


Semiconductors and Sustainability: Energy and Materials Use in Integrated Circuit Manufacturing

By

Matthew S. Branham
B.S. Mechanical Engineering
Stanford University, 2006



Submitted to the Department of Mechanical Engineering in Partial
Fulfillment of the Requirements for the Degree of

MASTER OF SCIENCE IN MECHANICAL ENGINEERING
AT THE
MASSACHUSETTS INSTITUTE OF TECHNOLOGY

September 2008

©2008 Matthew S. Branham. All rights reserved.

The author hereby grants to MIT permission to reproduce and to distribute publicly paper and electronic copies of this thesis document in whole or in part in any medium now known or hereafter created.

Signature of Author: _____

Department of Mechanical Engineering
Aug. 15, 2008

Certified by: _____

Timothy G. Gutowski
Professor of Mechanical Engineering
Thesis Supervisor

Accepted by: _____

Lallit Anand
Professor of Mechanical Engineering
Chairman, Committee for Graduate Students

ARCHIVES

Semiconductors and Sustainability: Energy and Materials Use in Integrated Circuit Manufacturing

By

Matthew S. Branham

Submitted to the Department of Mechanical Engineering
On 15 August 2008 in Partial Fulfillment of the
Requirements for the Degree of Master of Science

ABSTRACT

Semiconductors have propelled an incredible revolution in the way we generate, access, store, and communicate information; the effects of this revolution have transformed culture, society, and the economy. At the same time, there have been increasingly portentous signs of the extent to which people are overtaxing their planet and its resources. This study is conducted in the context of both of these upheavals with the goal of better understanding how and how much energy and material are used in semiconductor manufacturing. The results of a case study of the Analog Devices Micromachined Products Division (MPD) fabrication facility are presented in support. In it, energy consumption is examined both from a top-down (fab-level) and bottom-up (equipment-level) perspective. Total fab material use is also presented. We find that the facility currently uses roughly 1.53 kWh of electricity and 35 grams of chemicals per square centimeter of product wafer. Electrical energy consumption is further broken down by facility systems and then by process areas.

To give additional insight into the energy and materials intensity of semiconductor manufacturing, results from the monitoring of specific processes are presented. In order to provide structure to the method of evaluating the efficiency of these processes, a thermodynamic model of manufacturing is first introduced. The study concludes with a look at the fundamental reasons why the industry is so energy intensive and changes that Analog Devices could make to reduce their fab's energy intensity.

Thesis Supervisor: Timothy G. Gutowski
Title: Professor of Mechanical Engineering

Acknowledgements

I must begin by thanking my advisor, Professor Timothy Gutowski, who guided me in this project and helped me continue to grow intellectually over my time thus far at MIT. I am greatly in his debt for his willingness to help me identify and to support me in a research project that I would find interesting and engaging. I do not believe that many graduate students have the fortune of an advisor who is as amicable and accessible as Prof. Gutowski.

I would be remiss in not next thanking Craig Core, the former fab manager at Analog Devices Micromachined Products Division. He was exceedingly gracious in bringing me into his facility and granting me access to the equipment and information needed to produce a high-quality project. Before speaking with Craig, I spent five mostly fruitless months trying to find a host company; were it not for his kindness and openness to hosting academic research, I would certainly not be receiving my Masters degree at this time. I also owe a thank you to his successor, Jim Long, who allowed me to continue with my project at his facility after taking over from Craig. Additionally, he has fielded a great number of questions from me, and despite the awesome demands on his time, has always answered very promptly. Thank you both very much.

Sandy Annese and Paul Sessler – the Analog Devices facilities electricians and jacks-of-all-trades who worked with me in making untold electrical measurements – deserve more than what a thank you can convey. I hold them in the highest esteem for the professionalism and commitment to quality with which they carry out their work. More than that, they went out of their way (and practically risked their lives) to help me get the data I sought. I cannot venture a guess as to the number of hours we spent side-by-side working together over the past year, but in hindsight those were some of the most enjoyable times of this project. We commenced this undertaking as strangers, worked through it as colleagues, and finished as friends. Sandy and Paul, I cannot thank you enough for your time and effort on my behalf; I hope that you get a vicarious sense of satisfaction out of seeing this thesis in print, since it is practically yours as well.

To my other friends at Analog Devices who assisted me in various ways on this project, I thank you too: Dan Clark who from the beginning was instrumental in getting facility data and with whose blessing I was allowed to conduct the electrical measurements; Jack McCarthy who gave me my first fab tour and was always helpful; Regan Colm, Sharon Lu, and Dilipkumar Mahenthiran who made time in their packed schedules to help me run process tests; Ron Motta and Bruce Durgin who showed me the ins-and-outs of the process equipment; Bill Walker, Dave Kneedler, and Linda Swift for fielding questions and providing data; and Danielle and Eileen for greeting me with a smile every day. I have nothing but good things to say about Analog Devices.

Finally, I would like to express my appreciation to my parents for giving me every opportunity for which I could ask, and to my whole family – Dad, Mom, Lucas, Jordan, and Ben – for the unfaltering support they have always provided me, the past two years being no exception. I am so blessed to have the family that I do.

Table of Contents

Chapter 1: The Economic, Social, and Environmental Context of Semiconductors

1.1 Introduction.....	8
1.2 Integrated Circuit Manufacturing.....	10
1.2.1 The semiconductor industry.....	11
1.2.2 Semiconductor Manufacturing Overview.....	12
1.2.2.1 Lithography.....	12
1.2.2.2 Ion Implantation.....	13
1.2.2.3 Thermal Oxidation.....	13
1.2.2.4 Thin Film Deposition.....	13
1.2.2.4.1 Chemical Vapor Deposition.....	13
1.2.2.4.2 Physical Vapor Deposition.....	14
1.2.2.5 Etching.....	15

Chapter 2: Early Work on Energy and Materials Use in Semiconductor Processing: What Do We Already Know?

2.1 Introduction.....	17
2.2 Energy.....	18
2.3 Chemicals.....	18
2.4 Water Usage.....	19
2.5 Previous Work on Semiconductor Process Analysis.....	19
2.6 Contribution of this Thesis to the Field of Sustainability/Industrial Ecology.....	20

Chapter 3: Data Acquisition Method

3.1 Introduction.....	21
3.2 Material Inputs, Product Wafers, and Top-Down Electrical Data.....	21
3.3 Bottom-Up Fab Electrical Data.....	22
3.4 Process Analysis.....	23
3.5 Measurement Equipment.....	23
3.6 Power Measurements and Justification of the Steady-State Assumption.....	25

Chapter 4: Facility-Level Material and Energy Inputs to the Manufacturing System

4.1 Introduction.....	29
4.2 Top-Down Assessment of Fab Energy Use.....	29
4.3 Relationship to State-of-the-Art Fabs.....	33
4.4 Top-Down Assessment of Material Inputs to the Manufacturing Facility.....	35
4.5 Total Mass of Materials Required in the Front-End Manufacturing of a MEMS Chip.....	39
4.6 Bottom-Up Assessment of Fab Energy Use.....	40

4.7 Comparison of Bottom-Up and Top-Down Results.....	43
4.8 Seasonal Variation in Energy Consumption.....	44
Chapter 5: Analysis of Semiconductor Manufacturing Processes	
5.1 Introduction.....	45
5.2 A Thermodynamic Model of Manufacturing Systems.....	45
5.3 Semiconductor Manufacturing Process Analysis.....	50
5.3.1 Chemical Vapor Deposition.....	50
5.3.1.1 <i>Chemical Vapor Deposition of Silicon Dioxide and Silicon Nitride</i> <i>(Deposition Only)</i>	54
5.3.2 Sputter Deposition.....	55
5.3.2.1 <i>Sputtering of an Aluminum Copper Film (Deposition Only)</i>	57
5.3.3 Dry Chemical Etching.....	57
5.3.3.1 <i>Dry Oxide Etch</i>	58
5.3.3.1.1 <i>Dry Oxide Etch (Etch Time Only)</i>	59
5.3.3.2 <i>Dry Nitride Etch</i>	60
5.3.3.2.1 <i>Dry Nitride Etch (Etch Time Only)</i>	61
5.3.4 Wet Etching of Silicon Dioxide and Silicon Nitride.....	61
5.4 Process Rate vs. Specific Energy Consumption.....	65
5.5 Closing Thoughts: Efficiency and Sustainability.....	67
Chapter 6: Efficiency Improvement Recommendations for ADI's Micromachined Products Division Fab	
6.1 Introduction.....	68
6.2 Chiller Upgrade.....	68
6.3 Optimize Airflow and Exhaust Velocities.....	70
6.4 Throttle the DI Water Pump.....	71
6.5 Avoid Overlapping Heating and Cooling.....	72
6.6 General Recommendations.....	72
6.7 Tool Energy.....	73
Chapter 7: The Next Steps	
7.1 Semiconductors and Sustainability: Energy and Materials Use in Input Chemicals.....	74
7.2 Next Steps for Analog Devices.....	74
Appendix A: Measurements of Energy Use in Carbon Nanotube Growth	
A.1 Nanotube Growth Using a Radiative and Substrate-Heated Furnaces.....	76
Appendix B: Summary Power Data from Bottom-Up Tests.....	
77	
Bibliography.....	
80	

Chapter 1: The Economic, Social, and Environmental Context of Semiconductors

1.1 Introduction

Semiconductor devices pervade modern life. From the suite of sensors and microcontrollers in an automobile to the range of microchips controlling the computers, portable phones, cameras, video recorders, calculators, and planners that punctuate daily existence, the revolution brought about by integrated circuit technology is stunning. Since Jack Kilby engineered the first integrated circuit at Texas Instruments in 1958, the production of semiconductor devices has exploded. Whereas in 1968 some two billion (2×10^9) transistors were produced each year worldwide, by 2002 that number had grown to almost one quintillion ($1,000,000,000,000,000$ or 1×10^{18}) [1]. Such a phenomenal increase in the availability of computing power has been driven in part by the factor of 10^5 increase in the number of transistors that can be built on a chip (see Figure 1.1), coupled with the similarly dramatic growth in the number of chips produced by the industry. As a result, this technology has driven a remarkable expansion in productivity over the last thirty years that has radically elevated living standards (and, many would argue, quality of life) in industrialized nations and increasingly in developing countries. The innumerable benefits of semiconductors cannot be understated.

The development of integrated circuit technology has been concurrent with another explosion, that of humanity's imprint on the planet. In the same time period between 1958 and the present, earth's population has ballooned from three billion to 6.5 billion and has combined with other factors to ratchet up the stress humans exert on the planet [2]. For example, greater affluence combined with more sophisticated technologies now allows for the mining and mobilization of incomprehensible amounts of material (more than 3 million metric tons per year in the United States alone [3]). Material and resource acquisition, agricultural expansion, and real estate development have put such pressure on ecosystems that the effects (e.g., extinction of species, loss of biodiversity, and reduction in land productivity) are becoming increasingly acute. Worldwide power consumption now amounts to 450×10^{15} BTU/year (1.3×10^{14} kWh/year), up from 283×10^{15} BTU/year in 1980 [4]. Pollution of all sorts continues to foul waterways and land that are critical not only to human health but also to the health of the natural world. All of this is occurring against a backdrop of ever-increasing atmospheric greenhouse gas concentration of which the consequences are not certain, but that threatens extreme climate change with a

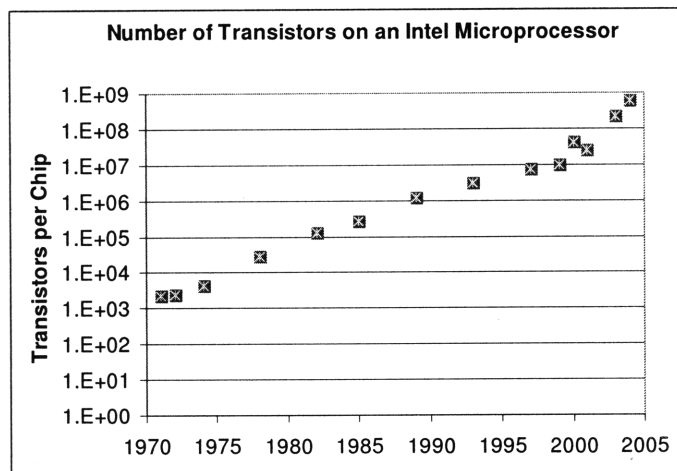


Figure 1.1: Growth in the number of transistors on Intel microprocessors [5].

host of consequences for the planet and its inhabitants. It is in this context – of the indisputable value and importance of semiconductor products coupled with an awareness of the strain exerted on the natural environment on which people inevitably depend – that the present study has been conducted.

This study of the energy and materials used in the production of semiconductors has three objectives. The first is to understand and characterize in detail how resources (particularly electricity) are utilized by a semiconductor fabrication facility, referred to in industry parlance as a “fab”. This includes a top-down assessment of the quantity of the fab’s energetic and material inputs and useful outputs on a yearly basis, as well as a bottom-up approach that seeks an understanding of how electricity is used once it enters a fab. The second goal is to use this information to identify ways in which wafer fabs can reduce their resource footprint. The final goal is to measure several specific processes from the semiconductor manufacturing chain to learn more about why semiconductor processes are so energy and materials intensive.

There are multiple reasons that the semiconductor industry is compelling to study from a sustainability perspective. Foremost is the industry’s impressive size and growth rate. Semiconductors form the basis of the \$373 billion computer and electronics sector of American manufacturing (total shipments), which as of 2006 trailed only four other sectors: transportation equipment manufacturing (\$698 billion), chemical manufacturing (\$658 billion), petroleum and coal products manufacturing (\$549 billion), and food manufacturing (\$538 billion) [6]. The importance of semiconductors to the global economy has paralleled their rise in the American economy, with the global sales of semiconductors alone reaching \$255.6 billion in 2007 (see Figure 1.2) [7]. Moreover, the global semiconductor industry has sustained a compound annual growth rate of 10.6% since 1990 (12.8% since 1982) [8]. Such a high growth rate in a developed industry implies not only that semiconductors and electronics will become an increasingly important link in the world’s economy but also that reducing its environmental footprint in absolute terms will be quite a challenge.

Underlying the growth rate of the semiconductor industry is an aggressive pace of change and innovation. The definition of cutting edge in semiconductor processing is typically rewritten less than every two years, meaning that a constant turnover of technology is occurring. This pace of change can be beneficial from the perspective of sustainable manufacturing since it allows for the fairly rapid incorporation of new ideas to reduce resource consumption into the semiconductor business. The time required to convert a traditional technology or practice in the semiconductor

World Semiconductor Revenue (Billion Dollars)




Figure 1.2: Semiconductor Industry Worldwide Revenue, 1970-2002 [9].

industry into a more sustainable alternative is logically much shorter than would be the changeover to a higher mileage vehicle fleet, for instance, owing to the longer lifetimes of automobile technologies. For example, recent advances in lowering the power draw of fab tools while idling is being incorporated relatively quickly as new fabs are built and older ones retooled. The possibility to effect change rapidly in this industry makes it an attractive target for sustainability research.

Finally, in spite of its stereotypically clean outward image, there are very real environmental concerns associated with the semiconductor industry. For example, integrated circuit manufacturing requires the input of a suite of toxic, corrosive, and otherwise harmful chemicals and gases that must be treated before disposal. Modern-day wafer fabs typically handle this waste responsibly, but the existence of 29 Superfund sites at former semiconductor facilities in the San Francisco Bay Area alone shows that is not always the case [10]. These cases also illustrate the potential for harm from the release of certain semiconductor chemicals into the environment. Moreover, improved waste management does not exculpate the industry from the impact of the production of these chemicals, which require large inputs of energy and raw materials to attain the purity required in the semiconductor industry. Wafer fabs also consume massive amounts of water in their production processes, which can strain local environments (particularly in water-starved regions like the desert Southwest of the United States or interior China). Lastly, the energy consumption of wafer fabs is sizable. Although the amount of energy consumed by the computer and electronic products sector is dwarfed by that used in chemical production, petroleum and coal products manufacturing, paper-making, primary metals production, food production, and the nonmetallic mineral products sector, the sector as a whole used some 200 trillion Btu in 2002 (slightly less than 1% of the United States's manufacturing energy consumption) [11]. As mentioned earlier, because wafer fabs consume a large amount of high-purity input chemicals, the industry is also responsible for energy use in other sectors of the economy. The environmental issues involved in semiconductor processing are quite complex, often involving tradeoffs between different environmental impact modes.

It should be mentioned that the present study is the final portion of a larger project within the Environmentally Benign Manufacturing Lab at MIT focusing on the environmental characterization of manufacturing processes. Earlier work analyzed the so-called traditional and advanced machining processes such as metal casting, injection molding, traditional machining (e.g., milling, turning), grinding, waterjet cutting, and electrical discharge machining [12-15]. This work deals with the micro and nano processes that typically operate at material processing rates slower than about $4 \cdot 10^{-3}$ cm³/sec.

1.2 Integrated Circuit Manufacturing

Before continuing on to the environmental assessment of semiconductor manufacturing, it is necessary to understand both the size and importance of the industry as well as the basics of semiconductor processing.

1.2.1 The semiconductor industry

In a little more than fifty years, semiconductors have evolved from a laboratory curiosity to a major component of the world economy. The economic importance of integrated circuits is underscored by the value of integrated circuits produced annually, which in 2007 amounted to \$256 billion or about 0.5% of world GDP [8]! This figure does not incorporate the value of ancillary industries such as semiconductor equipment manufacturing (\$41 billion in 2007 [16]) or semiconductor chemical manufacturing, nor does it include the broader electronics and computer industries of which semiconductors are the enabling technology. The inclusion of these sectors pushes the global cumulative economic importance of semiconductors into the trillions of dollars. The growth rate of the industry is as compelling as its sheer size. As mentioned in section 1.1, semiconductor revenues have averaged a 12.8% compound annual growth rate since 1982. A 13% industry-wide growth rate is admirable in its own right, but in light of the continuously falling price of semiconductors it is even more impressive.




Figure 1.3: Silicon wafer shipments in millions of square inches (MSI) [17].

The growth of semiconductor manufacturing can also be measured in terms of material output. There are several ways material output may be represented, such as by the number of transistors produced (see Figure 1.1), the quantity of microchips produced, or the number of wafers shipped. Probably the most reliable and convenient way to represent output product is using the area of wafers shipped, which normalizes for varying wafer sizes. Data from industry trade groups SEMATECH and SEMI show that silicon wafer shipments have grown from approximately 1.9 billion square centimeters in 1973 to 56 billion square centimeters in 2007, yielding a compound annual growth rate of about 10% (Figure 1.3) [17-18]. Note that this data does not include figures for other types of wafers used such as silicon-on-insulator (SOI), germanium, or other III-V semiconductors, which contribute a negligible fraction to the overall area of wafers shipped (<2%). One statistic hidden in the data on wafer shipments is the reality that microchips are continuously shrinking, such that many more chips are now produced per unit wafer area. This reality has allowed industry revenues to grow at a faster rate than that at which wafer shipments have grown in spite of falling chip prices.

Another interesting aspect of the integrated circuit industry is its distinct regionalism. Integrated circuits were first developed in the United States, and it is in the U.S. that the first major manufacturers were established (Fairchild, Intel, etc.). The U.S.

manufacturers dominated the global market through the 1970s, although by the late 1970s Japan's market share was growing quickly and surpassed that of the United States in 1986. In the time since then, the center of gravity in semiconductor manufacturing has shifted markedly to East and Southeast Asia (excluding Japan). As of 2005, the respective share of the market was roughly equal between Japan, Europe, and the United States at 17-19%; Asia/Pacific's share stood at 46%. The Asian market share has skyrocketed since the end of the 2001 recession, and it appears that Asia/Pacific will remain the dominant semiconductor manufacturing region for the foreseeable future.

1.2.2 Semiconductor Manufacturing Overview

The manufacture of semiconductors is an extraordinarily complex, capital-intensive undertaking; it is not uncommon for a microchip to require upward of 400 distinct process steps in going from bare wafer to functional device. Broadly speaking, there are five general classes of processes used in the front-end of chip fabrication: lithography, ion implantation, etching, thin film deposition, and thermal oxidation. This categorization is by no means a complete catalog of semiconductor processes, nor does it include the back-end processes (e.g., contact formation, packaging). Instead, it is meant to include the most fundamental front-end processes, as these are the focus of this thesis.

Each of the aforementioned classes has myriad variations that are implemented in various sequences to yield a final product. For a more thorough explanation of integrated circuit fabrication, see [19]. The following is a brief overview of the basic processes used in the manufacture of semiconductors, which is necessary background in understanding the analyses of the succeeding sections.

1.2.2.1 Lithography

Lithography is used to replicate patterns from a master copy (mask) on to a wafer substrate and is probably the process that first comes to mind when thinking of semiconductor fabrication. In actuality, lithography is not a single process but instead requires multiple steps. The first step in lithography is to spin a thin layer of a photoactive compound – termed photoresist – onto the surface of the wafer substrate. Photoresists come in two classes: positive resists (the dominant form used in modern manufacturing) which become more soluble in a developer when exposed to light, and negative resists that become more insoluble upon exposure. In the second step, a mask inscribed with the pattern to be transferred is positioned between a light source (of the wavelength required to interact with the photoresist) and the substrate with the photoresist layer. The mask pattern is composed of transparent and opaque regions, which expose the proper regions of the photoresist. After the photoresist has been exposed, the final step in a lithography sequence is to immerse the wafer in developer, which dissolves the soluble portions of the photoresist and leaves behind the insoluble ones. The wafer substrate is left with a three-dimensional photoresist pattern on its surface, which becomes the precursor for processes like ion implantation or etching.

1.2.2.2 Ion Implantation

A critical component of device production is the formation of the active regions, or areas of precisely controlled conductivity and charge that endow an integrated circuit component (e.g., a transistor) with much of its functional characteristics. Active region formation is accomplished first by implanting positively-charged P-type ions (such as boron) or negatively-charged N-type ions (such as arsenic). Prescribed doses of these ions are accelerated in massive tools using sufficiently high energies (typically 25 to 200 keV) to embed them in the substrate at an appropriate depth. The photoresist pattern described in the preceding section (or a thin film) shields certain portions of the substrate from the bombarding ions such that only the desired active regions are impinged upon. The process of implanting individual ions into a substrate damages the crystalline structure of the substrate. The damage is repaired by an annealing step in which the wafer is heated to allow dislocation strains to relax and the crystal structure to heal.

