

CO₂ CAPTURE FROM THE AIR: TECHNOLOGY ASSESSMENT AND IMPLICATIONS FOR CLIMATE POLICY

David W. Keith and Minh Ha-Duong

Department of Engineering and Public Policy,
Carnegie Mellon University, Pittsburgh, PA, 15213, USA.
keith@cmu.edu

INTRODUCTION

It is physically possible to capture CO₂ directly from air and immobilize it in geological structures. Today, there are no large-scale technologies that achieve air capture at reasonable cost. Yet, several strong arguments suggest that it will be comparatively easy to develop practical air capture technologies on the timescales relevant to climate policy [1, 2]. Moreover, we argue that air capture has important structural advantages over more conventional mitigation technologies which suggest that, in the long-run, air capture may be cheaper than some other methods of achieving deep reductions in net CO₂ emissions.

Air capture may be viewed as a hybrid of two related mitigation technologies. Like carbon sequestration in ecosystems, air capture removes CO₂ from the atmosphere, but it is based on large-scale industrial processes rather than on changes in land use, and it offers the possibility of near-permanent sequestration of carbon. It is also possible to use fossil fuels with minimal atmospheric emissions of CO₂ by capturing the carbon content of fossil fuels while generating carbon-free energy products, such as electricity and hydrogen, and sequestering the resulting carbon. We call this Industrial Carbon Management (ICM). Like ICM, air capture involves long-term storage of CO₂, but unlike ICM air capture removes the CO₂ directly from the atmosphere and so manipulates the global atmospheric concentration rather than the exhaust stream of large fixed-point sources such as power plants.

AIR CAPTURE: A BOUNDING ANALYSIS

An air capture system (as we define it) removes CO₂ from the air and delivers a pure CO₂ stream for sequestration. In general, air capture systems will use some sorbent material that selectively captures CO₂. The sorbent is then regenerated to yield concentrated CO₂ and fresh sorbent ready to be used for capture. A contacting system is necessary to expose the sorbent to fresh air. At the simplest, a pool of sorbent open to the air might serve as a contactor. Active contacting systems might look like water-air heat exchangers that are associated with power plants in which fans blow air past a system of flowing sorbent. Energy is required for regeneration, to pump sorbent and air through the contacting system, and to compress the CO₂ stream for pipeline transport or sequestration. The energy could be supplied by any source: fossil, solar, or nuclear systems are all plausible.

An air capture system that removed CO₂ from the air at the rate that it is emitted from a 1 GW coal fired power plant, might look much like a power plant. The physical plant would be about as large and not much more expensive. Instead of the one or two cooling towers required for a conventional power plant, the air capture plant might require about 10 contacting units of similar size and shape as cooling towers (~100 m in height and diameter) evenly distributed over a few square kilometers.

Extraction of pure gases from air is more than theory: oxygen, nitrogen, and argon are commercially produced by capture from air. The crucial questions about air capture are the cost and energy inputs. This section estimates first a lower bound on the inputs needed for air capture, describing the physical limits to the use of energy and land. It then compares the cost of air capture to the cost of capture from fossil-fuel energy plants and air capture systems that could plausibly be built with existing technologies.

Thermodynamics provides a lower-bound on the energy required for air capture. The minimum energy needed to extract CO₂ from a mixture of gases in which the CO₂ has an partial pressure p_0 and to deliver it as a pure CO₂ stream at final pressure p is set by the enthalpy of mixing, $kT \ln(p/p_0)$, where k is the Boltzmann constant (8.3 J/mol-K) and T is the working temperature. At a temperature of 300 K (27 C) that would be typical for extraction of CO₂ from the atmosphere, kT is about 2.5 kJ/mol. The minimum energy required to capture CO₂ from the air at a partial pressure of 4×10^{-4} atm and deliver it at one atmosphere is therefore about 20 kJ/mol or 1.6 GJ/tC. If we add the energy required for compressing the CO₂ to the 100 atm pressure required for geological sequestration (assuming a 50% efficiency for converting primary energy to compressor work) the overall energy requirement for air capture with geologic sequestration is about 4 GJ/tC.

