Capturing CO_2 from ambient air: a feasibility assessment

Thesis by

Joshuah K. Stolaroff

In Partial Fulfillment of the Requirements
for the Degree of
Doctor of Philosophy

Carnegie Mellon University
Pittsburgh, PA

2006
(Defended August 17, 2006)

2006, Joshuah K. Stolaroff

This work is licensed under the terms of the Creative Commons Attribution-NonCommercial-NoDerivs 2.5 License. You are free to copy and distribute this work in its original form for noncommercial purposes. View the full license at http://creativecommons.org/licenses/by-nc-nd/2.5/

For my family, friends, and everyone who will be hurt by climate change.

Acknowledgments

When I was deciding which graduate school to attend, I asked several Civil and Environmental Engineering students at Carnegie Mellon (CMU) about how the students related with each other, whether they were competitive, whether they tend to help each other out, and so forth. Then-graduate-student Joe Bushey replied that the general feeling is that "we're all trying to save the world here, so we might as well do it together". After four years at CMU, this is the best way I can describe the interactions I have had with colleagues in my departments and elsewhere on campus. I can't think of a time, in this wonderfully collaborative and cooperative environment, when I was denied assistance by anyone or when I declined to give assistance when asked. The seminars, reading groups, and informal talks with students and professors have stimulated my thinking, broadened my knowledge, and improved my research. I must thank my superlative colleagues for making my time at CMU enjoyable and deeply satisfying.

The spirit of cooperative assistance is exemplified by my friend and colleague Jeffrey Pierce, who took nights out of his vacation to write code for, and help me use, a model of drop coalescence in the late days of this thesis, handily navigating my complete inexperience with Fortran 77 and mounting deadline-related desperation.

I must thank Professor Granger Morgan for his unwavering enthusiasm (even, perhaps, when my own wavered) and support for my research and his wisdom in many subjects. I thank Professor Jay Apt, also for his enthusiasm, and for providing me with many opportunities to present my findings. I thank Peter Adams for some very key research help, particularly on modeling coalescence, but on other topics as well. I also thank Professor Edward Rubin for adamantly pointing out a weakness in my research which, in large part, pushed me to model coalescence in the first place.

I owe the success of the prototype experiments partly to two excellent, hard-working undergraduates at the University of Calgary: Kenton Heidel and Leif Menezes. Without their labors, the completion of those experiments would not have been possible in the short summer I spent in Calgary. The construction of that prototype was also aided by practical advice from Professor Larry Cartwright, who I thank for demonstrating in many subtle (as well as obvious) ways what it is to be an engineer.

Of course I owe the largest debt of gratitude to my advisors. I thank Professor Greg Lowry for the years of help and support, particularly his open-door policy and frequent guidance on practical matters and for his remarkable experimental insight. The same thanks go to Professor David Keith, particularly for challenging me and for having big ideas and inviting my opinions on them.

The research presented in this thesis was funded from a variety of sources: the Pittsburgh Infrastructure Technology Alliance (PITA), the Carnegie Mellon Seed Fund, the Canada Foundation for Innovation, and the Climate Decision Making Center (CDMC), which was created through a cooperative agreement

between the National Science Foundation (SES-0345798) and Carnegie Mellon University.

Abstract

In order to mitigate climate change, deep reductions in CO_2 emissions will be required in the coming decades. Carbon capture and storage will likely play an important role in these reductions. As a compliment to capturing CO_2 from point sources, CO_2 can be captured from ambient air ("air capture"), offsetting emissions from distributed sources or reducing atmospheric concentrations when emissions have already been constrained. In this thesis, we show that CO_2 capture from air is physically and thermodynamically feasible, discuss the various routes available, and explain why NaOH solution is a viable sorbent for large-scale capture. An example system using NaOH spray is presented. With experimental data and a variety of numerical techniques, the use of NaOH spray for air capture is assessed and an example contacting system developed. The cost and energy requirements of the example contacting system are estimated. Contactor estimates are combined with estimates from industry and other research to estimate the cost of a complete air capture system. We find that the cost of capturing CO_2 with the complete system would fall between 80 and 250 \$/t-CO_2, and improvements are suggested which reduce the upper-bound cost to 130 \$/t-CO_2. Even at the high calculated cost, air capture has implications for climate policy, however dedicated engineering and technological innovation have potential to produce much lower-cost systems.

Contents

Acknowledgments iv

Abstract vi

1 Introduction 1
- 1.1 Carbon Capture and Storage . . . 2
- 1.2 Air capture . . . 3
- 1.3 Thermodynamic and physical limits of air capture . . . 4
- 1.4 Routes to air capture . . . 5
  - 1.4.1 Organic carbon production . . . 5
  - 1.4.2 Metal-carbonate production . . . 5
  - 1.4.3 Capture with a regenerated sorbent . . . 5
  - 1.4.4 Metal hydroxide sorbents . . . 6
- 1.5 Research Objectives . . . 7

2 Example systems 8
- 2.1 Contacting with NaOH . . . 8
- 2.2 Caustic recovery . . . 10
- 2.3 Example air capture system . . . 11
  - 2.3.1 Overview . . . 11
  - 2.3.2 Caustization . . . 12
  - 2.3.3 Calcination . . . 13
  - 2.3.4 Integrated system . . . 14
  - 2.3.5 Energy requirements . . . 14

3 Contactor 17
- 3.1 Materials and Methods . . . 17
  - 3.1.1 Theoretical methods . . . 17
  - 3.1.2 Experimental methods . . . 19
- 3.2 Results . . . 21
  - 3.2.1 Mass Transfer . . . 21
  - 3.2.2 Energy requirements . . . 23

3.2.3 Water loss . . . 24
3.3 Scale-up of prototype results . . . 26
3.3.1 CO_2 depletion in air . . . 27
3.3.2 Changing drop size due to evaporation . . . 27
3.3.3 Spray droplet collision, coalescence, and breakup . . . 27
3.4 Contactor Cost . . . 37
3.4.1 Mass transfer . . . 37
3.4.2 Capital cost . . . 38
3.4.3 Operating cost . . . 40
3.4.4 Scenarios . . . 40
3.4.5 Total cost . . . 41
3.5 Contactor technology and sensitivity of future cost . . . 43
3.5.1 Spray technology . . . 43
3.5.2 Structural design . . . 43
3.5.3 Water loss . . . 46
3.5.4 Siting . . . 47
3.5.5 Materials and construction cost . . . 47
3.5.6 Solids formation – scaling and clogging . . . 49
3.6 Conclusions from contactor analysis . . . 49

4 Cost of air capture **51**
4.1 Lower bound . . . 51
4.2 Cost of example system . . . 52

5 Discussion **57**
5.1 Findings and implications . . . 57
5.2 Lessons for assessing of future energy technologies . . . 59

Bibliography **61**

A List of Symbols **66**

B Experimental details and procedure **69**
B.1 Physical apparatus . . . 69
B.1.1 Basic size and structural design . . . 69
B.1.2 Materials compatibility . . . 75
B.1.3 Air handling and air safety . . . 75
B.1.4 Liquid handling . . . 77
B.1.5 Measurement . . . 78
B.2 Experimental Procedure . . . 80
B.3 Data Analysis . . . 81

List of Figures

1.1 Global CO_2 emissions over time. . . . 2

2.1 Example air capture system (top level process diagram). . . . 11

3.1 Diagram and photograph of prototype contactor. . . . 20
3.2 CO_2 absorption by falling drops – comparison of theory and measurements. . . . 21
3.3 Outlet CO_2 concentration during a typical trial. . . . 22
3.4 CO_2 absorption for several solution concentrations of NaOH and two nozzles. . . . 23
3.5 Water loss measured in the prototype and calculated water loss in a full-scale system . . . 25
3.6 Size distribution of sprays used in coalescence modeling. . . . 30
3.7 Total surface area of a parcel of spray over time in the contactor. . . . 31
3.8 Size distributions over time for model calculation in Figure 3.7. . . . 32
3.9 Initial and steady-state size distributions of spray in CFSTR model of coalescence. . . . 33
3.10 Spray surface area as a function of liquid flow rate . . . 34
3.11 Spray surface area as a function of contactor height. . . . 35
3.12 Surface area as a function of liquid flow rate for a smaller drop size distribution. . . . 36
3.13 Simple diagram of a horizontal flow contactor . . . 46
3.14 Equilibrium speciation of carbonates. . . . 48

B.1 Dimensions, layout, and labeling of major components of the final prototype design. . . . 70
B.2 Photograph of completed prototype structure. . . . 71
B.3 Lining the inside of the Sonotubes with PVC sheets. . . . 72
B.4 Photograph of the reaction chamber being lifted by crane. . . . 73
B.5 Tower support structure . . . 74
B.6 Flow diagram of liquid and air systems in the final prototype design. . . . 76
B.7 Sampling points of CO_2 concentration in air. . . . 79
B.8 Measured CO_2 concentration during a portion of Trial 2. . . . 82
B.9 CO_2 absorbed during Trial 2: comparison of two measurement methods. . . . 83
B.10 Measured CO_2 concentration during a portion of Trial 11. . . . 84

List of Tables

2.1 Chemistry of example air capture system. 10
2.2 Chemistry of autocaustization process. 10
2.3 Chemistry of caustic recovery with titanates. 11
2.4 Energy requirements of the air capture system by component. 15

3.1 Energy requirements of contactor. 24
3.2 Capital cost of cooling towers. 39
3.3 Cost estimates for the contactor. 42

4.1 Input parameters for system cost estimate 53
4.2 Cost of example system and improved system by component 55
4.3 Sensitivity of cost estimates to changes in assumptions. 55

Chapter 1

Introduction

The climate is changing. For decades it was a matter of debate, a point of uncertainty, whether humans could fundamentally alter the earth's climate system. Now it is an observable fact: we can, we have, and we will continue to do so. Last year, 2005, was the hottest on record, continuing an accelerating warming trend that has brought the global mean temperature up 0.8°C over the past century (Hansen et al., 2006). But the changes we see are not only in temperature. Severe damage to ecosystems and harm to humans can now be attributed to climate change, including extensive bleaching of coral reefs (Sugden, 2005), massive loss of forests due to pine beetle infestation (Carroll et al., 2003; Gan, 2004), increased forest fire activity (Westerling et al., 2006), and increased intensity of hurricanes (Webster et al., 2005; Emanuel, 2005).

The principle cause of this climate change is the emission of carbon dioxide (CO_2) to the atmosphere when humans burn fossil fuel. The combustion of coal, oil, and natural gas has fueled two centuries of economic development across the world. Today, 80% of global energy use relies on fossil fuel (IPCC, 2005). Fossil fuel use is integral to the functioning of society at every level: the manufacturing of every manner of consumer good, the movement of people and products, the maintenance of shelter from the elements, the production of food, and the provision of water.

And yet, with every ton of carbon burned, the earth is altered in ways both predictable and unpredictable. The changes in climate to come will likely threaten many people's access to those basic services: shelter, food, and water. Additionally, the irreversible damage to the world's ecosystems, including the extinction of a large fraction of the earth's species, will be the most extensive civilized humans have ever experienced (Thomas et al., 2004).

This leaves us in a difficult position; industrial society is built around fossil fuels, and yet their use poses a grave and mounting threat. Fossil fuel use and the associated carbon emissions continue to grow exponentially (see Figure 1.1). Many nations have begun programs to limit CO_2 emissions, though current measures are not sufficient to achieve dramatic reductions. Notably, the two largest emitters, China and the United States, have no national regulation on CO_2. Nonetheless, there is a general expectation that further international agreements will be reached to limit CO_2 emissions, and that a global market for carbon emissions reduction will be developed. Such a market already exists in the European Union, and has the effect of putting a value on CO_2 emissions reduction.

There are various options available to meet regulatory demands. Energy conservation and efficiency improvements are the lowest-cost near-term measures. Improvements in average vehicle fuel efficiency

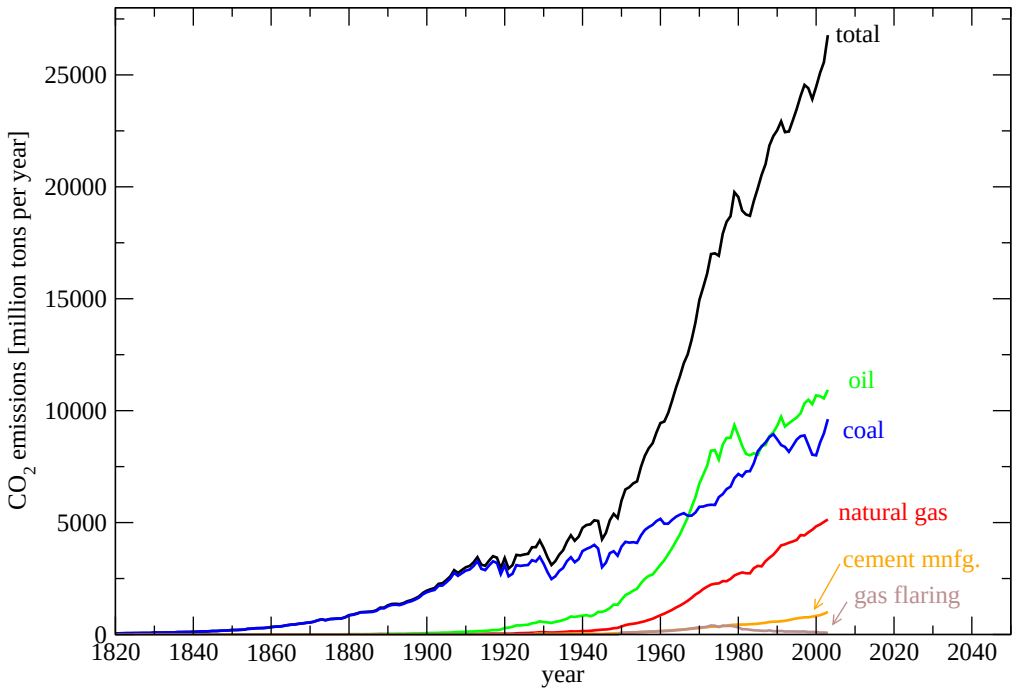

Figure 1.1: Global CO_2 emissions over time. Data from Marland et al. (2006). 2003 was the most recent year for which data were available.

are available at net negative cost (NRC, 2002). Improvements in industrial energy and material efficiency and reductions in building energy use are available at low or net negative cost (IPCC, 2001). We can also switch from fossil fuels to alternative energy sources: nuclear, wind, solar, biomass, and others. But they are comparatively expensive and so far limited in capacity. Especially challenging is the transportation sector, where solutions like battery-powered and hydrogen fuel cell cars are available but lacking in performance or still requiring substantial development. Biomass-based liquid fuels can work with existing vehicles but they are currently expensive to produce and are limited by agricultural capacity for growing feedstock. However, another option exists: we can retain fossil fuel as an energy source and capture the CO_2 emissions, preventing them from entering the atmosphere.

1.1 Carbon Capture and Storage

Carbon Capture and Storage (CCS) denotes the process of collecting CO_2, generally from fossil fuel combustion, and "storing" (sequestering), it outside of the atmosphere. It allows separation of the energy function of fossil fuel from the climate impact of combustion. It consists of three steps: (1) capture of the CO_2 from some source, (2) transport of the CO_2 in compressed form to a suitable storage site, and (3) injection of the CO_2 into the storage site and subsequent monitoring and management of the site.

In a CCS scheme, the CO_2 is captured from large point sources, mainly power plants. Several technologies exist to do this. CO_2 can be absorbed from the flue gas of a conventional plant, typically with a regenerable liquid solvent like monoethanolamine (IPCC, 2005). It is recovered from the solvent during

regeneration at a high enough concentration ($> 90\%$) for compression. The process is termed "post-combustion capture". One can also remove CO_2 from the fuel before combustion ("pre-combustion capture") by converting it to hydrogen and CO_2 by means of a water-gas shift reaction. CO_2 is collected after the shift then hydrogen is burned to generate electricity. The third option is to feed the plant with pure oxygen instead of air, an "oxyfuel" system. This way, the flue gas has a high enough CO_2 concentration to be compressed directly. Energy is saved by avoiding the solvent absorption and regenerating but expended for separation of oxygen from air.

Once compressed to a liquid, the CO_2 is ready for transportation to a sequestration site, which generally occurs via a pipeline like those currently used for oil and gas. Sequestration sites may include spent oil fields and unmineable coal seams, where the cost of storage can be offset by enhanced recovery of oil or extraction of methane, respectively. CO_2 can also be pumped into the ocean, though this may have harmful side-effects. The most secure type of sites are deep saline aquifers, which also have the largest estimated storage capacity: at least 40 years of CO_2 emitted at the current rate, and probably much more (IPCC, 2005).

There are three currently operating industrial scale CCS projects, with many more planned. IPCC expects that CCS will be widely deployed when the cost of CO_2 emissions (value of emissions reduction) reaches 25–30 \$/t-CO_2, and that CCS will account for a substantial share of carbon mitigation in future scenarios with carbon restrictions.

In the form described, CCS can only work with CO_2 from points sources, which currently account for about 40% of total emissions (and many of those point sources would be small, inaccessible by CO_2 pipelines, or otherwise not amenable to carbon capture). CCS can facilitate reductions in other sectors if vehicles, home heating, and other distributed sources are made to run on hydrogen which in turn is produced in facilities equipped with CCS. Uncertainty remains about the feasibility and cost of switching to this "hydrogen economy", and it may be high; Keith and Farrell (2003) calculate that the cost of switching to hydrogen fuel cell cars could be 300 \$/t-CO_2.

Even with a very successful CCS program and aggressive deployment of alternative energy sources, atmospheric concentrations of CO_2 will continue increasing throughout the century (IPCC, 2000a). Because CO_2 is long-lived in the atmosphere, cutting emissions does not reduce the concentration, but rather slows the rate of increase. If emissions are cut to zero, it would still take many centuries for natural removal mechanisms to bring CO_2 close to pre-industrial levels.

1.2 Air capture

Though nearly all current research on CCS focuses on capture from large point sources, it is also possible to capture CO_2 directly from the atmosphere, a process we call "air capture". Compared with point source capture, air capture has several advantages. CO_2 emissions from any sector can be captured, including emissions from diffuse sources such as automobiles, airplanes, agriculture, and home heating. The capture unit can be located at a favorable sequestration site, avoiding the need for extensive CO_2-transportation infrastructure. Consider a future climate scenario where society has been slow to adopt mitigation measures, but a sudden shift in the climate system dramatically raises concern and demand for action. Because

it is decoupled from the rest of the energy system, air capture can be deployed more quickly than other measures to reduce net carbon emissions.

Air capture also makes possible negative net emissions in the future. Consider a future climate scenario where climate sensitivity is on the high end of our expectations and climate change effects are quite severe, and in response we have achieved a highly carbon-constrained economy. Without air capture, CO_2 levels in the atmosphere would take centuries to approach pre-industrial levels, but with air capture society can choose the desired level of atmospheric CO_2 and, balanced against willingness to pay, how quickly to achieve it.

1.3 Thermodynamic and physical limits of air capture

The proposition of air capture is fundamentally one of concentrating CO_2 – taking from it from a dilute state (today roughly 380 ppm in the atmosphere) to a relatively pure gas ($>$ 90% for compression to pressures required for deep geologic sequestration). With a perfect mechanism for achieving this, we have only to overcome the free energy of mixing. It is given by $\Delta G = RT\ln(P/P_0)$, where in our case P is 1 atm and P_0 is 3.8×10^{-4} atm, for an energy requirement of 20 kJ/mol (0.5 GJ/t-C). This is quite modest compared with the energy liberated when burning fossil fuels, e.g. $\approx$ 600 kJ/mol-CO_2 (14 GJ/t-CO_2) for gasoline (DOE, 2005).

Relative to other technologies for generating carbon-neutral energy, the land-use requirements for air capture are potentially very small. Dubey et al. (2002) compare the energy available in a square meter of land collected by biomass (0.003 kW) or from sunlight (0.2 kW), and the kinetic energy passing through 1 m^2 for windmills (0.6 kW) with the fossil energy that can be generated in producing the quantity of CO_2 passing through that square meter (100 kW). The conclusion is that land requirements for air capture coupled with fossil energy generation are potentially orders of magnitude smaller than for these other options.

Given that it is possible to extract CO_2 rapidly from the atmosphere in a relatively small area, we may be concerned that the process would be limited by local atmospheric transport of CO_2. Johnston et al. (2003) have studied this problem with global atmospheric and chemical transport modeling. They conclude that the transport and circulation of CO_2 is such that the entire flow of anthropogenic CO_2 could be offset by a single global sink of no more than 75,000 km^2 in area, and with intelligent placement of sinks, a small fraction of that area would be needed. This is intuitively consistent with the observation that the atmosphere is relatively well-mixed with respect to CO_2. In general it seems that the issue of local CO_2 transport to a sink, and the related issue of depleted-CO_2 plumes that may be hazardous to plant life, can be easily resolved.

1.4 Routes to air capture

1.4.1 Organic carbon production

Many organisms naturally capture CO_2 through photosynthesis. One can effectively remove CO_2 from the atmosphere with land management and land-use changes which increase terrestrial biomass, such as growing a forest where there once was agricultural land. Estimates for the cost of these projects cover the range 0.03–8 \$/t-CO_2 (IPCC, 2000b). Though even the high end of this range is small relative to other carbon management options, this strategy is fundamentally limited to one-time reductions. Once a plot of forest has reached maturity, it is no longer compensating for CO_2 emissions; the CO_2 released as plant matter decomposes is in balance with the CO_2 absorbed as new plant matter forms.

Ocean flora has also been discussed as a means of carbon capture. It is suggested that adding key nutrients to some parts of the ocean will generate large blooms of plankton which will take up carbon and draw it down to the deep ocean. Large scale experiments on this method have been conducted, so far with limited success in sustaining the bloom (Buesseler et al., 2004; Buesseler and Boyd, 2003, for example).

Another means of capturing CO_2 with photosynthesis is to run a biomass-fueled power plant with a carbon capture system. When the biomass is grown, it extracts CO_2 from the air which is later captured from the power plant flue gas with an amine system or other point-source capture system. This scheme is renewable, since each new crop of plants further extracts CO_2. (Rhodes and Keith, 2005) have estimated the cost of CO_2 capture with one such system at 41 \$/t-CO_2.

1.4.2 Metal-carbonate production

For a chemical approach to air capture, some industrial waste streams are suitable for absorbing CO_2. In particular, steel slag and waste concrete, rich in calcium and magnesium oxides, readily react with CO_2 to form solid carbonates. Stolaroff et al. (2005) have analyzed a scheme for carbonating steel slag and waste concrete with CO_2 from ambient air, and estimate the cost as \$8/t-CO_2. Many other researchers have studied concrete carbonation not necessarily associated with air capture Fernandez et al. (2004). Iizuka et al. (2002), for instance, have estimated the cost at 2 \$/t-CO_2.

Another approach is to add these kind of waste materials or suitable virgin materials such as limestone ($CaCO_3$) or soda ash (Na_2CO_3) to the surface layer of the ocean. The aim is to increase the alkalinity of the ocean, thereby increasing its capacity to uptake CO_2, indirectly decreasing atmospheric CO_2 concentrations (Kheshgi, 1995).

1.4.3 Capture with a regenerated sorbent

None of the above routes has the potential to capture a large fraction of anthropogenic CO_2 emissions; each has fundamental limitations on its scope. Biomass, for instance, is limited by the land available for cultivating biomass crops and by the secondary impacts of agriculture. The annual U.S. production of concrete and steel slag could capture less than 1% of U.S. emissions (Stolaroff et al., 2005). Ocean sequestration is limited by area of suitable ocean surface and problems associated with altering the ocean's

chemistry.

To capture a large fraction of CO_2 emissions, for instance all of the emissions from the transportation sector (40% of total), it is natural to seek a sorbent which can capture CO_2 and then be regenerated and reused in a cycle. Most likely, this would be coupled with deep geological sequestration of CO_2. The regeneration can be accomplished with a temperature swing, such as when amine solvents are heated, a pressure swing, such as a solid sorbent which releases CO_2 when exposed to a vacuum, or with a chemical reaction.

An ideal sorbent would have a binding energy with CO_2 just larger than the 20 kJ/mol required to pull it from the atmosphere, would be inexpensive, abundant, and non-hazardous. Research on new and novel sorbents with these characteristics is ongoing. However, two well-known sorbents which, while not optimal on the 1st condition, satisfy the latter three well, are aqueous solutions of calcium hydroxide ($Ca(OH)_2$) and sodium hydroxide (NaOH).

1.4.4 Metal hydroxide sorbents

In order to support the feasibility of air capture, example systems with metal hydroxide sorbents have been proposed (Lackner et al., 2001; Dubey et al., 2002; Zeman and Lackner, 2004; Keith et al., 2006; Baciocchi et al., 2006). The proposals generally include a sodium hydroxide (NaOH) or calcium hydroxide ($Ca(OH)_2$) solution which absorbs CO_2 and is regenerated by "calcination" (see next chapter). In most of these proposals, the specific form of the system has remained vague.

A system using $Ca(OH)_2$ solution is outlined by Dubey et al. (2002). It consists of pools of water saturated with $Ca(OH)_2$ which absorb CO_2 as wind blows across the surface. Solid $CaCO_3$ is periodically collected from the bottom of the pool and calcined (see Chapter 2). Dubey et. al. estimate a cost of about 20 \$/t-CO_2 for this scheme.

A related proposal is to use NaOH solution. $Ca(OH)_2$ is still used in the regeneration of the NaOH, and the end product – $CaCO_3$ to be calcined – is the same. However NaOH solution has the advantages over $Ca(OH)_2$ that it can have lower vapor pressure to prevent water loss, it can contain higher concentrations of hydroxide for more efficient contacting systems, and is less prone to cause scaling (accumulation of solid carbonate minerals on surfaces). Zeman and Lackner (2004) have claimed the cost of operating a NaOH-based capture system can be as low as 7–20 \$/t-CO_2 , though it has been disputed whether the system could be operated at any cost and energy requirement low enough to make it feasible (Herzog, 2003). The component of the system which makes contact with air to extract CO_2 (which we term the "contactor") remains a particular point of contention, with suggested forms including large convection towers (Lackner et al., 2001), open, stagnant pools (Dubey et al., 2002), falling film channels (Zeman, 2006), and packed scrubbing towers (Herzog, 2003; Baciocchi et al., 2006). Herzog estimates that the energy requirement of a packed tower contactor similar to the ones used to capture CO_2 from power plant flue gas would be 6–12 times the energy produced when the fossil fuel was initially burned.

In addition to energy requirements, a natural concern for any large-scale aqueous contacting system is the quantity of water lost by evaporation. Water loss may be particularly large in an air capture system since the low concentration of CO_2 in the atmosphere requires a large amount of interaction between the gas and liquid phases. Water loss in proposed systems has not been previously calculated or addressed.

The NaOH approach is the primary subject of this thesis, and an example system is presented in detail. The system follows on the proposal in Keith et al. (2006) and similar systems are discussed by Zeman and Lackner (2004), Zeman (2006), and Baciocchi et al. (2006).

1.5 Research Objectives

This research is a feasibility assessment of air capture as a technology for mitigating climate change. We start from the preceding arguments that air capture is theoretically viable and turn to the cost and energy requirements of an actual system. To place an upper bound on this cost, we develop an example system comprised of current technology capable of large scale deployment. We then estimate the cost and energy requirements of this example system. Applying this result and other insights from the analysis, we estimate the possible range of future air capture costs.

Chapter 2 provides an overview of air capture systems that may be constructed with known technology. Where available, the energy requirements of these systems are given. A particular component of these systems, the contactor, is the largest source of uncertainty in feasibility and cost. This component is addressed in detail in Chapter 3, where an example contactor is developed. A variety of modeling strategies are used to explore the general features of a contactor based on sodium hydroxide (NaOH) spray. A prototype is constructed and analyzed in a series of experiments. Rates of CO_2 capture and energy use are measured. The results are used to estimate the cost and energy requirements of the contactor component. General issues related to spray-based contactor design are identified. The chapter concludes with details of the prototype and experimental procedure.