1.2.2.3 Thermal Oxidation

Oxidation is a specialized process used to grow a silicon dioxide layer from a silicon substrate and is typically an early step in device fabrication. Thermal oxide layers are grown by streaming oxygen (dry oxidation) or water vapor (wet oxidation) over a substrate in a furnace heated to very high temperatures (usually 700-1250°C). During thermal oxidation, oxygen diffuses through any SiO₂ layer already present on the wafer surface. The oxygen oxidizes silicon at the silicon/silicon dioxide interface and the SiO₂ layer grows outward from there using silicon cannibalized from the substrate. Oxide layers may also be deposited instead of grown, using CVD techniques discussed in the next section. Thermal oxides tend to be of much higher quality than deposited oxides, so they are used in the formation of transistor gates and other low-defect-tolerance applications. Additionally, thermal oxides are relatively thin since the rate of oxidation decreases with increasing oxide thickness.

1.2.2.4 Thin Film Deposition

In contrast to thermal oxides that are grown using the substrate as a precursor, other types of thin films are deposited. Thin film deposition techniques may be subdivided into chemical vapor deposition and physical deposition.

1.2.2.4.1 Chemical Vapor Deposition

Chemical vapor deposition (CVD) is used to deposit a large variety of thin films, such as polycrystalline silicon, silicon nitride, and silicon dioxide. In particular, CVD is used in the deposition of materials that have stable gaseous chemical precursors, mostly limiting its commercial application to thin films containing nonmetals like silicon, phosphorus, boron, and nitrogen (with the notable exception of tungsten) [19]. CVD comes in many different flavors but is typically classified as either atmospheric-pressure CVD (APCVD), low-pressure CVD (LPCVD), plasma-enhanced CVD (PECVD), and high-density plasma CVD (HDPCVD). As indicated by their names, APCVD and

LPCVD are distinguished by the pressure of the reaction chamber: whereas LPCVD takes place in an evacuated (sub-atmospheric) chamber, APCVD takes place at atmospheric pressure. PECVD is used to deposit films at lower temperatures than would otherwise be required, using a highly ionized plasma to supply the energy needed. PECVD also is done under vacuum conditions. The final technique, HDPCVD, is a fairly recent development that is primarily used to deposit oxide films. It combines PECVD and sputtering (see section 1.2.2.4.2) to achieve better deposition uniformity and gap filling than the other methods allow. As with PECVD, HDPCVD is conducted at very low pressures.

In a CVD process, chemical precursors of the type of film to be deposited, plus additional carrier gases such as nitrogen or argon, are flowed through a heated chamber in which the target wafers are situated. In some depositions, the decomposition of a single precursor gas results in the deposition of a film, as in the case of epitaxial silicon growth where silane undergoes the transformation $\text{SiH}_4 \rightarrow \text{Si} + 2\text{H}_2$ at 1000-1250°C (lower temperatures will form polycrystalline silicon). In most chemical vapor depositions, however, two gases react in the vicinity of the target wafers (or any other surface for that matter) to deposit the prescribed film. For example, silicon nitride films can be formed by reacting dichlorosilane and ammonia in the reaction $3\text{SiH}_2\text{Cl}_2 + 4\text{NH}_3 \rightarrow \text{Si}_3\text{N}_4 + 6\text{H}_2 + 6\text{HCl}$. The great majority of the mass of chemicals that enter a CVD chamber are exhausted, while much of the remaining gas ends up depositing on the walls of the chamber. Only a tiny fraction of the material entering the deposition chamber ends up on wafers. CVD is also energy-intensive, given the requirements for heating elements, vacuum pumps, and RF generators (for PECVD and HDPCVD).

1.2.2.4.2 Physical Vapor Deposition

As its name implies physical vapor deposition (PVD) uses physical techniques to mobilize a vapor for deposition on a wafer substrate. PVD is particularly useful for the deposition of metal films. The most basic PVD technique is evaporation, in which a source material is heated in a chamber that has been pumped down to a very low pressure. As the source material heats up, surface atoms vaporize from the bulk material. Since the chamber has been evacuated, the atoms rarely collide with other atoms and therefore travel in a line until they impact an opposing surface, ideally the target wafers which have been placed around the source with faces perpendicular to the angle of impingement of the evaporated species. Evaporation has difficulty depositing alloys and high-vapor pressure elements, and it typically has inadequate step coverage (uniformity of deposition on a three-dimensional surface), so it is not often used in semiconductor manufacturing today [19]. Sputter deposition has become the preferred PVD technique.

In a basic sputtering process, a plasma composed of argon atoms and ions is developed by applying a voltage across two electrodes. In a DC sputtering setup, the wafer(s) serves as the positively charged anode, while a target of the material to be deposited serves as the negatively charged cathode. As argon atoms are ionized in the plasma, they are accelerated by the electric field toward the target material where they impinge and eject atoms. These atoms are then transported through the plasma where they form a film on the wafers as well as the chamber walls. There are several variants to the sputtering setup described here, such as RF sputter deposition (which can be used to

deposit nonmetals, since DC sputtering can only handle metal films), reactive sputtering (a hybrid CVD/PVD process where the atoms ejected from the target undergo a chemical reaction before deposition on the wafer surface), and bias-sputter deposition (where the plasma sputters the wafer as well as the target, allowing for a more uniform film to be deposited). Among the films that are commonly deposited using sputtering are Al, Cu, AlCu, TiW, and TiN. Sputtering is particularly effective at depositing alloys and compounds of precise composition.

1.2.2.5 Etching

Etching is critically important to defining the three-dimensional structure of devices. It is used to pattern the wafer surface by removing portions of the original wafer substrate or of thin films that have been deposited. Etching is done either in liquid solution – wet etching – or using gases and plasmas, known as dry etching. Both wet and dry etches come in a veritable cornucopia of varieties depending on target material to be etched, selectivity requirements (preferential removal of one material over another), and directionality requirements (isotropic vs. anisotropic).

Wet etching is accomplished by immersing wafers in a liquid chemical bath (often heated). The chemical or combination of chemicals is formulated to attack a certain type of material on the wafer surface while leaving others relatively unaffected (this is referred to as *selectivity*). In wet etching, as with dry etching, areas of the wafer surface that are not to be etched are shielded using photoresist (following a lithography process) or a thin film that has been patterned in an earlier step and is impervious to the etchant. Although wet etches are typically quite selective, they often etch isotropically, meaning that the etch proceeds not only downward through the film but also laterally. In microelectronics manufacturing, anisotropic etching leading to nearly vertical sidewalls is typically preferred (Figure 1.4). It is for this reason that dry etching has become the preferred method to accomplish most etches in the semiconductor industry.




Figure 1.4: Example of an isotropic etch (left) and an anisotropic etch (right).

Etches that do not use liquid chemicals but instead use gaseous etchants in a plasma are referred to as dry etches. In the most common variety, plasma etching, a chemical etchant is fed into a chamber containing a plasma. The plasma produces highly reactive and energetic species from the gas, which then remove material from the wafer surface. Plasma etches thus have both a chemical component, owing to the reactivity of the plasma gases, and a physical component akin to sputtering. The ratio of chemical etching to physical etching can be controlled by the type of etch system used, the selection of etch gases, and so on. Generally, a dry etch that has a greater chemical component will produce a more selective but less directional etch, while a physical etch will be more directional but less selective. In ion-enhanced etching, for example, a

neutral species is introduced into the etch chamber which etches by a sputter mechanism, producing a highly directional (vertical sidewalls) though non-selective etch. As with PVD processes requiring a plasma, most dry etch processes are quite energy and materials intensive, which is particularly important since the materials are very reactive.

There are of course many other processes that go into semiconductor manufacturing, such as chemical mechanical polishing (for wafer planarization), cleaning processes (related to wet etching), metrology, and so on. Additionally, the processes discussed above have many subclasses for which there is insufficient space for discussion. But we are now equipped with a sufficient working knowledge of the fundamental, most common, and most energy- and resource-intensive processes to get the most out of the latter sections of this thesis.

Chapter 2: Early Work on Energy and Materials Use in Semiconductor Processing: What Do We Already Know?

2.1 Introduction

Issues of resource consumption in manufacturing have been entirely the domain of private industry only until recently. One can speculate on the reasons this has been the case. To begin, it is a basic tenet of manufacturing and commerce that individual economic players will seek to continuously reduce the cost of production in order to better compete in the marketplace. Such a drive for efficiency should serve to decrease the cost and thus – in the absence of substitution between different inputs – the quantity of material and energetic inputs to a manufacturing system. This argument implies that the burden of resource reduction be left to the manufacturers. Additionally, the environmental concerns associated with manufacturing have historically been the more visible and immediate: disposal of harmful compounds, release of toxic emissions, groundwater contamination, etc. Roughly in the last two decades, however, there has been a growing realization of the significance of material and energy use in manufacturing processes to the total environmental impact of a product. These more diffuse causes of environmental degradation provide fertile ground for academic research; by illuminating the magnitude and type of impact of energy and material use, academic research can lead industry toward a more sustainable model. The advent of sustainability science has greatly increased our understanding of where are the greatest environmental burdens in the life cycle of a product.

As revealed in a growing body of literature, semiconductors, unlike the great majority of consumer products, require the greatest input of energy and materials during the production stage of their lifetimes. Of the total energy consumed over the lifetime of a microchip, roughly 50% goes into the fabrication stage, while a little more than 25% goes into the use phase and the remainder into other components of the manufacturing process (materials production and assembly) [20]. This is in stark contrast to an automobile, for example, where energy required during the manufacturing phase is less than 20% of the total life cycle energy requirement. It is for this reason, and because semiconductor manufacturing is so materials intensive too, that we focus specifically on the manufacturing phase of the semiconductor life cycle here.

The foundational work on the life cycle of a computer chip – Williams, Ayres, and Heller's 2002 paper "The 1.7 Kilogram Microchip" – synthesizes a variety of sources to piece together the energy and materials required to manufacture a 2-gram memory chip [20]. Among the important findings and figures of the analysis of the front-end manufacturing process:

1. From sources reporting electrical energy use in semiconductor manufacturing, Williams et. al. take 1.5 kWh per square centimeter of silicon processed as a representative value.
2. Excluding bulk nitrogen and water use, data from a European IC manufacturer gives the total mass of chemical inputs as 45.2 kg/cm².

3. Summing the mass of input chemicals (minus water and nitrogen) and that of fossil fuels required to produce the input energy, the authors conclude that about 1.7kg of material are required for the production and use of a 2-gram chip. This result assumes that for every chip produced, 1.6cm² of silicon must be processed. Normalizing the Williams data to 1cm² and stripping out the mass of fossil fuels required for the non-fabrication stages of the lifecycle yields 537 grams of material required for the manufacture of 1 cm² of silicon (492 grams of fossil fuels and 45 grams of input chemicals).

Since this report constitutes the most thorough – and recent – analysis of the production of a computer chip, it is to the findings of “The 1.7 Kilogram Microchip” that the results of the present study will be compared.

As noted earlier, aside from the raw input material data, the Williams report presents data amalgamated from secondary sources. Primary source data on energy, input chemicals, and water consumption at wafer fabs is quite scarce; here we present an overview of that which is available.

2.2 Energy

One of the earliest studies on the environmental impact of semiconductors, the 1993 Microelectronic and Computer Corporation’s (MCC) evaluation of a model computer workstation found that an average six inch wafer required 285 kWh of electricity, or approximately 1.6 kWh per cm² [21]. Williams’s paper relates data from a 1997 study by the Japan Electronics Industry Development Association, which reported electricity use of 1.44kWh/cm². These data are somewhat dated now, since in the intervening time many facilities have implemented intensive energy reduction strategies (although at the same time, fab cleanliness and minimum feature requirements have tightened, resulting in competing effects on energy demand). In 2004, SEMATECH (a semiconductor manufacturers trade group) stated that the average power consumption amongst their member fabs was 777 kWh per cubic foot of cleanroom space [22]. The units of kWh per area of cleanroom space are inconvenient for the purposes of making a comparison based on manufacturing volume. Unfortunately, more recent data on fab energy use is tightly held and not publicly available. Since all available data show energy intensities varying between 1.4 and 1.6 kWh/cm², we follow the Williams example and use 1.5 kWh/cm² as a representative value.

Electricity is the dominant energy input into semiconductor manufacturing systems. Some facilities use minor amounts of fossil fuels for purposes such as gas abatement. Process heat supplied by steam is another energy input used by some facilities, primarily for heating and humidifying fab air and heating office areas. Sparse data is available on these inputs to fabs owing to the variability in their use.

2.3 Chemicals

Aside from the anonymous firm data presented by Williams et. al. – the most complete data set available – most other estimates of chemical usage are estimates based on national databases or data distilled from non-transparent studies. These estimates range from only 1.2 g/cm² to 610 g/cm² of silicon processed. Since these other estimates

are not transparent (in the sense that it is impossible to explain what and where specific chemicals and materials are being used), for the purposes of this paper we take the Williams data point (45.2 g/cm^2) as the reference value. See Chapter 4 for a detailed list of the chemical input data from the Williams study.

2.4 Water Usage

Wafer fabrication requires vast amounts of water. A straightforward internet search will yield many articles referring to the amount of water used in semiconductor processing, but it seems that these articles cite a relatively small number of primary sources. One of them, the 1993 MCC study, reports that 60L (~ 60kg) of water are used per cm^2 of silicon processed [21]. However, it is also noted in the MCC study that a state-of-the-art facility in 1993 was known to use only 21L/cm^2 (~ 21kg/cm^2). The results from an internal Intel study in 2004 shows that about 20L/cm^2 of city water is required to supply the DI water for the manufacture of an Intel Pentium® 4 [23]. This number almost certainly undercounts total water usage, but in general it seems that a number between 20kg/cm^2 and 40kg/cm^2 is reasonable for water consumption in a modern fab.

2.5 Previous Work on Semiconductor Process Analysis

There has been some relatively recent work out of the University of Texas that has focused on the development of parametric models for a number of standard wafer fabrication processes. These parametric models are intended to be used to anticipate resource requirements for specific processes, inform decisions on the tradeoffs between multiple process alternatives, and improve process resource requirements. The published paper presenting the development of these modules provides some data on thermal oxidation energy and oxygen requirements [24]. More useful is the final report submitted to the EPA on the project, which presents data on wet chemical cleans, oxide deposition, and thermal oxidation [25]. For example, the authors report that energy consumption in a wet oxidation process at 1050°C is given by:

$$E = 1.05 \cdot 10^{-4} \mu + 0.0915$$

where E is the energy requirement in kWh/cm^2 of silicon and μ gives the thickness of the oxide film in angstroms.

Boyd and Dornfeld's 2006 paper entitled "Life Cycle Inventory of a CMOS Chip" presents a valuable summary of the cumulative inputs into a single process step, in this case the chemical vapor deposition of undoped silicate glass (USG) [26]. As far as the author of this thesis is aware, this is the only publicly available data of its kind.

Lastly, some work on energy and materials consumption in specific semiconductor processes comes from industry sources and trade journals. [27] gives annual electrical consumption data for a variety of tools (although not as a function of production). The article also takes the tool energy analysis further by reporting figures for the embodied energy of the facility support services used by each tool as well. Using guidelines from the Semiconductor Equipment and Materials International (SEMI) trade group [28-29], Naughton applies conversion factors for the energy embodied in services

such as cooling water, vacuum, and dry air, to arrive at a value for the total annual energy required to drive a process including auxiliary inputs.

2.6 Contribution of this Thesis to the Field of Sustainability/Industrial Ecology

The overarching objective of this thesis is to present a complete assessment of the energy and material inputs into a semiconductor manufacturing facility. To the author's knowledge, no similar study has been conducted since the 1993 MCC report for an individual fab, so an update is earnestly needed. Furthermore, many studies focus on one type of input (only chemicals, for example) to the exclusion of others. This project with Analog Devices has allowed for the simultaneous assessment of energy and material inputs to give a more complete picture than has been available in the past. Furthermore, the data on how and where energy is used in the manufacturing facility and the data on specific manufacturing processes helps to round out our understanding of energy and materials use in semiconductor manufacturing.

Lastly, aside from just expanding on the paucity of data publicly available, one of the principle goals of this thesis is to present the data with greater transparency than previously available. Earlier studies lack important details on how the data was collected, processed, and normalized, making use of the data by third-parties quite challenging. With the blessing of Analog Devices, we attempt to be as lucid as possible in explaining the origins and processing of the data we have collected, with notable exceptions being specific production figures and composition of certain proprietary compounds like photoresists and developers. In this light, I invite any requests for access to the raw data or questions in general.

Chapter 3: Data Acquisition Method

3.1 Introduction

The data collection portion of this project took place between August 2007 and March 2008 at the Analog Devices, Inc. Micromachined Products Division (MPD) wafer fabrication facility in Cambridge, MA. This facility produces accelerometers, gyroscopes, and a variety of other chipsets that fall under the broader definition of MEMS. The first and second floors of this four-story building house class 100 cleanrooms and all of the fab's process tools (the second floor has some office space as well). In the basement is the facility support equipment such as the chillers, the DI water system, the wastewater treatment area, and so on. The third and fourth floors provide office space and were excluded from the data collection.

Five primary sets of data were collected from the MPD facility: 1.) information on total material inputs to the fab for the 2007 calendar year, 2.) information on the total electrical input to the fab for the 2007 calendar year, 3.) power consumption data for the fab broken down by process and tool type and including facility support equipment, 4.) the total number of wafers produced on a quarterly basis, and 5.) data on the inputs and useful outputs to specific process runs. Whereas the data for parts two and four were merely provided by the company, significant legwork was required to gather the information for parts one, three, and five.

3.2 Material Inputs, Product Wafers, and Top-Down Electrical Data

The material input data at the fab level come from a variety of different sources within the company: the process chemical information comes from purchasing records for the 2007 calendar year, the chemicals used in wastewater treatment and DI water production come from the facility's environmental consultant, and the bulk gas and water usage come from utility and bulk gas supplier bills. The difference between the amount of material purchased and that consumed by the fab is assumed to be very small in percentage terms, with the only notable cause for error being carryover of material from one year to the next.

With the exception of the data on water and bulk H₂, O₂, and N₂, the purchasing records provide information only on the total volume of each chemical or chemical mixture used annually. Determining the input mass of the various chemicals consumed in the facility required finding additional information on the concentration of the raw material and often on the composition of the compound (in the case of complex compounds like developers, etchants, and resists). For example, large canisters of phosphine actually contain only 15% phosphine or less, with nitrogen making up the balance. Such information is used in Chapter 4 to apportion material consumption by specific chemical as much as possible, with the goal of reporting just the mass of chemicals while removing that of diluents like water and nitrogen. It is also essential to note that data on the composition and density of some chemicals (e.g., certain photoresists, developers) could not be acquired owing to confidentiality concerns at

Analog Devices or the chemical manufacturers, so best estimates are used and noted accordingly.

As with input material data, managers at the Cambridge fab provided data on the number of product wafers per quarter. The data does not specify the type of product produced, number of chips produced per wafer, or yield; this information is confidential. Based on conversations with the fab managers, it is reasonable to assume that different product lines require roughly the same number of process steps, allowing all product wafers to be lumped together. Data on waste emissions and disposal were unattainable owing both to a lack of necessary instrumentation as well as the absence of a system and framework with which to make those measurements at this facility.

Total fab electrical input data for 2007 were taken from electric company billings. There are three switchgears that supply the Analog Devices fab (a switchgear is the transfer point for power from the utility grid to the end-user). Two of them primarily supply the ground, first, and second floors while the third switchgear supplies the third and fourth floors. Since the third and fourth floors consist of offices that are not directly involved in the manufacturing process, the data from that third switchgear is excluded in determining the electricity consumed by the fab. Although the first two switchgears power some non-fab loads, the size of those loads compared to the size of the fab-related loads is small (estimated to be on the order of 1%), justifying the use of switchgears 1 and 2 to determine the overall fab electricity consumption.

3.3 Bottom-Up Fab Electrical Data

The purpose of the bottom-up assessment of fab energy use was to track how and where electrical power is used in the fab with as much resolution as possible. This portion of the data collection process required the lion's share of the time invested in this project (85-90%). The original approach to monitoring power consumption within the fab was to measure every sizeable load independently. It quickly became evident that this approach would not work given that each measurement period was on the order of one day and there are hundreds of "significant" loads in the facility. Thus, the power consumption of individual tools and pieces of equipment was measured only for the largest loads in the facility. All other measurements were taken of electrical panels that power multiple pieces of equipment, "upstream" of the level of individual equipment.

For many of the facility electrical panels, the connected tools serve one primary purpose (e.g., all etch-related or all CVD-related). In panels where this is not the case and where the panel serves equipment dedicated to more than one type of process, an initial aggregate measurement was taken. At a different time, individual spot measurements were taken of the different loads in the panel to determine an approximate instantaneous power draw for each connected piece of equipment. These measurements were translated into percentages that were then used to allocate the power measured during the initial long test (the "aggregate" measurement) to the appropriate process areas. As an example, say that over a 24-hour period, Panel A has an average power draw of 10 kW. Furthermore, say that this panel has tools dedicated to dry etch, photolithography, and CVD. After the aggregate measurement is complete, Panel A is opened up and each individual tool is measured sequentially just long enough to determine its approximate instantaneous power consumption. During the spot tests, it is

found that 30% of the instantaneous power draw is going to dry-etch-related equipment, 30% to photolithography, and 40% to CVD. The power going to this panel is then allocated to the process areas as 3kW, 3kW, and 4kW respectively, for a total of 10kW. Note that during the spot test, the instantaneous power consumption would not sum to 10kW exactly. Instead it is the percentage of power consumed by each process area in the panel that is sought. This method introduces some small but unquantified error into the apportioning of power consumption within the fab to specific process areas, but it is more than sufficient for the purpose of an environmental characterization of the semiconductor manufacturing system.

