

Energy Demands and Efficiency Strategies in Data Center Buildings

Arman Shehabi

Environmental Energy Technologies Division
Indoor Environment Department
Lawrence Berkeley National Laboratory
Berkeley, CA 94720

September 2009

This work was supported by the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.

Energy Demands and Efficiency Strategies in Data Center Buildings

by

Arman Shehabi

B.A. (University of California, San Diego) 1997

M.S. (Stanford University) 2000

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy
in

Engineering Science –

Civil and Environmental Engineering

in the

GRADUATE DIVISION

of the

UNIVERSITY OF CALIFORNIA, BERKELEY

Committee in charge:

Professor William W Nazaroff, co-Chair

Professor Arpad Horvath, co-Chair

Professor Ashok J. Gadgil

Professor Edward Arens

Fall 2009

This dissertation of Arman Shehabi is approved:

co-Chair

Date

co-Chair

Date

Date

Date

University of California, Berkeley

Energy Demands and Efficiency Strategies in Data Center Buildings

© 2009

by Arman Shehabi

Abstract

Energy Demands and Efficiency Strategies in Data Center Buildings

by

Arman Shehabi

Doctor of Philosophy in Engineering Science –

Civil and Environmental Engineering

University of California, Berkeley

Professor William W Nazaroff, co-Chair

Professor Arpad Horvath, co-Chair

Information technology (IT) is becoming increasingly pervasive throughout society as more data is digitally processed, stored, and transferred. The infrastructure that supports IT activity is growing accordingly, and data center energy demands have increased by nearly a factor of four over the past decade. Data centers house IT equipment and require significantly more energy to operate per unit floor area than conventional buildings. The economic and environmental ramifications of continued data center growth motivate the need to explore energy-efficient methods to operate these buildings. A substantial portion of data center energy use is dedicated to removing the heat that is generated by the IT equipment. Using economizers to introduce large airflow rates of outside air during favorable weather could substantially reduce the energy consumption of data center cooling. Cooling buildings with economizers is an established energy saving measure, but in data centers this strategy is not widely used,

partly owing to concerns that the large airflow rates would lead to increased indoor levels of airborne particles, which could damage IT equipment. The environmental conditions typical of data centers and the associated potential for equipment failure, however, are not well characterized. This barrier to economizer implementation illustrates the general relationship between energy use and indoor air quality in building design and operation. This dissertation investigates how building design and operation influence energy use and indoor air quality in data centers and provides strategies to improve both design goals simultaneously.

As an initial step toward understanding data center air quality, measurements of particle concentrations were made at multiple operating northern California data centers. Ratios of measured particle concentrations in conventional data centers to the corresponding outside concentrations were significantly lower than those reported in the literature for office or residential buildings. Estimates using a material-balance model match well with empirical results, indicating that the dominant particle sources and losses – ventilation and filtration – have been characterized. Measurements taken at a data center using economizers show nearly an order of magnitude increase in particle concentration during economizer activity. However, even with the increase, the measured particle concentrations are still below concentration limits recommended in most industry standards.

The research proceeds by exploring the feasibility of using economizers in data centers while simultaneously controlling particle concentrations with high-quality air filtration. Physical and chemical properties of indoor and outdoor particles were analyzed at a data center using economizers and varying levels of air filtration efficiency.

Results show that when improved filtration is used in combination with an economizer, the indoor/outdoor concentration ratios for most measured particle types were similar to the measurements when using conventional filtration without economizers. An energy analysis of the data center reveals that, even during the summer months, chiller savings from economizer use greatly outweigh the increase in fan power associated with improved filtration. These findings indicate that economizer use combined with improved filtration could significantly reduce data center energy demand while providing a level of protection from particles of outdoor origin similar to that observed with conventional design.

The emphasis of the dissertation then shifts to evaluate the energy benefits of economizer use in data centers under different design strategies. Economizer use with high ventilation rates is compared against an alternative, water-side economizer design that does not affect indoor particle concentrations. Building energy models are employed to estimate energy savings of both economizer designs for data centers in several climate zones in California. Results show that water-side economizers consistently provide less energy savings than air-side economizers, though the difference in savings varies by location. Model results also show that conventional limits on humidity levels in data centers can restrict the energy benefits of economizers.

The modeling efforts are then extended to estimate national data center energy use. Different size data centers are modeled to represent the national variation in efficiency and operation of associated mechanical equipment. Results indicate increased energy efficiency opportunities with larger data centers and highlight the importance of temperature setpoints in maximizing economizer efficiency. A bottom-up modeling

approach is used to estimate current (2008) United States data center energy use at nearly 62-70 billion kWh annually. The model indicates that more about 65-70% of this energy demand can be avoided through energy efficient IT and cooling infrastructure design, equivalent to an annual energy efficiency resource of approximately 40-50 billion kWh available at a national level. Within the context of greenhouse gas emissions, benefits can be significantly increased by incorporating site location into energy-efficient design strategies.

The framework of this dissertation contributes to general building energy efficiency efforts by shifting the perspective of building design to address indoor and outdoor environmental impacts simultaneously, ensuring that one design goal does not eclipse the other. More specifically, the results presented here outline opportunities to temper the growing data center energy demand, so that IT can evolve into an energy efficient utility with the potential to facilitate a more sustainable expansion of goods and services.

Acknowledgements

This dissertation is the culmination of several years of effort, but it was not an individual effort. As I get ready to finally file my dissertation, I know I have reached this position because of the contributions of so many along the way.

Above all, I am grateful for the mentorship, enthusiasm, and inspiration from my advisors, Professors Bill Nazaroff and Arpad Horvath. Bill's prodigious editing has influenced every page of this dissertation. I will always strive to emulate the structured approach and self-imposed standard of excellence that Bill applies to all of his commitments. Arpad helped me find my way into Berkeley and his genuine concern for my development and success motivated me throughout my doctoral work. Bill and Arpad's mentoring styles compliment each other well and what I have learned from them will guide me well beyond graduate school. I hope future students will get to benefit from this advising dream team.

I also thank my dissertation committee members, Professors Ed Arens and Ashok Gadgil. Ed gave amiable support for my research ideas since my first year at Berkeley. Along with being on my dissertation committee, Ashok also provided valuable experience and insight while directing my research at Lawrence Berkeley National Laboratory. My first experience at the Lab was working under Phil Price, whose positive attitude toward research, and life in general, I will always remember and search for in myself. Bill Tschudi and Dale Sartor helped me through the data center work at the Lab.

Pacific Gas& Electric and the University of California Energy Institute funded much of the research contained in this dissertation.

My research was also influenced by mentorship I received before I arrived at Berkeley. Professor Gil Masters at Stanford University first piqued my interest in building energy issues. I repeatedly applied the consulting skills that I developed under Charles Eley at Architectural Engineering Corporation to my doctoral research.

Words cannot express the gratitude I feel towards my family. Reaching this milestone in my life is the result of their support and guidance. The love from my parents is the foundation from which all my successes are based. Soroush did a great job as a big brother and helped build the resolve in my personality to never give up. My grandparents have always represented a source of caring and love. Grandpa is really the first engineer in the family. Agha Joon was the initial inspiration for me to pursue a Ph.D.

So many friends made my time in Berkeley more enjoyable. I cannot image making it through Berkeley without my officemate and project partner, Bev. Tonya and Bryan often provided a much needed antidote to excessive academia. All the Airheads, Sharon, Dev, Seema, George, Josh, Juli, and Nasim made the cold grey concrete of Davis Hall a bit warmer and brighter every day. The Airheads before me, Priya and Rengie, provided valuable assistance and advice. Drew and Tim give me high hopes for the future of Airheads.

Finally, I thank Stephanie, whose love, support, and understanding helped me through my dissertation writing, and who was always there to remind me what really matters.

Arman Shehabi

December 17, 2009

Table of Contents

Abstract.....	1
Acknowledgements.....	i
Table of contents.....	iii
List of tables.....	vi
List of figures.....	x
Chapter 1: Introduction.....	1
1.1. Building energy use.....	1
1.2. Building energy and indoor environmental quality.....	2
1.3. Data center buildings.....	4
1.4. Growing data center demand.....	6
1.5. Dissertation scope.....	9
1.6. Dissertation format.....	10
Chapter 2: Particle concentrations in data centers.....	14
2.1. Introduction.....	14
2.2. Methods.....	16
2.2.1. Study sites.....	16
2.2.2. Experimental protocol.....	20
2.2.3. Modeling indoor particle concentrations.....	22
2.3. Results and discussion.....	27
2.3.1. Measured particle concentrations.....	27
2.3.2. Particle sources and sinks.....	43
2.3.2. Sulfate predictions.....	46

2.4. Conclusions.....	47
2.A. Appendix: Additional Data Center Particle Measurements.....	49
 Chapter 3: Combining economizers with improved filtration to save energy and protect equipment in data centers.....	 77
3.1. Introduction.....	77
3.2. Methods.....	79
3.2.1. Experimental design.....	79
3.2.2. Real-time measurements.....	80
3.2.3. Filter-based particle measurements.....	82
3.2.4. Energy calculations.....	89
3.3. Results and discussion.....	92
3.3.1. Particle concentrations.....	92
3.3.2. Energy use.....	112
3.4. Conclusions.....	116
 Chapter 4: Energy implications of economizer use in California data centers.....	 118
4.1. Introduction.....	118
4.2. Methods.....	120
4.2.1. Data center design scenarios.....	120
4.2.2. Energy modeling protocol.....	126
4.2.2. Power Use Effectiveness (PUE).....	130
4.3. Results and discussion.....	133
4.4. Conclusions.....	140

Chapter 5: National energy demand and potential energy savings in data centers	141
5.1. Introduction.....	141
5.2. Data and methods.....	142
5.2.1. IT equipment modeling procedure	142
5.2.2. IT equipment energy savings estimation.....	153
5.2.3. Non-IT equipment modeling procedure	155
5.3. Results and discussion	169
5.3.1. National energy estimates	169
5.3.2. PUE calculations	175
5.2.3. Carbon intensity	204
5.4. Conclusions.....	215
 Chapter 6: Conclusion.....	 219
6.1. Economizer implementation and energy savings potential	219
6.2. Opportunities looking forward.....	228
6.2.1. Expansion of data center operating conditions.....	228
6.2.2. Improvement of metrics	232
6.2.3. Sourcing of electricity and greenhouse-gas emissions.....	235
6.2.4. Embodied energy.....	235
6.3. An opportunity for sustainability	237
 References.....	 239

List of Tables

2-1.	Characteristics of three data center sites.....	21
2-2.	Average measured indoor particle concentrations at eight northern California data centers ($\mu\text{g}/\text{m}^3$).....	28
2-3.	Average indoor modeled and measured particle concentrations at three data center sites ($\#/ \text{m}^3$).....	44
3-1.	Average indoor and outdoor concentrations of OPC-measured particles, sorted by economizer activity and filtration efficiency rating ...	96
3-2.	Average indoor and outdoor filter-based particulate sulfate concentrations, sorted by economizer activity and filtration efficiency rating	96
3-3.	Average indoor and outdoor filter-based particulate nitrate concentrations, sorted by economizer activity and filtration efficiency rating.	96
3-4.	Average outdoor and indoor air temperatures ($^{\circ}\text{C}$) during the August 2008 study, sorted by filter type and according to whether the economizer was “on” or “off”.....	97
3-5.	Average indoor and outdoor filter-based particulate chloride concentrations, sorted by economizer activity and filtration efficiency rating	97
3-6.	Average indoor and outdoor filter-based particulate ammonium concentrations, sorted by economizer activity and filtration efficiency rating	102
3-7.	Moles of ammonium collected on citric acid impregnated cellulose filters compared to moles of nitrate and chloride collected on nylon filters	102
3-8.	Average indoor and outdoor quartz measured black carbon particle concentrations, sorted by economizer activity and filter type	103
3-9.	Average indoor and outdoor aethalometer measured black carbon particle concentrations, sorted by economizer activity and filter type scenarios.....	103
3-10.	Average indoor and outdoor of gravimetrically-measured particulate matter concentrations, sorted by economizer activity and filter type	103

3-11.	Average indoor and outdoor quartz measured total carbon particle concentrations, sorted by economizer activity and filter type	108
3-12.	Measurement time periods, average total building power, and the average power specific to IT equipment for each filter type	113
4-1.	Data center characteristics common to all design scenarios	128
4-2.	Data center fan properties	128
4-3.	Component peak power consumption for a typical computer server (Fan et al., 2007)	132
4-4.	Estimated PUE values specific to design scenario and location	134
5-1.	IDC defined space type categories used for tracking computer server sales and shipments (Bailey et al., 2007; Brown et al., 2007)	144
5-2.	Distribution of 2005 U.S. server stock by server class and data center space type, based on IDC server shipment and sales data (Brown et al., 2007)	145
5-3.	2007 U.S. stock of computer storage units distributed by data center space type, based on Seagate Technology external HDD shipment and sales data (Brown et al., 2007)	152
5-4.	Estimated market penetration of IT efficiency measures and calculated PUE values for the standard operations and energy efficient scenarios	154
5-5.	Regional allocation of computer servers located in buildings identified to have significant data center activity (derived from CBECS, 2003)	158
5-6.	U.S. metropolitan areas with largest concentration of existing data centers (Brown et al., 2007)	158
5-7.	Building and mechanical design parameters of the server room space type model for the standard operation and energy efficiency scenarios ..	163
5-8.	Building and mechanical design parameters of the localized data center space type model for the standard operation and energy efficiency scenarios	164
5-9.	Building and mechanical design parameters of the mid-tier data center space type model for the standard operation and energy efficiency scenarios	167

5-10.	Building and mechanical design parameters of the enterprise data center space type model for the standard operation and energy efficiency scenarios.....	168
5-11.	Current (2008) energy and energy efficiency potential of national data center energy use, by space type and equipment component	171
5-12.	Total data center energy use (IT + non-IT) separated by climate region for the Current Practices and Economizer Plus scenarios	175
5-13.	Annual energy modeling results for server room energy efficiency for a Baseline and Economizer scenario.....	178
5-14.	Annual energy modeling results for a localized data center in San Francisco, CA	183
5-15.	Annual energy modeling results for a localized data center in Seattle, WA.....	184
5-16.	Annual energy modeling results for a localized data center in Chicago, IL	185
5-17.	Annual energy modeling results for a localized data center in Dallas, TX	186
5-18.	Annual energy modeling results for a localized data center in Richmond, VA	187
5-19.	Annual energy modeling results for a mid-tier data center in San Francisco, CA	190
5-20.	Annual energy modeling results for a mid-tier data center in Seattle, WA.....	191
5-21.	Annual energy modeling results for a mid-tier data center in Chicago, IL.....	192
5-22.	Annual energy modeling results for a mid-tier data center in Dallas, TX	193
5-23.	Annual energy modeling results for a mid-tier data center in Richmond, VA	194
5-24.	Annual energy modeling results for an enterprise data center in San Francisco, CA	197
5-25.	Annual energy modeling results for an enterprise data center in Seattle, WA.....	198

5-26.	Annual energy modeling results for an enterprise data center in Chicago, IL	199
5-27.	Annual energy modeling results for an enterprise data center in Dallas, TX	200
5-28.	Annual energy modeling results for an enterprise data center in Richmond, VA	201
5-29.	Baseline PUE values for the standard operation scenario	202
5-30.	Current Practices PUE values for the standard operation scenario.....	203
5-31.	Economizer PUE values for the energy efficiency scenario.....	203
5-32.	Economizer Plus PUE values for the energy efficiency scenario.....	203
5-33.	Electricity resource mix for each of the modeled data center geographical regions	205
5-34.	Carbon intensity values associated with electricity generation in the United States compiled from national inventory data (EIA, 2009a; EIA, 2009b), Fthenakis and Kim (2007), and Pacca and Horvath (2002).....	206
5-35.	Greenhouse gas intensity (CO ² (e)/kWh) emissions associated with regionally specific electricity generation sources	206
5-36.	Hours of economizer activity for each climate region under the Economizer Plus scenario	209
5-37.	Hours of economizer activity for each month in San Francisco under the Economizer Plus scenario	210

List of Figures

1-1.	Breakdown of United States commercial and residential building energy use (EIA, 2008).....	4
1-2.	Breakdown of data center energy use from empirical data gathered at 22 data centers in California (Greenberg et al., 2006).....	6
1-3.	Documented (2000-2006) and projected (2006-2011) total United States data center energy use (Brown et al., 2007).....	8
2-1.	Schematics of airflow at the data centers.....	18
2-2.	Filter efficiency as a function of particle size.....	24
2-3.	Loss-rate coefficient for deposition to indoor surfaces as a function of particle size.....	24
2-4.	Cumulative probability distributions of 5-minute average measured mass particle concentration at three data-center sites.....	29
2-5.	Measured time-dependent particle mass concentrations at the Rocklin site during 15-20 September 2006.....	31
2-6.	Measured time-dependent particle mass concentrations at the Walnut Creek site during 14-20 October 2006.....	31
2-7a.	Time-dependent measured concentrations of particles 0.3-0.5 μm in diameter at the Rocklin data center during 15-20 September 2006.....	32
2-7b.	Time-dependent measured concentrations of particles 0.5-0.7 μm in diameter at the Rocklin data center during 15-20 September 2006.....	32
2-7c.	Time-dependent measured concentrations of particles 0.7-1.0 μm in diameter at the Rocklin data center during 15-20 September 2006.....	33
2-7d.	Time-dependent measured concentrations of particles 1.0-2.0 μm in diameter at the Rocklin data center during 15-20 September 2006.....	33
2-7e.	Time-dependent measured concentrations of particles 2.0-5.0 μm in diameter at the Rocklin data center during 15-20 September 2006.....	34
2-8a.	Time-dependent measured concentrations of particles 0.3-0.5 μm in diameter at the Walnut Creek data center during 14-20 October 2006.....	35

2-8b.	Time-dependent measured concentrations of particles 0.5-0.7 μm in diameter at the Walnut Creek data center during 14-20 October 2006	35
2-8c.	Time-dependent measured concentrations of particles 0.7-1.0 μm in diameter at the Walnut Creek data center during 14-20 October 2006	36
2-8d.	Time-dependent measured concentrations of particles 1.0-2.0 μm in diameter at the Walnut Creek data center during 14-20 October 2006	36
2-8e.	Time-dependent measured concentrations of particles 2.0-5.0 μm in diameter at the Walnut Creek data center during 14-20 October 2006	37
2-9.	Measured time-dependent particle mass concentrations at the Sunnyvale data center during 21-25 August 2006	40
2-10a.	Time-dependent measured concentrations of particles 0.5-0.7 μm in diameter at the Sunnyvale data center during 21-25 August 2006	41
2-10b.	Time-dependent measured concentrations of particles 0.7-1.0 μm in diameter at the Sunnyvale data center during 21-25 August 2006	41
2-10c.	Time-dependent measured concentrations of particles 1.0-2.0 μm in diameter at the Sunnyvale data center during 21-25 August 2006	42
2-10d.	Time-dependent measured concentrations of particles 2.0-5.0 μm in diameter at the Sunnyvale data center during 21-25 August 2006	42
2-11.	Modeled particle fates at each study site	45
2A-1.	Measured time-dependent particle mass concentrations at the Redwood City site during 29 September – 6 October 2006	51
2A-2a.	Time-dependent measured concentrations of particles 0.3-0.5 μm in diameter at the Redwood City site during 29 September – 6 October 2006.....	52
2A-2b.	Time-dependent measured concentrations of particles 0.5-0.7 μm in diameter at the Redwood City site during 29 September – 6 October 2006.....	52
2A-2c.	Time-dependent measured concentrations of particles 0.7-1.0 μm in diameter at the Redwood City site during 29 September – 6 October 2006.....	53
2A-2d.	Time-dependent measured concentrations of particles 1.0-2.0 μm in diameter at the Redwood City site during 29 September – 6 October 2006.....	53

2A-2e.	Time-dependent measured concentrations of particles 2.0-5.0 μm in diameter at the Redwood City site during 29 September – 6 October 2006.....	54
2A-3.	Measured time-dependent particle mass concentrations at the Berkeley site during 4-8 May 2006	56
2A-4a.	Time-dependent measured concentrations of particles 0.3-0.5 μm in diameter at the Berkeley site during 4-8 May 2006.....	56
2A-4b.	Time-dependent measured concentrations of particles 0.5-0.7 μm in diameter at the Berkeley site during 4-8 May 2006.....	57
2A-4c.	Time-dependent measured concentrations of particles 0.7-1.0 μm in diameter at the Berkeley site during 4-8 May 2006.....	57
2A-4d.	Time-dependent measured concentrations of particles 1.0-2.0 μm in diameter at the Berkeley site during 4-8 May 2006.....	58
2A-4e.	Time-dependent measured concentrations of particles 2.0-5.0 μm in diameter at the Berkeley site during 4-8 May 2006.....	58
2A-5.	Measured time-dependent particle mass concentrations at the Oakland site during 15-21 May 2006.....	60
2A-6a.	Time-dependent measured concentrations of particles 0.3-0.5 μm in diameter at the Oakland site during 15-21 May 2006.....	60
2A-6b.	Time-dependent measured concentrations of particles 0.5-0.7 μm in diameter at the Oakland site during 15-21 May 2006.....	61
2A-6c.	Time-dependent measured concentrations of particles 0.7-1.0 μm in diameter at the Oakland site during 15-21 May 2006.....	61
2A-6d.	Time-dependent measured concentrations of particles 1.0-2.0 μm in diameter at the Oakland site during 15-21 May 2006.....	62
2A-6e.	Time-dependent measured concentrations of particles 2.0-5.0 μm in diameter at the Oakland site during 15-21 May 2006.....	62
2A-7.	Measured time-dependent particle mass concentrations at the San Francisco site during 8-10 August 2006	64
2A-8a.	Time-dependent measured concentrations of particles 0.3-0.5 μm in diameter at the San Francisco site during 8-10 August 2006	64
2A-8b.	Time-dependent measured concentrations of particles 0.5-0.7 μm in diameter at the San Francisco site during 8-10 August 2006	65

2A-8c.	Time-dependent measured concentrations of particles 0.7-1.0 μm in diameter at the San Francisco site during 8-10 August 2006	65
2A-8d.	Time-dependent measured concentrations of particles 1.0-2.0 μm in diameter at the San Francisco site during 8-10 August 2006	66
2A-8e.	Time-dependent measured concentrations of particles 2.0-5.0 μm in diameter at the San Francisco site during 8-10 August 2006	66
2A-9.	Measured time-dependent particle mass concentrations at the Dublin site during 22-22 July 2006.....	69
2A-10a.	Time-dependent measured concentrations of particles 0.3-0.5 μm in diameter at the Dublin site during 22-22 July 2006.....	70
2A-10b.	Time-dependent measured concentrations of particles 0.5-0.7 μm in diameter at the Dublin site during 22-22 July 2006.....	70
2A-10c.	Time-dependent measured concentrations of particles 0.7-1.0 μm in diameter at the Dublin site during 22-22 July 2006.....	71
2A-10d.	Time-dependent measured concentrations of particles 1.0-2.0 μm in diameter at the Dublin site during 22-22 July 2006.....	71
2A-10e.	Time-dependent measured concentrations of particles 2.0-5.0 μm in diameter at the Dublin site during 22-22 July 2006.....	72
2A-11.	Measured time-dependent particle mass concentrations at the Dublin site during 1-8 December 2006.....	73
2A-12a.	Time-dependent measured concentrations of particles 0.3-0.5 μm in diameter at the Dublin site during 1-8 December 2006.....	73
2A-12b.	Time-dependent measured concentrations of particles 0.5-0.7 μm in diameter at the Dublin site during 1-8 December 2006.....	74
2A-12c.	Time-dependent measured concentrations of particles 0.7-1.0 μm in diameter at the Dublin site during 1-8 December 2006.....	74
2A-12d.	Time-dependent measured concentrations of particles 1.0-2.0 μm in diameter at the Dublin site during 1-8 December 2006.....	75
2A-12e.	Time-dependent measured concentrations of particles 2.0-5.0 μm in diameter at the Dublin site during 1-8 December 2006.....	75
3-1.	Indoor and outdoor sets of sample filters used for particle collection.....	83
3-2.	Denuder filter system for chemical speciation.	85

3-3.	Supply fan curve for one of 16 identical fans used to supply air at the Sunnyvale data center	91
3-4.	Indoor/outdoor concentration ratios for OPC-measured particles, sorted by HVAC filter type and economizer activity	98
3-5.	Indoor/outdoor concentration ratios for particulate sulfate, sorted by HVAC filter type and economizer activity	98
3-6.	Indoor/outdoor concentration ratios for particulate nitrate, sorted by HVAC filter type and economizer activity	99
3-7.	Indoor/outdoor concentration ratios for particulate chloride, sorted by HVAC filter type and economizer activity	99
3-8.	Indoor/outdoor concentration ratios for particulate ammonium, sorted by HVAC filter type and economizer activity	100
3-9.	The difference between measured ammonium concentrations (Table 3-5) and the difference between measured nitrate and chloride (Tables 3-3 and 3-4).....	104
3-10.	Indoor/outdoor concentration ratios for quartz filter collected black carbon particles, sorted by HVAC filter type and economizer activity...	105
3-11.	Indoor/outdoor concentration ratios for aetholometer-measured black carbon particles, sorted by HVAC filter type and economizer activity...	105
3-12.	Indoor/outdoor concentration ratios for quartz filter collected carbonaceous particles (organic and black carbon), sorted by HVAC filter type and economizer activity	106
3-13.	Indoor/outdoor concentration ratios for gravimetrically-measured particulate matter, sorted by HVAC filter type and economizer activity.....	106
3-14.	Measured time-dependent indoor and outdoor particle mass concentrations over a 24-hour period for each filter configuration	109
3-15a.	Outdoor particle concentrations measured by OPC during study period of 9-20 August 2008	111
3-15b.	Outdoor particle concentrations measured by OPC during study period of 21-28 August 2008	112
3-16.	Cumulative probability distributions of estimated non-IT power use for the economizer-on and economizer-off periods for each filter type ..	115

4-1.	Air and water flow schematic of data center cooling without any economizer use (base-case design)	121
4-2.	Air and water flow schematic of data center cooling with a water-side economizer	123
4-3.	Air and water flow schematic of data center cooling with an air-side economizer	125
4-4.	Data Center Chillers Part load efficiencies for a water-cooled centrifugal chiller with a capacity >1050 kW and an condenser water temperature of 26.7 °C (CEC, 2005)	130
4-5.	Evaluated Climate Zone Locations	130
4-6.	Typical electrical components in a data center that represent the total building load	132
4-7.	Disaggregated energy use for each design scenario (climate dependent values only)	136
4-8.	Hourly distribution of the outside air drybulb temperature for the 8760 hours throughout the year at the five California data center locations	137
4-9.	Hourly distribution of the outside air wetbulb temperature for the 8760 hours throughout the year at the five California data center locations	137
4-10.	Chiller and fan energy demand separated by design scenario and humidity restrictions	139
5-1.	Air Relationship between processor utilization and system power demand, both with and without power scaling (management) activated (Brown et al., 2007)	149
5-2.	Five climate regions used in PUE modeling	156
5-3.	Comparison of national data center energy use under the standard operation and energy efficient scenarios	170
5-4.	Total data center energy use, separated into IT and non-IT components, for each climate region	174
5-5.	Server room non-IT component energy use as a percentage of the IT energy demand	177

5-6.	Localized data center non-IT component energy use as a percentage of the IT energy demand	182
5-7.	Mid-tier data center non-IT component energy use as a percentage of the IT energy demand	189
5-8.	Enterprise data center non-IT component energy use as a percentage of the IT energy demand	196
5-9.	Total data center energy use, separated as IT and non-IT components, for each climate region.....	207
5-10.	Monthly energy use for enterprise data centers in San Francisco under the Economizer Plus scenario	212
5-11.	Daily variation in PUE value of an enterprise data center during a winter, spring, summer, and fall day in San Francisco under the Economizer Plus scenario	213
5-12.	Hourly data center energy use (MWh) using the hourly PUE values presented in Figure 5-11.	214

Chapter 1: Introduction

This chapter describes the motivation for and the objectives of this dissertation. Data centers are presented within the greater context of building energy, indoor air quality, and climate change. The role of energy efficient building design in reducing global greenhouse gas emissions is highlighted. The relationship between indoor air quality and energy efficiency is described. The need to balance these design goals in data centers is introduced as an important issue in addressing the rapid increase in energy demand from this building sector. The objectives of the dissertation are described and the research approach is outlined. The chapter concludes with a discussion of data center efficiency and the growth of information technology as an opportunity to advance sustainability.

1.1. Building energy use

Reducing the rate of anthropogenic carbon emissions and ultimately stabilizing atmospheric levels of CO₂ will require innovation and effort across a broad spectrum of disciplines. Successful strategies likely include a combination of developing low-carbon energy sources, sequestering carbon from fossil fuel combustion, using energy more efficiently to provide services, as well as creating transformative public policies and effective strategies for adaptation. Given the enormity of the challenge, reducing inefficiencies from all energy-intensive economic sectors is a necessary approach that can garner benefits on a relatively short timescale. The role for building designers is conspicuous, considering that buildings are the single biggest contributor to greenhouse gas emissions, accounting for 45% of worldwide energy consumption (Butler, 2008). Furthermore, technological leaps in building science are not required to conserve much of this energy, which is simply lost through inefficient design. Implementing established efficiency design strategies has been estimated to have the capacity to reduce carbon equivalent emission from the 2020 building stock by 29% relative to business-as-usual

practices (Ürge-Vorsatz et al., 2007). The cost of implementing strategies to stabilize atmospheric CO₂ levels is estimated to be about 1% of global gross domestic product (GDP) each year until 2050, while the business-as-usual trajectory will be equivalent to losing 5% of global GDP annually from the costs and risks of climate change (Stern et al., 2006). While the costs associated with climate change are stunningly large, many design measures to improve building efficiency represent carbon abatement strategies that have been identified to result in little or negative net cost (McKinsey and Company, 2007). Improving building energy efficiency therefore represents a relatively low cost opportunity to reduce global energy demand and significantly contribute climate change mitigation.

1.2. Building energy and indoor environmental quality

Using consumer interest to shift building practices through environmental rating systems is one strategy to advance building energy efficiency that has gained momentum. Green building rating systems attempt to better expose more traditionally opaque aspects of building design so that they can be incorporated into the decision processes that drive the building sector of the economy. Arguably the most prominent rating system is the Leadership in Energy and Engineering Design (LEED), which uses a primarily prescriptive approach to improve building design by focusing on five separate categories: site location, water use, material use, energy use, and indoor environmental quality (USGBC, 2009). The first four categories address how a building interacts with and affects different environmental resources. Indoor environmental quality (comprised of the thermal, air, and lighting quality) focuses on the health and comfort of building

occupants. These indoor environmental quality design goals can affect building energy use and the important role of indoor environmental quality in addressing climate change has been outlined before (Nazaroff, 2008). However, when energy and indoor environmental quality are presented as distinct categories in LEED, each promoting separate prescriptive design measures, these green building goals can seem unequal in priority. While energy efficiency measures appear to carry the weight of issues such as energy security and climate change, indoor environmental quality measures appear to be merely striving to increase the occupant experience and can be marginalized as an attempt to increase consumer appeal in the rating system. Accordingly, LEED has come under criticism for the prominence of indoor environmental quality in the rating system, with calls to give significantly more emphasis to energy performance (Nature, 2009). Discounting indoor environmental quality in favor of energy efficiency, however, overlooks the inherent interdependence between the two design goals.

As Figure 1-1 shows, heating, ventilation, and air conditioning (HVAC) energy accounts for a significant portion of the operational energy in both residential and commercial buildings (EIA, 2008). By regulating the temperature, humidity, and the amount of outside air entering the building, HVAC systems greatly influence overall indoor environmental quality; indeed, to do so is the primary purpose of this considerable energy use. HVAC design without concern for energy consumption can result in inefficiency, while HVAC design without concern for indoor environmental quality neglects a primary building function. Pursuing the benefits of either extreme is short sighted, with one limit missing the opportunity to reduce unnecessary energy consumption and the other ultimately hurting the progress and acceptance of energy-

efficient buildings. Simultaneous consideration of both design goals, however, would allow energy efficiency in buildings to continuously improve through innovation and new technologies, while insuring that the requirement for a healthy indoor environment is not eclipsed in the process.

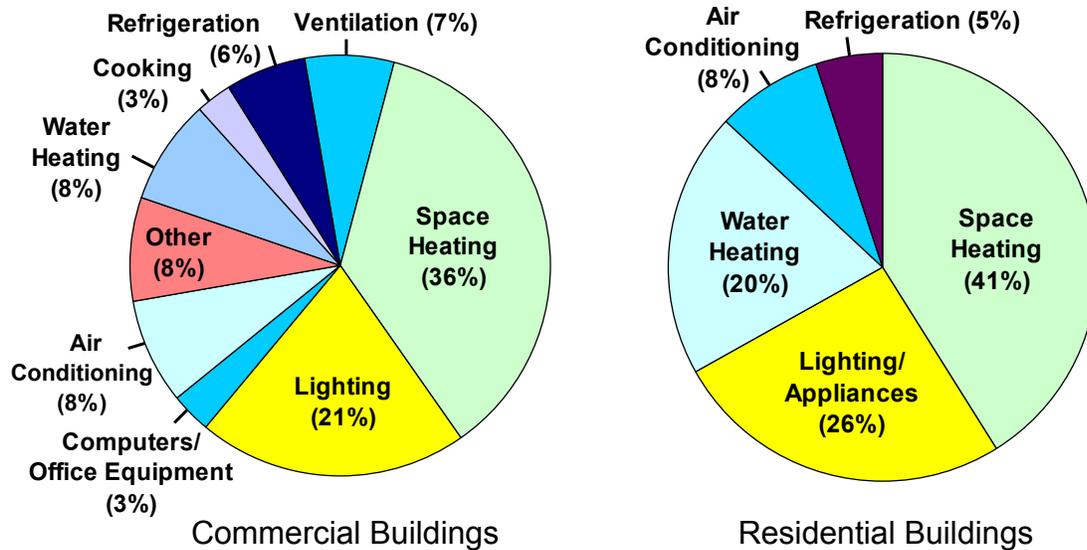


Figure 1-1. Breakdown of United States commercial and residential building energy use (EIA, 2008). Breakdown based on total 2003 energy use from commercial buildings (1,706 billion kWh) and total 2005 energy use from residential buildings (3,092 billion kWh).

1.3. Data center buildings

Data centers provide a clear illustration of the interdependence between energy and environmental quality, where concerns regarding building operational energy use and indoor air quality (IAQ) can have significant consequences. Data centers are buildings designed to contain information technology (IT) equipment used for data processing (servers), data storage (storage equipment), and communications (network equipment).

IT equipment is typically stacked approximately 2 meters high in multiple aisles of server racks. As Figure 1-2 shows, HVAC energy demand is comparable to the plug load generated by the IT equipment operating in data centers (Greenberg et al., 2006). Following the first law of thermodynamics that requires energy to be conserved, the electrical energy consumed by IT equipment is dissipated as heat, which must then be removed to prevent IT equipment from overheating. The high concentration of IT equipment and the complementary cooling systems in data centers combine to result in power density demands greater than 1 kW per m² of floor areas, orders of magnitude higher than conventional office buildings (Greenberg et al., 2006). Motivation for reducing this high energy demand reaches beyond environmental considerations. At large high density data centers, the annual cooling costs alone can be on the order of \$10 million (Patel and Shah, 2005) and these buildings have come under scrutiny because of the increasing amounts of energy they consume (Loper and Parr 2007). By the end of 2006, the U.S. Congress had requested an evaluation of data centers in part to address the economic damage that would result from an inability to meet future energy demand (U.S. Congress, 2006).

The IAQ concerns in data centers are also important. While IAQ in buildings such as offices and residences can be associated with a plethora of potential health impacts for occupants, ranging from subtle irritation to life threatening complications, the effects of IAQ on data centers are essentially binary. Data center operators are wary of any airborne pollutants entering the data center and damaging the IT equipment in a way that could cause operation failure within the 3-5 year equipment lifetime (Tschudi et al., 2004). Many data centers are considered “mission critical” and interruptions in service

can be financially catastrophic, with costs claimed to be as high as \$30 million per minute of operation downtime during peak periods (Brown et al., 2007). The types and concentrations of pollutants that could actually cause equipment failure are poorly understood. Because of the high cost of failure and the lack of detailed knowledge about risks, energy saving measures that involve using outside air and potentially altering IAQ conditions away from conventional practices are generally avoided.

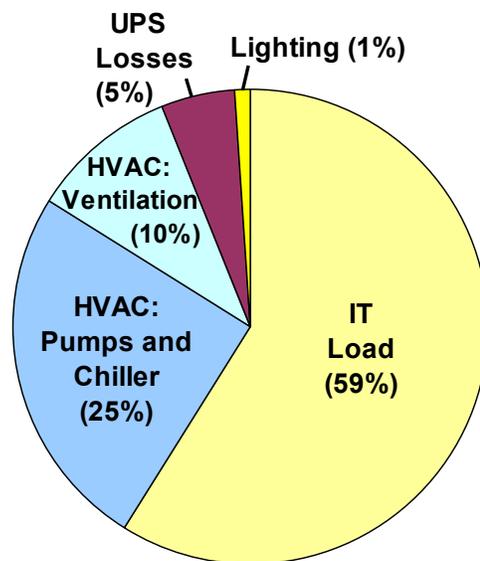


Figure 1-2. Breakdown of data center energy use from empirical data gathered at 22 data centers in California (Greenberg et al., 2006)

1.4. Growing data center energy demand

Properly navigating the precarious path between excessive energy use and degraded IAQ risks can yield energy efficiency benefits for many different building types. Data centers warrant specific evaluation as future growth in this sector may make imperative improved energy efficiency of these buildings. Data center buildings represent the backbone of the internet. Digital information processing, digital storage,

and digital communication are becoming increasingly integral to commerce and to the functioning of society as a whole. Investment in IT has grown from negligible in U.S. companies a generation ago to now representing as much as all other capital expenditures combined (Carr, 2003). The aggregate energy use for computer servers doubled between 2000 and 2005, both in the United States and worldwide (Koomey, 2007). Along with the overall increase in IT demand, increased centralization of IT services could place greater importance on the energy efficiency of large data centers. More businesses are outsourcing their IT activities, shifting servers, storage, and network equipment out of office closets into dedicated data center facilities (Carr, 2005). The advent of cloud computing (a neologism recently gaining traction) could further increase the IT burden of large data centers as the digital information processing and storage that currently occur in personal computers and office equipment become internet-based utilities (Fowler and Worthen, 2009). A global perspective indicates that data center growth may be further accelerated as IT is expanded in emerging markets. Many of the IT services associated with data centers are wireless, allowing the infrastructure needs to be geographically concentrated relative to other all-purpose technologies, such as electricity, telephones, and railroads. Accordingly, the growth rate of IT in developing countries has been rapid compared to earlier these all-purpose technologies (IMF, 2001). As shown in Figure 1-3, total U.S. data center energy use doubled to about 60 billion kWh annually between 2000 and 2006. This rapid increase in energy use is the result of the data center industry growing to meet the demand for more IT services, as businesses have automated more processes, data are stored with greater complexity, and rich media is being increasingly utilized (Brown et al., 2007). Brown et al. (2007) also estimated that growth in data

center energy demand would continue at a similar rate in the near future exceeding 100 billion kWh per year in the United States in 2011. It was also reported that current data center practices are layered with inefficiencies. Consequently, the projected increase in energy could be stemmed with the implementation of energy efficiency strategies. The different scenarios presented in Figure 1-3 represent varying levels of energy-efficient practices and technologies associated with both IT and non-IT equipment in data centers. The more efficient scenarios (“best practice” and “state of the art”) include design measures that change the ventilation rates and potentially the air quality in data centers. Given the emphasis on equipment reliability, realizing the benefits of the more energy efficient scenarios is contingent on understanding the associated IAQ repercussions in data centers.

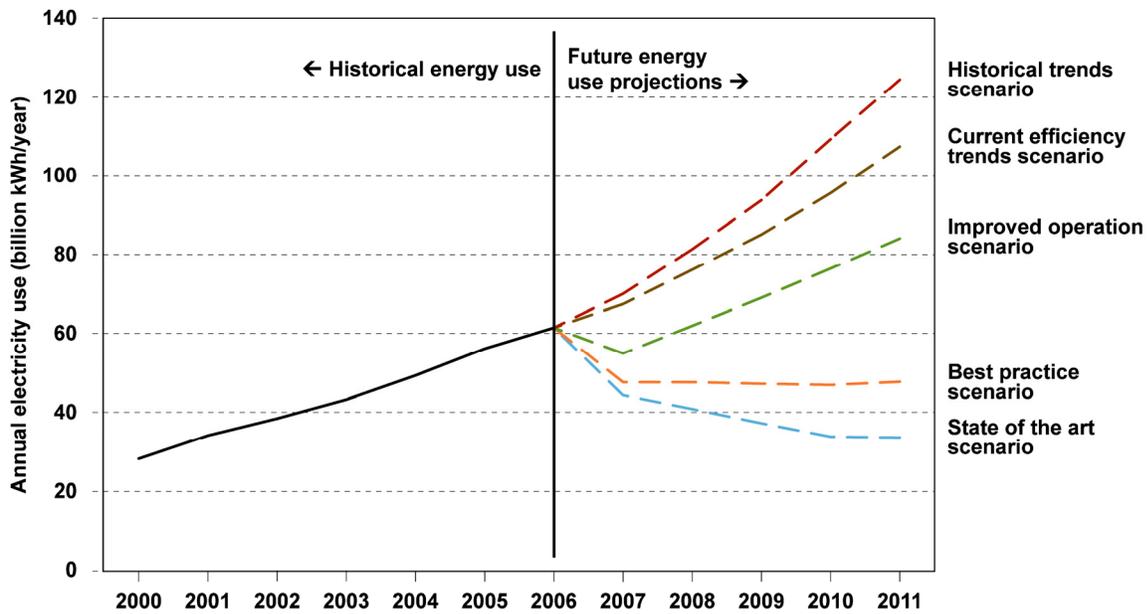


Figure 1-3. Documented annual (2000-2006) and projected (2006-2011) total United States data center energy use. Projected scenarios represent varying implementation of IT and non-IT energy efficient design measures (Brown et al., 2007)

1.5. Dissertation scope

This dissertation aims to better understand how building design and operation influence the relationship between energy use and IAQ in data centers. Indoor particle concentrations are addressed as a building performance issue, because of concerns that exposure to increased particle concentrations can compromise computer equipment reliability, which can hinder the implementation of more energy-efficient cooling practices. Specifically, this dissertation attempts to understand how the use of economizers can affect energy demand and indoor particle concentrations in data centers. Conventional data centers operate with essentially a closed-loop air cycle, where all hot air removed from the data center zone is directly returned to internal, compressor-based air-conditioning units to be cooled and again supplied to the zone. Economizers reduce the data center cooling energy by replacing the air being returned to the air-conditioning unit with outside air when the outside temperature is below that of the air being removed from the zone. Concerns that economizers would increase indoor levels of particles of outdoor origin that may damage electronic equipment have hindered widespread implementation of this technology (Tschudi et al., 2004). ASHRAE Technical Committee 9.9, the trade organization that traditionally addresses issues of data center design and operation, has avoided making any recommendations about the application of outside air economizers until more research can be provided to either support or reject their use for data centers (Syska Hennessy Group, 2007). Furthermore, incentives for the industry to shift towards more energy-efficient technologies are currently presented qualitatively, since the economic and energy benefits can depend on many parameters that are specific to site location and the HVAC system. Decision makers are left with a

vague description of economizer risks and benefits, which hampers motivation for technology shifts from within the industry and delays the demand for external regulation.

This dissertation first explores the IAQ impact of economizer use in data centers. A combination of monitoring and modeling efforts are used to establish particle concentrations in data centers under different mechanical cooling system design and operation. Concentrations of specific pollutants are then measured while applying filter-based mitigation during economizer use. The focus of the dissertation then shifts to quantifying the energy saving benefits of economizer use. Data center energy use is modeled under multiple conditions, including different mechanical designs, building types, and geographical climates. The modeled energy values are combined with results from Brown et al. (2007) to estimate current (2008) data center energy use. The potential energy saving available from implementing economizers is placed in the context of other prominent energy efficiency measures available to data centers. Overall, the results presented in this dissertation identify energy efficiency strategies that limit the risk of equipment damage from particulate matter and provide insight into the energy savings available from implementing those strategies.

1.6. Dissertation format

The British physicist and engineer Lord Kelvin is attributed with the quote, “if you cannot measure it, you cannot improve it.” Much of the research undertaken in this dissertation follows that philosophy by attempting to understand the costs and benefits of economizer implementation through the metrics of IAQ and operational energy use. These metrics are quantified under different operational modes, allowing an appropriate

balance to be maintained while working towards improved data center design.

Chapter 2 begins by establishing aspects of the IAQ landscape for current data center buildings. Eight different data centers are evaluated to determine both the differences and similarities in physical layout and mechanical design. The data centers include those conventionally operated as well as ones that employ economizers. Ionic particles, such as ammonium sulfate and nitrate, are identified as pollutants of special concern owing to their ability to deliquesce and bridge isolated conductors following deposition within the IT equipment. At each of the eight data centers, optical particle counters are used to simultaneously measure indoor and outdoor concentrations of particles 0.3-5.0 μm in diameter. The empirical results establish the magnitude of particle concentrations and indoor/outdoor particle ratios at data centers with and without economizer use. At three of the data centers, material balance modeling is applied to identify significant particle sources and sinks influencing data center indoor particle concentrations. This modeling procedure is also used to anticipate the indoor proportion of outdoor sulfate particles in a data center with and without economizer use.

With the effect of economizer use on indoor particle concentration measured in Chapter 2, methods to mitigate the associated particle increase in data centers while retaining the economizer energy benefits are the focus of Chapter 3. At a data center equipped with an economizer, particle concentrations are evaluated while the mechanical system operates with different air filters of increasing efficiency installed in the air handling units. In addition to again using optical particle counters to measure the concentration of particles 0.3-5.0 μm in diameter, an aethalometer is used to measure the concentration of black carbon inside and directly outside of the data center. Filter-based

particle measurements are also taken to measure concentrations of ammonium sulfate, nitrate, and chloride particles as well as concentrations of black and organic particulate matter. The sampling apparatus constructed for the filter-based measurements is designed to account for particle volatilization during collection, since ammonium nitrate and chloride particles reside in equilibrium with their gaseous constituents.

Indoor/outdoor ratios are calculated for each particle species during active and inactive economizer periods. The impact of enhanced filtration is evaluated against the particle concentrations measured in data centers without economizers. Fan energy attributable to the data center mechanical system is estimated and electricity use is monitored throughout the entire particle collection process to better understand the energy impact of different air filtration efficiencies. The data gathered in this chapter are used to propose utilizing enhanced air filtration as a strategy to gain the energy benefits of economizers while maintaining indoor particle concentrations comparable to those expected for a data center that is operated without economizers.

In Chapter 4, the emphasis shifts from measuring the IAQ impact of data center economizers to better understanding the energy benefits of economizers under different design strategies. Total data center energy use is disaggregated and an established metric, the Power Use Efficiency (PUE), is presented as a way to compare the efficiency of the non-IT portion of a data center. A building energy model is used to compare data center energy demand in different California climates. The model is developed to estimate energy use for a data center design with and without an economizer. A third design is also modeled that uses water-side economizers, which can be proposed as an alternative to traditional (air-side) economizers. Water-side economizers utilize cool

outside weather conditions to reduce chiller demand without requiring the ventilation increase necessary for traditional economizer design. Humidity restrictions are adjusted to determine the effect of these operational parameters on energy demand. The modeling results are evaluated to identify strategies to maximize energy efficiency.

Chapter 5 extends the modeling methods established in Chapter 4 to compare the efficiency of non-IT portions of a data center on a national level. Different size data centers are modeled to account for variation in equipment layout and efficiency. Temperature and humidity settings are adjusted in the model to determine their impact on mechanical energy demand. A bottom-up model developed in Brown et al. (2007) is used to estimate the total national energy use attributable to IT equipment in data centers. Prominent energy efficient IT measures identified in Brown et al. (2007) are incorporated into the IT energy estimate to compare the national energy use under current practice with a potential energy-efficient scenario. The modeled non-IT equipment efficiencies are then applied to the IT energy estimates to establish the current (2008) total United States data center energy use and to determine the energy savings available from this sector of the economy. The energy use and savings potential are presented in the context of greenhouse-gas emissions, based on regional variations in electricity generation sources.

Chapter 6 summarizes the findings from each chapter and places the results in the broader context of building energy efficiency and IAQ. Future areas of research are also proposed to address highlighted deficiencies in the research methods and to further the goals of this dissertation.

Chapter 2: Particle Concentrations in Data Centers

This chapter presents particle concentration measurements from multiple data centers. Concerns about particles in data centers are discussed and avoiding economizer use due to this concern is addressed as a loss in an available energy efficiency resource, the scale of which will increase as data center operations grow. The lack of published measurements and the variability of current data center IAQ standards are discussed, highlighting the need for measurements. The objective of this data center air quality study is explained. The measurement methods are described, along with data center layouts and cooling system approaches. Measurement results are presented discussing differences between particle concentrations measured in data centers with and without economizer use. Reproduced in part with permission from Atmospheric Environment 42, 5978-5990, 2008. Copyright 2008, Elsevier Inc. License Number: 2314471379968.

2.1. Introduction

Data centers house the vast amounts of equipment that provide the computational power, data storage, and global networking integral to modern information-technology systems. The high concentration of densely packed computers in data centers leads to floor-area-weighted power densities 15-100 times higher than those of typical commercial buildings (Greenberg et al., 2006). The operation of data center buildings in the United States consumes a substantial and rapidly increasing proportion of total national electricity demand. Data center energy use doubled during the first half of this decade and, in the US alone, accounted for about 45 TWh/y of electricity consumption, approximately 1.2% of total demand, in 2005 (Koomey, 2007). Under a business-as-usual trajectory, data center electricity use in the US has been projected to double again by 2010, although energy efficiency practices have been identified that could begin to significantly reduce this continued growth rate (Brown et al., 2007). A substantial

portion of the energy use in data centers is dedicated to cooling the computer equipment (Tschudi et al., 2004). The data-center cooling load can be reduced by a substantial fraction when large amounts of outside air are used to cool internal loads during favorable weather conditions (Sloan, 2008). However, many owners and operators are reluctant to use this cooling technique owing to concerns about the risk of equipment failure posed by introducing outdoor particulate matter into data center buildings.

Fine particulate matter can deposit on electronic circuit boards in the space between isolated conductors. When the humidity of the surrounding air rises above the deliquescence point, particles composed of water-soluble ionic salts can absorb moisture and dissociate to become electrically conductive (Weschler, 1991). Empirical results show that exposure to high sulfate concentrations at high humidity can cause electronic equipment failure (Litvak et al., 2000). However, the risk of failure under the environmental conditions typical of data centers is not well understood. Owing to the competitive nature and high economic value of businesses in this sector, failure data are not publicly shared. Furthermore, the effect of introducing greater flow rates of outside air (or any other design change) on equipment failure cannot be predicted with confidence, because little is known about the concentrations of particles in data centers, the sources of those particles, or their fate once introduced into the data center environment. This paper addresses these unknowns by measuring and modeling particle concentrations at operating data centers. The results provide a partial basis for assessing the equipment failure risk posed by particles for current data-center designs.

In the present study, time- and size-resolved particle concentration data were gathered over weeklong periods at eight different northern California data centers.

Building parameters for three of these data centers were documented and a material-balance model was employed to predict concentrations under various conditions so as to better understand the relative influence of potential sources and fates of airborne particles. Predicted indoor concentrations were compared against the measured results. The loss mechanisms of filtration, deposition, and ventilation were compared to assess particle fate. The model was also applied to estimate indoor concentrations of sulfate particles, which are of particular concern because of their ambient abundance and hygroscopicity.