In Chapter 4 we estimate the cost an example air capture system. Uncertainties in this estimate and potential improvements to the system are explored. Chapter 5 begins with a summary of our key findings. We then discuss some implications of these findings and implications of air capture in general. Finally, we outline the process we went through in this analysis, relating it to the general task of assessing a future energy technology. Appendix A provides a glossary of mathematical symbols used in this thesis and Appendix B gives details of the experimental design and procedures discussed in Chapter 3.

Chapter 2

Example systems

The simplest air capture scheme is that proposed by Dubey et al. (2002). They suggest constructing large, stagnant or gently-agitated pools of water saturated in $Ca(OH)_2$. As air blows naturally across the surface, CO_2 is absorbed into the solution and forms $CaCO_3$, which settles to the bottom of the pool. Periodically, the settled $CaCO_3$ is mechanically collected and regenerated in a process called calcination. In this process, the $CaCO_3$ mud is dewatered then heated in a kiln (the "calciner") to over 1000°C, at which point the CO_2 is driven off in a concentrated stream. With calcination the options for CO_2 capture are similar to the capture options for a power plant. Either a post-combustion absorbing system, like monoethanolamine scrubbing, or an oxyfuel system can be used to collect CO_2 at a concentration suitable for compression (pre-combustion systems do not work for calcination, since the process generates CO_2 in addition to that from the fuel). Following capture from calcination, the CO_2 is sequestered just as in conventional CCS.

We expect two difficulties with this scheme. First, the rate of CO_2 absorption is likely to be low. The low solubility of $Ca(OH)_2$ limits the hydroxide concentration ($[OH^-]$) in solution to about 0.015 M. In turn, this limits the rate of CO_2 absorption (see next chapter, especially Equation 3.3). The second and related problem is that this scheme would be subjective to very high evaporative water loss, relative to the quantity of CO_2 absorbed (see Equation 3.7 and related discussion).

Especially as a response to the first of these difficulties, other example schemes have turned to NaOH solution (also known as "caustic soda" or "caustic") as a sorbent. With unlimited solubility in water, reasonable working solutions of NaOH have $1 > [OH^-] > 10$ M. Also, the vapor pressure of water is lower in NaOH solutions, so that water loss from high molarity solutions can be small or zero.

2.1 Contacting with NaOH

Extraction of CO_2 from air with NaOH solution has been a well-known process for many decades. Even at ambient concentrations, CO_2 is absorbed readily by solutions with high pH. The most common industrial method of absorbing a gas into solution is to drip the solution through a tower filled with packing material while blowing the gas up through the tower (a "packed tower" design). Greenwood and Pearce (1953) and Hoftyzer and van Krevelen (1954) measured CO_2absorption by packed towers designed to produce air which is nearly free of CO_2 ($> 99\%$ capture efficiency). However, for the purpose of bulk CO_2 capture

there is no compelling reason to capture most of the CO_2 from any given parcel. Shorter towers and therefor lower pumping energies can be used with a lower capture efficiency of CO_2 passing through the system. We can extrapolate their data to a low capture efficiency by assuming CO_2 absorption with height follows first order decay, a standard assumption for packed-tower absorbers (Fair et al., 2001). Indeed, if we choose a capture efficiency of 50%, which is reasonable for our application, the combined gas and liquid pumping energy requirements of running such a unit are rather small – 0.3 GJ/t-CO_2, which corresponds to a tower 1.5 m tall.

Baciocchi et al. (2006) also investigate a packed tower system for CO_2 capture, this time specifically for the purpose of air capture. They model a system with several commercially available structured packings. In the optimal configuration they find the energy cost of a packed tower contactor of 0.7 GJ/t-CO_2 and a bed height of 2.8 m.

Zeman (2006) calculates the energy requirements for a packed tower contactor, again specifically for air capture, and gets 2 GJ/t-CO_2. He goes on to discuss a new technology in absorption tower packings – hollow fiber membranes – which have been shown in one empirical study to reduce energy costs by 33% over conventional packings. Applying this reduction, he gets 1.3 GJ/t-CO_2 for a contactor with hollow fiber membrane fill.

In all these cases the low air flow rate required to minimize pressure drop and the dilute concentration of CO_2 in air requires the “tower” to be short and very wide – perhaps hundreds of meters in diameter or set in an array of many hundreds of smaller towers – to capture CO_2 at a rate comparable to emissions from a medium-sized power plant. A contactor of these dimensions would be very different from conventional packed towers. The properties of this type of design are likely dictated by "edge effects" – the nature of the system at the top and bottom of the bed – and by the engineering of the distribution mechanism for air and water. While not intractable, these issues make it hard to estimate the cost of such a system.

A useful analogy may be the trickle-bed filter used in wastewater treatment plants. It consists of a wide, cylindrical basin, drafted from underneath, with a rotating arm which distributes wastewater over the top. The basin is filled with rocks or a plastic media on which a biofilm grows. The biofilm adsorbs and consumes organic matter from the passing solution and requires a steady stream of air in order to oxidize the material. Newer installations use plastic media in depths of 6–12 m and often have forced-air systems (Davis, 2000). Biofilms would not form in the caustic environment of a NaOH-based contactor so a different media geometry may be more suitable. But in all other respects, the cost and design features of trickle-bed filters seems a promising avenue of future research.

An alternate strategy to those above is to use a lighter packing and taller tower. In the limit, this becomes an empty tower with the solution sprayed through, much like a power plant evaporative cooling tower or an SO_2-scrubbing tower for combustion flue gas. For the purposes of this research, this strategy has the advantage that the costs may be easier to estimate because of the simplicity of the design and the analogy to industrial cooling towers. This design – an NaOH-spray-based contactor with a cooling-tower-like structure is elaborated in great detail in Chapter 3.

	Reaction			Enthalpy of reaction, ΔH°	
				kJ/mol-C	GJ/tCO$_2$
(1)	$CO_2(g) + 2OH^-$	$\rightarrow$	$CO_3^{2-} + H_2O$	-110	-2.5
(2)	$CO_3^{2-} + Ca^{2+}$	$\rightarrow$	$CaCO_3(s)$	12	0.27
(3)	$CaCO_3(s)$	$\rightarrow$	$CaO(s) + CO_2(g)$	179	4.0
(4)	$CaO(s) + H_2O(l)$	$\rightarrow$	$Ca^{2+} + 2\ OH^-$	-82	-1.9

Table 2.1: Chemistry of example air capture system. Values from Weast (2003).

(1)	$Na_2CO_3(s,l) + 2NaBO_2(s,l)$	$\leftrightharpoons$	$Na_4B_2O_5(s,l) + CO_2(g)$
(2)	$Na_2CO_3(s,l) + NaBO_2(s,l)$	$\leftrightharpoons$	$Na_3BO_3(s,l) + CO_2(g)$
(3)	$Na_4B_2O_5(s,l) + H_2O(l)$	$\rightarrow$	$2NaOH(aq) + 2NaBO_2(aq)$
(4)	$Na_3BO_3(s,l) + H_2O(l)$	$\rightarrow$	$2NaOH(aq) + NaBO_2(aq)$

Table 2.2: Chemistry of autocaustization process. From Lindberg et al. (2005).

2.2 Caustic recovery

After the NaOH solution has been through the contactor, it contains CO_2 in form of carbonate (CO_3^{2-}), so that the solution is mostly sodium carbonate (Na_2CO_3). The carbonate must be removed to regenerate the caustic (running the NaOH in a once-through mode is not practical). This is a problem shared with the pulp and paper, aluminum manufacturing, and other industries. The approach used overwhelmingly in these industries is called the kraft recovery process. The chemical reactions that comprise the process as applied to an air capture system are shown in Table 2.1. The Na_2CO_3 is sent to a batch reactor where $Ca(OH)_2$ is added. The result is that solid calcium carbonate ($CaCO_3$) precipitates. The recovered NaOH goes back to the contactor. Meanwhile, the $CaCO_3$ is dewatered and calcined in exactly the process described earlier.

Kraft recovery is a well-known process operating at a large scale. It lends itself to use in example systems because it is known to work and because analogies with existing industrial systems can help estimate the cost and energy requirements of the system applied to air capture. However, it is not an ideal process for air capture. Reaction 1 (Table 2.1) occurs at low temperature and so the heat released cannot be usefully recovered. Energy released during Reaction 4 may be available as medium-grade heat ($< 450^\circ$C) and so can only partly be recovered. In net we are left with a large energy demand per unit CO_2 captured, driven by Reaction 3. The heat of calcination, as the energy demand of Reaction 3 is referred to, of 179 kJ/mol is far larger than the 20 kJ/mol we theoretically need for an air capture sorbent. In addition, kraft recovery is a capital-intensive process, with phase separation, solids handling, washing, dewatering – in all, a typical kraft recovery plant has dozens of process units. There are, however, at least two alternatives to the kraft process for caustic recovery.

A process known as autocaustization is in early stages of commercial deployment in the pulp and paper industry. Sodium borates can be reacted with Na_2CO_3 at high temperature to release CO_2 as a gas and regenerate NaOH. Table 2.2 shows the chemistry of the autocaustization process. The Na_2CO_3 and $NaBO_2$ are both dissolved in the working solution. Then the solution is heated in a recovery boiler, driving off the water and melting the remaining solids. Reactions 1 and 2 occur in the molten salt phase, above 900°C. When the product is cooled and re-hydrated, the NaOH is recovered by Reactions 3 and 4

$7Na_2CO_3(l) + 5Na_2Ti_3O_7(s)$	$\rightleftharpoons$	$4Na_8Ti_5O_{14}(s) + 7CO_2(g)$
		$\Delta H = 63$ kJ/mol-CO_2
$4Na_8Ti_5O_{14}(s) + 7H_2O$	$\rightarrow$	$14NaOH(aq) + 5Na_2Ti_3O_7(s)$
		$\Delta H = 7.6$ kJ/mol-NaOH

Table 2.3: Chemistry of caustic recovery with titanates. From Richards et al. (2002).

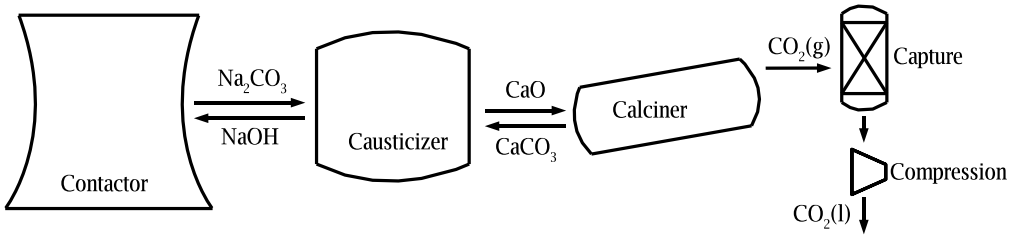

Figure 2.1: Example air capture system. CO_2 from the atmosphere is absorbed into NaOH solution in the contactor. The resulting Na_2CO_3 solution is regenerated to NaOH in a series of reaction vessels by addition of CaO in the causticizer. The resulting $CaCO_3$ solids are in turn regenerated to CaO by heating in the calciner. CO_2 in the calciner flue gas is captured and compressed by conventional means.

(Lindberg and Backman, 2004).

On the one hand, this process relies on a reaction (2) that has about the same heat of reaction as calcination (Richards et al., 2002) and must be run at about the same temperature, so this process does not appear to offer an advantage. On the other hand, it may be possible to keep all compounds in a liquid or aqueous phase throughout the reaction cycle. This would simplify caustic recovery by avoiding solids handling and dewatering steps. Also, the liquid phase allows for very efficient heat exchange; perhaps the heat released from reactions 3 and 4 can be recovered at a usefully high temperature to compensate for the demands of 1 and 2. Additionally, if the boiler is fired indirectly or with oxyfuel, it may be possible to extract CO_2 from the boiler in a pure stream. Perhaps we can run the reaction at positive pressure, saving energy in the final compression step. Industry experience demonstrates that autocaustization works, though at this point it is speculation whether the process would be favorable for air capture.

Another possible caustic recovery mechanism is by addition of sodium titanates. The chemistry of caustic recovery with titanates is shown in Table 2.3. The promising point in this case is that the heat of reactions is about a third of that in calcination, suggesting this process offers significant energy savings compared with calcium-based systems. Titanate-based caustic recovery is in early stages of development.

2.3 Example air capture system

2.3.1 Overview

We have now briefly discussed all the components necessary to construct an air capture system. To illustrate a complete system that could be built with available technology we will describe an example with a spray-based contactor and kraft recovery plant. The appeal of this system is that the chemicals involved

are all inexpensive, abundant, and relatively benign, and that almost all the processes are well-understood as current industrial-scale practices. We choose this system for ease of analysis and because of its similarity to current commercial technologies although we doubt it is the best or most likely means of achieving air capture. We expect it is a system which will convince skeptics that air capture is possible with current technology and perhaps point the way for research into improved systems.

A top-level process diagram is shown in Figure 2.1. In the contactor, the NaOH is brought into contact with atmospheric air and absorbs CO_2, forming sodium carbonate (Na_2CO_3). This carbonate-containing solution is then sent to the Causticizer. In the Causticizer, lime (CaO) is added to the solution, producing solid calcium carbonate ($CaCO_3$) and NaOH. The $CaCO_3$ is collected and sent to the Calciner while the NaOH is sent back to the Contactor. The Calciner heats the $CaCO_3$ until the CO_2 is driven off and CaO is re-formed. The CO_2 is collected and compressed for sequestration. The contactor is the least-understood component of the system in terms of costs and energy requirements available from analogous industrial units. It is discussed in detail in Chapter 3. The causticizer and calciner are discussed in detail below.

2.3.2 Caustization

In this step, the Na_2CO_3 solution from the capture unit is mixed with CaO from the Calciner. Two reactions occur. The hydrating of CaO (Reaction 4 in Table 2.1) is known as slaking. Then aqueous calcium ions combine with carbonate in the causticizing reaction (Reaction 2). Most of the slaking typically occurs in a separate batch reactor prior to Na_2CO_3 addition, but it may also occur in the same reactor as the causticizing. A near-perfect analogy can be drawn between this and the causticizing step in the kraft recovery process used in the pulp and paper industry. The kraft process takes spent pulping chemicals, primarily Na_2CO_3 and Na_2S, and regenerates them to NaOH and Na_2S with the same chemical reactions as above. The substantive differences between the kraft process and the proposed caustization process for air capture are as follows.

Sulfur content

The presence of sulfide aids the preparation of wood pulp, and so must be carried through the kraft recovery process. The process has been tested, however, without the addition of Na_2S, and the primary result is an improvement in the conversion efficiency of Na_2CO_3 to NaOH by a few percent. In general the sulfur only complicates the process. Since our proposed system doesn't require any sulfide, we expect it to run a bit more efficiently than the kraft equivalent.

Temperature

In the kraft process, the slaking and causticizing steps are typically performed with a solution temperature in the range of 70-100^{o}C. However, the solution entering this step in the proposed system will be at ambient temperature or cooler. The solution is heated by the slaking reaction; assuming a (typical) concentration of about 2 mol/l CaO added, the slaking reaction will increase the solution temperature by about 20C^o, but that would only bring the solution to, perhaps, 40^{o}C. While the equilibrium conversion efficiency of Na_2CO_3 to NaOH is higher at lower temperatures, the kinetics become prohibitively slow. Without

changing the process design to accommodate significantly longer residence times, we will have to add heat to the solution. Another 30C^o would bring us into the industrial range. We can do so with a liquid-to-liquid heat exchanger and a low grade heat input of (assuming the exchanger is 80% efficient) 14 kJ/mol-CO_2, or about 0.31 GJ/ton-CO_2.

Solids content

In the kraft process, the initial Na_2CO_3 solution contains organic particles and insoluble minerals ("dregs") in the part-per-thousand range. The dregs impair the performance of the process and so most must be removed in a clarifier. For the proposed system, the entire dreg-removal subsystem can probably be eliminated. The source of contamination most analogous to the dregs in the proposed system is fine particles captured from the air along with the CO_2. Assuming a particle concentration of 100 μg/m^3 and equal absorption efficiency with CO_2, the particle concentration in solution will be in the range of 10 parts per million.

2.3.3 Calcination

Calcination is the process where $CaCO_3$ is heated to make CaO. It is practiced at very large scale in the production of lime, cement, and in pulp and paper mills. Modern calciners typically take the form of a rotary kiln – a large brick-lined cylinder set at an angle which rotates as the lime works its way down the shaft by gravity. Fluidized bed versions are also used. They can be fired with natural gas or fuel oil. A calciner for lime production can operate with an energy input of 4.6 GJ/t-CO_2 – close to the thermodynamic limit of 4.1 GJ/t-CO_2. Calciners in the pulp and paper industry require more energy because they start with $CaCO_3$ mud instead of dry $CaCO_3$, and must drive off the water. A typical energy requirement for a paper-industry calciner is 8.7 GJ/t-CO_2 (Adams, 1989). As discussed in Section 2.3.5, much of this additional energy may be recoverable as high-temperature steam from the flue gas, but it is not generally recovered in pulp and paper mills.

The calciner in our example system is a close analogy with that paper industry since it is also starting with $CaCO_3$ mud produced during caustization. It is certainly possible to remove more of the water in the $CaCO_3$ mud before calcination. The pulp and paper industry has presumably optimized the trade-off between capital expenditure and mechanical energy for dewatering and energy cost for calcining for their circumstances. However, the higher energy costs we have assumed would probably elicit a more energy-efficient design.

In contrast to current industrial systems, the proposed system must, of course, capture CO_2 from the calciner. The most straightforward method for doing this would be to use a conventional high-efficiency fluid bed or rotary kiln calciner followed by an amine-based CO_2 capture system. If the calciner was fired by natural gas, the CO_2 concentration in the exhaust gases would approach 20% (dry basis) lowering the capital and energy costs of amine capture compared to existing estimates for capture from coal-fired power plants (14% CO_2). Furthermore the water in the lime mud becomes high-temperature steam in the calciner. This steam can be used as a heat source for regeneration of the amines. Analysis suggests that there is sufficient steam to supply much of the needed regeneration heat in the amine capture unit

(Rao and Rubin, 2002; Rao, 2004).

The capture of CO_2 from calcination might alternatively be achieved using oxygen blown combustion in a fluidized bed. Such a system would be a hybrid of existing fluid bed calciners and oxygen-fired coal combustion with CO_2 capture which has been studied (but not implemented) as a method for achieving CO_2 capture at electric power plants. Such a system might offer significant energy savings over the amine system but because it introduces significantly more new engineering and so we do not include it in our base example system. Other authors estimating the energy requirements of an example air capture system have assumed an oxyfuel system (Baciocchi et al., 2006; Zeman, 2006).

2.3.4 Integrated system

We argue that one may construct an air capture system essentially by stringing together the components described: a contactor, a kraft recovery system, an amine-based capture system, and a compressor. A few novel interconnections are required, like a heat exchanger between the contactor and the causticizer and heat exchange between the calciner flue gas and amine recovery boiler. We expect the capital and maintenance for this and some extra piping will be more than balanced by the capital equipment *not* required by an air capture system, compared with the caustic recovery system for a paper mill. Thus for the caustic recovery portion, we use a rule-of-thumb capital cost and energy requirements estimate for a "turn-key" caustic recovery system, obtained from an industry engineer (Flanagan, 2004). The capacity of this plant is 1000 ton-CaO per day, or about 300,000 ton-CO_2/yr. This is comparable to the throughput of a single large contacting tower as detailed in the next chapter. However, a full-scale air capture installation would likely include dozens or hundreds of such towers, allowing for a combined caustic recovery plant one or two orders of magnitude larger in scale. There may be significant savings of capital and maintenance cost at this larger scale, as well as opportunities for more energy-efficient operation. These benefits of economy-of-scale are not considered in our example system estimates.

The parameters for the amine capture plant are derived from the model used in Rao and Rubin (2002) with adjustments made for the higher CO_2 concentration in calciner flue gas. Compression is included in the amine plant. The amine system is assumed 90% efficient, so that all upstream values must be adjusted up to compensate for CO_2 leakage.

2.3.5 Energy requirements

Table 2.4 shows energy requirements of various components of the example system, along with comparable values from the analysis of Baciocchi et al. (2006) and Zeman (2006). They both assume more optimized systems than the pulp and paper mill caustic recovery system used in our example system. Baciocchi et al. have two cases; in case A they assume a conventional dewatering system and in case B they assume a newer pellet reactor system which removes water much more efficiently. This turns out to have a substantial effect, as the reduced demand for calciner heat multiplies through other components, reducing demand for oxygen and compression.

Another important factor determining total energy requirements is the handling of low-grade heat from the slaking reaction and from steam in the flue gas. Neither of these are typically utilized in a kraft mill.

Component / source	GJ/t-CO_2 electric	GJ/t-CO_2 thermal
Contacting		
Packed tower (Greenwood)	0.3	-
Packed tower (Baciocchi)	0.69	-
Packed tower (Zeman)	1.3	-
Spray tower (see Ch. 3)	1.4 (range: 0.71–3.2)	-
Caustic Recovery		
Example system	(included in thermal)	13
Baciocchi A	0.11	12
Baciocchi B	0.11	8.0
Zeman	-	5.13
CO2 capture		
Amine system in example system	0.121	(3.8)
Oxygen separation (Baciocchi A)	0.62	-
Oxygen separation (Baciocchi B)	0.49	-
Oxygen separation (Zeman)	0.29	-
CO2 compression		
Example system	0.43	-
Baciocchi A	0.42	-
Baciocchi B	0.36	-
Zeman	0.34	-
Totals		
Example system	(included in thermal)	16
Baciocchi A	(included in thermal)	17
Baciocchi B	(included in thermal)	12
Zeman	(included in thermal)	11

Table 2.4: Energy requirements of the air capture system by component. For totals, thermal energy is converted to electricity at 35% efficiency. Values are given as integrated with the indicated system, that is, per net ton CO_2 captured by the given system.

When the author toured the caustic recovery plant at a paper mill (the management of the mill asked that it not be identified), a process engineer there explained that heat recovery from the slaking reaction would be difficult because of scaling issues – any heat-exchange equipment with small openings and lots of surface area would be prone to clogging by mineral build-up. On the other hand, he says, heat recovery from the calciner flue gas would be possible but has not been implemented at the mill, which has limitations of space, capital, and no sense of scarcity of energy (the mill generates much of its own heat from burning lignin in the pulping process). Thus we do not assume recovery of the slaking energy but do use the calciner steam to run the amine recovery boiler in our example system. In contrast, Zeman assumes that heat from the slaking reaction (Reaction 4 in Table 2.1) can be recovered for pre-drying the $CaCO_3$ mud, and so does not include energy for dewatering. We would not argue that this is unrealistic, only that the practicality is uncertain. Baciocchi et al. take the most conservative stance, assuming no recovery of low grade heat at all.

Chapter 3

Contactor

In this chapter we propose a form of the contactor modeled after a power plant cooling tower and functioning similarly to a power plant sulfur-scrubbing tower, and then estimate the cost and energy requirements of that system. We build a prototype of the contactor which illustrates the feasibility of the process and assists estimation of the cost and energy requirements of the full-scale analogue.

In the Section 3.1 we first discuss the theoretical modeling and calculations that motivated our contactor and experimental design, then describe the prototype contactor we developed and methods of its testing. In Section 3.2 we present our findings from a combination of numerical techniques and experimental data on (1) the characteristics of CO_2 absorption in the contactor, (2) spray droplet collision and coalescence as it relates to scale-up of prototype results, (3) the energy requirements in the prototype contactor and scaling to a full-sized system, (4) the quantity of water lost to evaporation in the prototype and a full-sized system, and (5) an engineering-economic analysis of the cost of a full scale contactor. In Section 3.5, we identify the factors that may significantly affect the cost and feasibility of a NaOH spray-based contactor positively or negatively. Appendix B, "Experimental Details and Procedure", supplements the information in this chapter, giving additional description and documentation of the prototype experiments.

3.1 Materials and Methods

3.1.1 Theoretical methods

To determine the feasibility and inform the design of a NaOH spray system, we applied several theoretical models to estimate the mass transfer to a drop of NaOH solution falling through air at terminal velocity. In principle, mass transfer may be limited by gas-phase transport, liquid phase transport, liquid phase reaction, or a combination. We will find (in Section 3.2) that liquid phase transport and reaction should be limiting on theoretical grounds, and experimental data confirm this. This section describes the theoretical models that we applied.

We first assumed gas-phase limitation to mass transfer, so that the flux of CO_2 to a drop surface, J_{CO_2}, is given by (Seinfeld and Pandis, 1998):

$$J_{CO_2} = k_g(C_\infty - C_s) \tag{3.1}$$

where k_g is the gas-side mass transfer coefficient, C_∞ is the CO_2 concentration in the bulk air, and C_s is the equilibrium CO_2 concentration in air at the surface of the drop, which is effectively zero for high-pH solutions such as ours. We estimate k_g with an empirical correlation in terms of the Sherwood number, *Sh*, (Bird et al., 1960):

$$Sh = 2 + 0.6Re^{1/2}Sc^{1/3} \text{ and } Sh = \frac{k_g D}{D_g} \tag{3.2}$$

where *Re* is the Reynolds number and *Sc* is the Schmidt number for a drop falling at terminal velocity with diameter D, and D_g is the diffusion coefficient for CO_2 in air.

Next we assumed liquid phase limitation to mass transfer. For the dominant sizes of drops used in our prototype, $D \sim$100–400 μm, internal circulation does not occur (Seinfeld and Pandis, 1998), so that a simple reaction-diffusion model can be used to estimate the rate of CO_2 absorption.[1] Our system can be approximated by a pseudo-first-order, irreversible reaction at steady state, so the rate of absorption per unit surface area is given by (Danckwertz, 1970):

$$J_{CO_2} = C_0\sqrt{D_l k[\mathrm{OH}^-]} \tag{3.3}$$

where C_0 is the concentration of CO_2 in solution at the drop surface, assumed to be in equilibrium with the bulk air and D_l is the diffusion rate of CO_2 in water. Following convention, $[\mathrm{OH}^-]$ denotes the aqueous concentration of hydroxide ions in mol/L. The reaction constant, k, refers to the hydrolyzing of dissolved CO_2 (Danckwertz, 1970):

$$CO_2(aq) + \mathrm{OH}^- \longrightarrow \mathrm{HCO}_3^- \,,\, k = 8500M^{-1}s^{-1} \tag{3.4}$$

This is a second-order reaction, however the concentration of OH^- is so large in our system ($>$ 1M) that it is effectively fixed, and this can be treated as a first-order reaction in CO_2.

Equations 3.1 and 3.3 give CO_2 absorption per unit surface area. Multiplying by the surface area and dividing by the volume of a sphere gives absorption per unit of solution, Q. For the liquid-limited case, we have:

$$Q = \frac{6C_0}{D}\sqrt{D_l k[\mathrm{OH}^-]}\,. \tag{3.5}$$

Multiplying Q by the residence time of a drop in the contactor yields an estimate of CO_2 absorption per pass. We estimate the residence time, τ, by:

$$\tau = \frac{H}{v_{air} + v_t} \approx \frac{H}{v_{air} + k_v \cdot D} \tag{3.6}$$

where v_{air} is the average velocity of air in the contactor, H is the height of the contactor, and v_t is the terminal velocity for the volume-average drop size in m/s, which can be approximated as shown with the empirical constant $k_v = 4 \times 10^3\, s^{-1}$ for drops in this size range (adapted from Pruppacher and Klett, 1997). This simplified model ignores the transient velocity of spray as it escapes from the nozzle, effects of drop shadowing, and non-uniformity in the flow patterns, but it is useful as an order-of-magnitude estimate of

[1] A full time-dependent surface-renewal model was also applied but the results did not differ significantly from the simple reaction-diffusion model under the conditions we consider.

the CO_2 capture rate in a spray system.

