranges from about \$0.5 to \$2/GJ for raw coal to \$15 to \$30/GJ for electricity, (5 to 10¢/kWh).

is not much larger than those for flue gas scrubbers, because the sorbent binding energies in both cases are very similar.

This analysis, although oversimplified, allows for two general and important conclusions: First, for a successful air capture design the cost of contacting the air is not likely to drive the cost of the overall system. Instead, the cost of air capture is likely to be driven by the cost of sorbent recycling. Second, for a good design, this cost need not be larger than the sorbent recycling cost for today's flue gas scrubbers. Even if flue gas scrubbers were to improve greatly, the difference in cost would never be large, as thermodynamic requirements on the sorbents are very similar.⁹ These conclusions, reached first in 1999 [10], only depend on general physical principles. There are no particular assumptions made concerning the sorbent choices. Of course, it needs to be shown that specific implementations can be devised that approach the efficiencies assumed in this theoretical analysis.

3 The air capture collector

A passive, sorbent-based air collector can be viewed as a large filter standing in an airflow with the filter surfaces covered with or made from a CO₂ selective sorbent. Air that comes in contact with sorbent surfaces will relinquish some or all of its CO₂. The larger the surface area and the longer the contact time, the more CO₂ is removed from the air. However, more extensive contact with sorbent surfaces either increases the pressure drop across the filter, or it slows down the flow speed of the air moving through the filter. Also, CO₂ absorption rates on these surfaces drop as the concentration of CO₂ in the air stream is reduced. Therefore, the thickness of the filter should not be so large as to greatly reduce the CO₂ content of the air stream. In a wind-driven filter system, the pressure drop and flow speed will adjust, so that the partial stagnation of the air in front of the filter maintains the air flow speed through the filter. The goal of a well-designed air capture device is to optimize between flow speed and CO₂ depletion of the air stream.

A significant CO₂ reduction is feasible in a wind-driven filter system, because air drag and CO₂ uptake follow similar transport laws. In the absence of turbulence, the dominant momentum transfer is through viscous drag, which involves molecular diffusion of momentum. The diffusion constants for N₂, O₂ and CO₂ are similar. As long as the absorption is gas-side limited, i.e., the partial pressure of CO₂ directly on the sorbent surface is much smaller than in the bulk of the gas flow, the diffusion of momentum and the diffusion of CO₂ follow similar laws. However, there are differences. First, turbulent mixing can improve momentum and CO₂ transport, but it can also dissipate energy without contact to stationary surfaces. Hence, a collector that creates a strong turbulent wake wastes part of the pressure drop into kinetic energy that does not contribute to transporting CO₂ to the sorbent surfaces. Second, momentum transfer along pressure gradients has no counterpart in CO₂ transport, and thus momentum transfer tends to be more efficient. Finally, if the transport of CO₂ is not dominated by air side CO₂ gradients, but is limited in part by the ability of the surfaces to absorb the CO₂ from the gas stream, then the boundary concentration of CO₂ – in contrast to the momentum boundary – will not approach zero and the rate of CO₂ uptake will be less than the momentum uptake on the wall.

There are a variety of possible geometries for air filters. Many involve non-turbulent flow passing over sorbent surfaces. One possible design is analogous to a heat exchanger with a large number of flat surfaces and a flow that is tangential to these surfaces. A variation of this design is a honeycomb-like structure with narrow but straight passages as in the monolith structure of a catalytic converter for an automobile exhaust system. A second approach involves a filter mat made of loose fibers that are thin enough so that the Reynold number of the flow over their surfaces remains small. A good analog of such a design is a glass fiber-based furnace filter, which is designed to remove particles from an air flow in a duct.

⁹ A corollary of this observation is that air capture may prove to be competitive with retrofitting old power plants, particularly in situations where the power plant is far from a good disposal site.

The momentum loss due to viscous drag along the surfaces is made up by a constant pressure gradient pointing in the direction of flow through the filter. The fraction of transverse momentum lost in any layer of the filter is comparable to the fraction of CO₂ removed. The pressure gradient is maintained by partial stagnation of the air flow in front of the filter. For air flow speeds significantly smaller than the free wind speed, the pressure drop is close to the stagnation pressure, which is given by $\rho\nu^2/2$, where ρ is the density of air and ν the wind speed. The pressure gradient maintains a constant air flow through the filter. While the velocity remains constant, the CO₂ concentration gradually drops.

For a very conservative design capable of operating at low wind speeds, the effective flow speeds through the filter are on the order of 1 m/s even if wind speeds are measured at several meters per second. In that case, most of the kinetic energy and more than half the momentum have been lost to the drag resistance of the filter, but only a small fraction of the wind's CO₂ is collected.¹⁰ Assuming a collection efficiency of 30% for capture of the CO₂ that passes through the filter, the rate of collection is 0.25 g/s of CO₂ per square meter of frontal collector area.

Thus, an uptake of about 20 kg per day is a rough figure of merit for a square meter sized collector of this design. For a ton per day of CO₂ one requires a device with 50 m² of wind facing frontal area. This device is much larger than we estimated before, mainly because we now assume that the air flow through the apparatus is about six times slower than the wind speed in a good windmill area. The advantage of this approach is that air collectors can be installed in areas of low wind speed which would not be suitable to windmills.