The methodology applied to the bottom-up analysis required the critical assumption that the fab operates in a steady state. Although it is obvious that on short time scales this assumption is invalid, given sufficiently long measurement periods this assumption is quite reasonable. (The fab operates twenty-four hours a day and seven days a week virtually every day of the year.) Validation of this assumption will be provided later

3.4 Process Analysis

In addition to the effort to quantify the inputs and outputs to the entire manufacturing system, we are also interested in evaluating specific processes. Conducting process tests required greater coordination among Analog staff than needed for other portions of the data collection process. For each of the process runs, the power meter was first hooked up to the power supply for the process tool. A fab technician would then begin a process run with a predetermined number of wafers while the power meter logged electrical data. Material consumption data was gathered manually but in real time from the digital readouts of the various tools and corroborated with process recipe data, noting pauses in process runs so as not to overestimate material inputs. The useful output from the process (e.g., thickness of material deposited or etched) was taken from the process recipe. These data were then time-synced to provide a complete picture of the process of interest. As at the facility level, waste composition could not be measured, although total mass of waste could be estimated by subtracting material outputs from the inputs.

3.5 Measurement Equipment


Electrical measurements were made using an AEMC PowerPad® 3-phase power meter, model 3945-B. The PowerPad® is extremely capable, logging more than twenty different power metrics. In addition to basics used in this study such as current, phase-to-ground voltage, power, total energy, and power factor, it also records more rarely-used characteristics such as reactive energy, unbalanced current and voltage, apparent power, and current K factor and detects harmonics and transients. Sample data from the power meter



Figure 3.5: The AEMC PowerPad® 3945-B.

showing voltage, current, power, and power factor are below. The data is also recorded in Excel format to allow for quantitative evaluation.

The first step in an electrical recording session was to open the electrical panel associated with the tool or system to be measured and remove all protective covers. An Analog Devices electrician would then attach the proper current and voltage clamps to the current-carrying wires and live conductors, respectively. Finally, the meter would be set to recording mode, logging data points every twenty seconds to one minute depending on the length of the test. All tests were conducted without interrupting operation of the equipment, which made measurement sessions both more challenging (interlocking of tool control panels often required creative solutions to obtain readings) and dangerous. The Analog Devices lead electricians who worked with the author to set up these tests, Sandy Annese and Paul Sessler, deserve much approbation both for their professionalism and their willingness to work around obstacles to allow collection of a thorough and high-quality data set.







Figure 3.6: Sample data from the monitoring of a panel supplying primarily dry-etch equipment. From top to bottom: phase-to-phase voltage, current, power (top curve is sum of three phases), and power factor.

3.6 Power Measurements and Justification of the Steady-State Assumption

The bottom-up power analysis of the MPD fab faced an immediate quandary: how to measure all power flows within the facility simultaneously using a single power meter. To work around this hurdle, we started with the assumption that the Analog Devices fab operates in steady state. Then, over the following seven months, we systematically worked through the facility to measure all loads and build a profile of power consumption in the fab. Measurement periods lasted from four hours to longer than three days.

The validity of the steady state assumption can be established by examining a sampling of test results. Figures 3.3 through 3.6 illustrate some of the different operating power characteristics observed in tools and power panels. Figure 3.3 is a measurement of a recirculation fan that shows true constant power steady-state operation. Figure 3.4 is a measurement of an air compressor, showcasing oscillatory – but steady – behavior. Figure 3.5 demonstrates the somewhat less cyclic operation of a furnace stack containing four CVD/diffusion tubes, and figure 3.6 shows the variable power consumption of a panel that powers multiple fab tools. In each of these cases last two cases, the power characteristics appear unsteady. Yet measured over sufficiently long periods, the average power consumption of separate tests is almost always within 2% and at worst within 5%. For example, two tests of the furnace from Figure 3.5, 68.1 hours long and 23.5 hours




Figure 3.7: Recirculation fan sample illustrating constant operating power.




Figure 3.8: Air compressor sample showing cyclic, steady-state operation.




Figure 3.9: CVD/Diffusion furnace stack demonstrating quasi-cyclic operation.




Figure 3.10: Chase panel illustrating the variable operation characteristic of panels with multiple unassociated pieces of equipment attached.

long, yielded average powers of 29.008 kW and 28.796 kW respectively, a difference of 0.73%.

Table 3.1 gives average power data for each of the tests in Figures 3.3 through 3.6 divided into time intervals. The percentage error from the average value over the entire course of the test is also given. By inspection and from an examination of the rest of the data set, the minimum length of a recording required to attain a reliable long-term average power with less than 5% error is about 3-4 hours for all but the most irregular loads. It appears then that the steady state assumption steady state assumption is well justified.

Variation with Time from the Test-Long Average for a Sampling of Measurements								
Test Subject	Recirculating Fan		Compressor		CVD/Diffusion Furnace Stack		Chase Panel Serving Dry Etch	
Total Test Length (hours)	3.35		3.88		68.08		24.08	
Elapsed Time	W Avg. Over Elapsed Time	% Difference from Test Avg.	W Avg. Over Elapsed Time	% Difference from Test Avg.	W Avg. Over Elapsed Time	% Difference from Test Avg.	W Avg. Over Elapsed Time	% Difference from Test Avg.
1 Hour Avg.	43063.8	-0.023%	94552.4	-3.20%	22801.9	-21.4%	19790.5	-1.91%
2 Hour Avg.	43077.7	0.009%	96510.9	-1.20%				
3 Hour Avg.	43073.9	0.000%	97633.4	-0.05%	27609.6	-4.82%	19404.0	-3.82%
6 Hour Avg.					29347.4	1.17%	17522.4	-13.1%
12 Hour Avg.					30656.3	5.68%	18752.7	-7.05%
18 Hour Avg.					29900.7	3.08%	19683.7	-2.44%
24 Hour Avg.					29629.3	2.14%		
48 Hour Avg.					29047.7	0.14%		
Test Avg.	43073.8		97680.5		29008.1		20175.3	

Table 3.1: Examination of the Time Dependency of Measurement Power Consumption

The steady state assumption has one intractable flaw; fab power consumption varies seasonally, peaking in the humid summer months and bottoming out in the winter months. This pattern indicates that the sources of the variation are the fab systems that interact with the environment, primarily the chillers. In order to account for this variation, measurements of the components associated with the chiller system were taken at different times of year in dissimilar weather conditions. In addition to seasonal variation, the fab also shuts down for two weeks out of the year (once in June and again in December) and fab power consumption shows a corresponding dip in those months. Since the bottom-up analysis is concerned with steady-state operation, however, the fab shut downs are immaterial.

Chapter 4: Facility-Level Material and Energy Inputs to the Manufacturing System

4.1 Introduction

As stated earlier, the chief mission of this project is to characterize the inputs and useful outputs to semiconductor manufacturing as a first step to understanding the industry's environmental footprint. As opposed to other studies on semiconductor manufacturing that employ a top-down approach, the methodology used here combines both a top-down and a bottom-up analysis. The Analog Devices, Inc. (ADI) Micromachined Products Division (MPD) wafer fab in Cambridge, MA served as the test site for the data presented herein. In many respects, the MPD fab is not representative of the fabs producing the vast majority of the world's microchips. Whereas state-of-the-art wafer fabs have capacities ranging from 25,000 to 200,000 wafer starts per month (WSPM) and use 300mm (12-inch) diameter wafers, the MPD fab uses 150mm (6-inch) wafers and has a capacity of only about 3,000-4,000WSPM. The Cambridge site also lacks the extensive automated material handling systems employed in modern fabs, because production is low enough to allow for manual wafer handling. Additionally, the fab tolerates a "dirtier" clean room because MEMS features are larger than those of present-generation microprocessors and memory chips and can therefore tolerate a manufacturing environment with more particles.

In spite of these differences, the case study of the ADI MEMS facility provides valuable information about how resources are used in semiconductor manufacturing. To a first approximation, the manufacturing processes and inputs in MEMS manufacturing are equivalent to the ones used in microprocessor manufacturing. As will be discussed later, we can use the results of this study to place an upper bound on the energy required to manufacture an integrated circuit, which can then be extended to the national and international semiconductor industry.

4.2 Top-Down Assessment of Fab Energy Use

Chapter 3 summarized the two approaches to evaluate the electrical energy used in the manufacturing area of the MPD building. Data on overall fab energy use from April 2006 to December 2007 – gathered using the top-down approach – are presented in Figure 4.1. The data are presented in terms of average energy rate per month, found by dividing the total electrical energy used in the fab in a month by the amount of time in that billing cycle. The most striking feature of the graph is the seasonal variation in electricity consumption. The increase in electricity use seen during the summer is driven by the building's chiller system. During the warmer and humid summer months the chiller must work harder to maintain the temperature and humidity levels in the fab at acceptable levels. (Note also that the slight dips in July 2007 and Dec. 2006 are a consequence of fab maintenance shutdowns.) Aside from the sinusoidal variability, however, power consumption in the fab appears little changed from 2006 to 2007.

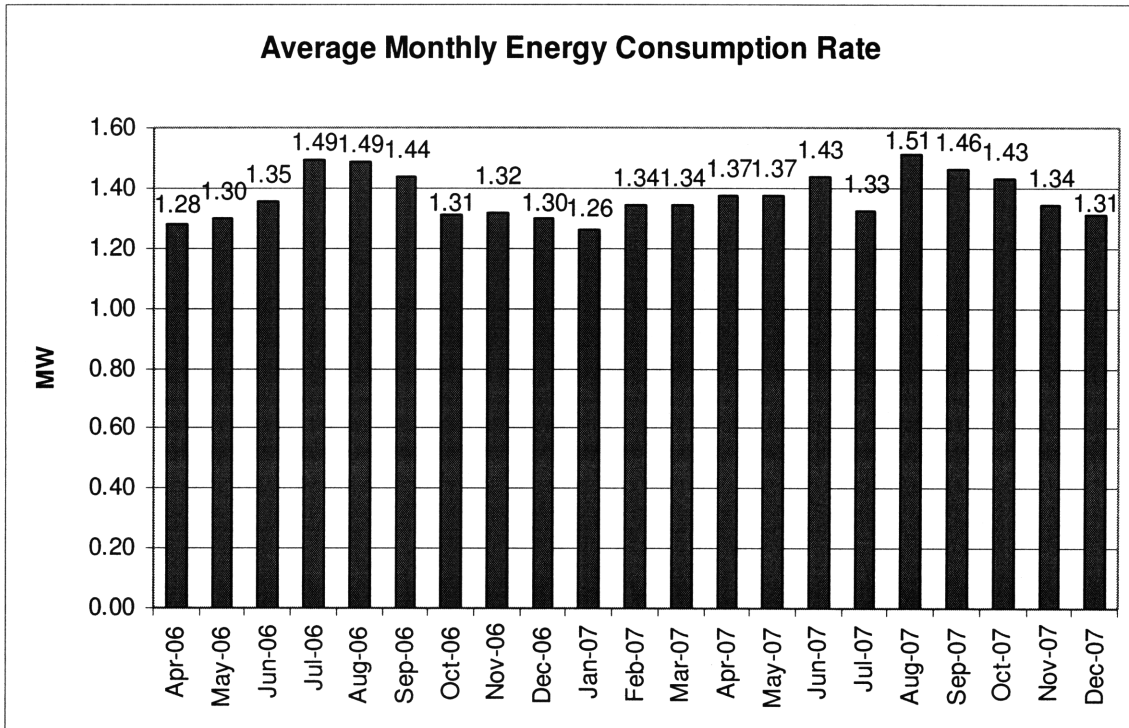


Figure 4.11: Average power consumed at the MPD fab on a monthly basis using utility data (includes lighting and some non-fab-related air conditioning)




Figure 4.12: Quarterly wafer output from the MPD fab. Although the scale has been removed to protect confidential information, relative change in wafer production is visible in this representation where the y-axis scale is linear and begins at 0.

At the same time that power consumption in the fab has remained relatively constant, it can be seen from Figure 4.2 that production has ramped considerably. The observation that absolute fab energy consumption is relatively independent of throughput is indicative of common phenomena associated with semiconductor manufacturing. The first, extrinsic to the manufacturing process, is the reality that processing equipment consumes power at a fairly constant rate regardless of whether it is or is not running product. This conclusion is borne out in the bottom-up power consumption data gathered in the fab. Typical tools idle at a very high fraction of the power consumed during processing. That is to say, the incremental increase in power consumption by a tool while it is running product is small in comparison with the baseline power consumption of the tool. Thus, as fab productivity increases and capacity utilization rates rise, the additional energy requirement is trivial.

Why is it that fab tools idle at such a high power? Many process tools operate at elevated temperatures and/or sub-atmospheric pressures. Process requirements dictate that those temperatures and pressures be kept within a very tight tolerance. Thus, manufacturers choose to maintain fab tools in a “ready” state between process runs to avoid cycling the temperature and pressure and the requirement to re-equilibrate before runs. At high utilization rates, this strategy can be sensible. If the time between process runs is too brief to allow the tool to power down without delaying the beginning of the following process, then powering down is unwise and would be even more costly than maintaining the tool in a ready state. However, fab tools often idle at a high fraction of run-time power regardless of utilization rate, meaning great opportunities for reduction in power consumption are lost. Until recently, there was little incentive for equipment manufacturers to build in the capability to carefully manage tool power consumption. In the mid-2000s, however, the desire of semiconductor manufacturers to decrease the operating cost of their tools combined with a burgeoning environmental awareness compelled Semiconductor Equipment and Materials International (SEMI) – an equipment manufacturers trade group – to develop standards for the reduction of electricity and materials use in fab tools. These standards, S22 and S23, have driven a substantial reduction in power consumption in new fab tools [28-29]. Although newer fabs have improved capability to actively control power consumption during processing, the great majority of wafer fabs use equipment that does not modulate power consumption effectively.

A second phenomenon illustrated by the apparent lack of correspondence between throughput and energy consumption, this one intrinsic to the manufacturing process, is the relatively small amount of energy actually required to activate a process in comparison to the amount of energy physically used in a process. In all manufacturing processes, the energy requirement for the process can be divided into two components: base power, which is a constant independent of throughput, and the variable power that is dependent on throughput. In a semiconductor manufacturing plant, the base power would correspond to that used by the building support systems (chillers, DI water, air compression, etc.) and that required by tools running full tilt but without any product, while the variable power would be given by the infinitesimal additional energy required to process each wafer (i.e., to heat an additional wafer in a diffusion furnace, supply the additional activation energy required for another wafer in a CVD process, etc.). Figure 4.3 shows a simple sketch of the throughput-energy relationship in semiconductor

processing compared to a higher material rate process such as casting. It is evident that as throughput increases, the specific energy consumption decreases according to the equation [14]:

$$SEC = \frac{P_0}{\dot{V}} + k \quad \text{Eq. 4.1}$$

SEC = Specific Energy Consumption (energy per unit processed (e.g. kJ/wafer))

P_0 = Base or auxiliary power (kW)

\dot{V} = Process rate (units processed per unit time (e.g. wafers/second))

k = Physically-determined constant giving rate of change of energy consumption with changing throughput (energy per unit processed)

In a manufacturing activity like casting, the variable power increases rapidly with throughput. In semiconductor manufacturing the variable power is negligible. It should be expected then that the energy intensity of semiconductor manufacturing (energy per unit output) should decline dramatically with increasing throughput, assuming that no additional tools or systems are necessary.

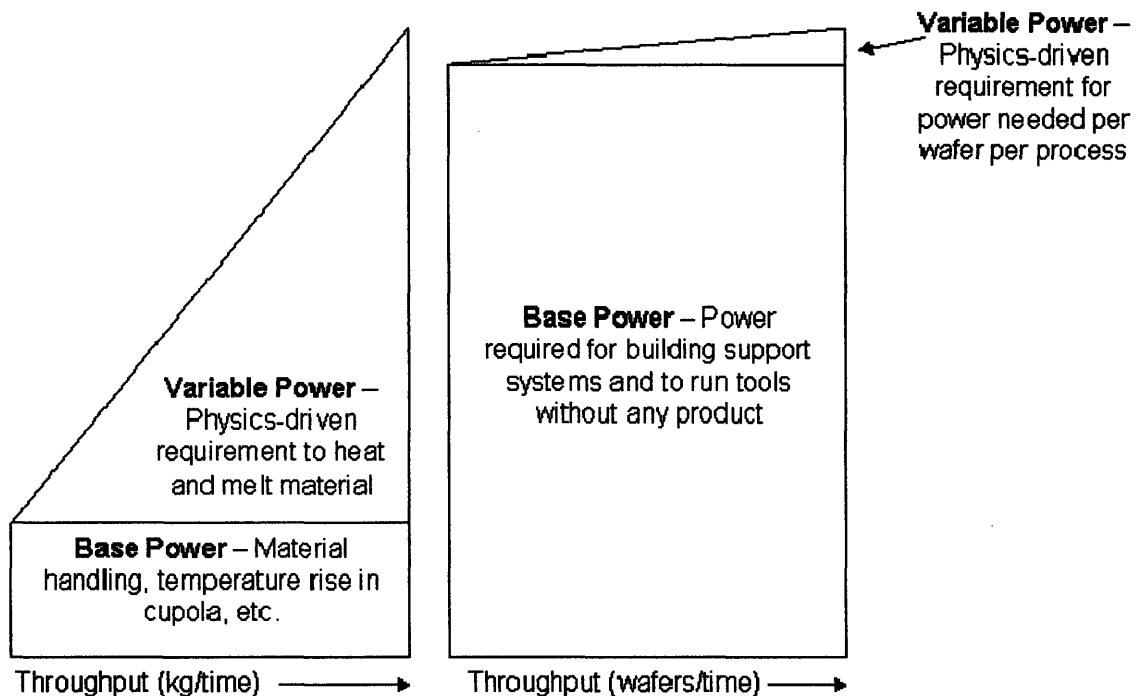


Figure 4.13: Power consumption comparison between casting (left) and semiconductor manufacturing (right).

Using the top-down data on power consumption, the energy intensity of the Analog Devices manufacturing setup can be established. In Table 4.1, this data is presented in terms of energy used per wafer processed, which can be converted to standard units of kWh/cm². The table clearly shows the anticipated decline in the energy required per wafer.

As hinted earlier, the origin of the decline in energy intensity is not so much a story of improving efficiency, but rather of increasing production. In the last year- and-a-half, production at the facility has grown at a 49% percent annual rate, or 10.4% per quarter (see Fig. 6.2). This gain in productivity has come almost entirely from increased tool utilization and not through capacity expansion. In fact, over the same time period, absolute energy consumption in the fab has grown only about 4 to 4.5% year over year, 1% per quarter. Of course, there are competing effects at play

Quarter	Electricity consumption per wafer (kWh)	Electricity consumption per square centimeter of product wafer (kWh/cm ²)
2Q06	477.4	2.70
3Q06	489.1	2.77
4Q06	360.1	2.04
1Q07	374.0	2.12
2Q07	357.2	2.02
3Q07	325.0	1.84
4Q07	271.0	1.53

Table 4.2: MPD specific electricity consumption from a top-down analysis (using utility data)

here. For example, some upgrades to more efficient pumps and motors have been made. More importantly, yearly variation in climate conditions can veil changes in tool energy consumption since such a large percentage of fab power use goes toward climate control. Nonetheless, the trend toward less energy-intensive production should not be understated.

Although the energy intensity of ADI's manufacturing plant was very high during the 2nd and 3rd quarters of 2006, since the plant began moving to higher volume production in the fourth quarter of 2006 the manufacturing energy intensity has been comparable to values reported in other studies. While the specific energy consumption has been dynamic at the MPD fab while production has increased, we will use the figure from the final quarter (1.53 kWh/cm²) as the representative energy intensity from the MPD fab. Since further planned increases in production will function to lower this figure even further, the author is confident in accepting this value as an upper bound on the energy intensity of the manufacturing system at the MPD fab.

As discussed in the introduction, the 1.53 kWh/cm² figure is only about 5% lower than the 1.6 kWh/cm² reported in the MCC study. Furthermore, in light of the results shown here, the value used by Williams et. al. in the LCA of a microchip, 1.5 kWh/cm², seems to be an appropriate upper-bound selection for the manufacture of six-inch wafers.

4.3 Relationship to State-of-the-Art Fabs

Previous studies on the resource intensity of semiconductor manufacturing have tended to be fairly opaque, particularly in comparison with the present one. Moreover, to the author's knowledge there are no assessments of the high-volume, 300 mm wafer fabs that are becoming increasingly mainstream. It is therefore rather difficult to identify specific differences in energy use patterns between those found in this study, those found in others, and what is likely to be found in the 300 mm fabs. Comparing the MPD fab with modern large volume memory and microprocessor fabs, however, suggests four overriding characteristics that would affect the energy intensity of large-scale fabs

compared to the MPD: production volume, number of processing steps, age of fab tools, and feature-size and cleanliness requirements.

The production volume of a fab obviously plays an important part in determining the energy intensity of its manufacturing setup. Once again, the opacity of previous studies makes it difficult to compare the production volume of the MPD fab with that of other studies, so it is not a simple task to estimate how energy intensity varies between fabs of different production volumes. Following the explanation of the impact of idle/base and variable power on specific energy consumption from Section 4.2 (see Figure 4.3), it follows that for semiconductor manufacturing, in which the variable energy component is so small compared to the base component, specific energy consumption should continue to decrease with increasing throughput. Of course, this relationship only holds up until capacity additions are required (more tools, larger air handling systems, more DI water capacity, etc.). Nonetheless, it is logical to assume that one of the important benefits of higher-volume production would be a decrease in the energy intensity per unit product, particularly in comparison with the energy intensity of the ADI fab.

A second factor suggesting that the energy intensity at the MPD fab is an upper bound for semiconductor manufacturing has to do with the type of products manufactured. Although feature sizes on MEMS devices do not rival the state-of-the-art seen in microprocessors, MEMS devices are very complex to manufacture owing to the integration of both mechanical and electronic components. As a result, the average sensor produced at the MPD fab undergoes some 450+ process steps, as compared to the 300 to 350 that are required by microprocessors and even fewer required by memory chips. The greater number of process steps certainly increases the energy intensity of MEMS manufacturing.

A third factor strongly affecting the energy intensity of the MPD fab is the fact that it uses older production methods – and older tools – than modern fabs. Since MEMS technology does not require the cutting edge, sub-100nm linewidths and 300mm wafers characteristic of state-of-the-art fabs, the MPD fab can accomplish its mission using vintage tools. Although this earlier-generation equipment is cheaper upfront, it does not often have the energy-saving features that modern tools have. The fact that the MPD fab handles 150mm wafers instead of 300mm wafers also reduces the apparent efficiency of the tools by reducing the area throughput capacity of the tools. The cumulative effect is to increase the measured energy intensity of their manufacturing plant.