The equations presented so far address mass transport of CO_2. In principle, numerical methods may be used to estimate evaporative water loss in a contactor as well. As discussed in Section 3.2, it turns out that Equation 3.1 is not the correct model for CO_2, which appears to instead follow Equation 3.3 (CO_2mass transfer is liquid-side limited). However, Equation 3.1 can be applied to the evaporation of water (Seinfeld and Pandis, 1998). If we define R_{H_2O/CO_2} as the molar ratio of water evaporated to CO_2 captured, then we can estimate it by the ratio of mass transfer equations:

$$R_{H_2O/CO_2} = \frac{k_g(C_{\infty,H_2O} - C_{s,H_2O})}{C_0\sqrt{D_l k[OH^-]}} \tag{3.7}$$

The calculation is complicated because the vapor pressure of water at the drop surface, C_{s,H_2O}, is a function of drop temperature, $[OH^-]$, and $[CO_3^{2-}]$ (the concentration of CO_2 already absorbed in mol/L). For a sense of scale, we plug in initial conditions for temperature, relative humidity, and apply other reasonable parameters. Equation 3.7 then evaluates to the order of 10^3, which is quite large (though it is clearly an overestimate because it does not account for changing temperature and other dynamic effects). It is large enough to suggest that the air leaving a full-scale contactor will be saturated with water vapor. Thus we can estimate evaporative water loss by assuming the air leaving the contactor has water partial pressure equal to the vapor pressure of water at the surface of the drops.

For simplicity we assume that air and liquid reach the bottom of the contactor at the same temperature, so that an energy balance yields the temperature at the outlet:

$$T_{out} = T_{in} - \frac{\Delta C_{H_2O}\Delta H_{vap}}{\rho_{air}c_{p,air} + \rho_l c_l} \tag{3.8}$$

where m_{vap} is the mass of water evaporated per units contactor volume, $c_{p,air}$ and c_l are the heat capacities of air and the liquid solution (assumed equivalent to water), and ρ_{air} and ρ_l are the bulk densities of air and suspended solution. The quantity of water evaporated, ΔC_{H_2O}, is the difference between inlet water concentration and outlet vapor pressure of the solution:

$$\Delta C_{H_2O} = C_{s,H_2O}(T = T_{out}) - C_{H_2O,in} \tag{3.9}$$

Equations 3.8 and 3.9 can be solved simultaneously by iteration to yield ΔC_{H_2O}. Assuming an overall capture efficiency for the contactor of CO_2 from air allows a calculation of R_{H_2O/CO_2}. The results of this calculation for reasonable conditions in the contactor are presented in the next section.

3.1.2 Experimental methods

We constructed a prototype contactor in order to demonstrate the feasibility of CO_2 capture by NaOH spray and to allow us to measure the energy requirements in a way such that the results could be scaled up to a full size contactor. Details of the design choices and experimental procedure are given in Appendix B along with additional photographs of the structure. We required (1) a tower diameter large enough to

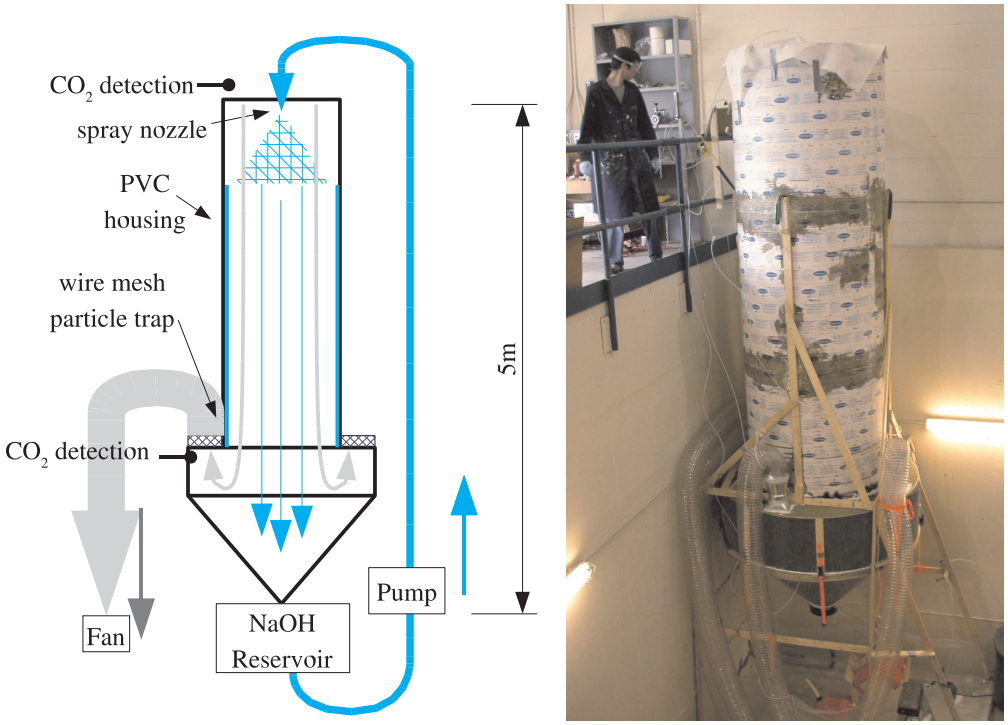

Figure 3.1: Diagram of prototype contactor and photograph of finished structure. Atmospheric CO_2 is absorbed by NaOH spray. CO_2 concentration at the air intake and outlet is measured and the rate of CO_2absorption is calculated.

accommodate a spray nozzle such that CO_2 absorption by the spray can be isolated from CO_2 absorption by the wetted wall, (2) a tower tall enough to assure that drops travel at steady-state velocity for most of the fall, and (3) a design in which solution loss by entrainment is minimal and small caustic droplets are filtered from the outlet air for safety.

The prototype contactor is diagrammed and pictured in Figure 3.1. NaOH solution is sprayed through the tower, collected, and recirculated from a 15-L reservoir while air is blown down through the tower co-currently at approximately 0.4 m/s. Though it may sacrifice some efficiency over an upward-flow design, the downward-flow, co-current design allowed for simpler construction and maintenance of the particle trap system. CO_2 concentration in inlet and outlet air is measured using a LiCOR infrared gas analyzer. Carbonate (CO_3^{2-}) concentration is measured in periodic liquid samples to corroborate the LiCOR measurements. Additionally, temperature, relative humidity, and pressure drops are recorded.

Two single-fluid spray nozzles producing uniform, full-cone spray patterns were used: a spiral-tip nozzle from Allspray with a higher flow rate and larger drop size (the “high-flow nozzle”), and a swirl-chamber nozzle from Delavan with a lower flow rate and smaller drop size (the “low-flow nozzle”). Several other nozzles were initially tested, but these two were observed to produce the smallest drops at flow rates appropriate to the system. Most nozzles have flowrate and some have drop size specifications provided by the manufacturer. Drop size data from the manufacturer were available for the high-flow nozzle and that is used in calculations for Figure 3.2.

Experiments were run with nozzle pressures ranging from 100 to 620 kPa. Higher pressures tend to

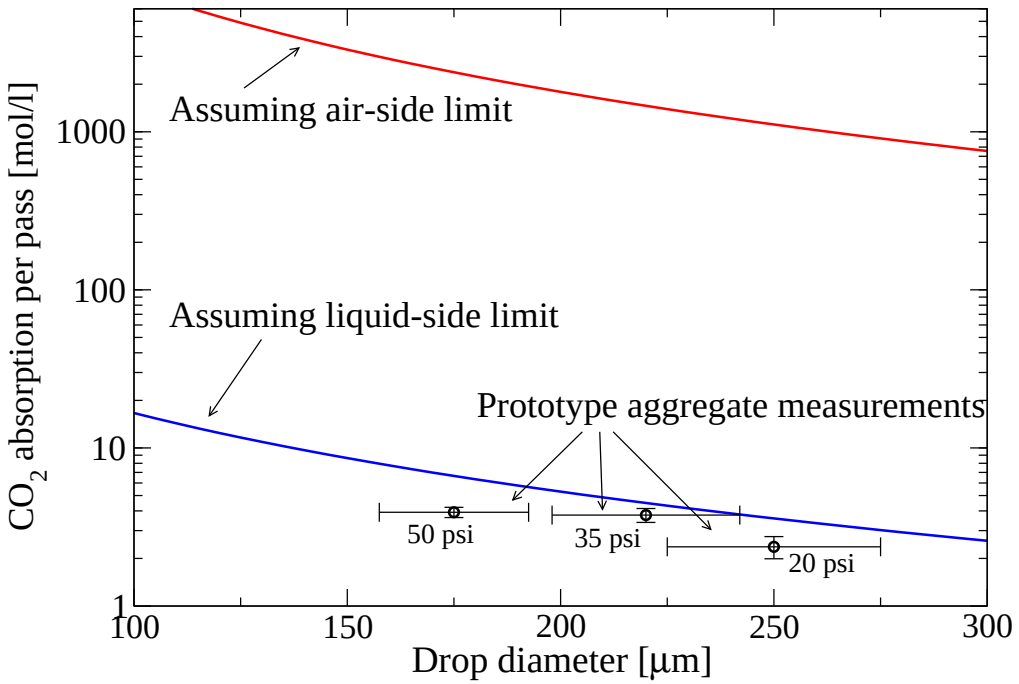

Figure 3.2: Comparison of CO_2 absorption by a single falling drop predicted assuming gas-side limitation to mass transfer (upper curve) and assuming liquid-side limitation (lower curve) with comparable prototype measurements. Results indicate liquid-side resistance is limiting. Circles represent measured aggregate absorption by a spray with the indicated mean drop diameter, as inferred from nozzle manufacturer reported data for flow rate and drop size as a function of nozzle pressure. Measured nozzle pressures are shown beneath each point. Error bars represent combined subjective uncertainties and standard error of repeated measurements.

produce smaller drops for a given nozzle, but of course require more energy. However this is only a trade-off in operation, not in design, since some types of nozzles produce very small drops at low pressures. While the nozzles chosen were most suitable among those tested, they are not optimal, and we expect that in a real system smaller drops would be produced at similar or lower pressures by careful choice of nozzle.

In order to investigate the effect of NaOH concentration on CO_2 absorption, we tested three concentrations. A 0.35 M solution represented a dilute state, with viscosity and vapor pressure about the same as water. A 5 M solution, about 20% NaOH by mass, represented a high-concentration solution which has a viscosity about 4 times that of water and is hygroscopic in some climatic conditions. 1.33 M served as an intermediate concentration (each concentration is separated by about a factor of 4).

3.2 Results

3.2.1 Mass Transfer

Figure 3.2 shows the expected CO_2 absorption assuming gas-phase limitation and assuming liquid-phase limitation to mass transfer. The results indicate that diffusion and reaction in the liquid phase limit mass

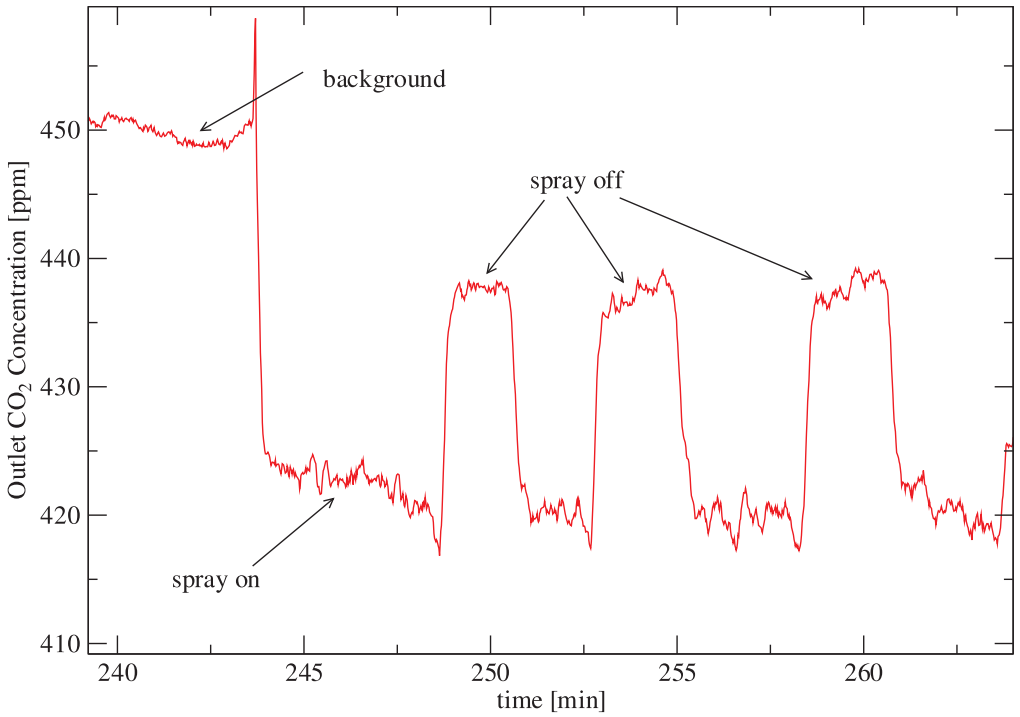

Figure 3.3: Outlet CO_2 concentration during a typical trial. Conditions: [NaOH] = 1.3M, liquid flowrate = 4 l/min, nozzle pressure = 55 psi

transfer.

In addition to estimating CO_2 absorption for a given set of parameters, Equation 3.5 has two consequences relevant to contactor design: (1) CO_2 absorption should increase with increasing NaOH concentration,

$$Q \propto [OH^-]^{1/2} \tag{3.10}$$

and (2) the rate of absorption should increase with decreasing drop size,

$$Q \propto \frac{1}{D} . \tag{3.11}$$

Experimental Results

Figure 3.3 shows the outlet CO_2 concentration during a typical trial. CO_2 is absorbed by NaOH spray and by NaOH solution on the walls, reducing the outlet CO_2 concentration compared with the background (inlet). By running the system to steady state then suddenly turning the spray off we can separate the absorption of the wall from the spray. Figure 3 shows several on-off cycles. The absorption rate is calculated from the average difference in peak heights. Under the conditions in this trial, 17 ppm of CO_2 is absorbed by the spray and 17 ppm by the wetted wall from the air passing through the system. Considering just the spray effect, that is an absorption rate of 3.7 mM per pass.

By adjusting nozzle pressure with other parameters held constant, we indirectly measured the effect of drop size. We combined our measurements with manufacturer-reported data for flow rate and mean

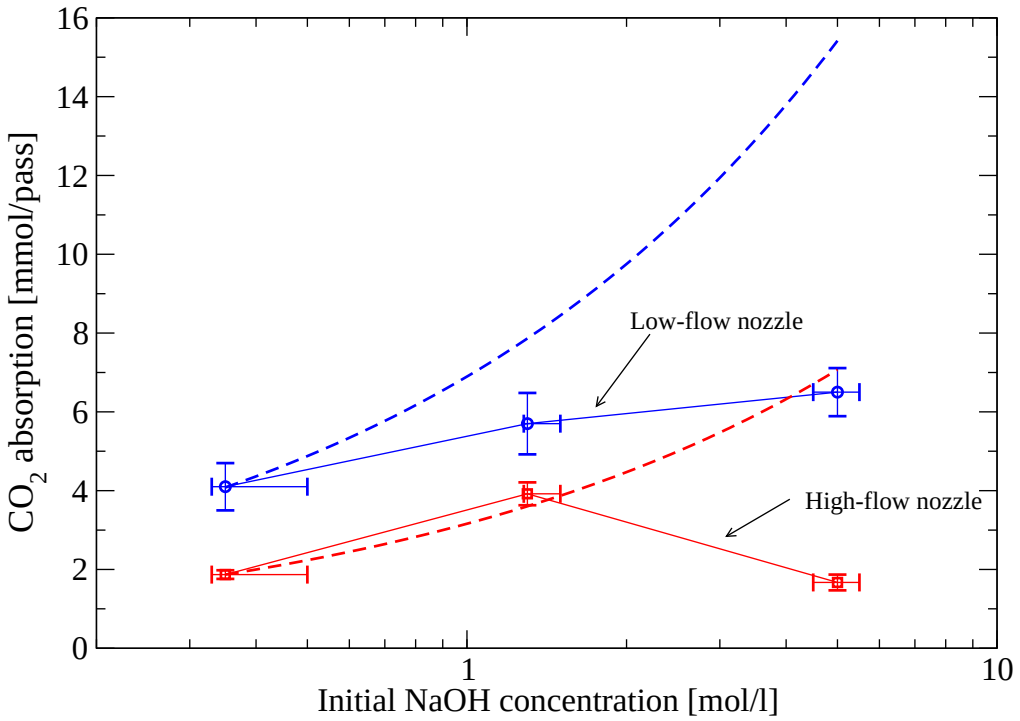

Figure 3.4: CO_2 absorption for several solution concentrations of NaOH and two nozzles. Nozzle pressure is held constant. Dotted lines indicate the theoretical effect of [NaOH] if all else were constant, i.e., Equation 3.10 fitted to the first point for each nozzle. Absorption does not rise as quickly with hydroxide concentration as would be expected in the ideal case. A likely explanation is that the higher viscosity of more concentrated solutions resulted in larger drops being formed at the nozzle.

drop size as a function of nozzle pressure (Allspray, 2002) to produce the circles in Figure 3.2. Higher nozzle pressures clearly improved absorption, with a trend and absolute value on the order of the model prediction discussed above.

With other parameters fixed, changing the solution concentration seemed to have competing effects. Higher concentration solutions should absorb better due to a faster reaction rate (Equation 3.10), but it is also known that higher viscosity solutions produce larger drops from typical nozzles (Lefebvre, 1989). Figure 3.4 shows absorption for 3 different NaOH concentrations. For the low-flow nozzle, absorption increases with concentration, but not as quickly as would be expected for sprays of constant drop size. The high flow nozzle actually peaks with the 1.3 M solution, and the 5 M solution absorbs even less than the 0.33 M solution. The effect of viscosity is highly dependent on nozzle geometry, and this may explain the discrepancy.

3.2.2 Energy requirements

The energy requirements of operating the contactor consist of mechanical work for pumping solution and for forcing air flow. Pumps must lift solution from the reservoir at the bottom of the tower to the nozzles at the top and overcome pressure at the nozzle (friction losses are comparatively small with sufficiently-sized

	Measurement	Prototype requirement	Full-scale equivalent[a]
Fraction CO_2 captured	$\Delta C = 5.3$ ppm	1.3%	35%
Solution lifting energy	$H = 3.7$ m	6.0 kJ/mol	6.7 kJ/mol
Nozzle pressure energy	ΔP_{nozzle} =550 kPa	54 kJ/mol	1.8 kJ/mol
Fan energy	$\Delta P_{air} = 70$ Pa	290 kJ/mol	10 kJ/mol

Table 3.1: Energy requirements of contactor: measurements of associated prototype parameters and log-linear extrapolation to a 120m tall contactor. Energy given per unit CO_2 captured.
[a] Log-linear scale-up of prototype results, ignoring some important effects that reduce the absorption efficiency of taller contactors (see Section 3.3).

piping). Fans work primarily to overcome the pressure drop across the particle trap and also to accelerate air. The pumping energy can be calculated from the height of the contactor and the pressure reading near the nozzle. Similarly, the fan energy is calculated from the pressure difference measured between the inlet and the duct following the mesh particle trap. Table 3.1 shows the results of these calculations for two cases: (1) our actual prototype, and (2) a contactor with the same conditions scaled up to full height. Conceptually, this latter case is achieved by stretching the prototype from a ≈4 m fall height to the 120 m that is possible for a full-scale unit. We assume that pressure drops and flow rates remain the same, while solution-lifting energy increases proportionally. We further assume that the CO_2 absorption follows a log-linear profile through the tower (reflecting first-order decay):

$$C_{out} = C_{in} e^{-Kh} \tag{3.12}$$

where C is the concentration of CO_2 in air and the decay constant, K, is calculated at $h = 3.7$ m using our empirical measurements, allowing extrapolation of C_{out} to $h = 120$ m. The validity of this log-linear extrapolation is discussed in the following two sections. We will find that coalescence of drops into larger drops diminishes the absorption efficiency of taller contactors compared with this extrapolation. The full scale values are given here only to illustrate that while the nozzle and fan energy requirements are large per unit CO_2 in the prototype, they become relatively less important in a larger system.

3.2.3 Water loss

Evaporative water loss was measured in two trials. At a solution concentration of 0.35 M, evaporation was 150 mol-H_2O per mol-CO_2 (60 m^3/ton-CO_2) and with a 1.3 M solution it was 70 mol-H_2O per mol-CO_2 (30 m^3/ton-CO_2). As expected, this is quite high, high enough that air leaving a full-scale contactor is likely to be saturated in water. Thus, to estimate full-scale water loss we apply Equations 3.8 and 3.9. Figure 3.5 shows the experimental data points along with the upper-bound estimate of water loss (assuming saturation) for the same climatic conditions as in the trials. A CO_2 capture efficiency from air of 40% was assumed for the full-scale contactor. Measured data are higher than full-scale predictions because the data reflect only the unsaturated regime of the contactor. In a full-scale system, CO_2 is also absorbed in the lower portion of the contactor where the air is already saturated, bringing down water loss per unit CO_2.

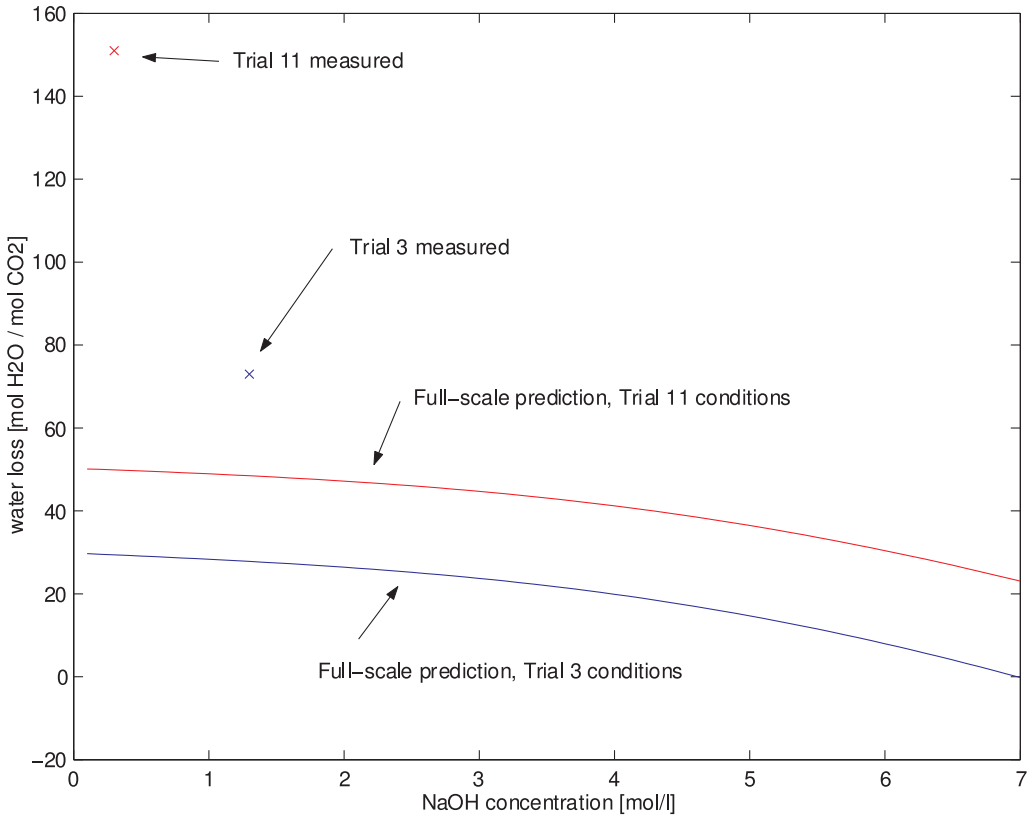

Figure 3.5: Water loss measured in the prototype during two experiments and calculated water loss in a full-scale system assuming saturation of outlet air. Ambient conditions for Trial 3: $T = 19°\text{C}$, $RH = 50\%$. Trial 11: $T = 20°\text{C}$,$RH = 23\%$.

Nevertheless, for low NaOH concentrations, water loss can be very large. However, NaOH solutions can be made at high concentrations and significantly lower vapor pressures. At sufficiently high concentrations, the solutions actually become hygroscopic and begin absorbing water from the atmosphere. The conditions for both these trials were warm and dry in global climatic terms. Air capture systems may easily be built in colder or more humid climates which cause much less evaporation or allow for lower loss-neutral NaOH concentrations. We expect that in a full-scale system water loss can be managed by adjustment of the NaOH concentration in temperate or humid climates. The cost of makeup water in lower-molarity solutions is considered in Section 3.4.

3.3 Scale-up of prototype results

We have already discussed how water loss and fan and pumping energy requirements become more favorable when moving from prototype scale to a full-scale system. We also expect capital and maintenance cost per unit CO_2 to be lower in larger-scale systems. Though the prototype was designed to represent key aspects of a full-scale contactor as realistically as possible, there are some important differences. In particular, the residence time of the spray (or, alternately, the fall height) and the spray density would be much larger in a full-scale system.

Residence time in the contactor is given by Equation 3.6. We can similarly estimate the spray density , ρ_{spray}, (volume of spray per unit volume of contactor) by:

$$\rho_{spray} = \frac{F \cdot \tau}{V} = \frac{F \cdot \frac{H}{v_{air}+v_t}}{V} = \frac{F}{A(v_{air}+v_t)} \tag{3.13}$$

where F is the liquid flow rate and V is the volume of the contactor, which can then be replaced by A, the cross-sectional area of the contactor.

The CO_2 absorption rate in the contactor is proportional to the total surface area of spray, which, neglecting complications that will be discussed below, is proportional to ρ_{spray} for a given spray distribution. In the prototype, ρ_{spray} was in the range 1–4$\times 10^{-5}$ m^3/m^3 and $v_{air} = 0.4$ m/s. These values were driven by practical limitations and sensitivity of the CO_2 measurements. While a full-scale system could be operated with these parameters (perhaps limited by the stability of such a low velocity air flow stream), the capital costs relative to the (small) quantity of CO_2 captured would be very high. We would like to use a higher ρ_{spray} to capture more CO_2 per contactor, and a higher v_{air} so that CO_2 is not depleted before reaching the outlet of the contactor. The desired values of ρ_{spray} and v_{air} should be the results of an optimization. Higher values of ρ_{spray} and v_{air} yield lower capital cost but higher pumping energy per unit CO_2 captured. We assume that values of v_{air} lower than 2 m/s would be impractical due to interference from wind. The trade-offs are discussed in more detail in the following sections. We will find that $5\times 10^{-5} < \rho_{spray} < 10^{-4}$ m^3/m^3 is reasonable for a full-scale contactor, and we consider this range in our discussion of scale-up.

3.3.1 CO_2 depletion in air

With longer residence times and denser sprays, the CO_2 in air becomes depleted toward the bottom of the tower, leading to lower absorption efficiency there. This is an effect common to a large class of absorbers, and is generally represented by first order decay (Fair et al., 2001). This effect was already accounted for in Table 3.1, and is relatively minor unless the system is tuned to capture much more than half of the CO_2 from process air.

3.3.2 Changing drop size due to evaporation

There is also an effect of changing drop size due to evaporation as the drops fall. Theoretically, absorption will decrease because of the smaller surface area as the drop evaporates, but increase because the OH^- becomes more concentrated. We know from Equation 3.5 and the preceding derivation that absorption is proportional to surface area, A_s, and to the square root of $[OH^-]$:

$$Q \propto A_s [OH^-]^{\frac{1}{2}}.$$

Replacing those terms with their relationships to drop volume,V_d, we get:

$$Q \propto V_d^{\frac{2}{3}} \cdot (V_d^{-1})^{\frac{1}{2}} = V_d^{\frac{1}{6}}.$$

As the drop evaporates, V_d gets smaller so absorption declines, but not very quickly. The change in volume for a single pass should not be very large to begin with, perhaps 20% for a high evaporation rate. This gives a maximum change in instantaneous absorption rate of about 4%. Thus, changing drop size due to evaporation does not appear to have a significant effect on scale-up.