The ~4 GJ/tC minimum may be compared to the carbon-specific energy content of fossil fuels: coal, oil, and natural gas have about 40, 50, and 70 GJ/tC respectively. Thus if the energy for air capture is provided by fossil fuels, then the increase in consumption of fuel required to capture the CO₂ emissions arising from fuel combustion can—in principle—be relatively small (<10%).

Now consider the requirement for land. Land-use is an important constraint for energy technologies in general, and is a particularly important constraint for biological methods of manipulating atmospheric carbon that may compete with air capture. An air capture system will be limited by the flux of CO₂ that is transported to the absorber by atmospheric motions; even a perfect absorber can only remove CO₂ at the rate at which it is carried to the device by advection (winds) and diffusion. At large scales (100's of km), CO₂ transport in the atmospheric boundary layer limits the air capture flux to roughly 400 tC/ha-yr [1].

This rate is more than 20 times the rate at which the most productive ecosystems can fix carbon. Assume air capture is used to offset emissions from fossil fuels as a means to provide energy with zero net CO₂ emissions. We can then divide the power provided by the fossil fuels by the land area required to capture the CO₂ emission resulting from the fuel combustion, in order to compute a power density. This is the effective density at which fossil fuels with air capture provide power with zero net emissions, it may be directly compared with the power densities of alternative CO₂-neutral energy systems.

If coal is used as the fuel, then an air-capture/coal system can provide a CO₂-neutral energy flux of 50 Wm⁻² (the value would be almost twice this for an air-capture/natural-gas system). Both wind power and biomass based systems can produce roughly 1 Wm⁻², and even solar power which is constrained much more strongly by cost than by land use can deliver only ~20 Wm⁻².

There is a strong analogy between the land requirements for air capture and for wind power. Both depend on the rate of turbulent diffusion in the atmospheric boundary layer: diffusion of CO₂ for air capture and diffusion of momentum for wind power. In each case, large scale processes limit the average flux. Large wind farms, for example, must space their turbines about 10 rotor-diameters apart to avoid "wind shadowing" by allowing space for momentum to diffuse downward from the fast moving air above. Similarly for an air capture system, the individual units must be spaced far enough apart to ensure that each receives air with near-ambient CO₂-concentration. In each case the footprint of the actual hardware can be limited to only a tiny fraction of the required land because the footprint of an individual wind turbine or air capture system can be small, and the land between the units can be preserved for other uses (e.g. farming). If computed using only the footprints of the individual air capture units, CO₂-neutral flux can be many 100's of Wm⁻².

The large effective power densities of air capture, however computed, make it implausible that land use could be a significant constraint on the deployment of air capture.

Capture from air vs capture at power plants

Almost all the literature on industrial CO₂ capture and sequestration has addressed the problem of capturing CO₂ from large centralized facilities such as electric power plants. It is therefore instructive to compare the cost of air capture with the cost of capture from these sources via Industrial Carbon Management (ICM).

First, siting issues are less acute for air capture facilities. The location of an ICM plant is constrained by three transportation requirements: fuel must be transported to the plant, CO₂ from the plant to a suitable storage site, and finally the carbon-free energy products—electricity or hydrogen—to users. The location of an air capture plant is less constrained: there is no final energy product, and the energy inputs per-unit-of-CO₂-output are less than half, and perhaps as little as a 1/10 of that needed for an ICM plant. Moreover, air capture plants will likely be located at CO₂ sequestration sites, eliminating the CO₂ transport cost.