The three previous factors – production volume, number of process steps, and tool age – should result in higher energy intensities for the MPD fab than state-of-the-art fabs. The fourth factor, dimensional and cleanliness requirements, should have the opposite effect. In general, the MPD fab would be termed “dirty” by modern semiconductor processing standards. The class 100 cleanroom at the MPD fab pales in comparison to the class 10 or class 1 cleanrooms typical of memory and microprocessor fabs. (The fab class number gives the maximum number of particles 0.5 μm or larger in size per cubic foot, although class 10 and class 1 cleanrooms also have limits on the number of smaller particles that can be accepted.) More stringent cleanliness requirements demand better air filtration and thus higher power fans and filters. The finer dimensional requirements for microprocessor and memory chips also add to the energy requirements of production

– particularly during lithography – as pattern definition requires higher energy light sources and more complex alignment systems.

In summary, the cumulative effect of larger production volumes, fewer process steps, and more efficient tools most likely outweighs the countervailing effect of finer dimensional and cleanliness requirements. Therefore, the Analog Devices data can be considered an upper-bound for energy intensity for semiconductor manufacturing. This result is quite useful since combined with an estimate of the number of wafers processed worldwide, we can then estimate an upper bound for the energy required by the global semiconductor manufacturing industry. In 2007, $\sim 51,000,000,000 \text{ cm}^2$ of silicon wafer was processed worldwide by the semiconductor industry [30]. Using the top-down result, it can be estimated that the world semiconductor industry requires some 8×10^{10} kWh of electrical energy per year in front-end semiconductor manufacturing alone. That is ~ 1.5 times as much energy as Massachusetts uses per year!

4.4 Top-Down Assessment of Material Inputs to the Manufacturing Facility

The collection of data on material inputs to the fab are discussed in Chapter 3. The material inputs to the MPD manufacturing system for 2007 are presented in Table 4.2, normalized per square centimeters of wafer surface processed during 2007. As much as possible, the data have been distilled from their original form to shield information that cannot be disclosed. The central column indicates the mass of a given chemical that enters the manufacturing facility including any diluting agent (water for liquids or nitrogen for gases). In this column, multiple concentrations of a single chemical type are lumped together. The far right column gives only the amount of pure chemical input. For example, if 11 kg of 10:1 HF were used over the course of 2007 for wafer processing, one kg of HF would be reported while the other ten kilograms would be recorded as water. All of the different sources of HF (i.e., HF at different concentrations) can then be summed to yield a cumulative HF use for the year. Similarly, most deposition gases (e.g., B_2H_3 , PH_3) are heavily diluted in nitrogen, so the total volume of etch gas used is determined by multiplying the mass fraction of the pure etchant by the total amount of the etch gas/nitrogen mixture used in 2007. One exception is in the reporting of complex compounds such as developers and photoresists for which compositions are not often available (compositions are considered proprietary information by the chemical manufacturer). These chemicals are lumped together and reported simply as “photoresist” or “developer,” which is necessary also to protect information that Analog Devices deems confidential.

A few additional notes: 1.) Following the example of Williams et. al., bulk gases (nitrogen, oxygen, argon, helium, and hydrogen) and water are recorded separately because they are used in such large volumes that their inputs swamp that of other material inputs to the fab. These ubiquitous inputs are commonly used as carrier gases or, more so in the case of water, for cleaning. Although there is certainly an environmental impact associated with the use of these “chemicals,” because they are prevalent in the environment the environmental burden imposed by their use is more a function of the energy used to purify them. (One prominent exception is the case of water usage in water poor areas.) 2.) Steam is primarily used in winter to heat and humidify fab air, but

instead of considering it as a material input into the fab, it is viewed instead as an energy input since very little steam is “consumed” in the manufacturing plant (by condensation).

Analog Devices 2007 Materials Usage		
Compound	g/cm ²	
Liquid Chemicals	w/ water balance	w/o water balance
Acetone		8.65E-02
Aqueous Ammonia	0.467	0.135
HCl	0.487	0.185
HF	0.896	0.106
H ₂ O ₂	3.553	1.066
Isopropanol		1.878
H ₃ PO ₄	0.184	0.157
HNO ₃	0.351	0.246
H ₂ SO ₄	1.161	1.114
HMDS		2.93E-03
Phosphorus Oxychloride		3.41E-04
1,2-Dichloroethylene		2.36E-02
FC-84 (perfluorocompound)		1.74E-02
FC-40 (perfluorocompound)		4.53E-03
Buffered Etchants	1.766	
Mixed Acid Etchants	0.588	
Post-Metal Etchants	1.753	
Developers	13.925	
Resists		0.668
Deposition Gases	w/ nitrogen balance	w/o nitrogen balance
NH ₃		1.37E-02
SiH ₂ Cl ₂		3.33E-03
N ₂ O		0.445
SiH ₄		7.00E-02
AsH ₃		6.06E-05
B ₂ H ₆	1.78E-03	1.02E-04
PH ₃	6.25E-03	4.52E-04
Etchant Gases		
CF ₄		3.85E-02
CHF ₃		1.92E-02
C ₄ F ₈		7.01E-02
SF ₆		0.187
C ₂ F ₆		0.163
HBr		4.12E-02
BF ₃		8.24E-05
Cl ₂		2.75E-02
DI Water Production and Wastewater Treatment	w/ water balance	w/o water balance
HCl	1.82	0.633
NaOH	4.03	2.016

Table 4.3: Material Inputs to the Analog Devices MPD fab for 2007. The central column gives the total mass of material including water or nitrogen balance weight, whereas the rightmost column is distilled to show only the mass of concentrated chemical used.

Bulk Gases		
N ₂		742.5
H ₂		6.02E-02
O ₂		11.94
Ar		7.36E-02
He		6.40E-03
Water		
H ₂ O		29710.0

Table 4.2 Cont.

4.4.1 Chemical Usage Comparison with Previous Studies

In the literature, the only relatively comprehensive previous account of the chemicals used in the manufacturing of a semiconductor was published in [20]. The total mass of material consumed in a fab can be assessed in different ways. The Williams analysis excludes water and bulk gas consumption, but uses the diluted weight of input chemicals and gases and then normalizes the results per square centimeter of product wafer. Table 4.3 shows a comparison between the values in the Williams study and those of the present study per Williams's methodology (using masses from Table 4.2 with water weight or nitrogen balance included).

The lack of explanation accompanying the Williams chemical data makes comparison difficult, but a few meaningful comparisons can be made. First off, it appears that much of the difference in the mass of chemicals comes from the disparity in sulfuric acid (H₂SO₄) consumption. The MPD fab has led a concerted effort to reduce sulfuric acid use in its manufacturing setup, which may be one reason why MPD uses much less than the European fab from which the Williams data comes. Similarly for the other acids, the Williams fab data indicates a greater use of aqueous etchants than the MPD fab while the opposite is true for dry etchants. This finding suggests that the fab from which the Williams data comes has a greater proportion of wet etch processes and a smaller proportion of dry etch processes than the MPD fab. Since liquid etchants are denser than dry etchants, it follows that the Williams data should lead to a larger mass of chemicals than the MPD data. The data for the Williams fab also indicates a greater mass of sodium hydroxide (NaOH) used in wastewater neutralization than in the MPD fab. One result which the author of this study cannot explain is the greater amount of developer used in the MPD fab but the smaller amount of resist; these results seem contradictory.

Comparison Between MPD Fab Materials Usage and the Williams Study Results

Compound	g/cm ²		Difference
	Williams	Analog Devices	
Liquid Chemicals			
Acetone	5.54E-01	8.65E-02	0.467
Aqueous Ammonia	3.616	0.467	3.149
HCl	3.031	0.487	2.544
HF	4.128	0.896	3.232
H ₂ O ₂	4.430	3.553	0.877
Isopropanol	2.02	1.878	0.142
H ₃ PO ₄	2.410	0.184	2.226
HNO ₃	1.190	0.351	0.839
H ₂ SO ₄	7.850	1.161	6.689
HMDS	2.20E-02	2.93E-03	1.91E-02
Phosphorus Oxychloride	NA	3.41E-04	
1,2-Dichloroethylene	NA	2.36E-02	
FC-84 (perfluorocompound)	NA	1.74E-02	
FC-40 (perfluorocompound)	NA	4.53E-03	
Buffered Etchants	NA	1.766	
Mixed Acid Etchants	NA	0.588	
Post-Metal Etchants	NA	1.753	
Developers	5.730	13.925	-8.195
Resists	1.48	0.668	0.812
slurry	0.286	NA	
Deposition Gases			
NH ₃	1.20E-02	1.37E-02	-1.74E-03
SiH ₂ Cl ₂	1.40E-03	3.33E-03	-1.93E-03
N ₂ O	0.072	0.445	-0.373
SiH ₄	7.80E-03	7.00E-02	-6.22E-02
AsH ₃	4.30E-06	6.06E-05	-5.63E-05
B ₂ H ₆	4.30E-06	1.78E-03	-1.77E-03
PH ₃	1.70E-05	6.25E-03	-6.23E-03
WF ₆	4.30E-04	NA	
Etchant Gases			
CF ₄	3.00E-02	3.85E-02	-8.48E-03
CHF ₃	3.10E-02	1.92E-02	1.18E-02
C ₄ F ₈	NA	7.01E-02	
SF ₆	6.50E-03	0.187	-0.180
C ₂ F ₆	5.00E-02	0.163	-0.113
HBr	2.20E-03	4.12E-02	-3.90E-02
BF ₃	3.50E-05	8.24E-05	-4.74E-05
Cl ₂	4.80E-03	2.75E-02	-2.27E-02
BCl ₃	8.70E-03	NA	
NF ₃	2.30E-03	NA	
DI Water Production and Wastewater Treatment			
HCl	NA	1.82	
NaOH	8.25	4.03	4.22
Total	45.2	34.8	10.5

4.5 Total Mass of Materials Required in the Front-End Manufacturing of a MEMS Chip

A driving motivation for this study has been to develop an estimate of the amount of material required in the production of a microchip, with the goal of corroborating the findings of Eric Williams, Robert Ayres, and Miriam Heller's paper "The 1.7 Kilogram Microchip." In it, they conclude that a 2-gram microchip requires about 1700 grams of material inputs over the course of its lifetime [20]. The particulars of their analysis are discussed in Chapter 2.

The 1.7kg value found in the Williams study applies to the inputs required for materials preparation, fabrication, and use of a microchip corresponding to a silicon wafer area of 1.6 cm^2 . The mass value is broken down into two components. The first and largest is the mass of fossil fuels required to supply the electrical and non-electrical energy inputs to the chip over its lifetime. This is found by multiplying the total electrical energy used over the lifetime of the chip by 320 grams of fossil fuels per kWh of electricity (a European average), yielding a total mass of 1640 grams of fossil fuel inputs. As noted earlier, the Williams analysis includes not only the fabrication process that takes place in the wafer fab, but also includes the energy needed to produce the silicon wafers and chemicals going into the manufacturing system, the energy used during assembly, and the power during use. The present study of the Analog Devices fab is concerned principally with the front-end manufacturing of the chip to the exclusion of other inputs to the life cycle. Thus, for the purposes of comparison, we take Williams' assessment that 48% of lifecycle energy use goes into the production phase, which yields about 787 grams of fossil fuels for the processing of 1.6 cm^2 of silicon wafer area. The second component of Williams's 1.7 kg value is the mass of chemicals used only in the fabrication of the integrated circuit (not in silicon wafer production or chemical production as is the case with the energy value). It is estimated that 72 grams of chemicals are used per chip (excluding nitrogen and water inputs). This leads to 859 grams of total materials consumed in the manufacture of a 1.6 cm^2 integrated circuit. Normalized to energy per square centimeter, the Williams estimate becomes: 492 grams of fossil fuels, 45 grams of material inputs, and 537 grams of total mass inputs required per square centimeter of silicon wafer processed.

Whereas the Williams study estimated that 45.2 grams of chemicals are required in the front-end manufacture of a microchip, the data gathered at Analog Devices gives a mass of chemicals of **34.8 grams/cm²**. Using the specific energy consumption determined in the top-down analysis (1.53 J/cm^2) and multiplying by the Williams energy-to-fossil-fuel conversion factor of 320g/kWh yields 491 grams of fossil fuels consumed per chip. (Although the 320 g/kWh factor is determined from the European energy mix, it is used here to allow a one-to-one comparison.) The total mass of materials consumed in the Analog Devices MPD manufacturing process is $491 + 34.8 = 525 \text{ grams/cm}^2$.

The results presented here correspond remarkably well with those of the Williams study, with only a 2% disagreement between the two. In the field of sustainability, such close agreement is vanishingly rare. Indeed, a closer look at the data suggests that the similarity of the two results might be deceptive. To begin, it is evident that the stated material usage is very sensitive to the estimated energy use, since the total mass is

dominated by that of fossil fuels for electricity generation. A small variation in specific energy use can lead to significant changes in the calculated mass. For example, if for the MPD data instead of using the 4Q07 figure of 1.53 kWh/cm² we used the entire 2007 output-weighted average of 1.851 kWh/cm², the mass of fossil fuels required becomes 592 grams and the total mass becomes 592+34.8 = 627 grams, or 17% higher than the result from the Williams study. Similarly, it appears that in future quarters the energy intensity of manufacturing at the MPD fab will continue to decrease as production ramps upward, driving down the amount of material required per square centimeter of chip. Additionally, an increase or decrease in the fossil fuel mass required per kWh of electricity produced will likewise affect the mass of material used in the manufacturing process. However, for the reasons stated in Section 4.3, it seems reasonable to accept the mass of materials required per square centimeter of silicon processed at the MPD fab as an upper bound for semiconductor manufacturing.

A second reason that the Williams/MPD comparison can be deceptive is because of the insensitivity to the mass of chemicals consumed in the production facility. For example, if one were to elect not to report the diluted weights of the input chemicals but instead only to report the concentrated value (i.e., exclude the weight of water and nitrogen diluents where possible), the mass of chemical used at the MPD fab becomes 27.5 g/cm². This 20% reduction in input chemical weight (from 34.8 g/cm²) translates only to a 1% reduction in total mass inputs (from 525 g/cm²). Furthermore, as is evident from Tables 2 and 3, the mass of gaseous chemicals consumed are negligible compared to the mass of liquid compounds. Therefore, it seems that reporting the mass inputs to semiconductor manufacturing as a lumped sum may actually be blurring what is a more important story about trends in chemical usage in the industry: increases or decreases in the use of specific chemicals.

4.6 Bottom-Up Assessment of Fab Energy Use

To this point in the discussion, the Analog Devices Micromachined Products Division fab has been treated as a black box, with material and energy inputs on one side and useful and waste outputs on the other. Such a broad view has limited utility in helping understand how and where energy is used – and how it can be conserved – within the facility. A comprehensive bottom-up analysis of fab power consumption was conducted in an attempt to tackle this issue of where and how energy is used in the fab. The fabwide electrical power consumption distribution is shown in Figure 4.4. (Please see Chapter 3 of this report for an explanation of how the data were gathered.) One of the more surprising results is that the majority of the electrical power used in the fab goes not to the tools and processing equipment, but instead to the so-called “facility support equipment” (e.g., the chiller system, air handling, deionized water production, air compression, ventilation, etc.). The chiller system is the biggest energy hog in the MPD fab – accounting for about a quarter of energy consumption at the facility – but the air handling, DI water production, and air compression systems are significant energy users as well. Suffice it to say that for all the attention given to the energy intensity of specific semiconductor processes, an equally promising starting point when assessing how to make semiconductor manufacturing more environmentally benign would be these support systems.

Analog Devices MPD Fab Power Consumption Distribution (kW)





Figure 4.14: Electricity Use Breakdown of the Analog Devices MPD Wafer Fab (Excludes Lighting, non-Fab-Related Air Handling)

Toshiba Fab Power Consumption Breakdown



LBNL Fab Power Consumption Breakdown




Figure 4.15: Power Consumption Break Down Data from a Toshiba Fab and from data collected at 12 different fabs by Lawrence Berkeley National Lab [31-32].

These results can be compared with statistics from two other sources that have reported on the electrical energy end-use of wafer fabs. The results of the present study correspond well with data from a Toshiba fab (Figure 4.5). The specifics of the study (e.g., product type, production volume) are not available, but the finding that between 43 and 49% of input energy goes to fab tools and related manufacturing equipment is in excellent agreement with the Analog Devices figure of 43%, particularly given the number of variables involved. Similarly, the air conditioning and refrigeration pieces of the Toshiba pie account for 32% of power consumption. The two corresponding sectors of the Analog Devices pie are the chillers and air handling which comprise 37% of the Analog power pie. Once again, the agreement is quite good, with much of the difference likely attributable to differences in the climate in which the fabs are located or to different methods of accounting.

The agreement between the MPD fab data and those from Lawrence Berkeley National Lab do not agree quite as well, until it is recognized that neither the Analog Devices fab nor the Toshiba fab have a dedicated liquid nitrogen plant (instead, nitrogen is delivered on a frequent basis by an external supplier). Also note that the energy required for cooling water in the MPD and Toshiba fabs is supplied by the chiller system, whereas that receives a separate heading in the LBNL data. Reapportioning the 11% of fab power dedicated to the liquid nitrogen plant and the process cooling water to the other areas of the fab brings the data from LBNL into better accord with that from the MPD fab. Although three data sets do not constitute a trend, the fact that this study and the other two available in literature correlate well suggests that the Analog Devices study is a sound assessment of fab electrical energy distribution, and can – with caution – be applied more globally.




Figure 4.16: Power breakdown by process area.

The manufacturing tools certainly account for a large portion of the energy consumed in a semiconductor plant and offer some of the more interesting targets for energy reduction during operation. Figure 4.6 shows the fab tools portion of Figure 4.4 expanded to detail how power is distributed amongst the various process areas. For more details on how the energy used in each process area was determined, please see Chapter 3.

The rather surprising result that a large fraction (nearly 1/3) of the energy used by fab tools goes to cap and trim seems to be unique to the MPD fab. Capping refers to a pre-back-end series of steps that prepare the wafer for shipping to another facility for final manufacturing and packaging. Although the equipment associated with capping uses some electrical energy, the lion's share of this portion of the pie comes from the "trim" section. (Unfortunately, owing to facility restrictions, the trim data cannot be parsed from the capping data.) Trimming is used to fine-tune the MEMS devices as one of the final steps in the manufacturing chain before the wafer is capped. It is a relatively slow, serial process in which a laser ablates small amounts of material from the wafer surface, hence the large energy use. Since the trim process is fairly unique to MEMS manufacturing, little work has been done to improve throughput and energy efficiency. These tools would make good targets for such an effort at energy reduction.

The energy intensity of CVD has been documented in the literature [33], so it is curious that dry etching accounts for a greater portion of energy use at the MPD fab than CVD or diffusion and almost as much as CVD and diffusion combined. In justification of this result, we again turn to a process rate argument. Whereas Analog Devices's CVD tubes can run better than a hundred wafers at a time, virtually all of the dry etch steps are serial processes in which only one wafer is operated upon at a given moment. Thus, although it is true that CVD furnaces operate at very high temperatures, vacuum pressures (in the case of LPCVD and PECVD), and use plasmas (in the case of PECVD), the much larger number of dry etch tools compared to CVD furnaces results in dramatically higher energy use by the dry etch process area.

4.7 Comparison of Bottom-Up and Top-Down Results


The bottom-up and top-down energy assessments yielded encouragingly similar results for total fab power. At 1.46MW, the bottom-up result is only about 3% off from the peak top-down monthly value of 1.51MW and is practically equivalent to the average power consumption during warm months. The author speculates that one reason that these results are so similar is that the chiller system was monitored during the summer at times when it was working very hard – that is, on days that were either quite humid or quite hot (but dry). Since the chiller system is not only the largest power consumer but also because its power draw varies substantially with changing weather conditions, taking measurements on the chiller during the humid and hot weather conditions amounts to taking an upper-bound measurement on this system and thus the fab as a whole.

Entering this project, the expectation was that the sum of the power draws measured in the bottom-up analysis would be lower than that for the top-down analysis. The bottom-up analysis did not include energy required for lighting, as lighting was considered to be intrinsic to the manufacturing process. Also, certain other extraneous loads, such as the basement air ventilation system, were not included in the bottom-up

analysis but contributed to the top-down numbers. Given these factors, the similarity of the results between top-down and bottom up suggests not only that the bottom-up study was thorough but also that the steady-state assumption is valid.

4.8 Seasonal Variation in Energy Consumption

As mentioned earlier in this chapter, the average power rate varies seasonally. The primary reason for this is that the chiller system need not work as hard in colder and drier months. As illustrated in Figure 4.7, the energy going to the chiller system in winter drops roughly by a factor of two, reducing its relative contribution to fab energy consumption to 17-20% and raising the relative contribution of the other areas in the fab. In this case too we see that the total power consumption is in the same range as that for the top-down data during the winter months.



Chapter 5: Analysis of Semiconductor Manufacturing Processes

5.1 Introduction

The energy intensity of semiconductor manufacturing processes, where intensity is defined as energy required per unit volume of material processed, has drawn increasing attention, particularly as the industry continues to grow at a rapid clip. It is not only the energy intensity of these processes that are of concern, however, but also the materials intensity. As illustrated in earlier works and suggested in the preceding chapter by the fab-wide data, semiconductor manufacturing is extremely wasteful compared with other industries in terms of the amount of material incorporated in the final product compared to the total amount of input material. “Wasteful” is a loaded term. The intention here is not to vilify semiconductor manufacturers, who make a highly ordered, fantastically complex product from much less valuable raw starting materials. Instead, it is to determine how the excess material and energy not being incorporated into the final product can be conserved, to the benefit of both the semiconductor manufacturers (through reduced costs) and the environment.

A method is needed for the evaluation of these processes, and manufacturing systems more generally, that can be used both to quantify the interaction of the process with the environment as well as to measure – or even better, to suggest changes for – process improvement. Life Cycle Analysis (LCA) and its many variations (Environmental Input-Output LCA (EIO LCA), Streamlined LCA, etc.) are useful for looking at the impact (environmental and otherwise) of man-made products from a global perspective. However, in this thesis we are more concerned with a single phase of the life cycle of semiconductors, so although the methods of LCA are useful we choose not to adopt the entire framework. Input-output analytical methods are another evaluation tool used in sustainability studies. I/O analysis accounts for all the inputs and outputs to a system; the Chapter 4 accounting of the fab-wide materials and energy use at Analog Devices is an example of I/O analysis, albeit lacking information on the waste output stream. What is needed is an analytical method that not only accounts for the inputs and outputs to a system or process, but also captures the quality or value of those inputs and outputs, and for that we turn to thermodynamics.