3.3.3 Spray droplet collision, coalescence, and breakup

As drops fall through the tower, they collide with each other. If we assume a roughly uniform flow field, then this process is driven by differential settling. That is, larger drops fall faster and strike smaller drops beneath them. If the distribution of drops were monodisperse, *i.e.* all drops were the same size, then no collisions would occur. With a very wide distribution, collisions tend to be dominated by the largest drops falling very fast and striking smaller drops along the way. When two drops collide, essentially one of two things can happen: they can coalesce into a larger stable drop ("coalescence"), or they can coalesce temporarily then break apart into many, smaller drops ("collisional breakup") due to internal turbulence created by the collision or entrained air causing instability. They can also bounce off one another, retaining their original sizes, but for purposes of calculation this is not considered a collision.

The probability of coalescence given a collision is denoted E_{coal}, and so the probability of breakup given a collision is $(1 - E_{coal})$. E_{coal} is a function of the drop sizes, fluid characteristics (surface tension, viscosity, ...), collision speed and angle. For water drops falling in air and colliding by differential settling, the only situation we will consider, E_{coal} is only a function of drop size and it is found empirically. Values range from close to 1 for both drops smaller than 50 μm in diameter, to 0.5 for a 300 μm drop striking a

200 μm drop, to zero for a 1.5 mm drop striking a 3 mm drop (Pruppacher and Klett, 1997, p. 597). The function is quite complicated, but on average for a rain drop distribution, about half of collisions result in breakup (Pruppacher and Klett, 1997).

There is also another mode of breakup; "spontaneous breakup" occurs when a drop is so large that it is hydrodynamically unstable at its settling velocity. For water in air, this occurs when $D \gtrsim 5$ mm. Such a drop will break into a distribution of hundreds of smaller drops, with a volume mean around 1 mm (Pruppacher and Klett, 1997, p. 415).

With coalescence, collisional breakup, and spontaneous breakup, rain drops tend to reach a steady-state distribution with a volume mean of 0.2–2 mm (Pruppacher and Klett, 1997). Collisional breakup is the more important mechanism for determining rain spectra. The characteristic time to reach steady state is fairly long, and often rain will not have reached a steady-state distribution after falling several km to the ground (Tzivion et al., 1989).

Modeling coalescence

For some simple drop distributions, the collision rate for a given drop size can be calculated analytically, however, for most realistic distributions analytical solutions are not possible. Since this is an important problem in geophysics – similar solution methods are used to model the evolution of aerosol size distributions, the collection of drops in clouds to form precipitation, and the evolution rain drop size distributions (Pruppacher and Klett, 1997) – numerical methods have been developed. We will apply a computer code and some solution methods developed for these applications to conditions in a contactor to try to estimate the importance of collision to overall mass transfer.

We begin our analysis with a code developed for predicting global aerosol size distributions in general circulation models (Adams and Seinfeld, 2002). It solves the Stochastic Collection Equation by the method of Tzivion et al. (1987). The drop distribution is divided into size bins (in our implementation, 50 bins, spaced log-linearly), and then for every possible pair of sizes a collision rate is calculated. The collision rate is a function of the two sizes and the number density in each bin. This rate can be calculated for the relatively simple distribution containing two sizes (the average bin sizes in the pair), each with known number density. The code makes the assumption that every collisions results in the two drops coalescing to form a larger drop. This is a good assumption for aerosols, for which the code was intended, but not for the larger drops in a spray distribution. Thus we can take the results from this model as an upper bound on the rate of coalescence, since it is overestimating this occurrence. The code redistributes the mass and number of drops in each bin according to the calculated collisions in each time step. The mass and number distributions are allowed to change somewhat independently so that the average drop size floats within each bin.

The probability per unit time of a drop of diameter D_1 colliding and coalescing with a drop of diameter D_2 is referred to as the "collection kernel", $K(D_1, D_2)$. The collision rate between two bins is then proportional to K. We replace the collection kernel for aerosols, which is driven by Brownian motion, with a polynomial approximation to the collection kernel for larger drops (Long, 1974):

$$K(D_1, D_2) = E_{coal} \frac{\pi}{4} (D_1 + D_2)^2 (v_{t,1} - v_{t,2})$$

where $v_{t,1}$ and $v_{t,2}$ are the terminal velocities of the drops. We retain the assumption $E_{coal} = 1$ so the probability of collision and coalesce are equivalent. Certain features of the code for interfacing with a general circulation model are stripped and the number of spatial cells in the model is reduced to one.

We are left with a pseudo-Lagrangian box model where the spray evolves in a single control volume traveling at the average spray velocity. This model is not strictly valid because it ignores the differential settling rates of the drops. Larger drops would leave the contactor sooner than smaller ones so total mass in a control volume is not conserved. Since the objective of this first calculation is an upper bound, this effect is ignored, leaving us with an overestimate of the prevalence of large drops and so an overestimate of the collection rate.

To run the model we need an initial drop distribution and number density. We have a sense for the volume-mean drop size for some of the nozzles used in the prototype, but not a fully characterized distribution. However, Spielbauer (1992) reports the "spread" of spray distributions by various measures, including the maximum, minimum, and average spread for full-cone pressure-swirl nozzles like the low-flow nozzle used in the prototype. Spray distributions are typically log-normal, so combining the measured spread with a mean drop size yields a distribution. We chose a volume mean of 150 μm, roughly the implied value for the low-flow nozzle in the prototype and on the small end of means we know to be achievable with single-fluid nozzles. Figure 3.6 shows two of the distributions used to initialize the model. The "average" distribution is fitted to the average spread of a full-cone nozzle reported by Spielbauer (1992), (represented by $\frac{D_{0.9}}{D_{0.1}}$=3.3 – the 90th volume percentile diameter divided by the 10th percentile diameter), *i.e.* $\sigma = 1.6$, which we reasoned would be a good representation of the distribution in the prototype. Considering that narrow distributions are less prone to coalescence and that a contactor designer would choose nozzles with a minimum distribution spread, we also ran a "narrow" distribution, fitted to the minimum measured spread reported by Spielbauer: $\frac{D_{0.9}}{D_{0.1}}$=2.4, giving $\sigma = 1.4$.

After imputing the initial distribution and running the model, we have mass and number distributions in the reference volume for each time step. Summing across the size bins and making an appropriate conversion, we calculate the total surface area at each time step. Results are shown in Figure 3.7. Two spray densities are shown: one matching the prototype conditions with the low flow nozzle, and five times that density, representing a desirable full-scale density. The denser spray starts off with more surface area, but coalesces more quickly, losing 92% of surface area averaged over a one minute residence time. CO_2 absorption should be proportional to surface area, so a contactor with these conditions would only absorb 8% of the CO_2 predicted from its initial spray distribution. Figure 3.8 shows the distribution at several time steps. The distribution moves rapidly toward larger drops and quickly populates the multi-mm size bins. Such a reduction in surface area would render a spray-based contactor infeasible. Thus the bound on the coalescence effect provided by this first calculation is too high to be useful. As noted above, this simplified model neglects some mechanisms that may significantly reduce coalescence. In reality, we would expect the distribution not get larger than a steady state distribution for rain, and we know that 1 cm drops would not exist.

In order to address these limitations, three changes are made to the model: (1) changing the control volume from a Lagrangian box model to a Continuous Flow Stirred Tank Reactor (CFSTR) model of the entire contactor, (2) including spontaneous breakup of large drops, (3) adjustment of the collection kernel so that breakup collisions are not counted as coalescence.

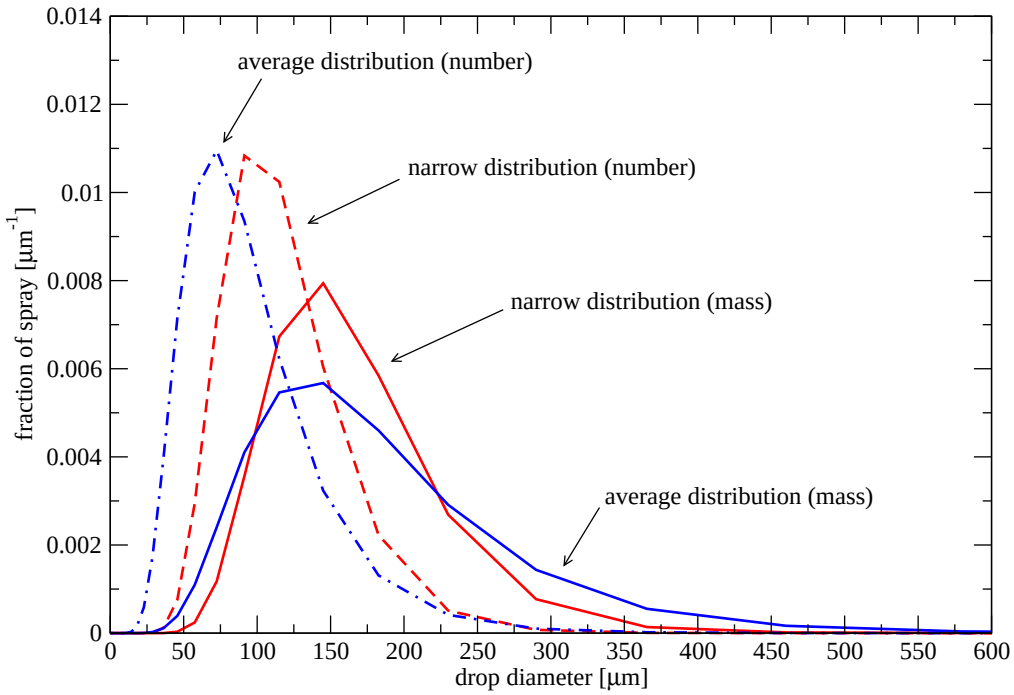

Figure 3.6: Size distribution of sprays used in coalescence modeling. Both distributions are log-normal with a volume-mean diameter of 150 μm. The "average" distribution has $\sigma = 1.6$ and the "narrow" distribution has $\sigma = 1.4$.

By running the model in a CFSTR mode, we can account for the smaller residence time of larger drops. The largest drops fall quickly enough that they have a residence time of about a fifth of the average size input drops, so this change should have a significant effect. In this mode, the model is seeded with the initial spray distribution as in the earlier calculations. Then at each time step, drops are removed from each size bin according to the settling velocity of that size and the air flow rate through the contactor. Drops are then added from the spray distribution according to the liquid flow rate. The contactor height is assumed to be 120 m and v_{air} is 2 m/s so that the overall change rate is 0.017 s^{-1}. The model is then run to steady state (about 2 minutes). The total surface area of the steady state distribution, S, represents the average efficiency of the contactor. It turns out that moving to the CFSTR mode moves S from 0.3 to 0.57 m^2/m^3-contactor – a significant improvement, but still only 15% of what would be expected in the absence of coalescence.

Next we add a provision to the code for spontaneous breakup. We can approximate the process well by assuming that every drop larger than a diameter of 5.1 mm breaks apart (Pruppacher and Klett, 1997). Thus, at each time step the mass of drops in the three largest size bins is removed and the fragments are redistributed according an empirical size table given by Pruppacher and Klett (1997, p. 415).

The third modification is to adjust the collection kernel so that it does not count breakup as coalescence. A full treatment of collisional breakup would require substantial rewriting of the core solution mechanism in the model, and is beyond the scope of this research, but we can make adjustments to the kernel to prevent some of the coalescence we know should not be occurring. In principle we would like to multiply the kernel by E_{coal}, and we would end up with about half as many coalescence events overall. E_{coal}, however

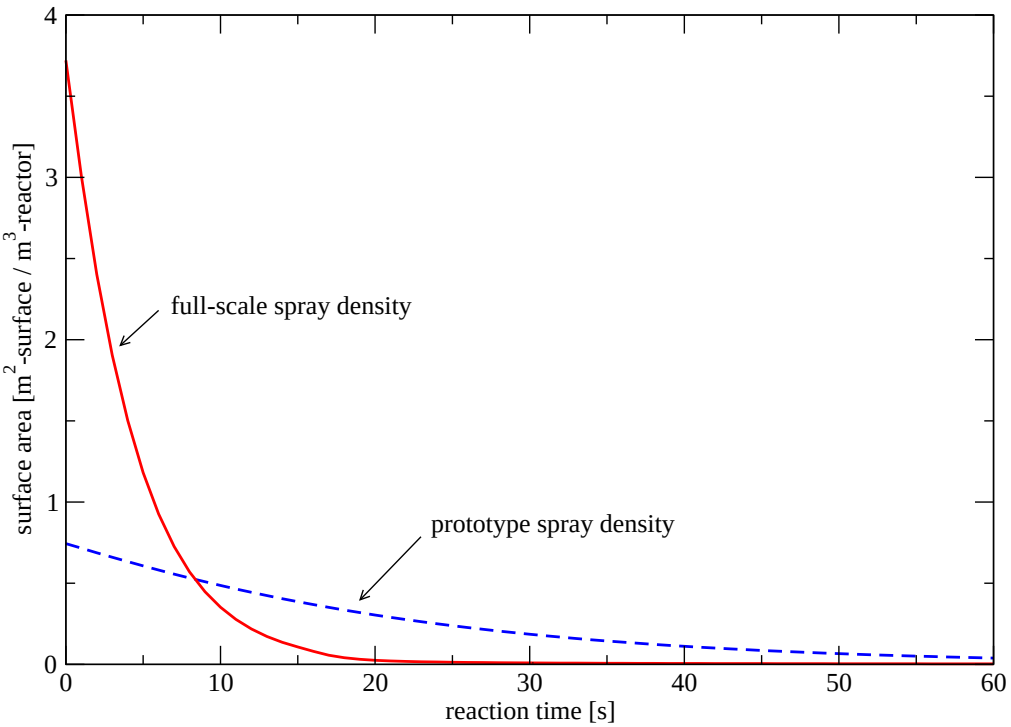

Figure 3.7: Total surface area of a parcel of spray over time in the contactor. Surface area decreases dramatically due to drops colliding and coalescing. Results are for the "average" input distribution in 3.6 with (1) prototype spray density, $(1.5 \times 10^{-5} \frac{m^3 spray}{m^3 reactor})$ and (2) a larger spray density (7.5×10^{-5}) more desirable for a full-scale system. Typical target residence times for the spray in full size contactor are 20–50 s.

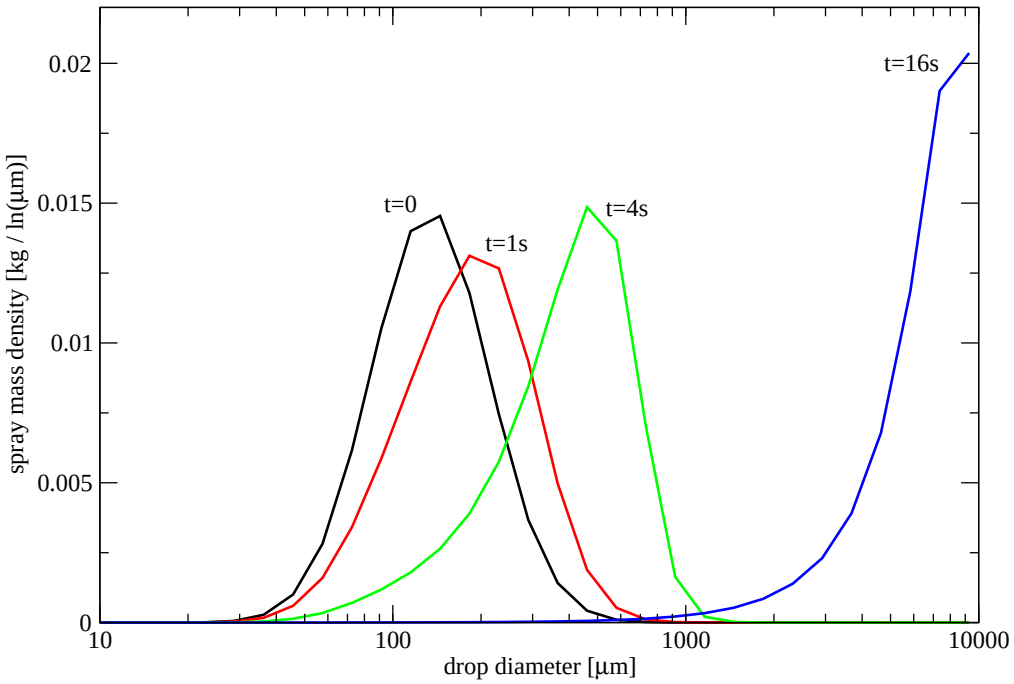

Figure 3.8: Size distributions over time for model calculation in Figure 3.7 (full scale density). Drops rapidly converge on the largest size bin ($D = 1$cm).

is a very complex, empirical function which is difficult to code precisely. Instead, we will approximate it by breaking the collision space (diameter of first drop by diameter of second drop) into 5 regions, and applying a constant which is the approximate average E_{coal} for each region.

The results of the CFSTR model with spontaneous breakup are presented in Figure 3.9 for the adjusted and unadjusted collection kernel. The steady-state distribution is bimodal, reflecting the small drops of the input spray combined with the aggregated drops which approach the largest stable size. The sharp downturn the edge of the steady state distributions is at the largest stable size. Addition of spontaneous breakup had a negligible impact on steady state surface area. This may be because the largest drops have about the same collection efficiency as the 0.5–2 mm size drops they break into when accounting for the difference in residence time (1 mm drops collide with less frequency but stay in the contactor longer). Adjustment of the kernel has an effect, but perhaps not as large as would be expected: S is increased by a little over 20%. Even a blanket reduction of the collection kernel by 0.5 (halving the collision rate) only increases S by 45%, to 0.84 m^2/m^3 (case not shown).

Parameter optimization

Given that coalescence appears to be a very serious problem, we consider contactor design parameters that may mitigate its effects. Clearly, if coalescence is as bad as the previous results suggest, a dense spray from the top of a 120 m contactor is not the optimal solution. To explore this issue further, we continue with the CFSTR model with spontaneous breakup and adjusted collection kernel, varying several contactor parameters.

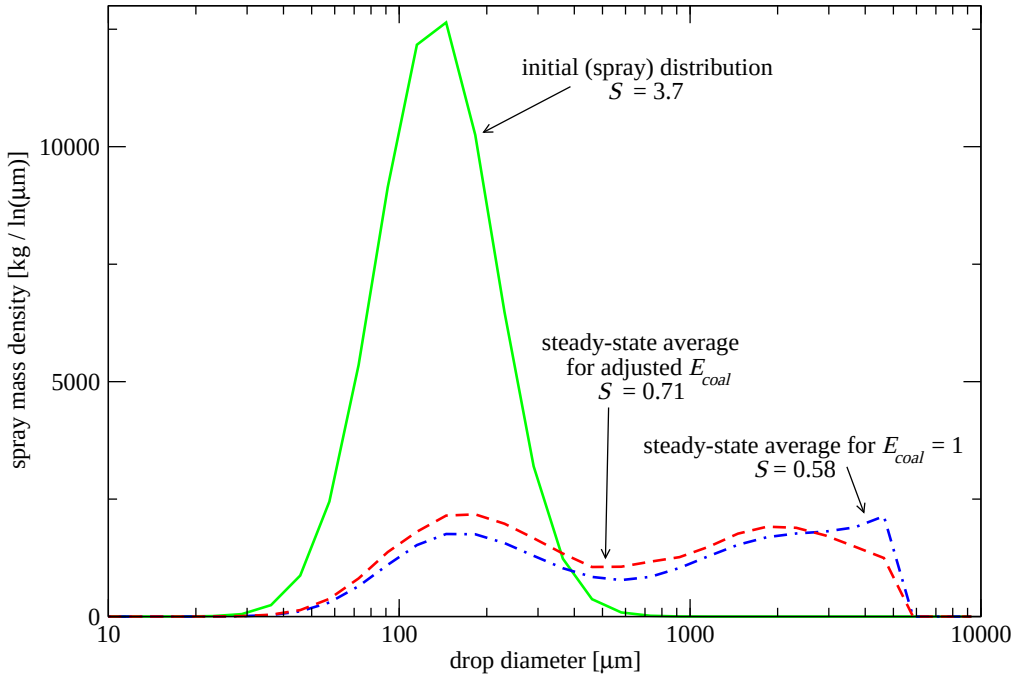

Figure 3.9: Initial and steady-state average size distributions of spray in CFSTR model of coalescence. Surface area S, and total mass drop dramatically from the input spray. E_{coal} is the probability of coalescence given a collision ($E_{coal} = 1$ implies every collision results in coalescence). S is the average surface of the spray per unit volume of contactor in m^2/m^3. Input spray is the "average" distribution from Figure 3.6 (same as previous two figures) and "full-scale" spray density from Figure 3.7 (same as previous figure).

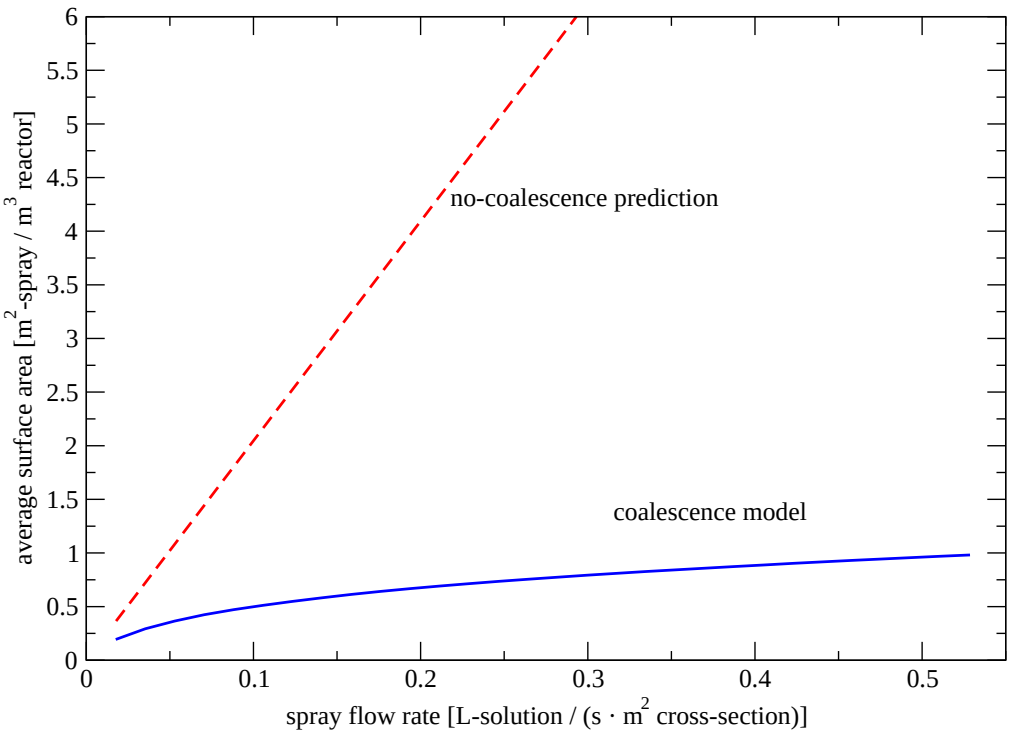

Figure 3.10: Spray surface area as a function of liquid flow rate in a hypothetical contactor. The deviation from a no-coalescence prediction increases with flow rate (*i.e.* spray density). Likewise, higher efficiencies (reactive surface area per unit liquid) are achieved at lower flow rates. Assumed contactor height is 120 m and air flow velocity is 2 m/s.

Perhaps the most obvious parameter to adjust is the spread of the spray distribution. However, moving from the average to the narrow distribution (see Figure 3.6) has no significant effect on surface area. This may be an artifact of the model. Since it is "well-mixed", large drops are distributed evenly throughout the volume. With the bi-model distribution in Figure 3.9, the width of the smaller mode may not matter much since most collection is probably due to drops in the large mode striking drops in the small mode. If the model were resolved with height, we would see the spray coalescing more slowly near the top of the tower, and bringing the average surface area up. This behavior was noticed in the Lagrangian model. The narrow distribution is used in the following calculations although, again, it doesn't seem to matter.

The way the curves cross in Figure 3.7 suggests there may be an optimal liquid flowrate for maximum surface area. Figure 3.10 shows S as a function of flowrate. With this model, a peak is not obtained, but the diminishing returns are obvious. In the absence of coalescence, we would expect S to increase proportionally to F.

Contactor height also has an effect on spray density. Without height resolution, the model only partly captures this effect through the larger change rates for shorter contactors. The results are shown in Figure 3.11. The effect is substantial, though only impractically short towers begin to approach the no-coalescence efficiency.

The last parameter we will vary in the model is mean drop size. Modest changes can be achieved with choice of nozzle, but dramatically smaller drops are only possible by moving to an air-assist nozzle. We

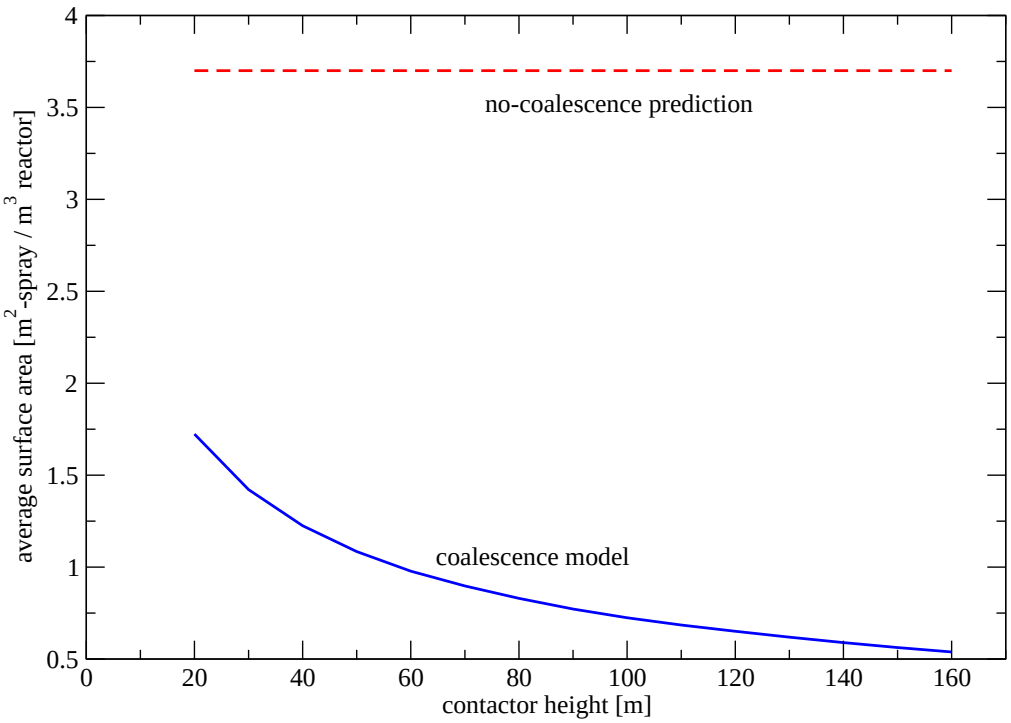

Figure 3.11: Spray surface area as a function of contactor height. Shorter fall distances produce higher efficiencies (reactive surface area per unit liquid or per unit contactor volume). Total absorption is still reduced in shorter towers because of reactor volume is also proportional to height.