2.2. Methods

2.2.1. Study Sites

Size-resolved particle concentrations were measured as a function of time at data centers in eight different northern California cities. With respect to ventilation and cooling, all of the data centers are conventional except for one in Sunnyvale, which was specifically designed to be energy-efficient and therefore has distinctive characteristics. This chapter presents detailed results from three of the monitored data centers — at Rocklin, Walnut Creek, and Sunnyvale. Appendix 2.A presents a summary of results from the remaining data centers. The Rocklin and Walnut Creek sites are both large buildings with multiple rooms designated for computer servers. Each of these rooms has characteristics common in data centers: rows of server racks, a raised-floor plenum, and computer-room air-conditioning (CRAC) units. The CRAC units are data-center-specific air-handling units (AHUs) that are situated on the data-center floor. By contrast, the data center in Sunnyvale is located in a single room within an office building, and is a

showcase for energy-efficient data center design and equipment. The room utilizes overhead air distribution and therefore contains no underfloor air plenum. The AHUs at this site are located separately in an adjacent room. In Sunnyvale, the amount of outside air entering the data center is controlled by an energy management and control system (EMCS). The EMCS implements the energy-efficient measure of cooling the computer equipment with large flow rates of outside air whenever climate conditions are favorable.

Figure 2-1 schematically displays the airflow configuration at each site. At Rocklin, outside air enters a rooftop AHU, passes through a 40% filter¹, and then enters the data center through a ceiling duct before mixing with the surrounding indoor air. Room air in the Rocklin data center enters the top of a CRAC unit, passes through another 40% filter, and is then cooled and discharged to the underfloor plenum. Perforations in the floor tiles in front of the server racks allow the cooled air to exit from the plenum into the data-center room. Fans within the computer servers draw the conditioned air upward and through the servers to remove heat generated by the equipment. After exiting the back side of the server housing, the warmed air then rises and is transported to the intake of a CRAC unit. The majority of air circulation at the Rocklin site is internal to the data-center zone. The Rocklin site has a single rooftop AHU to supply outside air to the room. This AHU supplies some outside air to positively pressurize the room and thereby limit infiltration. No air is mechanically removed from the room; rather, the mechanical supply air is balanced by air exfiltration across leaks in the data-center envelope.

¹ All filter efficiency specifications reported in this paper are based on the ASHRAE dust-spot test method (ASHRAE, 1992).

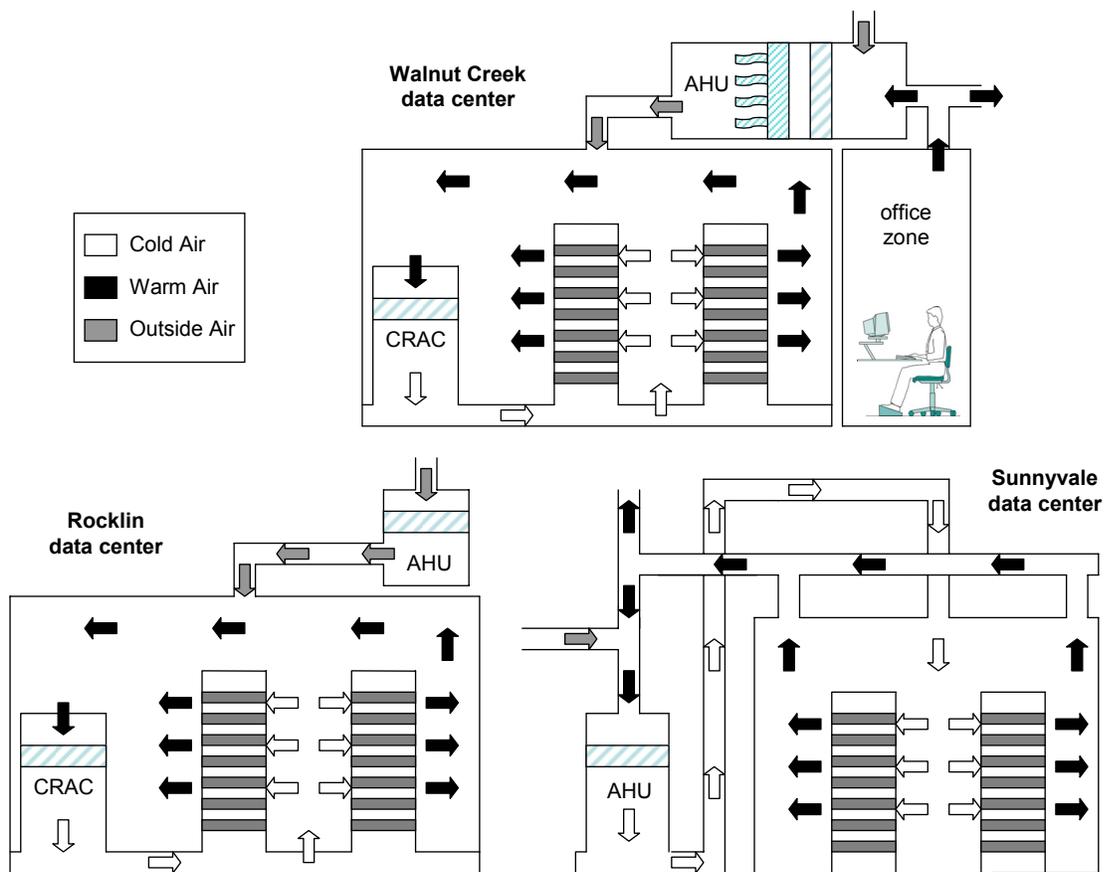


Figure 2-1. Schematics of airflow at the data centers. The Rocklin and Walnut Creek sites use an underfloor air distribution system. Air handling units (AHUs) are placed on the data center floor and air is thermally conditioned within the room. To maintain positive pressurization, a small amount of outside air is supplied from a separate rooftop AHU. The Sunnyvale site uses an airflow design common in office buildings. Air is supplied and removed through ceiling ducts and the AHUs are located outside of the data center zone.

Similar to Rocklin, the Walnut Creek site supplies only a small flow of outside air, as compared to the flow rate of air passing through the CRAC units. Supply air at Walnut Creek, however, is a combination of outside air and makeup air from other building zones. The rooftop AHU that supplies outside air to the data center also supplies air to office zones within the building. After entering the rooftop AHU, the outside air mixes with return air from the office zones. The ratio of outside air and office return air is automatically adjusted within the AHU, depending on the outside air temperature. This blend of outside and return air first passes through a 40% filter and then through an 85% filter before entering the data center and mixing with the surrounding indoor air.

Traditional CRAC units and the underfloor plenum are absent from the Sunnyvale site. Rather, air moves into and out of the room through ceiling-mounted air supply registers. These registers are connected via ducts to AHUs, which are located in a separate utility room adjacent to the data center. Ducts also connect the AHUs to the outside environment. Air from outside passes across adjustable dampers before being blended with return air from the data center. Once mixed, the air passes through a bank of 40% filters and is then thermally conditioned. The conditioned air is ducted into the data center and supplied through ceiling registers located between the server racks. As the cold supply air migrates toward the floor, fans draw air through the servers. After exiting the server rack, the warmed air is removed via ceiling return registers and ducted back to the AHUs. Before reaching the AHU, the air passes through another set of dampers. Some of the return air is exhausted while the rest is returned to the AHUs to be mixed with outside air before being conditioned and then returned to the data center. During the monitoring period, the EMCS at the Sunnyvale site was set to provide 85%

outside air whenever the outdoor temperature was below 15 °C. When the outdoor temperature increased above this set point, the amount of outside air was minimized to about 1% of the total flow. In each case, recirculation provided the remaining flow.

2.2.2. Experimental Protocol

Particle concentrations were measured both inside and outside of each site over a period of approximately one week. Size-resolved data were gathered using Met-One 237B optical particle counters (OPC), capable of detecting and sizing particles within the range 0.3-5.0 µm optical diameter with a maximum uncertainty of ±20% in particle counts for each size bin. Particle counts are separated into different size bins based on light scattering: 0.3-0.5 µm, 0.5-0.7 µm, 0.7-1.0 µm, 1.0-2.0 µm, and 2.0-5.0 µm. Mass concentrations were calculated from particle number counts by assuming a particle density of 1.5 g/cm³ (Pitz et al., 2003). A lognormal mass distribution of the particles within each size bin is assumed so that the geometric mean of that bin can represent the mass median diameter, which allows the particle mass concentration can be calculated as:

$$MassConc = \sum PC_i \left(\rho \frac{\Pi}{6} D_i^3 \right) \quad (2-1)$$

where the particle concentration is calculated as the sum of particle mass in each size bin, *i*. The mass for each size bin is calculated as the product of the particle count (PC), particle density (ρ), and mass median volume, where the particle diameter, *D_i*, represents the geometric mean for the size bin.

Table 2-1. Characteristics of three data center sites

	Sunnyvale	Walnut Creek	Rocklin
Parameters			
Floor Area (m ²)	616	360	1,208
Ceiling Height (m)	2.7	2.7	3.0
Volume (m ³)	1,690	931	3,681
Ventilation Flows (m ³ /min)			
outdoor supply	23*	10	9
recirculation	1,332	2,107	5,607
Monitoring Period			
start date	18 Aug. 2006	14 Oct. 2006	15 Sept. 2006
end time	25 Aug. 2006	20 Oct. 2006	20 Sept. 2006

*When in low outside air mode

Outdoor concentrations were measured by placing an OPC within the outside air intake that services the data center. Indoor concentrations were measured using a second OPC that was placed in front of a server aisle to measure the particle concentration in the air as it was about to pass through the server rack.

Measurements were taken for 5-minute intervals once every 25 minutes. Each OPC would draw air at a rate of 2.8 L/min for five minutes and then pause for 20 minutes before beginning the next particle-counting cycle. At the Sunnyvale site, the count for the 0.3-0.5 μm size range in the outdoor OPC reached the instrument limit for some sampling cycles, indicating that the true outdoor concentration was greater than the reported value. Consequently, data from this size range at the Sunnyvale site were not used in the analysis reported here.

The OPCs were factory calibrated prior to monitoring. Calibration was checked after monitoring by exposing both OPCs to the same conditions to ensure that each instrument produced consistent particle counts. During this calibration check, particle counts within each size category varied by less than 10% between the OPCs, and hence no corrections were applied to the analysis of data from the site measurements. Given the low concentrations measured at some of the data centers, the OPCs were also exposed to particle-free air, confirming that the monitors exhibited no lower-limit threshold.

2.2.3. Modeling Indoor Particle Concentrations

Indoor particle concentrations were predicted from time-dependent outdoor concentrations measured at each site. In the model, each data center zone was represented as a single, well-mixed chamber, using the parameters reported in Table 2-1. Assuming that the variation in particle concentration during each five-minute monitoring period is relatively small, the time-averaged, size-specific, mass-balance model is well represented by this equation:

$$\frac{C_{i,in}}{C_{i,out}} = \frac{\lambda_{out} (1 - \eta_{i,out})}{\lambda_{out} + \beta_i + \lambda_{rec} \eta_{i,rec}} \quad (2-2)$$

Equation (1) estimates indoor particle concentration as a size-specific proportion of the outdoor particle concentration. In the model, $C_{i,in}$ and $C_{i,out}$ are the indoor and outdoor concentration, respectively, for particles within size bin i . The parameter λ_{out} represents the outdoor air-exchange rate and λ_{rec} represents the recycled air-exchange rate, each defined as the respective airflow rate divided by the interior volume of the data center. The parameters $\eta_{i,out}$ and $\eta_{i,rec}$ are the respective size-dependent filter efficiencies for outside and recycled airflows. The coefficient, β_i , is the size-dependent deposition

loss rate for particle size section i . The terms in equation (1) represent time averages and assume uncorrelated ventilation rates and particle concentrations. These assumptions allow the dynamic time-averaged material balance to be represented by equation (1) without the need to assume steady-state conditions (Nazaroff and Klepeis, 2004). The model neglects resuspension, particle coagulation, or phase-change processes, based on the assumption that they have a relatively small influence as compared to the processes modeled. The data centers are positively pressurized and particle infiltration is designed to be negligible. The model assumes no unintended infiltration into these zones. Filter bypass, which reduces overall filter efficiency (Waring and Siegel, 2008) and merits investigation in data centers, is not addressed in this model.

Recycled airflow rates at the Walnut Creek and Rocklin sites are obtained from CRAC unit design specifications and are assumed to be constant throughout the monitoring period. An AccuBalance balometer was used to determine the supply airflow entering the data-center zone at the Rocklin and Walnut Creek sites, since design specifications for the outdoor air supply were not available. Balometers, commonly used within the HVAC industry for measuring air flows at registers, have been shown commonly to have errors of approximately 20% (Walker et al., 2001). This level of accuracy is adequate for the modeling analysis performed in this study.

Ventilation airflow at the Sunnyvale site depends on whether the HVAC system is in “low” (1% outside air) or “high” (85%) outdoor-air mode. Hourly data on the percentage of outside and recycled air entering the data center were gathered from the EMCS and then applied to the model calculations. As illustrated in Figure 2-2, particle removal efficiencies for the 40% and 85% filters used in the model are based on previous

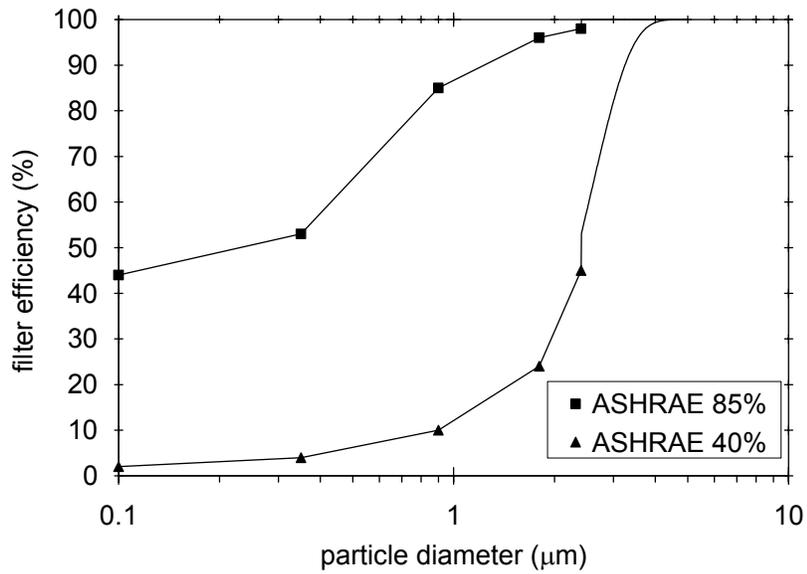


Figure 2-2. Filter efficiency as a function of particle size, from measured data (represented by squares and triangles) (Hanley et al., 1994). Linear interpolation provides estimates between measured data points. Fibrous-bed filter theory was used to extrapolate efficiency for particles larger than the measured particle sizes (Riley et al., 2002).

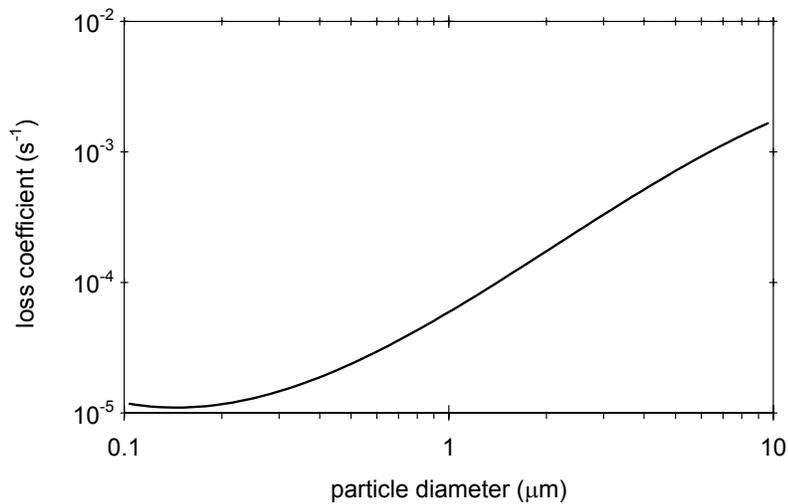


Figure 2-3. Loss-rate coefficient for deposition to indoor surfaces as a function of particle size. Line represents a least-squares cubic polynomial fit to logarithmically transformed data based on results compiled from six separate experimental studies (Riley et al., 2002).

empirical measurements of new filters for particle diameters of 0.35 μm , 0.9 μm , 1.8 μm , and 2.4 μm (Hanley et al., 1994). Each of the particle-size bins monitored by the optical particle counter was represented by its geometric median particle diameter for model calculations. Linear interpolation provided filter efficiency estimates for particle sizes between the measured data points. For particles larger than 2.4 μm , the filter efficiency was estimated from a best fit of the data of Hanley et al. to theoretical predictions of fibrous-bed filter efficiency (Riley et al., 2002). Since data are unavailable on the ratio of outside air and makeup air from other building zones that together comprise the supply air at the Walnut Creek site, additional particle measurements were taken at this site after the supply air had passed through the 40% and 85% filters. These post-filter particle measurements were used to represent the supply air entering the Walnut Creek data center. Size-dependent values for the indoor loss-rate coefficient (β_i) are based on six separate experimental studies that measured particle deposition rates across a range of particle sizes, ventilation conditions, and interior surface-to-volume ratios. The deposition loss coefficient, β_i , is equivalent to $\Sigma(v_{d,i,j}S_j/V)$, where $v_{d,i,j}$ is the size-dependent deposition rate for size section i onto surface j , S_j is the area of surface j , and V is the interior volume of the data-center zone. Figure 2-3 presents a least-squares cubic polynomial fit to the logarithmically transformed results from these six studies, as developed by Riley et al. (2002).

Particulate matter composed of water-soluble ionic salts present a special concern for data centers, owing to the ability of some of these salts to deliquesce and thereby conductively bridge isolated elements on circuit boards (Shields and Weschler, 1998). To investigate this concern, indoor sulfate concentrations were also specifically modeled.

Sulfate has been previously used to demonstrate current leakage attributable to particle deposition under conditions of high particle concentration and high humidity (Litvak et al., 2000). Sulfate, nitrate, and sea salt particles are the most common water-soluble ionic salts in ambient air and together represent a significant portion of urban particulate matter (McMurry et al., 2004). While each of the three particle types has the potential to cause equipment damage, sulfate was chosen for this study because its atmospheric abundance, size and thermal stability suggest that these particles may be of relatively greater concern than the other salts. Atmospheric sulfate is commonly found in the accumulation-mode size range (Milford and Davidson, 1987), which is expected to exhibit a relatively high indoor proportion of outdoor particles (IPOP) (Riley et al., 2002). By comparison, the IPOP of sea salt can be expected to be much lower, as sea salt particles are primarily found in the coarse mode (Seinfeld and Pandis, 2006), and so are efficiently removed by typical building filters and by settling onto room surfaces (Weschler, 1991). Sulfate is also likely to have a greater IPOP than nitrate (Sarnat et al., 2002). Nitrate particles, being volatile, can evaporate to their gaseous constituents when exposed to a warmer indoor environment (Lunden et al., 2003). The effects of nitrate particles on equipment risk in data centers appear to be worth investigating; however, to do so is beyond the scope of the present study.

Outdoor sulfate particle concentrations were estimated using data collected by the South Coast Air Quality Management District (SCAQMD) in central Los Angeles during a study conducted from January 1995 to February 1996 as part of the PM₁₀ Technical Enhancement Program (PTEP) (SCAQMD, 1996). SCAQMD used chemical mass-balance modeling to estimate that ammonium sulfate represented approximately 11% of

the average ambient PM₁₀ concentration of 48 µg/m³. For the present paper, this mass concentration, 5.3 µg m⁻³, was apportioned to a sulfate particle-size distribution using data compiled by Whitby (1978) from five studies of 15 urban sites. The mass-weighted sulfate particle size distribution is summarized as a single lognormal distribution with a geometric mean (GM) of 0.48 µm and a geometric standard deviation (GSD) of 2.0. The size distribution allowed the representative outdoor sulfate mass concentration to be segregated by particle diameter and applied to estimate indoor sulfate concentrations using equation 2-2.

2.3. Results and Discussion

2.3.1. Measured Particle Concentrations

Table 2-2 presents time-averaged, size-resolved, measured indoor particle concentrations for all eight data centers monitored. Average indoor concentrations for particles of diameter 0.3-5 µm are less than 1 µg/m³ in all conventional data centers and are substantially higher at the Sunnyvale data center with an energy-efficient design. A closer evaluation of the results from Rocklin, Walnut Creek, and Sunnyvale follows. Figure 2-4 presents the cumulative distributions of outdoor measured, indoor measured, and indoor modeled particle concentrations for these three sites during their respective monitoring periods. The average measured indoor concentrations at the Rocklin and Walnut Creek sites were 0.3 µg/m³ and 0.2 µg/m³, respectively, with indoor concentrations being approximately 1% of the corresponding outdoor values.

Table 2-2. Average measured indoor/outdoor particle concentrations at eight northern California data centers ($\mu\text{g}/\text{m}^3$)

	Data Center Location (CA)	Monitoring Period (2006)	Particle Size Range (μm)					Total
			0.3-0.5	0.5-0.7	0.7-1.0	1.0-2.0	2.0-5.0	
NetAps	Sunnyvale	18-25 Aug.	n/a	1.07	0.84	1.44	1.28	4.64
KP	Walnut	14-20 Oct.	0.06	0.02	0.03	0.07	0.05	0.22
	Creek	15-20 Sept.	1.20	0.45	0.45	1.27	6.63	10.0
GAP	Rocklin	15-20 Sept.	0.13	0.02	0.03	0.07	0.08	0.33
		29 Sept.-6 Oct.	1.74	0.84	1.10	2.87	7.72	14.3
Oracle	Redwood City	29 Sept.-6 Oct.	0.20	0.07	0.05	0.12	0.40	0.84
		22-24 July	2.04	1.00	0.69	1.33	5.19	10.3
Sybase	Dublin	22-24 July	0.14	0.03	0.03	0.07	0.03	0.30
		22 Nov.-5 Dec.	0.65	0.11	0.07	0.14	0.54	1.51
Sybase	Dublin*	22 Nov.-5 Dec.	1.34	0.78	0.60	0.74	0.20	3.67
		15-22 May	2.14	2.13	2.35	4.94	8.45	20.0
NERSC	Oakland	15-22 May	0.08	0.02	0.01	0.02	0.03	0.15
		8-10 Aug.	0.79	0.87	1.38	3.66	6.91	13.6
BofA	San Francisco	8-10 Aug.	0.33	0.12	0.07	0.13	0.30	0.95
		4-8 May	0.78	0.43	0.34	0.75	2.49	4.79
LBNL	Berkeley	4-8 May	0.08	0.04	0.03	0.05	0.11	0.31
		May	0.81	1.30	2.53	7.28	15.0	26.9

*The Dublin data center was monitored twice

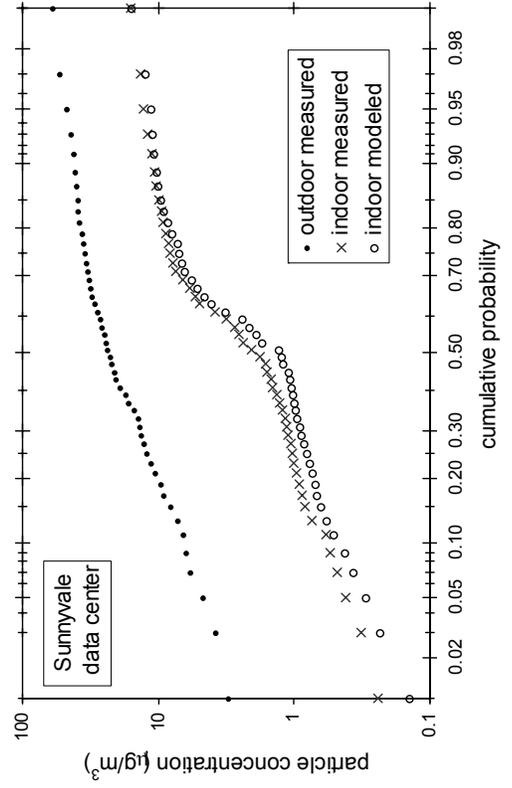
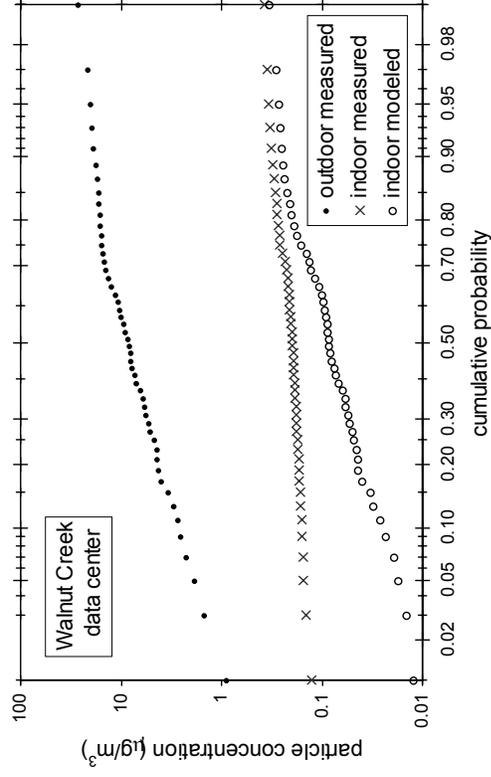
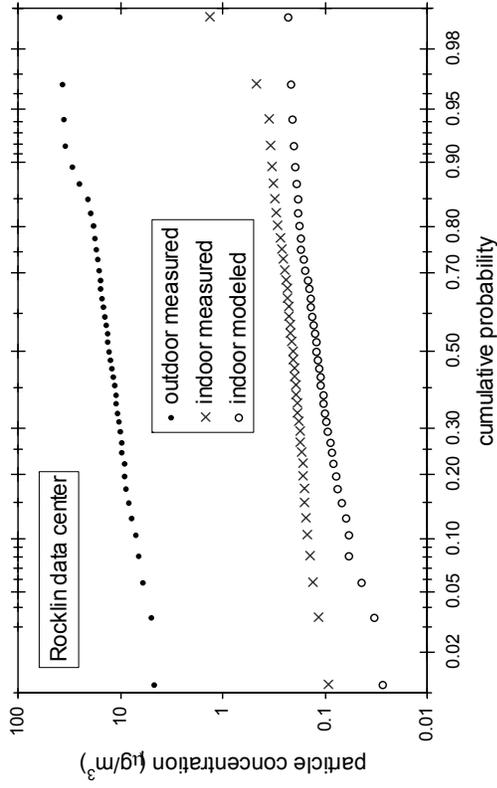


Figure 2-4. Cumulative probability distributions of 5-minute average measured mass particle concentration at three data-center sites. Sunnyvale concentrations only include particles 0.5-5.0 μm in diameter, while Rocklin and Walnut Creek include particles 0.3-5.0 μm in diameter. For clarity, only representative data points are displayed.

The median concentrations and IPOP from both of these sites are considerably lower than PM_{10} and $PM_{2.5}$ measurements previously reported for residential buildings (Ott et al., 2000; Long et al., 2001). The indoor concentration was significantly higher at the Sunnyvale site where the average measured indoor concentration was $4.6 \mu\text{g}/\text{m}^3$ and the IPOP was about 20%. This concentration remains lower than the indoor concentration limit for data centers suggested by ASHRAE for fine PM ($15 \mu\text{g}/\text{m}^3$). Particle guidelines for data centers vary widely among industry documents and some server manufacturers specify concentration limits that are orders of magnitude higher (ASHRAE, 2005). The average measured particle concentration at Sunnyvale is similar to previous measurements made in an office building across the same particle size range (Fisk et al., 2000). However, outdoor concentrations around the office building in the Fisk et al. study were much lower than the levels measured in Sunnyvale. High variability in indoor concentration is observed at the Sunnyvale site and is clearly associated with the proportion of outside air being toggled between 1% and 85% of the supply airflow. The indoor concentration between these two HVAC modes differs by an order of magnitude.

Low and steady indoor particle concentrations were measured at the Walnut Creek and Rocklin sites (Figure 2-5 and 2-6). The indoor concentration was less than $1 \mu\text{g}/\text{m}^3$ at almost all times, seemingly independent of fluctuations in the outdoor concentration. A few minor increases of short duration in indoor concentration are observed that do not correspond to any changes in outdoor concentration; these might be caused by occupants working or walking in the vicinity of the OPC.

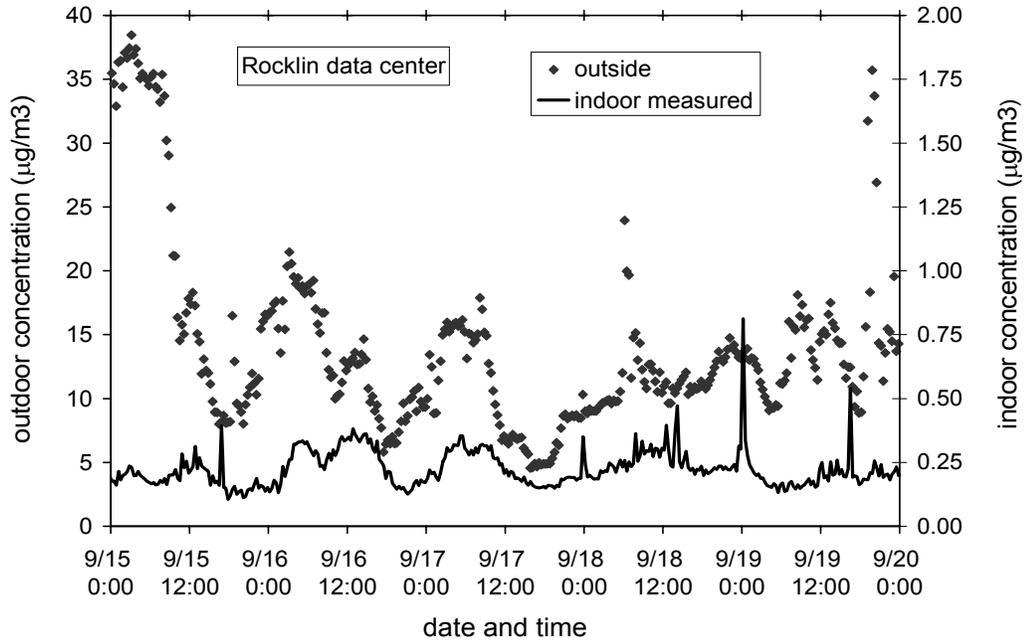


Figure 2-5. Measured time-dependent particle mass concentrations at the Rocklin site during 15-20 September 2006. Concentrations represent particles 0.3-5.0 µm in diameter.

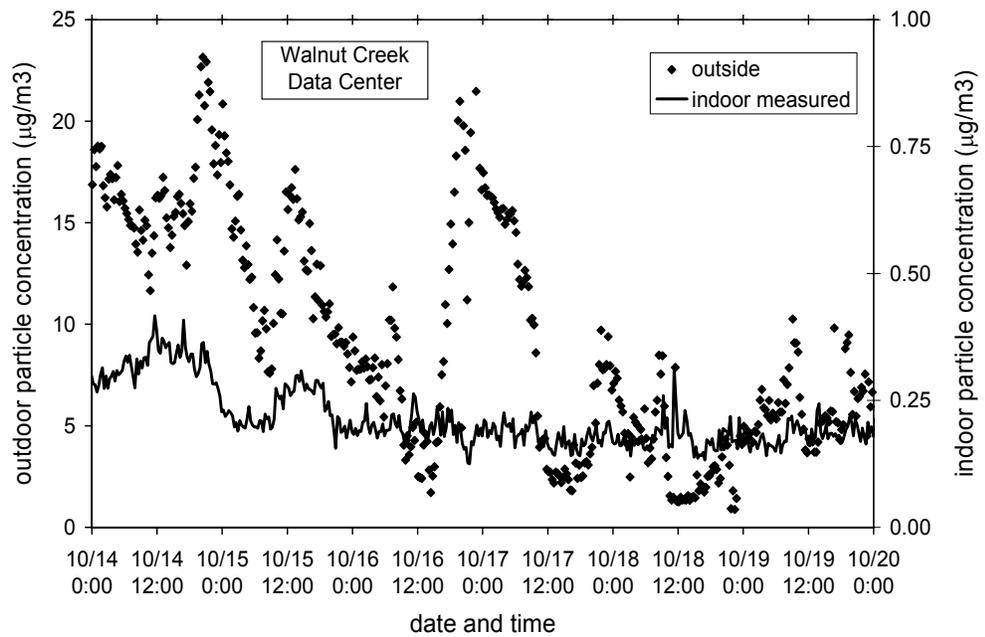


Figure 2-6. Measured time-dependent particle mass concentrations at the Walnut Creek site during 14-20 October 2006. Concentrations represent particles 0.3-5.0 µm in diameter.

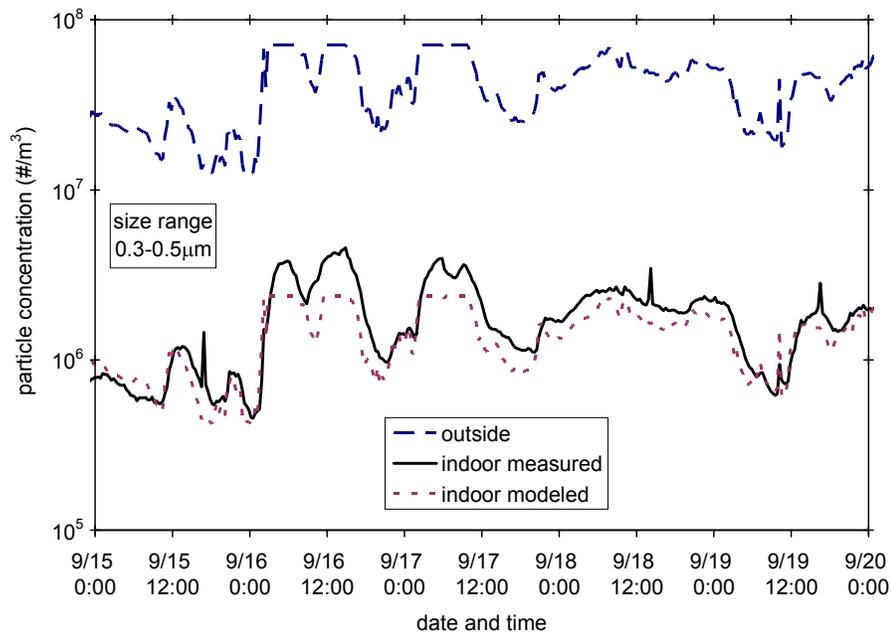


Figure 2-7a. Time-dependent measured concentrations of particles 0.3-0.5 μm in diameter at the Rocklin data center during 15-20 September 2006.

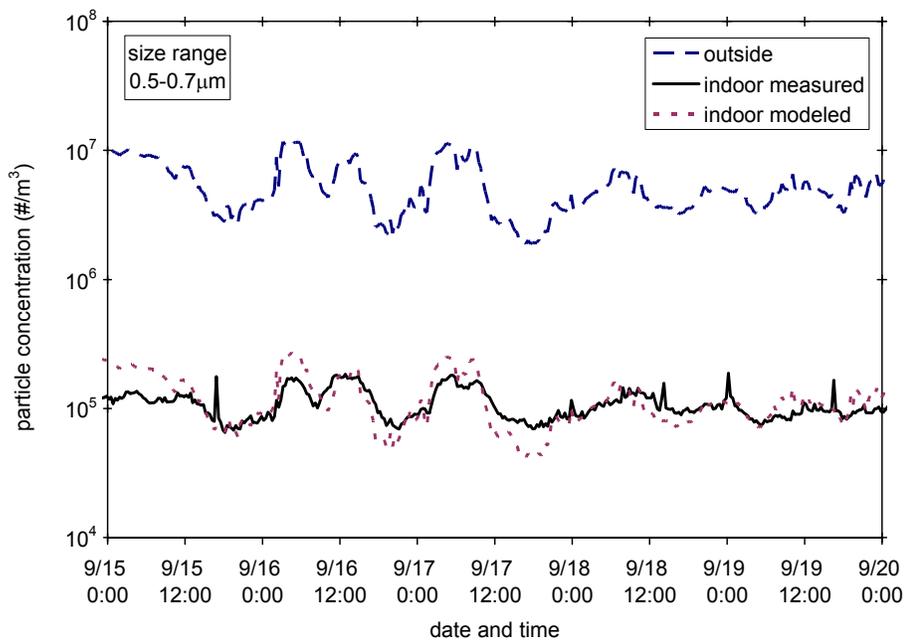


Figure 2-7b. Time-dependent measured and modeled concentrations of particles 0.5-0.7 μm in diameter at the Rocklin data center during 15-20 September 2006..

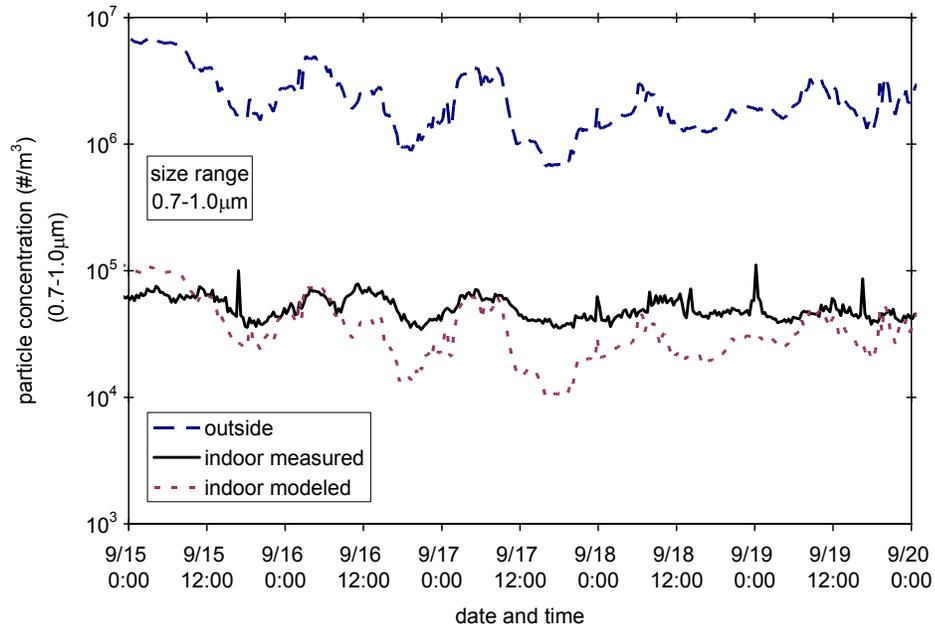


Figure 2-7c. Time-dependent measured and modeled concentrations of particles 0.7-1.0 μm in diameter at the Rocklin data center during 15-20 September 2006.

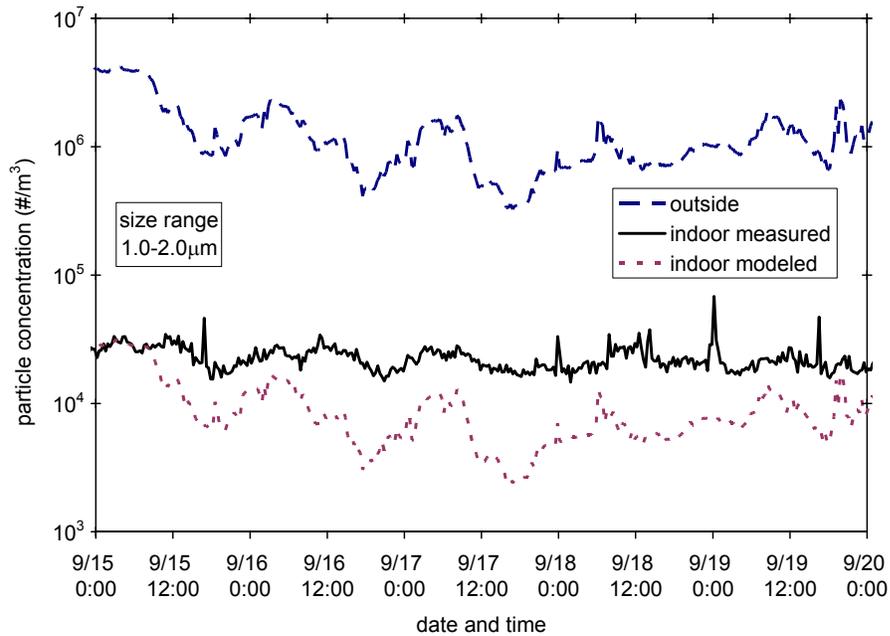


Figure 2-7d. Time-dependent measured and modeled concentrations of particles 1.0-2.0 μm in diameter at the Rocklin data center during 15-20 September 2006.

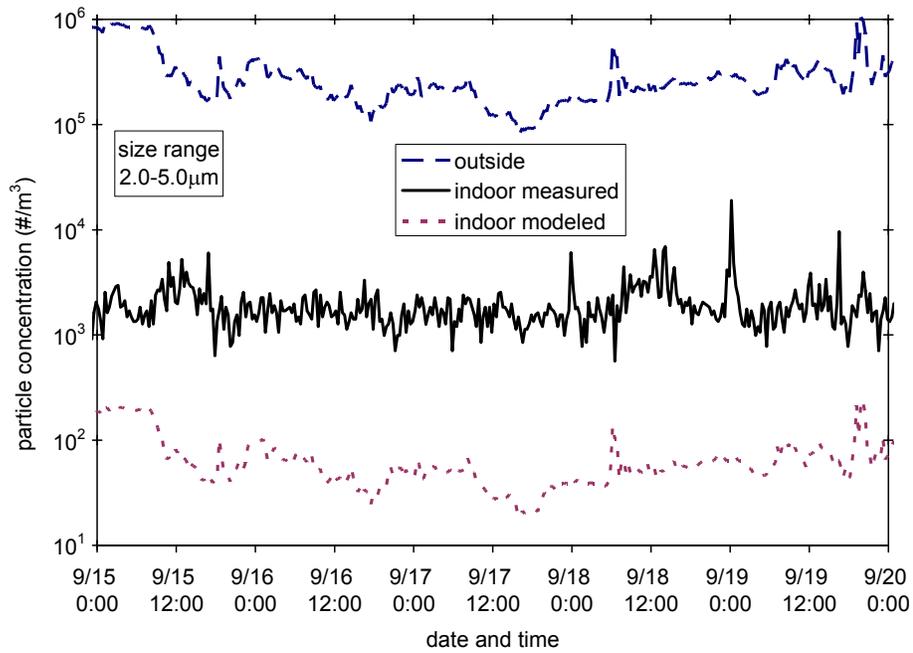


Figure 2-7e. Time-dependent measured and modeled concentrations of particles 2.0-5.0 μm in diameter at the Rocklin data center during 15-20 September 2006.

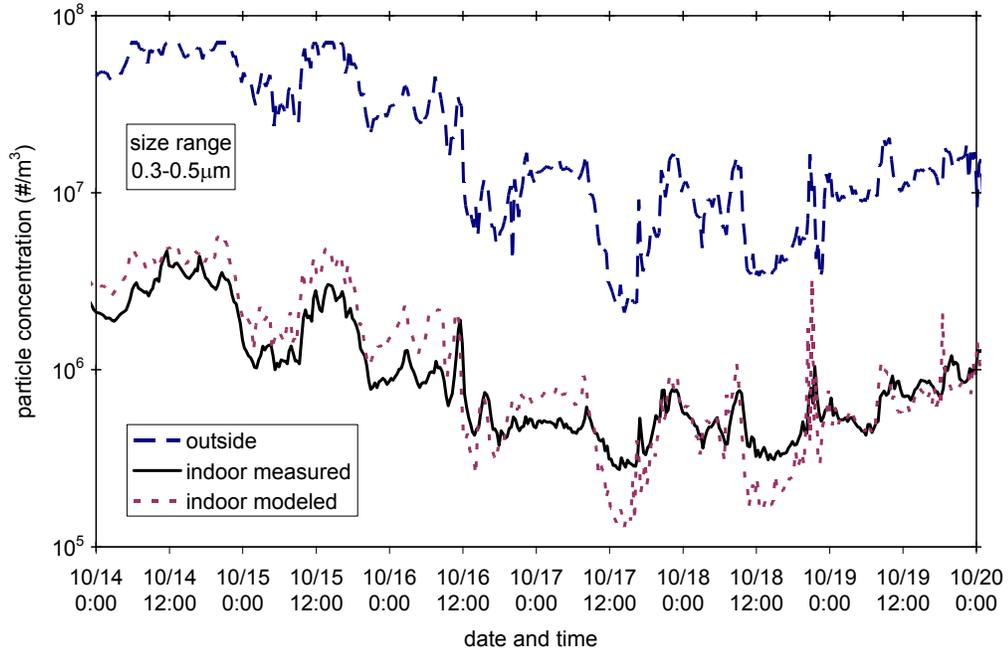


Figure 2-8a. Time-dependent measured and modeled concentrations of particles 0.3-0.5 μm in diameter at the Walnut Creek data center during 14-20 October 2006.

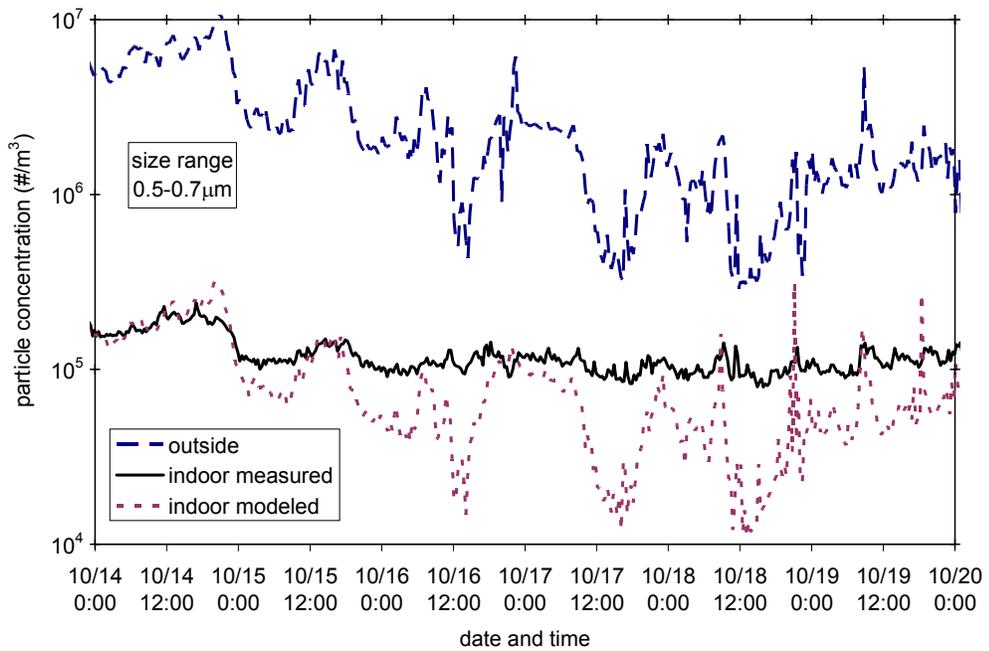


Figure 2-8b. Time-dependent measured and modeled concentrations of particles 0.5-0.7 μm in diameter at the Walnut Creek data center during 14-20 October 2006.

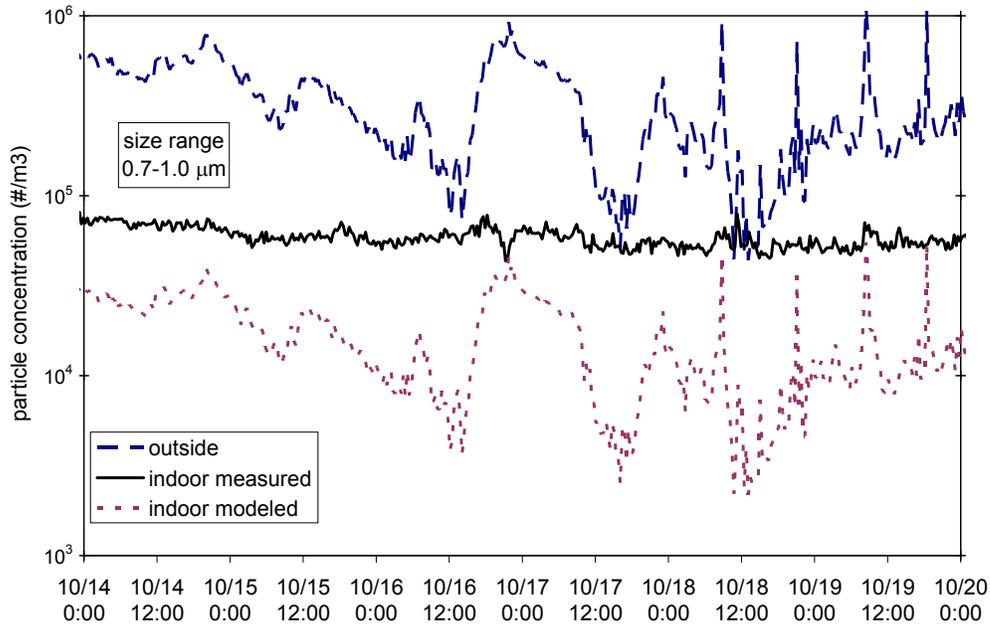


Figure 2-8c. Time-dependent measured and modeled concentrations of particles 0.7-1.0 μm in diameter at the Walnut Creek data center during 14-20 October 2006.

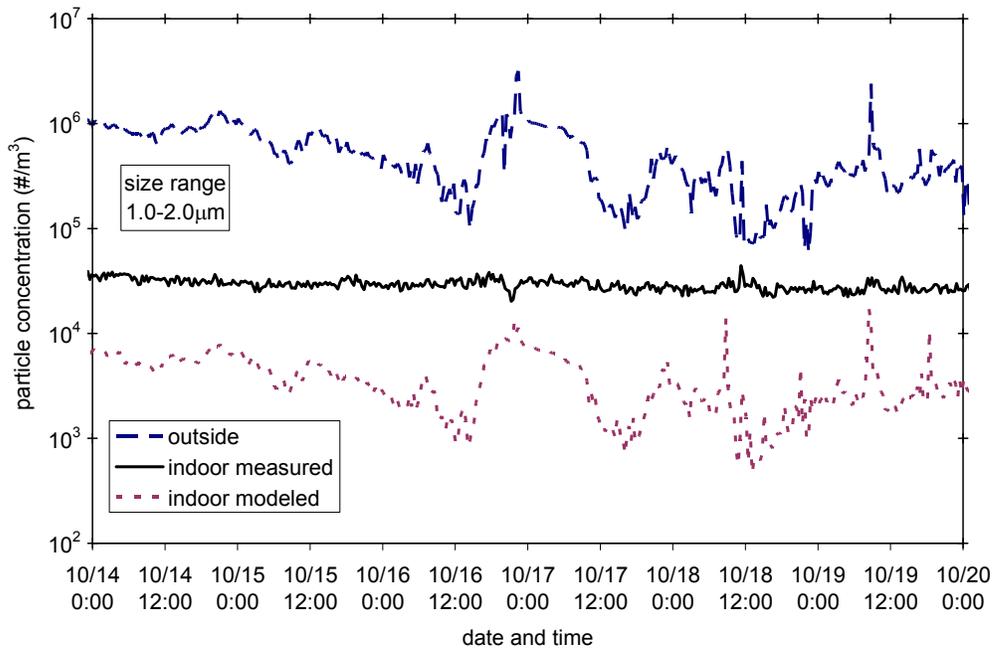


Figure 2-8d. Time-dependent measured and modeled concentrations of particles 1.0-2.0 μm in diameter at the Walnut Creek data center during 14-20 October 2006.