5.2 A Thermodynamic Model of Manufacturing Systems

Manufacturing can be conceived as a complex system featuring a sequence of thermodynamic processes in the same manner as proposed by Gyftopoulos for an energy conversion/materials processing system [34]. Each stage in the system has energy and mass inputs, the material inputs of which are either taken from the environment directly as raw materials or indirectly as material processed in an earlier stage. The useful output from a given stage is then passed on to become the input for a succeeding stage, until a final product is produced which is sold to consumers. Each manufacturing stage inevitably generates entropy and a waste stream. Waste material is discarded to the

environment (although in some cases, it can be recycled). In addition, low-quality energy flows (e.g., waste heat) are dissipated to the environment.

Figure 5.1 depicts this generalized model of a manufacturing system. The manufacturing subsystem (Ω_{MF}) is powered by a work rate fed from an energy conversion subsystem (Ω_{ECMF}). The input materials are shown to be manipulated upstream in the materials processing subsystem (Ω_{MA}), which also has its own energy conversion subsystem (Ω_{ECMA}). This network representation can be infinitely expanded to encompass ever more complex and detailed manufacturing systems and/or to include all elements of the life cycle of the product.

At every stage, the manufacturing system interacts with the environment (shown here at reference pressure p_0 and reference temperature T_0) either by dumping low-quality heat (accompanied by corresponding entropy flow) or by sourcing and sinking the materials used in the process. The characteristics of this environmental reference state are critical to the calculation of the thermodynamic interactions taking place in the process. By convention, the reference pressure is selected as 101.3 kPa and the reference temperature as 298.15K.

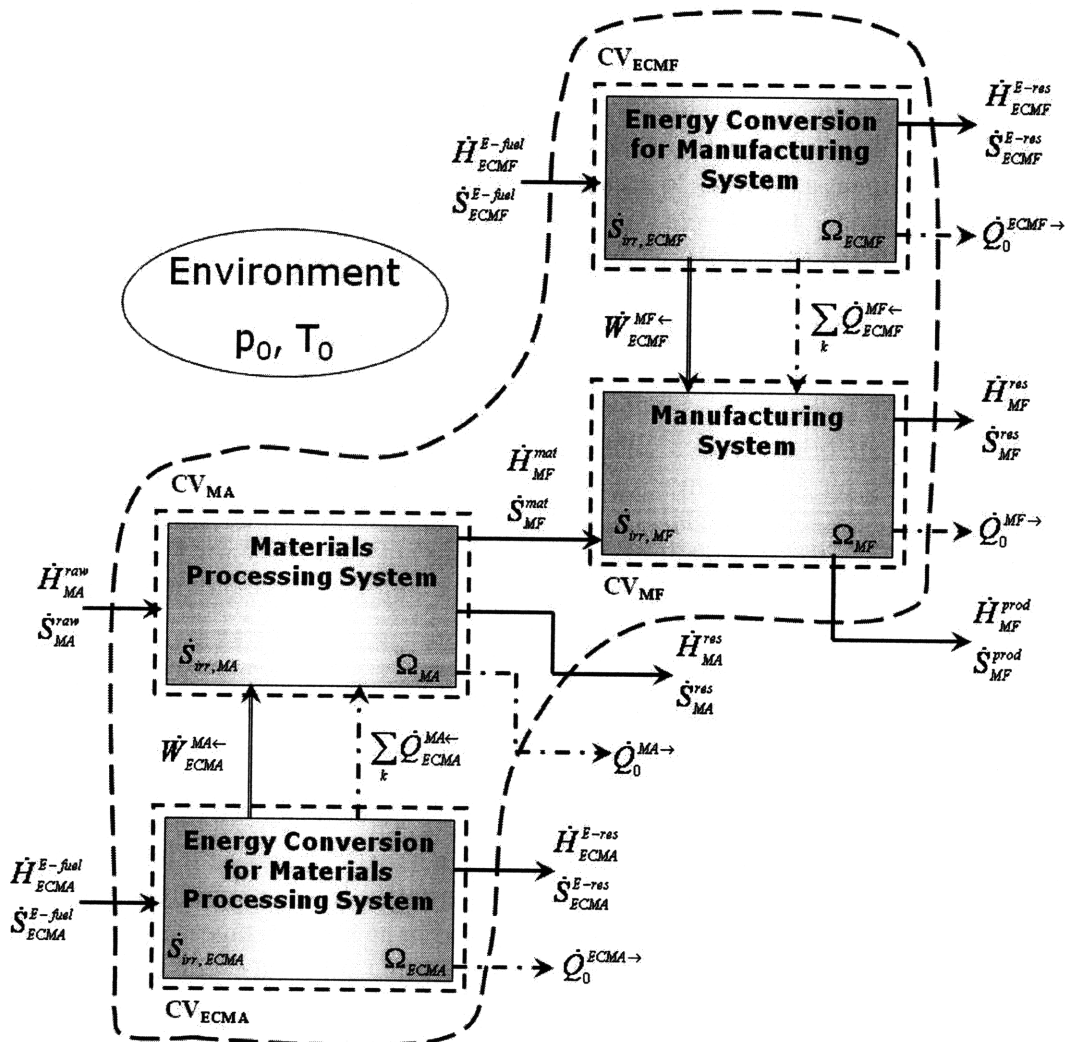


Figure 4.17: General Manufacturing Model

This set of subsystems can be completely described in thermodynamic terms by formulating mass, energy, and entropy balances [35]. Beginning with the manufacturing system Ω_{MF} , we have three basic equations:

Mass Balance:

$$\frac{dm_{MF}}{dt} = \left(\sum_{i=1} \dot{N}_{i,in} M_i \right)_{MF} - \left(\sum_{i=1} \dot{N}_{i,out} M_i \right)_{MF} \quad \text{Eq. 5.1}$$

where \dot{N}_i is the number of moles of the i^{th} component entering or leaving the system and M_i is the molar mass of that component.

Energy Balance:

$$\frac{dE_{MF}}{dt} = \sum_i \dot{Q}_{ECMF}^{MF\leftarrow} - \dot{Q}_0^{MF\rightarrow} + \dot{W}_{ECMF}^{MF\leftarrow} + \dot{H}_{MF}^{mat} - \dot{H}_{MF}^{prod} - \dot{H}_{MF}^{res} \quad \text{Eq. 5.2}$$

where \dot{H} gives the sum of the enthalpy rates of all materials, products, and residue bulk flows. Note that for reacting systems, $(\dot{H}_i / \dot{m}_i) = h_i = h_{f,i}^o + \Delta h_i$, where $h_{f,i}^o$ represents the specific enthalpy of formation and Δh_i accounts for the change in the enthalpy from the reference temperature and pressure to the actual temperature and pressure of the i^{th} component.

Entropy Balance:

$$\frac{dS_{MF}}{dt} = \sum_i \frac{\dot{Q}_{ECMF}^{MF\leftarrow}}{T_i} - \frac{\dot{Q}_0^{MF\rightarrow}}{T_0} + \dot{S}_{MF}^{mat} - \dot{S}_{MF}^{prod} - \dot{S}_{MF}^{res} + \dot{S}_{irr,MF} \quad \text{Eq. 5.3}$$

where \dot{S}_i gives the sum of the entropy rates of all components. Note that in the entropy balance, the entropy of each bulk flow (e.g., reacting systems) can be defined in terms of specific absolute entropy s_i^o and specific entropy difference evaluated at the given temperature and pressure and at the reference values, Δs_i .

In steady state, the terms to the left of the equality sign are zero, so the next step is to solve the energy and entropy balance equations for \dot{Q}_0 and then equate them. This yields an expression for the work requirement for the manufacturing process:

$$\dot{W}_{ECMF}^{MF\leftarrow} = ((\dot{H}_{MF}^{prod} + \dot{H}_{MF}^{res}) - \dot{H}_{MF}^{mat}) - T_0 ((\dot{S}_{MF}^{prod} + \dot{S}_{MF}^{res}) - \dot{S}_{MF}^{mat}) - \sum_{i>0} \left(1 - \frac{T_0}{T_i} \right) \dot{Q}_{ECMF}^{MF\leftarrow} + T_0 \dot{S}_{irr,MF} \quad \text{Eq. 5.4}$$

The quantity H-TS appears often in thermodynamic analysis and is referred to as the free energy. When related to the reference state (p_0, T_0) the free energy becomes equivalent to exergy (B). That is, $B = (H-T_0S)_{P,T} - (H-T_0S)_{P_0,T_0}$. Exergy represents the

maximum amount of work that could be extracted from a system as it is reversibly brought to equilibrium with a well-defined environmental reference state (hence the requirement that the exergy reference be p_0 and T_0). It is very important to note that the material-related terms in (4) may be split into contributions that account for the so-called physical and chemical exergies (see below). For a further discussion of exergy and its meaning, see [35].

The physical exergy is that portion of the exergy that can be extracted by bringing a given state to a restricted dead state at a reference temperature and pressure (T_0, p_0). Therefore, we generalize (5.4) with:

$$\dot{B}_i^{ph} = \dot{m}_i [(h_i - h_{i,o}) - T_0 (s_i - s_{i,o})] \quad \text{Eq. 4a}$$

is known as the *physical exergy* rate of the bulk flow system, defined as the availability between the system in its given state and the system in equilibrium with its surroundings (T_0, p_0).

Alternately, the terms $h_{f,i}^o - T_0 s_i^o$ are known as *chemical exergies* ($b_i^{ch} = \dot{B}_i^{ch} / \dot{m}_i$) and represent the additional potential that can be extracted by bringing the chemical potentials (μ_i) of a given system to equilibrium with its surroundings in the ultimate dead state (T_0, p_0, μ_0). Note that chemical exergies are calculated by using enthalpy of formation and absolute entropies. In addition to requiring a reference temperature and pressure, the definition of chemical exergies also requires a reference state with a specified chemical composition. This reference is typically taken to be representative of the compounds in the earth's upper crust, atmosphere, and oceans. In this article, exergy values are calculated using the Szargut reference environment [36].

We simplify (4) with:

$$\begin{aligned} \dot{W}_{ECMF}^{MF\leftarrow} = & ((\dot{B}_{MF}^{prod} + \dot{B}_{MF}^{res}) - \dot{B}_{MF}^{mat})^{ph} + \left(\sum_{i=1}^n b_{0,i}^{ch} \dot{m}_i \right)_{MF}^{prod} \\ & + \left(\sum_{i=1}^n b_{0,i}^{ch} \dot{m}_i \right)_{MF}^{res} - \left(\sum_{i=1}^n b_{0,i}^{ch} \dot{m}_i \right)_{MF}^{mat} - \sum_{i>0} \left(1 - \frac{T_0}{T_i} \right) \dot{Q}_{ECMF}^{MF\leftarrow} + T_0 \dot{S}_{irr,MF} \end{aligned} \quad \text{Eq. 5.5}$$

Using the same analysis for the system Ω_{ECMF} yields:

$$\dot{W}_{ECMF}^{MF\leftarrow} = (\dot{B}_{ECMF}^{E-fuel} - \dot{B}_{ECMF}^{E-res}) + \left(\sum_{i=1}^n b_{0,i}^{ch} \dot{m}_i \right)_{ECMF}^{E-fuel} - \left(\sum_{i=1}^n b_{0,i}^{ch} \dot{m}_i \right)_{ECMF}^{E-res} - \sum_{i>0} \left(1 - \frac{T_0}{T_i} \right) \dot{Q}_{ECMF}^{MF\leftarrow} - T_0 \dot{S}_{irr,MF} \quad \text{Eq. 5.6}$$

The rate of exergy lost during these two processes (excluding waste) is given by:

$$D_{ECMF+MF} = T_0 (\dot{S}_{irr,MF} + \dot{S}_{irr,ECMF}) \quad \text{Eq. 5.7}$$

Finally, we seek to define a ratio that can be used as a metric in the evaluation of the efficiency of these processes. As opposed to a second-law-style efficiency that compares an actual process to an ideal process, we instead set up a ratio between the desired output

from and the total input to a process. The “degree of perfection,” as this ratio is called, is given by Szargut as:

$$\eta_p = \frac{B_{Useful\ products}}{B_{Inputs}} \quad \text{Eq. 5.8}$$

The degree of perfection essentially relates the percentage of exergy that is preserved in the output product compared with the amount of exergy that went into the process. Thus, for a reversible process with no waste heat or materials, the degree of perfection goes to unity. On the contrary, a perfectly dissipative process with no useful output (e.g., idling a car) has a degree of perfection of zero. From (6) and the definition of the degree of perfection, the thermodynamic efficiency of a process is reduced 1.) through the generation of entropy and 2.) through the waste of input exergy (i.e., lack of incorporation of input materials or energy in a final product).

The degree of perfection ratio is useful for additive, chemical, and bulk forming processes, but its shortcomings are apparent when one attempts to apply it toward removal processes. How does one define the “useful products” when the end goal of the process is the removal of material and the concomitant destruction of exergy? Additionally, traditional measures of efficiency have difficulty treating the material that flows through a process but which is not operated on. For example, engraving a part number on a large aluminum work piece should not be considered a more efficient process than one in which a small aluminum work piece is engraved on the basis that such a large percentage of the input material is available in the output product. Sorin et. al. treat this dilemma by identifying two separate exergy streams: the so-called “transiting exergy” which flows through the process unaltered, and the “utilizable exergy” which is manipulated and active in the process [37].

To the author’s knowledge, the question of how to talk about the efficiency of removal processes thermodynamically is not resolved in the literature. It would be useful to have a metric for the evaluation of the efficiency (in exergy terms) of subtractive processes. This metric should be dependent on the amount of material removed in a process, as well as the amount of exergetic work that is input into the process to effect the material removal. It should be independent of the amount of the “transiting exergy” that enters a process, or raw material that is not operated upon. This metric should approach unity as the exergy loss approaches zero, and it should approach zero as the amount of exergy expended in the process becomes much greater than the exergy of the material removed by the process. I propose the following ratio, which fits these criteria and can be termed the exergetic efficiency of removal (η_R), as a way to evaluate removal processes on a basis that can be compared to additive, chemical, and bulk molding processes:

$$\eta_R = \frac{B_{Removed}}{B_{Removed} + B_{Inputs}} \quad \text{Eq. 5.9}$$

where $B_{Removed}$ gives the exergy of the material removed and B_{Inputs} gives the exergy of the input materials and energy required to effect the material removal. For example, in

the etching of Si_3N_4 using CF_4 , B_{Removed} would represent the standard chemical exergy of silicon nitride multiplied by the amount of material removed. B_{Inputs} would correspond to the standard chemical exergy of carbon tetrafluoride multiplied by the amount used in the process, plus the electrical energy used. In essence, the summation ($B_{\text{Removed}} + B_{\text{Inputs}}$) gives the “utilizable” exergy described in [37].

In addition to satisfying the mathematical requirements for a removal efficiency ratio, equation (9) makes logical sense as well. In the same way that the numerator in equation (8) captures the desired output of an additive process, the numerator in (9) gives the exergy embodied by the material removed, which corresponds to the desired output for a removal process. Similarly, the denominators in both (8) and (9) represent the utilizable exergy going into the process. The use of B_{Removed} in the numerator might seem unusual at first until it is recognized that the term accounts for that part of the input material that is operated upon during a process. For example, in the etching of a thin film on a silicon wafer, it is not the mass and exergy of the entire silicon wafer that is of interest, but instead only that part of the thin film on the wafer surface which is removed. The application of the exergetic efficiency of removal can be found in subsections 5.3.3 and 5.3.4.

With these equations and the appropriate ones for the other subsystems, we have completely described the manufacturing system in terms of loss of potential. We save the discussion of the interesting aspects of these equations for the following sections.

5.3 Semiconductor Manufacturing Process Analysis

A precise but general characterization of semiconductor manufacturing processes would be extremely challenging to achieve, particularly given the virtually infinite combination of variables that can be tweaked in a given process area. As an example, the power and material consumed in a CVD process is a function of the type of tool used, the auxiliary equipment required (e.g., PECVD requires an RF generator), the design of the deposition chamber and the number of wafers it holds, the type and thickness of the film to be deposited, and the process recipe (i.e., temperature, pressure, RF power, existence of cleaning steps, idle time, etc.). A rigorous analysis would isolate the effect of each of these variables. Such a task would require virtually unlimited access to CVD equipment, and in reality, is much more than is required for the environmental characterization of CVD. Instead, to obtain an understanding of several important semiconductor processes from an industrial ecological perspective, a small number of tests have been conducted to bring to the fore the major themes associated with each. The results from the monitoring and analysis of these specific processes (plasma-enhanced chemical vapor deposition, sputter deposition, dry chemical etching, and wet chemical etching) are presented below.

5.3.1 Chemical Vapor Deposition

Chemical vapor deposition (CVD) is one of the two principal additive techniques in semiconductor processing (excluding ion implantation and thermal oxidation). Please see Section 1.2.2.4.1 for a description of the CVD process and its variations.

The ability to deposit thin films with high dimensional, crystallographic, and compositional accuracy makes chemical vapor deposition an extraordinarily valuable

instrument in the process engineer's toolbox. Yet from an environmental standpoint, CVD is often viewed as the epitome of a wasteful process. Before CVD can be improved, however, it is necessary to understand well how energy and materials are used in CVD. To that end, tests were conducted at the Analog Devices Micromachined Products Division (MPD) facility in Cambridge, MA in which material and energy flows and useful product were measured during the plasma-enhanced CVD (PECVD) of silicon dioxide and silicon nitride.

The deposition of a 16,000Å undoped oxide layer was conducted on a PECVD tool for a run of twenty-five wafers. The deposition chamber holds seven wafers at a time, with each wafer rotating through the chamber in about three-and-a-half minutes. The entire test run lasted seventy-two minutes and included a pre-coat step, deposition, and two clean cycles. Gas flow rates were gathered from the digital readout on the tool. All components of this tool – the RF generators, vacuum pumps, heaters, robotics, and controls – are powered from a single circuit, so the power required during the process was monitored with high precision using the AEMC power meter described in Chapter 3.

Film	Undoped Silicon Dioxide (SiO ₂)	Silicon Nitride (Si ₃ N ₄)
Thickness	16,000 Å	5,000 Å
Batch Size	25 six-inch wafers	
Temperature	400°C	
Pressure	2.07T	2.8T
Total Time	72 minutes	49 min., 15 sec.

Table 5.4: Process parameters for test runs of oxide and nitride PECVD

PECVD of Silicon Dioxide (Full Process)				
Input Deposition Gases				
<i>Inputs</i>	<i>Mass (g)</i>	<i>Moles</i>	<i>Specific Chemical Exergy (kJ/mol)</i>	<i>Primary Exergy (kJ)</i>
N ₂	276.3	9.86	0.69	6.80
SiH ₄	8.57	0.267	1383.7	369.4
N ₂ O	440.6	10.01	106.9	1070.2
Input Cleaning Gases				
O ₂	69.09	2.16	3.97	8.57
C ₂ F ₆	298.0	2.16	962.4	2078.1
Input Energy				
Electricity				50516
<i>Total In</i>				54049
Output				
Undoped Silicon Dioxide Layer	1.555	2.59E-02	7.9	0.204
<i>Total Out</i>				0.204
Degree of Perfection (η_P)				3.78E-06

Table 5.5: Summary of oxide CVD (entire process). Thermodynamic data from [36, 38].

PECVD of Silicon Nitride (Full Process)				
Input Deposition Gases				
<i>Inputs</i>	<i>Mass (g)</i>	<i>Moles</i>	<i>Specific Chemical Exergy (kJ/mol)</i>	<i>Primary Exergy (kJ)</i>
N2	264.1	9.43	0.69	6.50
SiH4	12.34	0.384	1383.7	531.6
N2O	37.80	0.86	106.9	91.81
NH3	38.87	2.28	337.9	771.1
Input Cleaning Gases				
O2	25.12	0.79	3.97	3.12
C2F6	108.4	0.79	962.4	755.7
Input Energy				
Electricity				32174
	<i>Total In</i>			34333
Output				
Silicon Nitride Layer	0.685	4.88E-03	1917.9	9.36
	<i>Total Out</i>			9.36
Degree of Perfection (η_P)				2.73E-04

Table 5.6: Summary of nitride CVD (entire process). Thermodynamic data from [36, 39].

The deposition of a 5,000Å nitride layer was conducted on the same tool with many of the same parameters, although the total process time was only 49 minutes and 15 seconds and included an undercoat step before the pre-coat (both steps are used to prepare the deposition chamber for the next deposition process and avoid contamination from earlier runs by laying down an initial layer of material). A summary of the test parameters for this process and the oxide deposition are given in Table 5.1.

Tables 5.2 and 5.3 present the results for the entire deposition process of the oxide and nitride depositions (including cleaning steps). The thermodynamic modeling of the process takes all input and output materials at atmospheric pressure and temperature (i.e., the physical exergy component is zero). In essence, the system boundaries for the analysis are drawn large enough to allow the materials to reach mechanical and thermal equilibrium with the environment. We do this because we are principally interested in the chemical exergy component of the input materials, understanding that any physical exergy component would have been imparted to the raw material at an earlier stage of the manufacturing process (e.g., pressurization of process gases for bottling and shipping). Furthermore, the physical exergy component of the pressurized gases used in CVD and other semiconductor processes is small in comparison to their chemical exergy component (with the possible exceptions of nitrogen, oxygen, and argon). The useful output is quantified as the material deposited on the wafer surfaces only. Other outputs (waste heat, waste gases) are not accounted for in this analysis.

Probably the most eye-popping statistic is the ratio of the amount of material input into these processes compared to the amount of material incorporated into the final product; for both cases it is about 700:1! Yet the exergy of the input materials makes up only a fraction of the total input exergy. The great majority enters the process in the form

of electrical work. However, only an infinitesimal amount of the input electrical energy actually goes into effecting the reaction necessary to deposit the film. Most of the electrical energy goes into so-called auxiliary process: vacuum pumps for reducing the chamber pressure (necessary for the plasma to strike and for film purity), robotics and controls for moving the wafers, and heaters which dissipate most of their heat into gases that go straight to the exhaust. Even the RF generator, which ignites the plasma in the deposition chamber, ultimately converts very little of the input electrical power into energy for ionizing deposition gases.