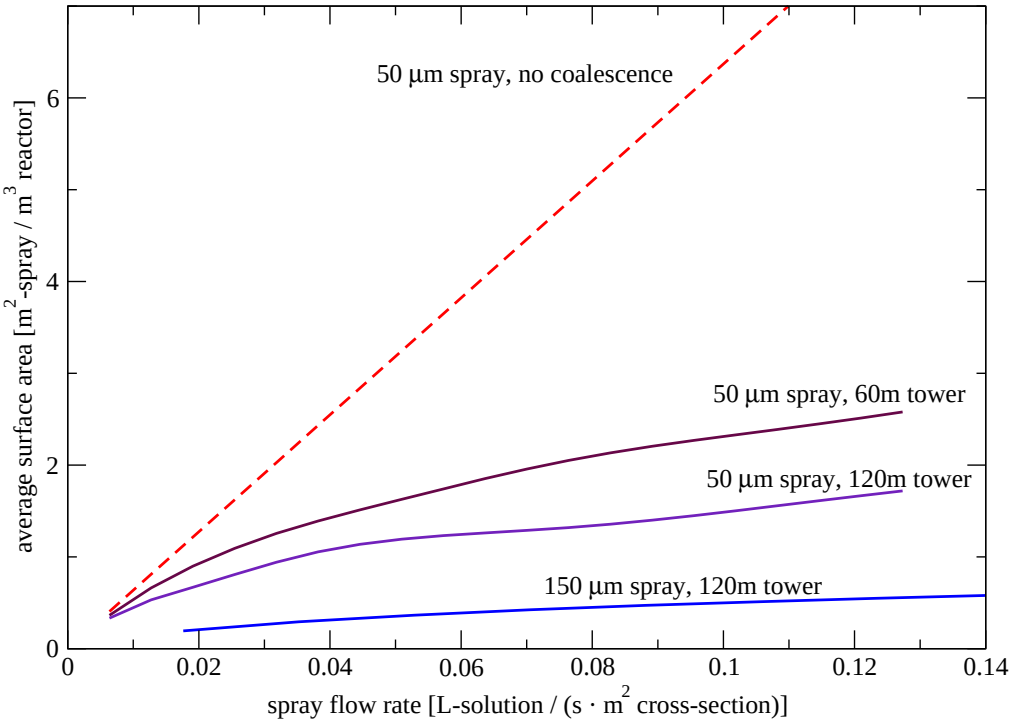

Figure 3.12: Surface area as a function of liquid flow rate for a smaller drop size distribution. Smaller drops allow more reactive surface area at smaller liquid flow rates, but do not escape the coalescence problem. Input volume mean diameter is 50 μm, typical of a dual-fluid nozzle, and σ=1.4. Contactor conditions as in Figure 3.10

did not experiment with these because single-fluid nozzles appeared to offer sufficient mass transfer with comparative technical simplicity, but coalescence problems may necessitate the smaller drops and more controlled distributions possible with air-assist nozzles (see Section 3.5.1 for a more detailed description). Figure 3.12 shows the results for a 50 μm distribution (typical of an air-assist nozzle). Surface area is substantially increased compared with the 150 μm distributions, but still scarcely reaches $S = 2$ m^2/m^3, compared with the target value in the no-coalescence base case of 3.5 m^2/m^3.

Limitations of the model and conclusions

There are two essential limitations of this coalescence model: it does not account for collisional breakup and does not have height resolution. Unfortunately, building a model which accounts for these effects is beyond the scope of this research. It is hard to say, *a priori*, how important these features would be, but it seems likely that they could substantially increase the predicted surface area in the contactor. The lack of height resolution, in particular, dampens the effects of any parameter adjustment on average surface area. Distribution spread, drop size, contactor height, and flowrate were all surprisingly ineffective at changing *S*. A height-resolved model may lead to a set of parameters yielding a sufficiently high *S*.

However, in the absence of any favorable model results or more detailed empirical results, we must conclude that coalescence may pose a serious threat to the feasibility of the suggested contactor design.

As a bound on the effects of coalescence, we will consider contactors where $S = 1\ \mathrm{m}^2/\mathrm{m}^3$ in a single fluid system, and $S = 2$ in an air-assist system.

3.4 Contactor Cost

Estimating the cost of a contactor is a dual problem. On the one hand, we try to estimate the capital and operating costs of a device to be built in the future to which no complete analogue currently exists. On the other hand, we must assume that the future engineers of the device will optimize the design to minimize costs, and they will have considerable leeway to do so. So the two sides of the problem, specifying the design and estimating the costs, feed back on each other.

Air capture only makes sense in a very large scale deployment, so we can expect that engineers designing such devices will not be limited to off-the-shelf technology or process experience from other industries, especially since here we are not concerned with the cost of early air capture plants, but of the average or "*n*th plant" cost. In order to estimate that cost today we have to make some informed judgment about what the optimal system will look like. We must do so under considerable uncertainty about some parameters – uncertainty that would be resolved for those future engineers. In particular, our uncertainty about the full effects of coalescence and breakup, and about the technical potential and costs of relevant technologies, makes specifying the contactor difficult.

Ideally we would like to build all of the relevant parameters and functions into a cost model and perform a nonlinear optimization to find the system which minimizes cost per unit CO_2 captured, as the engineers and operators of a real plant are likely to do. However, we do not have sufficient knowledge of the functional relationship between, for example, capital cost and contactor height, or spray nozzle type and operating cost, to complete such a model. Instead we will choose several scenarios, and for each scenario choose a set of reasonable (but not optimal) parameters to use in a simple cost model. The goal will be to choose parameters such that they are on the order of the likely optimums while being well withing the practical capabilities of known technology.

We know our model overestimates the rate of coalescence, so we will bound the effect of coalescence by considering no-coalescence cases and cases based on the model results. We will also consider systems using single fluid nozzles and using air-assist nozzles.

3.4.1 Mass transfer

We define the CO_2 absorption rate constant, k_{spray}, such that

$$\frac{dC}{dt} = S k_{spray} C(t) \tag{3.14}$$

where C is the CO_2 concentration in air in mol/m^3. Then k_{spray} is the absorption rate per unit CO_2 in air per unit drop surface area, with units:

$$k_{spray} \sim (\frac{\mathrm{mol}}{\mathrm{s}})(\frac{1}{\mathrm{m}^2})(\frac{\mathrm{m}^3}{\mathrm{mol}}) = \frac{\mathrm{m}}{\mathrm{s}}$$

The reactive surface area during trials in the prototype cannot be precisely known, but we can makes estimates using available information about nozzles, flow rates, and other parameters. Using our experimental data with these estimates we calculate values of k_{spray} of 1–3$\times 10^{-3}$. The range is from the low molarity solution to the high molarity solution, as expected. We may note by comparison with Equation 3.3, that theoretically:

$$k_{spray,theory} = \frac{C_0}{C}\sqrt{D_l k[\mathrm{OH}^-]}$$

Calculated values of k_{spray} using this equation are about twice those calculated from data, reflecting that absorption in the contactor was about half of what was expected from theory and available drop size information. The discrepancy may be due to some mechanistic inefficiency in mass transfer, less-than-expected drop surface area, or some other effect. We will use the highest empirical value in the cost calculations.

We approximate S as uniform with height (this is not strictly accurate, but the distribution of S will not affect the average mass transfer significantly), so that Equation 3.14 can be integrated to yield:

$$C = C_{in} e^{-S k_{spray} t}$$

To get the outlet concentration, C_{out}, we evaluate this at $t = \frac{H}{v_{air}}$:

$$C_{out} = C_{in} e^{-S k_{spray} H / v_{air}} \tag{3.15}$$

The rate of CO_2 capture of the contactor (denoted $\dot{M}$) is then:

$$\dot{M} = (C_{in} - C_{out}) \cdot v_{air} \cdot A = C_{in}(1 - e^{-S k_{spray} H / v_{air}}) \cdot v_{air} \cdot A \tag{3.16}$$

3.4.2 Capital cost

A contactor can be constructed by modest modification of a power plant cooling tower. This may not be the optimal form of an air capture contactor, but such structures are assured to be possible and we can draw capital cost estimates from the power industry. Other examples of spray-based reactors exist in industry, like SO_2-scrubbers for power plants, but cooling towers are the largest, most closely approaching the large scale desired for air capture.

There are two basic types of cooling towers: natural draft and forced draft. Natural draft towers often have a hyperbolic profile, are constructed of concrete, and can be very tall, as high as 160 m, though 90–120 m is more typical. They make use of the convective forces generated by their shape, height, and temperature gradient created by the spray to move air through without a fan (hence the "natural" draft). Structures include a foundation, spray collection basin, pumps and piping, and often particle filtering ("plume abatement") mechanisms. The main differences between a conventional cooling tower and a version used for air capture would be: (1) our liquid flow rate would be smaller by an order of magnitude, requiring fewer pumps and smaller piping, (2) we will add a fan or bank of fans to control the air flow, and (3) we will add a high-efficiency particle filter ("demister").

For simplicity, we suggest a single, large fan essentially identical to a wind turbine without the tower

Draft type	Height [m]	Cross-sectional area [m^2]	cost per unit [$millions]	cost with mods [$millions]	cost per cross-section [$/$m^2$]
Natural[1]	90	5100	36	41	8000
Forced[1]	40	3000	18	18	6000
Natural[2]	120	7900	25–75	31–81	4000–10000
Forced[2]	50	280	0.5–1	0.5–1	1800–3500

Table 3.2: Capital cost of cooling towers. Costs represent complete installed costs. EPA costs are from 1996, adjusted to 2006 dollars using the Construction Building Index. Upper bounds of ranges reflect towers with plume and noise abatement and unusually high site-specific costs.
[1] EPA (2002)
[2] Mykyntyn (2006)

and foundation. We take the cost of this addition as $4 million, typical in the wind industry. For the demister, our prototype demonstrated that a wire mesh filter constructed manually of stainless steel wool can be effective with an acceptable pressure drop. However, in the full-scale system a more sophisticated system is warranted. We obtained a price quote and product specifications from a commercial particle trap manufacturer (Amistico, 2006) and apply those directly, ignoring the substantial bulk discount that is likely for a project as large as even a single air capture tower. We get 500 $/$m^2$-demister. With a downward flow contactor, a reasonable placement of the demister would be as an annulus around the base of the tower. The total area of the demister can be adjusted by the height of the annulus. Demisters of this type collect drops more efficiently at higher air velocity but the pressure drop increases with air velocity. We expect that a total area of the demister of one half that of the tower cross section (demister velocity twice the tower velocity) makes a reasonable trade-off between these competing effects. This is what we assume for capital cost calculations.

Forced-draft cooling towers are more directly adaptable to air capture since they already have fans and demisters. They are smaller, however – typically 20–50 m high and arrayed in square cells 10–20 m on a side. They can be constructed of concrete or fiberglass for similar costs. Again, the liquid flowrate in an air capture version would be about a tenth that of a conventional version. Also, forced draft towers have some “splash fill” material which we will not require. Air flow velocities are similar. We will apply contactor costs for forced draft cooling towers directly to an equivalently-dimensioned contactor.

Table 3.2 shows capital cost estimates for power plant cooling towers and includes the cost of modifications to natural draft towers described above. Personal communication with industry experts and estimating documents from the EPA (2000, 2002) were used to arrive at the estimates. Cost is usually given per unit liquid flow, which is the primary figure for which cooling towers are generally sized. Typical flow rates and sizes were then used to calculate the implied cost per unit cross-sectional area and cost per physical structure.

The costs span a fairly wide range. This may reflect the quality of the data sources more than actual uncertainty in cooling tower construction; cost information is proprietary and industry tends to be loathe to share it. Some of the variation is due to inclusion versus exclusion of components, particularly noise and plume abatement. We will use EPA costs as the base case, and consider the lowest costs in the sensitivity

analysis.

3.4.3 Operating cost

The physical operating and maintenance cost of a cooling tower is typically estimated at 4% of capital per year. There is some argument that the fraction is 5% for smaller cooling towers and drops to 2% for the largest towers (EPA, 2002), but the flat 4% is a widely accepted rule of thumb. It includes purchase of makeup water at typical industrial prices ($0.13/m^3) – much more water than is likely to be used in an air capture system. It also includes periodic replacement of pumps of which an air capture system would require 1/5 as many. So this operating cost may be an overestimate, but neither water nor pump maintenance makes up the majority of operating cost, so it is retained.

Electricity for running pumps and fans is the other important operating cost. Electricity is used by pumps to lift solution up the height of the tower and to overcome the pressure drop across the nozzle, ΔP_{nozzle}. The fans use energy to overcome the pressure drop across the tower, which is dominated by the particle trap, and to accelerate the air. Frictional losses are considered comparatively small and are neglected. The total rate of energy use is then given by:

$$\dot{E} = \frac{F}{\varepsilon_{pump}}(\Delta P_{nozzle} + \rho_{soln} g H) + \frac{A}{\varepsilon_{fan}}(\Delta P_{air} v_{air} + \rho_{air} v_{air}^2) \tag{3.17}$$

where ρ_{soln} and ρ_{air} are the densities of air and of the solution and ε_{pump} and ε_{fan} are the mechanical efficiencies of the pumps and fans, respectively.

For the cost calculations we assume $\Delta P_{nozzle} = 350$ kPa for the single fluid case (a reasonable value, though smaller pressure-drops are possible). In the air-assist nozzle case, we use a nozzle pressure of 280 kPa and air/liquid volume ratio of 5. The latter figure is on the small side of typical values, so there has been some optimization of nozzle choice. We take the efficiency of the pump as 80% and the efficiency of the fan as 70%. We purchase carbon-neutral electricity at the price, p_{elec}, of $0.07/kW-hr, or $19/GJ. The reflects roughly the cost of base-load electricity from nuclear or CCS plants. In practice, the electricity could be generated on-site, potentially with excess sold to the grid. However, using a simple market rate for electricity allows us to decouple the costs of generator and contactor.

3.4.4 Scenarios

As previously mentioned, we consider scenarios based on both mechanical draft and forced-draft cooling towers, with mass transfer cases with and without coalescence, with the intent to bound the contactor cost. They are as follows:

1. Favorable conditions. In this scenario, we do not choose all parameters to be as favorable as possible or make any unrealistically hopeful assumptions, however several important parameters are set at the favorable side of their reasonable range. The low end of capital cost is assumed for the larger natural draft cooling tower. Dual fluid nozzles are assumed, and no coalescence takes place.

2. Median assumptions, no coalescence. In this scenario, the higher EPA capital costs for a natural-draft tower are used. Modest, single-fluid spray conditions are assumed (parameters are not optimized for a no-coalescence world, but rather left at low density where coalesce may not matter when all effects are accounted for).

3. Median assumptions with coalescence. This uses EPA costs for the forced-draft tower and modest, single-fluid spray conditions.

4. Dual-fluid nozzle with coalescence. Same as 3, but with dual-fluid nozzle conditions.

5. Coalescence with dual fluid, lower capital. This scenario uses the low capital estimate for a forced draft tower, illustrating the strong influence of capital cost in coalescence conditions.

6. Optimized. This scenario uses the lower cost estimate for a forced draft tower, dual-fluid nozzles, and no coalescence. Additionally, some parameters are tuned to slightly more favorable values based on what, in the author's subjective opinion, would be possible in an optimized system. The nozzle pressure is dropped from 280 kPa to 180 kPa, the pump efficiency is brought from 80% to 90%, and the spray constant is increased to $4 \times 10^{-3} \frac{\mathrm{m}}{\mathrm{s}}$, reflecting a higher molarity solution.

3.4.5 Total cost

We consider three main components of the contactor cost: capital, operation and maintenance, and electricity for operating pumps and fans. Though there are other costs, we expect these to dominate. Capital is amortized at a 15% capital charge rate. If we denote the amortized capital cost by Cap and the yearly operating cost, excluding electricity, by $O\&M$, then the total cost per unit CO_2 captured is:

$$\frac{\text{total cost}}{\text{CO}_2 \text{ captured}} = \frac{Cap + O\&M + p_{elec}\dot{E}}{\dot{M}} \tag{3.18}$$

where $\dot{M}$ is calculated from equation 3.16 and $\dot{E}$ is calculated from Equation 3.17.

Parameter choice

For each of the scenarios described in the previous section, we have a fixed geometry, spray distribution, capital, and maintenance cost. The two remaining tunable parameters are v_{air} and F. When these are chosen, $\dot{E}$ and S can be calculated and then $\dot{M}$. Both parameters effect the total cost in complex ways. Together they determine the spray density by Equation 3.13 which in turn determines S. With other parameters fixed, larger F will increase $\dot{M}$ (by increasing S), but also increase $\dot{E}$. As sprays become more dense, coalescence causes them to be less efficient and so energy cost overwhelms the total at high F. At low F, both $\dot{E}$ and $\dot{M}$ are small and so capital cost overwhelms the total. For the scenarios, we choose F such that spray densities remain low enough that the spray retains a reasonable efficiency, or, in the no-coalescence cases, low enough that the spray could be expected to retain a reasonable efficiency if coalescence were accounted for. It is a somewhat subjective choice. The alternative would be to perform

Scenario #	1	2	3	4	5	6
Tower type	natural	natural	forced	forced	forced	forced
Capital estimate	low	mid	mid	mid	low	low
nozzle type	dual	single	single	dual	dual	dual
coalescence model	no coal.	no coal.	coal.	coal.	coal.	no coal.
Capital cost	9	39	226	140	37	8
O&M	2	10	60	37	10	2
Pumping electricity cost	18	18	28	21	18	9
Fan electricity cost	6	6	15	9	8	4
Total cost	40	70	300	200	70	20

Table 3.3: Cost estimates for the contactor. Costs given in \$/ton-CO_2.

a multivariate nonlinear optimization to find F in each scenario, but this is both computationally difficult and not likely to yield realistic results.

The choice of F is tied to v_{air} by the spray density relationship. But if we assume F is adjusted to maintain constant spray density and thus constant S, we can look at the other effects of v_{air}. Since v_{air} controls the residence time of air in the contactor, it controls the outlet CO_2 concentration. For high v_{air}, the CO_2 depletion problem described in Section 3.3.1 becomes unimportant, increasing $\dot{M}$ slightly, but $\dot{E}$ becomes very high because of the needed fan energy. At low v_{air}, the inverse is true. Another constraint on v_{air} is that very small values may not be stable because of interference from natural wind. We take 2 m/s as a minimum plausible value. Otherwise v_{air} is set to maintain the captured fraction of CO_2 between 40 and 60%. The optimal value appears to be in this range, and variation within this range does not change costs drastically.

For example, in Scenario 2, we have a 90 m tall tower, single fluid nozzle, and no coalescence. For this scenario we will target a spray density yielding $S = 6\,\mathrm{m}^2/\mathrm{m}^3$. The top curve in Figure 3.10 illustrates the relationship between F and S for the no-coalescence, single-fluid case, and F is $0.3\,\mathrm{L}/(\mathrm{s}\cdot\mathrm{m}^2)$ where $S = 6$ for $v_{air} = 2\,\mathrm{m/s}$. Plugging the values into Equation 3.15 we find that CO_2 depletion is somewhat high. It turns out that a higher v_{air} reduces total cost, so we move to $v_{air} = 4$ and F increases proportionally to $0.6\,\mathrm{L}/(\mathrm{s}\cdot\mathrm{m}^2)$. We could have alternately chosen $v_{air} = 3$ resulting in lower energy use but higher relative capital for about the same total cost. Running these parameters along with the capital and maintenance costs described in previous sections through Equations 3.16–3.18 gives the total cost.

Results

A spreadsheet model was used to calculate costs for each of the six scenarios following the procedure in the example for Scenario 2 above. A summary is presented in Table 3.3.

We first see that even the lowest cost, for the "optimized" system, is significant on the scale of the total system, which has been previously estimated at 150 \$/t-CO_2 (Keith et al., 2006). Costs are overall large, even in the case of no coalescence. Energy costs are the majority in the two lowest-cost scenarios, elsewhere capital cost dominates. Particularly in the scenarios with coalescence, capital cost becomes huge so that, in two cases, the contactor costs are likely to overwhelm costs of the balance of the air

capture system. We can also see that some adjustable parameters have dramatic sway over the contactor cost, which varies by about a factor of 4 among the no-coalescence scenarios, and among the coalescence scenarios.

3.5 Contactor technology and sensitivity of future cost

We have attempted to estimate the cost of a NaOH-spray based contactor in a simple and transparent way, primarily as a proof-of-concept for a spray-based system. Any such estimate of future technology is inherently uncertain. If and when such devices are constructed in large numbers, unforeseen problems will likely drive costs up and, as well, clever engineering, parameter optimization, and new upstream technologies will tend to drive costs down. This sections describes factors that may significantly influence the future cost of a full-scale contactor.

3.5.1 Spray technology

There are basically two classes of nozzles commonly used to generate small drops in industrial applications: single-fluid, in particular, "pressure-swirl" type nozzles, and two-fluid, or "air assist" nozzles. The pressure-swirl nozzle generates turbulence by pushing the liquid through specially-designed "swirl chamber" before it exits through a small circular orifice and breaks apart. For a given nozzle, higher pressures at the nozzle generate smaller drops and higher flow rates. However, nozzle size and geometry have a much bigger effect on spray distribution and flow rates: mean drop sizes range in order from 100 μm to several mm and flow rates range in order from 0.1 L/min to 100 L/min or more in commercially available pressure-swirl nozzles. A more thorough optimization across nozzle type or engineering of a nozzle specifically for air capture can significantly improve the pumping energy requirements of the system, especially if smaller drops or narrower distributions of drops are produced than what was tested in the prototype, and also if a lower pressure at the nozzle head is required. All of these things appear possible.

The air-assist nozzles are known to generate smaller, more controlled drop distributions than liquid-only nozzles, but pressurizing the air adds significant energy cost per unit CO_2 and adds complexity and capital cost. Air to liquid volume ratios on the order of 30:1 are typical, which result in an energy requirement for air compression of roughly 30 times what would be required for pressurizing the liquid alone. However the appeal of air assist nozzles is that dramatically smaller drops – volume-mean diameters of 50 μm are typical in industrial applications – and narrower distributions of drops may be possible. And lower air to liquid ratios are certainly possible. If a suitable nozzle system with, for instance, 50 μm drops and an air to liquid ratio of 5:1 can be engineered, it may offer significant energy savings over a single-fluid system, especially considering the reduced occurrence of coalescence. With smaller drops, the surface area to volume ratio is higher so the mass density of liquid needed in the tower is lower.

3.5.2 Structural design

The basic considerations in contactor design are reflected in the terms of the total cost formula, Equation 3.18: capital, maintenance (which for simplicity we will consider tied to capital), energy use, and mass

transfer rate ($\dot{M}$). We can further break down energy use into the 3 most important quantities: lifting energy, $\dot{E}_{lifting}$, nozzle head energy, $\dot{E}_{nozzle}$, and fan energy, $\dot{E}_{fan}$. Following Equation 3.17 (which is general for most spray-based contactors), we can express these quantities in terms of the parameters which system designers have substantial control over:

$$\begin{aligned} \dot{E}_{lifting} &\propto HF \\ \dot{E}_{nozzle} &\propto \Delta P_{nozzle} F \\ \dot{E}_{fan} &\tilde{\propto} \Delta P_{air} v_{air} \end{aligned}$$

where "$\tilde{\propto}$" denotes "approximately proportional to", because we are neglecting (relatively unimportant) inertial effects.

Most design decisions trade off between lowering one type of energy use and raising another, or between energy use and capital cost. For example, a very tall contactor with very low spray density would operate very efficiently, but the capital cost per unit CO_2 captured would be very large. The goal is to work with is to work with the trade-offs to minimize total cost per ton.

We have considered designs based on power plant cooling towers because the cost of such structures is well known. A better design can clearly be achieved by designing the system from scratch with air capture specifically in mind, but data were not available for us perform an optimization of structural design. With a more detailed understanding of component-wise capital costs and spray technology, and a more complete model of drop collision and coalescence, a significantly different form may emerge. Shorter towers, taller towers, fiberglass skin towers, counter-current designs, and many other variations are possible. A description of some possible alternative designs follows.

Shorter tower

Most industrial spray towers are shorter than the power plant cooling towers we have considered. SO_2-scrubbing towers are typically on the order of 10 m high, for instance. As we saw in Figure 3.11, coalescence in shorter towers tends not to be as important so that the absorption by the spray per unit height remains high: $\dot{M}$ decreases but $\dot{E}_{lifting}$ decreases proportionally. However, $\dot{E}_{nozzle}$ and $\dot{E}_{fan}$ remain unchanged and so become relatively more important to total cost. This is the fundamental trade-off in setting contactor height: at high H, nozzle and fan energy become less important but coalescence drives up $\dot{E}_{lifting}$. At low H, $\dot{E}_{nozzle}$ and $\dot{E}_{fan}$ tend to dominate. Of course, H also affects capital cost, with shorter towers presumably being less capital-intensive. However, the ratio $Cap/\dot{M}$ is the important quantity, and it is not clear how that relates to H.

If three conditions can be met, then short towers may offer significantly lower costs and than the estimates for 50–120 m towers: (1) short towers can be constructed with much lower capital than tall ones, (2) nozzles can be used with relatively low ΔP_{nozzle}, and (3) ΔP_{air} or v_{air} can be adjusted to give a sufficiently low $\dot{E}_{fan}$. At least condition 2 appears likely from our knowledge of commercial spray technology.

Multistage spray

One possible solution to the coalescence problem is to add the spray in multiple stages along the tower, as is common in cooling towers. This has the effect of refreshing the spray distribution as it falls (and as drops from further up coalesce and become less efficient). The energetics are in many respects similar to a collection of short tower systems with height equal to the height of a stage. Except in this case the air is passed from one stage to the next so that the pressure drop across the particle trap need only be overcome once for all stages. That is, the problem of high $\dot{E}_{fan}$ is solved compared with a short tower system. A drawback is that solution must be pumped to the upper stages even though it only participates in absorption for a fraction of the total height, thus $\dot{E}_{lifting}$ can be larger in a multistage system compared with an equivalent collection of short towers. Perhaps a structure which collects spray between each stage can be included so that solution is mostly only lifted the height of one stage. In any case, the problem of high $\dot{E}_{nozzle}$ is shared with short tower systems. A multistage spray system would require a low ΔP_{nozzle}.

Upward flow

In forced-draft cooling towers, air flow is typically from bottom to top with spray nozzles at the bottom. This increases residence time of the drops since gravitational settling works against the flow rather than with it, (flow in our prototype was downward). With this design, the sign of v_t in Equation 3.6 changes to give a longer drop residence time of

$$\tau_{up} = \frac{H}{v_{air} - v_t}.$$

Following equation 3.13, we get a higher spray density for the same F or, alternately, we require a smaller F for the same spray density. For an idealized case with monodisperse spray and no coalescence, we can calculate the reduction in F and thus the (proportional) reduction in $\dot{E}_{lifting}$ and $\dot{E}_{nozzle}$ by switching to an upward draft system. Rearranging Equation 3.13 and substituting in τ_{up}, we have:

$$\frac{F_{down} - F_{up}}{F_{down}} = \frac{2v_t}{v_{air} + v_t}.$$

This is the fractional change in flowrate between upward and downward flow systems with all else equal. It would translate to the same fractional reduction in pumping energy. Assuming $v_t = 0.6$ m/s ($D = 150\,\mu$m) and $v_{air} = 2$ m/s, we would have an increase in drop residence time of 86%, reducing pumping energy by 46%. This is, however, an idealized case. With growing drop size (and growing v_t) due to coalescence, the change in τ and reduction in $\dot{E}_{lifting}$ between an upward flow and downward flow system would be more modest. No simple relationship can be described.

With upward flow there can also be a drop size sorting effect, where the largest drops settle faster than v_{air}, leaving smaller drops to continue upward. It is not clear how much this would mitigate the coalescence problem but sorting the largest drops out of the initial distribution would seem to help.

It appears that upward flow would provide at least modest improvement to the energetics of the contactor. It was not considered in the cost scenarios largely out of practical concerns about having the particle trap at the top of the structure, placing the fan in the path of the spray, and about applicability of prototype

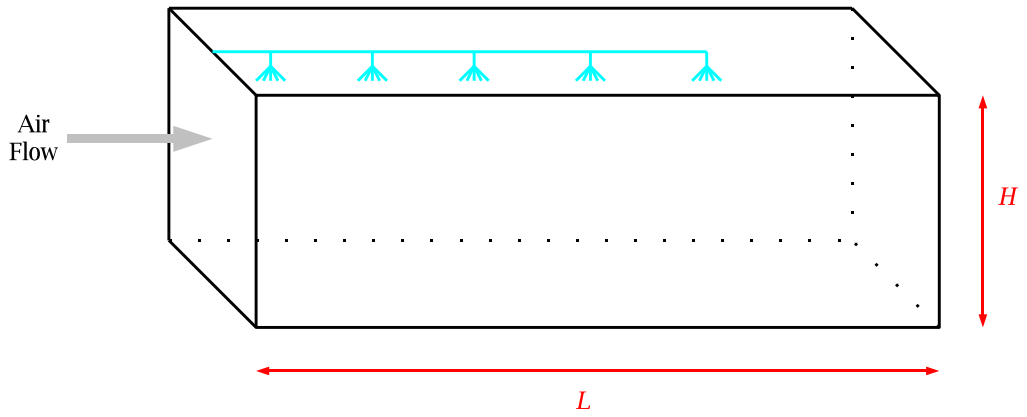

Figure 3.13: Simple diagram of a horizontal flow contactor

data, which was collected for a downward flow system.