Sorbent materials have to provide surface that is exposed to the flow. At a CO₂ concentration of 400 ppm, surface uptake rates that have been empirically measured for practical sorbents are in the range between 10 and 100 $\mu\text{mol m}^{-2}\text{s}^{-1}$. This will set the amount of surface required for a ton per day device (0.25 mol/s) to between 2,500 and 25,000 m². The specific uptake rate is proportional to the CO₂ concentration and the proportionality constant is largely determined by the properties of the resin. Only if the uptake rates are exceedingly fast will boundary layers in the air dominate the uptake rate. With a CO₂ diffusion constant in air of $1.39 \times 10^{-5} \text{ m}^2 \text{ s}^{-1}$, it would take a boundary layer thickness of 8 mm to interfere significantly with an intrinsic surface uptake rate of $25 \mu\text{mol m}^{-2} \text{ s}^{-1}$. At this boundary layer thickness, the flow rate through the boundary layer would also be $25 \mu\text{mol m}^{-2} \text{ s}^{-1}$ [10].

Unless the filter removes a significant fraction of the CO₂ in the air or is significantly air-side limited, the uptake rate of the device is independent of the flow speed. The surface area required is essentially independent of the wind speed, until the flow is so slow that the depletion of CO₂ from the air stream slows down the uptake rates at the collector surfaces.

For fresh air, the total surface area and hence the total amount of sorbent material determines the uptake rate. If the air is blowing fast, then it is possible to create a collector with a small frontal area and a long flow passage, i.e., a thick filter. Wherever the expected wind speed is small, the device will have a thin filter and a large frontal area, but the total surface area of the filter material and, therefore, the mass of sorbent material inside the filter are independent of these design choices.

These numbers more or less prescribe the dimensions of the collector. One will need between 50 and 500 m² of sorbent surface for every square meter of frontal area. This is easily achieved, for example with flat sheets that are one meter deep, spaced one centimeter apart and forming channels that run parallel to the intended air flow about one centimeter apart. A fiber mat with fibers 1 mm in diameter and 10% volume occupancy would require a thickness of 50 cm to create a surface area multiplier of 200. If the fiber thickness can be reduced to 0.1 mm, then the filter thickness could be as little as 5 cm.

The total amount of sorbent material will depend on the amount of surface one can create per unit of mass. Given a sorbent material, one can now calculate the total amount of material required, the time it takes to load and unload the sorbent, and hence, one has determined the major characteristics of the air capture device.

As an example, consider a mat of material with roughly one millimeter thick fibers, a specific surface area of 4 m²/kg, and an uptake rate of $25 \mu\text{mol m}^{-2} \text{ s}^{-1}$. Such a collector can absorb

¹⁰ The initial air velocity is assumed to be more than 2 m/s, but the actual velocity required will depend on the details of the design.

$100 \mu\text{mol m}^{-2}\text{s}^{-1}$, and therefore, needs 2,500 kg of resin to collect 0.25 mol/s or one ton per day. At a 10% volume occupancy and a resin density of 1 g/cm^3 , this requires 25 m^3 of filter boxes. In terms of filter boxes, this could be achieved with 30 filters, each 2.5 m high, 1 m wide and about 0.3 m thick. If a load cycle swings by 0.25 mol/kg of resin, it will take 2,500 seconds to load the sorbent with CO_2 .

4 Sorbent choice

The focus of research at GRT LLC in recent years has been on developing a better sorbent. Initial experiments with sodium hydroxide solutions convinced us that strong aqueous alkali solutions have many drawbacks beyond the very large binding energy which requires a very large effort in prying the CO_2 back off the sorbent. For example, we found it very difficult to maintain slow liquid flows that uniformly wet sorbent surfaces. To avoid local drying of sorbent surfaces, we needed to pump liquids so fast that the energy inputs into pumping became significant. We also had serious concerns over the corrosiveness of strong sodium hydroxide solutions. Finally, we considered the liquid droplets that occasionally get entrained into air flows a serious hazard for open air deployment.

The use of sodium hydroxide would invalidate the energy argument given in a previous chapter. While it is in principle possible to use a sorbent that has binding energies with CO_2 similar to those used in flue gas scrubbers, a strong hydroxide solution has a much higher binding energy. Sodium carbonate is bound ten times more strongly than is necessary. Even the best of all sodium hydroxide-based collectors will run against this thermodynamic limit.¹¹

Furthermore, we realized that the kinetic limitations are those of aqueous hydroxide chemistry and are not specific to sodium hydroxide. As a result, the uptake rate per unit area will always be dominated by the limited reaction kinetics of CO_2 transferring from the gas phase into the aqueous liquid.

This pointed us to solid sorbents with potentially different mechanisms to explore. Furthermore, solid sorbents, unlike fluids, have microscopic surface roughness that can improve the nominal uptake rate significantly. After building a successful sodium hydroxide-based collector that kept up with the CO_2 output of a single car, we began a study of different solid sorbents. Our goal was to operate with a sorbent with lower binding energy while maintaining an uptake rate that would not be worse than that of a 1-molar sodium hydroxide solution. In the end, we chose a strong-base ion-exchange resin, which combined a number of appealing properties.

The successful resin material is a composite material with a resin similar to Marathon A (Dow Chemicals) as the active ingredient. Small resin particles are embedded into an inert polypropylene sheet. The resin makes up about 60% of the weight of the material. The material we tested was originally produced as a specialty electrochemical membrane (Snowpure LLC, San Clemente California) and was delivered in the form of thin sheets. The material was a successful choice, because it combined the low binding energy of the carbonate to bicarbonate transition with the faster reaction kinetics known from sodium hydroxide solutions.