This mismatch between the quantity and quality of inputs and that of the outputs is captured in the vanishingly small values for the degree of perfection: 3.78×10^{-6} for oxide deposition and 2.73×10^{-4} for the nitride deposition. (The two-order of magnitude difference in the degree of perfection between the nitride and oxide runs is a function of the quality of the output product. Whereas the chemical exergy of silicon dioxide is only 7.9 kJ/kmol, that of silicon nitride is 1917.9 kJ/kmol.) In comparison, the degree of perfection for a manufacturing process like injection molding or the melting portion of metal casting, where most of the input material is incorporated into the final product, typically exceeds 0.40 and ranges upward of 0.75.

The results from these CVD process monitoring runs can be compared with the data of Boyd et. al. for the deposition of undoped silicate glass (USG), which is essentially an oxide film with a different crystalline structure. The data from that study, portrayed in Table 5.4, gives a degree of perfection for the entire process that is within a factor of three of that found in the MPD study. However, at about 9200:1, the ratio of the input mass of material to the deposited mass is an order of magnitude larger than for the

Input Deposition Gases				
<i>Inputs</i>	<i>Mass (g)</i>	<i>Moles</i>	<i>Specific Chemical Exergy (kJ/mol)</i>	<i>Exergy (kJ)</i>
SiH4	0.95	2.96E-02	1383.7	40.9
O2	0.49	1.53E-02	3.97	0.061
Ar	0.34	8.51E-03	11.69	0.099
N2	196.9	7.03	0.69	4.85
Input Cleaning Gases				
NF3	31.06	0.437	610.2	267
Input Energy				
Electricity				2,220
<i>Total In</i>				2,533
Output				
Undoped Silicate Glass Layer	2.48E-02	4.14E-04	7.9	3.27E-03
<i>Total Out</i>				3.27E-03
Exergy Destroyed				2,533
Degree of Perfection (η_p)				1.29E-06

Table 5.7: Published data on the CVD of undoped silicate glass (USG). 6,000 angstrom thickness. Data from [26].

MPD study and probably at least in part a function of batch size [26]. The similarity in the order of magnitude of the degree of perfection is indicative of the dominance of the electrical input in determining the resource requirements of the process.

5.3.1.1 *Chemical Vapor Deposition of Silicon Dioxide and Silicon Nitride (Deposition Only)*

The degree of perfection, and hence the thermodynamic “quality,” of a process is strongly influenced by the boundary selected for the process to be analyzed. In Section 5.3.1, we evaluate the entire CVD process, including steps like pre-coating and cleaning

PECVD of Silicon Dioxide (Deposition Only)				
Input Materials				
<i>Inputs</i>	<i>Mass Rate (g/sec)</i>	<i>Moles per second</i>	<i>Specific Chemical Exergy (kJ/mol)</i>	<i>Primary Exergy Rate (kW)</i>
N2	0.13	4.63E-03	0.69	3.20E-03
SiH4	0.004	1.31E-04	1383.70	1.81E-01
N2O	0.216	4.91E-03	106.9	5.25E-01
Input Energy				
Electricity				10.59
<i>Total In</i>				11.30
Output				
SiO2	1.38E-03	2.30E-05	7.90	1.82E-04
<i>Total Out</i>				1.82E-04
Degree of Perfection (η_p)				1.61E-05

Table 5.8: Oxide PECVD process-monitoring results for the deposition only.

PECVD of Silicon Nitride (Deposition Only)				
Input Materials				
<i>Inputs</i>	<i>Mass Rate (g/sec)</i>	<i>Moles per second</i>	<i>Specific Chemical Exergy (kJ/mol)</i>	<i>Primary Exergy Rate (kW)</i>
N2	1.21E-01	4.31E-03	0.69	2.97E-03
SiH4	6.13E-03	1.91E-04	1383.70	2.64E-01
NH3	2.09E-02	1.23E-03	337.9	4.15E-01
Input Energy				
Electricity				10.71
<i>Total In</i>				11.39
Output				
Si3N4	9.13E-03	6.51E-05	1917.88	1.25E-01
<i>Total Out</i>				1.25E-01
Degree of Perfection (η_p)				1.10E-02

Table 5.9: Nitride PECVD process-monitoring results for the deposition only.

during which energy and materials are being consumed but no useful output – thin-film deposition on a wafer surface – is taking place. Tables 5.5 and 5.6 give the statistics for the oxide and nitride depositions from the earlier section, but using only data for the deposition time. Instead of using cumulative mass and energy inputs for the process, the data in Tables 5.5 and 5.6 is given in rate form. As anticipated, the degree of perfection improves significantly (by about an order of magnitude) for the analysis of the deposition process only compared to the whole process. This is primarily because the electricity consumed by the tool during the rest of the operation is now ignored, but the absence of high-exergy cleaning gases also impacts the process quality.

5.3.2 Sputter Deposition

Sputtering is the most common physical vapor deposition technique, and a description of the process can be found in Section 1.2.2.4.2. The sputtering of a 10,000 Å aluminum copper film was monitored for a ten-wafer run at the Analog Devices MPD fab. This particular tool has four separate deposition/sputter-etch chambers, which allows the tool to run several parallel processes. Each individual chamber holds just one wafer, but the deposition time per wafer is fairly quick (~45 seconds for 10,000 Å). The entire test run for the ten wafers, including loading and unloading times, was twenty-one minutes. The argon flow rate was gathered from the digital readout on the tool; the input AlCu was calculated on a per wafer basis by weighing a used target to find the mass lost by the target over its lifetime and then dividing by the total number of wafers processed using that target.

Unlike the PECVD tool, this system’s components are powered by four separate circuits. The circuit that powers the DC voltage generators, referred to as ECLIPSE, varies irregularly and mostly powers components only during deposition times. For these reasons, we selected to measure this circuit directly during the process (see Figure 5.2).




Figure 5.18: Power record from the ECLIPSE circuit for the sputter deposition of AlCu (the spikes correspond to each wafer deposition step). Note that the process was initiated seven minutes before the first deposition; the lag is a result of time spent loading wafers.

Film	AlCu
Thickness	10,000 Å
Batch Size	10 six-inch wafers
Temperature	300°C
Pressure	2.4mT
Total Time	21 minutes

Table 5.10: Process parameters for an AlCu sputter deposition

The power consumption for the other circuits was taken from data gathered during the monitoring of building power consumption (Chapter 4) since these circuits run at fairly constant average powers. The largest, referred to as PCC, runs at a virtually constant 16.0 kW regardless of whether or not product is being run. This circuit powers the tool’s vacuum pumps and heaters, among other things. The two smallest circuits, referred to as Tower A and Tower B, also run at a fairly constant average (although cyclic) power of 0.35 kW and 1.84 kW, respectively.

The results of the process monitoring run are presented in Table 5.8. As in the case of the CVD processes, input and output materials are evaluated at the reference temperature and pressure of the environment. As is evident from the chart, the sputtering process differs from the CVD process in that the thermodynamic potential of the

Sputtering of an AlCu film (Full Process)				
Input Materials				
<i>Inputs</i>	<i>Mass (g)</i>	<i>Moles</i>	<i>Specific Chemical Exergy (kJ/mol)</i>	<i>Primary Exergy (kJ)</i>
Ar	3.43	0.09	11.69	1.00
AlCu	2.44	0.09	885.0	78.96
Input Energy				
Electricity				29909
	<i>Total In</i>			29988
Output				
AlCu Film	0.498	1.82E-02	885.04	16.13
	<i>Total Out</i>			16.13
Degree of Perfection (η_p)				5.38E-04

Table 5.11: Summary of AlCu sputtering for a ten-wafer run (entire process). AlCu input mass determined by weighing a used target to calculate the lost mass, dividing the mass lost by the total number of wafers processed using that target (which gives the mass removed per wafer processed), and finally multiplying by ten to get the mass lost from the target for a ten-wafer run.

materials input to the process is completely irrelevant in comparison to that of the input electricity. In spite of this difference, the degree of perfection is similar to that of nitride CVD since the process rates are relatively similar and electricity dominates the inputs in both cases.

It should be reiterated that the data for the sputtering test is not as “clean” as for the CVD tests. Whereas in the case of the CVD system, the entire tool is dedicated to a single process at a time, the sputter tool has multiple chambers each intended for a

specific process. For this process run, a recipe was chosen that required deposition in only one chamber in an attempt to minimize extraneous power readings. However, the energy required for vacuum and heating in the other chambers could not be extricated from the data for the whole tool. As discussed in earlier sections, this highlights the importance of tool type and design on process energy and material consumption.

5.3.2.1 Sputtering of an Aluminum Copper Film (Deposition Only)

Table 5.9 gives the statistics for the AlCu deposition from the earlier section, but only with data for the deposition time. Hence, instead of using total mass and energy inputs for the process, the data in Table 5.9 is given in rate form. The degree of perfection improves only by about a factor of three compared to that for the entire process. This is sensible because the only real difference in resource consumption between the two scenarios is in the additional amount of electricity required by the tool while it is idling or not running product.

Sputtering of an AlCu Film (Deposition Time Only)				
Input Materials				
<i>Inputs</i>	<i>Mass Rate (g/sec)</i>	<i>Moles per second</i>	<i>Specific Chemical Exergy (kJ/mol)</i>	<i>Primary Exergy Rate (kW)</i>
Ar	2.72E-03	6.82E-05	11.69	7.97E-04
AlCu	5.41E-03	1.98E-04	885.0	1.75E-01
Input Energy				
Electricity				32.08
<i>Total In</i>				32.26
Output				
AlCu Film	1.11E-03	4.05E-05	885.0	3.58E-02
<i>Total Out</i>				3.58E-02
Degree of Perfection (η_p)				1.11E-03

Table 5.12: AlCu sputtering process-monitoring results for the deposition time only. AlCu input mass determined by weighing a used target to calculate the lost mass, dividing the mass lost by the total number of wafers processed using that target (which gives the mass removed per wafer processed), and finally dividing by the length of the deposition of a single wafer (45 seconds) to get a mass rate.

5.3.3 Dry Chemical Etching

Etching is the counterpart to the deposition processes discussed in Sections 5.3.1 and 5.3.2. After a deposition process, lithography is used to define a pattern on the wafer surface. An etch process then removes the unprotected thin film as well as some of the photoresist. The focus of this subsection, dry etching, uses a plasma to liberate highly reactive species from precursor gases (usually fluorocompounds) which then remove the target material from the workpiece. More information on etching can be found in Section 1.2.2.5.

5.3.3.1 Dry Oxide Etch

The dry etching of a 4,000 Å silicon dioxide layer was monitored for a four-wafer run at the Analog Devices MPD fab. The tool etches one wafer at a time with each etch requiring about one minute. The entire process, including loading, unloading, and pump down times, was nine minutes. Gas flow rates were taken from the process recipe.

This tool is supplied by a labyrinthine network of eight separate circuits powering different tool components. With only one power meter, it is impossible to measure the power characteristics of all seven circuits during one test run. Two measurement runs were conducted, which allowed accurate and precise measurements of the main tool circuit and the RF generator circuit using the power meter. The tool has two temperature control units; power draw was measured on each of those by doing a short spot measurement at the conclusion of the test runs using the (well-founded) assumption that power draw was fairly constant throughout the test (confirmed by current measurements during the test). The current draw on the two vacuum pumps supplying the etcher was measured during the tests to determine an estimated power, but since these are three phase systems, such a calculus will induce some error into the results. That error is difficult to quantify without a point of reference, but the figures for power draw from these two pumps (2.56 and 1.78 kW) are expected to be an upper-bound. Another three-phase component, the gas abatement system which burns any remaining

Film	SiO ₂
Etch Depth	4,000 Å
Batch Size	4 six-inch wafers
Temperature	40°C
Total Time	12 minutes

Table 5.13: Process parameters for an oxide dry etch.

Dry Etching of a Silicon Dioxide Film (Full Process)				
Input Materials				
<i>Inputs</i>	<i>Mass (g)</i>	<i>Moles</i>	<i>Specific Chemical Exergy (kJ/mol)</i>	<i>Primary Exergy (kJ)</i>
Ar	4.18	1.05E-01	11.69	1.22
CHF ₃	0.378	5.40E-03	569.0	3.07
CF ₄	0.389	4.42E-03	454.1	2.01
Input Energy				
Electricity				5565
<i>Total In</i>				5571
Output				
Etched SiO ₂	0.067	1.12E-03	7.90	8.86E-03
<i>Total Out</i>				8.86E-03
Exergetic Efficiency of Removal (η_R)				1.59E-06

Table 5.14: Summary of a blanket oxide dry etch (entire process).

perfluorocarbons after the etch process, was measured in the same fashion. Its estimated average power consumption was 360 W. The eighth component was a very small 120V circuit drawing only 36W.

Table 5.11 gives the results of the oxide etch test run for the entire process. Once again, since it was not possible to measure all components simultaneously during the test, the quoted electricity consumption through the process is not as accurate as that given in the CVD tests. Based on the size of the various circuits and the likely error involved with each, the stated electricity consumption should be within 10-15% of the actual value.

As with the degree of perfection for chemical vapor deposition of silicon dioxide, the exergetic efficiency of removal for this etch process is on the order of 10^{-6} . The fact that this result is similar to that obtained for CVD indicates that the degree of perfection of a process might be correlated with the rate at which the process occurs, a topic which will be revisited in Section 5.4. Table 5.11 also shows that the ratio of the mass of input materials to the mass of material etched is about 75:1, an order of magnitude smaller than for CVD (though an etch process requires no materials-intensive cleaning step). It is evident that the chemical potential embodied in the input materials is a fraction of a percent of the potential going into the process in the form of electricity. Once again, the electrical inputs dominate the chemical inputs.

5.3.3.1.1 Dry Oxide Etch (Etch Time Only)

Table 5.12 gives the results of the oxide etch, excluding time spent loading the wafers and pumping down the chamber. The difference between the exergetic efficiency of removal for the full process and for the etch time is about a factor of three. As with the earlier processes, this difference is a result of the time spent idling, moving product from one chamber to the next, pressurizing or depressurizing chambers, or otherwise functioning in a way that is only indirectly related to the process objective.

Dry Etching of a Silicon Dioxide Film (Etch Time Only)				
Input Materials				
<i>Inputs</i>	<i>Mass Rate (g/sec)</i>	<i>Moles per second</i>	<i>Specific Chemical Exergy (kJ/mol)</i>	<i>Primary Exergy Rate (kW)</i>
Ar	5.45E-03	1.36E-04	11.69	1.59E-03
CHF3	1.05E-03	1.50E-05	569.0	8.53E-03
CF4	1.08E-03	1.23E-05	454.1	5.57E-03
Input Energy				
Electricity				8.74
	<i>Total In</i>			8.76
Output				
Etched SiO2	2.81E-04	4.67E-06	7.90	3.69E-05
	<i>Total Out</i>			3.69E-05
Exergetic Efficiency of Removal (η_R)				4.22E-06

Table 5.15: Oxide dry etch process-monitoring results for the etch phase only.

5.3.3.2 Dry Nitride Etch

The dry etching of 2,000 Å of silicon nitride was monitored for a three-wafer run at the Analog Devices MPD fab (different from that used in the oxide etch). As with the oxide etch, this piece of equipment etches one wafer at a time, but the etch time per wafer is fairly quick (~30 seconds per 2,000 Å). The length of the test run for the three wafers, including loading and unloading times, was four minutes. Gas flow rates were taken from the process recipe.

In addition to the main power supply to the tool, there are also two vacuum pumps and one dedicated water chiller that are on separate circuits. The power draw of the main supply was monitored with the AEMC meter during the test, while current readings were taken of the other three circuits to derive estimated power usage. Without the surety of the power meter, it is probable that there is considerable error in the estimated power for the vacuum pumps (1.49 and 2.59 kW) since those are three-phase systems (the chiller is 120V). The stated power consumption for the test should be considered an upper bound for this tool running the process described in Table 5.13.

Film	Si ₃ N ₄
Etch Depth	2,000 Å
Batch Size	3 six-inch wafers
Temperature	20°C
Total Time	4 minutes

Table 5.16: Process parameters for a nitride dry etch.

Dry Etching of a Silicon Nitride Film (Full Process)				
Input Materials				
<i>Inputs</i>	<i>Mass (g)</i>	<i>Moles</i>	<i>Specific Chemical Exergy (kJ/mol)</i>	<i>Primary Exergy (kJ)</i>
He	2.46E-02	6.13E-03	30.37	1.86E-01
SF6	2.688	1.84E-02	281.8	5.2
Input Energy				
Electricity				1178
			<i>Total In</i>	1184
Output				
Etched Si ₃ N ₄	0.033	2.34E-04	1917.88	0.45
			<i>Total Out</i>	0.45
Exergetic Efficiency of Removal (η_R)				3.79E-04

Table 5.17: Summary of a blanket nitride dry etch (entire process).

Table 5.14 shows that the exergetic efficiency of removal for this etch process is on the order of the degree of perfection for the chemical vapor deposition of nitride (2.73×10^{-4}). Just as in the dry etching of silicon dioxide, the electrical energy input dominates the input chemical potential, so η_R is primarily determined by the electricity

used during the process. Even so, the input mass of materials outweighs the mass of material etched by a ratio of about 80:1.

5.3.3.2.1 Dry Nitride Etch (Etch Time Only)

The thermodynamic summary of the dry nitride etch process is given in Table 5.15 and represents the rate of exergy inputs and useful outputs occurring instantaneously during the deposition process. There is a relatively small difference in the efficiency ratio between the etch-only analysis (7.14×10^{-4}) and the full process analysis (3.79×10^{-4}) because this particular etcher requires shorter pump-down and product handling times than do the other processes previously discussed.

Dry Etching of a Silicon Nitride Film (Etch Time Only)				
Input Materials				
<i>Inputs</i>	<i>Mass Rate (g/sec)</i>	<i>Moles per second</i>	<i>Specific Chemical Exergy (kJ/mol)</i>	<i>Primary Exergy Rate (kW)</i>
He	1.36E-04	3.41E-05	30.37	1.03E-03
SF6	1.49E-02	1.02E-04	281.8	2.88E-02
Input Energy				
Electricity				6.96
<i>Total In</i>				6.99
Output				
Etched Si ₃ N ₄	3.65E-04	2.60E-06	1917.88	4.99E-03
<i>Total Out</i>				4.99E-03
Exergetic Efficiency of Removal (η_R)				7.14E-04

Table 5.18: Nitride dry etch process-monitoring results for the etch phase only.

5.3.4 Wet Etching of Silicon Dioxide and Silicon Nitride

Although dry etching has become the more common method for removing material from a wafer – primarily owing to its better directional characteristics (that is, the ability to produce very straight sidewalls in an etched cavity) and ability to produce finer line widths – wet etching has an important place in semiconductor processing. Wet etches usually produce very selective etches (preference for etching one material over another), and even the isotropy (lack of directionality) common of wet etches is useful in certain processes. More information on wet etching can be found in Section 1.2.2.5.

In this section, results from the monitoring of oxide and nitride wet etches are presented and compared with their dry etch counterparts from Section 5.3.3. The data acquisition for these processes differed from that of the other process monitoring runs. Whereas gas-phase processes use the input chemicals in a one-off, easily measured fashion, wet chemical etching uses a single chemical bath that will often be used and reused for many batches of wafers before it is changed out. In order to get an estimate of

the materials input into these wet chemical etching processes, the volume of chemicals in a single chemical bath was amortized over the approximate number of wafers processed using a single batch of chemicals.

The electrical energy determination for the wet etch processes was also challenging. The oxide and nitride etches were conducted in different chemical baths, but the baths were each part of the same tool. The tool actually consists of four sinks, each with its own temperature requirements. The input electrical power to the entire tool was recorded with the AEMC meter during the test runs; unfortunately, the power draw to the individual sinks could not be parsed from this data. Furthermore, the power draw on the tool is not steady, but changes dramatically at irregular intervals as heaters turn on and off. It follows that for the short period of an etch process (<2 minutes), much of the thermal energy in a given etch bath is imparted before the technical start of the process. To determine as realistic a figure as possible for the electrical power used, knowledge of the temperature of the various sinks was combined with the power meter data to yield reasonable estimates for electricity consumption. The operating temperatures of the four sinks are 27°C, 151°C, 38°C, and 20°C. Obviously, the bulk of the electrical energy for bath heating goes to the second sink. With that sink (fortuitously) off for the oxide etch, which was conducted in the 27°C bath, the average power during the oxide etch process was taken directly from the power meter for the period in which the etch was conducted. Although that figure includes power going to the other two sinks as well, it is not unreasonably large and is therefore accepted as an upper bound for the oxide etch electrical consumption. Then, with the 151°C sink energized for the nitride etch, the power draw for that sink appeared as pronounced spikes in the meter data dwarfing the power going to the other sinks. The electrical energy for the nitride etch was then taken as the average power consumption between the beginning of the first spike and the end of the second (see Figure 5.3). Although it too includes some power going to the other sinks, the figure should be sufficiently accurate for the purposes of this characterization and at worst can be regarded as an upper-bound.




Figure 5.19: Power meter data for the main circuit for the VERTEQ quad sink. Note the three large spikes, corresponding to firing of the heater for the nitride sink.

The oxide etch was conducted for a 50-wafer run in a 6:1 buffered oxide etch (BOE) bath. As stated earlier, the bath temperature was 27°C. For the data presented in Table 5.17, the chemical composition of the 6:1 BOE etch was taken as the average given

Etched Film	Silicon Dioxide (SiO ₂)	Silicon Nitride (Si ₃ N ₄)
Thickness	3,000 Å	1,100 Å
Etchant	6:1 BOE	H ₃ PO ₄
Batch Size	50 six-inch wafers	
Bath Temperature	27°C	151°C
Etch Time	3 minutes	1 minute

Table 5.19: Process parameters for oxide and nitride wet etch monitoring tests.

Wet Etching of a Silicon Dioxide Film (Etch Only)				
Input Materials				
<i>Inputs</i>	<i>Mass (g)</i>	<i>Moles</i>	<i>Specific Chemical Exergy (kJ/mol)</i>	<i>Primary Exergy (kJ)</i>
H ₂ O	176.82	9.82E+00	0.9	8.83
NH ₄ F	78.587	2.12E+00	360.0	763.86
HF	25.260	1.26E+00	80.0	101.01
Input Energy				
Electricity				261.0
<i>Total In</i>				1135
Output				
Etched SiO ₂	0.602	1.00E-02	7.90	7.91E-02
<i>Total Out</i>				7.91E-02
Exergetic Efficiency of Removal (η_R)				6.97E-05

Table 5.20: Summary of a wet oxide etch.