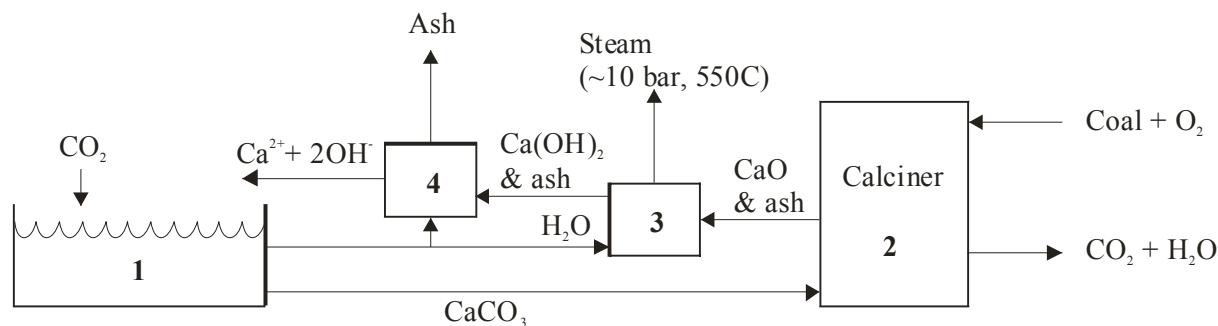
While the cost of transportation is hard to quantify, there is little doubt that reduction in transportation requirements lowers the cost of air capture compared to ICM. Perhaps more importantly, it lowers the barriers posed by the siting of new energy infrastructure. Public resistance to the construction of new energy transportation infrastructure such as gas and electric transmission systems is already a serious problem; the development of ICM would only increase these difficulties. It is plausible that the reduction of these barriers is one of the most important attributes of air capture.

Second, air capture systems can be built big, taking maximum advantage of returns-to-scale, because air capture need not be tightly integrated into existing energy infrastructures. One of the features that makes ICM so appealing is its compatibility with existing fossil energy systems: CO₂ capture may be thought of as a retrofit of the energy system. But this feature also limits the rate of its implementation and the scale of the individual ICM facilities. An power plant with CO₂ capture may need to be located and sized to replace and existing power plant. Too rapid implementation of ICM will raise adjustment costs. Air capture, in contrast, is more loosely coupled to the existing system: it is not an intermediate but a final energy use. This implies that the air capture facilities can be optimally sized to suit geology and technology, and can be constructed rapidly if required. In this respect air capture is more like geoengineering than it is like conventional mitigation [3].

Finally, air capture differs from ICM because it effectively removes CO₂ with equal ease or difficulty from all parts of the economy. This is its most important feature. In this section we have argued that the long-run cost of air capture may be quite close to the cost of capture from power plants. But even if the cost of abatement with air capture is considerably more than the cost of abatement at large centralized facilities, air capture still has the unique ability to provide abatement across all economic sectors at fixed marginal cost. Air capture operates on the heterogeneous and diffuse emission sources in the transportation and building sectors where the cost of achieving deep emissions reductions by conventional means are much higher than they are for centralized facilities.

Air capture with calcium hydroxide

Lackner et. al. and Elliott et al. [1, 2] have proposed a direct CO₂ capture system using calcium hydroxide. Figure 1 shows a schematic diagram of such a system. Aqueous calcium hydroxide (Ca(OH)₂) is used as a sorbent for CO₂ producing calcium carbonate CaCO₃. The carbonate is then heated in a closed vessel producing CaO and CO₂. Finally, the CaO is returned to the water regenerating the sorbent and producing heat.



		kJ/mol-C	GJ/tC
1.	$\text{CO}_2(\text{g}) + \text{Ca}^{2+} + 2\text{OH}^- \rightarrow \text{CaCO}_3(\text{s}) + \text{H}_2\text{O}(\text{l})$	-96	-8
2.	$\text{CaCO}_3(\text{s}) \rightarrow \text{CaO}(\text{s}) + \text{CO}_2(\text{g})$	179	15
3.	$\text{CaO}(\text{s}) + \text{H}_2\text{O}(\text{l}) \rightarrow \text{Ca}(\text{HO})_2$	-65	-5
4.	$\text{Ca}(\text{HO})_2 \rightarrow \text{Ca}^{2+} + 2\text{OH}^-$	-17	-1

Figure 1: Air capture with calcium hydroxide.

The numbers on the schematic diagram in Figure 1 correspond to the four chemical reactions listed in the table. At constant pressure and temperature no energy is required to complete the cycle, so the sum of the enthalpy changes is zero. In practice, the energy released in reaction (1)—the contactor—goes into heating the water in the absorption system and is therefore lost. The heat required for calcination (2) is the primary energy cost of the process. Most of the energy released in reaction (3)—called slaking—could, in principle, be recovered either in a separate reaction vessel (as shown schematically in the figure) or in the lower temperature section of the calcination reactor. Heat from the dissolution reaction (4) cannot be effectively recovered because of the temperature rise at dilution is small even for a saturated solution.