The ion exchange resin is composed of a polystyrene backbone with quaternary amine ligands attached to the polymer. These quaternary amine groups carry a permanent positive charge. They can be thought of as the analog to an NH_4^+ , in which each hydrogen has been replaced by an organic carbon chain attached to the polymer matrix. Because all four chains are anchored into the resin material, the transition back to the NH_3 analog is not possible. Therefore, the resin acts as a strong base. The positive ions fixed to the polymer backbone never release a proton (as NH_4^+ would). Thus, the resin is more akin to a solidified sodium salt brine than to an ammonia solution.

The resin is typically purchased in its chloride form. In this form, chloride ions provide the negative countercharge that balances out the positive charges embedded into the polymer matrix. Washing the resin in sodium hydroxide will replace the chloride ions with hydroxide ions. The resulting material is very much akin to a solidified sodium hydroxide solution.

¹¹ For a discussion of strong hydroxide solutions see [2, 19, 20, 22, 23].

It is not at all surprising that this material, like many other solid, amine-based resins [5], readily absorbs CO₂ from a gas mixture. These resins are capable of absorbing CO₂ from a gas stream, even if the CO₂ is far more dilute than in ambient air.

In order to measure the uptake rate of these resin materials, we put them into a small contained volume of air (approximately 20 liters) using a glass bell and a bottom plate through which electric leads and gas tubes enter the volume. We maintained air flow inside, either by blowing air with a small fan over the sorbent material or by moving the sorbent material in a circular motion at high speed, which breaks up any air-side boundary layers. We then continuously and in real time monitored the CO₂ and H₂O content of the gas by circulating a small stream of gas through an infrared gas analyzer. After analysis, the gas circulated back into the main volume, which remained completely isolated from the outside throughout the course of the experiment. However, the apparatus allowed for the injection of CO₂, and the removal or addition of water vapor.

Other experiments involved the single pass of air blown at a constant speed through a filter box, with CO₂ and H₂O concentrations in the gas stream measured before and after the sorbent filter. Finally, a destructive analysis allowed us to titrate out the anionic composition of the resin material. In this way we determined the hydroxide, chloride, carbonate and bicarbonate loading of the resin.

From these and similar experiments, we established that the resin material has a rapid uptake rate for CO₂ that exceeds that of a one molar sodium hydroxide film of equal nominal area. Uptake rates are typically between 10 and 500 μmol m⁻² s⁻¹ and depend on the level of prior saturation and the amount of moisture present.

We also established that in order to saturate the resin with CO₂, one has to load it with one mole of CO₂ per mole of positive charge attached to the resin surfaces. Hence, the absorption leads to a compound, which has the same CO₂ to charge ratio as sodium bicarbonate. We therefore refer to this state as the bicarbonate state. Similarly, we have the hydroxide state of the fresh material, and a half-filled state with one CO₂ for every two charges. We think of this state of the resin material as a carbonate.


The fundamental reactions are the direct bicarbonate formation on the resin



the formation of carbonate



and the formation of bicarbonate from carbonate



We were surprised to find that the uptake rate of the sorbent material drops only gradually as it loads up from the hydroxide state with more and more CO₂. Even with a CO₂ concentration of less than 400 ppm in air, the uptake does not stop until the material virtually reaches the full saturation of the bicarbonate state. In the range between carbonate and bicarbonate, the uptake rates are comparable to those of one molar sodium hydroxide solutions converting to sodium carbonate solutions, i.e. between 10 and 80 μmol m⁻² s⁻¹.

Through titration experiments we established that washing the resin in a 0.5 molar sodium carbonate solution rather than a 1 molar sodium hydroxide solution leaves the resin in the carbonate state, i.e., loaded with half the CO₂ that can be loaded onto the resin when exposed to air. It is possible to wash a resin that has been equilibrated against ambient air in a sodium carbonate solution and remove half of the CO₂ that has been absorbed.¹² During drying, the resin begins loading itself up again with additional CO₂ from the air until it again reaches the bicarbonate state.

With a carbonate wash, the swing of the resin sorbent is limited between the carbonate and the bicarbonate state and thus has effectively been cut in half. However, eliminating the use

¹² To demonstrate that the residual sodium carbonate is not responsible for the CO₂ uptake, one can simply wash the resin in de-ionized water prior to recycling the material. However, the presence of such sodium carbonate residue does not have a significant effect on the performance of the resin.

of sodium hydroxide greatly simplifies the overall process scheme, as the day-to-day operation would only require benign chemicals like sodium carbonate. Most importantly, the energy required to recover the CO_2 from the bicarbonate brine, which in the process is reconverted to a carbonate brine, is far less than the energy required to convert a carbonate into a hydroxide. Since the charge density is on the order of 1.7 mol/kg of resin, it is theoretically possible to absorb 0.85 mol of CO_2 per kg of resin. Experiments have shown that the resin can be regenerated many times. Even after prolonged use, we have not been able to measure a rate of degradation.