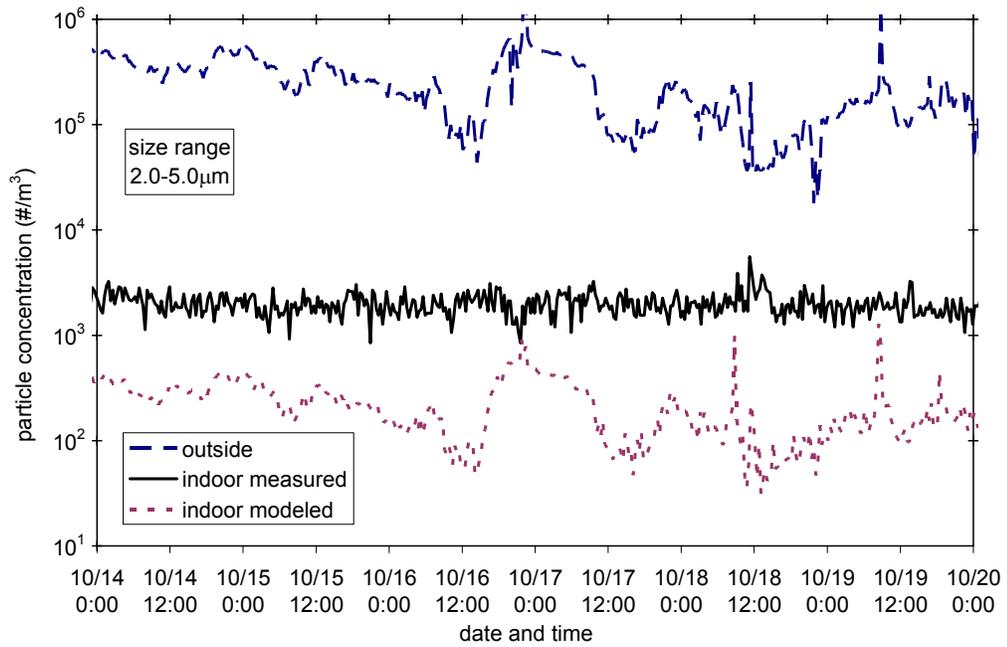


Figure 2-8e. Time-dependent measured and modeled concentrations of particles 2.0-5.0 μm in diameter at the Walnut Creek data center during 14-20 October 2006.

Figures 2-7a-e show that modeled indoor particle concentrations at the Rocklin site agree well with measurements in the smaller particle size bins, but particle concentrations in the larger size bins appear to be underpredicted by the model. Particles in the larger size bins also appear to be underpredicted by the model at the Walnut Creek site (Figures 2-8a-e). At both the Rocklin and Walnut Creek sites, the modeled indoor concentrations follow the fluctuations in the outdoor concentrations, while the measured indoor particle concentrations remained steady throughout the monitoring period for all size ranges except 0.3-0.5 μm . The influence of outdoor concentration fluctuations on indoor particle measurements appears to decrease with increasing particle size. The steady indoor particle concentrations measured in the larger size bins, and underpredicted by the model, suggest the presence of a weak, yet stable indoor source of particles, probably mechanically generated. Conceivably, this particle source might be worn or misaligned fan belts in the CRAC units, which has been previously suggested as a possible source of particles in data centers (ASHRAE, 2005).

As expected, indoor particle concentrations are strongly related to the rates at which outdoor air enters the building. Time-averaged indoor concentrations are approximately an order of magnitude lower at the two sites that use minimal outside air than at the Sunnyvale site, where a high percentage of outside air was used during a portion of the monitoring period (Figure 2-9). The indoor concentration responds rapidly to changes in the HVAC system setting between “low” and “high” outside-air modes. When in the “low” mode, results were similar to those at the other two study sites. During this mode of operation, the measured indoor concentrations were approximately 1 to 2 $\mu\text{g}/\text{m}^3$ for nearly all times, regardless of outdoor concentrations. During the “low”

outside-air period, the IPOP was about 3%, which is comparable in magnitude to values at the other two sites (~ 1%).

A sudden increase in particle concentration is apparent in Figure 2-9 whenever the HVAC system switches to the “high” outside-air mode. The increase in indoor particle concentration begins toward the end of the day, around midnight, and then typically ends late in the morning. During the “high” outside air mode, the indoor concentration increases by nearly an order of magnitude (as compared with the “low” outside air mode) and varies more directly in response to changing outdoor concentrations. The indoor concentration shifts from approximately 3% to 36% of the outdoor concentration. The higher indoor concentration is sustained until the HVAC returns to the “low” outside-air mode.

Figure 2-10a-d shows modeling results for each of the particle size categories measured at the Sunnyvale site (0.5-5 μm). The modeled indoor particle concentrations agree well with measurements during both “low” and “high” outside-air modes, except for the particle size range 2.0-5.0 μm , which was slightly underrepresented by the model for “high” mode operation.

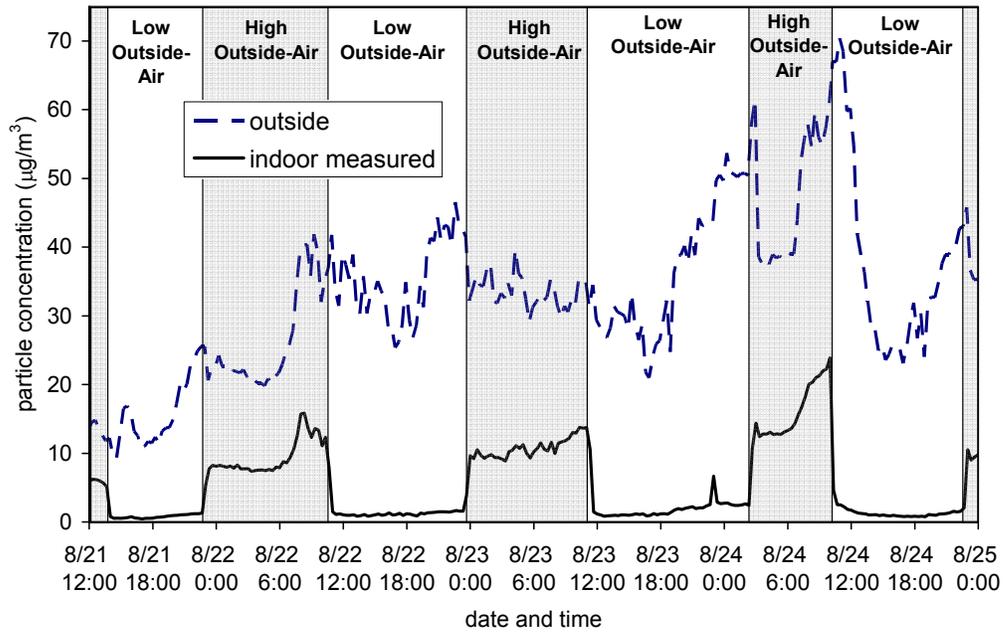


Figure 2-9. Measured time-dependent particle mass concentrations at the Sunnyvale data center during 21-25 August 2006. Particle concentration represents 0.5-5 μm particulate matter.

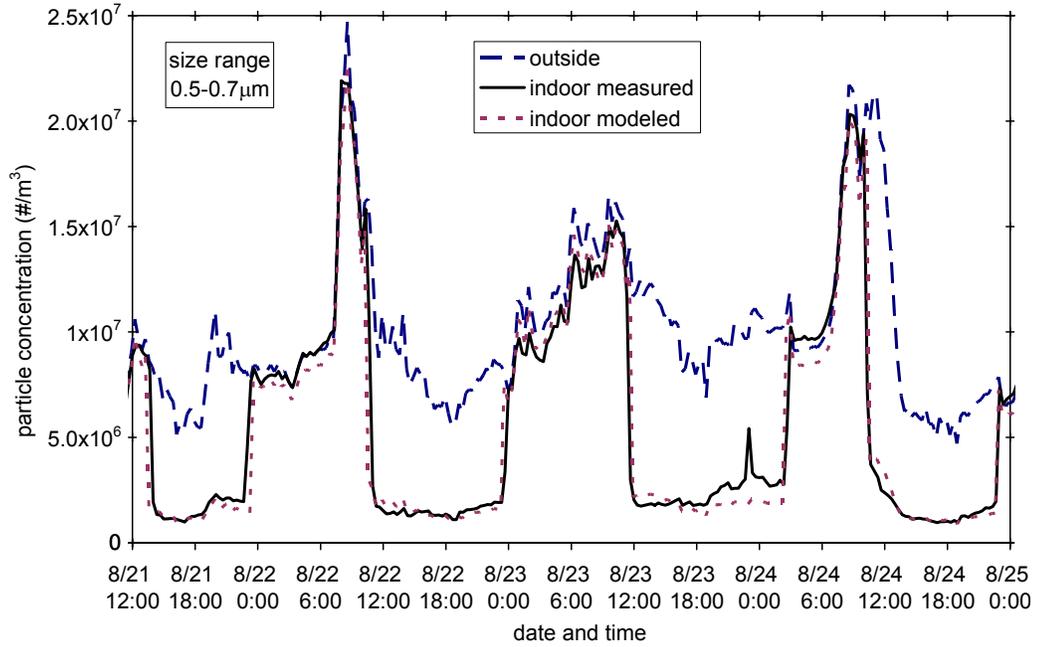


Figure 2-10a. Time-dependent measured concentrations of particles 0.5-0.7 μm in diameter at the Sunnyvale data center during 21-25 August 2006.

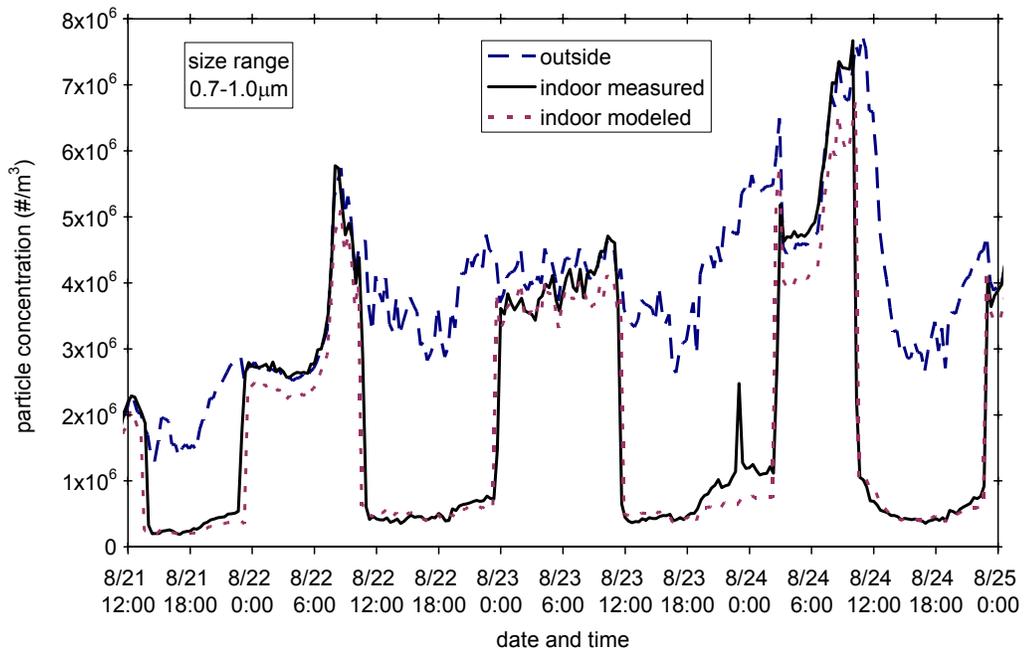


Figure 2-10b. Time-dependent measured and modeled concentrations of particles 0.7-1.0 μm in diameter at the Sunnyvale data center during 21-25 August 2006.

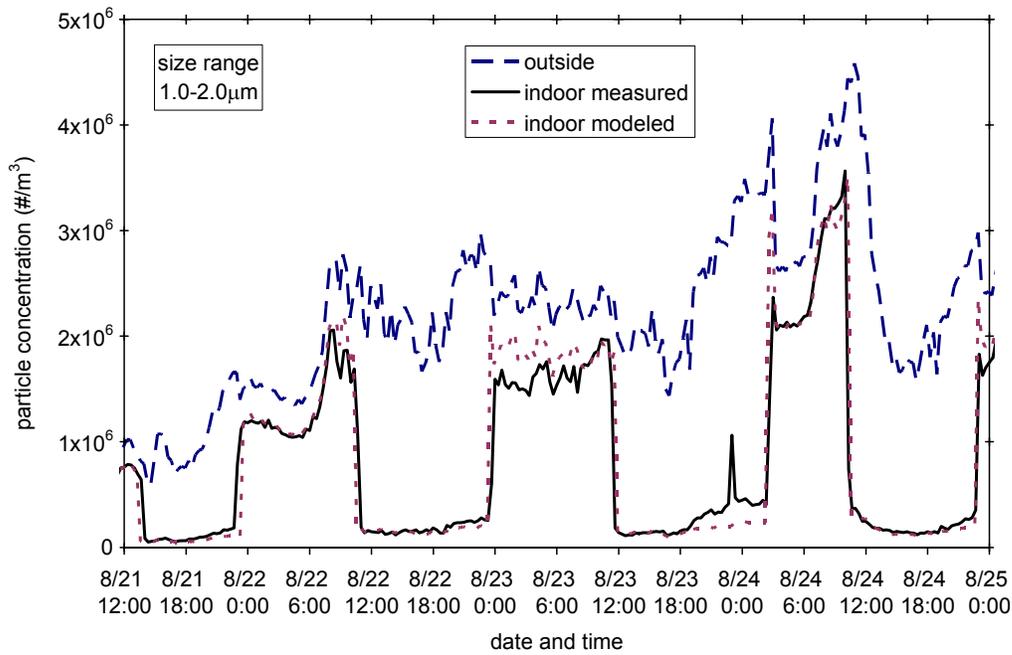


Figure 2-10c. Time-dependent measured and modeled concentrations of particles 1.0-2.0 μm in diameter at the Sunnyvale data center during 21-25 August 2006.

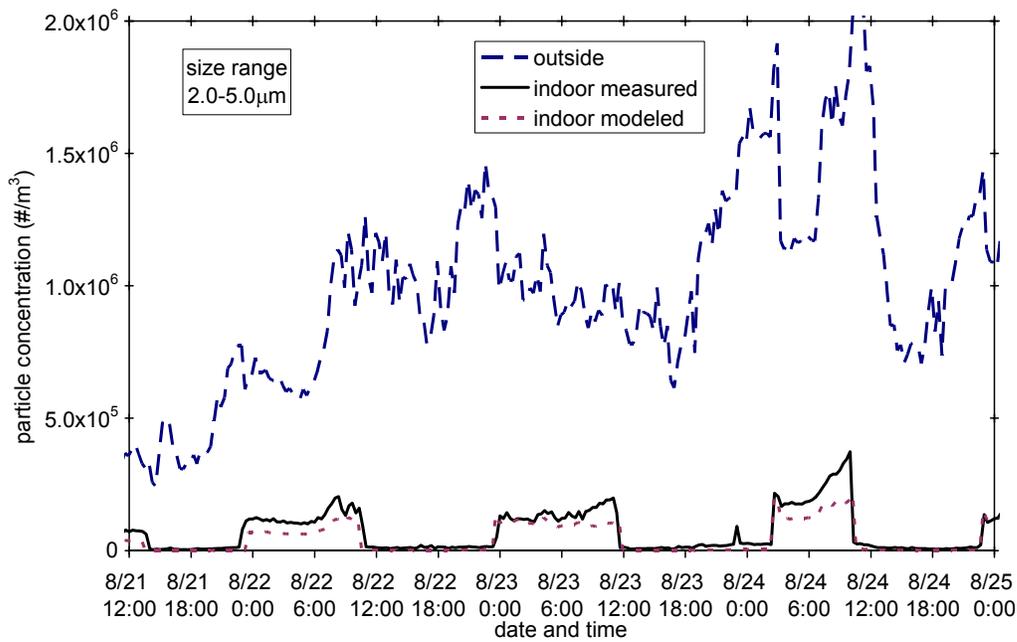


Figure 2-10d. Time-dependent measured and modeled concentrations of particles 2.0-5.0 μm in diameter at the Sunnyvale data center during 21-25 August 2006.

2.3.2. Particle Sources and Sinks

Outdoor air appears to be the main source of airborne particle mass in all three data centers. Additional potential indoor sources of particles in data centers include occupant activities, fan belt wear, and resuspension from occupant activities (Shields and Weschler, 1998; Brusse and Sampson, 2004; Roth, 2005). While indoor particle generation may contribute to the particle concentrations in data centers, modeled indoor mass concentrations assuming no indoor-generated particles match well the indoor measurements. When comparing the measured indoor concentrations relative to the measured outdoor concentrations, the mean absolute deviation in IPOP between model and measurement is 1%, 1%, and 3% for the Walnut Creek, Rocklin, and Sunnyvale sites, respectively. This level of agreement indicates that any indoor source of particles during the monitoring periods was small in relation to the supply of particles from outdoor air. Indoor measurements show a fairly steady indoor particle concentration with few aberrant increases or decreases, indicating that any sporadic indoor particle source, such as that from occasional occupant activities, has little impact on time-averaged indoor concentrations. Data centers typically have air filters for both outdoor and recirculated air. Because of the importance of outdoor air as a source of indoor particles, the results of this study suggest that data center particle mitigation efforts might benefit from focusing filtration more heavily on the entering outdoor air.

The difference between measured particle concentrations at the Rocklin and Walnut Creek sites, summarized for number concentration in Table 2-3, is a mass concentration of approximately $0.1 \mu\text{g}/\text{m}^3$. While this concentration is small relative to ambient concentrations, the discrepancy is clearly detectable against the low indoor

concentrations measured at these conventional data centers that supply minimal outside air. The time-series and size-dependent discrepancies between model and measurement presented in Figures 2-7a-e and 2-8a-e suggest that there is a stable, yet weak mechanical source of indoor airborne particles. A potential source is the CRAC-unit fan belts. Reconciling model predictions to measurement results suggests an indoor emission source of approximately 1 mg/h per fan belt at each of the two study sites, which would correspond to a 1-5% loss of fan-belt mass over the typical fan-belt lifetime of six months.

Once particles enter the data center, their possible fates are (i) to be exhausted with the ventilation, (ii) captured during filtration, or (iii) deposited onto an interior surface. The sum of these three potential loss terms make up the denominator in equation (1), with λ_{out} representing the ventilation loss rate coefficient, β_i representing surface deposition, and the product of $\eta_{rec,i} \lambda_{rec}$ representing removal via filtration of recirculated air. The relative contribution of these particle sinks varies with particle size and among the data centers.

Table 2-3. Average indoor modeled and measured particle concentrations at three data center sites ($\#/m^3$)

particle size range	Walnut Creek		Rocklin		Sunnyvale	
	measured	modeled	measured	modeled	measured	modeled
0.3-0.5 mm	1.2×10^6	1.6×10^6	2.7×10^6	1.3×10^6	n/a	n/a
0.5-0.7 mm	1.2×10^5	8.9×10^4	1.4×10^5	1.1×10^5	6.6×10^6	6.2×10^6
0.7-1.0 mm	6.0×10^4	1.7×10^4	6.7×10^4	3.7×10^4	1.8×10^6	1.6×10^6
1.0-2.0 mm	2.9×10^4	4.0×10^3	3.0×10^4	9.5×10^3	6.5×10^5	6.6×10^5
2.0-5.0 mm	2.0×10^3	2.3×10^2	3.3×10^3	7.1×10^1	5.2×10^4	3.2×10^4

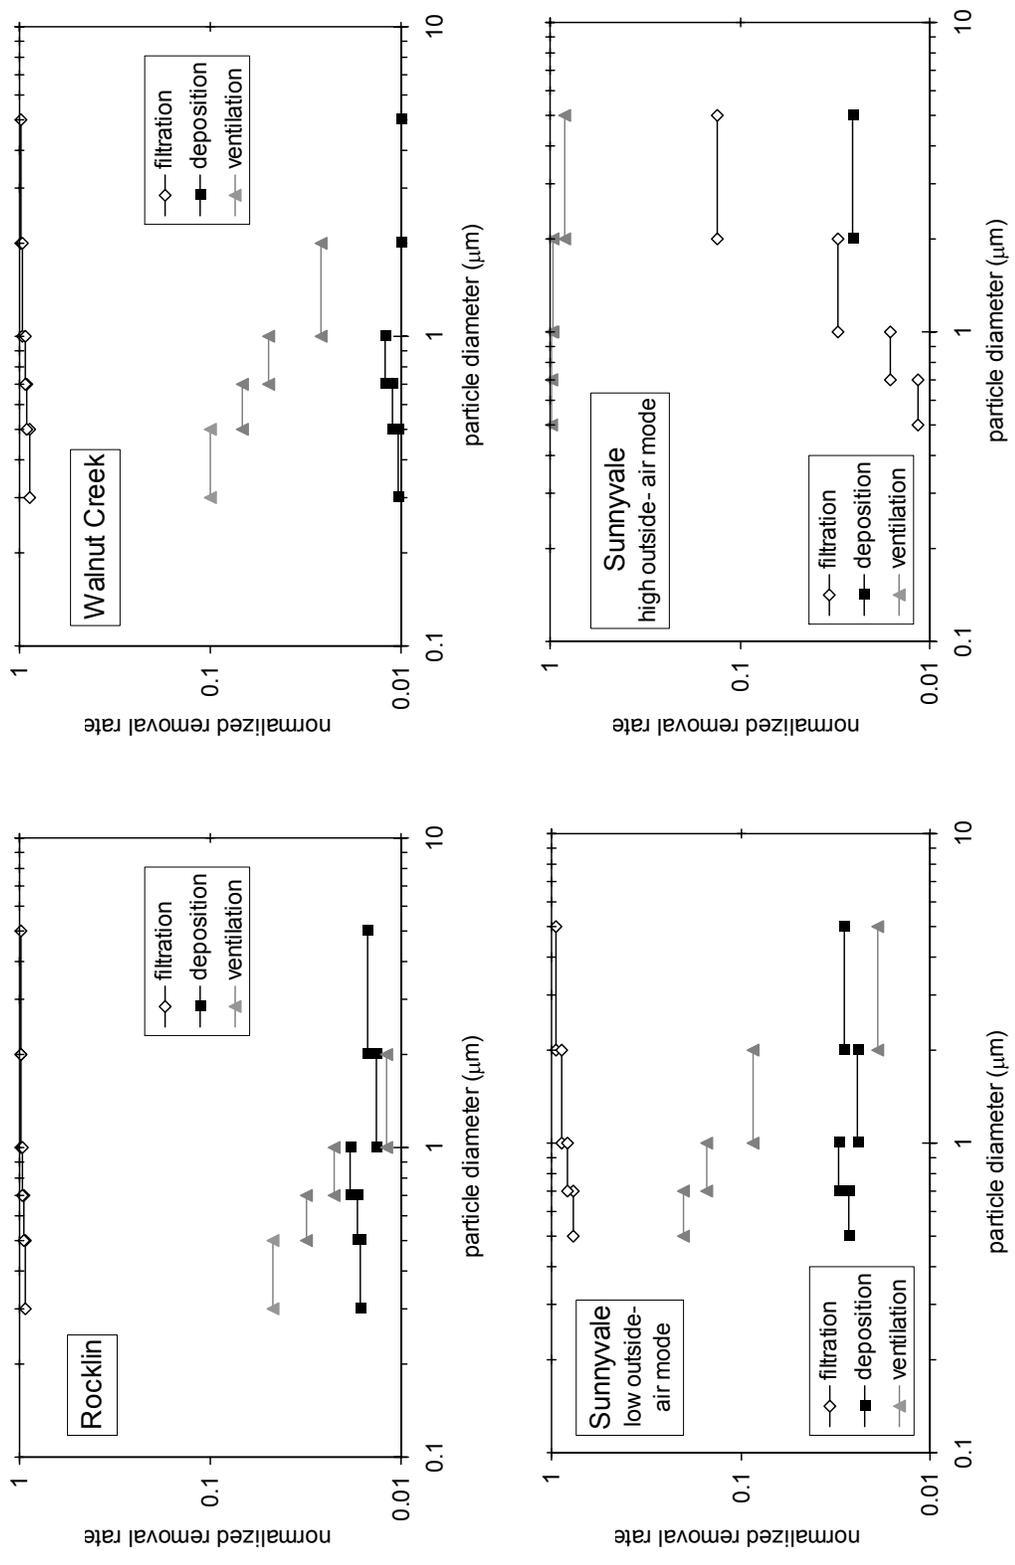


Figure 2-11. Modeled particle fates at each study site.

The normalized rate of particle removal by each loss mechanism is presented in Figure 2-11 for each particle size range at each study site. Filtration dominates particle removal at the Rocklin and Walnut Creek sites. ASHRAE (2005) recommends 40% filters in data centers that use minimal outside air. This type of filter was observed in the CRAC units and in most of the outside air handlers at the data centers monitored in this study. Even though the CRAC units have filters with modest efficiency, the large rate of recirculating flow through the CRAC units relative to the amount of outside air introduced into the data center results in high relative particle removal by this means. At the Sunnyvale site, when the HVAC system is in the “high” outside-air mode, ventilation is the dominant removal mechanism owing to the relatively high proportion of indoor air exhausted from the data center. Filtration dominates during the “low” outside-air mode at the Sunnyvale site and the relative contribution of the loss terms is similar to that found at the other two sites.

2.3.3 Sulfate Predictions

The modeled indoor particle concentration and corresponding IPOP values depend on the size distribution of outdoor particles. Within the particle size range studied (0.3-5 μm), outdoor concentrations that have greater proportion of their mass in larger particles will result in lower modeled IPOP values, since larger particles are more efficiently removed by filtration and by surface deposition. Conversely, a greater contribution of total mass from smaller particles would reduce interior loss rates, resulting in a higher IPOP value. The size distribution of outdoor particles varies by time and location and also by particle chemical composition. Since sulfate represents a particle type of particular concern for equipment reliability, its mass distribution was

applied to the model to predict the IPOP of sulfate at the Rocklin and Sunnyvale sites. At the Rocklin site, the modeled IPOP increases from less than 1% for total outdoor particle mass to about 2% for sulfate mass. At the Sunnyvale site, the modeled IPOP increases from approximately 3% (for total mass) to about 19% (for sulfate) in the “low” outside-air mode and from 36% (total mass) to 88% (sulfate) for the “high” outside-air mode. Since the particle size range studied at the Sunnyvale site was curtailed to 0.5-5 μm particle diameter, the sulfate size distribution includes more mass from smaller particles (<0.5 μm as opposed to <0.3 μm at the other site), which is not represented in studied particle size range. This results in a greater increase from total outdoor particle mass IPOP to sulfate mass IPOP at the Sunnyvale site, especially during the “low” outside air mode when indoor particle mass is predominantly represented by smaller particles.

2.4. Conclusions

Prudent implementation of energy-saving measures that would expose data center equipment to more outside air requires two tiers of investigation: first, understanding how these design measures would change indoor particle concentrations, and second, understanding how such changes in concentration would influence equipment reliability. This study contributes to the former goal by presenting the first published measurements of particle concentrations in operating data centers. The data and their interpretation provide baseline information for conditions in typical data centers, revealing significantly lower particle concentrations than typically found in offices or residential buildings. Estimates using a parsimonious material-balance model match fairly well with the empirical results. This agreement indicates that the dominant particle sources and losses

have been identified and are being appropriately described, increasing the basis for confidence in one's ability to predict particle concentrations in data centers under different scenarios. Measurements taken at the Sunnyvale site, where high flow rates of outside air are already deployed to save energy, show nearly an order of magnitude increase in particle concentration during "high" outside-air periods as compared to the "low" outside air periods. Sulfate modeling results indicate that this increase may be even greater when including particles smaller than the size range measured in this study. While these data confirm and quantify the increase in particle concentrations caused by using more outside air, the equipment risk associated with such concentration increases remain unknown. We note that average indoor particle concentrations at Sunnyvale still were well below particle limits recommended by some server manufacturers and were less than the limit suggested by ASHRAE. The results presented here provide a partial foundation for future work to investigate the risk to data center equipment posed by expected particle levels. A more thorough understanding of the equipment reliability risks associated with supplying greater outside air in data centers will help determine what conditions are safe for this energy-saving measure. One can also explore mitigation alternatives, such as enhanced filtration, that aim to improve energy efficiency while simultaneously minimizing risk to electronic equipment from the deposition of particulate matter. Overall, such efforts can help temper the growing energy demand of data centers and thereby allow the expansion of information technology to proceed in a more sustainable fashion.

2.A. Appendix: Additional Data Center Particle Measurements

Size-resolved particle concentrations were measured in eight northern California data centers in 2006 using a pair of optical particle counters. Measurements from three of these data centers — Rocklin, Walnut Creek, and Sunnyvale — were discussed in the main part of Chapter 2. This appendix presents the measured particle data from the remaining five data centers: Redwood City, Dublin, San Francisco, Oakland, and Berkeley. These remaining data centers were operated without economizers and the results presented in this appendix provide further support that non-economizer data centers maintain relatively very low indoor particle concentration that are not strongly influenced by fluctuations in outdoor particle levels. Two separate obstacles arose during the data acquisition process that hindered an in-depth analysis of the measurements from the remaining data centers, resulting in their exclusion from the main body of the chapter. First, the complexity of the floor plans in these data centers and the inability to confirm HVAC parameters, such as building infiltration rates, prevented the modeling of indoor concentrations at these locations. Second, the measurements at some of these data centers were taken early in the experimental process and improper sampling times resulted in some loss of data.

Measurements at most sites were taken in five-minute intervals, with the OPC drawing in air at a rate of 2.8 L/min and providing a size-specific count of the particle in the sample airstream. The counter would then pause for 20 minutes before beginning the next particle counting cycle. Particle counting cycles were initially longer. At the first two sites monitored, Berkeley and Oakland, particles were counted for 20 minutes (rather

than five minutes). The Met-One 237B OPCs used in this experiment are limited to a six-digit count (i.e. 999,999) for each particle size bin. Under these longer counting times the particle count for 0.3-0.5 μm particle size range reached the counter limit for many of the outdoor measurements. This outcome indicates that the measured outdoor concentration underrepresents the actual outdoor particle concentration. On a mass basis, this underrepresentation is expected to not be important since the 0.3-0.5 μm particle size range is a small contributor to the total outside particulate matter mass concentration (typically less than 10%), relative to the other size bins, for all of the data centers monitored in this study. All mass concentrations presented here represent particles 0.3-5.0 μm in diameter and are calculated using the same methods presented in Chapter 2.

Figure 2A-1 shows particle concentrations measured at the Redwood City site from 29 September to 6 October 2006. The data center is located within a building that is part of a corporate office park located less than one mile from a major freeway, placing the site in close proximity to significant vehicular traffic. Air within the data center is conditioned by passing through CRAC units located on the data center floor. A small percentage of outside air enters the data center through ceiling vents. The entering air is a mixture of outside air and return air that has passed through a filter in a rooftop AHU. The measured outdoor particle concentration averaged about $10 \mu\text{g}/\text{m}^3$, with concentrations typically increasing in the afternoon hours until about midnight at which point they would begin to decrease. The indoor particle concentration averaged about $0.8 \mu\text{g}/\text{m}^3$ and rarely increased beyond $2.0 \mu\text{g}/\text{m}^3$. The indoor particle concentration essentially toggles from $1.0 \mu\text{g}/\text{m}^3$ during the day to about $0.5 \mu\text{g}/\text{m}^3$ in the evening and early morning hours. Figures 2A-2a-e show a slight, yet consistent, fluctuation of indoor

particle concentrations; a sudden increase and then decrease of smaller particles at approximately 6:00 AM and 8:00 PM respectively. This fluctuation is clearly visible for particles in the 0.3-0.5 μm range, but decreases incrementally with each larger particle bin and is undetectable in the 2.0-5.0 μm size bin. The elevated concentrations during common working hours are most likely a consequence of changes in HVAC operations and infiltration (due to the opening and closing of doors) during occupant activity.

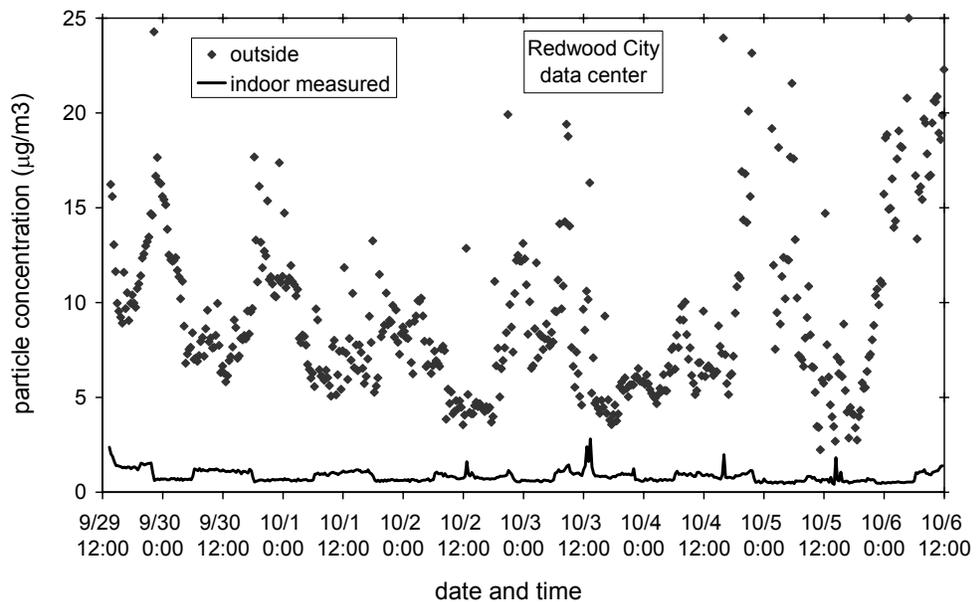


Figure 2A-1. Measured time-dependent particle mass concentrations at the Redwood City site during 29 September – 6 October 2006. Concentrations represent particles 0.3-5.0 μm in diameter.

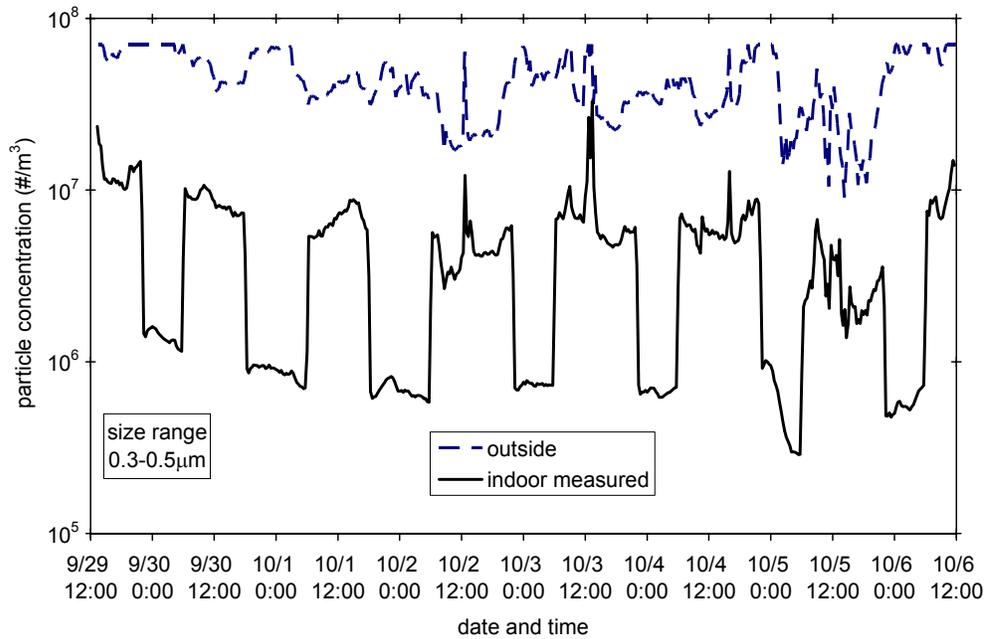


Figure 2A-2a. Time-dependent measured concentrations of particles 0.3-0.5 μm in diameter at the Redwood City site during 29 September – 6 October 2006.

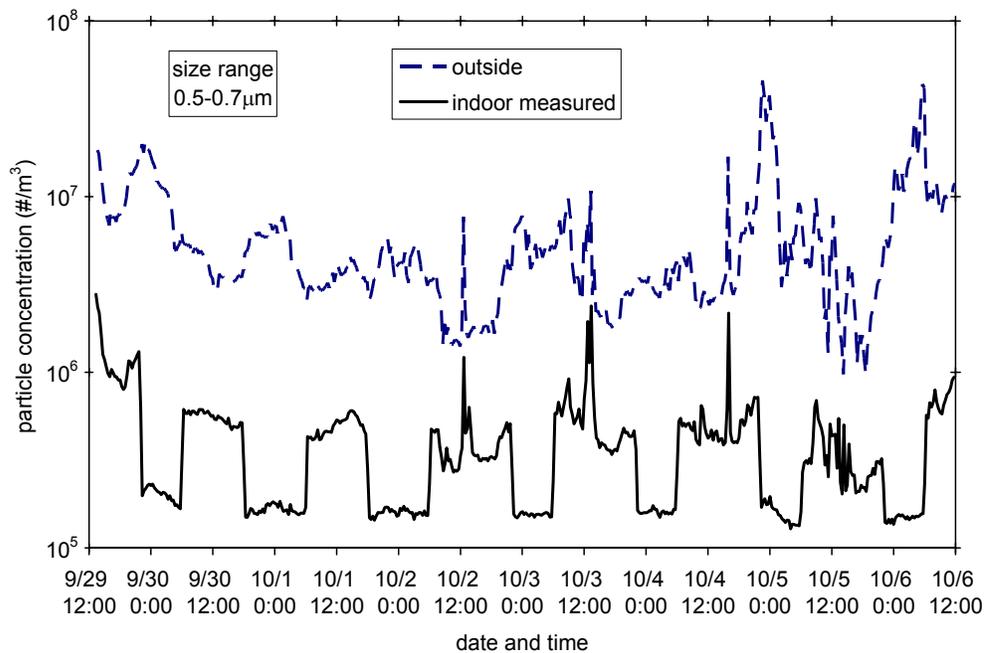


Figure 2A-2b. Time-dependent measured concentrations of particles 0.5-0.7 μm in diameter at the Redwood City site during 29 September – 6 October 2006.

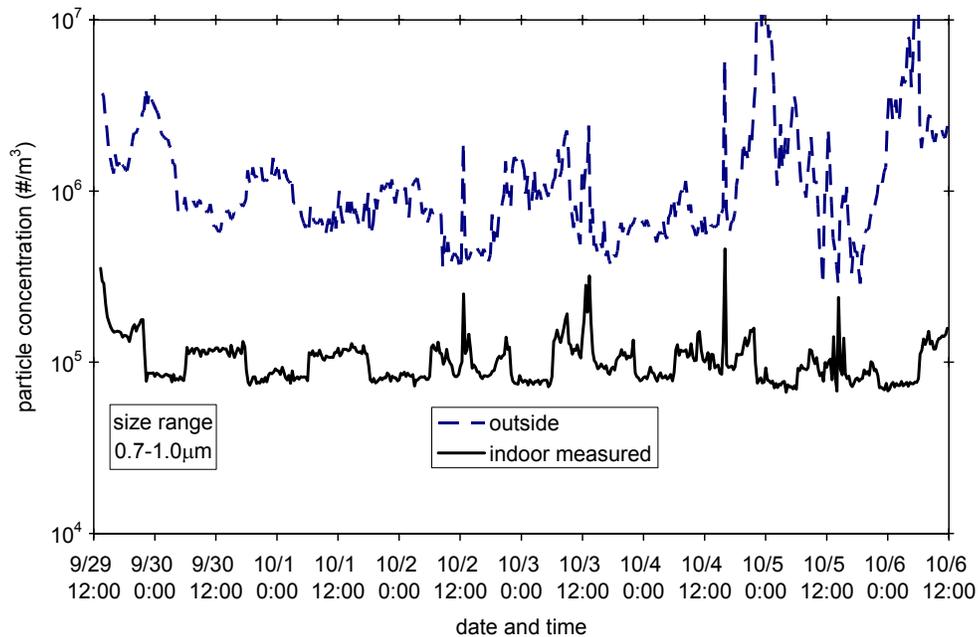


Figure 2A-2c. Time-dependent measured concentrations of particles 0.7-1.0 μm in diameter at the Redwood City site during 29 September – 6 October 2006.

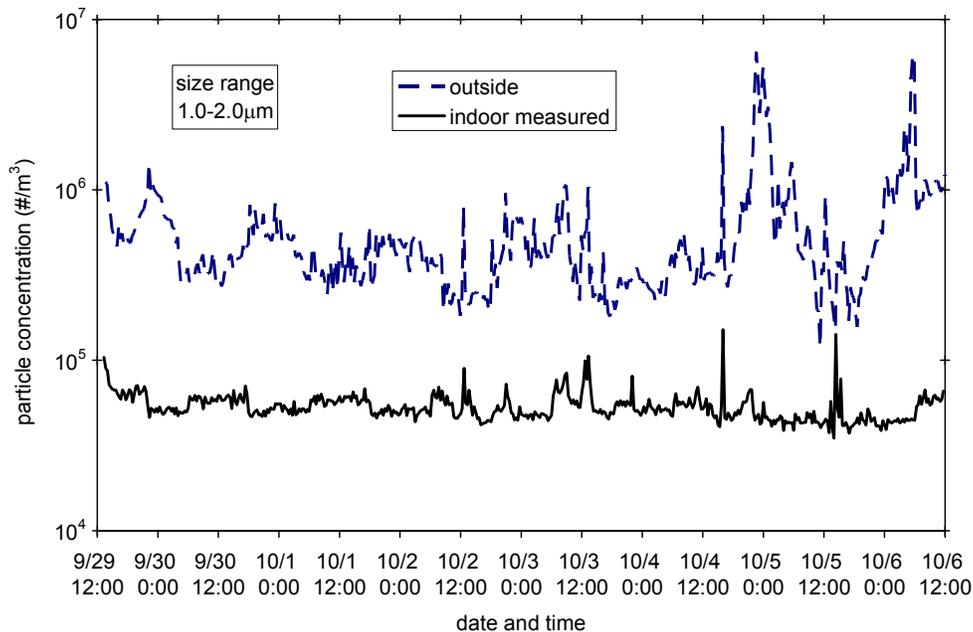


Figure 2A-2d. Time-dependent measured concentrations of particles 1.0-2.0 μm in diameter at the Redwood City site during 29 September – 6 October 2006.

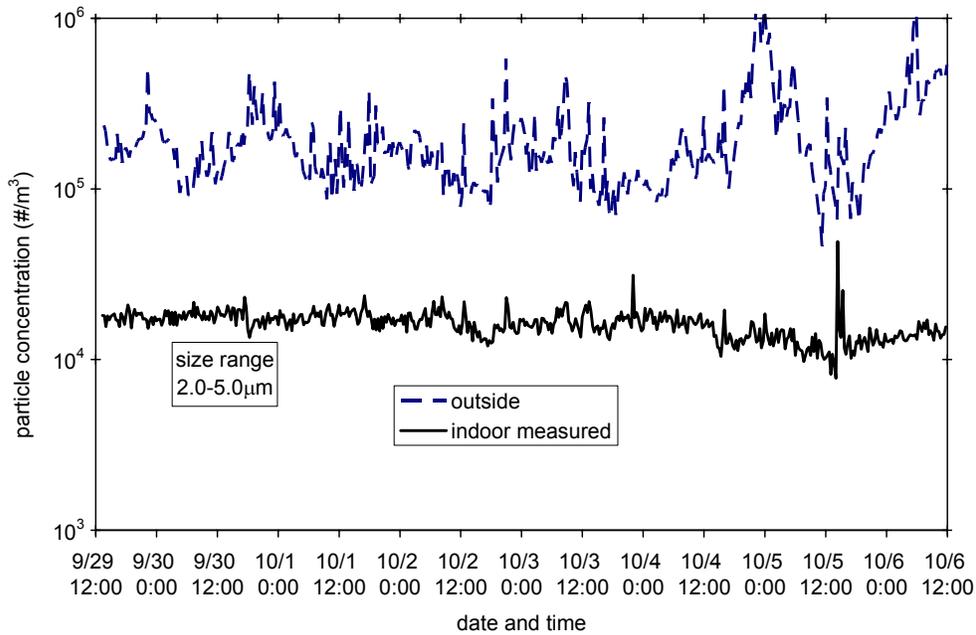


Figure 2A-2e. Time-dependent measured concentrations of particles 2.0 -5.0 μm in diameter at the Redwood City site during 29 September – 6 October 2006.

Particle measurements that were taken at the Berkeley data center site 4-8 May, 2006 are presented in Figure 2A-3. This data center is situated within an office building and uses an unconventional cooling design that incorporates both floor CRAC units and the building's general HVAC system. A small percentage of outside air mixed with return air is directly ducted into the data center room. Air within the room entered the floor CRAC units and passes through an underfloor distribution system. A small percentage of air within the room is exhausted through a single large vent on one end of the data center. Exhausted air is returned to a mechanical room, where it is mixed with outside air that enters the mechanical room from an air intake located in a semi-enclosed parking area. The mixture of filtered outside air and return air then passes through a set

of filters before being ducted back into the data center. Figure 2A-3 shows that particulate matter concentration measured in the outside air averaged about $27 \mu\text{g}/\text{m}^3$, with the highest concentrations typically measured during the early morning hours of the monitoring period. Figure 2A-4a shows that while monitoring the outside air at the Berkeley data center the particle count for 0.3-0.5 μm particle size range reached the counter limit for all of the measurements. These results indicate that the measured outdoor concentration underrepresents the actual outdoor particle concentration. However, the 0.3-0.5 μm particle size range contributes less than $1 \mu\text{g}/\text{m}^3$ to the total $27 \mu\text{g}/\text{m}^3$ for recorded conditions, indicating that even a large proportional increase in particle count, the 0.3-0.5 μm size range would only make a small contribution to the total particulate matter mass concentration. Figure 2A-3 shows that the indoor concentration at the Berkeley site is low and steady during the monitoring period. Indoor particle concentrations averaged about $0.3 \mu\text{g}/\text{m}^3$ and rarely fluctuated by more than $0.1 \mu\text{g}/\text{m}^3$ during the monitoring period. Figures 2A-4a-e show steady indoor particle concentrations with minimal fluctuations for all particle size bins.

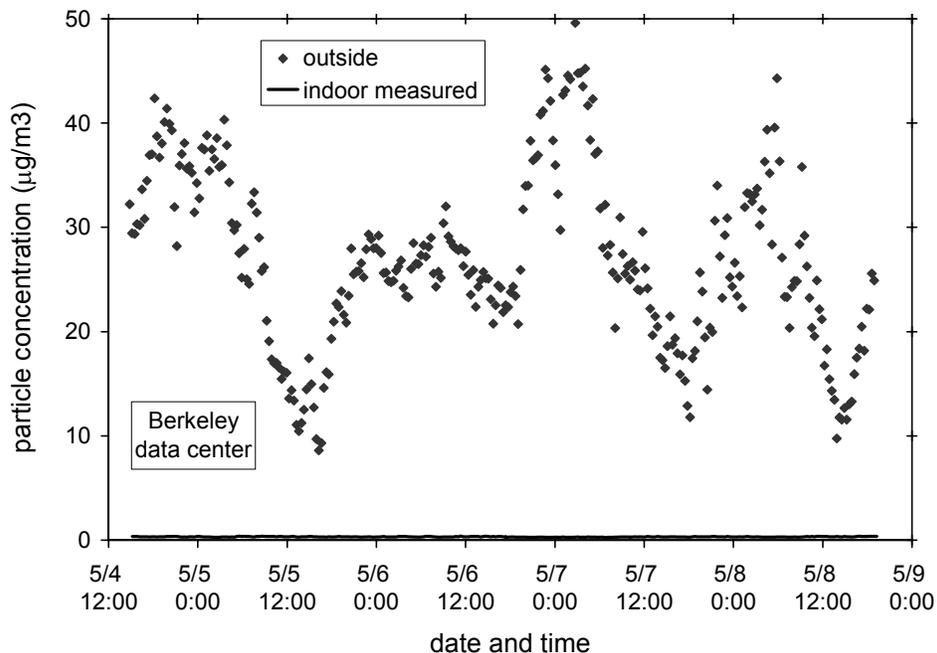


Figure 2A-3. Measured time-dependent particle mass concentrations at the Berkeley site during 4-8 May 2006. Concentrations represent particles 0.3-5.0 μm in diameter.

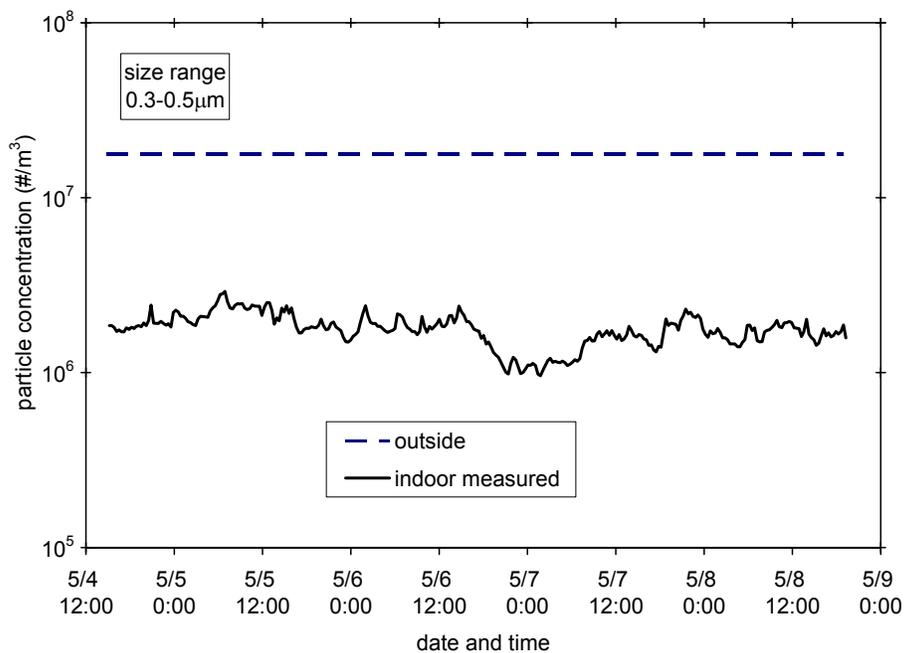


Figure 2A-4a. Time-dependent measured concentrations of particles 0.3-0.5 μm in diameter at the Berkeley site during 4-8 May 2006. The horizontal line representing the outdoor particle concentration indicates the OPC researched the particle counting limit and that the actual particle concentration is greater than what is presented here.

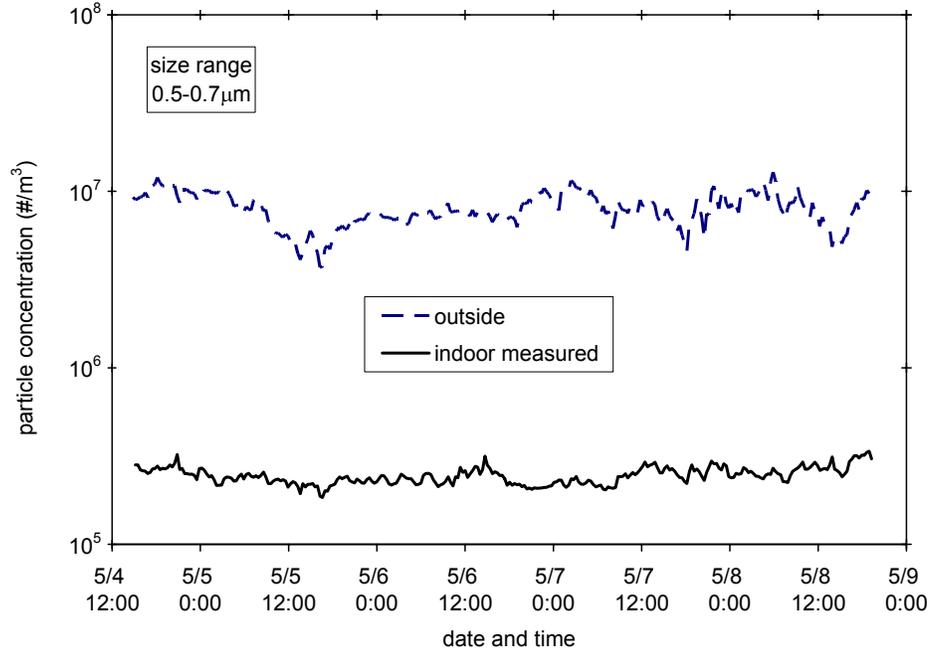


Figure 2A-4b. Time-dependent measured concentrations of particles 0.5-0.7 μm in diameter at the Berkeley site during 4-8 May 2006.