The contacting system is limited by the diffusion of CO₂ through the air to the water surface, and for some designs (such as the passive pool shown here) by diffusion in the water. For simplicity, the absorption system is shown as a passive pool in the figure, however, forced-draft systems may be preferred.

The advantage of the calcium hydroxide system is its technological simplicity. All the required chemicals are cheap, common, and environmentally benign. The recovery of CO₂ is accomplished by calcination, the same reaction that is used to make cement from limestone. This process is well understood: as the central step of cement manufacture, calcination is responsible for approximately 3% of anthropogenic CO₂ emissions. Strong calcium hydroxide solutions are easily handled in the collection system by conventional materials.

The disadvantage of the calcium hydroxide sorbent is the high energy of the CaO-CO₂ bond. It requires 15 GJ/tC of energy at 1000 C to efficiently break the bond in the calcination step. If none of this energy can be recovered the system thus requires about 10 times the theoretical minimum energy input describe above. This is about 37% of the energy available in coal per unit carbon. Thus, if the system was run with coal each ton-carbon worth of coal would capture only 2.7 tC of CO₂.

The high energy cost of the CaO system, may make it unlikely to succeed as a means of air capture. Better sorbents can likely be devised. Nevertheless, the CaO system serves as a convenient benchmark. There is little doubt that a working system could be designed today. What is uncertain is its cost and performance.

While a detailed study of the cost of capture using such a system had not yet be performed, there is sufficient commonality between the required components and existing technologies to permit a rough estimate. Calcination could likely be accomplished in a modified pressurized fluidized bed combustor (PFBC) burning coal in a CO₂ rich atmosphere with oxygen supplied by an air separation unit. PFBC systems have been demonstrated at capacities of over 100 MW. Such systems already include CaCO₃ injection for sulfur control, and operate at suitable temperatures and pressures for calcination.

The calciner is therefore much like a PFBC power plant using pure oxygen with CO₂ recycle (oxy-fuel combustion), except that air capture plant does not require a boiler, steam turbine, or generator since little net heat is produced, but given the 2:1 ratio of captured CO₂ to CO₂ generated from the coal the system requires a CO₂ drying and compression system, and a solids handling system, that is three times larger than the reference oxy-fuel power plant. Existing design studies of pulverized coal power plants with oxy-fuel combustion and CO₂ capture provide many of the necessary component costs.

The cost of the CO₂ absorber system is much more uncertain. Based on the cost of forced draft water/air heat exchangers, it appears that cost of the absorber system would be only 1/4 to 1/10 the cost of the calciner. A rough estimate based on these assumptions suggests the capital cost of the calciner is around 1000\$ for each kg-CO₂/hr of capacity. (Using consistent assumptions, the an oxy-fuel power plant with CO₂ capture would cost around 2000 \$/kWe.)

Combining these systems, and assuming that some heat can be recovered from CaO dissolution (reaction 2), we judge that such a system might plausibly provide air capture at an over all cost of under 200 \$/tC within a few decades.

AIR CAPTURE CHANGES OPTIMAL CLIMATE POLICY

The two previous examples suggest that air capture is possible, and that—in the long run—it may be achieved at cost of order 150\$ per ton of carbon. The cost of air capture is uncertain, but not necessarily more uncertain than the cost of more conventional emissions abatement technologies if we extend our view many decades into the future. In the remainder of this paper we assume that air capture will eventually be available at this cost and explore the consequences for climate policy.

In this study a stylized integrated assessment model, DIAM, is used to compare optimal global CO₂ strategy with and without air capture. The details of the DIAM model are described by Ha-Duong et al [4]. Details about the exploration of air capture using DIAM are described in a manuscript available from the authors. Only a brief description is provided here.