Horizontal air flow

Another design variation that may reduce contactor cost is a system with horizontal air flow. A diagram is shown in Figure 3.13. Here H would be relatively short, perhaps 20 m, but L may be much longer, perhaps 100 m. A fan would likely be placed at the inlet and a particle trap at the outlet. Spray nozzles would be spaced along the length, leaving some distance at the end for drops to settle before reaching the outlet.

The energetics are similar to a short tower system but with two advantages. The first advantage is that the residence time of the spray is longer compared with a vertical system, given simply by $\tau = \frac{H}{v_t}$. Longer τ for a given height reduces pumping energy as described in the previous section. The other advantage is that a new parameter is introduced, L, which determines the residence time of the air independent of the residence time of the spray. Thus L can be adjusted to keep $\dot{E}_{fan}$ small. The system shares the drawback with short towers of high $\dot{E}_{nozzle}$, though this is somewhat dampened by the effect of longer residence time. A low ΔP_{nozzle} would be required.

A horizontal system may offer reduced capital cost compared with a vertical system. Although the area of wall per unit contactor volume required is larger than for a very large cylindrical cooling tower, the walls bear a much lighter load and so less material may be required overall. On the other hand, more land area and a sealed roof would be required. It is not obvious how the capital costs would compare to an equivalent vertical system.

3.5.3 Water loss

We have measured and calculated a rate of water loss that in volume terms is quite high. We have discussed how the rate of water loss can be managed with high NaOH concentrations, but it may be desirable for other reasons to run at lower NaOH concentrations. In that case, water loss is highly dependent on the meteorology of the site. Paying for water at a rate typical of power plant cooling towers does not raise overall costs significantly, however in a world where air capture is widely deployed, the demand for water would be large on the scale of developed use, and could upset already overburdened water systems.

Although, the cost of water loss may be managed even in warm, dry conditions if an inexpensive source of water is nearby. It is not obvious, for instance, why seawater would not work in the system.

3.5.4 Siting

Since CO_2 can be captured from anywhere on the globe, air capture systems have enormous siting flexibility. A large number of factors may be greatly influenced by siting, and will compete for importance when site decisions must be made. When an optimal site is chosen, the total cost may be pushed significantly up or down from our estimates. Some important considerations for siting a contactor are:

- Climatic conditions, and availability of water, for reasons discussed above.
- Local construction cost and resources. Construction cost varies by location, and this can have a significant impact on capital cost. Remote locations may be desirable for many reasons, but necessitate long transportation distances of, e.g., concrete and steel. If many towers for a large air capture plant are constructed at once (which is likely), construction may be managed like a large dam project, where a dedicated cement plant and other production facilities are constructed on site.
- Availability of inexpensive energy. Clearly access to inexpensive carbon-neutral electricity, such as hydro-power, geothermal, or "stranded" natural gas would improve the economics of the system.
- Proximity to sequestration site. Although CO_2 transportation may not be a large component of total cost, there may be situations where pipeline construction is legally difficult or practically infeasible. An air capture plant located at a sequestration site would avoid these problems.
- Availability of land. A very large scale deployment of air capture will require a lot of land, and choosing sites that contested by other uses may be important.

3.5.5 Materials and construction cost

A large portion of contactor cost is initial capital for construction of the structure. Construction costs for standard cooling towers have increased rapidly in the last several years, as much as doubling since 2000 (Mykyntyn, 2006). Much of this increase is due to high steel and concrete prices. In turn, some of that increase is due to high global energy prices, and some is probably due to a rapid increase in global demand for these and other raw materials, driven largely by extraordinary Chinese economic growth. The cost estimates in this chapter are based on current commodity prices, which one may argue are artificially high. By the time air capture is actually deployed, world supply may have caught up with demand bringing prices down. On the other hand, a long-term trend toward higher energy prices or continued global shortage of concrete and steel could drive the real capital cost up significantly in the future.

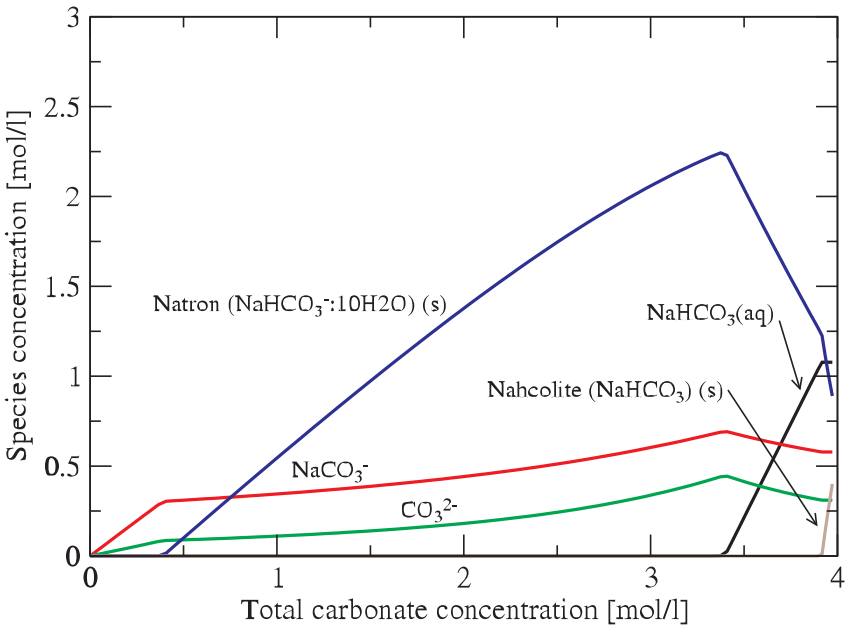

Figure 3.14: Equilibrium speciation of carbonates added to a 5M NaOH Solution at 20C. Formation of the solid Natron is predicted at CO_3^{2-} concentrations as low as 0.4 M, but not observed experimentally. Another solid, Nahcolite, is predicted at about 4M CO_3^{2-}, which is out of the range tested in the prototype.

3.5.6 Solids formation – scaling and clogging

At higher carbonate concentrations, solids formation may become a problem. To explore this issue, calculations with the chemical equilibria modeling software Chess were performed. The results are displayed in Figure 3.14. This analysis predicts that solids will form at carbonate concentrations larger than about 0.4 M in a solution with an initial NaOH concentration of 5 M. However, this was not observed in the prototype. Also, Apelblat and Manzurola (2003) report the solubility of sodium carbonate as more than 2 M at 20°C. The formation of this species may be kinetically limited. Still, if very high NaOH concentrations are used to combat water loss and the solution is recirculated to collect a high concentration of CO_3^{2-}, solid formation is likely.

If solids are present, they can be managed, although this may add complexity and capital cost. The main concern is if solid particles larger than the minimum free passage of a nozzle get into the spray supply line they will clog the nozzle. This problem is typically solved with inlet screens, which we used in the prototype to keep the line clear of debris and foreign particles pulled in at the air inlet.

The long term scaling (formation of a layer of solids adhered to a surface) of contactor walls, pipes, and equipment can probably be managed with a periodic water wash. Unlike calcium and magnesium compounds, the typical cause of scaling problems, sodium compounds are very soluble in water. Trace elements in air and process water may cause scaling not easily managed with a water wash, but this would be a similar problem as experienced in cooling towers and other industrial operations.

3.6 Conclusions from contactor analysis

The prototype demonstrates four key features of a potential NaOH spray-based contactor: (1) off-the-shelf single-fluid spray nozzles can produce a spray which efficiently absorbs CO_2 from ambient air (in terms of energy required for lifting the solution), (2) such nozzles can produce such a spray at pressures which are not prohibitive, (3) The pressure drop across a particle trap which controls entrainment of small drops from such a spray is not prohibitive (in terms of energy required for blowing air), and (4) materials compatibility and safety concerns in handling NaOH do not pose significant challenges to the design and operation of a contactor.

However, substantial uncertainties remain about the cost of a full-scale contactor. In particular, scaling up the mass transfer process observed in the prototype to meet the needs of global carbon mitigation scheme is a complex engineering challenge. Cost estimates for a full scale contactor scaled up from prototype observations came out high and highly variable. Overall, this seems to suggest that the current approach to contactor design and cost estimation is inadequate. The costs are high but do not appear to run up against any absolute limits to improvements. On the contrary, the results suggest drastic improvements can be made with both modest and radical redesign. On the other hand, the results suggest serious pitfalls that may easily render spray-based contactors infeasible. No absolute conclusion can be drawn, but the way is pointed for further investigation. Specifically, the feasibility and cost of spray-based contactors could be established with (1) a more exhaustive investigation of basic designs including counter-current flows, dual-fluid nozzles, and multi-stage spraying, (2) a better understanding of capital cost for potential contactors and how they scale with various design parameters, and (3) a precise model of, or empirical

data on, coalescence and breakup.

This chapter has explored spray-based contactors as opposed to packed or wetted-surface contactors. This approach was chosen because the simplicity and similarity to industrial cooling towers of the example design allowed for relatively simple cost estimates. Additionally, the success of large cooling towers offers a convincing proof-of-concept for large spray-based contactors. However, some of the experimental data suggest a filled or wetted-surface contactor would be a better approach. In the prototype, about half of the CO_2 absorption was by the wetted walls of the contactor. This is not surprising, considering that the inside wall area of the prototype was comparable to the calculated total surface area of suspended spray. But since the walls received only 10–20% of the liquid flow, CO_2 absorption there was much more efficient with respect to pumping energy. This is a general feature of packed or wetted-surface contactors: liquid pumping energy is lower than spray-based contactors and tends to be unimportant compared with fan energy. The geometry of the contactor, if extended to a full scale height, would yield a low air pressure drop. However, just as in the spray case, we would want a higher rate of CO_2 absorption per unit cross-sectional area in order to have low capital cost per unit CO_2. In a spray system, we proposed to do this by increasing liquid flowrate and thus spray density. In a wetted surface system, we would do this by adding walls or fill material. This, in turn, would impede the flow of air and increase the air pressure drop. The challenge in designing a wetted surface contactor is to maintain both a low air pressure drop and low capital cost per unit CO_2 absorbed. We recommend further investigation of wetted-surface contactors.

Chapter 4

Cost of air capture

Although a handful of researchers have described example air capture systems and two other researches have estimated the energy requirements of a system (those results are reproduced in Table 2.4), serious end-to-end cost estimates of a well-specified system are scarce. Such an estimate is attempted here. We limit ourselves to calcination-based systems. Clearly the cost of air capture with such a system cannot fall below the cost of calcination alone in a well-optimized, large scale industrial system. Thus in the Section 4.1, we calculate a lower bound cost based on industrial lime manufacturing.

We also expect that the cost of air capture at the time it is deployed, likely decades in the future, would not exceed the cost of a system that could be built today with well-known technology. And so in Section 4.2 we estimate the cost of an air capture system using current technology. We also make an estimate using some newer technology that is likely to be available in the near term, like an oxyfuel fluidized-bed calciner and pellet reactor for the caustization reaction.

4.1 Lower bound

Theoretically, the energy demand of the proposed system is dominated by the Calciner, where $CaCO_3$ is heated to release the captured CO_2, Reaction 3 in Table 2.1 (this dominance is borne out by cost estimates of the components of the total system). This reaction has a large thermodynamic energy requirement which must be overcome in even the most advanced calcining system, and significant capital requirements in terms of a large, high-temperature reactor which must be precisely tuned to maintain calcination efficiency.

The calcining operation in its simplest form is performed by the lime manufacturing industry at large scale and using long-established technology. In the production of high-calcium quicklime (CaO) in particular, crushed Calcite ($CaCO_3$) is heated in a kiln to form the product. Through the industry's long experience, it has developed a very efficient process, with energy requirements close to the thermodynamic limit ($\approx 85\%$). One can make arguments for how the costs of the other components of the air capture system can be greatly reduced by advances in design or economies of scale, but it is unlikely that a calcination-based air capture system could be less expensive than the current industrial calcining process, which has been optimized through decades of industrial experience.[1]

[1] A possibility which might be argued as a way to reduce the energy requirements below the requirement of Reaction 4 (and

The market price of high-calcium quicklime should reflect the industrial cost of calcination, inclusive of operation and capital recovery. One major adjustment is needed, however, because this price will include the cost of the raw material – crushed limestone. This is not needed in the air capture system, aside from a small amount of make-up lime, since the material is reused. Miller (2003) reports that the average price of this product in the United States is 78 \$/t-CO_2 (converted to CO_2 terms). Subtracting the average price of crushed limestone sold for lime manufacturing of 12 \$/t-CO_2 (Tepordei, 2002), this gives 65 \$/t-CO_2.

4.2 Cost of example system

In Section 2.3, we described an example air capture system making maximum use of existing, well-known industrial components, but costs for a contactor were not available. From Chapter 3, we now have a detailed cost analysis of one type of contactor. The analysis, unfortunately, gives a very large range of costs for spray-based contactors, some of which are very high. However, Chapter 3 also makes the case that there is no fundamental limitation to a low-cost contactor and that many engineering parameters have a large effect on the cost. No type of extremely costly contactor would be built. Considering that designers of an air capture system would select parameters to yield the lowest-cost contactor, we argue that one of the middle-assumption scenarios can serve as a reasonable upper bound on contactor cost. The "optimized" cost is also considered. We take the mid-level cost as 70 \$/t-CO_2, reflecting both the middle-assumptions no-coalescence scenario and the high-coalescence scenario with favorable assumptions. The optimized cost is 20 \$/t-CO_2.

Now that we have costs for the contactor, we can complete a total system estimate. Two systems are considered: the base system, and an improved system. The base system is just as described in Chapter 2: a spray tower, a conventional caustic recovery system, and an amine capture system. It is fired with gas, either natural, coal-derived, or bio-derived. We assume a price of thermal energy, p_{therm}, of 6 \$/GJ. Upstream carbon emissions from gas production are ignored. The CO_2 from fuel combustion is captured in the amine plant along with the calcined CO_2. In many ways this is a highly suboptimal system, but, as discussed in Chapter 2, it is the most valid way to make use of available cost estimates and industrial experience.

In the improved system we calculate the effects of moving to several more efficient components. The energy requirements for caustic recovery are matched to the results from Baciocchi et al. (2006) for the case of advanced dewatering technology. Additionally, an oxyfuel capture scheme is employed instead of amines, and the optimized contactor scenario is used. We don't have a strictly valid way of estimating capital costs for this new scheme. We will leave them the same as for the base case except for addition of capital and operating costs for an oxygen separation unit. Parameters for the unit are taken from Singh et al. (2003).

of industrial calcination) is to construct a system which captures the energy from Reaction 3 for heat or useful work, something which cannot be accomplished in quicklime manufacturing. This appears quite possible, but the captured energy would likely be used for several purposes which aren't part of quicklime manufacturing: solvent regeneration (or oxygen production in an oxyfuel system), dewatering and drying of $CaCO_3$ mud, and electricity for pumps and fans.

Component	Capital [$millions]	Capacity [t-CO_2/yr]	Electric[a] [GJ/t-CO_2]	Thermal[b] [GJ/t-CO_2]	Multiplier[c] ($R_{C/C_{net}}$)
Base System					
Contactor[d]	40	150,000	1.3	0	1.2
Caustic recovery[e]	60	290,000	0	11	1.2
Amine capture[f]	36	640,000	0.12	0	1.6
CO_2 compression[f]	10	640,000	0.4	0	1.6
Improved System					
Contactor[g]	0.5	10,000	0.6	0	1
Caustic recovery[h]	60	290,000	0	8	1
Oxygen separation[i]	180	6,600,000	0.4	0	1
CO_2 compression[f]	10	640,000	0.4	0	1.4

Table 4.1: Input parameters for system cost estimate.
[a] Electrical energy requirement per unit CO_2 processed in this component.
[b] Thermal energy requirement per unit CO_2 processed in this component.
[c] Tons of CO_2 processed in this component per net ton captured in total system.
[d] Scenario 2 contactor from Chapter 3.
[e] 1000 ton-CaO/day caustic recovery plant for a paper mill with industry rule-of-thumb parameters (Flanagan, 2004)
[f] Amine-based CO_2 scrubbing and compression system with parameters from the Integrated Environmental Control Model (Rao and Rubin, 2002).
[g] Scenario 6 contactor from Chapter 3.
[h] Same as in ([e]), with energy requirements reduced to match the improved system in Baciocchi et al. (2006).
[i] Cryogenic oxygen separation plant with 10,400 ton-O^2/day capacity; parameters from Singh et al. (2003)

A summary of the assumptions about capital cost and energy requirements used to calculate total costs is presented in Table 4.1. To reach the total cost per ton CO_2 we sum the cost of each component per net ton CO_2 captured by the system. We assume that capital and energy requirements scale linearly with plant capacity (a conservative assumption which ignores economies of scale) so only the unit cost matters. This allows us to estimate total costs using source data for components of different capacities. To adjust the unit costs to refer to net tons of CO_2 captured, we introduce a CO_2 multiplier, $R_{C/C_{net}}$, defined as the number of tons of CO_2 processed in the component per net ton captured in total system. In the base system, the amine capture plant is assumed to be only 90% efficient, so 10% of the captured carbon and 10% of the calciner fuel carbon is lost to the atmosphere during operation. Consequently, the contactor and caustic recovery plant must process about 18% "extra" CO_2 for each net ton captured ($R_{C/C_{net}} = 1.18$). The amine plant processes fuel carbon in addition to atmospheric carbon, giving $R_{C/C_{net}}$ of about 1.6. Since capture in an oxyfuel system is nearly 100% efficient, the multiplier is 1 for contacting and caustic recovery. In both systems, atmospheric and fuel CO_2 must be compressed, bringing $R_{C/C_{net}}$ for compression to 1.6 in the base system and 1.4 in the improved system. Immediately we can see that substantially more CO_2 must be sequestered than is captured from the air. We can view $(1 - R_{C/C_{net}})$ for compression as a sort of "carbon penalty" of air capture. 60%, in the base case, is the extra carbon that must be burned to capture and compress CO_2 from air. It depends on the fuel used for thermal energy and the method of electricity generation. If the calciner were fired with fuel oil, for instance, the penalty would be larger. If the electricity were generated by fossil fuel combustion with CCS, the penalty including electricity generation would be higher.

Even a carbon penalty above 100% does not invalidate air capture. Since the extra carbon is being sequestered, there is still net capture. However there is significant added burden on the fossil fuel supply with associated upstream and non-carbon impacts of fossil fuel use. Also, compared with point-source sequestration costs, the cost of sequestration (after the compression step) will be higher per unit of CO_2 captured, since the air capture system must also sequester carbon from the fuel. The cost of CO_2 transport and injection is not included in the figures given here, although it is expected to be comparatively small (McCoy and Rubin, 2005; IPCC, 2005).

We calculate the unit cost of each component in analogy with Equation 3.18 for the contactor cost where, again, Cap and $O\&M$ are the amortized capital and maintenance costs. We introduce $\dot{E}_{therm}$, the rate of thermal energy use by a component. We then sum across the components adjusting the unit costs by the carbon multiplier:

$$\frac{\text{total cost}}{\text{CO}_2\text{ captured}} = \sum_{components} \frac{Cap + O\&M + p_{elec}(\dot{E}_{elec}) + p_{therm}(\dot{E}_{therm})}{\text{capacity}} \times R_{C/C_{net}}$$

The same financial assumptions are made as for Section 3.4: 15% capital charge rate, 4% operation and maintenance, and carbon-neutral electricity for 19 \$/GJ.

The results of the cost estimation are given in Table 4.2. The total cost is 250 \$/t-CO_2 (900 \$/t-C) for the base case and 130 \$/t-CO_2 (500 \$/t-C) for the improved system. In the base system, capital and operating costs comprise half the total with energy cost as the other half. In the improved system energy becomes slightly more important, making up 60% of the total.

Component	Capital + O&M	Energy cost	Unit cost
Base System			
Contactor	60	30	90
Caustic recovery	50	80	120
Amine capture	20	4	20
CO_2 compression	5	10	20
Total			250
Improved system			
Contactor	10	10	20
Caustic recovery	40	50	90
Oxygen separation	5	8	10
CO_2 compression	4	10	20
Total			140

Table 4.2: Cost of example system and improved system by component. All costs in \$/t-CO_2.

	Base ΔE [GJ/t-CO_2]	Base $\Delta\$$ [\$/t-CO_2]	Impr. ΔE [GJ/t-CO_2]	Impr. $\Delta\$$ [\$/t-CO_2]
Switch to oxyfuel	-4	-54		
Packed tower with low capital cost[a]	-2	-62	+.06	-5
Fuel cost up to 8 \$/GJ		+24		+16
Fuel is stranded natural gas at 3 \$/GJ		-36		-24
Spray constant increased $\times 2$[b]	-2	-34	-2.5	-5
Capital charge rate is 12%		-19		-9
Economy of scale: capital cost$\times 0.5$		-61		-29

Table 4.3: Sensitivity of cost estimates to changes in assumptions. "ΔE" refers to change in energy requirement compared with results in Table 4.2 and "$\Delta\$$" refers to change in total cost. "Base" and "Impr." refer to the base and improved systems in Table 4.2.

[a] Using the packed tower energy requirement from Baciocchi et al. and assuming a per-ton capital cost half that of the optimized spray tower.

[b] k_{spray} becomes $6 \times 10^{-3} \frac{\text{m}}{\text{s}}$, which is only a 50% increase in the improved system, where it was already increased.

Table 4.3 shows how the estimated costs change with some reasonable changes in assumptions. We consider switching the base system to an oxyfuel mode, a hypothetical packed tower contactor with low capital cost, changes in fuel prices, changes in capital assumptions, and a spray constant that is closer to the theoretical expectation. All of these have significant effect on total cost. Note that since the system was constructed to be a defensible upper bound, most reasonable alternative assumptions tend to bring the cost down.

Chapter 5

Discussion

Overview

In this thesis, we have analyzed the potential performance and cost of technology for CO_2 capture from ambient air. We began by briefly describing the climate change problem and carbon management options and identified the unique features of air capture that could help address this problem. We argued that the fundamental physics of air capture is favorable and so a low-cost air capture technology can likely be developed. Still, many energy experts are skeptical that air capture is a realistic pursuit, and so proponents of air capture (including the author) have turned to existing technology to demonstrate the feasibility of the process, outlining several example systems. The most convincing of these examples uses an aqueous NaOH sorbent and calcination-based regeneration system. We developed this concept into a complete system, based on current industrial processes, for which the capital costs, operating costs, and energy requirements can be estimated. Because no ready analogy with cost information was available, we developed an example contactor in detail, and built a pilot-scale prototype. Using insight from this experimental work, we estimated the cost of a complete system.

5.1 Findings and implications

In the prototype contactor, we demonstrated the core features of a spray-based contactor, suggesting it is feasible in some form. We identified the design constraints on a successful full scale contactor. Small drops, low capital cost, and low spray density seem to be the routes to a low-cost contactor. On the other hand, coalescence of drops challenges the efficient operation of a contactor and high capital costs can drive up the overall cost substantially. Given current evidence, we can not say that contactor costs are small compared with the total air capture system; they may make up a large fraction of the total. However, there is enough room for improvement by redesign that we expect the contactor not to be the dominant cost. With energy requirements for packed towers and spray towers all tending toward the range of 0.5–1.5 GJ/t-CO_2, it seems the contactor would have an irreducible cost of 10–30 \$/t-CO_2 for the electricity.

Our analysis strongly suggests that a spray-based contactor is feasible in some form. Though there remains considerable uncertainty in this analysis, it is not dramatically different than that one obtains in the assessment of many future energy technologies. The cost estimate of 250 \$/t-CO_2, is very high. That

value is based on a system built with existing technology and no optimization for the purpose of being most defensible on feasibility terms. It is an inelegant and inefficient design by any other measure. To illustrate this point, we applied a few obvious improvements to get a more efficient, if less defensible, design. This reduces the total cost by 40%. We would expect that, with the decade or more of serious technological development that would likely precede the deployment of air capture, much more drastic improvements would be made.

But this analysis also suggests caution against exuberance for low-cost air capture. Even the improved system, at 140 \$/t-CO_2 is far more costly than most conventional mitigation options. CCS from power plants is more likely around 30 \$/t-CO_2, and CO_2 emissions credits on the European market are currently trading for less than that. Air capture still may be a significant savings over switching to hydrogen cars (Keith and Farrell, 2003), but perhaps not over cellulosic ethanol (Morrow et al., 2006) as a means of removing carbon from the transportation sector. Some early claims about the cost of calcination-based air capture were even lower than our lower-bound estimate of 60 \$/t-CO_2, which is partly why we included it. We don't expect air capture will ever be competitive with capture from point sources. If it were, our lower bound calculation suggests it would have to be based on a process other than calcination.

Yet the unique features of air capture – the potential for negative emissions and the backstop it provides on the cost of carbon mitigation – mean that even at a high price like 140 \$/t-CO_2, it is important for climate policy. Keith et al. (2006) demonstrated this with a long run global economic model, and the results make sense: air capture reduces the total social cost of mitigation. It guards against high carbon costs, and to some extent, against damage in extreme climate scenarios. By acting as a hedge against low probability, high-damage climate scenarios, it also shifts relatively more emissions reduction to the future. In a sense, fewer cautionary controls are needed on emissions when we know that air capture can be used to reduce atmospheric concentrations in case of emergency. As Parson (2006) points out, it is the classic "moral hazard" of insurance. But the danger is not so much that society, in maximizing total discounted social welfare, chooses to emit slightly more CO_2 in the near future, as that opponents of CO_2 regulation will use air capture as a political argument against action on the climate problem. If successful these arguments could lead to near term emissions far higher than optimum from a social welfare perspective.

This is indeed a point of which air capture proponents should be wary. And the results of this analysis should serve as a reminder that, while air capture is feasible, the cost could be very high. At 250 \$/t-CO_2, air capture is best viewed as an option for use in only the direst of climate emergencies and should be little comfort for present-day emitters.

In any case, this analysis supports the view that air capture is a serious option worthy of continued research. We have suggested many avenues for improving the design and finding lower cost versions of a NaOH-contactor-based air capture system. Additionally, we suggest that longer-term research into new recovery mechanisms and new sorbents – to break the thermodynamic constraints of calcination or hydroxide solutions, respectively – may yield the really dramatic cost reductions that would bring air capture into widespread use.

5.2 Lessons for assessing of future energy technologies

Assessing the future of energy technologies has long been an important task. But with climate change, greenhouse gas regulation, high oil prices, and strife in the Middle East, there is arguably more pressure than ever to radically change the energy system. The climate change issue in particular means that policy-makers are making decisions with century-scale consequences. Understanding the potential of particular nascent energy technologies can help direct R&D funding to effective projects and can inform various climate and energy forecasting models. Here we briefly consider the steps taken in this thesis – an assessment of a future energy technology – as they may apply to such assessments in general. We do not claim that this is an optimal strategy, but this is what we used.

We assume that we are trying to answer questions like "Could this technology be feasible at large scale?" and "Could it be made less expensive than alternatives?" in order to answer more immediate questions like "Should further research on this technology be funded?" and "How should the cost of this technology be represented in this forecasting model?" The general approach is one of bounding. We try to set an upper limit on the performance of the technology by considering the fundamentals of the process. If, because of some fundamental physical or economic constraint, the upper limit is not satisfactory, then the answer to the first two questions is "no", and we are done. If not, then we try to set a lower limit on performance (where low performance is bad) based on current knowledge and technology. If the lower limit is high enough, then we can answer "yes" and, though we may not be done, we have provided useful information to the decision maker. If neither answer is possible, then the door is open for further analysis and technology development.