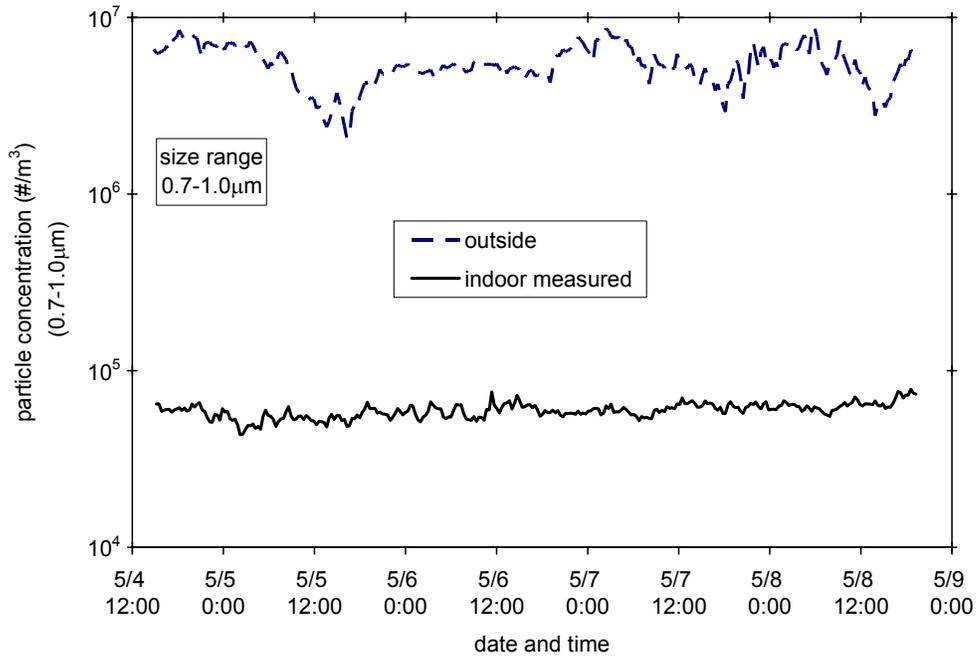


Figure 2A-4c. Time-dependent measured concentrations of particles 0.7-1.0 μm in diameter at the Berkeley site during 4-8 May 2006.

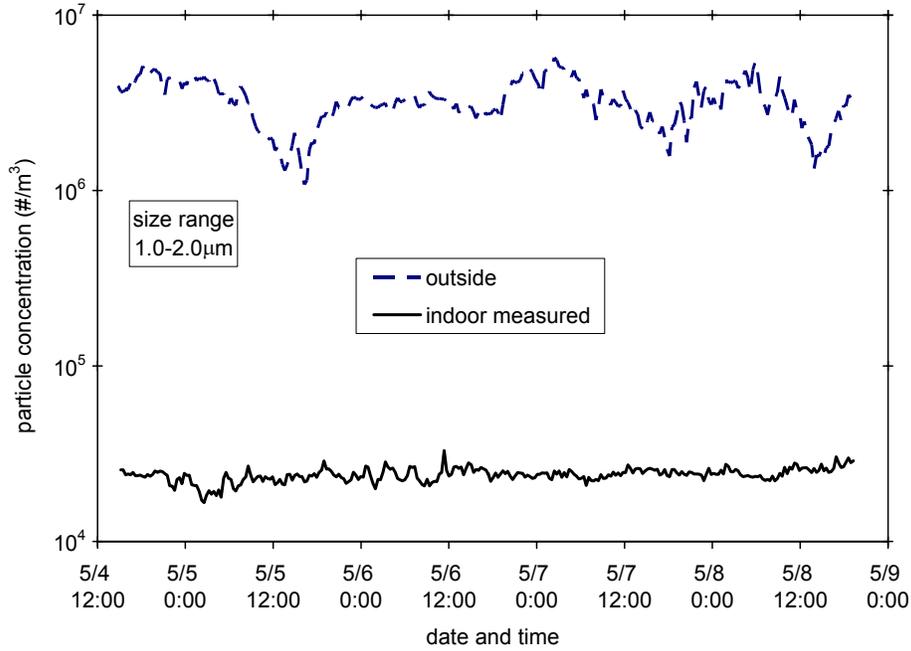


Figure 2A-4d. Time-dependent measured concentrations of particles 1.0-2.0 μm in diameter at the Berkeley site during 4-8 May 2006.

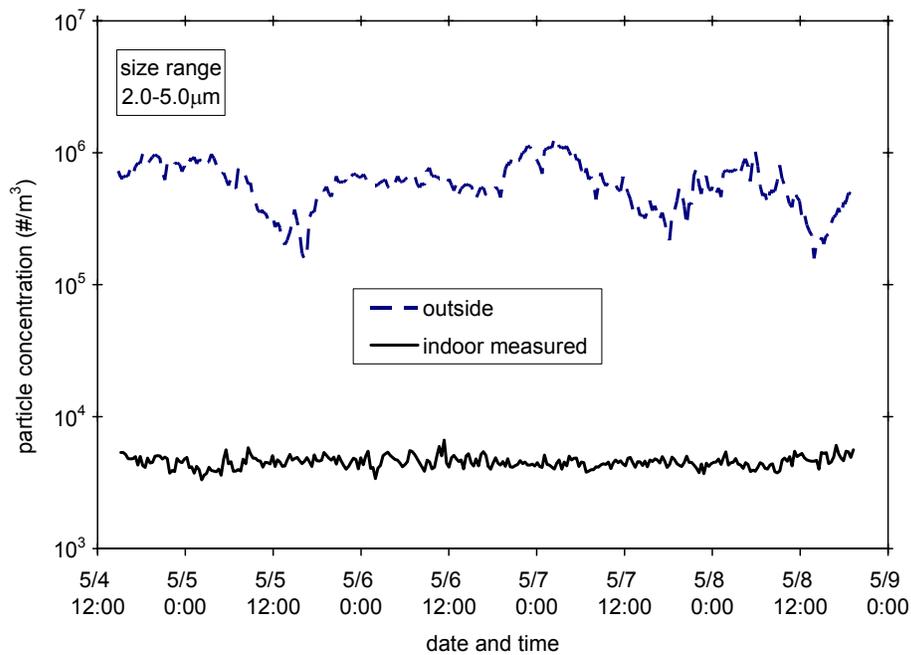


Figure 2A-4e. Time-dependent measured concentrations of particles 2.0-5.0 μm in diameter at the Berkeley site during 4-8 May 2006.

Particle measurements were taken at the Oakland data center 15 - 22 May 2006. At this site, no direct outside air is ducted into the data center zone. Infiltration from adjacent office zones is the only source of outside air entering the data center. Within the room, air is circulated and conditioned with CRAC units, which supply air to an underfloor distribution system. Figure 2A-5 shows that the average outside particle concentration was approximately $14 \mu\text{g}/\text{m}^3$, with concentrations typically increasing during the morning hours. Much less of a morning increase occurs during the final two days of the measurement period, which coincides with the weekend, indicating that the morning particle increases may be a consequence of increased vehicular traffic emissions near this downtown Oakland location. As with the Berkeley site monitoring, the particle count for the $0.3\text{-}0.5 \mu\text{m}$ particle size range reached the counter limit for most of the outdoor measurements. This outcome is clearly apparent in Figure 2A-6a and indicates that the measured outdoor concentrations is less than the actual particle concentration. Again, though, the contribution of this size bin to the total particle mass is expected to be small. Similar to the other data centers monitored with minimal infiltration, the indoor particle concentration at the Oakland site was significantly lower and steadier than the outdoor measurements. The average indoor particle concentration was less than $0.2 \mu\text{g}/\text{m}^3$ and rarely increased above $0.3 \mu\text{g}/\text{m}^3$.

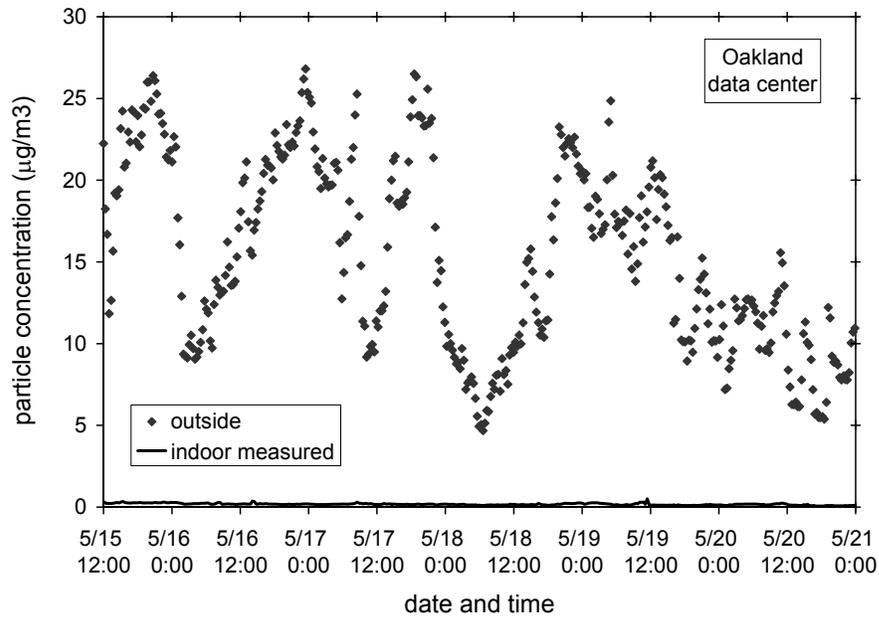


Figure 2A-5. Measured time-dependent particle mass concentrations at the Oakland site during 15-21 May 2006. Concentrations represent particles 0.3-5.0 μm in diameter.

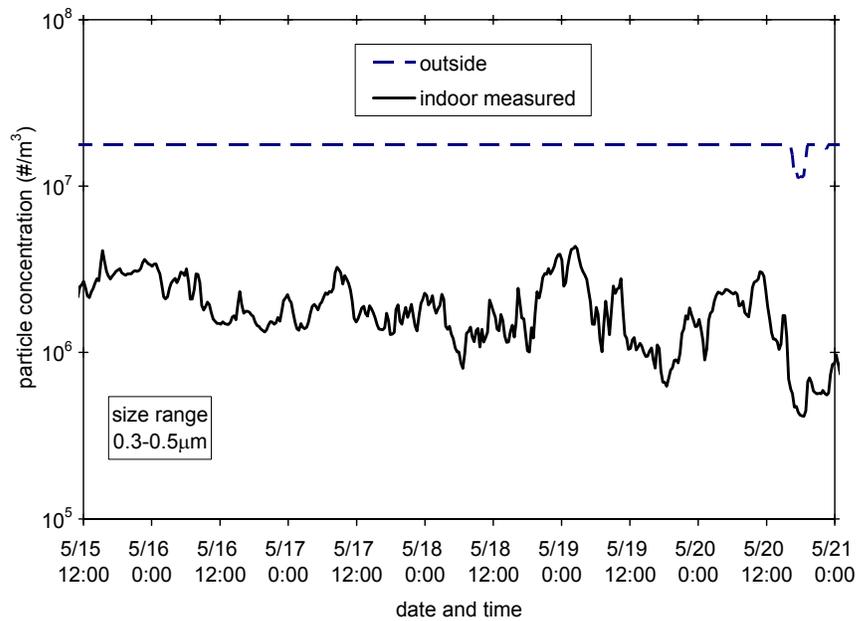


Figure 2A-6a. Time-dependent measured concentrations of particles 0.3-0.5 μm in diameter at the Oakland site during 15-21 May 2006. The horizontal line representing portions of the outdoor particle concentration indicates the OPC researched the particle counting limit and that the actual particle concentration is greater than what is presented here.

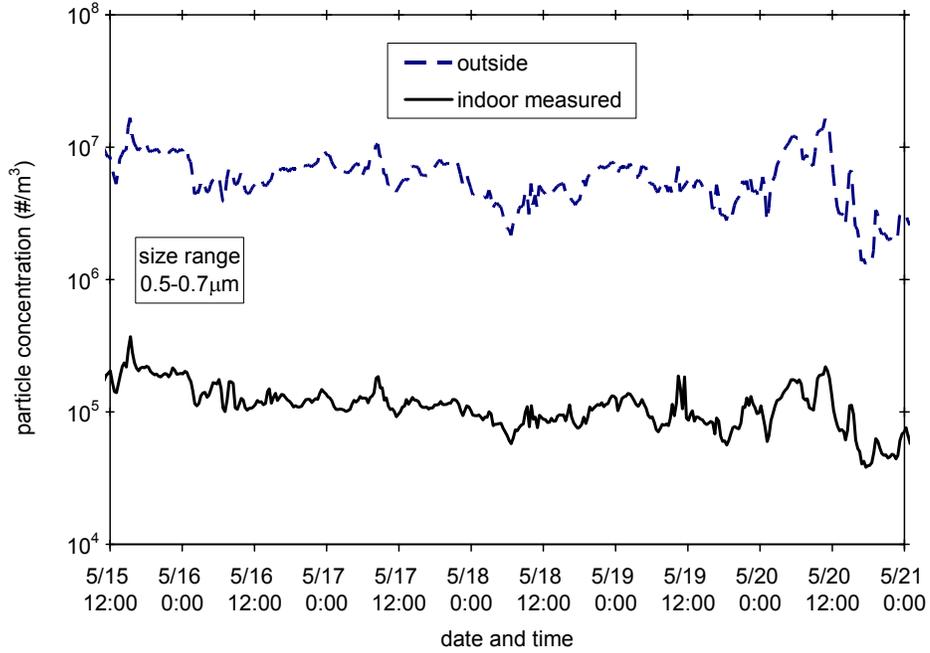


Figure 2A-6b. Time-dependent measured concentrations of particles 0.5-0.7 μm in diameter at the Oakland site during 15-21 May 2006.

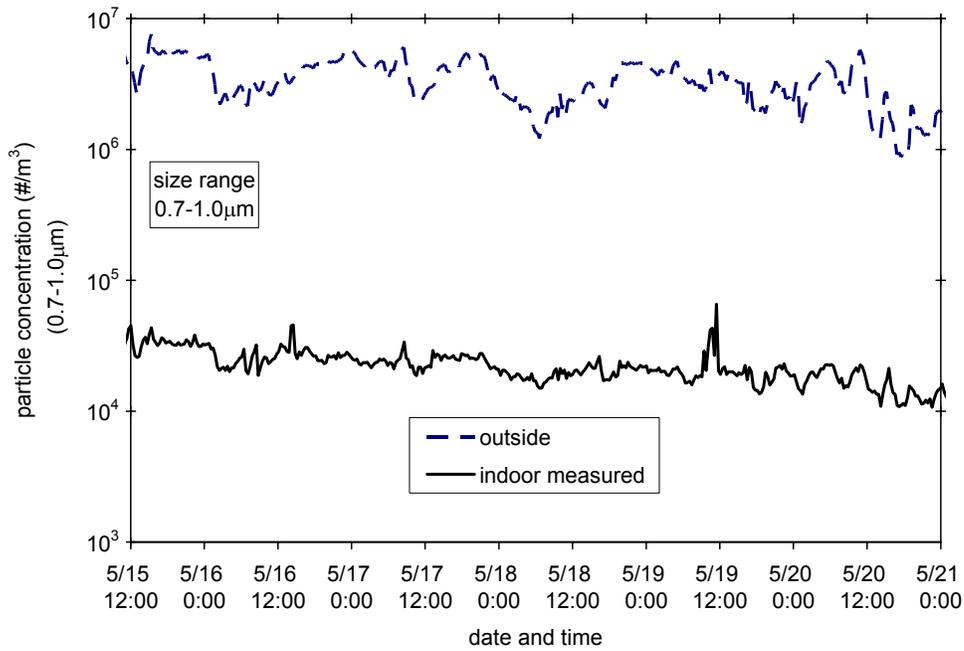


Figure 2A-6c. Time-dependent measured concentrations of particles 0.7-1.0 μm in diameter at the Oakland site during 15-21 May 2006.

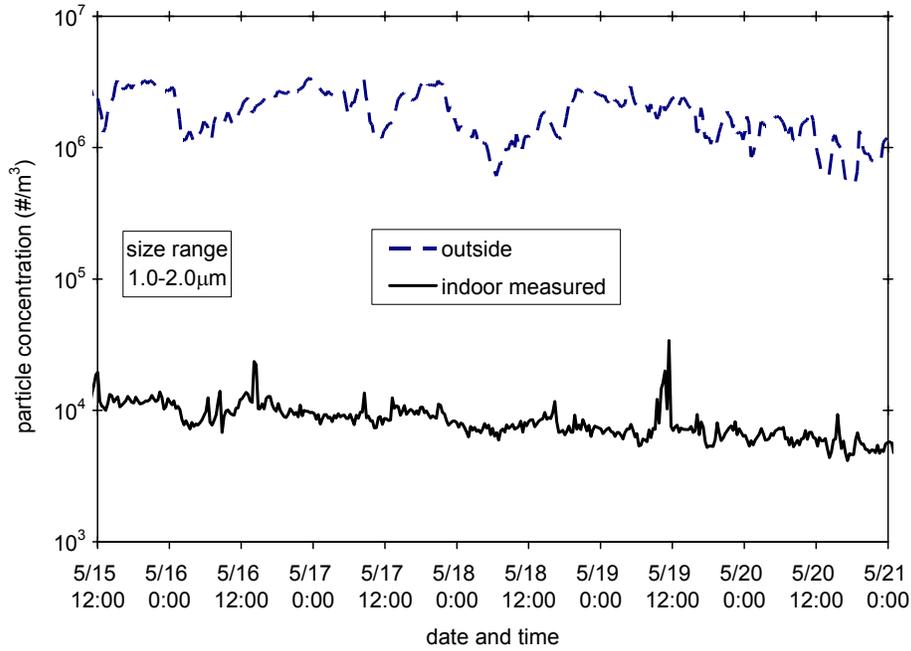


Figure 2A-6d. Time-dependent measured concentrations of particles 1.0-2.0 μm in diameter at the Oakland site during 15-21 May 2006.

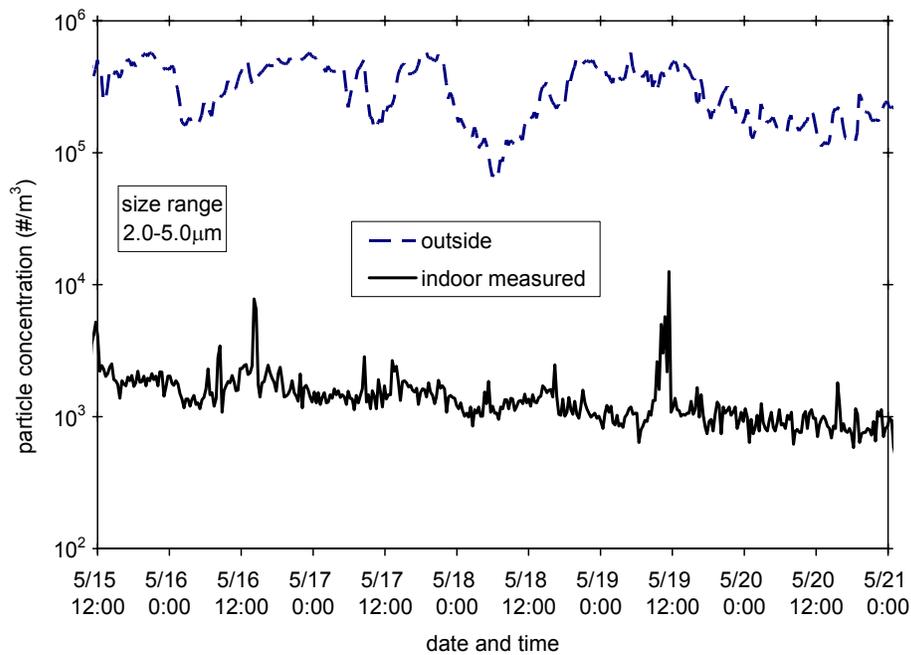


Figure 2A-6e. Time-dependent measured concentrations of particles 2.0-5.0 μm in diameter at the Oakland site during 15-21 May 2006.

Particle measurements were taken at a data center located in downtown San Francisco, CA on 8 - 10 August 2006. The data center zones are located throughout various floors within a high-rise building. The outside-air intake for the building is located on the eighth floor. A small percentage of outside air enters some of the data center rooms through ceiling vents, while other data center rooms receive no directly ducted outside air. The data center room where the indoor OPC was located receives no direct outside air and outside air can only enter this area via infiltration from adjacent zones. Air within the data center rooms is internally circulated and conditioned through CRAC units located on the data center floor and an underfloor plenum. Figure 2A-7 shows that the average outside particle concentration measured at the San Francisco site was approximately $5 \mu\text{g}/\text{m}^3$ with no measurements greater than $10 \mu\text{g}/\text{m}^3$ recorded during the monitoring period. The low measured outdoor particle concentration may be due to the significant height above ground level where the outside air intake for the building is located. While the outdoor concentrations at this location were among the lowest outdoor concentrations measured in this study, the indoor concentration was among the highest for data centers not using economizers, with an average concentration of about $1.0 \mu\text{g}/\text{m}^3$. Additionally, Figures 2A-8a-e show that fluctuations in the indoor concentrations matched fluctuations in the outdoor concentrations, indicating that possibly significant unintended air infiltration or filter bypass may be influencing conditions in the data center zone.

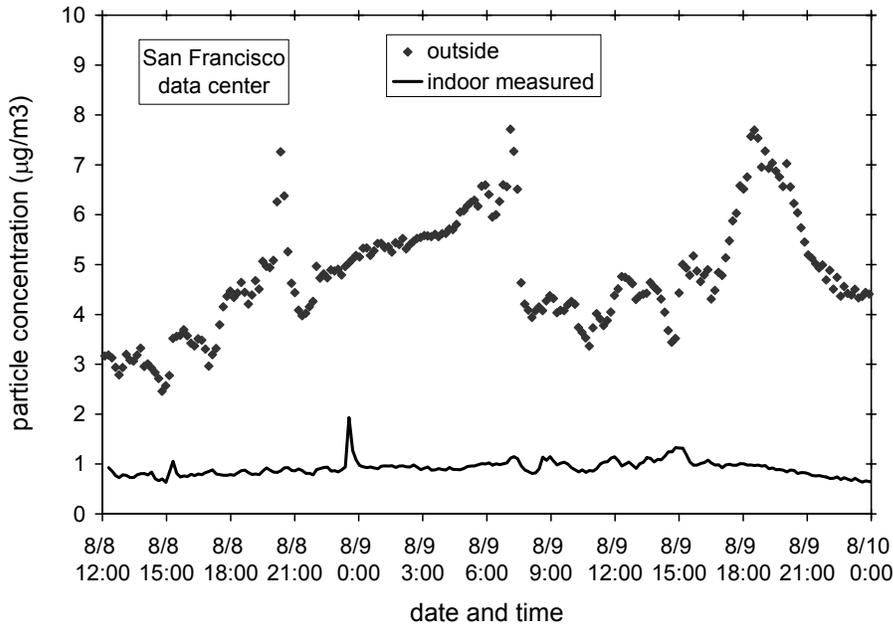


Figure 2A-7. Measured time-dependent particle mass concentrations at the San Francisco site during 8-10 August 2006. Concentrations represent particles 0.3-5.0 μm in diameter.

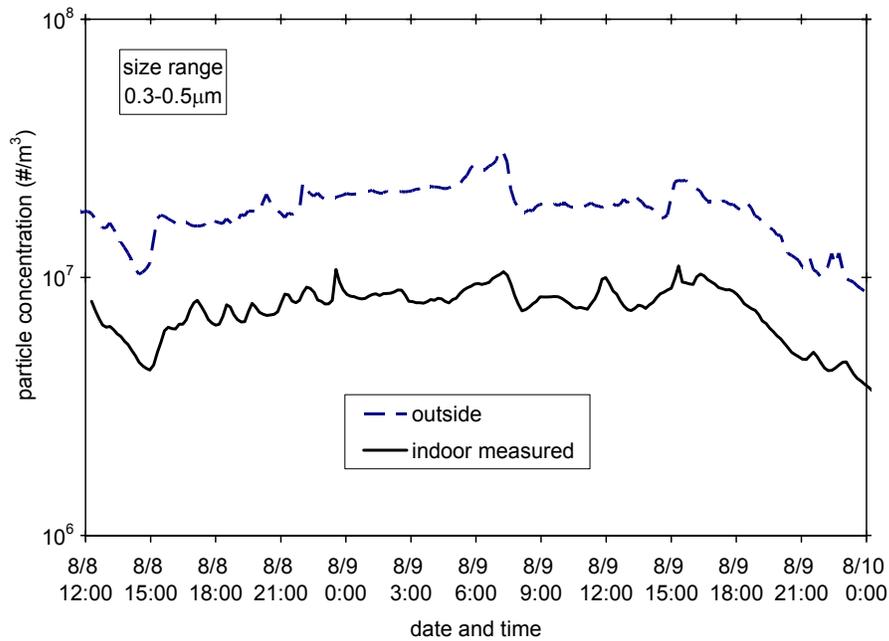


Figure 2A-8a. Time-dependent measured concentrations of particles 0.3-0.5 μm in diameter at the San Francisco site during 8-10 August 2006.

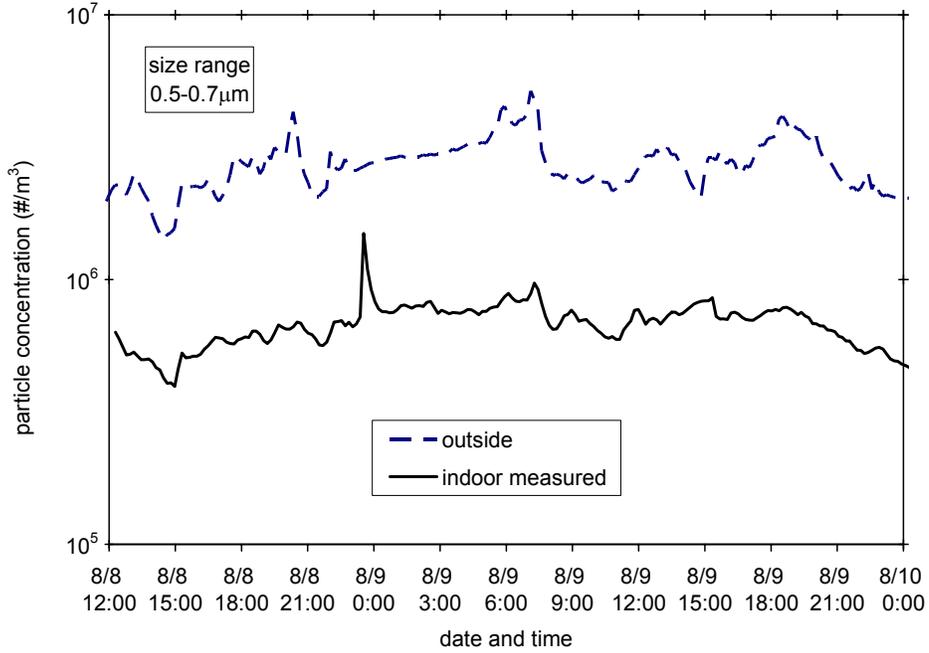


Figure 2A-8b. Time-dependent measured concentrations of particles 0.5-0.7 μm in diameter at the San Francisco site during 8-10 August 2006.

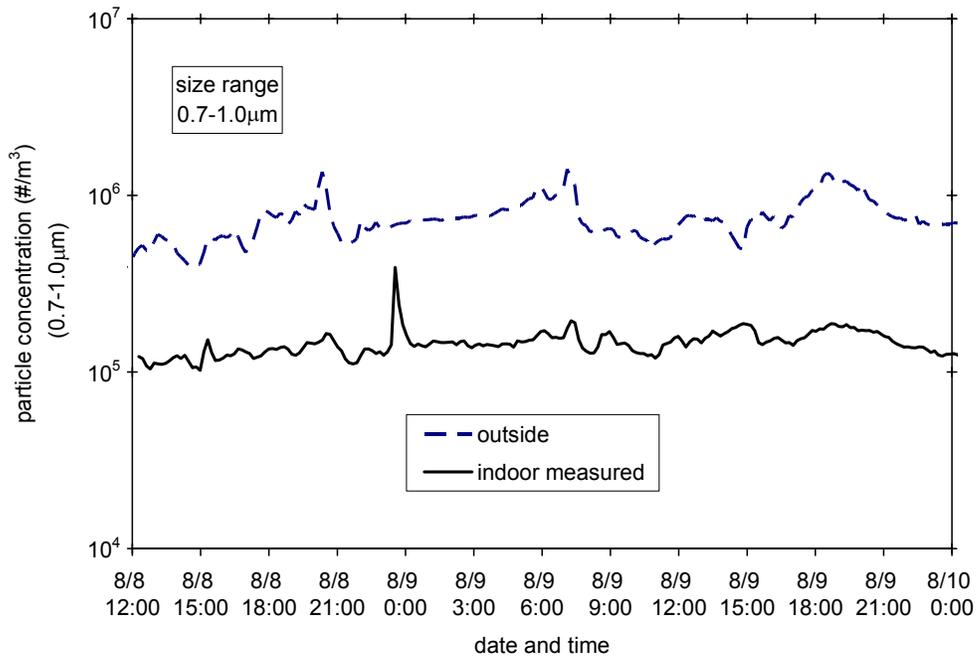


Figure 2A-8c. Time-dependent measured concentrations of particles 0.7-1.0 μm in diameter at the San Francisco site during 8-10 August 2006.

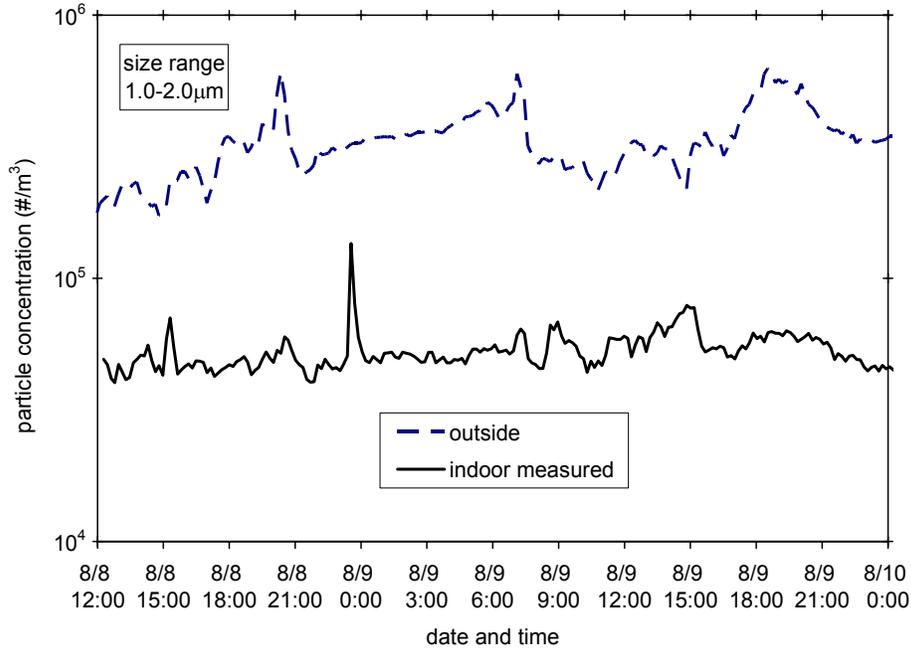


Figure 2A-8d. Time-dependent measured concentrations of particles 1.0-2.0 μm in diameter at the San Francisco site during 8-10 August 2006.

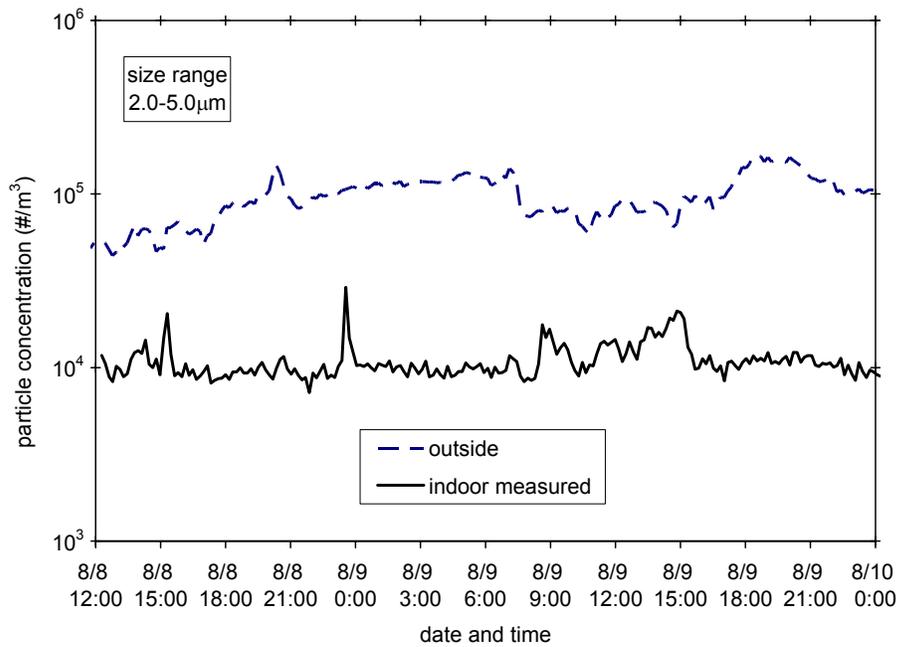


Figure 2A-8e. Time-dependent measured concentrations of particles 2.0-5.0 μm in diameter at the San Francisco site during 8-10 August 2006.

Particle measurements were taken at the Dublin site during 22 – 24 July 2006. The data center is located in a five-story office building. The area surrounding the building is undeveloped, but it is approximately one mile from a major freeway, placing the site in proximity with potential vehicular emissions. Air within the data center is conditioned by CRAC units and distributed through an underfloor plenum. During the July measurements, no direct outside air was ducted into the data center zone. The only source of outside air entering the data center came from infiltration through adjacent office zones. Figure 2A-9 shows that the average measured outside particle concentration was approximately $2 \mu\text{g}/\text{m}^3$. The two days of measurements were taken over a weekend period and the outdoor particle concentration drops significantly from Saturday, which averaged about $4 \mu\text{g}/\text{m}^3$ during the day, to Sunday when the particle concentration averaged below $1 \mu\text{g}/\text{m}^3$. Figures 2A-10a-e indicate that this drop in outdoor particle concentrations on Sunday is primarily attributable to a reduction in larger particles. The indoor concentrations appear to be generally independent of outdoor concentrations and no significant change in indoor particle concentration is seen in the data from Saturday to Sunday. Overall, the indoor particle concentration was fairly steady, with an average concentration of about $0.3 \mu\text{g}/\text{m}^3$ and rarely increasing above $0.5 \mu\text{g}/\text{m}^3$. Indoor particle concentration fluctuations are more exaggerated during the Saturday daytime hours. Possibly, fluctuations in the interior concentrations are caused by occupants in the data center. A retrofitted economizer/ventilation system was installed at the Dublin site during November 2006. This retrofitted system is designed to supply outside air to the data center through separate ductwork during cool-weather periods. Outside air is delivered from multiple locations along the ceiling of the data

center and provides, at maximum, an airflow equivalent to 25% of all the air delivered from the CRAC units. The delivered outside air passes through a conventional filter system in the air handling units before entering from the ceiling and mixing with the warm air in the room. Since the outside air is supplied from multiple ceiling locations, some of this air comes into immediate contact with the server equipment, while the rest of the outside air travels into the CRAC units. Air is removed from the data center zone through a large single exhaust vent positioned in the wall in the middle of the room. This retrofitted economizer system was designed to supplement the CRAC units in the data center, but the efficiency was restricted by the unconventional layout of air distribution. Rather than supplying directly to CRAC units or to an underfloor plenum, cool outside air mixes with the hot air exiting the servers, thereby reducing the efficiency of the system. Once the retrofitted economizer system was operating, particle measurements were taken again at the Dublin site from 22 November to 5 December 2006. The average outside particle concentration measured during this winter period was higher than the outside measurements taken during the summer weekend. Figure 2A-11 shows that the average outdoor measured particle concentration is $20 \mu\text{g}/\text{m}^3$, with many measurements approaching $70 \mu\text{g}/\text{m}^3$. Indoor particle concentrations at the server averaged about $3.7 \mu\text{g}/\text{m}^3$, with some measurements exceeding $15 \mu\text{g}/\text{m}^3$. Similar to the indoor particle levels at the Sunnyvale data center, the concentration changed significantly depending on the position of the economizer dampers. When the dampers were closed (no outside air supplied) the results are similar to the other data centers monitored that do not use economizers: the average particle concentration at the server was less than $1 \mu\text{g}/\text{m}^3$ and appeared to be independent of the fluctuations in outdoor particle concentration. With

the dampers opened so as to supply outside air, the concentration at the server quickly increased to approximately 5 to 15 $\mu\text{g}/\text{m}^3$ and appeared to be influenced by outdoor particle concentrations. For example, when the outdoor concentration was about 10 $\mu\text{g}/\text{m}^3$ the concentration at the server was about 2 $\mu\text{g}/\text{m}^3$, but when the outdoor concentration increased toward 70 $\mu\text{g}/\text{m}^3$ the concentration at the server increased to about 14 $\mu\text{g}/\text{m}^3$.

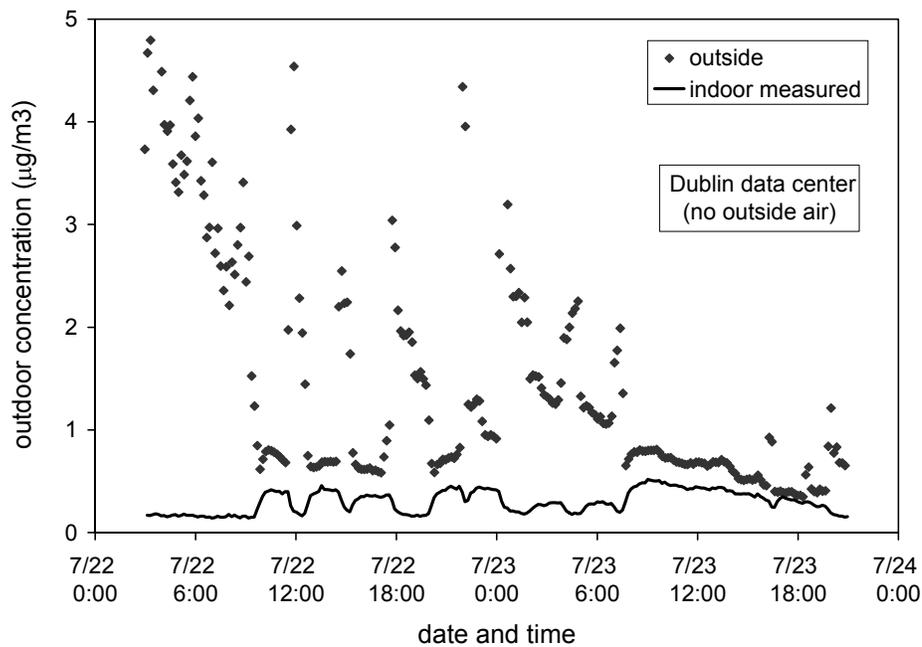


Figure 2A-9. Measured time-dependent particle mass concentrations at the Dublin site during 22-22 July 2006. Concentrations represent particles 0.3-5.0 μm in diameter.

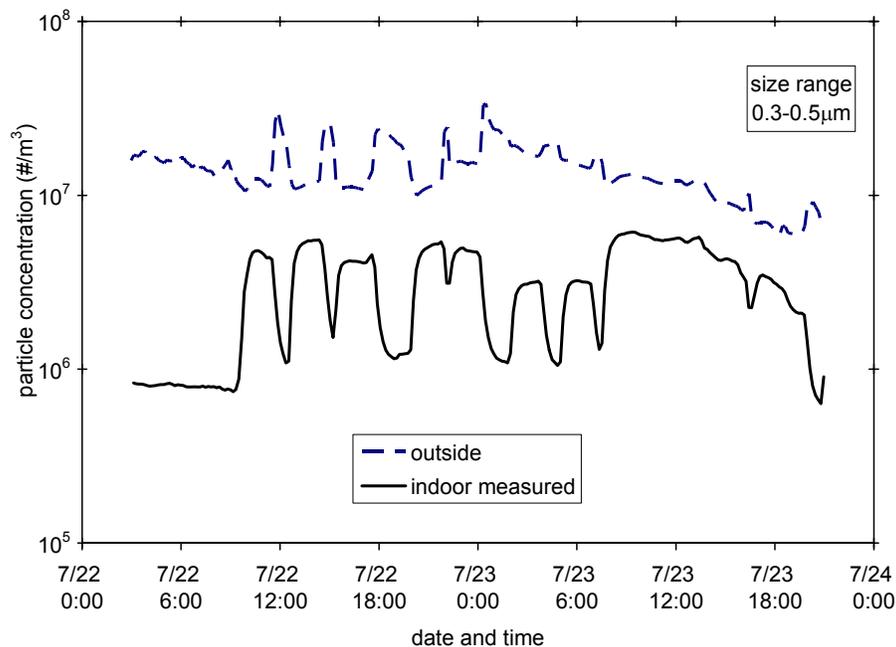


Figure 2A-10a. Time-dependent measured concentrations of particles 0.3-0.5 μm in diameter at the Dublin site during 22-22 July 2006

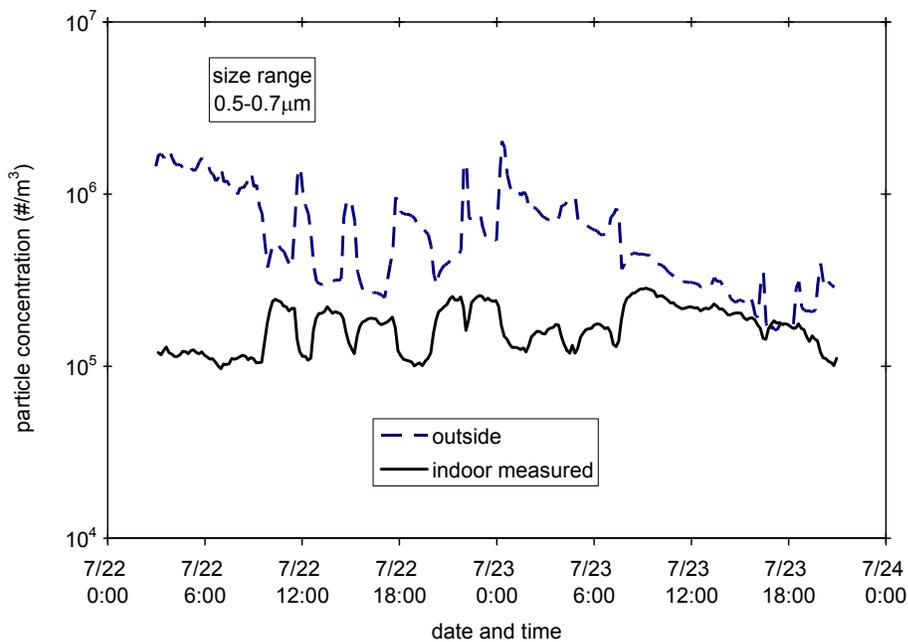


Figure 2A-10b. Time-dependent measured concentrations of particles 0.5-0.7 μm in diameter at the Dublin site during 22-22 July 2006

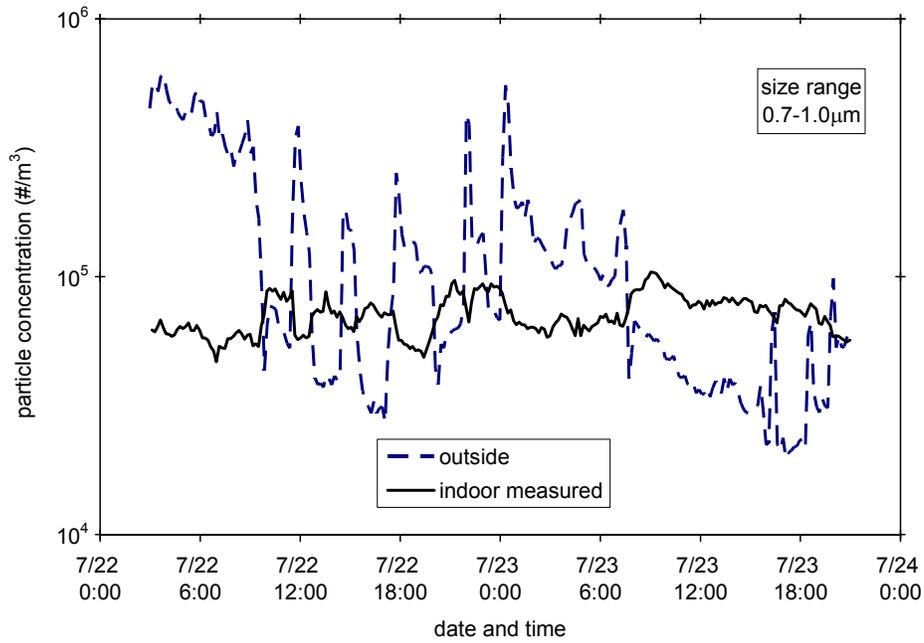


Figure 2A-10c. Time-dependent measured concentrations of particles 0.7-1.0 μm in diameter at the Dublin site during 22-22 July 2006

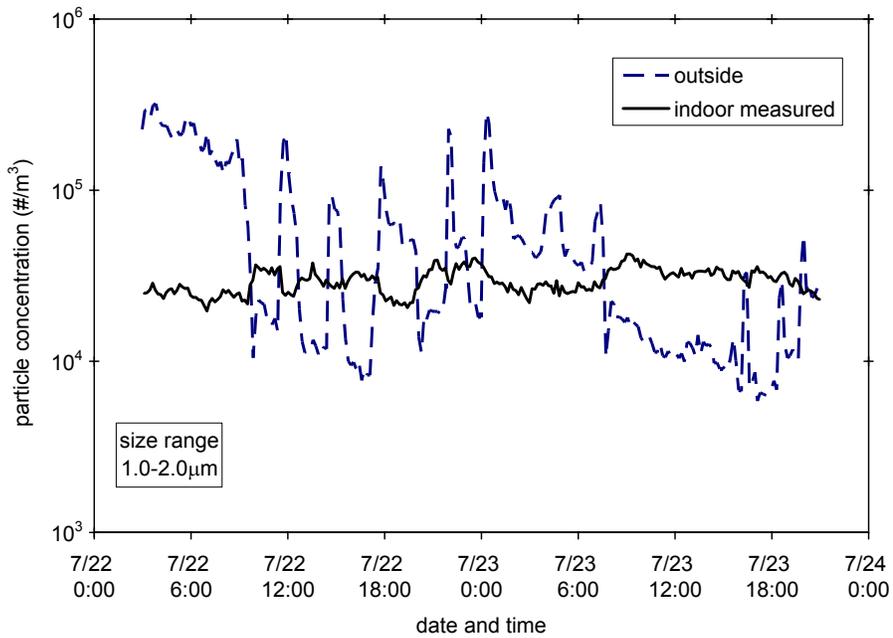


Figure 2A-10d. Time-dependent measured concentrations of particles 1.0-2.0 μm in diameter at the Dublin site during 22-22 July 2006

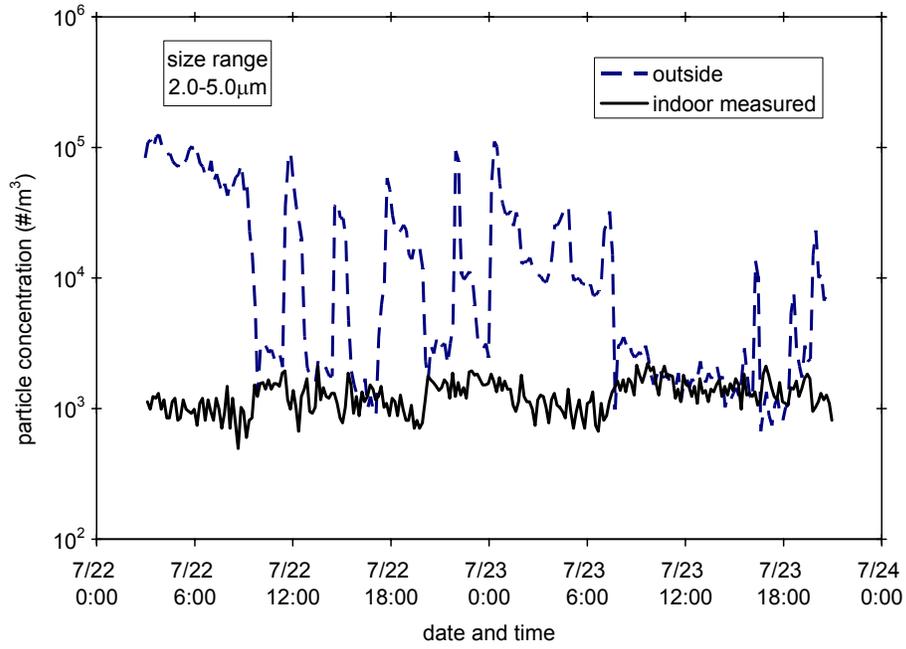


Figure 2A-10e. Time-dependent measured concentrations of particles 2.0-5.0 μm in diameter at the Dublin site during 22-22 July 2006

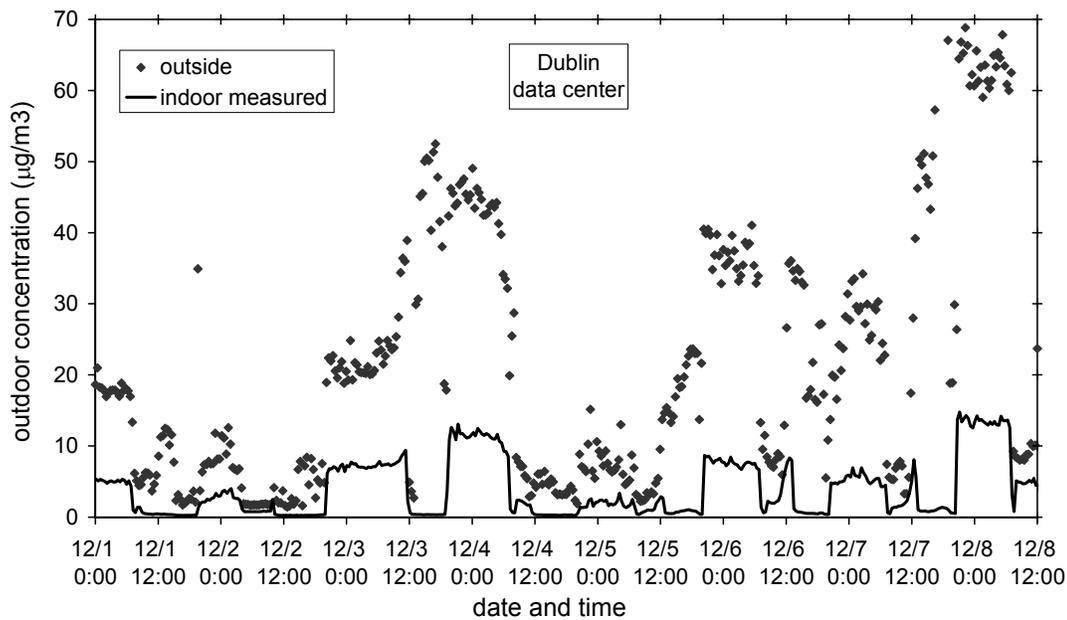


Figure 2A-11. Measured time-dependent particle mass concentrations at the Dublin site during 1-8 December 2006. Concentrations represent particles 0.3-5.0 μm in diameter.