Wet Etching of a Silicon Nitride Film (Etch Only)				
Input Materials				
<i>Inputs</i>	<i>Mass (g)</i>	<i>Moles</i>	<i>Specific Chemical Exergy (kJ/mol)</i>	<i>Primary Exergy (kJ)</i>
H ₃ PO ₄	252.82	2.58	104.00	268.32
H ₂ O	72.432	4.02	0.90	3.62
Input Energy				
Electricity				525.6
<i>Total In</i>				797.5
Output				
Etched Si ₃ N ₄	0.301	2.15E-03	1917.88	4.12
<i>Total Out</i>				4.12
Exergetic Efficiency of Removal (η_R)				5.14E-03

Table 5.21: Summary of a wet oxide etch.

in the MSDS (28% ammonium fluoride (NH₄F), 9% HF, 63% H₂O) since the precise composition is proprietary.

A “hot-phos” bath – phosphoric acid at 151°C – was used to etch the nitride film. Summary data for this process are given in Table 5.18.

The fact that input chemicals make up a substantial percentage of the input potential in wet etch processes (77% for the oxide etch and 34% for the nitride etch) distinguishes wet etching from other semiconductor processes. Since wet etching is a primarily chemical removal process, this result is sensible. Nonetheless, electricity consumption plays a non-negligible role in the resource intensity of these processes. Particularly in the case of these wet etch sinks, the use of high-quality-energy electricity to provide low-quality-energy (low temperature) heat is a quintessential example of a mismatch between the exergy required to accomplish a process and that actually used to accomplish the task.

The comparison between dry and wet chemical etching is a perfect example of the utility of thermodynamic (exergy) analysis, as the former is rather energy intensive while the latter is more materials intensive. Other metrics, such as Life Cycle Analysis, must either resort to treating energy and materials separately or develop an equivalent mass for the energy used. The real strength of thermodynamic analysis lies in being able to treat the potential of energy and matter simultaneously. From the data in Tables 5.12 and 5.15 (for the etch phase only of dry etch processes), it is apparent that the exergetic efficiency of removal for dry etch processes (4.22×10^{-6} and 7.14×10^{-4} for oxide and nitride, respectively) is much lower (by about an order of magnitude) than that for wet etch processes (6.97×10^{-5} and 5.14×10^{-3} for oxide and nitride). With these results, one could posit that wet etch processes are more sustainable – less exergy intensive – than dry etch processes. That might be misleading, however, because in order to understand the true energy cost of the different processes, it would be necessary to expand the analysis beyond the manufacturing process alone (Ω_{MF} from Figure 5.1) to the electricity production system (Ω_{ECMF}) and the chemical production system (Ω_{MA}). Expanding the boundaries of these analyses might give an altogether different result. Yet we can say with a high degree of confidence that at least at the process level, dry etching is more resource intensive than wet etching.

The difference in the input-to-output material ratio between the two types of etches should be emphasized as well. Whereas in the dry etch processes, the ratio of input mass to the mass of material etched is between 70:1 and 85:1, in the wet etch processes it ranges from 465:1 for the oxide etch to 1080:1 for the nitride etch. (It should be reiterated that the mass intensity ratios are affected by the estimated number of wafers that can be etched before the chemical bath must be replaced.) The roughly order of magnitude difference in the mass intensity of the dry and wet etch processes is not surprising given that one process uses liquid chemicals and the other gaseous chemicals. The reality that the more materials intensive process is still more efficient in thermodynamic terms illustrates how important electrical energy is to the determination of process efficiency.

5.4 Process Rate vs. Specific Energy Consumption

The possible relationship between the rate at which a manufacturing process occurs and the energy/resource intensity of the process has been alluded to earlier in this chapter and Chapter 4 on multiple occasions. In the following section we explore how the specific energy consumption of semiconductor processes relates to other manufacturing processes, and what that relationship might suggest about the origin of the very low efficiencies seen in these processes.

In Chapter 4, we present a model for the electrical consumption of manufacturing processes that consists of a base or idle component that is independent of process rate and a variable component that is process-rate dependent. In the simple case where absolute energy consumption increases linearly with increasing throughput, the specific energy consumption (energy per unit volume or mass processed) varies according to Equation 4.1:

$$SEC = \frac{P_0}{\dot{V}} + k$$

where P_0 gives the idle power, k is the constant of proportionality between throughput and total energy consumption, and \dot{V} is the process rate or throughput. Gutowski et. al. suggests that P_0 is highly constrained in most manufacturing activities by machine size and electrical current limitations. If that is the case, then it follows that most manufacturing processes should fall on roughly the same trend line on a graph of process rate vs. specific energy consumption as shown in Figure 5.4. In it, the data from the present Analog Devices study are combined with those for eight non-semiconductor manufacturing processes from [14].

The two different regimes visible in this plot are explicable in terms of Equation 4.1. In the constant region at the far right side of the graph (at throughputs greater than about $1 \cdot 10^1$ kg/hr), energy consumption is dominated by the proportionality constant (k) between throughput and energy intensity. In other words, in this regime the rate of a process is sufficiently high that the base, or idle, power is small in comparison to that which is determined by the physics of the process. It makes sense that the very high rate processes would be the most efficient, too, since manufacturers processing such massive amounts of material inefficiently would find themselves with a very expensive product. In the linear regime below about 10kg/hr, the energy intensity of the process becomes a strong function of throughput. It is in this regime where significant improvements in energy intensity can be made with an increase in throughput.

There are two major implications of this result:

1. The extraordinarily low degree of perfection or exergetic efficiency of removal for the semiconductor processes discussed in Section 5.3 is driven by the low process rates at which these processes operate and may be a fundamental obstacle to their improvement. Since the degree of perfection for all but the wet etch processes is governed by electricity consumption, the specific *exergy* consumption of each process is practically identical to the specific *energy* consumption shown in Figure 5.4. Despite the very small amount of material that is operated upon in semiconductor processes, energy is supplied in large quantities. In essence, there is a mismatch between the

scale of the process and the scale of the equipment. This inability to reduce the magnitude of material and energy requirements with decreasing dimensional scale is one of the primary drivers behind the lack of decoupling between resource use and economic growth as the world moves toward a knowledge-based economy [40].

2. Because the energy intensity of a manufacturing process decreases with increasing throughput, it follows that one of the simplest ways to improve the efficiency of a manufacturing process or entire system is to increase the throughput. As discussed in Chapter 4, this is precisely what has been witnessed at Analog Devices as production has increased. The energy intensity of individual processes can be bettered by multiplying the number of wafers that can be processed at once, increasing the size of the wafers being processed, and/or decreasing cycle and idle times. Likewise, with all else equal, the energy intensity per product chip can be improved by shrinking the size of the chips, another ongoing trend at Analog Devices and other IC manufacturers.




Figure 5.20: Comparison of the electricity requirements of various manufacturing systems as a function of throughput.

5.5 Closing Thoughts: Efficiency and Sustainability

Unfortunately, sustainability is not defined in terms efficiency but in terms of absolute consumption of resources. Efficiency improvement in and of itself is an admirable goal, but efficiency alone does not guarantee a sustainable manufacturing paradigm. As indicated in the example of the Analog Devices MPD fab, increasing throughput has resulted in a dramatically more efficient manufacturing system. Nonetheless, absolute resource consumption has not declined but has even increased slightly. (Dahmus and Gutowski) find this to be the case across a wide range of industries; increasing production virtually always outpaces increases in efficiency.

This reality might sound a bit dour, but the important lesson is that we must look beyond efficiency to attain the endpoint of sustainability in manufacturing. To be sure, efficiency is an important piece of the sustainability puzzle. But without supply side solutions, the goal of sustainability may continually recede from our grasp.

Chapter 6: Efficiency Improvement Recommendations for ADI's Micromachined Products Division Fab

6.1 Introduction

With a knowledge of the amount and cost of energy and materials being used in the MPD fab and an improved understanding of the end uses for those resources, the task remaining is to provide recommendations for improving the efficiency of this manufacturing facility. In truth, a thorough assessment of strategies for energy and materials use reduction at a manufacturing facility constitutes a consulting effort worthy of an additional Masters thesis or potentially a PhD. Additionally, the intent is not only to improve the efficiency of the MPD fab, but to provide recommendations that can lead to efficiency improvements across the industry. The strategy taken here is thus to provide two recommendations specific to the MPD fab backed by data and calculations, provide some more general recommendations that are suggested by the data acquired during this project, and finally to synthesize general recommendations from other sources.

6.2 Chiller Upgrade

The air conditioning system at the MPD fab demands between 20 and 30% of the electrical power consumption at the facility at any given time, making it an obvious place to start when searching for efficiency improvements. Of this percentage, the majority goes into the fab's two massive 400-ton chillers, making them the largest single users of electricity in the Analog Devices fab. This is characteristic of all semiconductor plants, since control of humidity and temperature in the fabrication facility is critical for minimizing run-to-run variation. (It should be noted that at the time of this writing, Analog Devices had recently added two more chillers to the building roof which are not treated in this study. These chillers will supplement the work done by the two 400-ton basement units.)

The first recommendation put forward by this study would be to replace the two chillers in the basement. The chillers at ADI Cambridge are around twenty-five years old and are not as efficient as modern cooling systems. Furthermore, it was discovered during this study that in the winter, when the facility runs only one chiller at a time and idles the other, the idle chiller draws power at a rate of about 60kW! That power is apparently used to prevent liquid refrigerant from pooling in the oil sump pump and fouling its operation. Yet that means that for approximately eight months out of the year, the idle Analog Devices chiller consumes a large amount of electricity without producing any cooling effect.

Replacing chillers is not a small capital investment by any means. A new 400-ton centrifugal chiller from Carrier runs \$89,000 (Model 19XR, also available with a VFD drive for \$112,000). Alternatively, a new absorption chiller from Carrier, in which the electrically-driven compressor is replaced by a steam input, costs \$210,000. These sort

of price tags often discourage companies from looking at replacing inefficient machinery, but is it worth it?

It's worth it. An estimate of the payback time on this investment can be made by assuming a load scenario and calculating the variable cost of operating the new chillers compared with the old. In this simplified model, several costs are neglected: installation costs (which would increase the capital cost of the new chillers), maintenance costs (which would probably weigh more heavily on the old chillers than the new ones), and the cost of capital (i.e, the expected return of the capital invested in the replacement chillers if it were instead invested elsewhere).

Let's begin by assuming a duty cycle based on the data collected in the facility and a knowledge of the local climate: for 100 days per year (~3.5 months corresponding to summer), two chillers run at 75% load 45% of the time, 50% load 50% of the time, and 25% load 5% of the time, while for the other 265 days one chiller idles while the other runs at 75% load 60% of the time and 50% load 40% of the time. Electrical power data for the Carrier 19XR are given in Table 6.1 at various loads. Power vs. load data information is not available for the current chillers, so it is assumed that the old chillers are on the order of 8% less efficient than the new chillers (this is predicated on the assumption that the chiller was running at 100% load when its highest-ever load was measured at 273 kW [which occurred on a very hot day]). The Carrier 16TJ absorption chiller uses only a small amount of electricity (6.8 kW) but large amounts of steam (7440 lb/hour) [42]. In the following calculations, it is assumed that the absorption chillers run constantly at these inputs. Analog Devices pays an average of \$0.122/kWh of electricity and \$24.44 per million pounds of steam (Mlbs). Note that a ton of cooling corresponds to a heat rate of 12,000 BTU/hr.

Centrifugal Chiller (Carrier Model 19XR)				
Load	100%	75%	50%	25%
Power (kW)	254	167	111	77

Table 6.22: Power vs. load characteristics for a replacement 400-ton centrifugal chiller [41].

Cost of Ownership			
Year	Old Chillers	New Centrifugal Chillers	New Absorption Chillers
Initial Cost	\$0	\$178,000	\$420,000
1	\$252,586	\$368,805	\$437,708
2	\$505,172	\$559,609	\$455,417
3	\$757,758	\$750,414	\$473,125
4	\$1,010,344	\$941,219	\$490,834
5	\$1,262,931	\$1,132,024	\$508,542
6	\$1,515,517	\$1,322,828	\$526,250
7	\$1,768,103	\$1,513,633	\$543,959
8	\$2,020,689	\$1,704,438	\$561,667
9	\$2,273,275	\$1,895,242	\$579,376
10	\$2,525,861	\$2,086,047	\$597,084

Table 6.23: Comparison of total cost of ownership (excluding maintenance, installation, and cost of capital (interest). See paragraph above for description of calculations.

Over the course of 10 years, replacing the two antiquated chillers with new centrifugal chillers yields a savings of almost \$500,000. Note that, were it not for the extra 60kW drawn by the idle chiller during cold months, the break-even point would be twelve years instead of the three years indicated in this calculation. More astonishing is the potential savings gained by transitioning to an absorption chiller. Although it has a very high initial cost, the cost of the steam required to power it is so low that after ten years it would yield a savings of almost \$2 million over the old chillers. Furthermore, since the steam to power an absorption chiller is produced in co-generation plants, the environmental benefit of using an absorption chiller is even more pronounced than the economic benefit.

6.3 Optimize Airflow and Exhaust Velocities

From Figure 4.4, it is evident that the air handling and circulation systems also consume a significant amount of fab power (~135 kW average). The bulk of this power (~85 kW) goes to the two large recirculation fans that provide the pressure drop to drive air through the first floor fab area at a rate of 90 feet per minute. The remainder of the electricity drives the make-up air fan and miscellaneous air-handling units and return fans for both the first floor and second floor fabs. The scrubber fans, which provide the pressure drop for the 3,045 ft/min exhaust flow from the wafer fab, are each ~20kW. Both of the air circulation and the exhaust systems flow rates are set at what is regarded as the industry standard.

Earlier studies have suggested that cleanroom air and exhaust velocities are not optimized, but are typically oversized for the task required [43]. Furthermore, the set point for air and exhaust flow rates seems to lack a scientific basis (i.e., rules of thumb are used in lieu of an assessment of the necessary flow rate required to maximize yield). It has been suggested that both air and exhaust flow rates can be reduced by 15% or more while still maintaining necessary cleanliness and safety standards.

The second recommendation put forward in this paper is to determine empirically the necessary flow rates of air to maintain the requisite particle count and chemical removal rates in the fab and to adjust flow rates accordingly by lowering fan velocities. Reducing the cleanroom air flow rate, for example, allows the circulation fans to run at significantly lower power (fan power decreases with the cube of the air velocity). Additionally, any reduction in power consumption by the circulation fans implies an incremental reduction in cooling requirements for the building chiller. Likewise, reducing the exhaust flow rate yields multiple benefits at the system level: the exhaust (scrubber) fan requires less power, less make-up air is needed which means the make-up air fan can be throttled, and the lower volume flow rate of make-up air results in reduced cooling requirements during humid or hot days.

As an example of the potential energy and cost savings, let's just look at the benefits of reducing the cleanroom air and exhaust velocities by 10%, which will directly impact the power required for the two recirculation fans, the building scrubber fans, and the make-up air fan. In order to throttle these fans, each one will need to have a VFD drive installed at a cost of about \$2,000 each, which includes a rebate from the power company N-Star. Table 6.3 summarizes the calculation results.

Fan	Current Power Draw (kW)	Power Draw with 10% Velocity Reduction (kW)	Energy Savings per Year (kWh)	Cost Savings per Year
2 Recirculation Fans	86.2	62.8	204635	\$24,945
2 Scrubber Exhaust Fans	40.1	29.2	95196	\$11,604
Make-up Air Fan	12.7	9.3	30149	\$3,675

Table 6.24: Energy and cost savings from a 10% reduction in cleanroom exhaust and air flow. Savings based on \$0.122/kWh cost of electricity.

These rough calculations indicate a savings of about \$40,000 per year for a 10% reduction in cleanroom air and exhaust flow rates, excluding corollary benefits such as reduced chiller load and maintenance costs. As is evident in the results, the cost of the VFD drives would be recouped in less than a year. There are many other fans in this facility that are also good candidates for velocity reduction without impacting the quality of service delivered. Throttling these fans would generate additional cost savings.

6.4 Throttle the DI Water Pump

Discussions with Analog Devices staff have indicated that the pump that supplies the fab with DI water is, at 17-18kW, substantially larger than is necessary, meaning much of the water it pumps to the fab either enters the wastewater stream or returns to the DI water reclamation system without having been utilized. The apparent reason the pump has not been previously throttled is that several pieces of fab equipment are set to throw an alarm if their pressure requirements are not met. Thus, in the effort to ensure that these few systems receive sufficient DI flow, other systems receive far more than necessary owing to oversized piping. A rebalancing of the system and the installation of proper valving would allow the VFD drive on the DI pump to throttle its output. It is estimated that greater than 50% of the outgoing DI water goes unused; if the current pump power were reduced by 50%, ADI would save roughly \$9,500/year at the current \$0.12/kWhr price paid for electricity. The third recommendation is thus to redesign the fab's DI water distribution system, possibly by installing a computer-controlled valve system to allow real-time and intelligent distribution of DI water within the fab.

This recommendation follows a general theme in improving fab efficiency, which is to design right-sized systems. In general, the MPD fab and others suffer from a tendency to oversize certain facility support systems, which subsequently leads to the oversizing of other systems. For example, oversized recirculation fans lead to larger cooling requirements and an oversized chiller system. Right-sizing facility support equipment not only reduces capital cost, but also reduces system operating cost. The actual execution of a right-size paradigm requires significant forethought and design and is beyond the scope of this report. When constructing new fabs, however, designing right-sized systems can lead to dramatically reduced upfront and variable costs.

6.5 Avoid Overlapping Heating and Cooling

One of the most common sources of inefficiency in semiconductor plants is overlapping heating and cooling, humidification and dehumidification. A great example of this occurs in wintertime. Cold air is drawn in from the outside and heated and humidified using steam supplied (at cost) from a nearby power plant. This make-up air circulates through the fab where it picks up heat from various tools. It is then run through a heat exchanger to cool the air back to nominal conditions. The heat absorbed in the heat exchanger is subsequently removed by the chiller system and rejected to the environment in an evaporator on the building roof. If one takes a step back and looks at this cycle from a facility-level perspective, it becomes apparent that the heat coming into the facility in the form of steam is promptly exhausted out the roof. Essentially, the chiller is being used to cool the steam input! Instead of using steam to heat the incoming air stream, a superior design would be to use the waste heat from the chiller to heat the incoming air. This could be achieved by adding a heat exchanger or evaporator system into the chiller loop at the point of entry of the make-up air.

The waste heat from the chiller also has other potential uses in the fab that could save cost and energy use. Relatively low temperature hot water (e.g., for use in wet etch baths) could be heated using the exhaust heat stream. Some applications would not merit the cost of the additional piping and a heat exchanger; however, producing low-temperature (and thus low-quality) heat with higher quality sources like steam or electricity is inherently inefficient and should be avoided if at all possible.

6.6 General Recommendations

In addition to the above recommendations, there are a multitude of additional opportunities for energy, material, and cost savings in semiconductor manufacturing. The following is a sampling of the more accessible and cost-effective ideas:

- When replacing old pumps, motors, fans, compressors, or other turbomachines, always purchase the most efficient option. The reduced operational cost will almost always cover any increase in initial capital cost.
- When designing a fab, heat removal from tools and equipment should be done as much as possible using tempered water as opposed to allowing the equipment to dissipate heat to cleanroom air. Roughly five times as much energy is required to remove a unit of heat from cleanroom air as opposed to the tempered water system as a result of the greater efficiency of liquid-liquid heat exchangers and the higher heat capacity of water [31].
- Design water and air supply systems with as little head loss (pipe bends, friction loss) as possible. This is most easily done during initial construction phases and results not only in operating cost savings but typically in construction cost savings [45].
- Purchase efficient tools. As indicated in Chapters 4 and 5, a large portion of fab energy use goes into maintaining tools at the ready during idle time. In the past few years, pressure from semiconductor manufacturers has pushed equipment manufacturers to develop tools that idle at lower power and are generally more

efficient than earlier tools. This is often as simple as turning off pumps when nothing needs to be pumped. When purchasing new or used tools, we recommend factoring the operating cost into the overall cost of the tool. It is likely that new tools will be cheaper over the long run in spite of their higher initial cost owing to reduction in power use (let alone maintenance and repair).

For additional literature on energy saving opportunities in semiconductor fabs, see [27, 31, 32, 43, 46].

6.7 Tool Energy

Historically, semiconductor tools have been designed around increasing yield and throughput with good reason; these factors have the largest and most direct impact on profitability. The inclusion of environmental criteria into tool design is a necessary extension in the quest for a more sustainable manufacturing paradigm. These environmental criteria need not conflict with the goals of increasing yield and throughput. Moreover in an environment where certain chip types are becoming increasingly commoditized (e.g., memory), cutting the variable cost of chip production by making processes more efficient can give a manufacturer the competitive edge necessary to remain profitable.

Over the past few years, equipment manufacturers (represented by the Semiconductor Equipment and Materials International (SEMI) trade group) have begun designing tools that do a better job of managing power consumption. Tool optimization for environmental performance requires cooperation between the tool manufacturers and the tool consumers (i.e., the semiconductor manufacturers) and to that end, we recommend joining with industry trade groups such as SEMATECH or the Semiconductor Industry Association (SIA) to work with semiconductor equipment manufacturers in the design of new equipment. By demanding tools that utilize less energy and materials but that still meet yield and throughput specifications, chip manufacturers can: 1.) cut their cost of production, 2.) hedge against future increases in energy and materials costs, and 3.) strengthen the company's environmental reputation and move further down the path to sustainability.

Chapter 7: The Next Steps

7.1 Semiconductors and Sustainability: Energy and Materials Use in Input Chemicals

With this work and that of previous studies, the energy and materials required used in semiconductor processing is better understood. However, there still exists a major gap in our understanding of the life cycle analysis of semiconductors: to the best of the author's knowledge, no one has yet measured and quantified the energy and material inputs into the production of the chemicals going into semiconductor processing. It is often assumed that the production of these materials is quite energy intensive, yet the breadth of chemicals used and the opaqueness of chemical companies has made it difficult to measure any but the most basic materials.