Steps we applied:

1. **Calculate thermodynamic limits.** Assuming an ideal process is developed, what is the potential of this technology? There are many examples of technologies that, when mature, approach their thermodynamic limits of efficiency. The calciner in lime production is one. Were one assessing the feasibility of natural gas turbines for electricity today, one would calculate the Carnot efficiency and multiply by the fuel cost. One might conclude immediately that the technology is not worth pursuing. For air capture, we presented this calculation in Section 1.3

2. **Assess natural and practical limits to large-scale deployment.** Are rare materials necessary? Is there excessive land use? Are other scarce resources involved? Is siting limited or difficult? We answered some of these questions for air capture in Section 1.3. Also a survey of industrial sodium carbonate and calcite production was performed. Had we found that global deployment of air capture would use all of the world's sodium, we would have considered the example system impractical. These are the kinds of questions that would keep biodiesel from restaurant waste oils from being a key part of the national energy agenda.

3. **Look for industrial analogies – processes that are similar or work with similar principles.** Do these analogous systems work and why or why not? No technology appears in isolation. There are typically major components that have been previously engineered or have evolved from similar

technologies already operating. The status of those technologies is instructive. If we just knew the reactions in Table 2.1, we wouldn't necessarily be convinced that a system based on that chemistry is practical at a large scale. However, knowing it is practiced in the pulp and paper industry indicates it is. And the argument can go in the other direction: knowing that the kraft process is used in the pulp and paper industry and that they do not recover heat from the slaking reaction begs the question: why not? It may indicate that heat recovery from slaking would also not be possible in other caustic recovery applications, like air capture.

4. **Consider the state of the art.** How much of the system can be assembled with available components? How much does this cost? This is the hard part. Given that we have come this far, we may be able to make an estimate of performance based on known technology (or near-term technology assessed for other applications) that gives a meaningful answer to the decision maker's questions. Note that there is a fundamental asymmetry here: an estimate of performance that reaches "good enough" on some measure supports the feasibility of the technology because real future performance is likely to be better than a present-technology version. But an estimate of very low performance is not necessarily meaningful. It could be that the system used for analysis was poorly chosen or that a key supporting technology doesn't exist but may still be developed. Herzog (2003) falls into this trap when attempting a cost estimate for air capture. He uses a current-technology version of a contactor and finds the energy costs to be two orders of magnitude higher than what we found in this paper, from there concluding that the concept is not feasible. He uses an apparently-plausible industrial analogy, but fails to optimize it sufficiently for the new purpose. This is related to what was referred to in Section 3.4 as the "dual problem" of assessing a system and developing a version of the system for assessment. It is difficult to do the latter without actually engineering the system in a detailed way and attempting to optimize it. However this is not the goal of a technology assessment, so some middle ground must be found. Some degree of cleverness must be used to perform the analysis on a favorable version of the technology.

5. **What technological gains must be made for the technology to work?** If we have reached this step with no definitive answer, than we must start thinking in terms of likelihood rather than absolute bounds. The question is, "given the performance calculated in Step 4", or, "given the need for supporting technology *X*", "how much innovation is necessary for this technology to be useful?" And then consider the likelihood of that innovation. Note that supporting technologies may come from outside the field, as has all of the technology we have discussed for air capture. If uncertainty about some physical process, or about the structure of the market, or about some other fundamental factor is impeding the analysis, we might ask what answer, when the uncertainty is resolved, would kill the technology. And then suggest further research on those points.

6. **Consider the positive and negative effects this technology would have, if realized.** Is support of the technology worthwhile? If it became technically and economically viable, what direct and indirect consequences would arise? This step is probably under-emphasized in technology assessment. When considering the history of technology and unintended consequences, its importance is obvious.

Bibliography

Adams, P. J. and Seinfeld, J. H. (2002). Predicting global aerosol size distributions in general circulation models. *Journal Of Geophysical Research-Atmospheres*, 107(D19):4370.

Adams, T. N. (1989). *Lime Reburning*, volume 1. 5 Alkaline Pulping of *Pulp and Paper Manufacture*, chapter XXII, pages 590–608. Joint Textbook Committee of the Paper Industry.

Allspray (2002). Product Specifications. Technical report, Allspray Inc.

Amistico (2006). Mesh & Vane Mist Eliminators (product brochure). Technical report, Amistico Separation Products, Inc, Alvin, Texas.

Apelblat, A. and Manzurola, E. (2003). Solubilities and vapour pressures of saturated aqueous solutions of sodium tetraborate, sodium carbonate, and magnesium sulfate and freezing-temperature lowerings of sodium tetraborate and sodium carbonate solutions. *Journal of Chemical Thermodynamics*, 35:221–238.

Baciocchi, R., Storti, G., and Mazzotti, M. (2006). Process design and energy requirements for the capture of carbon dioxide from air. *Chemical Engineering and Processing*. Available online, DOI:10.1016/j.cep.2006.03.015.

Bird, R., Stewart, W. E., and Lightfoot, E. N. (1960). *Transport Phenomena*. John Wiley and Sons, New York.

Buesseler, K. O., Andrews, J. E., Pike, S. M., and Charette, M. A. (2004). The effects of iron fertilization on carbon sequestration in the southern ocean. *Science*, 304:414–417.

Buesseler, K. O. and Boyd, P. W. (2003). Will ocean fertilization work? *Science*, 300:67–68.

Carroll, A. L., Taylor, S. W., Régnière, J., and Safranyik, L. (2003). Effects of climate change on range expansion by the mountain pine beetle in British Columbia. In Shore, T., Brooks, J., and Stone, J., editors, *Mountain Pine Beetle Symposium: Challenges and Solutions.*, Kelowna, British Columbia. Natural Resources Canada.

Danckwertz, P. V. (1970). *Gas-Liquid Reactions*. McGraw-Hill, New York.

Davis, W. T., editor (2000). *Air Pollution Engineering Manual*. John Wiley & Sons, New York.

DOE (2005). Properties of Fuels. Technical report, Alternative Fuels Data Center, United States Department of Energy. http://www.eere.energy.gov/afdc/pdfs/fueltable.pdf.

Dubey, M. K., Ziock, H., Rueff, G., Elliott, S., and Smith, W. S. (2002). Extraction of carbon dioxide from the atmosphere through engineered chemical sinkage. *ACS – Division of Fuel Chemistry Reprints*, 47(1):81–84.

Emanuel, K. (2005). Increasing destructiveness of tropical cyclones over the past 30 years. *Nature*, 436(7051):686–688.

EPA (2002). *Economic and Engineering Analyses of the Proposed §316(b) New Facility Rule.*, chapter Appendix B: Technology Cost Curves. Number EPA 821-R-02-003. United States Environmental Protection Agency.

Fair, J. R., Steinmeyer, D. E., Penney, W. R., and Crocker, B. B. (2001). *Gas absorption and gas-liquid system design*, chapter 14, pages 14–1–14–98. Perry's Chemical Engineering Handbook. McGraw Hill.

Fernandez, B. M., Simons, S., Hills, C., and Carey, P. (2004). A review of accelerated carbonation technology in the treatment of cement-based materials and sequestration of CO2. *Journal Of Hazardous Materials*, 112(3):193–205.

Flanagan, P. (2004). Email communication. Engineer with Groupe Laperrière and Verreault.

Gan, J. (2004). Risk and damage of southern pine beetle outbreaks under global climate change. *Forest Ecology and Management*, 191(1-3):61–71.

Greenwood, K. and Pearce, M. (1953). The removal of Carbon Dioxide from Atmospheric Air by Scrubbing with Caustic Soda in Packed Towers. *Transactions of the Institution of Chemical Engineers*, 31:201–207.

Hansen, J., Ruedy, R., Sato, M., and Lo, K. (2006). Global Temperature Trends: 2005 Summation. Technical report, NASA Goddard Institue for Space Studies.

Herzog, H. (2003). Assessing the Feasibility of Capturing CO2 from the Air. Technical report, MIT Laboratory for Energy and the Environment.

Hoftyzer, P. and van Krevelen, D. (1954). Applicability of the Results of Small-Scale Experiments to the Design of Technical Apparatus for Gas Absorption. *Transactions of the Institution of Chemical Engineers, Supplement (Proceedings of the Symposium on Gas Absorption*, 32:S60–S67.

Iizuka, A., Fujii, M., Yamasaki, A., and Yanagisawa, Y. (2002). A novel reduction process of CO2 fixation by waste concrete treatment. *Kagaku Kogaku Ronbunshu*, 28(5):587–592.

IPCC (2000a). Emissions Scenarios. Special report, Intergovernmental Panel on Climate Change.

IPCC (2000b). Land Use, Land-use Change, and Forestry. Technical report, Intergovernmental Panal on Climate Change.

IPCC (2001). Climate Change 2001: Mitigation. Summary for policy makers, Intergovernmental Panel on Climate Change.

IPCC (2005). Carbon Dioxide Capture and Storage. Special report, Intergovernmental Panel on Climate Change.

Johnston, N., Blake, D., Rowland, F., Elliott, S., Lackner, K., Ziock, H., Dubey, M., Hanson, H., and Barr, S. (2003). Chemical transport modeling of potential atmospheric CO2 sinks. *Energy Conversion and Management*, 44(5):681–689.

Keith, D., Ha-Duong, M., and Stolaroff, J. K. (2006). Climate strategy with CO_2 capture from the air. *Climatic Change*, 74(1-3):17–45. to appear.

Keith, D. W. and Farrell, A. E. (2003). Rethinking hydrogen cars. *Science*, 301(5631):315–316.

Kheshgi, H. S. (1995). Sequestering atmospheric carbon dioxide by increasing ocean alkalinity. *Energy*, 20(9):915–922.

Lackner, K. S., Grimes, P., and Ziock, H. J. (2001). Capturing carbon dioxide from air. In *Proceedings of the First National Conference on Carbon Sequestration*, Washington, DC.

Lefebvre, A. H. (1989). *Atomization and Sprays*. Taylor and Francis.

Lindberg, D., Perander, L., Backman, R., Hupa, M., Kochesfahani, S., and Rickards, H. (2005). Borate autocausticizing equilibria in recovery boiler smelt. *Nordic Pulp and Paper Research Journal*, 20(2):232–237.

Lindberg, D. K. and Backman, R. V. (2004). Effect of Temperature and Boron Contents on the Autocausticizing Reactions in Sodium Carbonate/Borate Mixtures. *Industrial Engineering Chemical Research*, 8:6285–6291.

Long, A. B. (1974). Solutions to the droplet collection equation for polynomial kernels. *Journal of Atmospheric Sciences*, 31:1040–1052.

Marland, G., Boden, T., and Andres., R. J. (2006). Global, Regional, and National CO2 Emissions. Technical report, Carbon Dioxide Information Analysis Center, Oak Ridge National Laboratory, Oak Ridge, TN, USA.

McCoy, S. T. and Rubin, E. S. (2005). Models of CO2 Transport and Storage Costs and Their Importance in CCS Cost Estimates. In *FOURTH ANNUAL CONFERENCE ON CARBON CAPTURE AND SEQUESTRATION, May 2-5*.

Miller, M. M. (2003). Lime. In *Minerals Yearbook 2003*, pages 40.1–40.15. United States Geological Survey.

Morrow, W., Griffin, W., and Matthews, H. (2006). Modeling Switchgrass Derived Cellulosic Ethanol Distribution in the United States. *Environmental Science & Technology*, 40:2877–2886.

Mykyntyn, H. (2006). Personal correspondence. Mechanical Estimator for Bechtel Power Corp.

NRC (2002). Effectiveness and Impact of Corporate Average Fuel Economy (CAFE) Standards. Technical report, Transporation Research Board, National Research Council, Washington, DC.

Parson, E. (2006). Reflections on Air Capture: the political economy of active intervention in the global environment. *Climatic Change*, 74(1 - 3):5–15.

Pruppacher, H. and Klett, J. (1997). *Microphysics of Clouds and Precipitation*. Kluwer, Boston.

Rao, A. B. (2004). Personal communication.

Rao, A. B. and Rubin, E. S. (2002). A technical, economic, and environmental assessment of amine-based CO_2 capture technology for power plant greenhouse gas control. *Environmental Science and Technology*, 36(20):4467–4475.

Rhodes, J. S. and Keith, D. W. (2005). Engineering economic analysis of biomass IGCC with carbon capture and storage. *Biomass and Bioenergy*, 29:440–450.

Richards, T., Nohlgren, I., Wamqvist, B., and Theliander, H. (2002). Mass and energy balances for a conventional recovery cycle and for a recovery cycle using borates or titanates. *Nordic Pulp & Paper Research Journal*, 17(3):213–221.

Seinfeld, J. H. and Pandis, S. N. (1998). *Atmospheric Chemistry and Physics*. John Wiley and Sons, New York.

Singh, D., Croiset, E., Douglas, P., and Douglas, M. (2003). Techno-economic study of CO2 capture from an existing coal- red power plant: MEA scrubbing vs. O2/CO2 recycle combustion. *Energy Conversion and Management*, 44:3070–3091.

Spielbauer, T. M. (1992). *The Stability and Disintegration of Radially Thinning Liquid Sheets*. PhD thesis, The Institute of Paper Science and Technology.

Stolaroff, J. K., Lowry, G. V., and Keith, D. W. (2005). Using CaO- and MgO-rich industrial waste streams for carbon sequestration. *Energy Conversion and Management*, 46(5):687–699.

Sugden, A. M. (2005). Bleaching in Hot Water. *Science*, 309(5742):1791a–.

Tepordei, V. V. (2002). Stone, Crushed. In *Minerals Yearbook 2002*, pages 72.1–72.36. United States Geological Survey.

Thomas, C. D., Cameron, A., Green, R. E., Bakkenes, M., Beaumont, L. J., Collingham, Y. C., Erasmus, B. F. N., de Siqueira, M. F., Grainger, A., Hannah, L., Hughes, L., Huntley, B., van Jaarsveld, A. S., Midgley, G. F., Miles, L., Ortega-Huerta, M. A., Townsend Peterson, A., Phillips, O. L., and Williams, S. E. (2004). Extinction risk from climate change. *Nature*, 427(6970):145–148.

Tzivion, S., Feingold, G., and Levin, Z. (1987). An efficient numerical solution to the stochastic collection equation. *Journal of the Atmospheric Sciences*, 44(21):3139–3149.

Tzivion, S., Feingold, G., and Levin, Z. (1989). The evolution of raindrop spectra. Part II: Collisional collection/breakup and evaporation in a rainshaft. *Journal of the Atmospheric Sciences*, 46(21):3312–3327.

Weast, R. C., editor (2003). *CRC Handbook of Chemistry and Physics*. CRC Press, Boca Raton, FL.

Webster, P. J., Holland, G. J., Curry, J. A., and Chang, H.-R. (2005). Changes in tropical cyclone number, duration, and intensity in a warming environment. *Science*, 309(5742):1844–1846.

Westerling, A. L., Hidalgo, H. G., Cayan, D. R., and Swetnam, T. W. (2006). Warming and earlier spring increases western U.S. forest wildfire activity. *Science*, page 1128834.

Zeman, F. S. (2006). *Air Extraction: The Feasibility of Absorbing CO_2 from the Atmosphere*. PhD thesis, Columbia University.

Zeman, F. S. and Lackner, K. S. (2004). Capturing Carbon Dioxide Directly from the Atmosphere. *World Resources Review*, 16(157–171):62–68.

Appendix A

List of Symbols

A Cross sectional area of the contactor [m^2]

C Concentration of CO_2 in air

C_0 Concentration of CO_2 in solution at the drop surface [mol/m^3]

Cap Amortized capital cost [US$/yr]

C_{in} Concentration of CO_2 in air at the contactor inlet

C_{out} Concentration of CO_2 in air at the contactor outlet

$[CO_3^{2-}]$ Concentration of carbonate ions in solution [mol/L]

c_l Heat capacity of solution $\left[\frac{J}{K \cdot kg}\right]$

$c_{p,air}$ Heat capacity of air at constant pressure $\left[\frac{J}{K \cdot kg}\right]$

D Drop diameter [m]

D_g Diffusivity of CO_2 in gas (air)

D_l Diffusivity of CO_2 in liquid (water)

$\dot{E}$ Rate of energy use by the contactor

$\dot{E}_{elec}$ Rate of electrical energy use by a component

$\dot{E}_{fan}$ Rate of energy use to run the fan in the contactor

$\dot{E}_{lifting}$ Rate of energy use to lift solution in the contactor

$\dot{E}_{nozzle}$ Rate of energy use to overcome the pressure drop across the nozzles in the contactor

$\dot{E}_{therm}$ Rate of thermal energy use by a component

F Liquid flow rate in the contactor [m^3/s]

g Gravitational constant [m^2/s]

H Height of the contactor [m]

h Vertical coordinate (height) within contactor [m]

K Collection kernel for colliding drops [cm^3/s]

k_g Gas-side mass transfer coefficient

k_{spray} Empirical spray constant: CO_2 absorption per unit surface area per unit time $\left[\frac{mol}{m^2 s}\right]$

k_v Empirical constant for calculating the terminal velocity of a drop

$\dot{M}$ Rate of CO_2 absorption by the contactor

$O\&M$ Operating costs excluding energy [US$/yr]

$[OH^-]$ Concentration of hydroxide ions in solution [mol/L]

p_{elec} Price of carbon-neutral electricity [US$/GJ]

p_{therm} Price of thermal energy (natural gas equivalent) [US$/GJ]

Q Rate of CO_2 absorption by a drop [mol/s]

R Gas constant

$R_{C/C_{net}}$ Ratio of carbon processed by a component to net carbon captured by the total system []

R_{H_2O/CO_2} Ratio of water lost by evaporation to CO_2 absorbed by NaOH solution, molar basis []

RH Relative humidity of air

S Average surface area of spray per unit bulk volume in the contactor $\left[\frac{m^2\ surface}{m^2\ contactor}\right]$

t time since a drop leaves the nozzle or since a parcel of air entered the contactor [s]

T Temperature of air

V Volume of the contactor [m^3]

V_d Volume of a drop

ΔC_{H_2O} Difference between inlet and outlet water vapor concentration in the contactor

ΔG Change in free energy

ΔP_{air} Pressure drop in air between the inlet and outlet of the contactor

ΔP_{nozzle} Pressure drop across the nozzle in the contactor

ε_{fan} Mechanical efficiency of fans in the contactor [ ]

ε_{pump} Mechanical efficiency of pumps in the contactor [ ]

ρ_a Density of air [kg/m^3]

ρ_l Bulk density of liquid spray suspended in contactor [kg/m^3]

ρ_{soln} Density of contactor solution [kg/m^3]

ρ_{spray} Volume of suspended spray per unit volume of contactor $\left[\frac{m^3}{m^3}\right]$

σ Geometric standard deviation of the spray drop size distribution

τ Residence time of a drop, or average residence time of the spray in the contactor [s]

τ_{air} Average residence time of air in the contactor [s]

Appendix B

Experimental details and procedure

This section documents many experimental decisions and design details related to the construction and testing of the prototype contactor. It is meant to allow future interested researchers to reproduce or critique the techniques.

B.1 Physical apparatus

B.1.1 Basic size and structural design

In setting out to build a prototype contactor based on NaOH drops, the nature of drop generation was central to determining the overall form. Though there are various means of generating small drops at very small flow rates (electrostatic augmentation of flow from an orifice, sonic disruption of falling streams or sheets), we reasoned that any full-scale contactor would rely on spray nozzles. To realistically simulate the conditions in the full-scale system, the prototype would need to accommodate a spray nozzle. Additionally, since the prototype is meant to simulate the average conditions in a very wide, very tall tower, edge effects should be minimized. That is, distortions to the rate of mass transfer caused by liquid on the walls, transient spray characteristics (escape velocity and angle from the nozzle), and the collection mechanism at the bottom should be minimized. Without significant edge effects, the prototype results can more easily be compared with theory and more justifiably scaled up to a full size contactor.

So the use of a spray nozzle introduces several primary constraints on the structural form of the prototype: (1) The tower should be wide enough that most of the spray falls through without hitting the walls, (2) the length of the fall should be long enough that initial transient conditions as fluid leaves the nozzle (escape angle and velocity) are relatively unimportant compared with time drops spend falling at terminal velocity in a uniform pattern, and (3) since some spray inevitably hits the walls, this portion of the flow should be measurable for analytical reasons.

Though nozzles with narrow (15-30°) spray patterns are available, wider patterns are much more common, starting with 60° patterns (both of the nozzles ultimately used for data collection were 60° nozzles). All told, we seemed to require a tower diameter of at least 1 m and fall distance of several meters. Additionally, smooth, straight walls would aid measurement of flow hitting the walls.

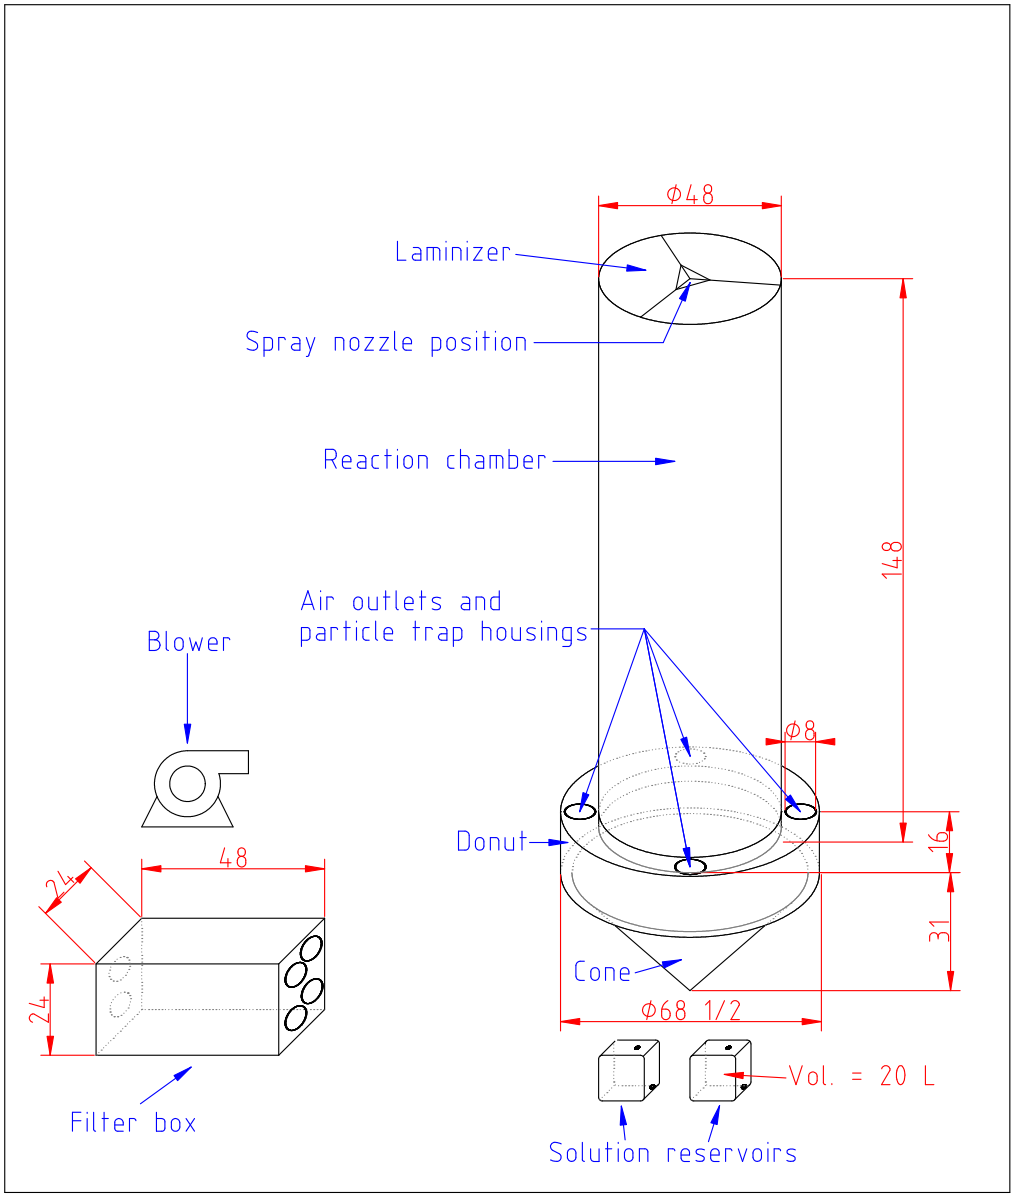

Figure B.1: Dimensions, layout, and labeling of major components of the final prototype design. Dimensions given in inches except where noted.

Figure B.2: Photograph of completed prototype structure.

Figure B.3: Lining the inside of the Sonotubes with PVC sheets.

The completed prototype (depicted in figures B.1 and B.2) includes a long, cylindrical "reaction chamber" capped at the top by a permeable fabric sheet (the "laminizer") and at the bottom by a short, larger-diameter cylinder (the "donut") and conical spray collection assembly (the "cone"). Liquid enters through a spray nozzle mounted just under the laminizer and is collected to a reservoir from the bottom of the cone. Air enters though the laminizer and leaves through four ducts attached to the top of the donut. The ducts lead to a filter box which is in turn attached to a blower. The entire unit is relatively well-sealed against air and liquid escaping and the blower keeps it all at negative pressure, further insuring against fugitive emissions.

For the reaction chamber we considered various wood-frame assemblies with plastic film or plastic sheet skin and prefabricated plastic tanks, but ultimately chose to construct the main body out of Sonotubes, heavy-duty cardboard tubes used as forms for pouring concrete columns. The Sonotubes provided stiff, precisely cylindrical walls that would stand up to negative pressure and support the top-mounted spray apparatus structurally.

The largest Sonotubes available had a 4 ft inner diameter and 4 ft length. Three lengths were fastened together with a fiberglass and epoxy wrap to produce a 12 ft total height of the reaction chamber. The Sonotubes were lined with 1/8 in-thick PVC sheets which were affixed to the waxed cardboard interior surface with epoxy (see Figure B.3). Silicon sealant was used to fill the seams. Once the Sonotubes were

Figure B.4: Two cranes lift the reaction chamber up and then orient it vertically to be dropped on to the support structure. The cone, resting on the ground, is also visible.

Figure B.5: Tower support structure

fastened together and lined, the reaction chamber was lifted by crane and placed onto the frame, a sparse truss-like structure of 1x3 in (nominal) pine boards (see Figure B.5). The donut, also suspended from the frame, is plywood and pine-board frame with 1/8 in PVC interior walls. The cone was constructed of two custom-cut PVC sheets folded into shape and fixed together with stainless steel bolts and sealant. The very tip of the cone is a standard polyethylene funnel fitted with a flexible 1 in diameter tube by hose clamp.

The bottom of the reaction chamber is positioned 4 in below the top of the donut so that air leaving the reaction chamber is forced to make a sharp U-turn and travel upwards a short distance, shedding most of the entrained drops, before entering the particle trap. The cone is set at a 45° angle with respect to the horizontal to minimize splashing (and creation of fine mist which may bias the rate of CO_2 uptake) at the bottom of the reaction chamber. Also, since the flow lines of process air pass into the donut above the cone, we reason that edge effects as the spray hits the bottom will have a negligible impact on mass transfer.

A lip on the bottom of the inside wall of the reaction chamber collects solution running down the walls and channels it to a separate return-flow tube that exits through a hold in the cone. The lip is a 1 in diameter flexible tube cut in half to form an open-top channel, and slightly sloped so that collected fluid drains to a single dedicated return-flow tube attached to the bottom of the channel. The separate return flow for fluid hitting the walls allows for measurement of the relative fraction of spray hitting the walls as opposed to remaining as drops through the reaction chamber.