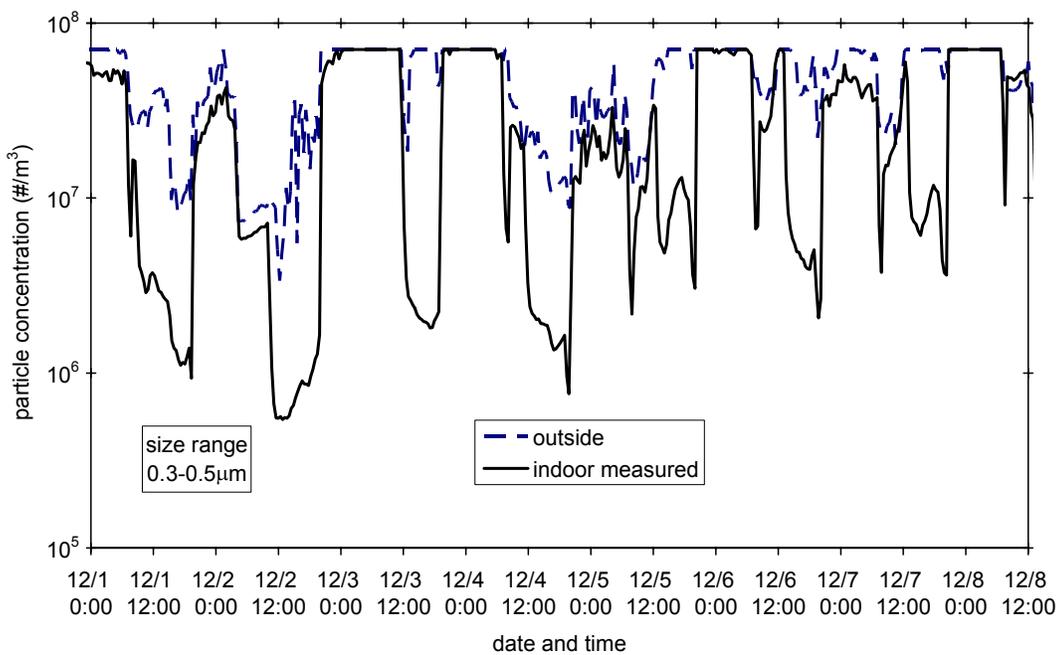


Figure 2A-12a. Time-dependent measured concentrations of particles 0.3-0.5 μm in diameter at the Dublin site during 1-8 December 2006. The horizontal line representing the portions of the indoor and

outdoor particle concentration indicates the OPC researched the particle counting limit and that the actual particle concentration is greater than what is presented here.

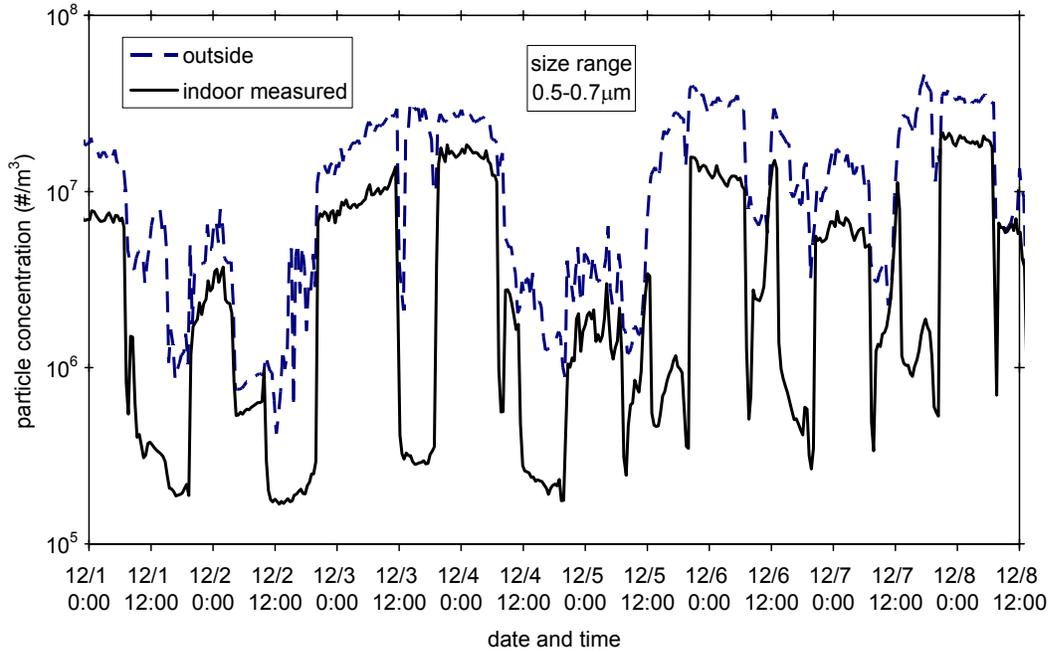


Figure 2A-12b. Time-dependent measured concentrations of particles 0.5-0.7 μm in diameter at the Dublin site during 1-8 December 2006.

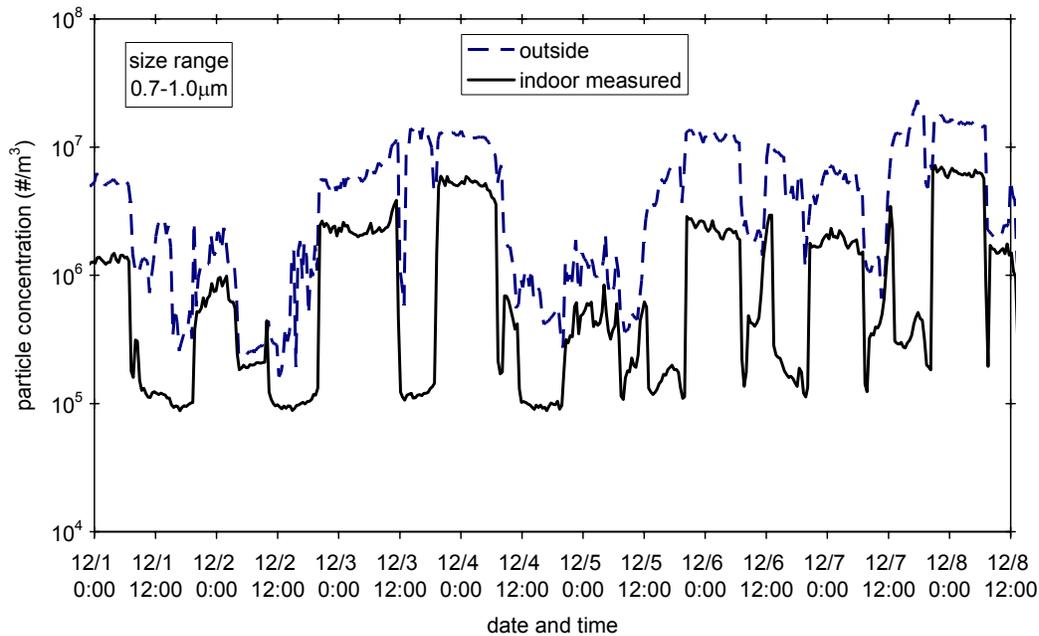


Figure 2A-12c. Time-dependent measured concentrations of particles 0.7-1.0 μm in diameter at the

Dublin site during 1-8 December 2006.

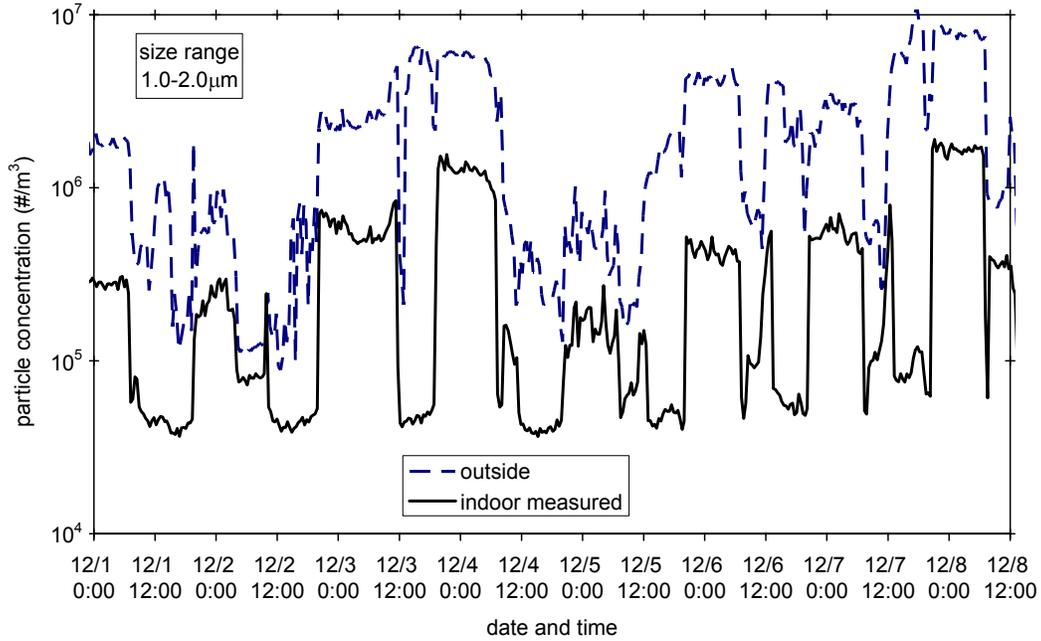


Figure 2A-12d. Time-dependent measured concentrations of particles 1.0-2.0 μm in diameter at the Dublin site during 1-8 December 2006.

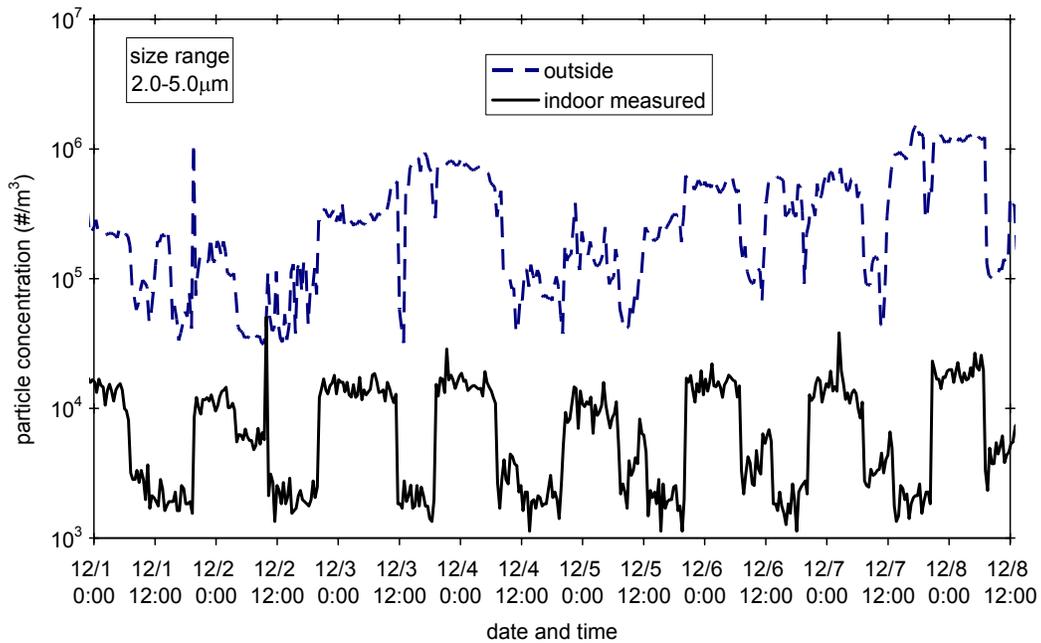


Figure 2A-12e. Time-dependent measured concentrations of particles 2.0-5.0 μm in diameter at the Dublin site during 1-8 December 2006.

The particle measurements from the five data centers presented in this Appendix support many of the findings observed from the three data centers evaluated in the main body of Chapter 2. Indoor particle concentrations measured at the data centers using conventional CRAC units (no economizer) were consistently below $1 \mu\text{g}/\text{m}^3$. These low indoor particle concentrations were observed even though the outdoor particle concentrations at these data centers varied by more than an order of magnitude, from a low of $2 \mu\text{g}/\text{m}^3$ measured during a weekend period at the Dublin site to a high of $27 \mu\text{g}/\text{m}^3$ at the Berkeley site that was heavily influenced by the mass proportion of larger particles. Similar to the Sunnyvale site, when large volumes of outside air were introduced into the Dublin data center the indoor particle concentration increased by an order of magnitude to nearly $4 \mu\text{g}/\text{m}^3$. These results provide additional empirical evidence that IT equipment in conventionally operated data centers are exposed to low particle concentrations ($<1 \mu\text{g}/\text{m}^3$). The results also indicate that economizer use will increase indoor particle concentrations, though even with these increases, the particle concentrations still appear to be below most IT equipment air quality guidelines.

Chapter 3: Combining economizers with improved filtration to save energy and protect equipment in data centers

This chapter presents concentrations and chemical analysis of particles measured in a data center under different filter configurations. The challenges involved with allaying data center particle concerns are discussed and the potential for high-quality filtration to circumvent particle problems with economizers is described. Particle concentration and speciation results are presented with meter electricity values from the data center. The chapter concludes that economizer use combined with improved particle filtration can reduce data center power demand while simultaneously maintaining indoor particle levels similar to those using conventional (non-economizer) ventilation system design and operation. Reproduced in part with permission from Building and Environment, 45, 718-726, 2010. Copyright 2009, Elsevier Inc. License Number: 2314400662208

3.1. Introduction

Economizers provide large amounts of outside air for cooling internal heat loads during favorable weather conditions, thereby reducing the air-conditioning energy demands in buildings. Economizer use is standard practice for many building types and even required by building codes in some regions (CEC, 2005). The use of economizers in data centers, however, has been hindered by perceived potential equipment reliability concerns associated with exposing IT equipment to outdoor particulate matter (PM) (Tschudi et al., 2004). Apprehension towards economizer implementation is exacerbated by the lack of an established level of particulate matter exposure considered safe for data center equipment. Results presented in Chapter 2 show that economizer use increases particle concentrations relative to conventional non-economizer data center cooling design. Although the higher concentrations observed are still below most air quality guidelines, particle concentration guidelines for data centers vary widely (ASHRAE,

2009), which reduces confidence that abiding by any particular standard will protect IT equipment from contamination damage. Any increase in particle concentration relative to conditions occurring under conventional practice can dissuade economizer implementation, especially given the strong emphasis in this sector on performance reliability. Establishing particle concentrations associated with standard practice creates an engineering opportunity: if data centers with economizers can operate with particle concentrations at or below those measured in conventional non-economizer data centers, then the concern surrounding increases in particle induced IT-equipment damage can be removed from the decision-making process regarding economizer implementation. The previous chapter presented data showing low indoor particle concentration levels ($<1 \mu\text{g}/\text{m}^3$) in conventional non-economizer data centers. Such a low indoor particle concentration may later prove to be more stringent than needed. The level is well below any published data center air quality guidelines (ASHRAE, 2009); however, this level establishes a particle limit during economizer operation that, by definition, creates no greater risk to IT-equipment reliability than conventional non-economizer data centers.

This chapter explores the feasibility of using economizers in data centers to save energy while simultaneously controlling indoor particle concentrations with improved air filtration. Time- and size-resolved particle concentration data were gathered at an operating northern California data center while using air filtration of varying levels of efficiency. Along with measuring the size- and time-resolved indoor and outdoor particle concentrations, chemical analyses of particulate sulfate, nitrate, chloride, and carbon were performed under each filter and economizer configuration. Metered electricity data were also gathered to compare overall operational energy use for each condition. Fan power

requirements were calculated, using the mechanical design specifications of the data center, to disaggregate the potential increase in fan energy associated with improved filtration from the energy use associated with compressor-based cooling. The measured particle characteristics and energy use are evaluated to explore the potential energy savings associated with a shift from conventional cooling and filtration practices in data centers to a system with economizer use, to save energy, combined with improved filtration, to ensure protection of equipment from particle contamination.

3.2. Methods

3.2.1. Experimental Design

Particle concentrations were measured at a data center in Sunnyvale, CA during August 2008. Real-time measurements were made and two- or three-day integrated filter-based samples were collected during 8-29 August for subsequent analysis. The mechanical system at this data center is designed with an economizer, allowing the amount of outside air entering the zone to be adjusted depending on outside temperature and humidity conditions. The amount of outside air entering the data center is controlled by an energy management and control system (EMCS). This data center is the same Sunnyvale site discussed in Chapter 2, which presents details of the building layout and mechanical design. Briefly summarizing, economizer dampers in the air-handling units (AHU) modulate the ratio of outside air to return air that enters the AHU. Once in the AHU, this blend of outside and return air passes across a bank of filters, is thermally conditioned, and is then ducted to the data center zone for the purpose of removing heat generated by the racks of operating IT equipment. During the monitoring period of this study, the economizer system was manually controlled to be in an “economizer-off”

mode from noon to 6:00 PM. In this condition, outside air accounts for only about 1% of the air entering the AHU with the rest being return air from the data center. During the remaining 18 hours of each day, the economizer system was set to an “economizer-on” mode, where all return air was exhausted from the building and the data center was cooled with 100% outside air, augmented as necessary with compressor-based air conditioning. When in economizer-on mode, the data center air-exchange rate reaches nearly 50 per hour. The period of economizer activity was chosen to match typical temperature- and humidity-controlled economizer use in data centers during summer months at this northern California location.

Heating ventilation and air-conditioning (HVAC) filters with three different minimum efficiency reporting value (MERV) ratings were installed during the monitoring period. A higher MERV rating indicates greater particle-removal efficiency by the filter (ASHRAE, 1999). Immediately before monitoring began, new HVAC filters with a rating of MERV 7 were installed. MERV 7 filters are commonly used at this data center and previous studies have reported that the use of MERV 7 filters is consistent with normal industry practice (Brown et al., 2007). During other phases of the monitoring period, the MERV 7 filters were removed and replaced with more efficient MERV 14 filters, and later the MERV 14 filters were replaced with MERV 11 filters. All the filters are from a single manufacturer and represent commonly sold filter models within their respective efficiency category. Each filter type remained in place for approximately one week of operation.

3.2.2 Real-Time Measurements

Size resolved particle concentrations were measured using two Met-One 237B

optical particle counters (OPC). These particle counters are capable of detecting particles within the range 0.3-5.0 μm optical diameter and categorizing the particle counts into different size bins, with an uncertainty of $\pm 20\%$ in particle counts for each size bin. The fine particle mass concentration was estimated from these particle counts assuming a particle density of 1.5 g/cm^3 (Pitz et al., 2003) and using the method described in Chapter 2. The indoor particle counter was placed on top of the server rack, with the intake tube sampling from a position in front of the servers. The outdoor particle counter was placed within the air handling unit (AHU) at the outside air intake, prior to any filtration, to monitor the outdoor particle concentration entering the AHUs that serve the data center. Measurements were taken for three-minute periods at 10-minute intervals. Each OPC would draw air at a rate of 2.8 L/min for three minutes and then pause for seven minutes before beginning the next particle-counting cycle. The two OPCs used in this study were tested prior to the monitoring period to ensure that each particle counter produced similar results (less than 10% difference) under the same conditions.

Black carbon was measured using a Magee Scientific AE22 aethalometer. The chemical structure of black carbon, or soot, results in high electrical conductivity (Andreae and Gelencser, 2006) and has been associated with higher than usual failure rates of electronic equipment (Morawska et al., 2009). The aethalometer, which measures light beam attenuation by particles collected on a quartz filter tape with an uncertainty of $\pm 10\%$, was programmed to calculate black carbon concentrations in one-minute intervals. The aethalometer was placed on the data center floor, with two lengths of 12.7 mm diameter \times 8 m copper tubing to collect air from both inside and outside of the data center. Sample air traveling through the copper tubing was first drawn through a

cyclone with a PM_{2.5} cutoff using an external pump at a flow rate of 25 l/min. A portion of this sample air then entered the aethalometer (4.5 l/min) while the rest was discarded.² Automated 2-way solenoid valves were used upstream of the aethalometer to switch between drawing air from the inside and outside tubing in 20-minute intervals. For each 20-minute interval, the average concentration from only the last 10 minutes was used to represent that period. Measurements from the first 10 minutes were excluded to ensure that sudden changes in relative humidity, which could occur when switching sample air locations, would not disturb the aethalometer measurements (LaRosa et al., 2002; Wallace, 2005).

3.2.3. Filter-Based Particle Measurements

Chemical constituents of indoor and outdoor particles were measured using two sets of sample filters, one for indoor and one for outdoor measurements. Each filter set consisted of four filter holders: two for teflon (Teflo, 25 mm diameter, 3.0 µm pore size, Pall p/n R2PI025) and two for quartz filters (Pall 2500QT). The teflon filters were mounted in Savillex (teflon) filter-holders downstream of honeycomb denuders (described below). The holders for the quartz filters were stainless steel with each holding two quartz filters in series. One sample filter set was placed on the data center floor and the other set was located in an adjacent auxiliary room with the inlet extending outdoors to a position in front of the AHU outside air intake. Figure 3-1 shows a schematic of the airflow through the sample filter sets. For each filter set, air was drawn through 12.7 mm

² Originally the experimental apparatus was designed to provide indoor and outdoor sample air for four different measuring devices: the Magee Scientific AE22 aethalometer, as well as a TSI model 3321 aerodynamic particle sizer (APS), a TSI model 3022A condensation particle counter (CPC), and a PMS Lasair 1003 OPC. No data was gathered from the latter three devices. The APS was not available for use during the monitoring period, while the CPC and Lasair OPC experienced irreparable mechanical failures immediately prior to the monitoring period.

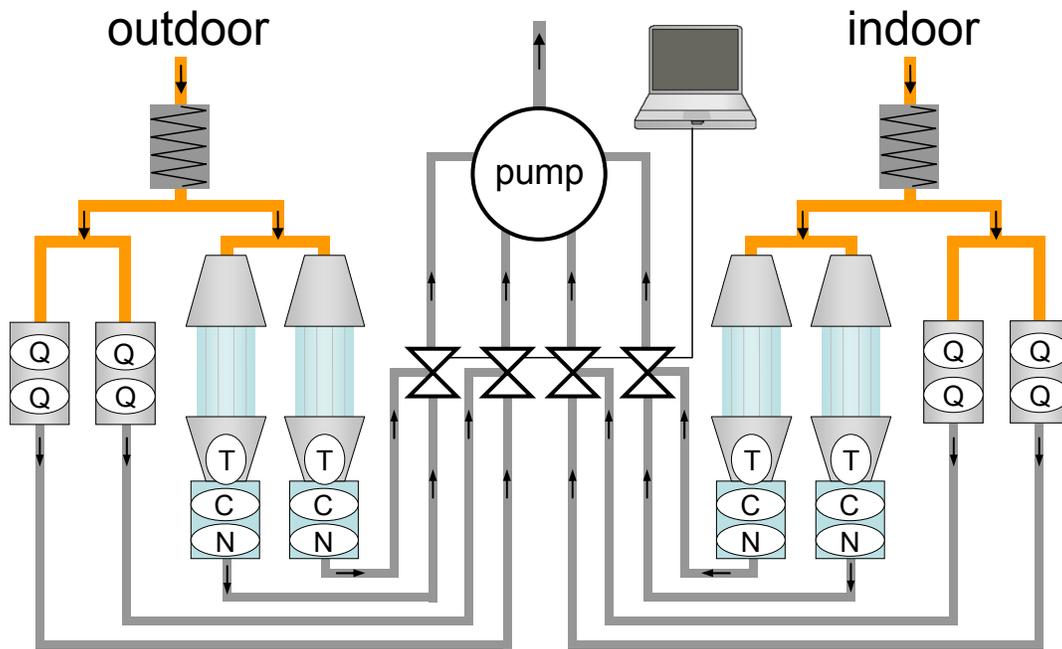


Figure 3-1. Indoor and outdoor sets of sample filters used for particle collection. For each filter set, a pump is used to draw air into copper tubing and through a cyclone before branching to four filter holders: two for teflon/cellulose/nylon filters and two for quartz filters. Teflon, cellulose, and nylon filters (labeled T, C, and N) are mounted in filter holders downstream of glass honeycomb denuders. Pairs of quartz filters (labeled Q) are held in stainless steel filter holders. Solenoid valves are programmed to direct air flow to one pair of filter holders (one with teflon/cellulose/nylon filters in series and one with two quartz filters in series) during the economizer-off periods, and then switch to the other matched pair of identical filters during the economizer-on periods.

diameter \times 2 m of copper tubing. The air passed through a cyclone with a PM_{2.5} cutoff at a flowrate of 25 l/min before branching to the four filter holders. Automated 3-way solenoid valves located downstream of the filters directed the air flow equally to one pair of filter holders (one teflon and one with two quartz filters in series) during the economizer-off periods (six hours per day, noon - 6:00 PM), and then switched to the other matched pair of identical filters during the economizer-on periods (18 hours per day, 6:00 PM - noon). Filters were collected and replaced at 2-3 day intervals. Additional teflon, nylon, and cellulose filters were transported to and stored at the monitoring site during the study period, but not used in the experiment. These filters were used to provide for blank correction during filter analysis. The split of sample airflow between teflon and quartz filters was controlled with needle valves and the airflow rates were measured for each filter set using a BIOS DryCal DC-2 calibrator. Swagelok openings were added to the end of the 2 m copper tubing of each filter set and also placed after the split between the teflon and quartz streams, allowing the calibrator to be connected at these locations to measure airflow. After each filter replacement, the calibrator would be connected after the split and the airflow would be adjusted with the needle valves to measure equal flow through both streams (one needle valve per stream). Airflow entering the 2 m of copper tubing (before the split) was then measured to confirm that 25 l/min of air was entering the filter set. Since air resistance could be different at each location where the calibrator was connected, the measurement process was iterative and continued until the calibrator measured 25 l/min entering the filter sample set and the flow measured after the split was equal in each branch.

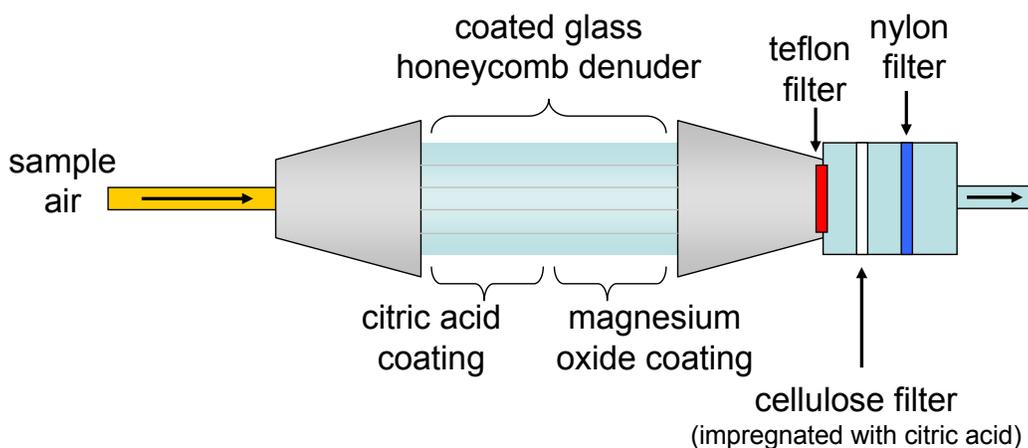


Figure 3-2. Denuder filter system for chemical speciation. Entering sample air passes through a glass honeycomb denuder, coated with citric acid on the upstream end to remove gas-phase ammonia, and then coated with magnesium oxide on the downstream end to remove nitric and hydrochloric acid from the airstream. A 47-mm cellulose filter and a 47-mm nylon filter were placed in series immediately downstream of each teflon filter to account for the volatilization of collected ammonium nitrate and ammonium chloride particles. The cellulose filter (impregnated with citric acid) collects ammonia while the slightly alkaline nature of the nylon filter collects nitric acid and hydrochloric acid.

Particulate matter was sampled on teflon filters to determine sulfate, nitrate, and chloride concentrations, since these particles present a potential concern for data centers owing to the ability of PM containing water-soluble ionic salts to deliquesce at high relative humidity and thereby conductively bridge isolated elements on circuit boards (Weschler, 1991). Sulfate has been used previously to demonstrate current leakage attributable to particle deposition under conditions of high particle concentration and high humidity (Litvak et al., 2000). Each teflon filter used to collect sulfate, nitrate, and chloride was preceded by a glass honeycomb denuder (see Figure 3-2). The denuder was coated in citric acid on one end and magnesium oxide on the other using a protocol adapted from an EPA method (Chow and Watson, 1998; Lunden et al., 2003); the purpose was to remove gas phase ammonia, nitric acid, and hydrochloric acid from the airstream before PM collection. A 47-mm cellulose filter and a 47-mm nylon filter were placed in series immediately downstream of each teflon filter to account for the volatilization of collected ammonium nitrate and ammonium chloride particles. The cellulose filters were impregnated with 70 mg of citric acid to collect ammonia while the slightly alkaline nature of nylon filters was used to collect nitric acid and hydrochloric acid.

The teflon and nylon filters were analyzed for anions using a protocol from IBM labs with an estimated $\pm 5\%$ measurement uncertainty (Christensen, 1996). Each teflon filter was extracted by first pipetting 200 μl of ethanol onto the filter surface (to overcome the hydrophobicity of teflon) and then placing the filter into a precleaned sealable Kapak™ plastic bag containing 5.8 ml of deionized water. The bags were heat-sealed and placed into an oven maintained at 60 °C for one hour. The bags were then sonicated before transferring the contents to an autosampler vial. The nylon filters were

extracted similarly, except that the extraction solvent was the eluent used for ion chromatography (2.7 mM sodium carbonate and 0.3 mM sodium bicarbonate). The extracts were analyzed by ion chromatography using a Dionex 2020 Ion Chromatograph utilizing an AS12A-series separatory and standard 4 mm guard column.

The citric-acid impregnated cellulose filters were extracted in deionized water with sonication and then analyzed for ammonium using an Orion model 95-12 ammonia gas sensing electrode. The electrode, in combination with a voltmeter, measures the increase in voltage associated with an exponential decrease in ammonium – as nitrogen – in solution. Standard concentrations of 200 μg , 100 μg , 50 μg , and 10 μg ammonium solutions were used to calibrate the relationship between measured voltage and ammonium concentrations each day before evaluating the unknown solutions extracted from the cellulose filters. Each standard also included 70 mg of citric acid as a precautionary measure to account for the approximately 70 mg of citric acid impregnated on each cellulose filter, which would end up in the unknown solution after filter extraction. Since the relationship between voltage and ammonium concentration is linear on a semi-log plot, a linear regression equation was then used to determine the nitrogen concentrations in the unknown solutions by individually measuring the voltage of that solution. A calibration check was performed each day at the end of analysis to insure that no significant shift in the calibration had occurred. This point was included in the fit for the calibration curve.

Mass concentrations of $\text{PM}_{2.5}$ were gravimetrically measured on the teflon filters using a Sartorius SE-2 microbalance, with the filters being placed in a temperature and humidification equilibration chamber at 35-40% RH for 24 h prior to weighing, pre- and

post-loading. However, inductively coupled plasma (ICP) analysis confirmed the visual appearance of minute brass filings on some of the filters after sampling. These filing may have been caused by the slight stripping of the Swagelok threading, located immediately upstream of the sample filters, during the repeated opening and closing associated with the filter replacement and airflow recalibration. Consequently, filter-based PM_{2.5} mass concentration measurements were deemed unreliable and are reported here (Table 3-10 and Figure 3-13) only to document these results.

Particulate matter was collected on quartz filters to determine particle carbon levels. Condensed-phase organic pollutants may be a concern as they can contribute to the failure of data center equipment by promoting arcing between relay contacts, which results in increased contact erosion, or by forming polymeric films that can increase electrical resistance on contact surfaces (Shields and Weschler, 1998). The carbon content of the particles collected on quartz filters was determined using thermal optical analysis (TOA), as described in Kirchstetter and Novakov (2007). The TOA method has an estimated $\pm 5\%$ measurement uncertainty. Filters were heated at a constant rate of 40 °C/min from 50 to 700 °C in a pure oxygen atmosphere. The carbon evolved from the filter was passed over a platinum-coated ceramic catalyst in an oxygenated atmosphere at 800 °C, causing it to fully oxidize to CO₂. The resulting CO₂ was measured with a nondispersive infrared analyzer over the entire temperature range. The intensity of light transmitted through the sample was continuously monitored during analysis using a spectrometer to determine when the light-absorbing carbon evolved from the sample. The majority of organic carbon is expected to evolve from the quartz filter at a lower temperature than black carbon. In this study, all particulate carbon was classified as

either organic or black, and carbon that evolved at temperatures above 400 °C was classified as black.

Each quartz filter sample consisted of two 25-mm diameter quartz filters in series. This tandem method is used to correct positive sampling artifacts caused by organic gases that sorb to the quartz filter matrix (Turpin et al., 1994; Kirchstetter et al., 2001), since both carbonaceous particles and sorbed organic gases can contribute to the measured carbon during TOA. While only the upstream filter collects particles, organic gases are assumed to sorb to saturation on both upstream and downstream filters. The amounts of organic gases collected on the two filters in series are assumed to be similar; hence, the difference in measured CO₂ between the filters is assigned to carbon generated from particles. Prior to use, all of the quartz filters were baked at 800 °C for 6 h to remove any carbonaceous impurities.

3.2.4 Energy Calculations

A combination of measured and theoretical power consumption calculations was used to estimate the energy loads for the three HVAC filter types and two ventilation conditions evaluated in this study. Output from the main electricity meter monitoring the entire data center building was continuously gathered to observe fluctuations in overall power demand. Separate sub-metering was performed for all electricity leaving the uninterruptible power supply (UPS) units, which represents electricity exclusively used to operate the IT equipment in the data center, such as server, network, and storage devices. The difference between these two metered electricity values represents the power demand for all non-IT equipment in the data center, which primarily consists of the HVAC chiller and fans, as well as UPS losses, lighting, and auxiliary uses (Brown et al., 2007). Power

use data were collected in five-minute increments from both the main facility meter and the UPS sub-meter. Increase in fan power associated with improved supply air filtration was estimated using fan curves specific to the HVAC design at the data center (Figure 3-3) and standard fan laws (ASHRAE, 2005):

$$P_i = \frac{Q_s p_s}{\eta_m \eta_f \eta_d \eta_b} \quad (3-1)$$

In equation (1), P_i is the input power to the supply fan, Q_s is the volumetric supply airflow rate, and p_s is the fan static pressure. The parameters η_m , η_f , η_d , and η_b represent efficiencies for the motor, fan, variable frequency drive (VFD), and fan belts respectively. Input power to the fan for the baseline (MERV 7) filtration case was calculated from the building fan curves assuming a motor efficiency of 0.9, which is representative of large commercial HVAC systems (Fisk et al., 2002). The fan efficiency was calculated from the building fan curves and the VFD efficiency was estimated at 0.95 based on manufacturer specifications. The data center fans are direct drive, corresponding to a fan belt efficiency of 1. Figure 3-3 presents the design fan curve for one of 16 identical fans used to supply air at the Sunnyvale data center. According to the building fan curves the design brake horsepower (BHP) is 3.45 kW (4.62 hp). Since BHP incorporates the fan curve efficiency, this corresponds to an input power of 4.03 kW after accounting for efficiencies. Assuming that the system efficiencies are constant under small changes in pressure, the input fan power is directly proportional to the fan static pressure.

$$P_i \propto p_s \quad (3-2)$$

The increase in static pressure caused by more efficient HVAC filters was then used to

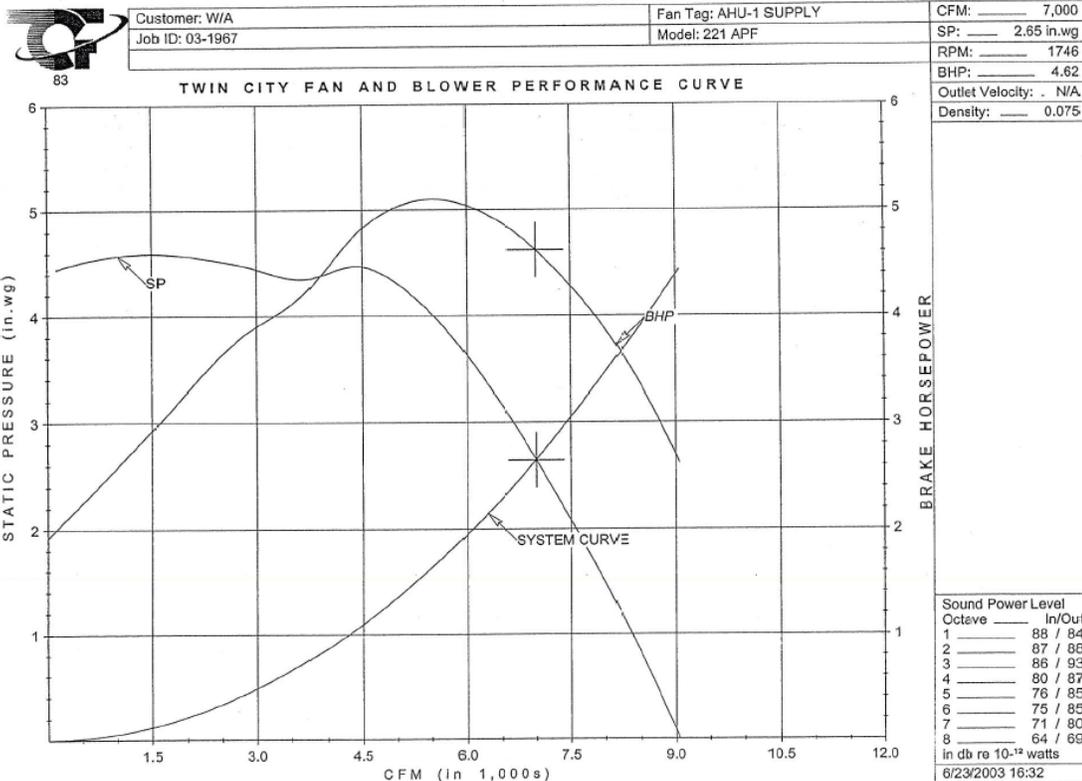


Figure 3-3. Supply fan curve for one of 16 identical fans used to supply air at the Sunnyvale data center. The fan curve indicates a design brake horsepower (BHP) of 3.45 kW (4.62 hp), which corresponds to an input power of 4.03 kW after accounting for the motor, drive, and VFD efficiencies. The fan curve also indicates a design static pressure of 660 Pa (2.65 in. wg), which accounts for a typical MERV 7 filter resistance. Increase in fan power associated with improved filtration is calculated by estimating the increase in static pressure caused by the increase in filter initial resistance.

estimate the corresponding increase in fan power associated with that filter type. Static pressure increase was estimated by adding the differences in the manufacturer's published initial resistance values for each filter type used in this study to the static pressure indicated by the building fan curves. The initial resistance values correspond to 12.5 Pa, 19.9 Pa, and 39.9 Pa for MERV 7, 11, and 14, respectively (Airguard, 2009), given that the dimensions of the AHUs at the Sunnyvale data center approximates a filter face velocity of 122 m/min (400 ft/min). Since the building fan curve indicates a design static pressure of 660 Pa (2.65 in. wg), which accounts for a typical MERV 7 filter resistance, the static pressure with MERV 11 and 14 filters can be estimated as 668 Pa and 689 Pa, respectively.

3.3. Results and Discussion

3.3.1. Particle Concentrations

Table 3-1 presents time-averaged particle mass concentrations, in the size range 0.3-2.0 μm diameter, as measured with the OPCs at the Sunnyvale data center operating with different HVAC filter types. For each measurement period, both the indoor and outdoor mass concentrations are averaged separately for the hours with 100% outside air ventilation (economizer-on) and 1% outside air ventilation (economizer-off). Previous analysis, presented in Chapter 2, has shown that fine particle mass concentrations measured in data centers are primarily of outdoor origin. This finding allows the indoor/outdoor (I/O) ratio of particles to be interpreted as an indicator of the performance of the HVAC system in protecting the indoor air in the data center from particles of outdoor origin. The I/O ratios for particle mass concentrations based on the OPC measurements during economizer-on and economizer-off periods for the three HVAC

filter types are presented in Figure 3-4. Two key qualitative findings are as expected: (a) the I/O ratio is higher when the economizer is *on* compared to when it is *off*, and (b) increased filter MERV ratings result in reduced I/O ratios. A key quantitative result illustrated in this figure is that the I/O ratio for the MERV 14 filters when the economizer is *on* is similar to the I/O ratio for the conventional MERV 7 filters when the economizer is *off*. In other words, this evidence suggests that a data center with an economizer using MERV 14 filters can expect similar indoor fine-particle mass concentrations to those for a conventional non-economizer data center using MERV 7 filters.

Indoor and outdoor concentrations of sulfate ions associated with airborne particles are presented in Table 3-2 and the ratios are depicted in Figure 3-5. Particles containing these ions are of special concern in data centers due to their ability to absorb water (deliquesce) and create conductive bridging between isolated conductors within computer servers. Sulfate-bearing particles tend to penetrate into the data center and persist with a higher I/O proportion than the PM fine-particle mass concentration measured by optical particle counting. Qualitatively, such a finding might be expected, since atmospheric sulfate is concentrated in the submicron portion of the accumulation mode (Milford and Davidson, 1987) and is consequently expected to exhibit a relatively high indoor proportion of outdoor particles (Riley et al., 2002). Furthermore, while sulfate in particles is thermodynamically stable, PM measured by the OPCs may include particles that volatilize once indoors, contributing to a lower I/O ratio. These results highlight an important point: if sulfate is a primary particle constituent of concern, then OPCs with the particle-size limitations used in this study may not be an ideal proxy for estimating the proportion of potentially harmful particles present in a data center.

However, for sulfate, as for OPC-determined fine-particle mass, the I/O ratios during the MERV 14 period when the economizer is on are comparable to the I/O ratios during the MERV 7 period when the economizer is off.

Indoor and outdoor concentrations of nitrate ions associated with airborne particles are presented in Table 3-3 and the ratios are depicted in Figure 3-6. The I/O ratios measured for particulate nitrate are much lower than the counterpart sulfate ratios under most conditions. Nitrate ratios can be influenced by gas-to-particle conversion processes that occur inside the data center. Ammonium nitrate particles can easily shift between the condensed phase and their gaseous constituents with changes in temperature or in the concentrations of gaseous constituents (Lunden et al., 2003). Table 3-4 shows the average indoor and outdoor temperatures measured at the data center. Indoor temperatures tended to be warmer than outdoor temperatures during the economizer-on period and cooler than outdoors during the economizer-off period but to extents that varied among the different filtration periods. The differences between the indoor and outdoor temperatures could have affected the extent of nitrate volatilization and formation. The gas-phase constituents of ammonium nitrate, specifically ammonia and nitric acid, can also be lost by interactions with indoor surfaces. The loss of these gaseous constituents, particularly nitric acid, to indoor surfaces is another important driver influencing the loss of indoor particulate nitrate. Indoor sources of gaseous constituents might have also influenced the measured nitrate concentrations and corresponding I/O ratios. Although cleaning product use is typically minimal in data centers, one interior wall of the Sunnyvale data center is glass that is routinely cleaned with an ammonia-based product. The consequent indoor release of ammonia may have

influenced the measured I/O nitrate ratios by shifting the balance between the gaseous constituents — ammonia and nitric acid — and particulate ammonium nitrate.

Indoor and outdoor concentrations of chloride ions associated with airborne particles are presented in Table 3-5 and the ratios are depicted in Figure 3-7. The measured particulate chloride concentrations are similar in magnitude to previous chloride PM_{2.5} measurements from the California coast and significantly higher than particulate chloride concentrations measured in other regions of California (Chow et al., 1996). Sodium chloride from marine aerosol is primarily found in coarse particles, but can react with acidic gases and ammonia to form submicron particulate ammonium chloride (Harrison and Pio, 1983). The Sunnyvale data center is located near the San Francisco bay and less than 2 km from a wastewater treatment plant and numerous salt ponds, all of which may have contributed to the relatively high particulate chloride concentrations measured. Ammonium chloride is similar in volatility to ammonium nitrate (Pio and Harrison, 1987). Consequently, particulate ammonium chloride concentrations can be influenced by temperature and by the concentrations of the gaseous constituents, ammonia and hydrochloric acid. Table 3-5 shows that outdoor particulate chloride concentrations are consistently lower during the economizer-off periods than when the economizer is on. This finding may reflect warmer temperatures during the afternoon hours, when the economizer is off, causing these thermodynamically unstable particles to volatilize into their gaseous constituents. Figure 3-7 shows minimal change in the particulate chloride I/O ratio with improved HVAC filter efficiency. Given the large number of data centers located in coastal regions (Brown et al., 2007), the potential impact of particulate chloride on IT equipment reliability warrants further investigation.

Table 3-1. Average indoor and outdoor concentrations of OPC-measured particles, sorted by economizer activity and filtration efficiency rating.

HVAC MERV rating	OPC PM (0.3-2.0 μm) ($\mu\text{g m}^{-3}$)			
	Economizer ON		Economizer OFF	
	Outdoor	Indoor	Outdoor	Indoor
7	10.2	2.5	5.3	0.43
11	13.2	2.3	6.5	0.38
14	3.3	0.22	0.85	0.03

Table 3-2. Average indoor and outdoor filter-based particulate sulfate concentrations, sorted by economizer activity and filtration efficiency rating.

HVAC MERV rating	Sulfate ($\leq 2.5 \mu\text{m}$) ($\mu\text{g m}^{-3}$)			
	Economizer ON		Economizer OFF	
	Outdoor	Indoor	Outdoor	Indoor
7	1.8	1.4	1.5	0.56
11	1.2	0.92	1.6	0.37
14	0.94	0.37	0.35	0.03

Table 3-3. Average indoor and outdoor filter-based particulate nitrate concentrations, sorted by economizer activity and filtration efficiency rating.

HVAC MERV rating	Nitrate ($\leq 2.5 \mu\text{m}$) ($\mu\text{g m}^{-3}$)			
	Economizer ON		Economizer OFF	
	Outdoor	Indoor	Outdoor	Indoor
7	1.1	0.48	1.2	0.11
11	1.4	0.60	2.5	0.38
14	0.20	0.04	0.16	0.04

Table 3-4. Average outdoor and indoor air temperatures (°C) during the August 2008 study, sorted by filter type and according to whether the economizer was “on” or “off.”

HVAC MERV rating	Outdoor		Indoor	
	ON	OFF	ON	OFF
7	18.0	30.4	21.4	21.4
11	21.2	33.6	22.1	22.1
14	17.5	26.2	22.2	22.2

Table 3-5. Average indoor and outdoor filter-based particulate chloride concentrations, sorted by economizer activity and filtration efficiency rating.

HVAC MERV rating	Chloride ($\leq 2.5 \mu\text{m}$) ($\mu\text{g m}^{-3}$)			
	Economizer ON		Economizer OFF	
	Outdoor	Indoor	Outdoor	Indoor
7	0.75	0.23	0.13	0.09
11	1.3	0.28	0.38	0.12
14	0.12	0.03	0.19	0.15

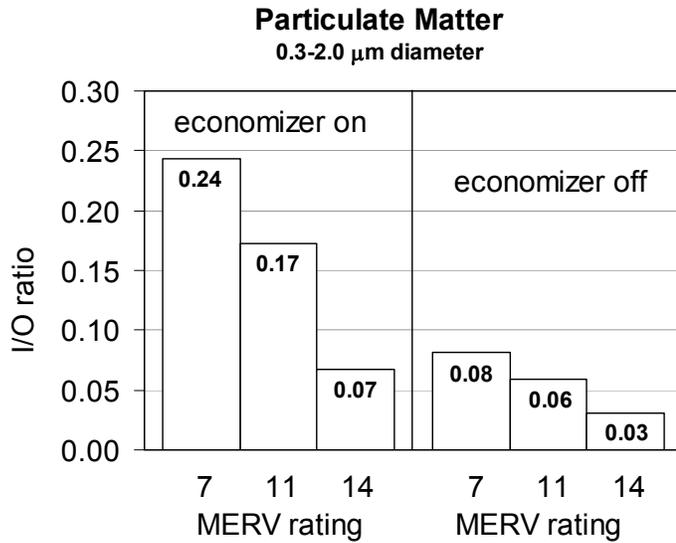


Figure 3-4. Indoor/outdoor concentration ratios for OPC-measured particles, sorted by HVAC filter type and economizer activity.

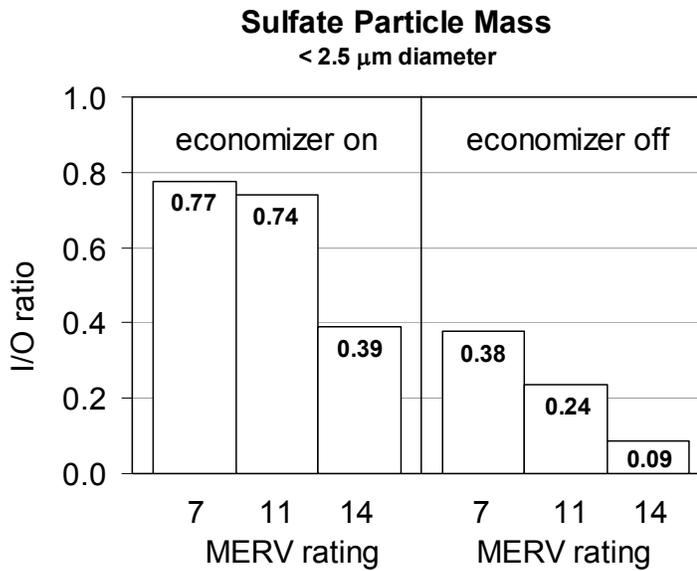


Figure 3-5. Indoor/outdoor concentration ratios for particulate sulfate, sorted by HVAC filter type and economizer activity.

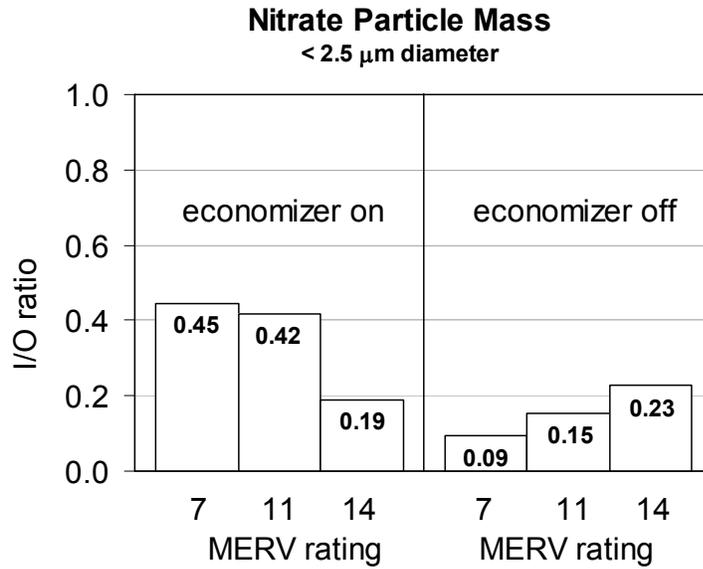


Figure 3-6. Indoor/outdoor concentration ratios for particulate nitrate, sorted by HVAC filter type and economizer activity.