An interesting and viable research subject would be to do for semiconductor chemical production what has been done here for semiconductor manufacturing, namely to quantify the energy and materials going into the production of various chemicals. Particularly intriguing, and not yet studied empirically (see [47] for a theoretical study), is the impact of purity (or concentration) on energy requirements for semiconductor chemicals. It is assumed that energy requirements scale with chemical purity, but how it scales has not been studied experimentally. These questions are particularly applicable to the LCA of semiconductors given the large amounts of high-purity chemicals used in the industry.

Such a project would not be elementary given the hurdles faced in finding cooperative companies. Additionally, in facilities that manufacture multiple chemicals, separating data on the inputs to each specific chemical might prove challenging. Such a project is possible, however; during the project described in earlier chapters, major inroads were made with Matheson Tri-Gas, a dominant player in the semiconductor chemical market. We also would recommend looking into semiconductor-specific chemicals such as resists and developers (Rohm & Haas is a prominent manufacturer, for starters). There is virtually no information on the resource requirements of the production of these materials; a better understanding of these and other semiconductor chemicals would be quite beneficial to the field of sustainability.

7.2 Next Steps for Analog Devices

With data in hand and a basic idea of problem areas to attack in Analog Devices's fabs, the next step is to make those improvements in existing fabs that make the most economic and environmental sense. The next step would be to experiment with additional improvements that might not show the same immediate economic dividends but will have an environmental payback; in this way, the company is reducing its exposure to the risk that emissions or hazardous waste disposal will become more costly. Finally, during the planning and design of new fabs (long before construction begins), it is imperative to work with a clean room consultant with credentials in energy efficiency.

The long term benefits of a more efficient fab make fantastic economic sense and should not be foregone in haste to erect new facilities.

Additionally, the essence of reducing resource consumption and improving efficiency comes down to two words: monitoring and measuring. Before any changes are made to a facility, a system should be put in place to allow for the monitoring of the entire facility and as many pieces of equipment as possible. In addition to energy usage, other metrics such as flow rates, pressure drops, and cooling output should also be measured. For the small cost of this monitoring equipment, it becomes a mere exercise to understand where power and materials are being used in a facility and how much they are costing the host company. Such knowledge makes it a cinch to determine whether or not to replace equipment. Without such knowledge, it becomes very difficult to determine whether or not the replacement or upgrading of a given piece of equipment has had the desired economic and environmental benefit.

Analog Devices has shown the will to lead an industry down the path toward a more sustainable production system. It is a safe bet that such an effort will benefit both the environment and the bottom line.

Appendix A: Measurements of Energy Use in Carbon Nanotube Growth

A.1 Nanotube Growth Using a Radiative Furnace and a Substrate-Heated Furnace

At an earlier stage in the author's field work, data on the energy and materials used in the growth of carbon nanotubes were gathered from the set-up of John Hart (<http://www-personal.umich.edu/~ajohnh/>). The first test run was of the growth of 10 nm diameter single-walled carbon nanotubes using an ethylene precursor. The laboratory set-up included a radiatively-heated glass tube into which a silicon substrate up to 30cm² in area was placed. A mixture of helium, hydrogen, and ethylene was flowed through the tube over the course of a 46 minute growth time that included a 10-minute standby period, a 20-minute ramp-up and temperature stabilization, a 6-minute growth, and a 10-minute cool-down.

The second test was conducted on a substrate-heated furnace. In this setup, instead of heating all of the gases flowing through the deposition tube radiatively, the substrate itself was resistively heated. This design is dramatically more efficient; not only is it no longer necessary to heat all of the gas flowing through the chamber to the same high temperature, but the ramp-up period is much shorter (only 1 minute). Instead of a 30cm² substrate, only a 1.5 cm² substrate was used. The essential statistics from this test and the radiative furnace test are given in Table A.1.

It should be understood that these data are from a laboratory set-up. The set-up therefore has a much lower energy overhead than would a large-scale manufacturing set-up. For starters, there is no rate requirement in this setup, meaning that the nanotubes can be grown as slowly as needed. Rate requirements are often one of the factors that drives industrial processes to be less efficient [48]; oxidation of silicon could be done at lower temperatures, for example, but because it is desired that the oxidation be relatively fast, higher temperatures must be used. Secondly, there is no automation associated with this set-up, another factor that often drives energy requirements higher. Nonetheless, given the scarcity of data on energy requirements in carbon nanotube growth, these data are valuable.

	Energy Intensity (J/kg)	Growth Rate (kg/s)
Radiative Furnace	$1.97 \cdot 10^{10}$	$1.1 \cdot 10^{-7}$
Substrate-Heated Furnace	$5.05 \cdot 10^{10}$	$5.5 \cdot 10^{-9}$

Figure A.21: Energy intensity and process rate for two different laboratory-scale carbon nanotube growth setups. Parameters used in these calculations: volume fraction = 0.01, density = 2.2 g/cm³, growth rate = 0.1 mm/min.

Appendix B: Summary Power Data from Bottom-Up Tests

Table B.1 presents the data gathered on power consumption by tools and panels power panels, collated and reduced for a condensed presentation. The specific application and tool manufacturer has been concealed to protect Analog Devices. Similarly, the identity of the various tools and components in the power panels has been omitted; instead, the function of the tools in each power panel has been generalized by process area. Table B.2 gives the collated data for the facility support systems.

Tool/System/Panel	Test Length	Process Area (Power in kW)										Total
		Dry Etch	Wet Etch	CVD	Diffusion	Photo	Metallization	Ion Implantation	Metrology	Trim & Cap	Other	
PECVD Tool	2h, 37min	0	0	8.35	0	0	0	0	0	0	0	8.35
Dry Etch Tool	24h	33.58	0	0	0	0	0	0	0	0	0	33.58
Sputter Deposition Tool	23h, 15min	0	0	0	0	0	21.76	0	0	0	0	21.76
Vertical Furnaces (2 CVD Furnaces)	21h, 50min	0	0	7.53	0	0	0	0	0	0	0	7.53
Furnace A (3 CVD Furnaces, 1 Diffusion Furnace)	91h, 37min	0	0	21.72	7.24	0	0	0	0	0	0	28.95
Furnace E (1 CVD Furnace, 3 Diffusion Furnaces)	24h, 10min	0	0	5.56	31.48	0	0	0	0	0	0	37.03
Furnace F (4 Diffusion Furnaces)	31h, 20min	0	0	0	30.27	0	0	0	0	0	0	30.27
Ion Implanter 1	20h, 20min	0	0	0	0	0	0	20.84	0	0	0	20.84
Ion Implanter 2		0	0	0	0	0	0	20.84	0	0	0	20.84
Photolithography Tool 1	4h, 58min	0	0	0	0	16.77	0	0	0	0	0	16.77
Photolithography Tool 2	20h	0	0	0	0	6.30	0	0	0	0	0	6.30
Machine Panel 300-01	46h, 29min	22.63	0	0	0	0	0	0	0	0	0	22.63
Machine Panel 118-02	22h, 27min	0	0	0	0	10.00	0	0	1.11	0	0	11.11
Machine Panel 00-01	24h, 5min	20.18	0	0	0	0	0	0	0	0	0	20.18
Machine Panel LAB-01	66h, 55min	1.90	0	0	0	1.14	13.68	0	0	0	2.28	19.01
Machine Panel 500-01	24h	0	0	9.87	0.61	0	0	0	1.71	0	0	12.18
Machine Panel 400-01	24h	0	5.29	0	0	0.79	0	0	0	0	0	6.08
Machine Panel 400-02	24h, 5min	0	12.30	0	0	0	0	0	0.65	0	0	12.95
Machine Panel 200-02	3h	0	0.76	0	0	1.34	0	0	0.82	0	0	2.91
Machine Panel 600-01	65h, 5min	0	3.90	1.99	0	0	0	0.000	0	0	1.76	7.64
Machine Panel 500-02	43h, 35min	4.49	0	5.49	0	0	0	0	0	0	0	9.98
Machine Panel 300-02	24h, 5min	14.73	0	1.82	0	0	0	0	0	0	0	16.55
Machine Panel A&T-01	28h	0	0	0	0	3.16	0	0	0	0	0	3.16
Machine Panel 100-02	26h, 55min	23.00	0	0	0	3.14	0	0	0	0	0	26.14
Machine Panel 1903A	65h, 35min	0	0	0	0	0	0	0	0	61.43	0	61.43
Machine Panel 1903B	26h, 5min	0	0	0	0	0	0	0	0	35.56	0	35.56
Receptacle Panel 1305	24h, 10min	13.49	0	0	0	0	0	0	0	0	0	13.49
Receptacle Panel 1906	22h, 35min	0	0	0	0	0	0	0	0	23.56	0	23.56
Receptacle Panel 1902A	27h, 10min	0	0	0	0	0	0	0	0	5.89	0	5.89
Receptacle Panel 1303	Spot Test	0	0	2.25	0	0	0	0	0	0	0	2.25
Transformer 204	20h, 10min	0	0	0	0	0	0	0	0	26.09	0	26.09
Transformer 3	20h, 40min	0	0	0	0	0	0	0	0	5.55	0	5.55
Transformer 102	6h, 55min	0	0	0	0	0	0	0	0	1.54	0	1.54
Transformer 106	17h	0	0	0	0	0	0	0	0	8.59	0	8.59
Transformer 201	22h, 5min	0	0	0	0	0	0	0	0	25.64	0	25.64
Totals		134.0	22.2	64.6	69.6	42.6	35.4	41.7	4.3	193.8	4.0	612.32

Figure 1: Summary power data from tools and control panels.

System	Test Length	Power Consumption (kW)	# of Systems	Total (kW)
Scrubber Fans	6h, 13min	20.04	2	40.07
Scrubber Pump	17h, 50min	4.28	1	4.28
Recirculation Fans ¹	3h, 21min	43.07	2	86.15
RO Pump #1 ²	24h, 10min	31.12	1	31.12
RO Pump #2	23h, 25min	31.75	1	31.75
DI Pump #1	21h, 55min	17.88	1	17.88
Booster Pump for DI Water	25min	5.07	1	5.07
Reclaim Pump ³	40min	1.06	1	1.06
Water Room Panel 1	25h, 15min	33.99	1	33.99
Water Room Panel 2	21h, 55min	2.15	1	2.15
Tempered Water Pumps	4h	5.20	2	10.40
Air Compressor	3h, 51min	97.96	1	97.96
First Floor Chiller Booster Pump	32h, 45min	7.35	1	7.35
Chiller Primary Pumps ⁴	2h, 31min	13.34	2	26.68
Chiller Secondary Pumps ⁵	2h, 48min	12.09	2	24.18
Chiller Tower Pump ⁶	23h, 35min	32.40	2	64.80
Cooling Tower Fan	5h, 20min	17.08	2	34.16
Chiller	140h, 54min	129.09	2	258.18
VCS Hot Water Heater	18h, 13min	12.67	1	12.67
Heating Water Pump	6h	5.39	1	5.39
Receptacle Panel 1303	Spot Test	0.75	1	0.75
Air Handling Unit 207	8min	0.55	4	2.20
Return Fan 3	20h, 5min	3.58	2	7.17
Air Handling Unit 211	8min	1.47	4	5.87
Make-Up Air Fan	5h, 11min	12.69	1	12.69
Air Handling Unit 202	16h, 50min	13.73	1	13.73
Return Fan (F-202)	6h, 30min	4.46	1	4.46
Air Handling Unit 205	15h, 35min	2.95	1	2.95

Figure 2: Summary power data from facility support equipment.

¹ The recirculation fans provide the pressure drop which drives air through the cleanroom.

² The RO (reverse osmosis) pumps are used in DI water production.

³ The reclaim pump is used in the building's wastewater reclamation system, which salvages a portion of the water that would be discarded, cleans it, and sends it back to the DI system.

⁴ The chiller primary pumps are located just before the chillers in the cooled water loop and "pull" the cooled water from the chillers through the facility.

⁵ The chiller secondary pumps are located just after the chillers in the cooled water loop and "push" the cooled water from the chillers through the facility.

⁶ The chiller tower pumps drive the water loop that is used to remove heat from the chiller and dissipate it on the roof.

Bibliography

1. Cullen, Steve. *Semiconductor Market Outlook*. In-Stat. Apr. 2004.
2. U.S. Census Bureau. "Total Midyear Population for the World: 1950-2050." *International Data Base*. 16 July 2007. 5 Mar. 2008 <<http://www.census.gov/ipc/www/idb/worldpop.html>>.
3. Matos, Grecia and Lorie Wagner. "Consumption of Materials in the United States, 1900-1995." *Annual Review of Energy and the Environment*. 1998. Vol. 23: 107-122. 4 Mar. 2008 <<http://arjournals.annualreviews.org/doi/pdf/10.1146/annurev.energy.23.1.107>>.
4. Energy Information Administration (EIA). "World Primary Energy Consumption (BTU), 1980-2005." *International Energy Annual 2005*. 21 Sep. 2007. 5 Mar. 2008 <<http://www.eia.doe.gov/pub/international/iealf/tablee1.xls>>.
5. Intel Corporation. "Moore's Law 40th Anniversary." 22 Jan. 2008 <http://www.intel.com/pressroom/kits/events/moores_law_40th/index.htm?iid=tech_mooreslaw+body_preskit>.
6. U.S. Census Bureau. "2006 Annual Survey of Manufactures." *American Fact Finder*. 31 Dec. 2006. 11 Mar. 2008 <http://factfinder.census.gov/servlet/IBQTable?_bm=y&_skip=100&ds_name=AM0631GS101&_lang=en>.
7. Semiconductor Industry Association (SIA). "Global Chip Sales Hit \$255.6 Billion in 2007." 1 Feb. 2008. 11 Mar. 2008 <http://www.sia-online.org/pre_release.cfm?ID=464>.
8. Semiconductor Industry Association (SIA). "Total Semiconductor World Market Sales & Shares." 11 Mar. 2008 <<https://www.sia-online.org/downloads/shares.pdf>>.
9. Moore, Gordon E. "No Exponential is Forever: But "Forever" Can Be Delayed!" 2003 IEEE International Solid-State Circuits Conference, Session 1 Plenary. 22 Jan. 2008 <<http://ieeexplore.ieee.org/xpl/multimedia.jsp?arnumber=1234194&isnumber=27661>>.
10. Mazurek, Jan. *Making Microchips: Policy Globalization, and Economic Restructuring in the Semiconductor Industry*. Cambridge, MA: The MIT Press, 1999.
11. Energy Information Administration. "Consumption of Energy (Site Energy) for All Purposes (First Use) for Selected Industries, 1998 and 2002." *Manufacturing Industry Trend Data 1998 and 2002 (NAICS)*. May 2006. 9 Sep. 2006 <http://www.eia.doe.gov/emeu/efficiency/mecs_trend_9802/mecs9802_table1a.html>.
12. Jones, Alissa. *The Industrial Ecology of the Iron Casting Industry*. M.S. Thesis. Massachusetts Institute of Technology, Department of Mechanical Engineering, Cambridge, Massachusetts. 2007.
13. Thiriez, Alexandre. *An Environmental Analysis of Injection Molding*. M.S. Thesis. Massachusetts Institute of Technology, Department of Mechanical Engineering, Cambridge, Massachusetts. 2006.
14. Gutowski, Timothy, et. al. "A Thermodynamic Characterization of Manufacturing Processes." *Proceedings of the 2007 IEEE International Symposium on Electronics and the Environment*. 7-10 May 2007: 137-142.
15. Dahmus, Jeff. *Applications of Industrial Ecology: Manufacturing, Recycling, and Efficiency*. PhD Dissertation. Massachusetts Institute of Technology, Department of Mechanical Engineering, Cambridge, Massachusetts. 2007.

16. Cooke, Mike. "Equipment Sales Consensus Gives Marginal 2007 Increase." *Semiconductor Fabtech*. Q3 2007: 34.
17. "Advantages and Challenges Associated with the Introduction of 450mm Wafers." International Technology Roadmap for Semiconductors (ITRS). White Paper. June 2005. 19 Mar. 2008 <<http://www.itrs.net/papers.html>>.
18. "Silicon Wafer Shipments Experience Growth for the Sixth Year in a Row." 5 Feb. 2008. SEMI. 19 Mar. 2008 <http://wps2a.semi.org/wps/portal/_pagr/103/_pa.103/248?dFormat=application/msword&docName=P043260>.
19. Plummer, James D., Michael D. Deal, and Peter B. Griffin. *Silicon VLSI Technology: Fundamentals, Practice, and Modeling*. Upper Saddle River, NJ: Prentice Hall, 2000.
20. Williams, Eric D., Robert U. Ayres, and Miriam Heller. "The 1.7 Kilogram Microchip: Energy and Material Use in the Production of Semiconductor Devices." *Environmental Science and Technology*. Vol. 36, No. 24 (2002): 5504-5510.
21. Microelectronics and Computer Technology Corporation (MCC). *Environmental Consciousness: A Strategic Competitiveness Issue for the Electronics and Computer Industry*. Austin, March 1993.
22. Ammenheuser, Jim and Phil Naughton. "Facilities Council: History and Benefits of Membership." ISMI: International SEMATECH Manufacturing Initiative. December 2004.
23. Yao, Marissa, et. al. "Comparative Analysis of the Manufacturing and Consumer Use Phases of Two Generations of Semiconductors." *Proceedings of the 2004 IEEE International Symposium on Electronics and the Environment*. 10-13 May 2004: 97-103.
24. Murphy, Cynthia Folsom and David T. Allen. *Development of Life Cycle Inventory Modules for Semiconductor Processing*. EPA Report. July 2004.
25. Murphy, Cynthia et. al. "Development of Parametric Material, Energy, and Emission Inventories for Wafer Fabrication in the Semiconductor Industry." *Environmental Science and Technology*. Vol. 37, No. 23 (2003): 5373-5382.
26. Boyd, Sarah, David Dornfeld, and Nikhil Krishnan. "Life Cycle Inventory of a CMOS Chip." *Proceedings of the IEEE International Symposium on Electronics and the Environment*. (2006): 253-257.
27. Naughton, Philip. "Measurement of Conservation of Energy by Semiconductor Manufacturing Equipment and Setting of Targets for Improvements." *Semiconductor Fabtech*. Q1, 2006.
28. "SEMI S22-0706a: Safety Guideline for the Electrical Design of Semiconductor Manufacturing Equipment." SEMI. 23 Aug. 2007 <www.semi.org>.
29. "SEMI S23-0705: Guide for Conservation of Energy, Utilities and Materials Used by Semiconductor Manufacturing Equipment." Mar. 2005. SEMI. 23 Aug. 2007 <www.semi.org>.
30. *World Wafers Markets*. Frost & Sullivan. 9 May 2007. (by subscription only)
31. Osborne, Mark. "Making the Most of Energy Conservation." *Semiconductor Fabtech*. 26th Ed. Q2 2005. Available at: <http://www.fabtech.org/component/option,com_remository/Itemid,93/func,select/id,201/>.
32. "Energy Efficient Cleanroom Information Site." 24 March 2000. Lawrence Berkeley National Laboratory. 23 Feb. 2007 <<http://ateam.lbl.gov/cleanroom/technical.html>>.

33. Isaacs, J. A., A. Tanwani, and M. L. Healy. "Environmental Assessment of SWNT Production." *Proceedings of the 2006 IEEE International Symposium on Electronics and the Environment*. 8-11 May 2006: 38-41.
34. Gyftopoulos, Elias P. and Gian Paolo Beretta. *Thermodynamics: Foundations and Applications*. Mineola, NY: Dover Publications, Inc., 2005.
35. Bejan, Adrian. *Advanced Engineering Thermodynamics*. New York: John Wiley & Sons, 1988.
36. Szargut, J., Morris, D.R., & Steward, F.R. *Exergy Analysis of Thermal, Chemical, and Metallurgical Processes*. New York: Hemisphere Publishing Corporation and Springer-Verlag, 1988.
37. Sorin, M., Lambert, J., and Paris, J. "Exergy Flow Analysis in Chemical Reactors." *Transactions of the Institution of Chemical Engineers* 76. 1998: 389-395.
38. Poling, Bruce E., John M. Prausnitz, and John P. O'Connell. *The Properties of Gases and Liquids*, 5th ed. New York: McGraw-Hill, 2001.
39. Pehlke, David Robert. High Temperature Thermodynamics of the Silicon, Nitrogen, Silicon-Nitride System. Masters of Science in Materials Science Thesis, Massachusetts Institute of Technology, Cambridge, MA. May 1958.
40. Perraton, Jonathan. "Heavy Constraints on a "Weightless World"?: Resources and the New Economy." *The American Journal of Economics and Sociology*. Vol. 65, No. 3. 2006: 641-691.
41. Carrier Corporation. *19XR Evergreen Centrifugal Chillers*. Product Description. 20 June 2008 <http://www.commercial.carrier.com/commercial/hvac/product_description/0,,CLI1_DIV12_ETI434_PRD1,00.html>.
42. Carrier Corporation. *16TJ Single-Effect Steam-Fired Absorption Chiller*. Product Description. 20 June 2008 <http://www.commercial.carrier.com/commercial/hvac/product_description/0,,CLI1_DIV12_ETI434_PRD1226,00.html>.
43. Shamshoian et. al. "High-Tech Means High Efficiency: The Business Case for Energy Management in High-Tech Industries." Nov. 2005. Lawrence Berkeley National Laboratory, Paper LBN-59127 15 Oct. 2007 <<http://repositories.cdlib.org/cgi/viewcontent.cgi?article=3852&context=lbnl>>.
45. Robertson, Chris. "Design for Energy Efficiency Adds Value to Semiconductor Company Shareholders." *Semiconductor Fabtech*. 10th Ed., June 2001.
46. "Energy and Water Efficiency for Semiconductor Manufacturing." Pacific Northwest Pollution Prevention Resource Center. Feb. 2000. 30 April 2008 <<http://www.p2pays.org/ref/04/03271/>>.
47. Gutowski, Timothy, et. al. "Bayesian Material Separation Model with Applications to Recycling." *Proceedings of the 2007 IEEE International Symposium on Electronics and the Environment*. 7-10 May 2007: 233-238.
48. Baumgartner, Stefan and Jakob de Swaan Arons. "Necessity and Inefficiency in the Generation of Waste: A Thermodynamic Analysis." *Journal of Industrial Ecology*. Vol. 7, No. 2 (2003): 113-123.