DIAM does not represent explicit individual technologies or capital turnover, but instead the inertia related to induced technical change. The inertia of the worldwide energy system induces adjustment costs, related to

the rate of change of abatement. The model maximizes the expected discounted inter-temporal sum of utility under the risk of abrupt adverse climate-change impact.

We acknowledge that this cost-benefit optimization framework lacks a special claim to universality, and that it obscures many of the distributional issues that drive climate politics. We adopt it nevertheless, to explore the effect of removing the irreversibility of CO₂ emissions accumulation via air capture, because of its importance as a reference for considering long-term climate policy.

DIAM represents the uncertainty in the benefits of avoiding climate change, or alternatively the cost of climate impact, a non-linear damage function. The non-linearity appears as a smooth step in the damage function where that the step starts at about 600 ppmv. Rather than imposing a fixed stabilization target, this formulation allows a cost-benefit tradeoff. But the kink in the damage function serves as a soft concentration ceiling, and the location of the kink is therefore a critical parameter.

The impact is a function of atmospheric carbon dioxide concentration. While it is measured in monetary units, it represents a global willingness to pay to avoid the given level of climate change, including non-market values. The impact at any date is defined as a fraction of wealth at this date. Therefore it scales over time with the size of the economy.

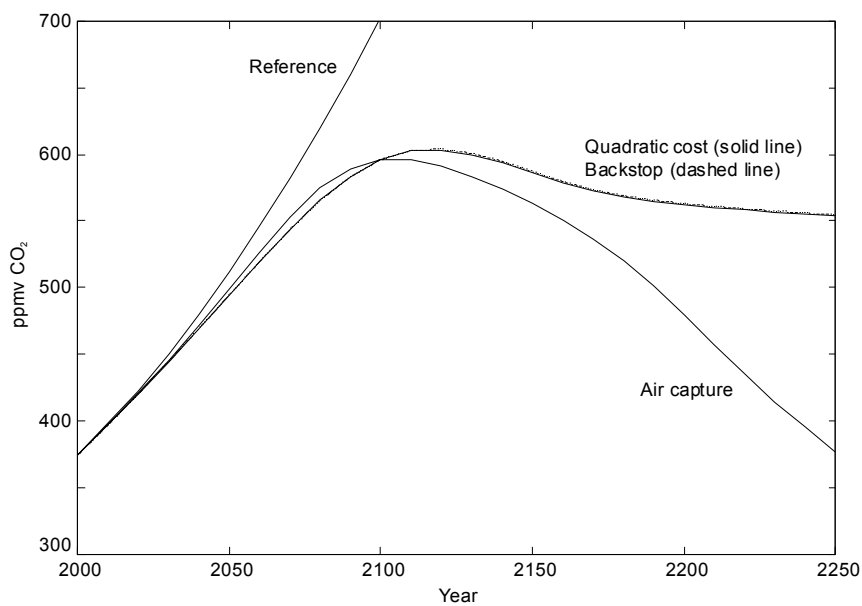


Figure 2: Optimal CO₂ concentration trajectories.

The model was used to compare three scenarios for the cost of mitigation (Figure 2). In the first, “Quadratic abatement cost,” the marginal cost curve is a simple linear ramp culminating at 250 \$/tC for complete emission abatement. Average abatement costs are therefore quadratic. In the second scenario, “backstop,” a backstop energy technology is available but not direct capture, so that the marginal cost ramp is capped at 150 \$/tC. Finally, the “air capture” scenario looks like the backstop scenario, except that it allows the 150 \$/tC plateau on the marginal cost curve to extend to beyond complete emissions abatement, allowing total emissions to be negative.

The backstop technology has almost no effect on the emission trajectory. The most important reason is that non-linearity in the impact function serves as a soft constraint on concentration, as discussed above. Therefore, even if ICM lowers the cost, the effect on optimal emissions remain small. Considering the optimum as the intersection of marginal cost and benefit curves, this is a situation where the marginal benefit of abatement is steep, so the optimal abatement quantity is relatively insensitive to changes in cost.