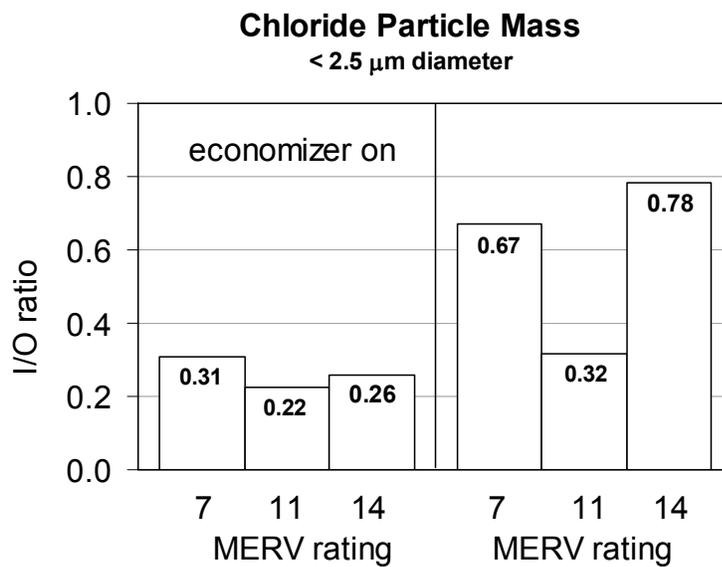


Figure 3-7. Indoor/outdoor concentration ratios for particulate chloride, sorted by HVAC filter type and economizer activity.

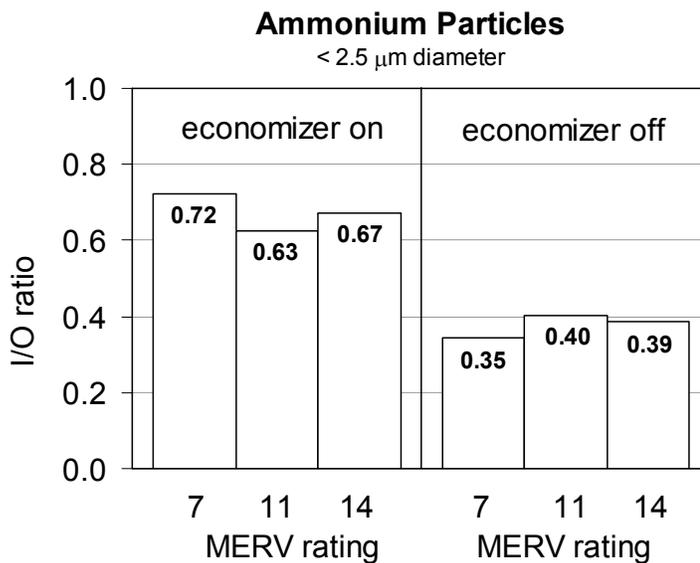


Figure 3-8. Indoor/outdoor concentration ratios for particulate ammonium, sorted by HVAC filter type and economizer activity.

Table 3-6 presents average indoor and outdoor filter-based particulate ammonium concentrations, sorted by economizer activity and filtration efficiency rating. The nitrate and chloride concentrations presented in Table 3-3 and 3-5 include ions measured on the teflon and nylon filters, where species on the nylon filters represent deposited ammonium nitrate and ammonium chloride particles that had volatilized from the teflon filters during the measurement period. Accordingly, molar concentration of ammonium measured on the citric acid impregnated cellulose filters are expected to match the molar concentration of nitrate and chloride measured on the nylon filters. Table 3-7, however, shows significantly higher molar concentration of ammonium ions. This imbalance may in part be due to gaseous ammonia in the air sample not being fully removed after passing

through the citric acid coated denuders, and then depositing onto the citric acid impregnated cellulose filters. A magnitude analysis using the dimensions of the denuders indicates a very low penetration of gaseous ammonia molecules through the denuders, though saturation of the citric acid coating could potentially increase ammonia breakthrough. Indoor and outdoor ratios of ammonium are depicted in Figure 3-8. The I/O ratios show little variation under different filter efficiencies and a higher I/O ratio is observed when the economizer is active, which would be expected if the ammonium concentration were being influenced by outdoor concentrations of gaseous ammonia. The excess measured ammonium, represented as gaseous ammonia, is presented in Figure 3-9 and shows an approximate doubling in ammonia concentration, to about 11 ppb, from the beginning of study (MERV 7 monitoring) to the end of the study (MERV 11 monitoring). Atmospheric concentrations of gaseous ammonia typically range from 0.1 to 10 ppb (Seinfeld and Pandis, 2006), but the outside air intake at the Sunnyvale data center is located approximately 150 m from a heavily trafficked freeway, and average on-road concentrations of ammonia have been measured in the range of 21-51 ppb at the entrance to a roadway tunnel in northern California (Kean et al., 2000)

Tables 3-8 and 3-9 show average concentrations for quartz filter measured and aethalometer measured black carbon particles under each filter condition, again sorted between economizer-on and economizer-off periods. Although the absolute values for the aethalometer are somewhat lower (the mean ratio of aethalometer to quartz-filter determination for black carbon for the different filter/economizer configurations is 0.65), Figures 3-10 and 3-11 show that the quartz-filter-based black carbon measurements provide similar I/O ratios to the aethalometer black carbon measurements.

Table 3-6. Average indoor and outdoor filter-based particulate ammonium concentrations, sorted by economizer activity and filtration efficiency rating.

HVAC MERV rating	Ammonium ($\leq 2.5 \mu\text{m}$) ($\mu\text{g m}^{-3}$)			
	Economizer ON		Economizer OFF	
	Outdoor	Indoor	Outdoor	Indoor
7	4.1	3.0	3.2	1.1
11	8.7	5.4	8.4	3.4
14	3.6	2.4	4.6	1.8

The ammonium concentrations presented here indicate more ammonium than what would be predicted according to the corresponding nitrate and chloride ion concentrations (Tables 3-3 and 3-4).

Table 3-7. Moles of ammonium collected on citric acid impregnated cellulose filters compared to moles of nitrate and chloride collected on nylon filters.

HVAC MERV rating	$\mu\text{mol/m}^3$ ammonium / ($\mu\text{mol/m}^3$ nitrate + $\mu\text{mol/m}^3$ chloride)			
	Economizer ON		Economizer OFF	
	Outdoor	Indoor	Outdoor	Indoor
7	227 / 10.4	164 / 8.1	176 / 8.9	61 / 3.4
11	481 / 12.1	301 / 5.0	465 / 16.3	186 / 6.3
14	200 / 4.1	134 / 1.3	256 / 3.2	99 / 4.1

Measurements are intended to represent teflon filter collected ammonium nitrate and ammonium chloride particles that have volatilized (ammonium sulfate particles are assumed to be stable and not volatilize after depositing on the teflon filters). The molar imbalance indicates a potential additional gaseous ammonia source in the air sample.

Table 3-8. Average measured indoor and outdoor quartz measured black carbon particle concentrations, sorted by economizer activity and filter type.

HVAC MERV rating	Black carbon ($\leq 2.5 \mu\text{m}$) ($\mu\text{g C m}^{-3}$)			
	Economizer ON		Economizer OFF	
	Outdoor	Indoor	Outdoor	Indoor
7	0.34	0.32	0.32	0.18
11	0.76	0.57	0.91	0.20
14	0.26	0.13	0.28	0.05

Table 3-9. Average measured indoor and outdoor aethalometer measured black carbon particle concentrations, sorted by economizer activity and filter type.

HVAC MERV rating	Black carbon ($\leq 2.5 \mu\text{m}$) ($\mu\text{g C m}^{-3}$)			
	Economizer ON		Economizer OFF	
	Outdoor	Indoor	Outdoor	Indoor
7	0.15	0.15	0.16	0.05
11	0.93	0.67	0.65	0.15
14	0.20	0.08	0.17	0.01

Table 3-10. Average indoor and outdoor gravimetrically-measured particulate matter concentrations, sorted by economizer activity and filter type.

HVAC MERV rating	Particulate Matter ($\leq 2.5 \mu\text{m}$) ($\mu\text{g m}^{-3}$)			
	Economizer ON		Economizer OFF	
	Outdoor	Indoor	Outdoor	Indoor
7	11.0	8.6	8.7	3.6
11	12.9	9.1	17.1	9.5
14	4.4	1.09	11.5	5.5

Brass filings had collected on some of the teflon filters used for these measurements, possibly due to slight stripping of the Swagelok threading caused by the repeated opening and closing associated with the filter replacement and airflow recalibration. Consequently, these data are considered unreliable and are only reported here to document these results

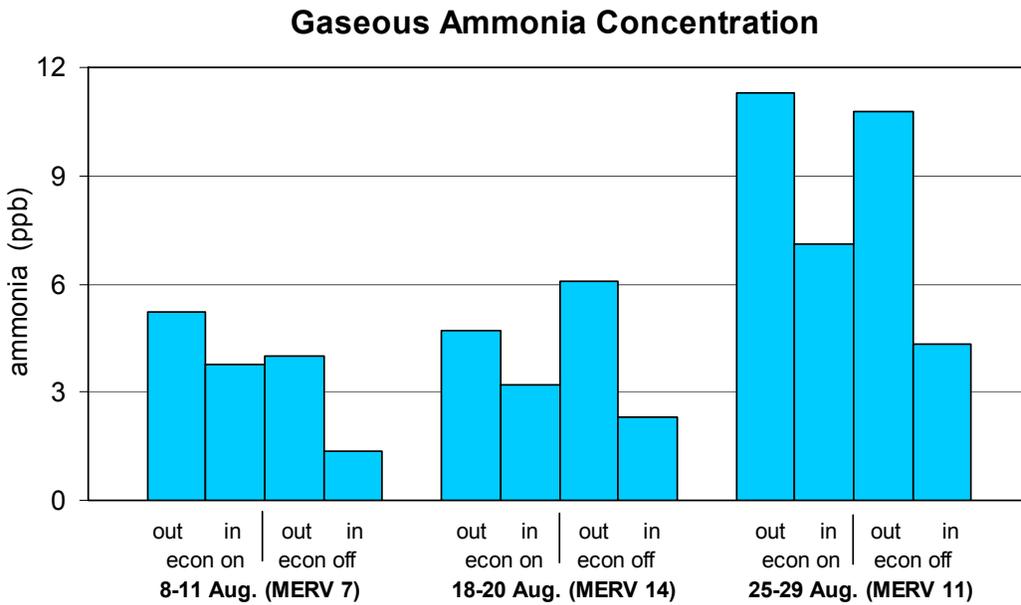


Figure 3-9. The difference between ammonium concentrations measured on the citric-acid coated cellulose filters (Table 3-6) and the difference between nitrate and chloride measured on the teflon filters (Tables 3-3 and 3-5), represented as moles of gaseous ammonia. Ammonium sulfate particles are assumed to be stable and not volatilize after depositing on the teflon filters. The difference gradually increases during the monitoring period, indicating a possible increase in ammonia breakthrough from the citric-acid coated denuders.

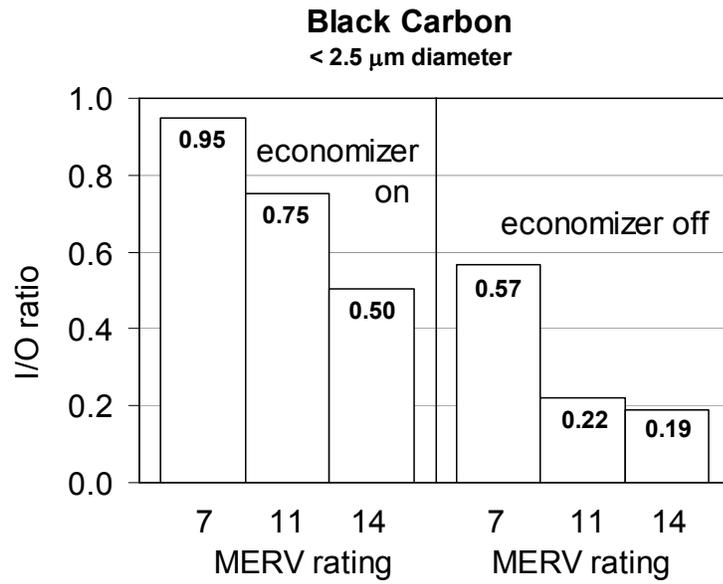


Figure 3-10. Indoor/outdoor concentration ratios for quartz-filter collected black carbon particles, sorted by HVAC filter type and economizer activity.

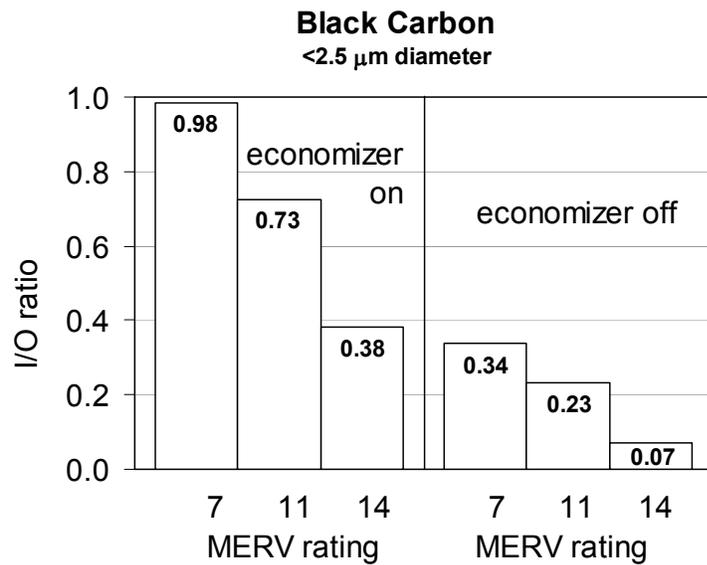


Figure 3-11. Indoor/outdoor concentration ratios for aetholometer-measured black carbon particles, sorted by HVAC filter type and economizer activity.

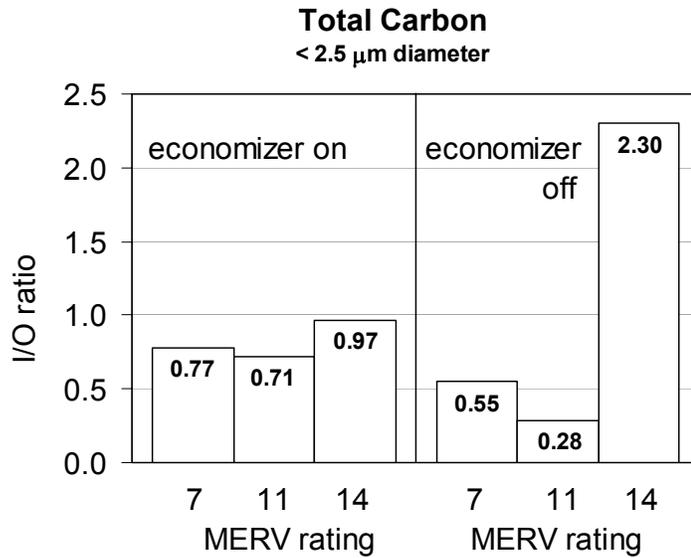


Figure 3-12. Indoor/outdoor concentration ratios for quartz filter collected carbonaceous particles (organic and black carbon), sorted by HVAC filter type and economizer activity.

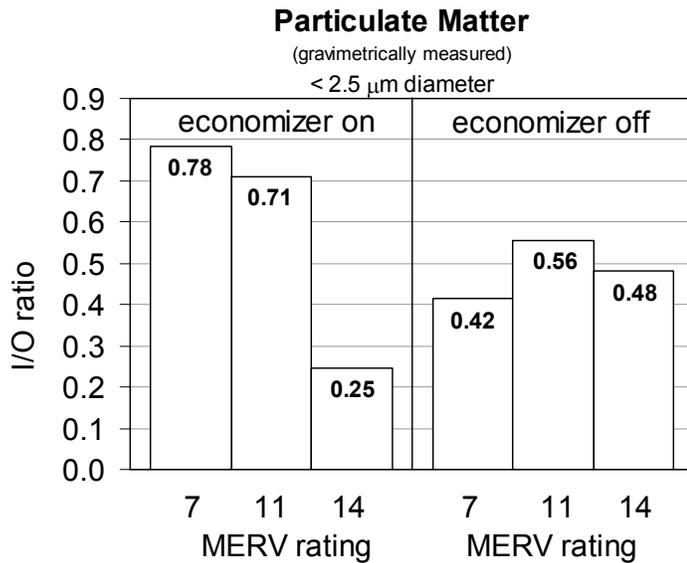


Figure 3-13. Indoor/outdoor concentration ratios for gravimetrically-measured particulate matter, sorted by HVAC filter type and economizer activity. Brass filings had collected on some of the teflon filters used for these measurements, possibly due to slight stripping of the Swagelok threading caused by the repeated opening and closing associated with the filter replacement and airflow recalibration. Consequently, these data are considered unreliable and are only reported here to document these results.

The black carbon I/O ratio trends are similar to those observed for the OPC-based particle mass concentration measurements. Improved filtration reduces the black carbon I/O ratio, so that the MERV 14 I/O ratio during economizer-on periods is comparable to that for MERV 7 filtration during economizer-off periods (Figure 3-4). However, across all filter types and both economizer modes, the I/O ratio of black carbon is greater than that indicated by the OPC data. The black carbon I/O ratios correspond closely to the ratios measured for sulfate under most conditions. Similar I/O ratios for sulfate and black carbon might be expected, since both are thermodynamically stable and both particle types can be dominated by the 0.1-1 μm size range (Seinfeld and Pandis, 2006), which are the most difficult to remove by fibrous filtration (Hinds, 1998). However, during the MERV 7 economizer-on period, the black carbon I/O ratios determined by both measurement methods are near unity and significantly higher than the corresponding sulfate ratios. A contributing factor to this observation might be the measurement uncertainty associated with the low outdoor black carbon concentrations during the MERV 7, economizer-on, monitoring period (Tables 3-8 and 3-9), which was conducted during a weekend. The data center may have been exposed to significant vehicular traffic owing to the close proximity to a heavily trafficked freeway. Direct tailpipe emissions from vehicular traffic contain a high proportion of black carbon and near-freeway concentrations can be lower during weekends than on weekdays owing to relatively lower weekend traffic (Morawska et al., 2002).

The results for total carbon (Table 3-11) do not show a clear trend with increasing HVAC filter efficiency or economizer activity. Measurements of total carbon might have been affected by indoor gaseous emissions of semivolatile (SVOCs) or volatile organic

compounds (VOCs) that sorb to or are otherwise captured by the sample filters. These chemical classes may have significant indoor sources in data centers, such as the constituents of the large amounts of synthetic organics associated with the IT equipment, building materials, and furnishings. High concentrations of VOCs have been measured in data centers relative to other building types and VOC concentrations have been observed to be high in data centers with minimal ventilation, indicating the presence of significant internal sources (Shields et al., 1996).

Table 3-11. Average measured indoor and outdoor quartz measured total carbon particle concentrations, sorted by economizer activity and filter type.

HVAC MERV rating	Total carbon ($\leq 2.5 \mu\text{m}$) ($\mu\text{g C m}^{-3}$)			
	Economizer ON		Economizer OFF	
	Outdoor	Indoor	Outdoor	Indoor
7	1.4	1.1	1.3	0.70
11	2.9	2.1	4.0	1.1
14	0.62	0.60	0.65	1.5

Except for total carbon and chloride, all species measured show significantly reduced indoor concentrations for all measured species with the MERV 14 filters installed as compared with the results for MERV 7 and MERV 11 filters. This observation is partly attributable to the lower outdoor concentrations during the MERV 14 monitoring period. For quality assurance, the outdoor total PM concentrations measured in this study were compared to regionally available outdoor $\text{PM}_{2.5}$ particle concentration data reported by the Bay Area Air Quality Management District (BAAQMD) during these measurement periods (BAAQMD, 2009). The BAAQMD data

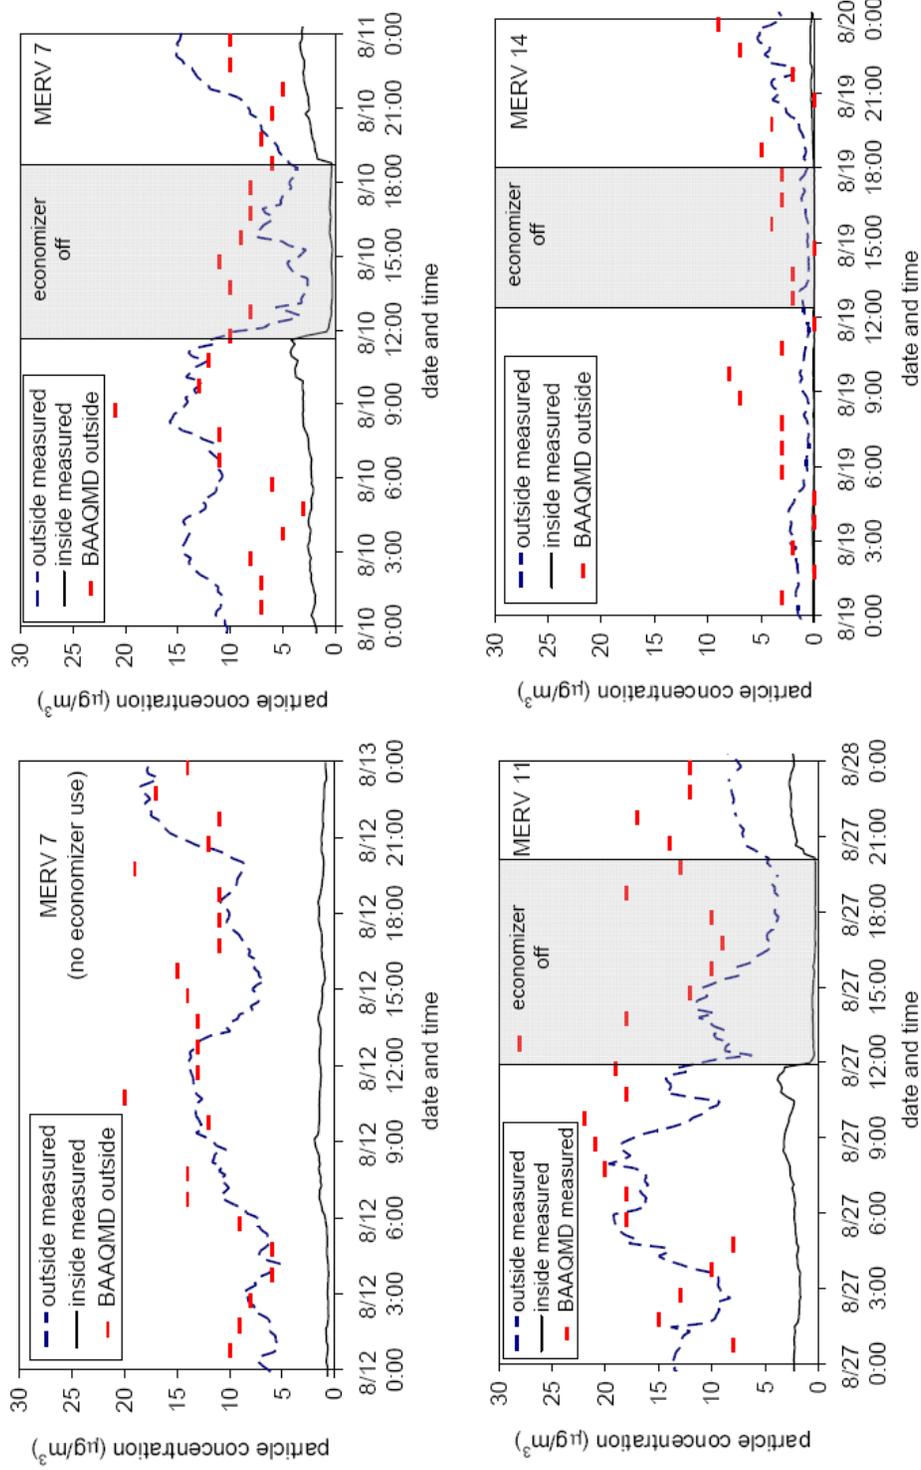


Figure 3-14. Measured time-dependent indoor and outdoor particle mass concentrations over a 24-hour period for each filter configuration. Concentrations represent particles 0.3–2.0 μm in diameter. Upper right graph represents a period when the economizer was inactive for all 24 hours. For the other three graphs, the economizer was inactive approximately between the hours of noon and 6:00 PM and active during the remaining hours. Also indicated are PM_{2.5} measurements from the Bay Area Air Quality Management District (BAAQMD) at a sampling site approximately 25 km from the data center.

represent measurements from a sampling site approximately 25 km away from the Sunnyvale data center. A decrease in outdoor concentrations during the MERV 14 sampling period is observed in both data sets, indicating that the change in outdoor particle concentration is a consequence of changes in ambient conditions, e.g. because of shifting meteorological factors, rather than anything specific to the experimental conditions of this study. A comparison of the measured outdoor concentrations and the BAAQMD data can be seen in Figure 3-14.

Scrutiny of Table 3-1 indicates that, within each measurement period, the OPC-determined outdoor particle concentrations were consistently higher when the economizer was operating (6:00 PM to noon) compared to the hours when the economizer was off (noon to 6:00 PM). This finding may result from several factors. The economizers are active during nighttime hours, when the mixing height of the atmosphere is generally lower, resulting in increases in ambient particle concentrations associated with proximate ground-level emissions (Seinfeld and Pandis, 2006). Also, the economizers are active during morning commute hours, when nearby emissions from vehicular traffic are expected to be high. The close proximity of the Sunnyvale data center's outside air intake to a heavily trafficked freeway may have contributed to these higher outdoor particle concentrations during economizer-on periods. The close proximity of the study site to a freeway may also account for the rapid changes in outdoor particle concentrations observed at the data center during some economizer-on periods that are not seen in the BAAQMD data.

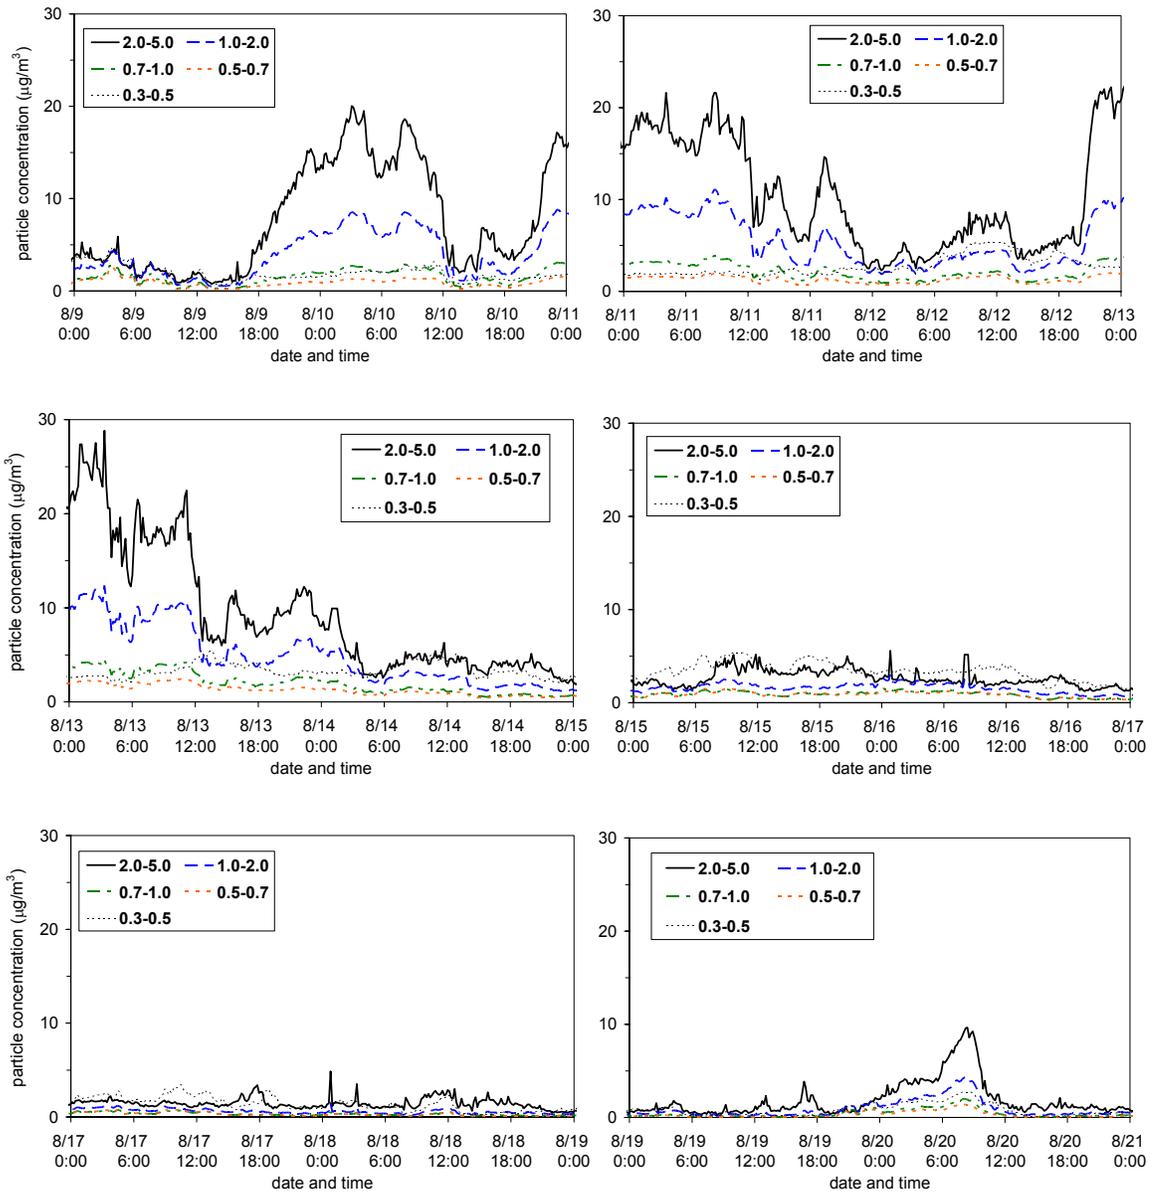


Figure 3-15a. Outdoor particle mass concentrations estimated from OPC data during study period of 9-20 August 2008. Particle concentrations are separated into five particle diameter size bins: 0.3-0.5 μm , 0.5-0.7 μm , 0.7-1.0 μm , 1.0-2.0 μm , 2.0-5.0 μm .

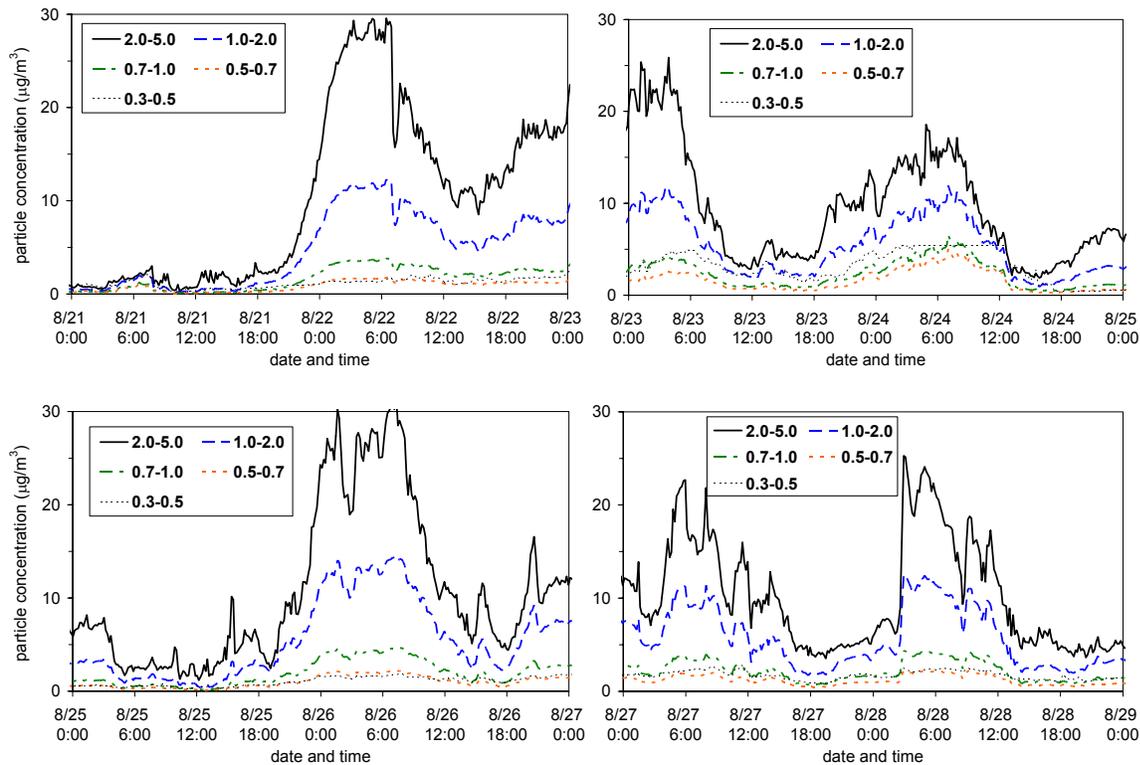


Figure 3-15b. Outdoor particle concentrations measured by OPC during study period of 21-28 August 2008. Particle concentrations are separated into five particle diameter size bins: 0.3-0.5 µm, 0.5-0.7 µm, 0.7-1.0 µm, 1.0-2.0 µm, 2.0-5.0 µm.

3.3.2. Energy Use

Table 3-12 presents the average total building power and the average power specific to IT equipment for each filter and economizer operating condition. Average power values are separated into the two time categories representing the economizer-on (6:00 PM - noon) and economizer-off (noon - 6:00 PM) periods for each filter configuration, except for the first data line presented in Table 3-12 for which the economizer was off during all times. As expected, the average IT loads are nearly

Table 3-12. Measurement time periods, average total building power, and the average power specific to IT equipment for each filter type.

Measurement period (2008)		HVAC MERV rating	Power use measurement results (kW)					
			Total building		IT Equipment		Non IT use	
Start	End		18:00- 12:00	12:00- 18:00	18:00- 12:00	12:00- 18:00	18:00- 12:00	12:00- 18:00
12 Aug 0:00	14 Aug 0:00	7	1047 *	1081	680	680	367 *	402
8 Aug 13:00	11 Aug 12:30	7	937	1078	681	679	256	398
25 Aug 15:00	29 Aug 13:00	11	987	1105	687	686	299	418
18 Aug 19:30	20 Aug 11:30	14	930	1071	690	688	240	383

* The economizer was not operated during 12-14 August.

constant across all conditions, whereas the average load for the entire building decreases for the periods when the economizer is active. A smaller decrease in the main building load is also observed during the period when the economizer remained off for both time categories. The decrease is probably attributable to lower outdoor temperatures, which would affect the amount of cooling required. Reduced lighting and other auxiliary demands during the evening and early morning hours when few occupants are in the building also may contribute to the observed difference. Cumulative probability distributions of the estimated non-IT power use for the economizer-on and economizer-off periods for each filter type are presented in Figure 3-16. Lower power demand during the economizer-on periods can be attributed primarily to reduced chiller operation. After adjusting for the 35 kW difference (402-367 kW) during the configuration when the economizer was off for both time categories, economizer savings can be estimated from the differences between the non-IT power use during each filter and economizer configuration. Applying this method to the power values in Table 5, the estimated

average power savings associated with economizer use for the building were 107, 84, and 108 kW for the MERV 7, 11, 14 filter cases, respectively.

The lower savings during the MERV 11 period can be attributed to higher outdoor temperatures during that period, as shown in Table 3-4. The higher ambient temperatures decreased the number of hours during which the chiller was completely shut off. When the economizer is active, the system operates in either “full” or “partial” economizer mode. During full-mode operation, the outdoor temperature is below the supply air set point and no chiller power is required. When the outdoor temperature is above the supply air set point, but below the return air temperature, the system operates in partial economizer mode, during which the chiller is active, but at a reduced level compared to when the economizer is off.

The economizer-on non-IT power values for all three filter conditions presented in Figure 3-16 show two distinct distributions, which represent conditions during full and partial economizer modes. During the full economizer mode, the non-IT power consumption is reduced to approximately 150 kW; this level is indicative of the base building power demand when the chiller is off. During the partial economizer mode, the building power load shifts to about 300-400 kW. A subtle bimodal distribution can also be seen in the economizer-off values, a result of differences in afternoon temperatures between different days within the same measurement period. .

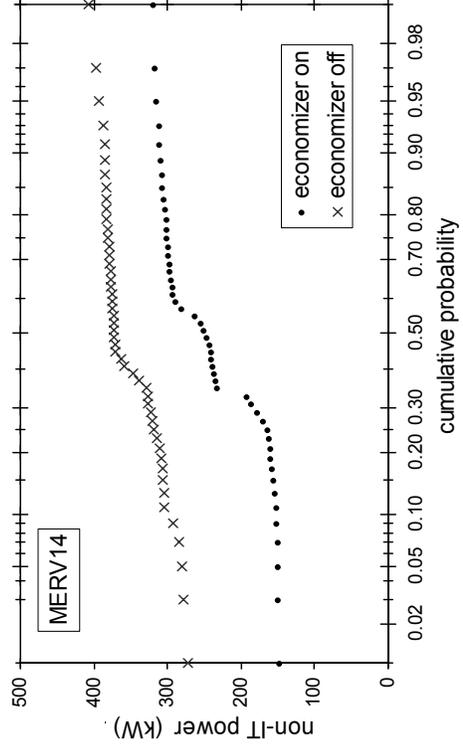
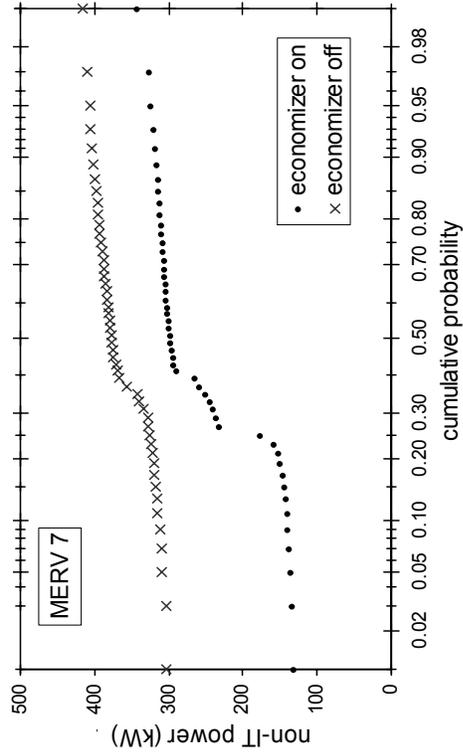
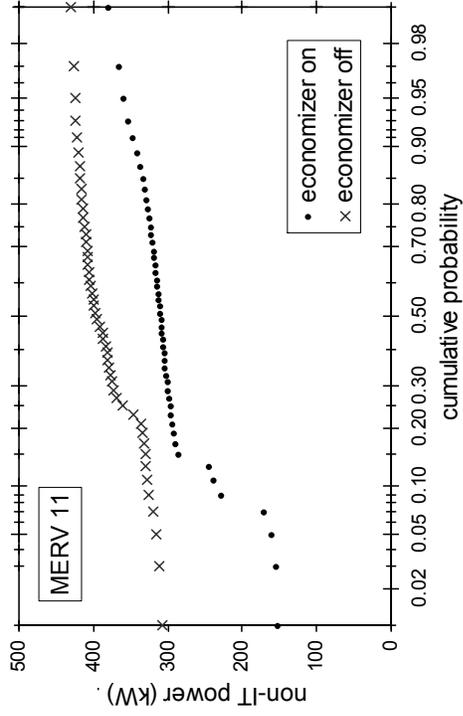


Figure 3-16. Cumulative probability distributions of estimated non-IT power use for the economizer-on and economizer-off periods for each filter type. Each data point represents a 5-minute average power use, through for clarity only representative data points are displayed. Economizer-off values have been adjusted for the 35 kW difference measured when the economizer was off for both time categories. Results indicate an average power savings of 107, 84, and 108 kW for economizer use for the MERV 7, 11, 14 filter conditions, respectively.

Evidence in Figure 3-16 reveals that the power savings observed in August is primarily attributable to periods when the economizer is in partial operation. During the MERV 11 measurements, full economizer operation accounted for less than 10% of the economizer-on period. Economizer savings can be expected to be much larger, possibly more than double, during cooler months when the chiller can be completely off for longer periods.

Improved filtration might be associated with an energy penalty from higher fan power to overcome larger pressure drops. However, estimates indicate that fan power increases would be relatively minor. Building design fan curves indicate a total supply fan power demand of 64 kW during the baseline (MERV 7) filter condition. Supply fan power with the MERV 11 and 14 filters is calculated to increase to 67 and 70 kW, respectively.

Hence, the MERV 14 filters are expected to increase fan power by about 10%. The absolute increase of 6 kW is much smaller than the ~100 kW of expected chiller power savings during economizer use. Furthermore, the increase in fan power is constant throughout the year whereas chiller savings during economizer use would be expected to increase during cooler periods relative to the August period studied here.

3.4. Conclusions

In conventional practice, economizers are often not implemented in data centers. A key reason is to limit the exposure of IT equipment to particles of outdoor origin. This aspect of data center design contributes to their high rate of energy use. The research reported here shows that economizer use combined with high-quality particle filtration can reduce data center power demand while simultaneously maintaining indoor particle

levels similar to those using conventional (non-economizer) ventilation-system design and operation. Specifically, for most of the particle types measured in this study, when MERV 14 filters are used along with economizers, the I/O ratio is near levels that occur when using conventional MERV 7 filters without economizers. Energy analysis of the data center investigated in this study revealed that, even during a warm summer month (August) in northern California, chiller energy savings from economizer use greatly outweighed the increase in fan energy associated with improved filtration.

Investigating economizer use combined with improved filtration for other climates and during other seasons is needed to generalize the findings from this case study to a larger scale. Future work should also explore whether the protection provided by improved filtration is necessary to ensure IT equipment reliability. The particle concentrations measured at this site during economizer use with conventional MERV 7 filters, while higher than concentrations during non-economizer periods, were still below many IT equipment and data center guidelines. A more detailed understanding of how particle concentrations might adversely influence equipment reliability may obviate the need for improved filtration and the associated increase in fan energy and material cost associated with this protective measure. However, even without such understanding, the results reported in this chapter indicate that improved filtration is a viable mitigation alternative. High quality filtration can be combined with economizer use to simultaneously protect electronic equipment from outdoor particles while achieving significant energy savings in the operation of data centers.

Chapter 4: Energy implications of economizer use in California data centers

This chapter presents energy modeling results of economizer savings in California. This energy modeling marks a shift in the dissertation away from understanding the IAQ issues that hinder economizer implementations and towards quantifying the potential energy saving benefits of using economizers in data centers. An energy modeling approach is presented to estimate data center efficiency. Power Use Efficiency (PUE) values are calculated for a modeled data center in various climates. PUE values are calculated for a baseline (no economizer) case and compared to air-side and water-side economizer cases. Results for each scenario are presented, showing significantly greater savings with air-side economizers, but also showing that these savings can be lost with strict humidity restrictions. Reproduced in part with permission from American Council for an Energy-Efficient Economy, Proceedings of the 2008 ACEEE Summer Study on Energy Efficiency in Buildings, Scaling Up: Building Tomorrow's Solutions. Asilomar, CA, August 17 – 22, 2008, Paper 251, Copyright 2008, ACEEE

4.1. Introduction

The decision to implement economizers in data centers should be based on weighing the relative costs and benefits of such design changes using the best information available. The primary cost of economizer use is the increased equipment reliability concern raised by changes in indoor air quality, especially associated with the increased introduction of particulate matter from outdoor air. Benefits to economizer implementation include the cost savings and reduced environmental impact associated with a more energy efficient cooling design. Chapters 2 and 3 have sought to advance understanding of the air quality concern to improve industry information on this economizer cost. Monitoring in eight different data centers revealed that these buildings typically operate at particle concentration levels well below data center air quality guidelines. A detailed experimental evaluation at the Sunnyvale data center showed that

the observable increase in the indoor/outdoor particulate ratio during economizer use can be eliminated through the use of better particle filters without significantly compromising the economizer energy savings. The metered electricity use at the Sunnyvale site suggested substantial electricity savings from economizer use, with as much as 60% savings being observed in non-IT energy. However, although based on robust experimental evidence, the energy benefits observed at the Sunnyvale data center may be specific to conditions at that site. The energy savings potential from a broad, industry-wide implementation of economizers are not known. A better understanding of the potential benefits associated with economizer use, such as identifying when and where these benefits are substantial, can help decision makers overcome the hesitation associated with providing high volume flow rates of outside air to data centers.

This chapter takes a closer look at how electricity is used in a data center and how more efficient cooling systems that employ low outside air temperatures can reduce chiller loads in different locations. Two approaches for economizer use are analyzed here. Air-side economizers, which represent the economizer style discussed in Chapters 2 and 3, involve directly providing outdoor air for cooling whenever the temperature of outside air is lower than the set-point for return-air temperature in the data center. An alternative economizer design, known as a water-side economizer, employs cooling towers that use ambient air to directly cool or precool the chilled water in the cooling system. Water-side economizers are typically more complex and expensive in design and require more frequent maintenance compared to air-side economizers. On the other hand, water-side cooling systems represent a form of economizer that avoids direct contact with excess outside air of the IT equipment in the data center. Water-side economizers can be

viewed as an appropriate design choice for data center owners wary of exposing IT equipment to outside air, but with desire to improve the energy efficiency of their data center operation. While the results from Chapter 3 should allay some of the fears associated with air-side economizer use, this chapter attempts to better understand the relative benefits of these two economizer designs. In this chapter, the energy implications for a data center using a CRAC system are quantified and compared with alternative cooling systems using air-side or water-side economizers for five different California climate zones. The modeling results and discussion focus on understanding the energy implications for both type of economizers and their effectiveness in different climate zones.

4.2. Methods

4.2.1. Data center design scenarios

Energy-use simulations were performed for three different data center HVAC design scenarios. The baseline case considers a data center using conventional Computer Room Air Conditioning (CRAC) units. In this scenario (Figure 4-1), CRAC units are placed directly on the computer room floor. Air enters the top of a CRAC unit, passes across the cooling coils, and is then discharged to the underfloor plenum. Perforations in the floor tiles in front of the server racks allow the cool air to exit from the plenum into the data-center room. Fans within the computer servers draw the conditioned air upward and through the servers to remove equipment-generated heat. After exiting the backside of the server housing, the warm air rises and is transported to the intake of a CRAC unit.

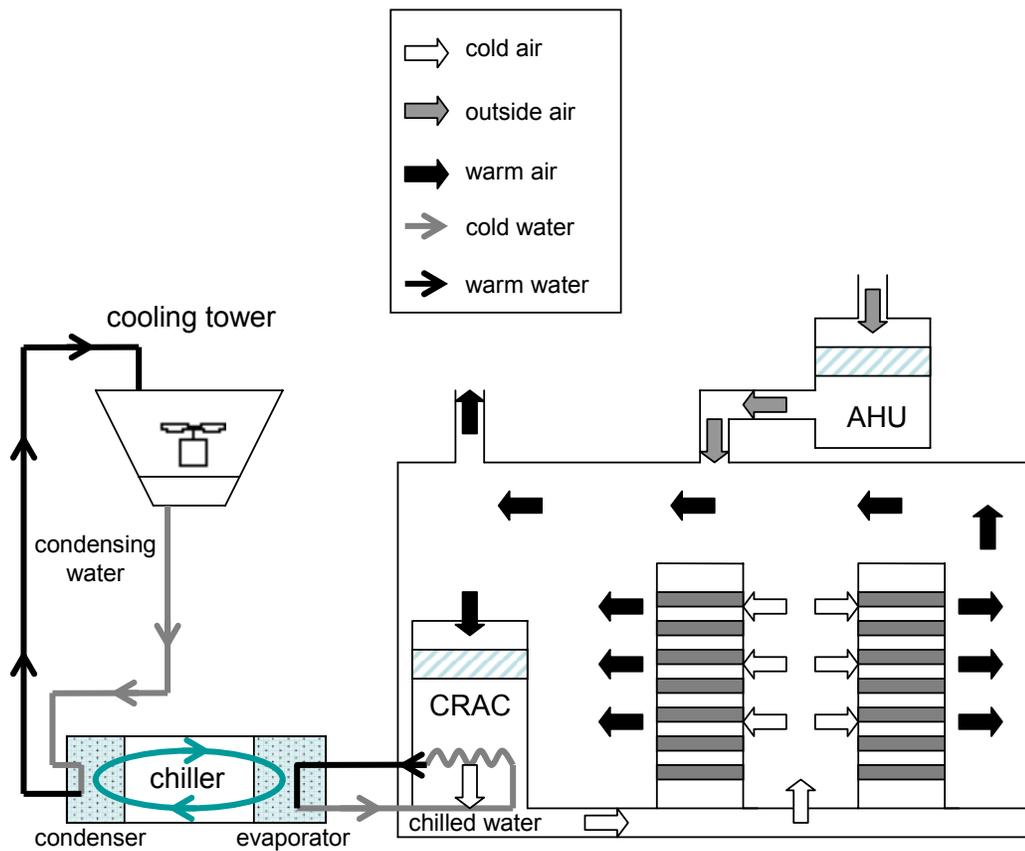


Figure 4-1. Air and water flow schematic of data center cooling without any economizer use (base-case design).

Most air circulation in the baseline scenario is internal to the data center. A small amount of air is supplied through a rooftop AHU to positively pressurize the room and to supply outside air for occupants. Cooling is provided by a water-cooled chiller plant. Refrigerant in the chillers is used to cool water through heat exchangers at the evaporator. The chilled water is then piped to the CRAC units on the data center floor. Waste heat from the chiller refrigerant is removed by water through heat exchangers in the condenser. Condenser water is piped from the cooling towers, which cools the water through interaction with the outside air. This baseline design is common to most mid- to large-size data centers (Tschudi et al., 2004; Rumsey Engineers, 2005; Syska Hennessy Group, 2007).

The water-side economizer scenario assumes a CRAC unit layout similar to that of the baseline case, except that additional heat exchangers are installed between the condenser water in the cooling towers and the chilled water supplied to the CRAC units (Figure 4-2). Under appropriate weather conditions, the cooling towers can cool the condenser water enough to cool the chilled water in the CRAC units directly, without operating the chiller plant. The CRAC units and chiller plant are assumed to be the same as in the baseline scenario.

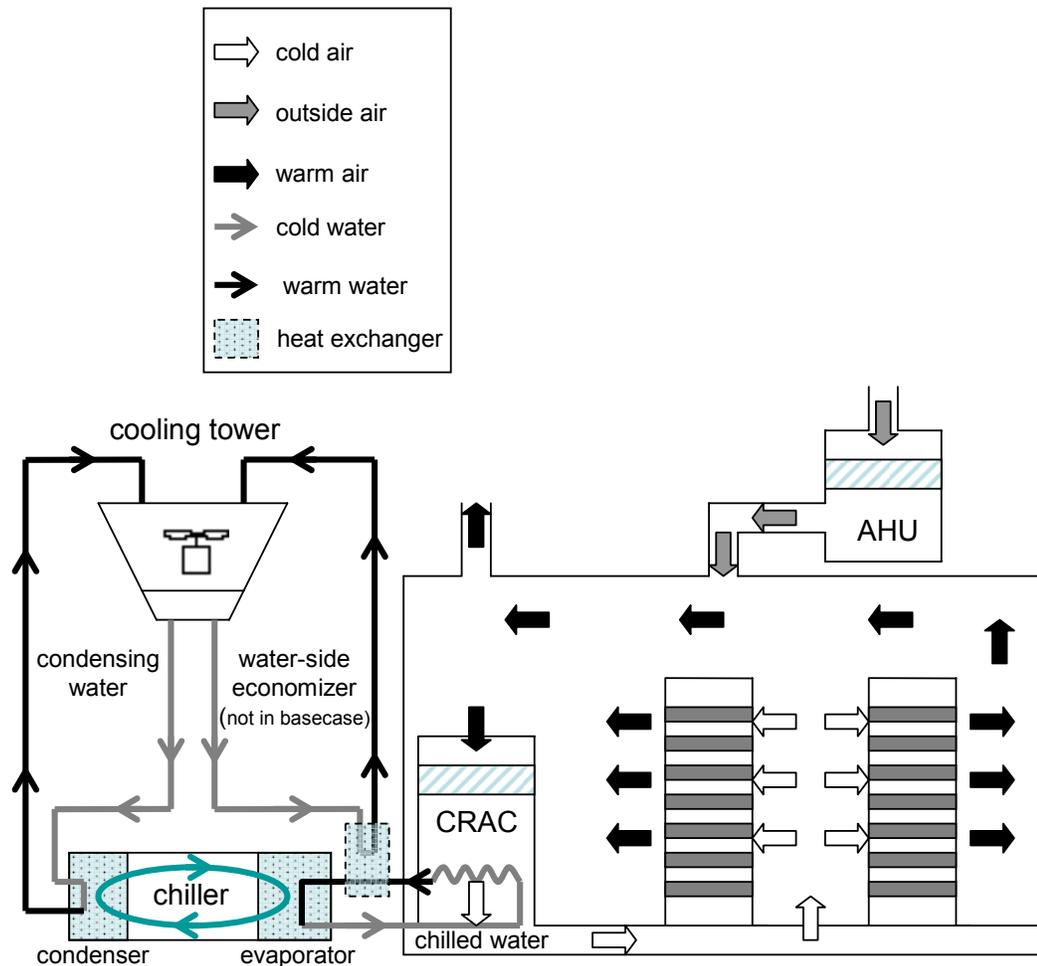


Figure 4-2. Air and water flow schematic of data center cooling with a water-side economizer.