B.1.2 Materials compatibility

Since the working solution would be strong caustic (up to 20 wt%, pH 14.7), we required materials that would stand up to caustic for several cumulative days of running time. Fortunately, most plastics have excellent caustic compatibility, including the most common and inexpensive varieties, like polyvinyl chloride (PVC) and polyethylene (PE). Most of the wetted surfaces in the prototype were constructed of PVC. Portable PE gasoline cans served as convenient solution reservoirs and storage containers. Most tubing was PE or nylon, and fittings were stainless steel Swagelok. The wire mesh particle trap was constructed of stainless steel wool and the spray nozzles and some miscellaneous fasteners were also stainless steel. A chemical-resistant pump head and an off-the-shelf silicon-based sealant completes the list of wetted materials (the sealant was not explicitly resistant to caustic but it was tested and did not visually degrade after several days of submersion).

B.1.3 Air handling and air safety

As in a full-scale contactor, the prototype has a forced-air system which is meant to move air uniformly through the tower co-current with the spray and employs a particle trap to keep a significant fraction of the working solution (in the form of very fine droplets) from leaving the system with the outlet air. See Figure B.6 for a diagram. The prototype had the additional constraint that experimenters working and breathing close to the prototype should not be exposed to hazardous levels of caustic particles. OSHA standards for airborne concentration of caustic solution were used as a guide.

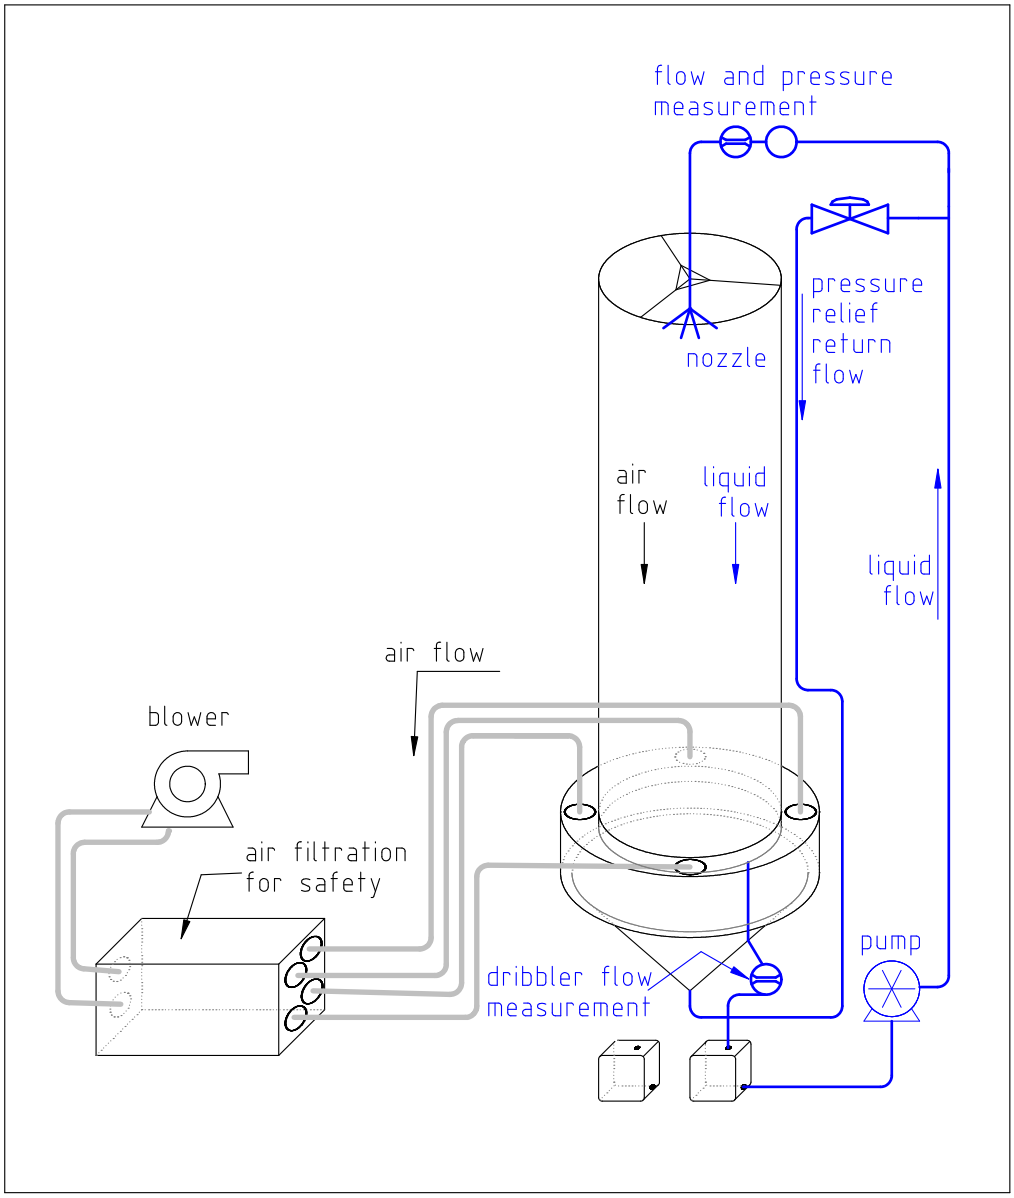

Figure B.6: Flow diagram of liquid and air systems in the final prototype design.

The prototype has three particle filters. Process air leaves the main body of the prototype through four 4x8 in rectangular openings in the top of the donut, spaced evenly around the perimeter. The primary particle trap is housed in each opening. Most of the volume of entrained liquid is caught here and drips down to the cone to join the bulk of the return flow. This filter is considered analogous in function to the particle trap that would be required in a full-scale system.

Past the primary particle trap, air enters 8 in diameter flexible ducts attached to each opening and then a custom-built filter box. The filter box houses two additional filters for safety reasons. The first is a standard fabric home furnace filter which acts as an inexpensive pre-filter. The second is a micro-glass media, mini-pleat duct filter rated for $>$ 95% removal of $_{PM_{2.5}}$ (Filter Group catalog # 40102). Air leaving the filter box travels through a pair of ducts to the blower, which vents to the room. Considering the efficiencies of the particle traps and a conservative estimate for the concentration of entrained fine particles, air leaving the blower should have caustic particle content well below (by perhaps an order of magnitude) OSHA standards.

B.1.4 Liquid handling

The design goal of the liquid system provides the necessary flowrate of NaOH solution to the nozzle at sufficient and adjustable pressure, allowing measurement of the flowrate and pressure at the nozzle. Also, since CO_2 absorption by the portion of fluid hitting the walls of the reaction chamber is not counted in most calculations, the flow on the walls should be measurable so that it can be subtracted out.

A schematic of the liquid system is shown in Figure B.6. For most trials, the liquid handling system was arranged as shown in a simple loop with a reservoir at the bottom, so that the working solution is recirculated. Reservoir residence times were 3–15 minutes. Most of the system is comprised of 3/8 in nylon tubing with stainless steel Swagelok valves and fittings. Additional branches of the system (not shown) allowed for in-line sample collection, rinsing, and solution transfer between reservoirs.

A pump was required to provide up to 6 L/min of flow at a total head of up to about 650 kPa (lifting 6 m, a target maximum nozzle pressure of 550 kPa, and frictional losses). A centrifugal pump was first tried. Though it was rated for far higher flows, it could not reach pressures at the nozzle above 300 kPa. A rotary-vane pump was then used and achieved higher pressures. However, pump performance was not consistent, and with the high-flow nozzle, the highest pressure reliably achieved was 380 kPa. This limited the range of pressures tested on the high-flow nozzle, although the pump was rated for more than 800 kPa. Air entrained in solution seemed to cause vibration and reduced performance. Measures such as arranging the return flow line to minimize splashing in the reservoir improved this.

Eight different nozzles were tested for use in the prototype. Nozzles were visually observed spraying water from a height of 3 m, and seven of the eight were also characterized with a Malvern laser diffraction spray analyzer. The two with the smallest mean drop size were selected for use in experiments. Both nozzles have full-cone spray patterns.

B.1.5 Measurement

The primary goal of the experiments was to calculate pumping energy, fan energy, and water loss per unit CO_2 captured. Accordingly, the quantities we meant to measure were:

1. CO_2 uptake
2. liquid flow rate
3. liquid line pressure
4. air flow rate
5. air pressure drops
6. inlet temperature
7. inlet humidity
8. outlet temperature
9. outlet humidity.

CO_2 uptake was primarily measured by real-time monitoring of the CO_2 concentration in air, alternately at the inlet and outlet of the reaction chamber. Measurement points are shown in Figure B.7. Air was pumped continuously from the end of a 1/4 in PE tube inserted at the measuring point through a LiCOR infrared CO_2 analyzer. A computer connected to the LiCOR recorded the CO_2 reading in ppm every second. This yields a ($\sim$ 3 second delayed) time series of CO_2 concentration at the sampling point. Sampling background at the tower inlet was straightforward. Some concerns, however, were raised about whether outlet concentration at any given sample point is representative of the average rate of CO_2 capture. Hence, six different points at the outlet were tried. Sampling directly in the reaction chamber proved problematic as drops would get pulled into the sampling tube and bias the measurement. Sampling in the donut appeared to be the most reliable approach. Agreement between the two sampling points and among different depths of tube insertion (5, 10, and 20 cm) was good (the differences were not distinguishable from noise). Agreement between the donut and duct measurements was also good. Measurements from the box and blower tended to run higher (indicating less CO_2 absorption) than donut measurements in some trials. This is probably explained by dilution of the process air with room air through leaks in the box and duct connections. Overall, the data used in final calculations appear to be valid representations of average outlet concentration. As a double-check on the LiCOR readings, the quantity of CO_2 absorbed into solution for some trials was measured also with a Total Organic Carbon analyzer so that results could be compared.

Total liquid flow rate was measured with an inline digital turbine flowmeter for some trials. When this was not available, total flow was estimated from manufacture specifications for the nozzle at the given pressure. Flow of the active spray was taken to be the difference between total flow and dribbler flow. Dribbler flow was measured for some trials by diverting the flow to a 2 L graduated cylinder for

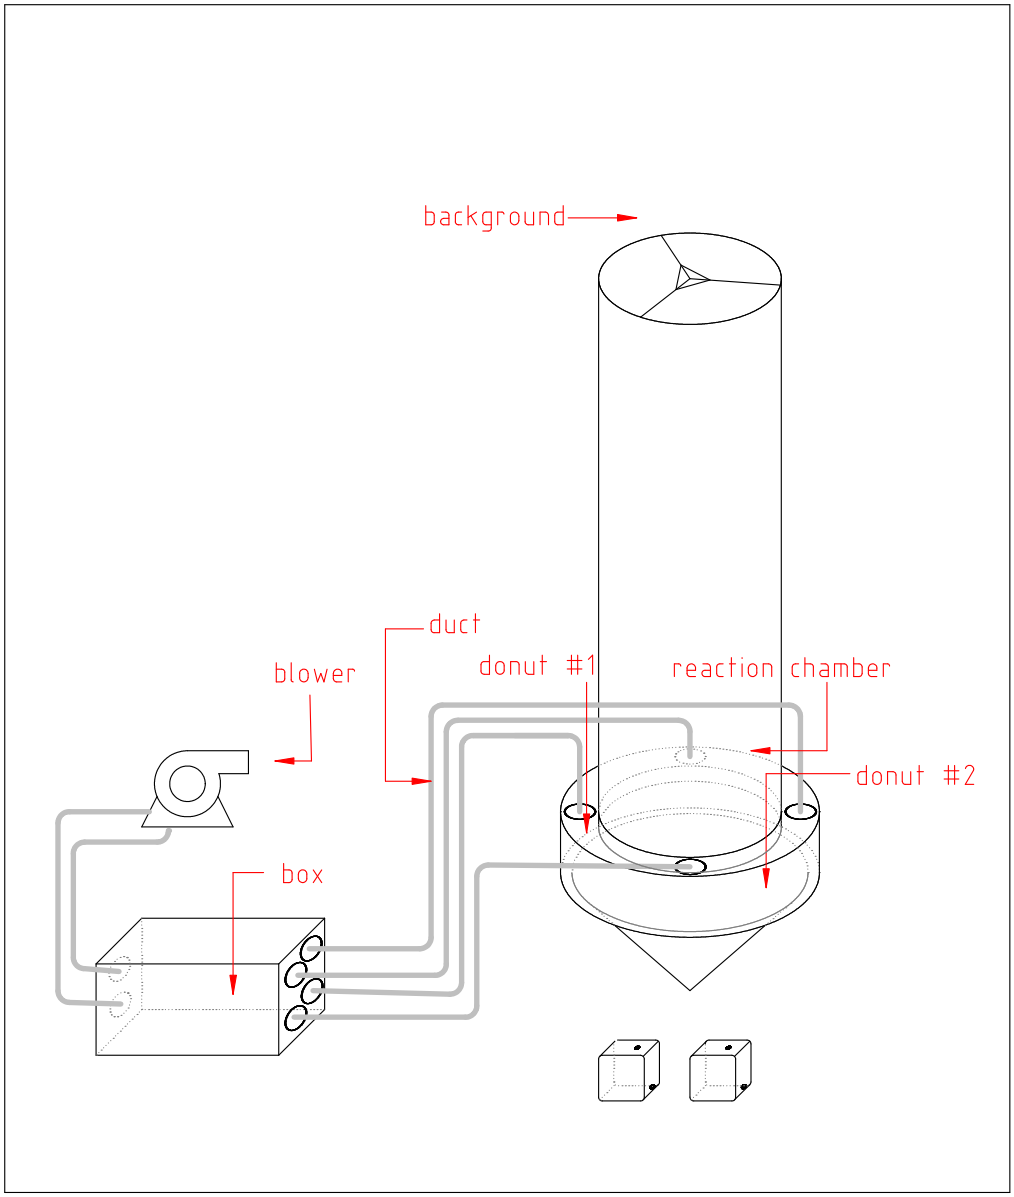

Figure B.7: Sampling points of CO_2 concentration in air. Arrows indicate position of the end of the sampling tube. Multiple locations for outlet concentration were tested and compared.

a measured time. Dribbler flow for other trials with the same nozzle was extrapolated from these data. Liquid line pressure was read from an inline pressure gauge.

Air flow velocity was measured with a hot-wire anemometer inserted into the ducts between the donut and the filter box. Velocity measurements were made at 2-5 cross-sectional depths in each of the 4 ducts. The values were stable over time and across spray conditions but varied by up to twenty percent among ducts and depths. Using the arithmetic average velocity and cross-sectional area of the ducts, the volumetric flow rate was calculated.

The air pressure drop across the tower was measured with a simple water-in-tube manometer open on one side to the air and on the other side inserted in a duct just after the primary particle trap. Thus the measured value includes pressure drop across the particle trap, frictional losses through the tower, and the (small) pressure drop across the laminizer.

Inlet and outlet temperature and relative humidity were measured with a handheld digital temperature and relative humidity meter.

B.2 Experimental Procedure

In all, we ran 12 trials with NaOH solution spraying and CO_2 monitoring. Variations among the trials included:

1. Spray nozzle: either high-flow or low-flow nozzle.
2. Solution concentration: 0.33, 1.33, or 5 M.
3. Liquid pressure at the nozzle: 550 kPa for the low-flow nozzle, 140, 240, 340, or 380 kPa for the high-flow nozzle.
4. Spray mode: continuous spraying or periodic switching (spray on, spray off).
5. Refinements of the apparatus and measurement techniques.

The first 5 trials were run in a continuous mode, where first the air flow was turned on, and then the spray was turned on and left running for 10–30 minutes. The outlet CO_2 appeared to reach steady state after 5–15 minutes. Once at steady state, inlet and outlet temperature and humidity were measured. The solution was recirculated to allow long run times with a manageable volume of solution. Periodic samples of the solution were taken in order to measure the carbonate concentration in liquid at various times during the course of the trial. The continuous mode allows comparison of the CO_2 absorbed as indicated by the outlet CO_2 concentration with the CO_2 measured in liquid samples. It does not, however, allow separation of CO_2 absorbed by spray from CO_2 absorbed on wetted walls of the reaction chamber or other surfaces.

Trials 6–12 were run in a switching mode, which allowed precise isolation of absorption by spray. After an initial period with the spray on to reach steady state, we would turn the spray of for 1-3 minutes then back on for 1-3 minutes. Typically we would run several such cycles and then change some conditions, such as the CO_2 measurement point or nozzle pressure, and then run another set of cycles.

In order to calculate CO_2 absorbed, outlet CO_2 concentration must be subtracted from background. This posed some challenges because background CO_2 sometimes varied dramatically, with spikes and drift on a time scale of 10 seconds to many minutes. These variations were often similar in magnitude to the absorption signal (4–50 ppm). Much of this appeared to be due to local variation from the ventilation system and from experimenters breathing near the inlet. We reasoned that connecting the inlet to the outside air would stabilize the CO_2 background. In fact outside concentration was at least as variable as inside concentration, perhaps due to signals from cars, plants, and drift from diurnal variation. The most stable background was achieved in trials 8–12 when the inlet was placed in a corner of the room and all experimenters stayed at the base of the tower, separated by about 3 m in height and 5 m in lateral distance from the inlet. Repeated measurements of the same conditions over time and across trials were largely consistent, so we were generally able to overcome background variation.

Variation of absorption by NaOH concentration was measured by running separate trials, changing the working solution but holding other conditions constant. The main confound to this approach is that different concentrations have different densities and viscocities, so the drop distribution was different among the trials. The measured differences in CO_2 absorption are thus due to both changes in the reaction rate of CO_2 hydrolysis and to differing drop size.

Variation of absorption by nozzle pressure was measured by changing the pressure (by adjustment of a needle valve) within the same trial. Using manufacturer data for changes of drop size and flow rate with pressure, the effect of drop size on absorption was inferred. The calculation is distorted by the use of 1.3 M NaOH solution instead of water, which is what the manufacturer data apply to. However, since the solution is constant for the three pressures, the trend should be similar as for water.

B.3 Data Analysis

In the continuous mode trials, data were relatively easy to analyze when significant drift was not present. Figure B.8 shows the CO_2 concentration over a such a trial. The difference between the average background concentration and average steady-state outlet concentration gives the rate of CO_2 absorption. Some subjective judgment is required to choose the bounds of each period to average over. Periods were chosen to exclude transient regions between measuring states and regions which appear excessively noisy.

We have samples of the solution taken periodically through some of the trials, including the one shown in Figure B.8, Trial 2. The CO_3^{2-}concentrations in those samples were measured with a Shimadzu Total Organic Carbon (TOC) analyzer after dilution and partial neutralization of the remaining NaOH. The results are shown in Figure B.9. If the rate of CO_2 uptake is taken as roughly constant, we would expect carbonate (CO_3^{2-}) concentration in solution to increase linearly with operating time, assuming a constant solution volume. But the solution volume changes due to evaporation (and entrainment and solution stuck on surfaces, but we take these to be comparatively minor). Using the measured evaporation rate from Trial 3 (which had similar conditions) to adjust the solution volume over time, we can make a prediction of CO_3^{2-}concentration based on the LiCOR measurements. This prediction is shown along with the TOC measurements in Figure B.9.

For Trials 6–12, a more nuanced method of analysis was needed. The procedure in these trials involved

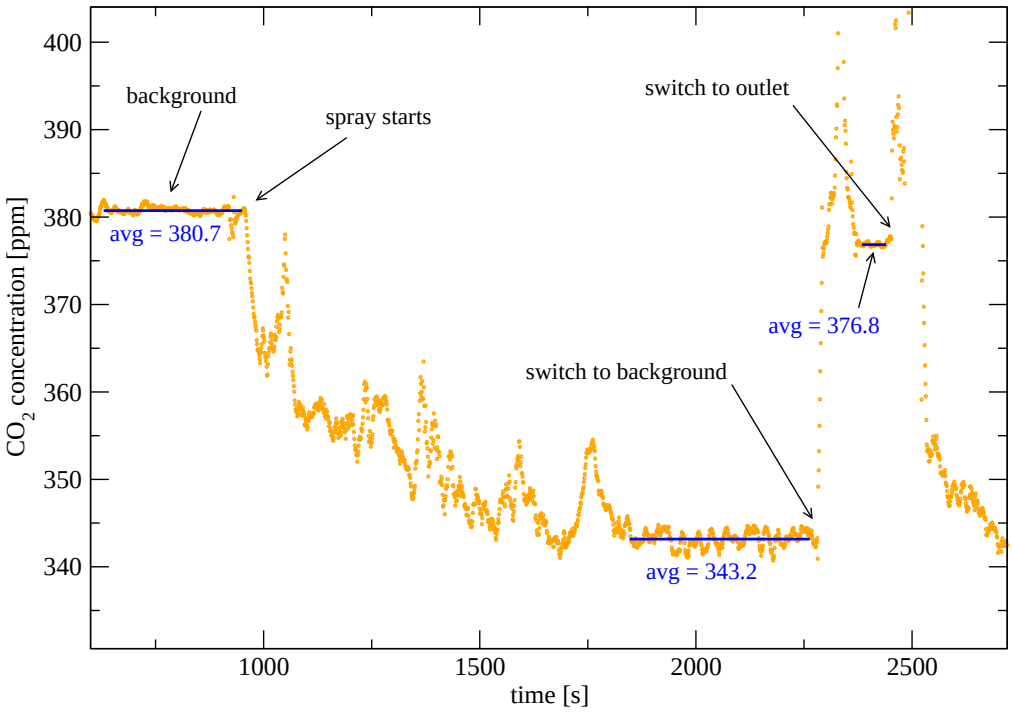

Figure B.8: Measured CO_2 concentration during a portion of Trial 2. Horizontal lines indicate the average value over the time span of the line. The difference between background and steady state outlet concentration represents the rate of CO_2 absorption into solution. "Switch to..." refers to changing the point of measurement, which begins in the outlet position.

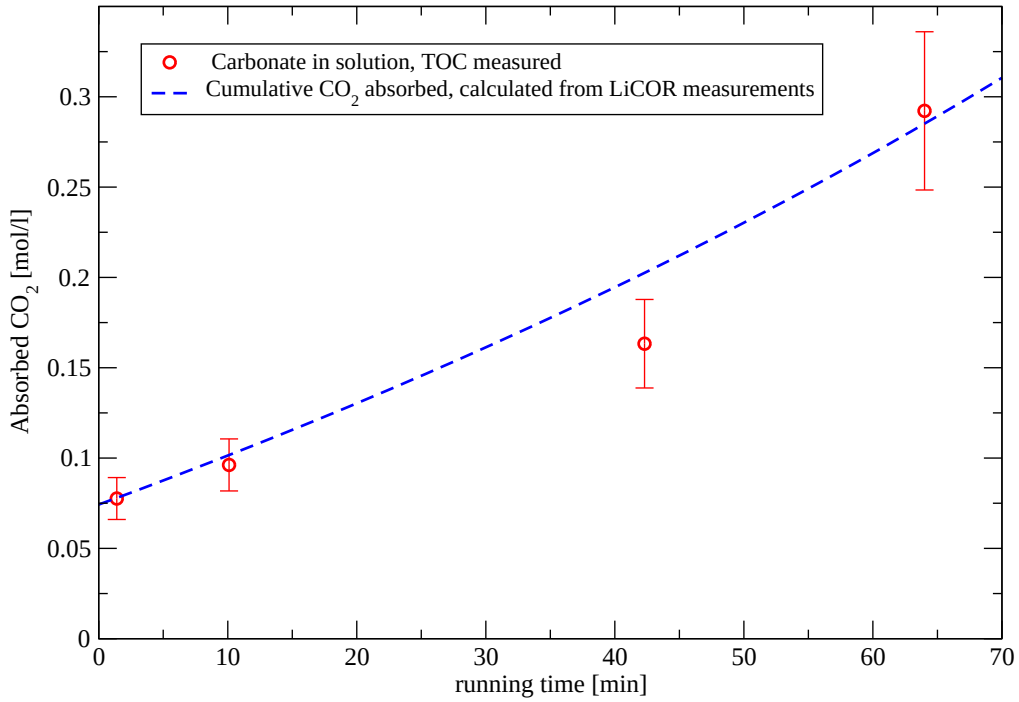

Figure B.9: CO_2 absorbed during Trial 2: comparison of two measurement methods. The solution starts with a nonzero carbonate concentration because of absorption during the previous trial. The Y-intercept of the LiCOR calculation is adjusted to match the first TOC sample in order to account for this initial carbon content. Error bars indicate subjective standard error of measurement.

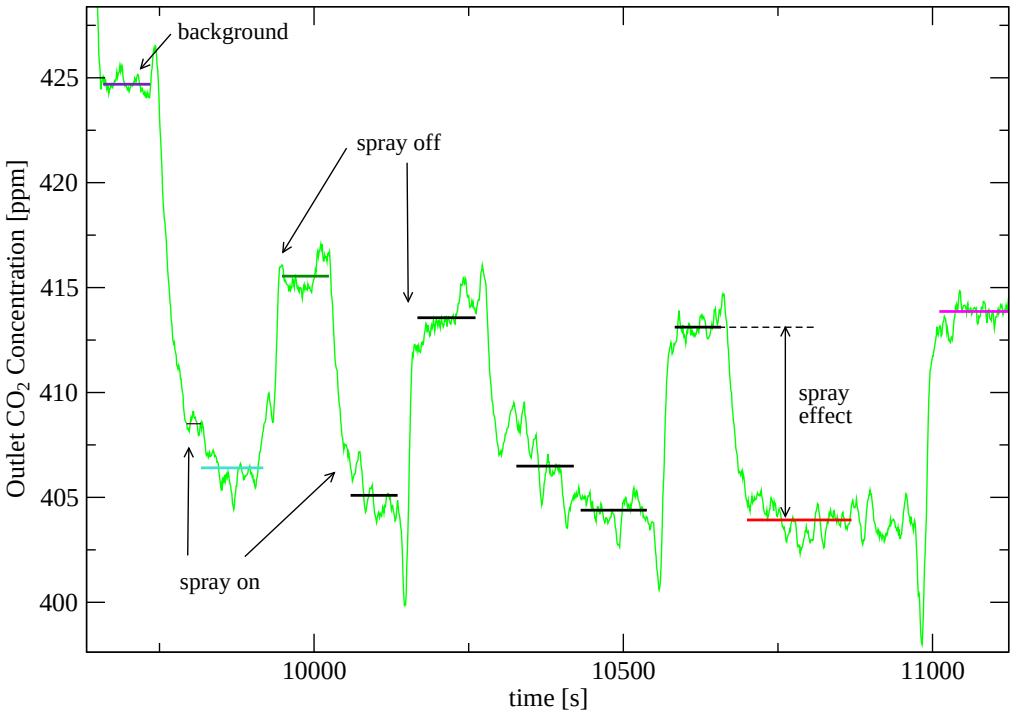

Figure B.10: Measured CO_2 concentration during a portion of Trial 11. Horizontal lines indicate the average value over the time span of the line. Valleys are periods with spray turned on while peaks (except the peak labeled "background") are periods with spray turned off. Difference between averages of adjacent peaks and valleys is taken as a measurement of absorption by the spray.

periodic switching of the spray on and off. This allowed better separation of the spray signal from background variation and from absorption by the wetted walls. Additionally, the magnitude of the spray signal in some of these trials was much smaller, 4–15 ppm instead of the 20–40 ppm of earlier trials. Multiple measurements were required to separate the signal from background noise. Figure B.10 shows data from a typical trial. Substantial subjective judgment is used in choosing the bounds of each averaging period. Bounds are chosen to exclude transient effects of when the spray is switched or point of measurement is changed and to exclude regions of noise or otherwise incoherent signal. Each switching cycle is analyzed by an average over at least one "spray on" region and one "spray off" region. The difference between a each pair of averages is taken to be an independent measurement for purposes of estimating uncertainty bounds. Sometimes multiple averages are drawn for one region because significant drift is visible. Averages are paired to be close in time and to avoid effects of visible drift. Data where noise or drift renders peaks incoherent are discarded.

This document, including figures and associated data analysis, was created entirely with free software, notably LyX, LaTeX, Grace, and Linux.