The carbonate wash provided the first step toward an improved design, but it turned out that further improvements were possible. The wash cycle proved to be more than a passive removal of CO_2 from the resin. By washing the resin with a sodium carbonate solution, we produce a sodium carbonate/bicarbonate brine with an equilibrium partial pressure of CO_2 over itself significantly higher than that of ambient air. This can be observed easily, as the partial pressure of CO_2 inside the collector rises dramatically to several thousand ppm during every wash cycle. This suggests that the aqueous brine does more than just act as receptacle for CO_2 . It actively changes the resin material into a state that does not bind CO_2 . In effect, we created a CO_2 pump that removes CO_2 at low ambient pressure from air and then collects it into a brine which passes it back into a gaseous phase at elevated partial pressures. As the resin dries, it returns back to the original state in which it is able to absorb CO_2 once again. We demonstrated that one can remove CO_2 from a saturated resin until the brine approaches the bicarbonate state. At this point the equilibrium partial pressure of CO_2 over the system reaches several percent of an atmosphere.

Another observation confirmed that water plays a role in this change in state. We showed that the equilibrium pressure of CO_2 over a partially loaded resin is a strong function of the H_2O partial pressure in the air. Raising the water vapor pressure induces the resin to release CO_2 , whereas removing water vapor from the air induces additional CO_2 uptake. Based on simple considerations of the mass action law for equation 4, this seems counterintuitive, as the mass action law should result in the opposite behavior. However, a large number of measurements have demonstrated the effect.

For example, it is possible to provide a sufficient amount of CO_2 to a sorbent that is exposed to 30 parts per thousand (ppt) of water vapor, so that it eventually comes to an equilibrium with the surrounding air at a partial pressure of CO_2 of 1000 ppm. If one now takes this resin and exposes it to dry air, it will absorb additional CO_2 rather than release CO_2 . On the other hand, if the resin is equilibrated at 200 ppm of CO_2 and 5 ppt of H_2O , it will release CO_2 into ambient air, if the partial pressure of H_2O is raised to 300 ppt.

Even more dramatic is the response of the resin to exposure of liquid water, which results in a rapid release of CO_2 . Since pure water cannot hold onto the CO_2 , the CO_2 evolves into the gas volume above the water. This release could be maintained against a back pressure of 5,000 to 10,000 Pa. The release stops at a CO_2 loading of the resin that is close to its carbonate state. The highest partial pressures of CO_2 are achieved with fully loaded resin, at slightly elevated temperatures of 45 °C.

Inspired by these results, we demonstrated that simply exposing the resin to water would release CO_2 from the resin, even against a noticeable back pressure of CO_2 in the surrounding gas volume. It appears that water vapor of 45 °C is sufficient to drive most of the CO_2 off the resin and have it revert to the carbonate state.¹³ Since the resin only comes in contact with water vapor and not the liquid water source, it is possible to use a salt brine as the source of water.

It is now possible to create a cycle where CO_2 is loaded onto dry resin, driven off the resin by moisture, and once the resin has dried, it is again able to absorb CO_2 . In order to make contact with the air we considered flat sheets, thin strands of resin in furnace filter like packages, and thin strands attached to twisted wires as in a bottle brush.¹⁴ Once loaded with CO_2 from air flowing over the material, the resin is then exposed to moisture inside an enclosure from which

¹³ It is possible that this effect is due to condensation of liquid water on the surface on the resin. In any event, the presence of water vapor regenerates the resin.

¹⁴ If one considers the collector as a synthetic tree, than these bottle brush designs resemble pine branches.

the air has been removed. The water drives the CO₂ off the resin, creating a gas stream that can be dried and compressed to pipeline pressures.

In one test, we inserted a thin tube filled with strands of resin material into a sturdy tube that acted as vacuum chamber. Once the air was removed, we flowed water through the chamber and then extracted CO₂ from the system, by collecting the gas flow in a chamber kept at low pressure by external cooling with liquid nitrogen. The gas flow was first dried, then the CO₂ in a second chamber was condensed out. By closing off this chamber and letting it warm up to room temperature, we were able to create an internal pressure of approximately 0.2 MPa.

5 Design of a prototype

The considerations and experimental results summarized in the previous two chapters lead to a specific design for a CO₂ collector. Our goal is to create a device that can be built based on the current state of the art and that could be easily packaged into a standard cargo shipping container. These containers, which are approximately 12 m × 2.5 m × 3 m, can accommodate a device that can collect one ton of CO₂ per day. By building a unit that can be easily transported, we take advantage of the fact that air capture can eliminate the need for transporting CO₂ to the site of use or to the disposal site. Instead the capture apparatus can be installed at this site. Also, a one-ton-a-day unit provides an excellent scale for a first commercial demonstration of the technology.

The current plans envision a modular design involving a set of air filters, 2.5 m tall, 1 m wide, and roughly 30 to 40 cm thick. We design for a flow speed through the filters of approximately 1 m/sec. A continuous deployment of thirty such air capture units would collect one ton of CO₂ in the course of a day. Since the time for recovery and partial drying takes equally long, we design the entire system to contain 60 such units. The total volume of all of these units is easily stored inside a standard shipping container.

It will take a single filter unit approximately one hour to load up with CO₂. After the filter unit has been saturated with CO₂, it will be removed from the air collector and moved on an automated conveyor system into a chamber designed to be evacuated. There will be six such chambers built into the container, and each will hold five air filters. Once a chamber is filled with filter units, all air will be pulled out creating a low grade vacuum. Then moisture will be injected. This could be done by exposing the inside of the chamber to a reservoir of warm briny water, or by spraying clean water into the chamber. In either case, moisture will come in contact with the resin. In response to the exposure, the resin will release CO₂ and will do so for an extended period of time.