The air-side economizer scenario requires a different type of air delivery than typically found in a data center with conventional CRAC units (Figure 4-3). AHUs are placed outside of the data center room, commonly on the rooftop, and air is then sent to and from the computer racks through ducts. A ducted air delivery system creates greater air resistance than a conventional CRAC unit layout, though this system better prevents cold and warm air from unintentionally mixing within the data center. When the outside air temperature is equal to or below the temperature of the air supplied to cool the server, the AHU can directly draw outside air into the data center and exhaust all of the return air after it has passed across the computer servers. The movement of 100% outside air through the system can require more fan energy than the baseline case due to additional exhaust fans used to remove circulated air from the data center space. The fan energy is also increased in the air-side economizer design due to increased air resistance through the system from the additional ducting. However, during this 100% outside air mode the cooling is provided without operating the chiller, chilled water pumps, condenser water pumps, or the cooling tower fans. Outside air is also provided instead of recirculated air whenever the outside air temperature is greater than the supply air temperature but lower than that of the return air. Under this condition the chiller must operate, but the cooling required from the chiller is less than in a case with complete recirculation.

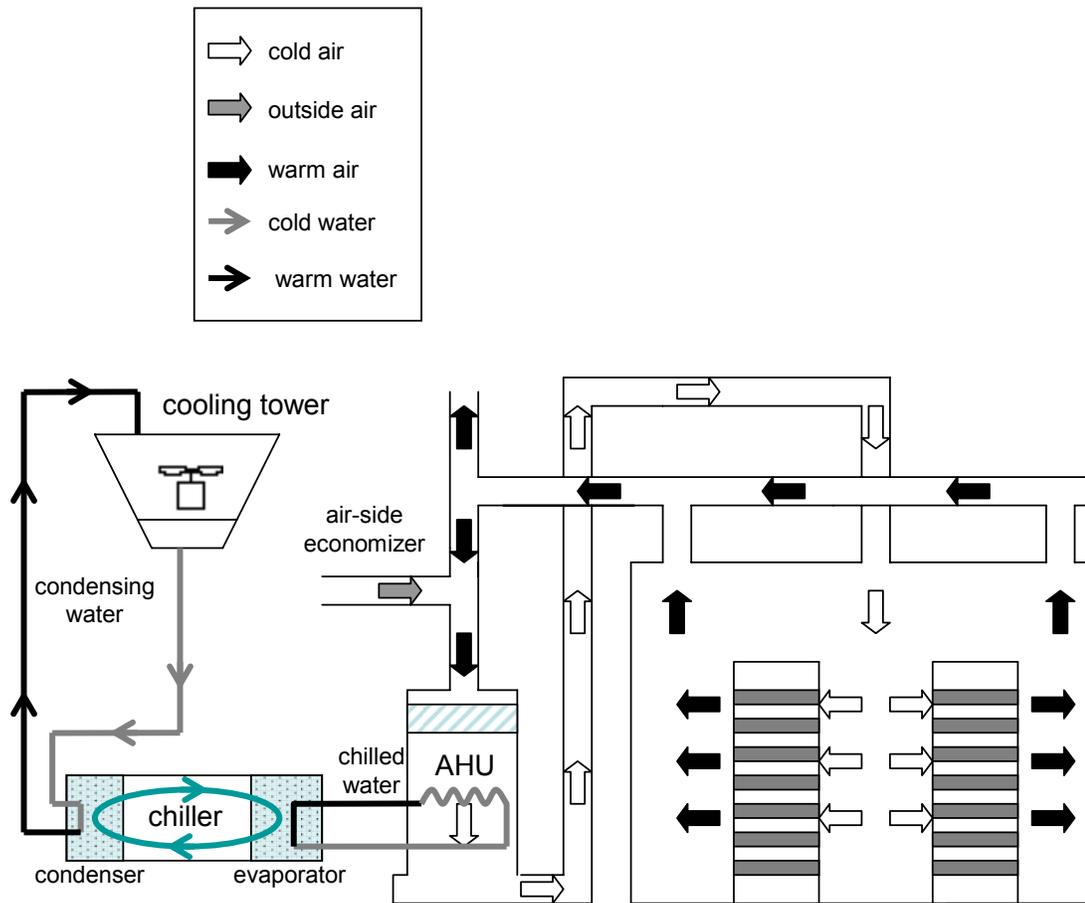


Figure 4-3. Air and water flow schematic of data center cooling with an air-side economizer.

4.2.2 Energy modeling protocol

For each design scenario, the model calculations assume a 2,800 m² (30,000 ft²) data center with an internal heat density of approximately 0.72 kW/m² (67 W/ft²) or 2.4 MW total. This size and power density are characteristic of data centers evaluated in previous studies (Greenberg et al., 2006; Tschudi et al., 2003). The size of data centers varies greatly; 2,800 m² is within the largest industry size classification, which is responsible for most servers in the United States (IDC, 2007). Power density in data centers is rapidly increasing (Uptime Institute, 2000) and a power density of 0.72 kW/m² is currently considered to be of low- to mid-range of industry practice (Rumsey Engineers, 2008).

Basic properties of the modeled data center for all three scenarios are summarized in Table 4-1. The base case and water-side economizer scenarios assume conventional humidity restrictions recommended by ASHRAE (ASHRAE, 2005). The air-side economizer scenario assumes no humidity restriction, which is an adjustment required to gain the energy efficiency benefits in California data center design (Rumsey Engineers, 2008). Air-side economizers also require a different air distribution design. The fan parameters associated with each design scenario are listed in Table 4-2. The properties of other pumps and fans throughout the HVAC system remain constant for all three scenarios. Values are from previous data-center energy analyses (Rumsey Engineers, 2008; Rumsey Engineers, 2005).

The energy modeling approach presented in this chapter applies a previously used protocol (Rumsey Engineers, 2008; Rumsey Engineers, 2005) and is based on a combination of fundamental HVAC sizing equations that apply equipment size and

efficiencies observed through professional experience. Building energy modeling is typically performed using energy models such as DOE-2, which simultaneously models heat sources and losses within the building and through the building envelope (CEC, 2005). However, models such as DOE-2 are not designed to incorporate some of the HVAC characteristics unique to data centers. For example, DOE-2 is not able to model partial water-side economizing in a data center, in which a cooling tower connected to a heat exchanger partially cools the return chilled water before the chiller provides the remaining cooling. Also, data centers have floor-area-weighted power densities that are 15-100 times higher than those of typical commercial buildings (Greenberg et al., 2006). This feature allows for accurate modeling of data-center energy use that focuses exclusively on internal heat load and the thermal properties of outdoor air entering the building, and is the approach taken in this chapter. Specifically, heat generated by data center occupants and heat transfer through the building envelope are negligible relative to the heat produced by IT equipment. The building envelope may influence the cooling load in low-density data centers housed in older buildings that have minimal insulation. Evaluating this building type is worthwhile, but the required analysis is more complex and outside the scope of this dissertation.

Table 4-1. Data center characteristics common to all design scenarios

Data Center Parameters	
Floor Area	2,800 m ²
UPS Waste Heat	326 kW
Data Center Lights	30 kW
Total Rack Load	2000 kW
Total Internal Load	2 MW
Average Internal Load Density	0.72 kW/m ²
Minimum Ventilation	2 m ³ /s
Supply Air Temperature	13 °C
Return Air Drybulb Setpoint	22 °C
Chiller Capacity	1750 kW
Number of Chillers	3

Table 4-2. Data center fan properties

Fan System Parameters	Baseline and WSE			ASE	
	MUAH	Exhaust	CRACs	Supply	Relief
Total Air Flow (m ³ /s)	2	2	234	207	207
Fan Motor Size, Nominal (kW)	5.6	2.2	7.5	22.4	37.3
Number of Fans	1	1	30	10	5
Fan Efficiency	53.3%	44.0%	55.6%	63.8%	67.5%
Fan Drive Efficiency	95%	95%	95%	95%	95%
Fan Motor Efficiency	89.6%	86.2%	90.1%	92.5%	93.2%
VFD Efficiency	n/a	n/a	n/a	98%	98%
Total Static Pressure Drop (Pa)	872	249	398	498	249

Both air-side and water-side economizers are designed to allow the chiller to shut down or to reduce chiller energy load under appropriate weather conditions. Less overall energy is required for operation when the chiller load is reduced, but chiller efficiency is compromised. Chiller efficiencies used in this analysis are shown in Figure 4-4, representing a water-cooled centrifugal chiller with a capacity > 1050 kW and condenser water temperature of 27 °C (80 °F). A chilled water temperature of 7 °C (45 °F), which is standard practice for data center operation, is used in the base case and air-side economizer scenarios. The water-side economizer scenario uses a chilled water temperature of 11 °C (52 °F), which is common when using water-side economizers. The higher temperature increases the required airflow rates but allows for greater use of the water-side economizers. The efficiency curves presented in Figure 4-4 are based on the DOE2.1E software model and apply coefficients specified in the Nonresidential Alternative Calculation Method (ACM) Approval Manual for the 2005 Building Energy Efficiency Standards for Residential and Nonresidential Buildings (CEC, 2005).

Annual data center energy use is evaluated for each of the three configuration scenarios assuming that a data center building is located in one of five different cities distributed throughout California (Figure 4-5). Weather conditions at each city are based on hourly DOE2.1E weather data for California climate zones (CEC, 2005).

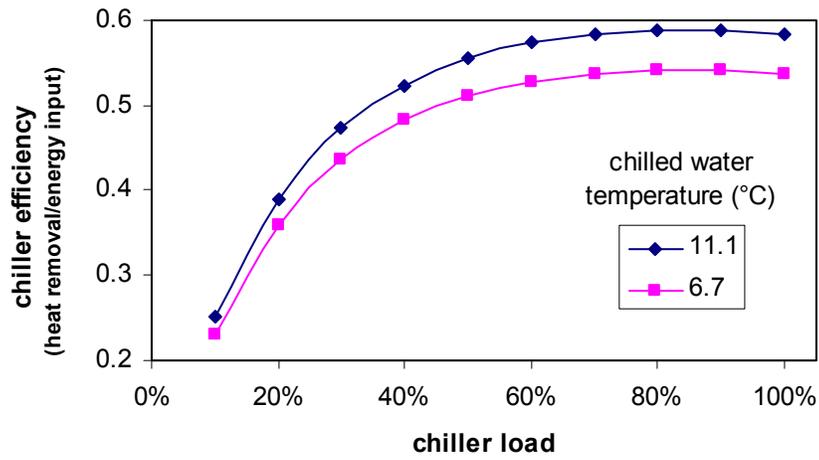


Figure 4-4. Data Center Chillers Part load efficiencies for a water-cooled centrifugal chiller with a capacity > 1050 kW and a condenser water temperature of 26.7 °C (CEC, 2005)



Figure 4-5. Evaluated climate zone locations.

4.2.2. Power Use Effectiveness (PUE)

Figure 4-6 outlines the many components of a data center that contribute to building energy demand during operation. The information processing services are provided by the rows of IT equipment racks that contain servers, storage devices, and network equipment. Data centers include power delivery systems that provide backup power, regulate voltage, and make necessary alternating current/direct current (AC/DC) conversions. Before reaching the IT equipment rack, electricity is first supplied to an uninterruptible power supply (UPS) unit. The UPS acts as a battery backup to prevent the IT equipment from experiencing power disruptions, which could cause serious business disruption or data loss. In the UPS the electricity is converted from AC to DC to charge the batteries. Power from the batteries is then reconverted from DC to AC before leaving the UPS. Power leaving the UPS enters a power distribution unit (PDU), which sends power directly to the IT equipment in the racks. Electricity consumed by inefficiencies in this power delivery chain is considered part of the overall building load.

Electricity entering servers in the IT equipment rack is converted from AC to low-voltage DC power in the server power supply unit (PSU). The low-voltage DC power is used by the server's internal components, such as the central processing unit (CPU), memory, disk drives, chipset, and fans. The DC voltage serving the CPU is adjusted by load specific voltage regulators (VRs) before reaching the CPU. Typical power levels for these various server components are shown in Table 4-3 (Fan et al., 2007). Electricity is also routed to storage devices and network equipment, which facilitate the storage and transmission of data. The electricity used to operate all the components within the servers, storage, and network equipment is considered the IT load.

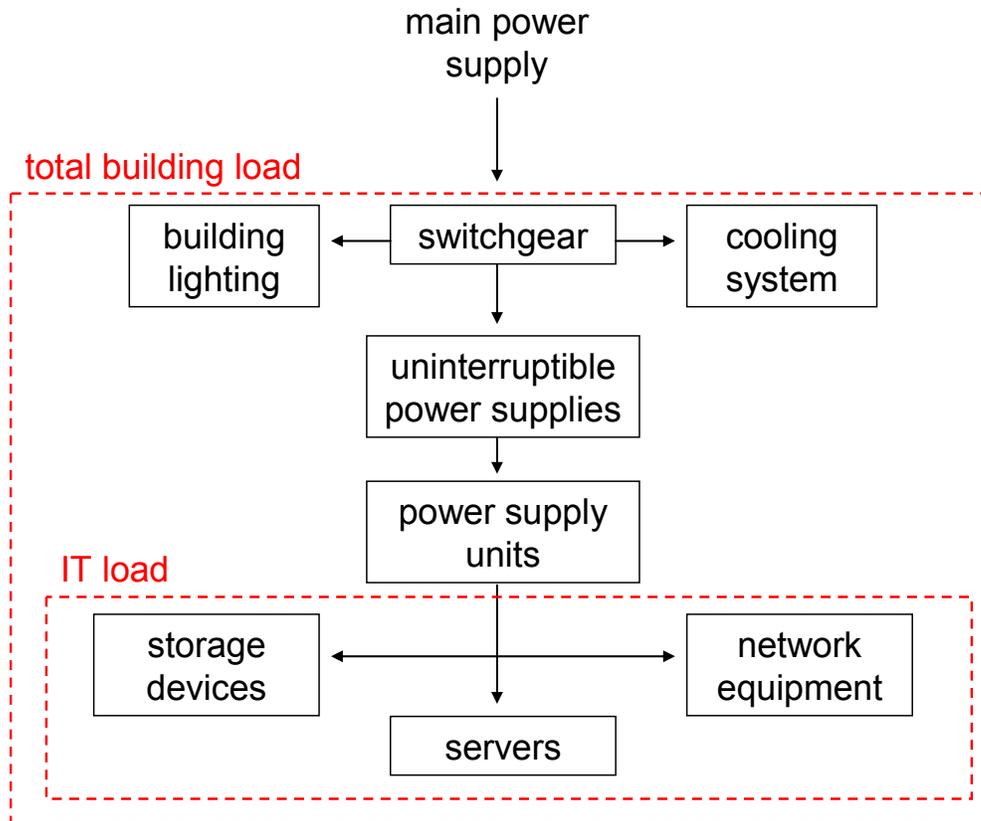


Figure 4-6. Typical electrical components in a data center that represent the total building load. The IT load represents the electricity dedicated to operate the computer servers, storage devices, and network equipment.

Table 4-3. Component peak power consumption for a typical computer server (Fan et al., 2007)

Component	Peak Power (watts)
CPU	80
Memory	36
Disks	12
Peripheral Slots	50
Motherboard	25
Fan	10
PSU Losses	38
Total	251

The overall building energy demand in the model is calculated as the sum of the loads generated by IT equipment, chiller use, fan operation, transformer and UPS losses, and building lighting. The chiller encompasses coolant compressors, chilled water pumps, condensing water pumps, humidification pumps, and cooling-tower fans. Energy demand for servers, UPS, and lighting are constant, unaffected by the different design scenarios, but are included to determine total building-energy use. The energy efficiency of a data center can be characterized by using a simple metric called the Power Use Effectiveness (Stanley et al., 2007), which is defined as the ratio of the total data center building load to the data center IT equipment load.

$$\text{PUE} = \frac{\text{Total Building Electricity Load}}{\text{IT Electricity Load}} \quad (4-1)$$

This metric can quickly indicate how much more electricity is required to operate a data center relative to the electricity that directly operates the IT equipment; where an ideal PUE of 1.0 would indicate that all electricity consumed at the data center is dedicated to IT equipment and a PUE of 2.0 indicates that the electricity required for the cooling, lighting, and power distribution losses is equivalent to IT electricity demand. The PUE performance metric is used here to compare the results from each design scenario evaluated in this chapter.

4.3. Results and Discussion

Results from each modeled scenario are presented in Table 4-4 as PUE values to represent the energy utilization of the HVAC system. The PUE ratio for the base case is 1.55 and, as expected, is the same for all the cities analyzed, since the operation of this

design is practically independent of outdoor weather conditions³. The base case PUE ratio is better than the current stock of data centers in the United States because the base case represents newer data centers with water-cooled chillers. The PUE ratio of data centers is generally assumed to be about 2.0 (Brown et al., 2007; Koomey, 2007) and recent industry data gathered by the EPA estimates the national PUE at 2.04 (Sullivan, 2009). Water-cooled chillers are more efficient than the air-cooled chillers and direct expansion (DX) cooling systems found in older data centers.

Table 4-4. Estimated PUE values specific to design scenario and location

	San Jose	San Francisco	Sacramento	Fresno	Los Angeles
Baseline	1.55	1.55	1.55	1.55	1.55
Air-side Economizer	1.44	1.42	1.44	1.46	1.46
Water-side Economizer	1.53	1.54	1.53	1.53	1.54

The PUE ratios for the air-side and water-side scenarios show that air-side economizers consistently provide savings relative to the base case, though the difference in savings between the two scenarios varies. It is important to note that even small changes in the PUE ratio indicate significant savings, given the large amount of energy used in data centers. For example, given a 2.0 MW internal load, reducing the performance ratio at the model data center in San Jose from 1.55 to 1.44 represents a savings of about 1.9 million kWh/y, which corresponds to a cost savings of more than \$130,000/y (assuming \$0.07/kWh).

³ Cooler outdoor temperatures will improve the efficiency of heat removal from the compressor, but these changes in efficiency are too subtle to be captured in the PUE metric.

Figure 4-7 shows the disaggregation of the cooling systems' annual energy use, normalized by floor area, for each modeled data center by location and design scenario. The annual energy use dedicated to the IT equipment, UPS, and lighting is 6.3, 1.0, and 0.1 MWh/m², respectively. These energy values are independent of the climate and HVAC design scenario and are not included in the graphs in Figure 4-7. Economizer use can potentially be controlled by a combination of outside air temperature, humidity, and enthalpy; however, results shown in Figure 4-7 are for economizer use controlled by outside air temperature only, which is common design practice in California climates. Results show that the air-side economizer scenario provides the greatest savings in San Francisco owing to the consistently cooler temperatures in this region. The warm temperatures in Fresno resulted in the least favorable air-side economizer savings. Sacramento benefited the most from the water-side economizer scenario while minimal savings were realized with this configuration in Los Angeles and San Francisco. The San Francisco water-side economizer scenario, where significant gains would be expected because of the cool climate, is hindered by the conventional setpoints for activating full water-side economizer use. Figure 4-8 and 4-9 present distributions of drybulb and wetbulb temperatures, respectively, for the 8760 hours throughout the typical meteorological year at each of the data center modeling locations. Figure 4-8 shows that the drybulb temperature in San Francisco primarily resides in the partial air-side economizers region where energy efficiency benefits can be realized. However, Figure 4-9 shows that while San Francisco has few hours with very high wetbulb temperatures, the majority of the hours have a wetbulb temperature slightly higher than the water-side economizer setpoint, owing to the relatively high moisture content in this location.

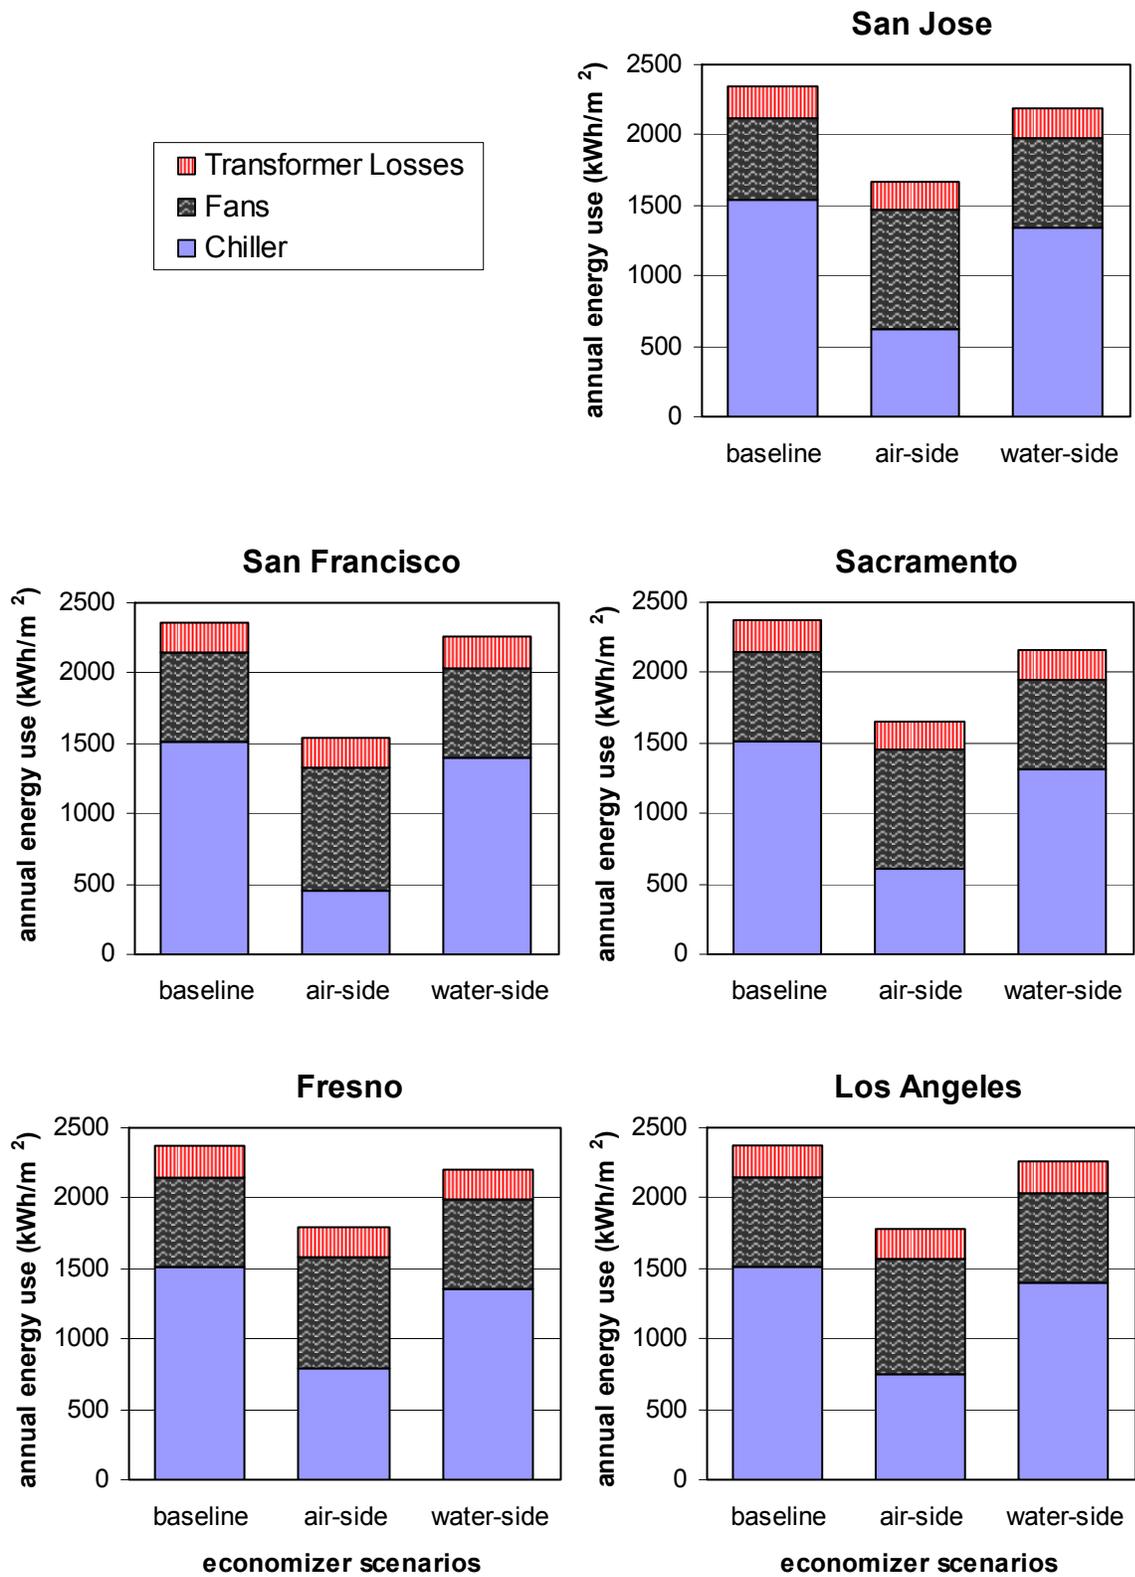


Figure 4-7. Disaggregated energy use for each design scenario (climate dependent values only)

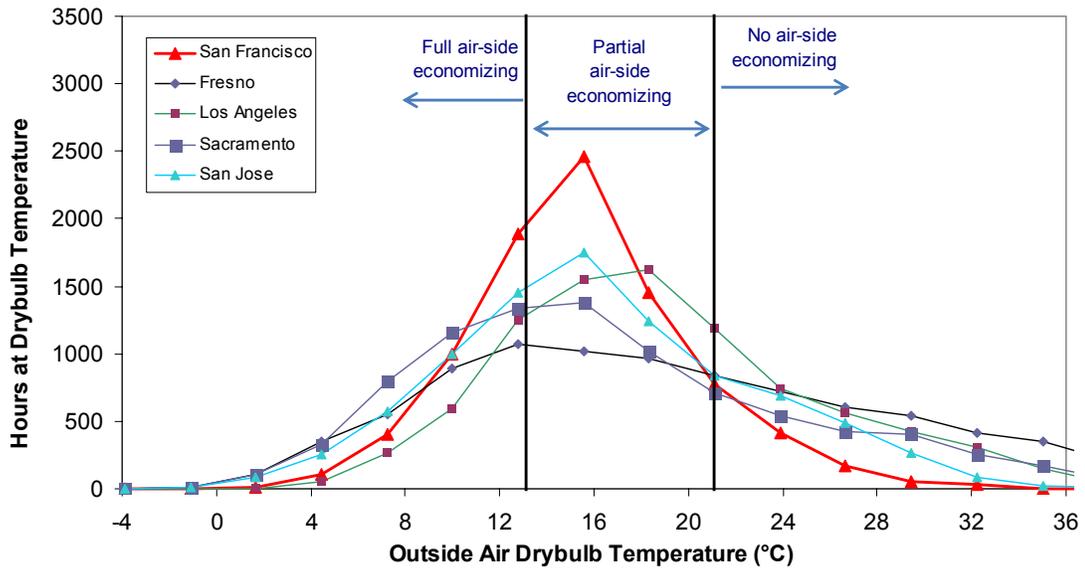


Figure 4-8. Hourly distribution of the outside air drybulb temperature for the 8760 hours throughout the year at the five California data center locations

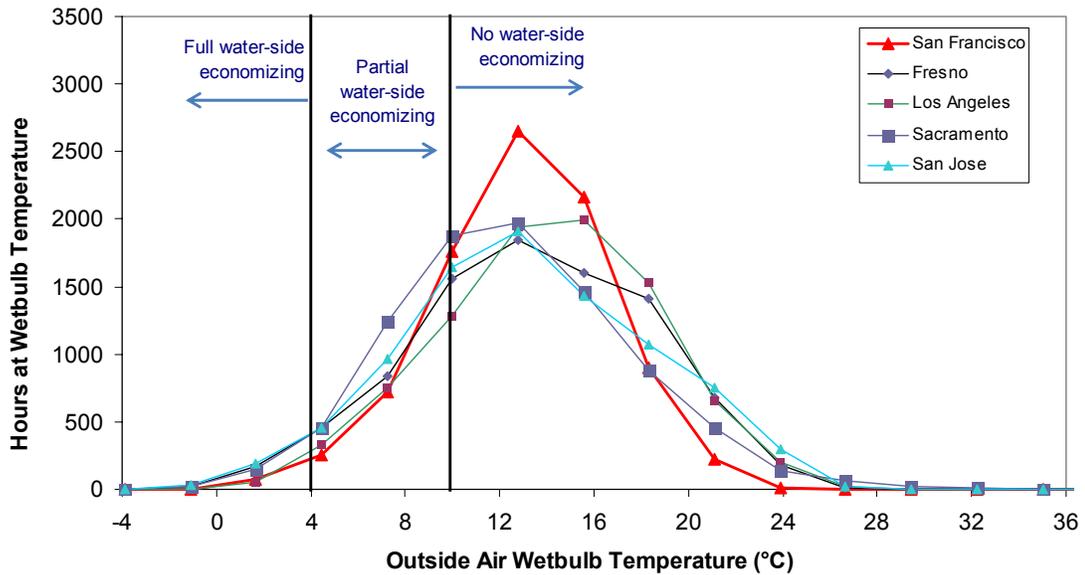


Figure 4-9. Hourly distribution of the outside air wetbulb temperature for the 8760 hours throughout the year at the five California data center locations

Furthermore, the moisture in the air increases latent cooling demand in the model and causes the chiller plant to often reach the capacity limit of the first chiller, activating a second chiller. The second chiller shares the cooling load equally with the first, resulting in a transition from one chiller at a high load factor (efficient operation) to two chillers at slightly above half the load factor (less efficient operation). The results from the water-side economizer scenario in San Francisco emphasize the need for engineers to model the hour-by-hour load, rather than just the peak load, and to size chillers and adjust setpoints to optimize chiller performance.

Figure 4-10 shows that relaxing the humidity restrictions that are commonly applied to data centers is necessary to achieve large air-side economizer energy savings in California. As the RH range is narrowed, energy use from the chiller begins to sharply increase, surpassing the equivalent baseline energy in most of the cities studied. This outcome results from a common suboptimal control algorithm. Most air-side economizers in California are controlled by a drybulb temperature setpoint that allows the economizer to remain active during periods when the drybulb temperature is lower than the setpoint, but the moisture in the air is either low enough to require humidification or high enough that the latent cooling demands increase the electricity demand from the chiller. Dewpoint and enthalpy sensors can be used to better identify when to shut down economizers to prevent excessive humidification and high latent cooling loads, but there is extra initial investment and maintenance associated with this equipment and the use of such sensors is currently not common industry practice. Humidity levels are often restricted in data centers to minimize potential IT equipment failure risks. ASHRAE's guidelines for data centers provide a "recommended" RH range between 40-55% and an

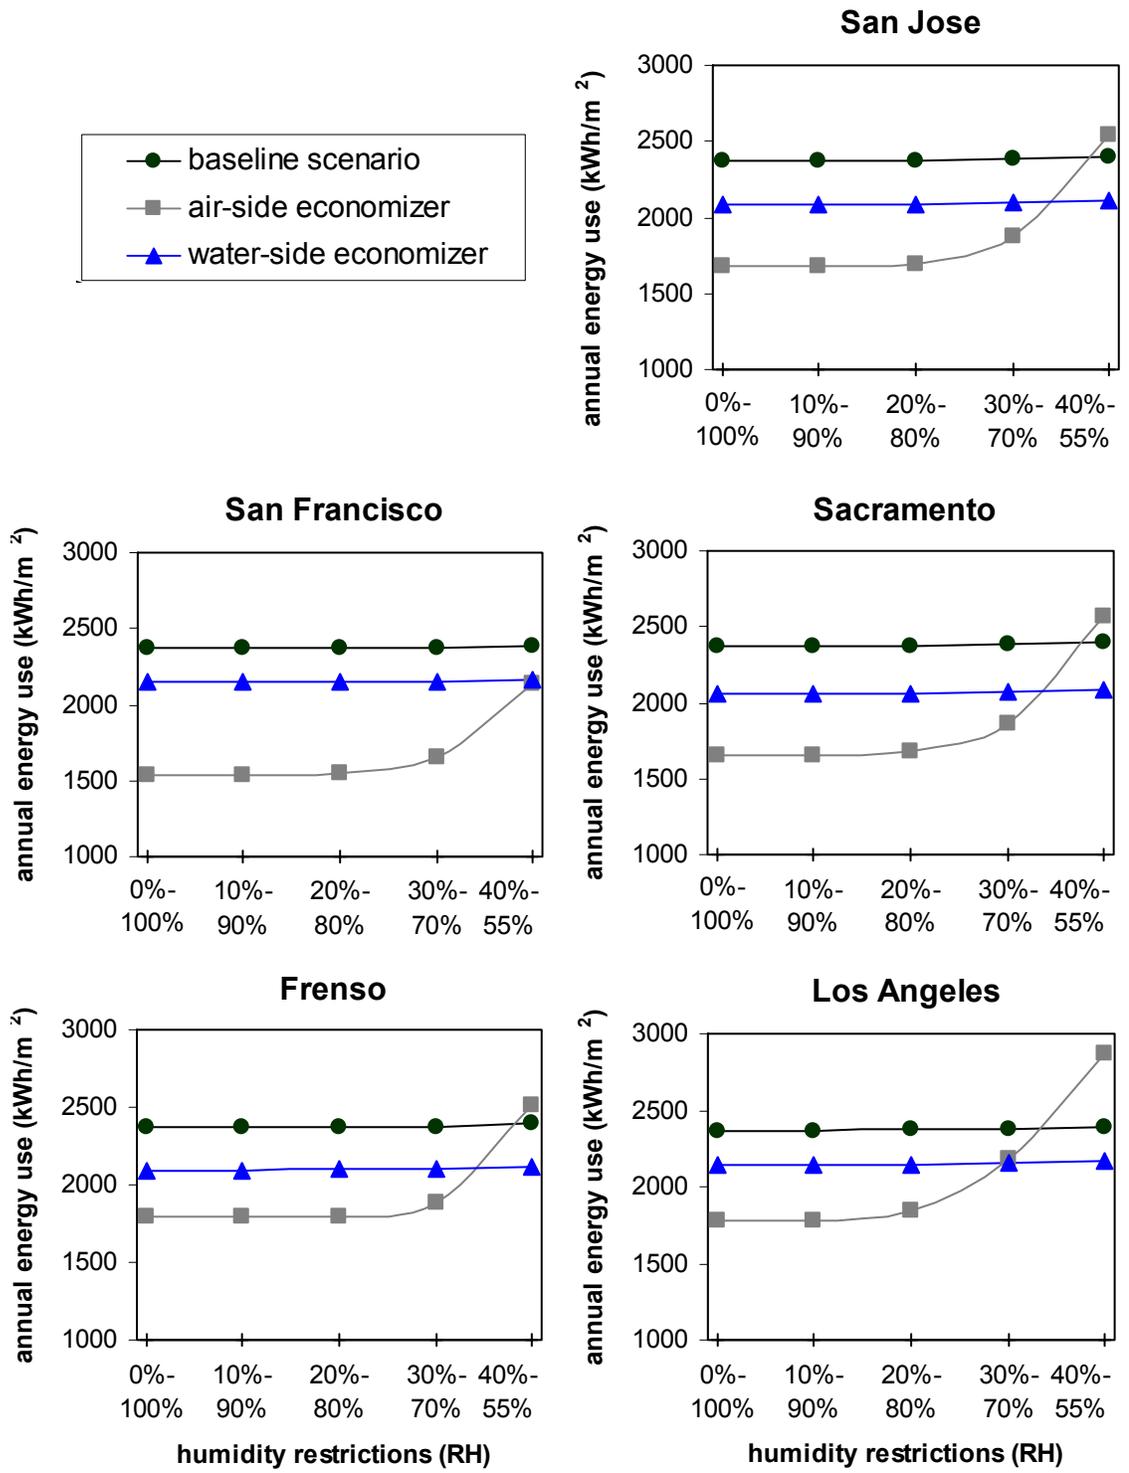


Figure 4-10: Chiller and fan energy demand separated by design scenario and humidity restrictions

“allowable” range between 20-80% (ASHRAE, 2005). There is minimal cost in applying the more conservative ASHRAE RH restrictions in conventional data center design, such as the baseline scenario in this study. The influence of humidity on server performance, however, is poorly documented and the need for humidity restrictions is increasingly being questioned (Fontecchio, 2007). The energy saving difference between adhering to ASHRAE’s *recommended* RH range versus the *allowable* RH range is substantial and warrants further investigation to determine when additional sensors are required and whether the RH restriction should be relaxed or removed altogether.

4.4. Conclusions

Employing the energy-saving measures evaluated in this chapter would require a shift in conventional data center design and operation. Various operational concerns must be addressed before widespread adoption of these technologies could be expected in data-center buildings. This analysis contributes to the informed implementation of air-side and water-side economizers by assessing the energy benefits of adopting these efficiency improvements. Air-side economizers are shown to consistently outperform water-side economizers in California, though the difference in performance varies by the climate conditions of the locations evaluated. Furthermore, the models show that without additional dewpoint or enthalpy controls conventional humidity restrictions must be relaxed or removed to substantially realize the energy benefits of air-side economizers. As the data center economy continues to rapidly grow, energy efficiency will continue to emerge as an important financial and environmental concern. The results presented here contribute to our understanding of different design implications and should assist decision makers in the implementation of energy-efficient data centers.

Chapter 5: National energy demand and potential energy savings in data centers

This chapter presents an estimate of the current (2008) national energy demand for data centers. A previously developed bottom-up modeling approach to estimate the IT energy associated with data centers is summarized. Prominent IT efficiency measures are described. IT energy is geographically distributed. Power Use Efficiency values specific to climate and data center space type are developed and applied to estimate the current total national data center energy use. The technical savings potential of economizer use is evaluated and compared along with other potential energy saving measures. Differences in PUE values between space type, location, and time of year are discussed. Greenhouse gas emissions associated with data centers are evaluated by accounting for changes in the energy mix of different geographical regions.

5.1. Introduction

In 2007, growing concern about data center energy demand and interest in energy efficiency opportunities led to Congress passing Public Law 109-431, which directed the U.S. Environmental Protection Agency (EPA) to study data center energy use, equipment, and opportunities for improving energy efficiency (U.S. Congress, 2006). One result from the EPA study was the development of a bottom-up model that can be used to estimate data center energy use and to quantify the benefits of various combinations of potential efficiency measures. The details of this model are described in Brown et al. (2007). The methods and results from Koomey (2007, 2008) provide the foundation for the Brown et al. (2007) model. Koomey's work succeeded several peer-reviewed data center energy estimates published several years earlier (Kawamoto et al., 2001; Mitchell-Jackson et al., 2002; Roth et al., 2002) and specifically improved on the analysis of Roth et al., which used aggregate server data and measured power data to estimate energy use.

This chapter builds on these previous modeling efforts, applying updated IT equipment stock data to the Brown et al. (2007) modeling procedure and integrating additional complexity into the non-IT calculations, to estimate energy use associated with data centers and the savings potential available through efficiency measures. Updated equipment stock data were gathered to provide a national IT energy demand estimate in data centers under current (2008) standard operational practices. IT energy demand is also estimated under a more energy efficient scenario that implements IT efficiency from Brown et al. (2007). The updated IT stock estimates are distributed across different geographical regions and data center building types. Energy use associated with the non-IT operations of these building is then estimated from modeled PUE values specific to climate and space type, using an approach similar to that presented in Chapter 4. Greenhouse gas emissions associated with data center operation are evaluated based on regionally specific carbon intensity values associated with electricity generation. The analysis evaluates measures that can provide significant reduction in energy demand and highlights energy-intensive aspects of data centers where improvements are especially warranted.

5.2. Data and Methods

5.2.1. IT equipment modeling procedure

The IT data center energy use estimates presented in this Chapter are calculated using a bottom-up modeling approach that was developed and applied in Brown et al. (2007). Energy use estimates have been updated to match greater data availability. This chapter section (5.2.1) summarizes the methods used to calculate IT equipment energy in the Brown et al. model (2007).

The number of operating servers in the U.S. is first estimated from market research data and then separated by class type and end-use building space type to account for differences in equipment power demand and environmental operating conditions. Server energy use is estimated after including the market penetration of prominent energy efficiency measures for servers. The energy use of data storage and network equipment in data centers is also estimated, which – together with the server energy – comprise the total IT energy demand. Storage equipment energy use is based on shipment data of enterprise or external hard disk drive (HDD) arrays. Typical network support levels associated with server and storage equipment are used to estimate the network equipment energy.

The quantity of servers installed in the United States and operating in 2008 is based on data from the International Data Corporation (IDC), which is a market research and analysis firm that annually tracks computer server sales from major computer manufacturers (IDC, 2009). IDC server data categorizes servers into three class types: volume servers (<\$25,000 per unit), mid-range servers (\$25,000 to \$500,000 per unit), and high end servers (>\$500,000 per unit). The end-use location for each class of servers in the United States is estimated based on additional IDC data for 2005 United States installed servers organized by CPU type and space type (Brown et al., 2007; Bailey et al. 2007, IDC 2007). Five space types are used to categorize potential end-use locations: server closets, server rooms, localized data centers, mid-tier data centers, and enterprise-class data centers. These space types are defined by IDC (Bailey et al., 2007) and assumptions about the major differences among these five types of spaces are listed in Table 5-1.

Table 5-1: IDC defined space type categories used for tracking computer server sales and shipments (Bailey et al., 2007; Brown et al., 2007).

Space type	Typical size	IT equipment characteristics	Typical site infrastructure system characteristics
Server closet	<19 m ²	1-2 servers No external storage	Typically conditioned through an office HVAC system. To support VOIP and wireless applications, UPS and DC power systems are sometimes included in server closets. Environmental conditions are not as tightly maintained as for other data center types. HVAC energy efficiency associated with server closets is probably similar to the efficiency of office HVAC systems.
Server room	<47 m ²	A few to dozens of servers No external storage	Typically conditioned through an office HVAC system, with additional cooling capacity, probably in the form of a split system specifically designed to condition the room. The cooling system and UPS equipment are typically of average or low efficiency because there is no economy of scale to make efficient systems more first-cost competitive.
Localized data center	<93 m ²	Dozens to hundreds of servers Moderate external storage	Typically use under-floor or overhead air distribution systems and a few in-room CRAC units. CRAC units in localized data centers are more likely to be air cooled and have constant-speed fans and are thus relatively low efficiency. Operational staff is likely to be minimal, which makes it likely that equipment orientation and airflow management are not optimized. Air temperature and humidity are tightly monitored. However, power and cooling redundancy reduce overall system efficiency.
Mid-tier data center	<465 m ²	Hundreds of servers Extensive external storage	Typically use under-floor air distribution and in-room CRAC units. The larger size of the center relative to those listed above increases the probability that efficient cooling is used, e.g., a central chilled water plant and central air handling units with variable speed fans. Staff at this size data center may be aware of equipment orientation and airflow management best practices. However, power and cooling redundancy may reduce overall system efficiency.
Enterprise-class data center	>465 m ²	Hundreds to thousands of servers Extensive external storage	The most efficient equipment is expected to be found in these large data centers. Along with efficient cooling, these data centers may have energy management systems. Equipment orientation and airflow management best practices are most likely implemented. However, enterprise-class data centers are designed with maximum redundancy, which can reduce the benefits gained from the operational and technological efficiency measures.

Table 5-2. Distribution of 2005 U.S. server stock by server class and data center space type, based on IDC server shipment and sales data (Brown et al., 2007).

Server class	Percentage of installed servers					Total
	Server closets	Server rooms	Data centers			
			Localized	Mid-tier	Enterprise-class	
Volume	17%	20%	17%	15%	30%	100%
Mid-range	0%	5%	16%	14%	65%	100%
High-end	0%	0%	16%	14%	71%	100%

Table 5-2 summarizes the server distribution by space type derived by Brown et al. (2007) and indicates that a significant fraction of U.S. servers are located in smaller sized data center rooms, which can have significantly different IT equipment and building infrastructure characteristics than larger data center buildings. The following equation was developed to express the national IT energy use, E_{IT} , when accounting for the different server class and space types defined above:

(5-1)

$$E_{IT} = \left[\sum_{c,s} (SC_{c,s} \times PSRR_{c,s}^{-1} \times ASUE_{c,s}) + (HDD \times ADUE \times (1 + f_{OH})) \right] (1 + f_{NW})$$

where for each server class c , installed in space type s , $SC_{c,s}$ is the number of installed servers (server count), $PSRR_{c,s}$ is the physical server reduction ratio associated with implementing virtualization, and $ASUE_{c,s}$ is the average server unit energy after accounting for different energy efficiency measures. The sum of this array presented in term 1 represents the total energy use associated with servers. Term 2 presents storage equipment energy, where HDD is the number of installed hard disk drives, $ADUE$ is the

average unit energy for each hard disk drive, and f_{OH} represents an overhead factor for all other auxiliary equipment associated with disk drive operation. Term 3 presents network equipment energy, which is represented by f_{NW} as a fraction of the server and disk drive energy use.

While still in the nascent stage of implementation, many new server microprocessors are designed to facilitate hardware virtualization (Brown et al., 2007). Virtualization uses software to allow multiple servers that operate at low average processor utilization levels to be replaced with a single host server that provides the same services and operates at a higher average utilization level. The result is computer demand being met through fewer physical servers (IDC, 2007). Virtualization may offer significant energy savings for volume servers because these servers typically operate at an average processor utilization level of only five to 15 percent, indicating a potentially high level of consolidation potential (Brown et al., 2007). The typical U.S. volume server will consume from 60 to 90% of its maximum system power at such low utilization levels (AMD, 2006; Bodik et al., 2006). The physical server reduction ratio (PSRR) is defined by Brown et al. (2007) is to characterize the effects of server consolidation associated with virtualization on the installed base and is defined as:

$$PSRR = \frac{\text{pre - consolidation installed server base}}{\text{post - consolidation installed server base}} \quad (5-2)$$

The $ASUE_{c,s}$, is calculated by estimating the typical server unit power demand (UPD) and then adjusting those estimates to account for the impact of prominent IT energy efficiency measures. All servers are assumed to operate constantly throughout the

year. UPD estimates for each server class, absent IT energy-efficiency measures, are extrapolated to 2008 values from Koomey (2007), which presents measured data and estimates of power use by the most popular models in each server class during years 2000-2005. The UPD values extrapolated from Koomey (2007) are assumed to represent servers absent the energy-efficiency measures evaluated in this study since these measures had insignificant market penetration prior to 2005 (Brown et al., 2007). Each server class UPD is then adjusted to account for (1) the penetration of Energy Star volume servers in the installed base, (2) increase in processor utilization due to virtualization, (3) the use of power scaling on applicable servers. These server energy efficiency measures are incorporated into the 2008 UPD estimates to calculate a weighted average server UPD as shown in Equation 5-3:

$$ASUP = \frac{(f_{ES}) \times [(f_{PS})(UPD_{ES,PS}) + (1 - f_{PS})(UPD_{ES})] + (1 - f_{ES}) \times [(f_{PS})(UPD_{PS}) + (1 - f_{PS})(UPD_B)]}{(f_{ES}) \times [(f_{PS})(UPD_{ES,PS}) + (1 - f_{PS})(UPD_{ES})] + (1 - f_{ES}) \times [(f_{PS})(UPD_{PS}) + (1 - f_{PS})(UPD_B)]} \quad (5-3)$$

where for each server class c , installed in space type s , $ASUP$ is the average server unit power use; f_{ES} is the fraction of Energy Star rated servers across the installed server base; f_{PS} is the fraction of servers with power scaling utilization across the installed server base; $UPD_{ES,PS}$ is the average power demand for Energy Star servers with power scaling activated; UPD_{ES} is the average power demand for energy star servers without power scaling activated; UPD_{PS} is the average power demand for non-Energy Star servers with power scaling activated. UPD_B is the average server power for non-energy star servers without power scaling activated, which represents servers absent of the energy efficiency measures as extrapolated from Koomey (2007), and is estimated as 235, 789, and 9,292

W/server unit for volume, mid-range, and high-end servers, respectively. The calculated ASUP is converted to the ASUE value applied in Equation 5-1 by assuming constant server operation throughout the year (i.e. an ASUP of 1 kW would correspond to an ASUE of 8,760 kWh).

While virtualization reduces overall computational energy use by meeting computing demand with fewer physical servers, the processor utilization level of the remaining host servers will increase due to running multiple virtual servers as well as a small processor utilization overhead associated with virtualization software. Industry data (AMD, 2006) regarding the relationship between server energy use and processor utilization were used to account for the energy impacts of power scaling activation as well as the increased host processor utilization due to virtualization. Figure 5-1 shows this trend in energy use to processor utilization for servers both with and without power scaling capabilities activated. Server power scaling represents dynamic frequency and voltage scaling features, which allow microprocessor frequency or voltage to ramp up or down to better match the computational demands (Brown et al., 2007). This server power scaling decreases microprocessor activity when utilization is low, which reduces energy consumption and heat dissipation. Frequency and voltage scaling are done automatically, and constantly adjust to changes in computational demand, continuously minimizing processor energy consumption. In the absence of virtualization, the average processor utilization for volume and midrange servers is assumed to be 10% and 20%, respectively, based on estimates compiled from industry experts (Brown et al., 2007). The relative difference between the two trends depicted in Figure 5-1 is used to estimate the energy savings of activated power scaling for a given processor utilization. For

example, Figure 5-1 indicates that a processor utilization of 10% without power scaling corresponds to an average system power of about 70% of maximum power, which for volume servers is assumed to represent the conditions for the estimated UPD_B of 235 W/server. When power scaling is activated, the average system power is reduced to about 55% of maximum power for a processor utilization of 10%. The *relative* difference between 70% system power and 55% system power corresponds to 20% decrease in system power (i.e. $1 - 55/70$), which is then used to estimate the UPD_{PS} as 80% of the UPD_B , so that a UPD_B of 235 W/server would indicate a UPD_{PS} of about 188 W/server. Activated power scaling is assumed to be a default property of high-end servers and therefore already incorporated into the baseline UPD of this server class.