Carbon dioxide is a long-lived stock pollutant. Technical progress, the decreasing energy intensity of the economy and the effect of adjustment costs make it comparably cheaper to reduce emissions in the future. These factors explain why the optimal trajectory overshoots. Also, an inverted Malthus effect is at play. On one hand, climate change impact is a fraction of wealth. Therefore the willingness to pay to solve the problem grows exponentially. On the other hand, the costs of abatement are proportional to the amounts of pollution generated in the business-as-usual scenario in which emissions are assumed to grow only up to 2100 and linearly.

When air capture is available concentration declines rapidly toward preindustrial levels. Without air capture the dynamics that dictate a return toward a low concentration target remain as described above, but the rate of CO₂ concentration decline is determined by natural removal. With air capture, the rate is determined by the tradeoff between its costs and the benefits of reducing climate change.

This idea is overlooked in the existing literature. Admittedly, the decrease is mostly happening in the twenty-second century, while most scenario studies reasonably focus on this one. Yet the factors driving the return toward preindustrial concentration—the inverted Malthus effect—are general, as is the environmental Kuznets curve. This result suggests that if technology exists to reverse climate damage it will be employed.

CONCLUSIONS

No air capture systems exist today, yet we argued that they could plausibly be built with adaptations of existing technology, and that in the long run, the ability to capture CO₂ from the air fundamentally alters the dynamics of climate policy. We do not claim that air capture systems are trivial to build today, nor that they will play a quantitative role in the next few decades.

Air capture differs from conventional mitigation in three key aspects. First, it removes emissions from any part of the economy with equal ease or difficulty. Consequently, its price caps the cost of mitigation with a scope unmatched by any other kind of abatement technology. Second, because air capture allows the removal of CO₂ after emission it permits reduction in concentrations more quickly than can be achieved by the natural carbon cycle. Third, because it is weakly coupled to the energy system, air capture may offer stronger returns-to-scale and lower adjustment costs than conventional mitigation options.

Air capture is (arguably) a form of geoengineering because it directly modifies the biosphere and would be implemented with the aim of counterbalancing other human actions [3]. Like geoengineering, its availability reduces our vulnerability to some high-consequences low-probability events. In an optimal sequential decision framework the consequence is a decrease in the need for precautionary short-term abatement. Because air capture may provide some insurance against climate damages, it presents a risk for public policy: the mere expectation that air capture can be achieved reduces the incentive to invest in mitigation. Yet, while air capture removes irreversibility in CO₂ concentration increase, it does not protect against irreversibilities in the climate system's response to forcing.

While air capture may reduce the amount of mitigation in the short run, it can increase it on longer time-scales. If air capture is possible, even at comparatively high cost, and if the willingness to pay for climate change mitigation grows with the economy, then the optimal trajectory follows an environmental Kuznets curve. At some point the optimal target will be to return atmospheric greenhouse gases concentration to lower levels. These may be even lower than present-day levels. Air capture changes the temporal dynamics of mitigation by making this response possible.

ACKNOWLEDGMENTS

Research supported by the Centre National de la Recherche Scientifique, France and by the Center for Integrated Assessment of Human Dimensions of Global Change. This Center exists via a cooperative agreement between the National Science Foundation (SBR-9521914) and Carnegie Mellon University. We thank Hadi Dowlatabadi, Granger Morgan, Ted Parson, Hans Ziock, Klaus Keller and Alex Farrell for useful comments.

REFERENCES

1. Elliott, S., et al., Compensation of atmospheric CO₂ buildup through engineered chemical sinkage. *Geophysical Research Letters*, 2001. **28**, 1235–1238.
2. Lackner, K.S., P. Grimes, and H.-J. Ziock., *Carbon dioxide extraction from air?* 1999, Los Alamos National Labs. www.lanl.gov/energy/ecology/carbon/docs/AirExtraction.pdf.
3. Keith, D.W., Geoengineering the Climate: History and Prospect. *Ann. Rev. Energy and Environ.*, 2000. **25**, 245-284.
4. Ha-Duong, M., M.J. Grubb, and J.-C. Hourcade, Influence of socioeconomic inertia and uncertainty on optimal CO₂-emission abatement. *Nature*, 1997. **390**, 270–274.