In order to maximize the CO₂ partial pressure in the collector chamber, we will increase the partial pressure in stages. First, the water vapor will come in contact with a nearly depleted resin filter and then be pumped to the next chamber, where the resin still contains more CO₂. The gas stream will move in a counter-flow manner from chamber to chamber, gradually building up the partial pressure of CO₂. In the last chamber, the CO₂ content will be highest and will be maximized by letting the gas get in contact with resin filters that are still fully loaded. These filters will be able to release some of this CO₂ against a substantial back pressure of CO₂.

Rather than moving resin filters from one chamber to the next, one will change the valving between the chambers which are logically arranged into a circle. At any given time, one of the chambers will be open to the outside; fresh air filters will be removed and spent air filters loaded back in. The other five chambers will be in various stages of CO₂ extraction from the resin. Gas will be transferred from the chamber that is closest to completion, to the next chamber upstream and is eventually pumped out of the last chamber which contains the highest concentration of CO₂. Once the open chamber is refilled, it will be valved into the circle, acting as the last chamber in the chain. The first chamber will then be removed from the cycle, opened and emptied out.

Our goal is to build up a CO₂ partial pressure between 5 and 10 kPa. In a subsequent step, we will pump this CO₂ out of the exit chamber and compress it up to pipeline pressure. In the first compression stages, a significant amount of water will condense out, because the inlet gas

is saturated with water vapor. During condensation, heat is released, which will be used to keep the recovery system at a temperature above ambient. In effect, the compression system acts as a heat pump for maintaining a slightly elevated temperature in the regeneration chamber.

Analysis of all energy inputs into the regeneration system reveals that the system can be driven exclusively with mechanical energy. All heat needed will be produced as an incidental by-product of the compression [8].

The energy consumption can be broken into the three major components. First, there is the removal of air from the regeneration chamber. For a volume fill of 10% and a CO₂ loading of 0.25 mol/kg (which is a small swing in sorbent loading), we require approximately 100 kJ of energy to evacuate a volume of 1 m³ which contains 100 kg of resin and approximately 25 mol of CO₂. Thus, we require 4 kJ/mol of CO₂. The second energy consuming step is the compression of CO₂ from 5 kPa to 6.7 MPa. At this pressure and at a temperature of 300 K, CO₂ turns liquid, so any further compression requires very little energy. This compression is the largest energy consumer, and it requires theoretically 19 kJ/mol. Lastly, during the initial compression stages one compresses water vapor against a constant partial pressure. This condensation requires 2.5 kJ/mol of water. If we are operating at 45 °C, the ratio of water vapor to CO₂ gas is roughly 1/1. Hence this adds another 2.5 kJ/mol. In addition, the condensation of the water vapor releases a substantial amount of heat, 40 kJ/mol of water. This heat source, together with heat release in the compression of CO₂, enables us to maintain the regeneration chamber at 45 °C.¹⁵

Inefficiencies in the pumps and compressors raise the actual energy consumption, and there is some additional energy consumption related to the mechanical operation of the air collector. In total, we estimate the energy consumption to be nearly 50 kJ/mol of CO₂ or 1.1 MJ/kg of CO₂ [8]. Even in a carbon-intensive economy like China's, the CO₂ captured would be nearly four times as much as the CO₂ released in the generation of the electricity required to operate the air capture device.

6 Future designs

The current design is ready for a first prototype implementation. However, one can expect that future advances will lead to further improvements. Much of the research will focus on the chemistry of the sorbent, which so far takes advantage of an off-the-shelf resin without any modifications to adapt it to this new purpose. Desirable changes would include a higher CO₂ pressure in the regeneration stage, possibly at the cost of a higher minimum partial pressure for absorption. If possible, the temperature and humidity conditions for capture and release should be tuned to the climate in which the device operates. The current resin choice is best for a desert climate, but can be made to work in cool climates regardless of the average relative humidity. This is possible because the humidity swing is sensitive to the absolute humidity rather than the relative humidity [8]. At low temperatures, even water saturated air has a low water vapor content.

Extremely cold temperatures, and a tropical climate with its high absolute humidity, both limit the operation of the present air capture device. Current designs can be used down to temperatures of about -5 °C before the reaction kinetics becomes too slow. In extremely hot and humid conditions, the absolute humidity is large enough to limit the load capacity of the resin. However, it is worth pointing out that the current design also does not take advantage of the full range of loadings from bicarbonate to carbonate. In cool climates the CO₂ collection works, but the rate of drying of the resin filters is much slower than in a desert climate in which drying is fast and efficient. It is, however, possible in a cool climate to force the condensation of the water vapor against surfaces that are kept cold by ambient temperatures.

Material development is likely to emphasize more surface area per unit of resin. The present material has a surface area of approximately 4 m²/kg; we aim to increase this by roughly a factor of ten. At this point, a typical fiber would have a thickness of 0.1 mm. This would reduce

¹⁵ It is worth noting that the evaporation of water inside the regeneration chamber is matched by a similar amount of condensation or adsorption of the water on the resin material. Therefore the humidification of the resin does not require low grade heat from external sources.

the amount of resin required by at least a factor of ten. The improvement is likely to be even larger, as the uptake rate appears to be limited by the rate at which CO₂ diffuses into the interior of the resin particles. Thus, reducing the average thickness of the particle will increase the uptake rate per unit of surface area, which is an improvement not accounted for in the above mentioned factor of ten that only accounts for the increase in specific surface area for a fixed mass of resin.

Unloading times would also be reduced by a factor of ten, leading to a reduction in the size of the regeneration chamber. In effect the apparatus runs faster, on a smaller volume of material, and with smaller regeneration chambers. However, the effort going into pumping and compressing remains unchanged.

Because of the smaller size of the filter boxes, and the smaller size of the regeneration chambers, the amount of equipment that can be packaged into a standard cargo-shipping container will be able to collect more than one ton of CO₂. However, the size of the frontal area of the collector installed on top of the container unit would be larger. If volumes shrink roughly by a factor of ten, a single unit may be able to collect as much as 10 tons of CO₂ per day and would require a wind facing frontal area of 500 m². Thus, it would turn into a tall structure.

At present 70% of the cost of the device is accounted for by the resin and the regeneration chambers. A reduction by a factor of ten in these parts of the equipment would make the compressors and pumps the largest capital expenditure.

Based on the current costs of compressors and pumps, and the amount of resin required, plus the cost of the conveyer system that moves filters around, we expect the cost of the system to drop from about \$200,000, which is our current estimate for the present design to about \$20,000 for a future design. The operational costs will also come down, as maintenance costs are reduced and one-of-a-kind pieces of equipment enter into mass production. At 5¢/kWh the cost of 50 kJ/mol CO₂ of electricity amounts to \$16 per ton of CO₂. Some further reduction in energy consumption certainly appears feasible, and thus it seems quite plausible that the long term price of air capture could drop as low as \$30/ton of CO₂.

Capturing a few tons of CO₂ on a daily basis is not sufficient for changing the dynamics of the world's CO₂ emissions. The world emits 29 Gt/yr of CO₂. CO₂ from oil alone amounts to 11 Gt/yr.¹⁶ One approach to scaling up would be to retain the modularity and build many more units. One would start with large air capture parks and move from there to huge collector fields. It would take on the order of ten million units to make a significant impact on the world's CO₂ emissions. Ten million units would collect 3.6 Gt/yr of CO₂ if the unit size stays constant. If these modules increase their collection capacity, ten million units could remove 10 to 20 Gt of CO₂ on an annual basis.

Ten million units may sound like a very large number; however, the recent world production of cars and commercial vehicles is approximately 73 million per year,¹⁷ and the number of shipping containers shipped per year is measured in hundred's of millions of units.¹⁸ The number of collectors that will need to be built is large, but not prohibitively large.

A more conventional approach would scale up individual units so that they become much larger. We expect that there are similar limitations to the size of these units as there are for windmills. This would limit the likely size scale of individual units to less than 1000 ton per day. However, we see significant advantages from a mass production paradigm that allows the production of smaller units in facilities that look more like automobile factories than power plant construction sites. Certainly, in the early stages of development there is a significant advantage in being able to readily ship a collector to a site where CO₂ is used or can be readily disposed of, but where demand may only be temporary. Air capture performed by small, transportable units in effect provides a virtual pipeline for CO₂ delivery to remote locations.

¹⁶ Energy Information Administration (EIA) Data: <http://www.eia.doe.gov/environment.html>, last accessed February 1, 2009.

¹⁷ Organisation Internationale des Constructeurs d'Automobiles (OICA), the numbers are from their website for the year 2007. <http://oica.net/category/production-statistics/> (last accessed February 1, 2009).

¹⁸ In 2002, 266 million twenty foot equivalent containers units (TEUs) were shipped across the ocean [14].

7 Discussions and conclusion

We conclude that capture of carbon dioxide from ambient air is physically possible, that it is possible to create collection devices with CO₂ uptake rates several orders of magnitude larger than those of trees, and that the cost of CO₂ capture is not prohibitive.¹⁹ While air capture technology will not be immediately economic for CO₂ sequestration, the technology has already reached the point at which some commercial and industrial CO₂ demands could be met by CO₂ captured from the atmosphere. Commercial use of air capture in niche markets would provide a path for the technology to be implemented that requires little government support. Niche markets include CO₂ use in food production, in refrigeration, or as dry ice. The US consumption for these purposes is 8 million tons per year. A much larger market can be found in enhanced oil recovery, where CO₂ facilitates the production of oil. This strategy works particularly well during times of high commodity prices and high oil prices. As the cost of capture comes down, the size of the potential markets increases, and eventually applications for carbon management come within reach.

An accelerated path toward the large scale implementation of air capture for carbon management would begin with large demonstration projects. For example, there currently are many CO₂ storage projects under way that lack access to readily available CO₂.²⁰ Air capture could provide CO₂ to virtually any storage site and thus make it possible to try out these concepts at remote sites, even if there are no power plants nearby that could deliver CO₂.

We expect that the first prototypes could break even at \$200/ton of CO₂. With learning-by-doing and further advances in the technology, the cost of CO₂ capture could drop much lower. A reasonable target is \$30 per ton of CO₂. To set the scale, \$30 per ton of CO₂ would add 7¢ to the cost of a liter of gasoline. The initial cost is derived from GRT LLC internal cost estimates that are based on a ten year life time, the cost of resins and other materials, the costs of pumps and compressors, the cost of electric power and water, and estimates of the cost of maintenance and fabrication. The long term estimate is based on the observation that the amount of resin required can be drastically reduced and that ultimately the cost will be bounded by the cost of electricity and the cost of compressors and pumps which are already being produced in large quantities and are therefore not likely to yield to substantial cost reductions. 50 kJ of electric energy per mole of CO₂ at 4¢/kWh translates into \$12.60 per ton of CO₂.

Air capture in connection with carbon dioxide storage provides a powerful tool for carbon management. It provides a conceptually simple, readily accountable way of collecting CO₂ [15]. It can collect CO₂ anywhere and at any time. Internal energy consumption of the capture process is small and can be easily accounted for. The life-cycle based indirect emissions are very small compared to the life time collection of CO₂ from an air collector.

Air capture takes advantage of the atmosphere as a huge temporary storage buffer that can readily handle a few years of current world emissions. The atmosphere also provides cost-free means of transporting CO₂ from the source to the sink. CO₂ emissions, as well as CO₂ depletions, rapidly mixes with the surrounding air, and will be transported around the globe in a matter of months (Elliot et. al 2001). Even the air over the southern and northern hemispheres mixes well within a few years.²¹ Therefore air capture need not operate near sources of emissions. It is best deployed at the location the CO₂ is to be used or stored.

A powerful argument for air capture is that there is no need to compare air capture to any baseline. Reductions are not computed based on a comparison to a scenario that was never tried

¹⁹ A large tree has been estimated to build up biomass of several tons of CO₂ over its life time of 30 to 50 years, forests can hold on the order of 400 tons of CO₂ per hectare in above ground biomass (see, for example, [13] and [6]). By contrast an air collector with a footprint of a less than 1% of a hectare can capture a ton of CO₂ a day.

²⁰ Wallace Broecker suggested in a private communication to use air capture as a way of facilitating experimental carbon dioxide storage.

²¹ According to the Scripps Institution of Oceanography Network, the atmospheric CO₂ concentration in Antarctica lags roughly two years behind the CO₂ concentration at the Mauna Loa Station in Hawaii. Since the vast majority of the world's emissions are in the northern hemisphere, this time lag is indicative of the mixing time. CDIAC website: <http://cdiac.ornl.gov/trends/co2/sio-keel.html> (last visited February 13, 2009).

out, but they simply reflect the net amount of CO₂ that has been collected. CO₂ reductions can make any process carbon neutral.

Air capture does not prevent efficiency improvements, fuel switching, or a transition to non-fossil fuels. If it proves more efficient to collect CO₂ at the site of an emission, or to simply avoid the emission altogether, the air capture effort is simply redirected to another emitter. Thus, air capture becomes the capture of last resort, which can even be applied to past emissions. The fraction of the market it would likely take on will directly depend on its cost.

Air capture will first be deployed for emissions that are otherwise hard to avoid. These include the emissions from mobile sources like airplanes and automobiles. However, it can also play a role in the carbon management at existing power plants, many of which will be very difficult to retrofit, either because of the layout of the plant or because of a lack of a good CO₂ disposal site nearby. Even for plants that are fitted with carbon capture technology, the residual emissions could be eliminated by air capture. In principle, air capture can keep up with the entire world emissions, and could even be used to reduce the CO₂ content of the atmosphere.

Gasoline, diesel and jet fuel are such powerful and convenient energy carriers that it is hard to see how they might be displaced by other energy carriers, unless environmental concerns make their use entirely undesirable. Batteries, ultra capacitors, flywheels etc. all struggle to approach 1 MJ/kg of energy storage, whereas liquid hydrocarbon fuels have energy densities of 40–50 MJ/kg.

In the face of severe carbon constraints, air capture provides a way for oil producers to protect the value of the oil in the ground. The main reason why oil receives a high premium over coal is that it can be readily transformed into a transportation fuel. If carbon constraints were to eliminate this use, the value of oil would drop back to the value of coal. By allowing the continued use of liquid fuels in the transportation sector, air capture maintains the value of oil in the ground.

Air capture also offers an opportunity for renewable energy to enter large-scale into the transportation sector. If the cost of non-fossil electricity drops to about 2¢/kWh, it becomes feasible to use H₂O and CO₂ as feedstock to produce synthetic hydrocarbon fuels like methanol, dimethyl-ether or long chained alkanes like gasoline or diesel.²² Most electricity is more expensive, but for example in Norway and Iceland electricity costs are already sufficiently low.

Bio-fuels already provide an example of air capture combined with the production of synthetic fuels. However, technical air capture has advantages. It requires far less land, and does not compete for land with agriculture for food production. A single air collector that can handle a ton a day and fits in a standard shipping container would roughly match agricultural rate biomass production that take up a hectare of land.

Introduced at a sufficiently large scale, air capture can be used to drive the CO₂ content of the air back down to earlier levels. Once it becomes possible to compensate for all emissions in a hydrocarbon based transportation sector, then it should also be possible to increase the size of the collection effort sufficiently to create a noticeable reduction in CO₂ in the atmosphere. Since it is very unlikely that political changes in the world at large will manage to stop the increase in atmospheric CO₂ concentrations rapidly, this feature of air capture will become of increased importance in the future.

In the public press, air capture has been lumped with geo-engineering approaches to climate change.²³ Air capture technology would indeed turn into geo-engineering, if in the future it would be used to hold the CO₂ content of the atmosphere at a desired level against natural driving forces. However, as long as air capture is only used as a means of compensating for various CO₂ emissions, it is a carbon management strategy like any other. Its advantage is that it is not bound to a particular location and time and thus can compensate for any emission. Air capture technology, once successfully developed, provides a valuable option in the fight against climate change. It can be built up and deployed without interfering with the existing energy infrastructure. For it to be ready in time requires policy decisions now.

²² The energy content of gasoline is about 10 kWh/liter, hence at 2¢/kWh and 50% conversion efficiency, the electricity cost is about 40¢/liter.

²³ See for example the Guardian, September 1, 2008, <http://www.guardian.co.uk/environment/2008/sep/01/climatechange.endangeredhabitats> (last visited February 17, 2009).

I would like to thank Allen Wright for starting and running a research and development company that made this great adventure possible. My deepest gratitude goes to the late Gary Comer who cared enough about climate change to fund the company and its research. Lastly, I would like to thank Wally Broecker for his enthusiastic support and sage advice.

References

1. G. Astarita, *Mass Transfer with Chemical Reactions* (Elsevier Publishing Company, Amsterdam, London, New York, 1967)
2. R. Baciocchi, G. Storti, M. Mazzotti, Process design and energy requirements for the capture of carbon dioxide from air, *Chem. Eng. Proc.* **45**, 1047–1058 (2006)
3. A. Bandi, M. Specht, T. Weimer, K. Schaber, CO₂ recycling for hydrogen storage and transportation - Electrochemical CO₂ removal and fixation, *Energy Conv. Manag.* **36**, 899–902 (1995)
4. S. Elliott, K.S. Lackner, H.-J. Ziock, M.K. Dubey, H.P. Hanson, S. Barr, N.A. Ciszkowski, D.R. Blake, Compensation of Atmospheric CO₂ Buildup through Engineered Chemical Sinkage, *Geophys. Res. Lett.* **28**, 1235–1238 (2001)
5. M.L. Gray, K.J. Champagne, D. Fauth, J.P. Baltrus, H. Pennline, Performance of immobilized tertiary amine solid sorbents for the capture of carbon dioxide, *Int. J. Greenh. Gas Contr.* **2**, 3–8 (2008)
6. R. Joosten, J. Schumacher, C. Wirth, A. Schulte, Evaluating tree carbon predictions for beech (*Fagus sylvatica* L.) in western Germany, *Forest Ecol. Manag.* **189**, 87–96 (2004)
7. D. Keith, M. Ha-Duong, J. Stolaroff, Climate Strategy with CO₂ Capture from the Air, *Climat. Change* **74**, 17–45 (2006)
8. K.S. Lackner, *Thermodynamics of the Humidity Swing Driven Air Capture of Carbon Dioxide* (GRT LLC, Tucson, AZ, 2008)
9. K.S. Lackner, H.-J. Ziock, P. Grimes, *Carbon Dioxide Extraction from Air: Is it an Option?* Proceedings of the 24th International Conference on Coal Utilization & Fuel Systems (Clearwater, Florida, 1999a)
10. K.S. Lackner, P. Grimes, H.-J. Ziock, *Carbon Dioxide Extraction from Air?* (Los Alamos National Laboratory, LAUR-99-5113, Los Alamos, NM, 1999b)
11. K.S. Lackner, C.H. Wendt, Exponential Growth of Large Self-Reproducing Machine Systems, *Math. Comp. Modell.* **21**, 55–81 (1995)
12. B. Metz, O. Davidson, H.d. Coninck, M. Loos, L. Meyer, *IPCC Special Report on Carbon Dioxide Capture and Storage* (Cambridge University Press, New York, 2005)
13. F. Montagnini, P.K.R. Nair, Carbon sequestration: An underexploited environmental benefit of agroforestry systems, *Agroforest. Syst.* **61**, 281–295 (2004)
14. T.E. Notteboom, Container Shipping and Ports: An Overview, *Rev. Network Econom.* **3**, 86–106 (2004)
15. D. Sarewitz, R. Nelson, Three rules for technological fixes, *Nature* **456**, 871–872 (2008)
16. M. Specht, A. Bandi, C.U. Maier, J. Schwarz, Energetics of solar methanol synthesis from atmospheric carbon dioxide compared to solar liquid hydrogen generation, *Energy Conv. Manag.* **33**, 537–543 (1992)
17. N.A. Spector, B.F. Dodge, Removal of carbon dioxide from atmospheric air, *Trans. Amer. Inst. Chem. Eng.* **42**, 827–848 (1946)
18. M. Steinberg, V.-D. Dang, Production of synthetic methanol from air and water using controlled thermonuclear reactor power—I. technology and energy requirement, *Energy Conv.* **17**, 97–112 (1977)
19. J.K. Stolaroff, Capturing CO₂ from ambient air: a feasibility assessment. Pittsburgh, Carnegie Mellon University, Doctor of Philosophy **85** (2006)
20. J.K. Stolaroff, D.W. Keith, G.V. Lowry, Carbon Dioxide Capture from Atmospheric Air Using Sodium Hydroxide Spray, *Environm. Sci. Technol.* **42**, 2728–2735 (2008)
21. S. Stucki, A. Schuler, M. Constantinescu, Coupled CO₂ recovery from the atmosphere and water electrolysis: Feasibility of a new process for hydrogen storage, *Int. J. Hydr. Energy* **20**, 653–663 (1995)
22. F.S. Zeman, *Air Extraction: The Feasibility of Absorbing CO₂ directly from the Atmosphere. Earth and Environmental Engineering*, Columbia University, Fu Foundation School of Engineering and Applied Sciences, New York, Doctor of Engineering Science **179** (2006)
23. F.S. Zeman, K.S. Lackner, Capturing Carbon Dioxide Directly from the Atmosphere, *World Res. Rev.* **16**, 157–172 (2004)