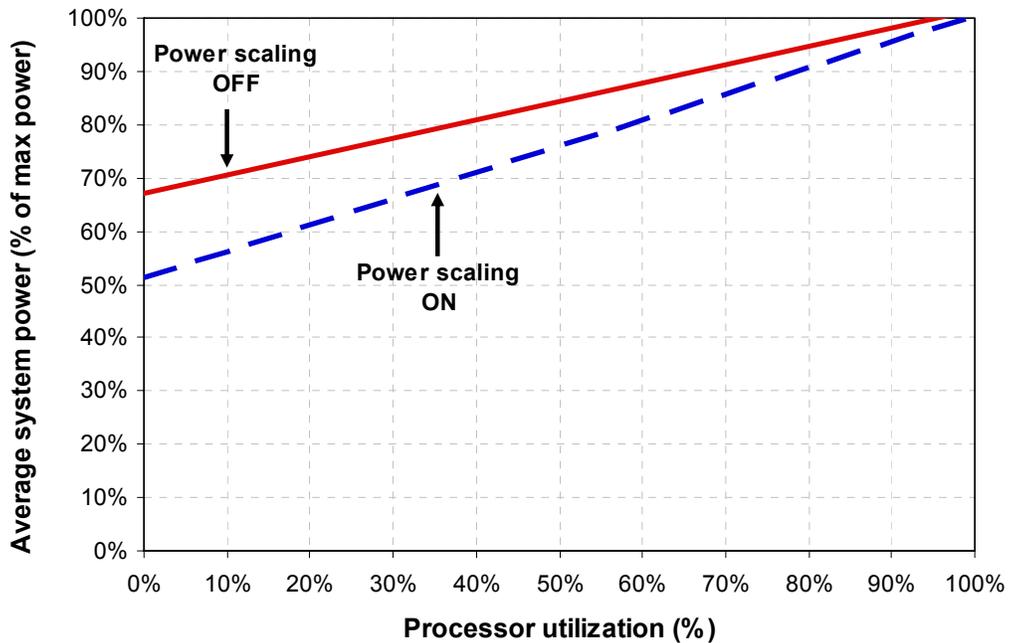


Figure 5-1. Relationship between processor utilization and system power demand, both with and without power scaling (management) activated (Brown et al., 2007).

Data presented in Figure 5-1 are also used to adjust volume server UPD values to account for changes in processor utilization due to virtualization. For volume servers subject to virtualization, the average server processor utilization level after server reduction, represented as U , is calculated as:

$$U = (U_o * PSRR) * (1+H) \quad (5-4)$$

where U_o is the average processor utilization prior to virtualization and H accounts for the software overhead associated with running virtualization software on the host machine. The virtualization overhead is assumed to be 10% based on industry feedback (Brown et al., 2007), but can vary depending on the percentage of installed volume server base that serve as host servers. For example, assuming a PSSR of 4 and an average processor utilization for volume servers prior to virtualization, U_o , of 10%, the average server processor utilization level after server reduction, U , would be 50%. According to Figure 5-1, this increased processor utilization level corresponds to an average system power of about 85% of maximum power. The relative increase in average system power, from 70% to 85% (i.e. about a 20% increase) is then used estimate the increase in UPD_B from 235 to about 280 W/server due to virtualization. In other words, in this example virtualization allows a single volume server consuming 280 watts to replace the operation of four servers each consuming 235 watts. For mid-range and high-end servers, it was assumed that virtualization would not be applicable since the average processor utilization level is already maximized for much of this equipment type (Brown et al., 2007).

The increasing penetration of Energy Star volume server models will tend to decrease the average UPD across the entire volume server base in the United States. Energy Star servers include more efficient microprocessors, cooling fans, and power supplies, which combined can account for 50% to 80% of total server energy use (Eubank et al., 2003; Patterson et al. 2006). The Environmental Protection Agency (EPA) developed the first Energy Star requirements for computer servers in 2009 and this rating represents servers that are, on average, 30% more energy efficient than similar volume server models across a range of typical processor utilization levels (EPA, 2009).

Energy Star has only been developed for the volume server market, an appropriate prioritization given their enormous market share, the large percentage of typical data center energy use they represent, and the trend toward increasing their power density (e.g., blade servers). Recent, more efficient, low-voltage multiple-core microprocessors have been offered almost exclusively for the volume server market (Brown et al., 2007). Moreover, mid-range and high-end servers already typically employ high-efficiency power supplies (Koomey, 2007). For these reasons, the UPD savings associated with the energy efficient Energy Star servers are applied only to volume servers and not to mid-range and high-end servers in this study.

Energy used by storage devices and network equipment are included in this analysis to present a more complete picture of the total energy used by IT equipment and associated power delivery and cooling systems. The energy use of external storage devices can vary widely depending on the need for digital storage in a particular data center. Estimates of external storage devices were developed using data on the installed U.S. base and energy use of external HDD storage devices (Brown et al., 2007). Table 5-

3 summarizes the distribution of these estimates by space type. This storage distribution assumes that the total enterprise storage system energy use could be allocated across localized, mid-tier, and enterprise-class data centers in a proportional manner, based on the installed number of servers in each respective space type. The number of external storage devices installed in U.S. server closets and server rooms is assumed to be negligible, as the necessary storage capacity in these space types is assumed to be typically provided by internal server HDDs. The ADUE is assumed to be 14 watts per drive and the overhead factor to account for all other auxiliary equipment associated with storage disk drive operation, f_{OH} , is assumed to be 100%, based on estimates provided by storage industry experts (Brown et al., 2007).

Table 5-3. 2007 U.S. stock of computer storage units distributed by data center space type, based on Seagate Technology external HDD shipment and sales data (Brown et al., 2007).

US Installed Base External HDDs (millions)	
Server closet	0
Server room	0
Localized DC	4.4
Mid-tier DC	4.0
Enterprise-class	8.0
Total	16.4

Minimal public data are available on the average energy use of network equipment in data centers and other server installations in the U.S. Estimates of the energy use of network equipment in server rooms, localized data centers, mid-tier data centers, and enterprise-class data centers are made by assuming that current network

equipment consumes on average 8 watts per port and that the typical data center will have, on average, three installed network ports per installed volume server (Brown et al., 2007). This corresponds to a network equipment energy factor, f_{NW} , of 10%, which agrees well with industry IT equipment energy use breakdown data for typical data centers (Dell, 2007). The network equipment energy factor was reduced to 5% for server closets based on the expectation that fewer ports (one to two per server) would be required in server closets because these spaces are typically designed for small workgroup support (e.g., file and print server applications). This estimate, however, excludes the network equipment energy use attributable to office equipment (e.g., personal computers) that may be connected to network equipment in server closets.

5.2.2. IT equipment energy savings estimation

The potential of energy savings from IT equipment is estimated by adjusting the operational parameters for servers in the modeling equations from standard operations to a more energy-efficient scenario. Assumed differences between standard operations and the energy-efficient scenario are presented in Table 5-4. The current market penetration of Energy Star servers is estimated at 5%, based on industry data compiled during the development of the server ratings (EPA, 2009). The fraction of volume and mid-range servers that currently utilize dynamic frequency and voltage scaling is assumed to be 10%, based on industry estimates of the current utilization rate for microprocessor power-management features (Brown et al., 2007). In the energy-efficient scenario, all applicable servers are assumed to meet Energy Star efficiency standards and utilize power scaling. Under standard operations, a PSRR of 1 (indicating no virtualization) is assumed for the

2008 installed base estimate of volume servers. Given the nascent state of virtualization implementation (Brown et al., 2007), the impact of virtualization on the average utilization rate of volume servers in the installed base is assumed to be negligible. The amount of virtualization available in a data center is limited by the maximum desired utilization level of the host volume servers. While host volume server utilization can theoretically be maintained at 100%, a maximum utilization of 50% is assumed in the energy-efficient scenario (corresponding to an average PSRR of 4), since data center operators prefer to reserve some processing capacity in case of temporary demand increases. Additionally, virtualization is applied to only half of the volume servers located in server closets (corresponding to an average PSRR of 2), based on the expectation that many server closets will only host one local workgroup server and are thus not candidates for virtualization.

Table 5-4. Estimated market penetration of IT efficiency measures and calculated PUE values for the standard operations and energy efficient scenarios.

	Standard Operation Scenarios		Energy Efficiency Scenarios	
	Current Practices	Baseline	Economizer	Economizer Plus
EnergyStar	Represents 5% of volume servers		Represents 100% of volume servers	
Power Scaling	Enabled on 10% of applicable servers		Enabled on 100% of applicable servers	
Virtualization	No virtualization relative to current stock		PSRR of 4 to 1 (2 to 1 for server closets)	
PUE Values	Scaled to 2.0 from Baseline PUE values	Data center without economizer use	Economizer use with Baseline temp/humidity	Economizer use with expanded temp/humidity

5.2.3. Non-IT equipment modeling procedure

The energy use of non-IT data center components is estimated in this chapter to obtain a more complete understanding of the total energy demand associated with data center operation. Energy from the non-IT components of the data center include compressor cooling, air-handling, power distribution losses, and building lighting. National energy use associated with total data center operation is calculated using the equation

$$E_{\text{total}} = \sum_{s,r} (E_{IT})_{s,r} (PUE)_{s,r} \quad (5-5)$$

where E_{total} is the total operational energy used for data centers, $(E_{IT})_{s,r}$ is the operational energy use associated with IT equipment for space type s in region r . Total data center building energy use is calculated as a function of the IT energy by using the PUE ratio of total data center energy demand to the direct energy demands of the IT equipment, which is typically calculated to represent the energy efficiency of the data center (see Chapter 4). When estimating the energy demand of non-IT data center components, the PUE ratio can vary widely among individual buildings depending on the non-IT equipment configurations and efficiencies, time of year, and local climate (Brown et al., 2007). $PUE_{s,r}$ represents an annual average energy use performance of the non-IT equipment associated with data center space type s and climate region r .

The total calculated IT equipment energy associated with each space type s is equally distributed into five different representative U.S. cities to address the effect of climatic variations on building operation among prominent data center locations (Figure 5-2). The five cities chosen, San Francisco, Seattle, Chicago, Dallas, and Richmond, are based on the analysis of two data sets. First, commercial buildings identified to have

significant data center activity from the Commercial Building Energy Consumption Survey (CBECS, 2003) are analyzed to disaggregate the number of installed computer servers documented in that data set by regional census division (Table 5-5).

Second, a list of U.S. metropolitan areas with large concentration of existing data centers previously compiled from U.S. Department of Energy data (Brown et al., 2007) is used to identify specific cities with significant data center activity (Table 5-6). The cities from the Brown et al. (2007) list are then used to represent each of the CBECS regions.

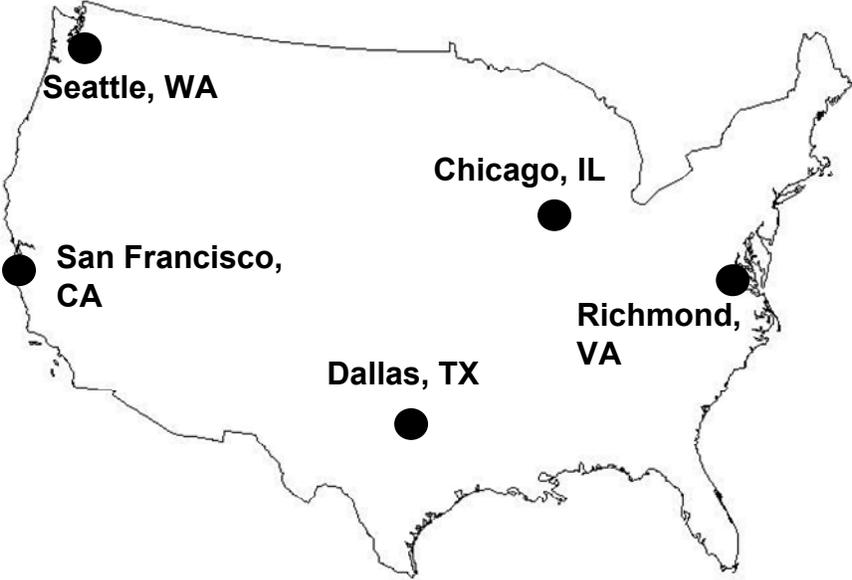


Figure 5-2: Five cities used to represent different climate regions in PUE modeling. Each location represents a climatically distinct region with significant data center activity. Together, these five cities are suitable for representing the majority of all data center activity in the U.S.

The chosen cities each represent regions of the country with an approximately equivalent number of computer servers and together represent the majority of all computer servers documented in the CBECS data. Space type and climatic region specific PUE values were developed using the modeling procedure outlined in Chapter 4. Separate models were developed to account for the different IDC-defined space type described in Table 5-1, and then each of these models were evaluated using annual hourly climate data for each of these cities in Figure 5-2.

As outlined in Table 5-4, space type specific PUE values were developed for two standard operation scenarios that did not include air-side economizer use and two energy efficient scenarios with an economizer installed in the mechanical system. Water-side economizers were not modeled in this analysis since results from Chapter 4 revealed that air-side economizers provided significantly greater energy savings than water-side economizers. The PUE values calculated for the standard operation scenarios were classified as “Baseline” PUE values and “Current Practices” PUE values. The Baseline PUE values represent the actual PUE value generated from the model for data centers without economizers. For the Current Practices PUE values, the Baseline PUE values are scaled to correspond to an overall average PUE value of 2.0 when weighted by the space type distribution presented in Table 5-3.

Table 5-5. Regional allocation of computer servers located in buildings identified to have significant data center activity (derived from CBECS, 2003)

Census Division	Percentage
New England	3.2%
Middle Atlantic	15.7%
East North Central	17.8%
West North Central	5.2%
South Atlantic	16%
East South Central	2.3%
West South Central	6.6%
Mountain	5.6%
Pacific	27.6%

Table 5-6. U.S. metropolitan areas with largest concentration of existing data centers (Brown et al., 2007)

U.S. Metropolitan Area
New York City / Northern New Jersey
San Francisco Bay Area CA
Chicago IL
Dallas TX
Washington DC area
Austin TX
Los Angeles CA
Atlanta GA
Miami FL
Seattle WA
St. Louis MO
Denver CO
Boston MA
Houston TX
San Antonio TX
Phoenix AZ
Kansas City MO
Sacramento CA

The current stock of data centers is generally assumed by the data center industry to have an average PUE ratio of approximately 2.0 (Brown et al., 2007; Koomey, 2007). This PUE estimate is supported by the small amount of empirical data available from Greenberg et al. (2006) and Belady and Malone (2007), by an industry consensus reported in Brown et al. (2007), as well as by recent industry survey data gathered by the U.S. Environmental Protection Agency (Sullivan, 2009). Modeled baseline data center energy estimates generally indicate a PUE performance better than 2.0. For example, a 1.55 baseline PUE ratio was estimated for the analysis reported in Chapter 4. While this lower baseline PUE ratio is partially due to the relatively efficient equipment selected in that model, part of the discrepancy between the modeled and measured PUE values may be due to inefficiencies in operation and airflow management that are not captured in the model. The Current Practices PUE values are developed to match the best available data on current data center performance while still providing variation between data center space type and climate region. Since much of the infrastructure equipment associated with the economizer use is not compatible with closet data centers, a PUE of 2.0 is used throughout the standard operation and energy efficient scenarios of this space type.

The PUE values calculated for the energy scenarios were classified as “Economizer” PUE values and “Economizer Plus” PUE values. The Economizer PUE values represent data centers using economizers while maintaining the indoor temperature and humidity at the levels in Baseline PUE design. Economizer implementation has been hindered by IT equipment reliability concerns associated with exposing IT equipment to large volumes of outside air. Results from Chapter 2 showed that economizer use can increase indoor particle concentration by an order of magnitude.

While Chapter 3 showed that improved filtration can eliminate this particle concentration increase with minimal additional energy use, reliability concerns could still remain since other pollutant levels (e.g. gaseous contaminants) may still be relatively higher during economizer use. The Economizer PUE values are calculated to determine the energy saving benefits associated with accepting or addressing this potential reliability risk and allowing large volume of outside air into the data center space.

Results from Chapter 4 showed an increase in economizer energy savings as wide ranges of relative humidity were allowed in data center. Allowing higher temperatures in the data center could also increase energy savings since economizers could operate for more hours of the year. A wider range of relative humidity and increased operating temperature could potentially impact IT equipment reliability, through the extent of potential disruption to typical data center operations is not clearly understood. The Economizer Plus PUE values are calculated to better understand if the energy savings benefits from these expanded environmental conditions warrant exploring strategies that would allow exposing IT equipment to broader ranges of indoor temperature and humidity while still maintaining acceptable IT equipment reliability. In calculating the Economizer Plus PUE values, the model was evaluated under three different humidity ranges and three temperature setpoints to determine how incremental changes to these parameters would affect economizer performance. The Economizer Plus PUE values represent a humidity range of 1-100% (no restriction) and temperature setpoints of 17.8 °C (64°F) / 28.9 °C (84 °F) supply/return.

The size and efficiencies of mechanical system components assumed in the models are based on a combination of manufacturer design guidelines, fundamental

HVAC sizing equations, and observations gained through professional experience that have been documented and applied in previous modeling analyses (Rumsey Engineers, 2008; Rumsey Engineers, 2005). The capital costs required for the mechanical equipment associated with the Economizer and Economizer Plus scenarios are not addressed in this chapter, but previous analysis has shown that the financial gains from energy savings in data centers greatly outweigh capital cost, with payback periods on the order of months (Rumsey Engineers, 2008; Rumsey Engineers, 2005).

Closet data centers were assumed to be impractical candidates for installing an economizer and no PUE was calculated for this space type. Closet data center spaces are defined as only including 1-2 computer servers and being thermally conditioned solely through the central HVAC system of a building without any dedicated cooling equipment. Additionally, closet data centers may be located within the interior portion of an office building, making the ductwork needed for a dedicated economizer to the closet prohibitively costly for the relatively small internal heat load. A PUE of 2.0, the estimated performance of current data centers, was used to represent the non-IT equipment performance in closet data centers.

Server rooms are modeled as a 23 m² (250 ft²) room with an internal IT load of 0.43 kW/m² (40 W/ft²). No underfloor air distribution or humidity controls are assumed to be present owing to the small size of the room. In the standard operation scenarios, the room is cooled with a single dedicated air-cooled direct expansion (DX) CRAC unit placed within the room. The CRAC unit contains a constant speed fan that draws air through the unit directly from the room without the use of any ductwork. This cooling system is converted to a dedicated outdoor package DX air-conditioning unit with an on-

board air-side economizer in the energy-efficient scenarios. A constant-speed fan in the outdoor unit supplies air to the room through a ceiling duct system and a second fan exhausts air from the room during economizer activity. The supply duct increases the pressure drop experienced by the supply fan relative to the CRAC system, though the exhaust fan experiences a much lower pressure drop, since little or no ducting would be installed downstream of that fan. Table 5-7 presents details of the mechanical equipment used in the model for the server room space type.

The localized data center space type is modeled as a 47 m² (500 ft²) room with an internal IT load of 0.65 kW/m² (60 W/ft²). The standard operation scenarios use three air-cooled DX CRAC units with constant speed fans placed in the data center that supply air through an underfloor plenum. The CRAC units are equipped with active humidity control and include electric humidifiers. In the energy-efficient scenarios, the CRAC units are replaced with two outdoor DX package air-conditioning units with on-board air-side economizers. Air is ducted from the outdoor air handlers to an underfloor plenum using constant speed fans. The package units are equipped with active humidity control, but use more efficient adiabatic humidifiers to address the wider range of humidity that the system must condition owing to the high rate of outside air entering the data center. Adiabatic humidifiers use the heat from incoming air to vaporize water, while conventional electric humidifiers require additional energy to vaporize water through electric resistance heating. Table 5-8 presents details of the mechanical equipment used in the model in the localized data center space type.

Table 5-7. Building and mechanical design parameters of the server room space type model for the standard operation and energy efficiency scenarios

<i>General Characteristics</i>	Scenario	
	Baseline	Energy Efficient
Floor area (m ²)		23.2
IT load density (kW/m ²)		0.43
Total load (kW)		10
<i>Fan System</i>		
Supply airflow rate (m ³ /s)		2.64
<i>Supply Air Delivery</i>		
Static pressure (kPa)	0.4	0.5
Number of fans	1	1
Fan size (kW)	2.2	3.7
Fan efficiency	44.0%	50.0%
Drive efficiency	95.0%	95.0%
Motor efficiency	86.5%	87.5%
VFD efficiency	n/a	n/a
<i>Exhaust Air Delivery</i>		
Static pressure (kPa)	n/a	0.2
Number of fans	n/a	1
Fan size (kW)	n/a	1.5
Fan efficiency	n/a	34.5%
Drive efficiency	n/a	95.0%
Motor efficiency	n/a	85.5%
VFD efficiency	n/a	n/a
<i>Cooling System</i>		
DX unit size (kW)		19.0
Number of DX units		1
DX unit efficiency (EER)		10.3

Table 5-8. Building and mechanical design parameters of the localized data center space type model for the standard operation and energy efficiency scenarios

<i>General Characteristics</i>	Scenario	
	Baseline	Energy Efficient
Floor area (m ²)		46.5
IT load density (kW/m ²)		0.65
Total load (kW)		30
<i>Fan System</i>		
Supply airflow rate (m ³ /s)		7.9
<i>Supply Air Delivery</i>		
Static pressure (kPa)	0.4	0.5
Number of fans	3	2
Fan size (kW)	5.6	7.5
Fan efficiency	47.0%	55.6%
Drive efficiency	95.0%	95.0%
Motor efficiency	89.5%	90.2%
VFD efficiency	n/a	n/a
<i>Exhaust Air Delivery</i>		
Static pressure (kPa)	n/a	0.2
Number of fans	n/a	2
Fan size (kW)	n/a	3.7
Fan efficiency	n/a	50.0%
Drive efficiency	n/a	95.0%
Motor efficiency	n/a	87.5%
VFD efficiency	n/a	n/a
<i>Cooling System</i>		
DX unit size (kW)	56.3	63.3
Number of DX units	3	2
DX unit efficiency (EER)	9.5	9.7

A 232 m² (2,500 ft²) room with an IT load of 0.86 kW/m² (80 W/ft²) is used in the model to represent mid-tier sized data centers. The cooling system in the standard operation scenarios consists of seven CRAC unit air-handlers placed on the data center floor that receive chilled water from two air-cooled chillers located outside of the data center. Constant speed fans in the CRAC units supply conditioned air through an underfloor plenum. A typical minimum ventilation requirement of 2.7 m³/h per m² of floor space (ASHRAE, 2005) as well as humidity controls are provided through a separate makeup air-handler that supplies only outside air and uses an electric humidifier. The separate air-handler also receives chilled water from the air-cooled chiller water plant. The makeup air-handler is balanced with an exhaust fan of equal size. In the energy-efficient scenarios, four air-handling units with an on-board air-side economizer are placed outdoors and receive chilled water from two air-cooled chillers. The outdoor air-handlers use variable speed fans to supply air to the data center via ducts through an underfloor plenum and are equipped with active humidity controls that use adiabatic humidifiers. An exhaust fan in each of the outdoor air-handlers is activated when the economizer is used to balance the influx of outdoor air. Table 5-9 presents details of the mechanical equipment used in the model for the mid-tier data center space type.

The enterprise data center space type is represented in the model as a 465 m² (5,000 ft²) data center with a 1.1 kW/m² (100 W/ft²) internal IT load. The cooling and air-handling systems for the standard operation and energy-efficient scenarios match the respective cooling and air-handling systems of the mid-tier data center, except that more air-handling units are used to meet the increased internal heat load. The standard operation scenarios use 17 CRAC unit air-handlers placed on the data center floor while

the energy-efficient scenarios use 10 air handling units placed outdoors with on-board air-side economizers. Two water-cooled, rather than air-cooled, chillers with a cooling tower system are assumed to be deployed for this larger space type for all scenarios. In the energy-efficient scenarios, the economizer system is assumed to be custom designed for the enterprise data center, resulting in a lower supply-fan pressure resistance within the ducted air-handling system. Table 5-10 presents details of the mechanical equipment used in the model for the enterprise data center space type.

Along with the mechanical systems, other significant non-IT energy demands include lighting energy and UPS energy losses. Lighting energy density is assumed to be about 10 W/m^2 , which is a typical value for many commercial buildings and data centers (Rumsey Engineers, 2008). The lighting energy density is constant for both the standard practice and the energy efficiency scenarios across all data center space types. Losses from the UPS systems are estimated from empirical data of UPS efficiency relative to the load factor (Greenberg et al., 2006), which indicates an efficiency of approximately 85% for any load factor greater than 0.2. The UPS load factor is calculated by first selecting UPS modules that provide adequate redundancy specific to each space type and then determining the ratio of IT load of total UPS capacity. Both the lighting energy and UPS losses are added to the overall mechanical cooling load in the model.

Table 5-9. Building and mechanical design parameters of the mid-tier data center space type model for the standard operation and energy efficiency scenarios

<i>General Characteristics</i>	Scenario	
	Baseline	Energy Efficient
Floor area (m ²)		232
IT load density (kW/m ²)		0.86
Total load (kW)		200
<hr/> <i>Fan System</i> <hr/>		
Supply airflow rate (m ³ /s)		56.2
<i>Supply Air Delivery</i>		
Static pressure (kPa)	0.4	0.5
Number of fans	7	4
fan size (kW)	7.5	11.2
Fan efficiency	60.0%	58.7%
Drive efficiency	95.0%	95.0%
Motor efficiency	91.7%	90.2%
VFD efficiency	n/a	98.0%
<hr/> <i>Makeup Air Delivery</i> <hr/>		
Makeup airflow rate (m ³ /s)	0.2	n/a
Static pressure (kPa)	0.7	n/a
Number of fans	1	n/a
fan size (kW)	1.5	n/a
Fan efficiency	34.5%	n/a
Drive efficiency	95.0%	n/a
Motor efficiency	85.5%	n/a
VFD efficiency	n/a	n/a
<hr/> <i>Exhaust Air Delivery</i> <hr/>		
Exhaust airflow rate (m ³ /s)	0.2	n/a
Static pressure (kPa)	0.2	0.2
Number of fans	1	4
fan size (kW)	0.7	3.7
Fan efficiency	14.5%	50.0%
Drive efficiency	95.0%	95.0%
Motor efficiency	80.0%	87.5%
VFD efficiency	n/a	98.0%
<hr/> <i>Cooling System</i> <hr/>		
Chiller capacity (kW)		351.7
Number of air-cooled chillers		2
Avg chiller efficiency (kW/ton)		0.82

Table 5-10. Building and mechanical design parameters of the enterprise data center space type model for the standard operation and energy efficiency scenarios

<i>General Characteristics</i>	Scenario	
	Baseline	Energy Efficient
Floor area (m ²)		465
IT load density (kW/m ²)		1.1
Total load		500
<i>Fan System</i>		
Supply airflow rate (m ³ /s)		136.2
<i>Supply Air Delivery</i>		
Static pressure (kPa)	0.4	0.5
Number of fans	17	10
fan size (kW)	7.5	14.9
Fan efficiency	60.0%	60.8%
Drive efficiency	95.0%	95.0%
Motor efficiency	91.7%	91.0%
VFD efficiency	n/a	98.0%
<i>Makeup Air Delivery</i>		
Makeup airflow rate (m ³ /s)	0.4	n/a
Static pressure (kPa)	0.7	n/a
Number of fans	1	n/a
fan size (kW)	1.5	n/a
Fan efficiency	34.5%	n/a
Drive efficiency	95.0%	n/a
Motor efficiency	85.5%	n/a
VFD efficiency	n/a	n/a
<i>Exhaust Air Delivery</i>		
Exhaust airflow rate (m ³ /s)	0.4	n/a
Static pressure (kPa)	0.2	0.2
Number of fans	1	10
fan size (kW)	0.7	7.5
Fan efficiency	14.5%	55.6%
Drive efficiency	95.0%	95.0%
Motor efficiency	74.0%	90.2%
VFD efficiency	n/a	0.98
<i>Cooling System</i>		
Chiller capacity (kW)		879.2
Number of water-cooled chillers		2
Avg chiller efficiency (kW/ton)		0.63

5.3. Results and Discussion

5.3.1. National energy estimates

Figure 5-3 compares 2008 data center energy use estimates for the standard operation and energy efficient scenarios. The Current Practices scenario estimates that data centers consumed nearly 70 billion kWh of energy in 2008, with volume servers and HVAC components contributing the majority of this demand. The total represents a 56% increase in data center energy use compared to the 2005 estimate (Koomey, 2007) and matches well a published prediction of 2008 data center energy use (Brown et al., 2007). Applying the Baseline PUE values instead of the Current Practices PUE values reduces the total data center energy use to approximately 62 billion kWh/y by removing about 8 billion kWh/y from the estimated energy demand for non-IT data center components. The combined IT and HVAC efficiency measures included in the energy-efficient scenarios reduce the data center energy use to approximately 23 billion kWh/y for the Economizer scenario, and relaxing the humidity and temperature settings saves an additional 2 billion kWh/y in the Economizer Plus scenario. These estimated energy savings indicate a potential savings of between 40 and 50 billion kWh/y. Much of the savings is realized through better efficiencies in volume servers, which is consistent with the findings from Brown et al. (2007). The energy savings from volume servers is the result of IT efficiency measures such as virtualization and Energy Star that exclusively target the large energy demand of this server class.

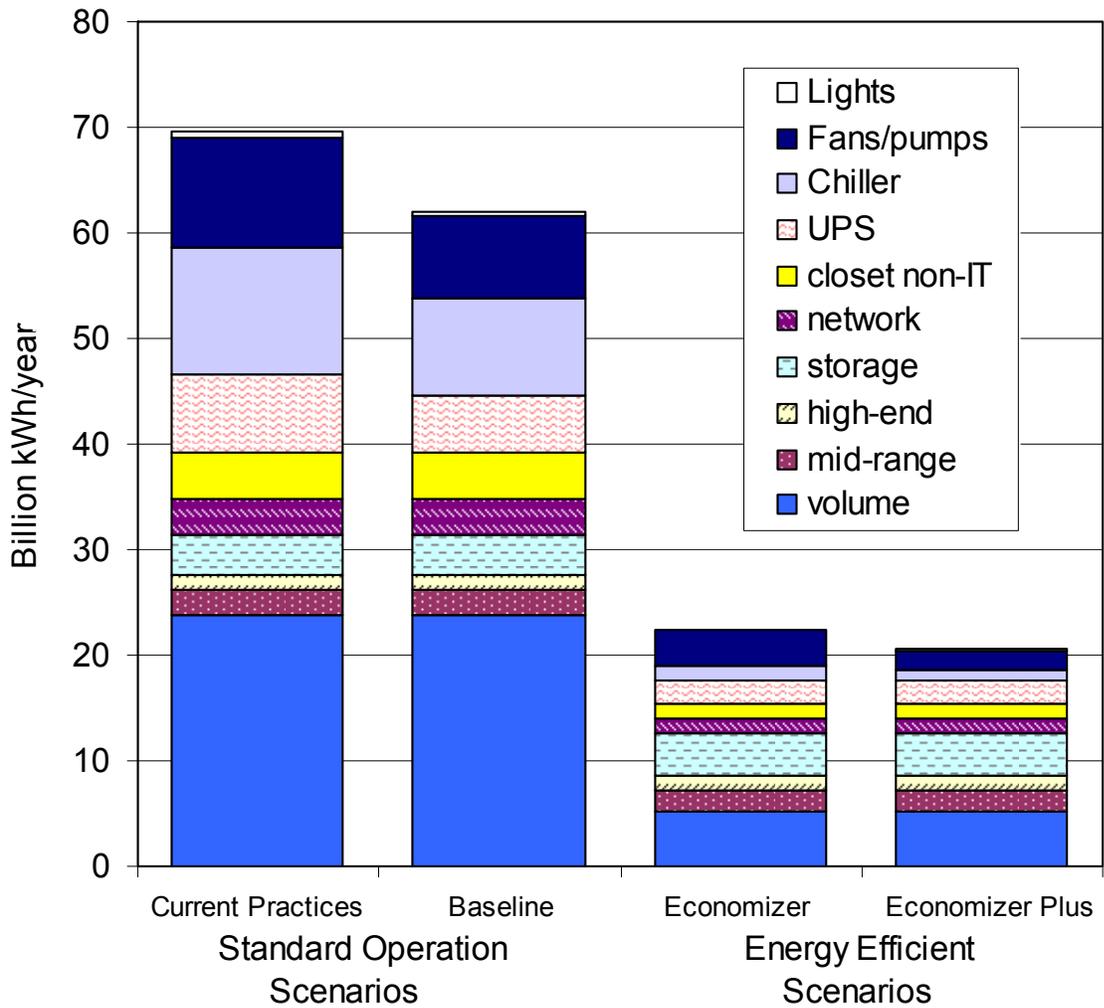


Figure 5-3. Comparison of national data center energy use under the standard operation and energy efficient scenarios. The difference between the two scenarios represents the technical potential energy savings available from implementing the identified IT and cooling efficiency measures into national data center operations.

Table 5-11. Current (2008) energy and energy efficiency potential of national data center energy use, by space type and equipment component.

U.S. data center energy use (billion kWh/year) by space type							
	Server Closet	Server Room	Localized	Mid-tier	Enterprise	Total	% of Total
Current Practices							
Volume	4.1	4.7	4.0	3.7	7.2	23.7	34%
High-end	0	0	0.2	0.2	1.0	1.4	2%
Mid-range	0	0.1	0.4	0.4	1.6	2.5	4%
Storage	0	0.0	1.0	0.9	1.9	3.9	6%
Network	0.2	0.7	0.6	0.6	1.3	3.4	5%
Server closet non-IT	4.4					4.4	6%
UPS losses		1.2	1.5	1.4	3.2	7.3	11%
Chiller		2.2	2.9	2.5	4.5	12.1	17%
Fan/pumps		2.1	2.2	1.9	4.1	10.3	15%
Lights		0.2	0.1	0.1	0.2	0.6	1%
Total	8.7	11.2	13.1	11.6	24.9	69.6	100%
% of Total	13%	16%	19%	17%	36%	100%	
Economizer Plus							
Volume	1.4	0.9	0.8	0.7	1.4	5.2	25%
High-end	0	0	0.2	0.2	1.0	1.4	7%
Mid-range	0	0.1	0.3	0.3	1.3	2.1	10%
Storage	0	0.0	1.0	0.9	1.9	3.9	19%
Network	0.1	0.2	0.3	0.2	0.6	1.4	7%
Server closet non-IT	1.5					1.5	7%
UPS losses		0.2	0.5	0.4	1.1	2.3	11%
Chiller		0.1	0.3	0.2	0.3	1.0	5%
Fan/pumps		0.4	0.3	0.3	0.7	1.7	8%
Lights		0.03	0.04	0.03	0.1	0.2	1%
Total	2.9	2.0	3.8	3.3	8.4	20.5	100%
% of Total	14%	10%	19%	16%	41%	100%	

The savings potential is also seen in the HVAC energy demand, resulting from a combination of air-side economizer use reducing the PUE and energy-efficient IT measures reducing the overall heat load. As expected, energy demand remains constant through all scenarios for the data center components that are unable to benefit from the energy efficient measures identified in this study, resulting in an increased relative contribution to the total data center energy use from these components. For example, the absolute storage equipment energy use is stagnant throughout the analysis, causing the contribution of this energy-use component to increase from 6% to 18% of total data center energy use across the Current Practices and Economizer scenarios.

A comparison of the energy use estimates for the Current Practices and Economizer Plus scenarios, disaggregated by data center component and space type, is presented in Table 5-11. While total energy use for each space type drops significantly in the Economizer Plus scenario, the relative contribution to the overall data center energy use shifts towards the smallest and largest space types; increasing from 13% to 14% and from 36% to 41% for server closets and enterprise class data centers, respectively. The small increase in relative contribution of energy demand from server closets is a result of fewer servers being candidates for virtualization as well as energy-efficient HVAC equipment that is available to larger data centers not being designed for this smaller space type. The PUE remains at 2.0 for server closets in both scenarios. This increased contribution in server closets is tempered by this space type not hosting any high-end servers or storage equipment, components for which energy demand remains constant between the standard operation and energy-efficient scenarios. While enterprise class data centers benefit from more efficient HVAC equipment, the increase in relative energy

consumption in this space type highlights that these large data centers contain the majority of high-end servers and storage devices. Conversely, server rooms show the greatest relative decrease in energy demand between the two scenarios, since this space type is small and contains no high-end servers or storage equipment, but is large enough to benefit from more efficient HVAC equipment.

Figure 5-4 and Table 5-12 present the total IT and non-IT energy use distributed among the five climate regions for the Baseline and Economizer Plus scenarios. Variations in the regional energy use result from climate-associated differences among cities in the calculated non-IT energy loads, since IT energy is assumed to be equal for each region. Under both the baseline and energy-efficient scenario, a less-than-500 million kWh/year difference is predicted to occur between any pair of results from the climate regions. The regional similarity in energy use for the Baseline scenario results from similarities in the regional PUE values, which is expected for data centers with mechanical designs that employ minimal outside air. The Baseline PUE values are nearly identical for the larger data centers that use chilled water systems, rather than less efficient DX cooling. The corresponding PUE values in the Economizer Plus scenario, where outside air is used to reduce chiller operation when the outside air temperature is less than the return-air temperature setpoint, show a greater dependence on climate. This increase in PUE variation results in some divergence in energy use among different regions. For example, non-IT data center energy use in Dallas is 5% greater than Seattle in the Baseline scenario, but this difference increases to 30% in the Economizer Plus scenario. The impact of the PUE variations in the energy-efficient scenario is muted, however, owing to the significant IT energy reductions that result from the volume server

efficiency measures. Since non-IT energy use is a function of the heat generated from IT energy use, IT efficiency measures reduce the need for non-IT energy, thus minimizing the absolute differences among regions in total data center energy use.

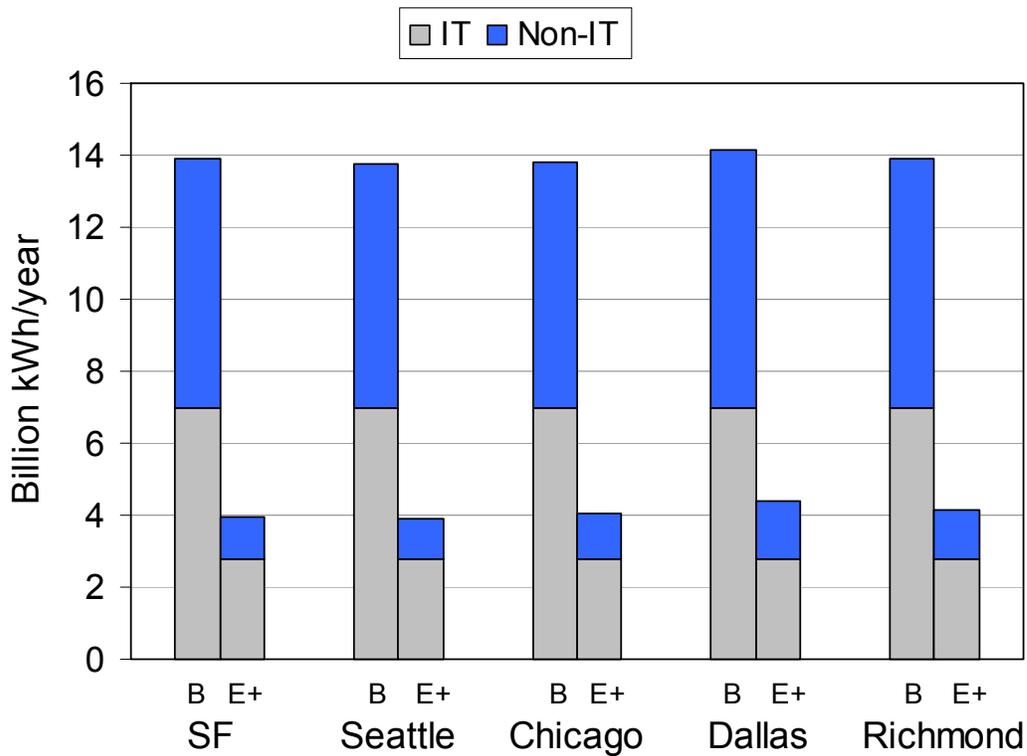


Figure 5-4. Total data center energy use, separated into IT and non-IT components, for each climate region. Each climate region compares energy use under the Baseline (B) and Economizer Plus (E+) scenarios.

Table 5-12. Total data center energy use (IT + non-IT) separated by climate region for the Current Practices and Economizer Plus scenarios.

Region	Scenario	IT	Non-IT	Total
San Francisco, CA	Current Practices	7.0	6.9	13.9
	Economizer Plus	2.8	1.2	4.0
Seattle, WA	Current Practices	7.0	6.8	13.8
	Economizer Plus	2.8	1.1	3.9
Chicago, IL	Current Practices	7.0	6.9	13.8
	Economizer Plus	2.8	1.3	4.1
Dallas, TX	Current Practices	7.0	7.2	14.2
	Economizer Plus	2.8	1.6	4.4
Richmond, VA	Current Practices	7.0	7.0	13.9
	Economizer Plus	2.8	1.4	4.2

5.3.2. PUE calculations

The comparison of total data center energy use under the standard operation and energy-efficient scenarios is calculated using PUE values modeled specifically to data center space type and climate region. Table 5-13 presents the modeled PUE values for server rooms in each climate region. Under both Baseline and Economizer scenarios, no humidity restrictions are imposed on the small data center space and systems are designed to supply 17.8 °C (64 °F) air to the computer servers at an airflow rate that causes the return air temperature to increase to 22.2 °C (72 °F). Higher temperature setpoints were assumed impractical for server rooms and therefore the Economizer Plus scenario was not modeled because of greater temperature fluctuation inherent to such a small space type. The temperature fluctuations result from the room lacking space to adequately separate the hot and cold airstreams through aisle containment (alternating

rows of hot and cold aisles), which is typical for larger data centers. The use of only one CRAC unit in the server room also contributes to temperature fluctuations as the compressor in the DX unit cycles on and off, leading to possible periods when the supply air temperature reaching the IT equipment is greater than intended. The economizer use in the Economizer scenario reduces the energy use associated with the DX cooling. The savings vary by climate region. A 31% reduction in DX cooling is observed in the Dallas climate while DX cooling is reduced by nearly 86% in Seattle. A portion of the cooling savings is lost by increases in fan energy in the Economizer scenario, which results from the increased air resistance caused by the HVAC ducting, relative to the ductless CRAC systems used in the Baseline scenario. The use of the exhaust fan during economizer periods also increases the fan energy and results in a greater overall fan energy use in climates where the economizers can be used for more hours in the year. The increased fan energy, combined with modest DX cooling savings, results in very small improvements in the PUE in the warmer climates. The effect of increased fan energy on the PUE reduction is clearly observable in Figure 5-5, which compares the server room HVAC component energy use in the Baseline and Economizer scenarios, for each climate region, as a percentage of the IT energy demand.

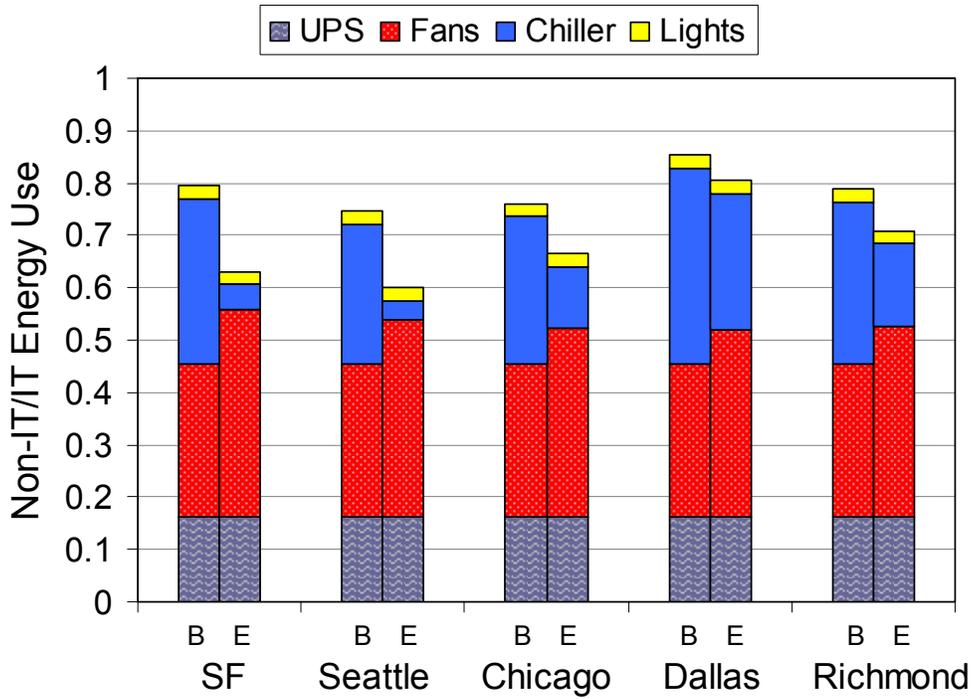


Figure 5-5. Server room non-IT component energy use scaled in proportion to the IT energy demand. Each climate region compares energy use under the Baseline (B) and Economizer (E) scenarios.

Table 5-13. Annual energy modeling results for server room energy efficiency for a Baseline and Economizer scenario.

Scenario	Climate Zone	Server Load	Lights	UPS Waste Heat	Fans	DX Cooling	Total Annual Energy Usage	Peak Electric Demand	Power Usage Effectiveness (PUE)
Baseline	San Francisco, CA	88	2.2	14	26	28	157	20	1.79
Economizer					35	4.3	143	21	1.63
Baseline	Seattle, WA	88	2.2	14	26	23	153	20	1.75
Economizer					33	3.3	140	21	1.60
Baseline	Chicago, IL	88	2.2	14	26	25	154	20	1.76
Economizer					32	10	146	21	1.66
Baseline	Dallas, TX	88	2.2	14	26	33	162	20	1.85
Economizer					31	23	158	21	1.80
Baseline	Richmond, VA	88	2.2	14	26	27	157	20	1.79
Economizer					32	14	150	21	1.71

Under both Baseline and Economizer scenarios no humidity restriction are imposed and the supply/return temperature setpoints are 18/22 °C. The Economizer Plus scenario is not applied to server rooms. The difference in PUE values indicates a potential savings of 21%, 19%, 13%, 6%, and 10% for data center non-IT energy use in San Francisco, Seattle, Chicago, Dallas, and Richmond, respectively.

Tables 5-14 through 5-18 present PUE modeling results for localized data centers in each climate region. Overall, greater savings are observed, relative to the savings calculated for the server rooms, by combining economizer use with different temperature setpoints. The Baseline scenario complies with the ASHRAE recommended humidity range of 40-55% (ASHRAE, 2005) and the cooling systems are designed to supply 17.8 °C (64 °F) air to the computer servers at an airflow rate that causes the return air temperature to increase to 22.2 °C (72 °F). Energy performance is modeled in the energy efficient scenario under seven different combinations of humidity and temperature settings. While maintaining the Baseline scenario temperature settings, the humidity restrictions are first maintained at the ASHRAE recommended range of 40-55%, which represents the Economizer scenario. The humidity restrictions are then expanded to the ASHRAE allowable humidity range of 20-80% RH, and for a third case energy use is estimated with the humidity restrictions removed altogether (represented as a humidity range of 1-100% RH). The energy use consequences for these three different humidity ranges varies by climate region. The economizers result in significant DX cooling energy saving in the cooler climates of San Francisco and Seattle and these DX cooling savings increase as the humidity restrictions are relaxed as a result of increased hours of economizer operation. In the climates of Chicago, Dallas, and Richmond, which experience more periods of increased temperature and humidity, saving from reduced DX cooling DX cooling are still achieved while maintaining the 40-55% RH restriction, though to a lesser extent than in San Francisco or Seattle. These savings are increased as the humidity restriction is relaxed to 20-80%. Savings from reduced DX cooling needs are ever higher in these climate regions when the humidity restrictions are completely

removed owing to economizer operation during high humidity periods when the outside temperature is below the return air setpoint (22.2 °C), which results in significant latent cooling demand. As expected, greater DX cooling savings are observed in all climate regions when the temperature setpoints are increased to supply 24.4 °C (76 °F) air to the computer servers. The airflow rate remains constant, causing the same temperature increase in the return air, which now reaches 28.9 °C (84 °F). Also, at the higher temperature set point, DX cooling savings continue to increase or remain unchanged for all climate regions as the humidity restrictions are relaxed, indicating that, even in the more humid regions, savings for these higher temperature settings are not compromised by increased latent cooling. In the five climate regions, humidification energy only nominally varies as a consequence of differences in the humidity restrictions, which is due to the use of more efficient adiabatic humidifiers. Previous studies have indicated that economizer energy savings can be eclipsed by humidification energy when electric humidifiers are used with economizers while maintaining 40-55% RH (Hydeman and Tschudi, 2009).

Like the server room space type, the fan energy in the Economizer scenario of the localized data centers increases owing to the addition of exhaust fans and the increased air resistance associated with the ducted air delivery system. The effect is smaller for localized data centers than for server rooms because more efficient fan and motor equipment are available for the larger air-handling equipment. An additional modification is made to the temperature parameters in the Economizer Plus scenario to prevent an increase in fan energy from compromising the cooling energy savings through economizer operation. The return air temperature is increased from 22.2 °C (72 °F) to

28.9 °C (84 °F) while the supply air temperature remains at 17.8°C (64 °F). This increased difference between the supply and return air temperatures allows for a reduction in airflow rate, resulting in less fan energy. Under these temperature settings economizer activity is similar to the 24.4/28.9 °C supply/return air setting, though to achieve the lower supply air temperature setpoint many of the full economizer mode hours shift to partial economizer mode. The increased chiller activity combined with the reduced airflow rate results in a net energy savings in all climate regions except for Dallas where the efficiency does not change.

The Economizer Plus PUE values associated with the increased difference in temperature setpoints and no humidity restriction represent the greatest potential savings calculated in the chapter that is available to localized data centers. Figure 5-6 compares the localized data center HVAC components under the Economizer Plus PUE scenario with the Baseline scenario for each climate region, as a percentage of the IT energy demand. Reductions in both the fan and chiller energy contribute to the improved PUE values in the Economizer Plus scenario.

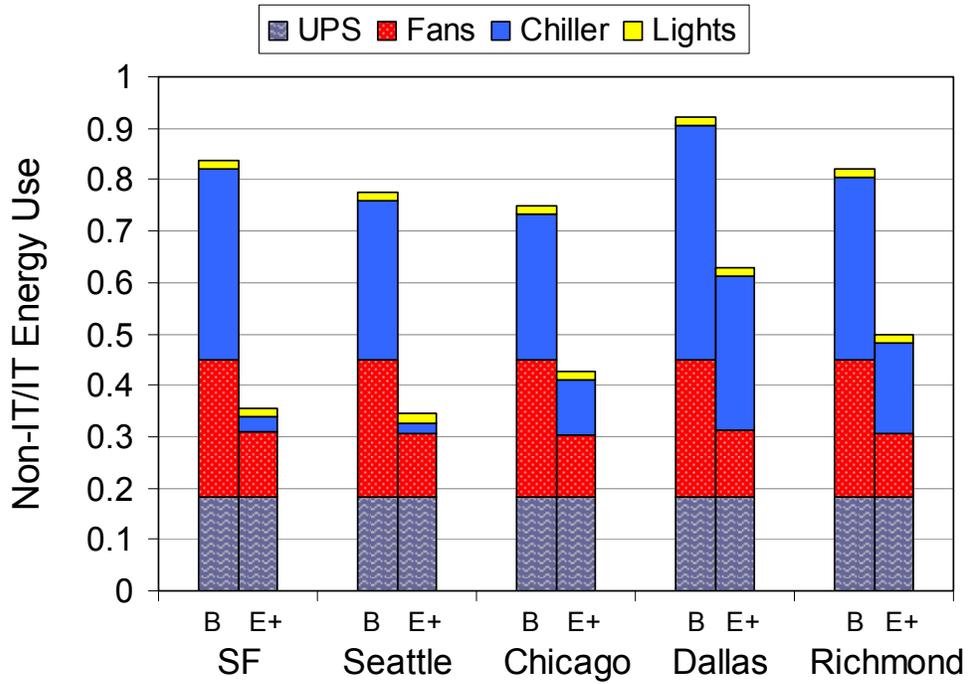


Figure 5-6. Localized data center non-IT component energy use scaled in proportion to the IT energy demand. Each climate region compares energy use under the Baseline (B) and Economizer Plus (E+) scenarios.

Table 5-14. Annual energy modeling results for localized data center energy efficiency in San Francisco, CA.

Scenario	Supply/ Return Air Setpoint (°C)	Humidity Restrictions		Server Load	Lights	UPS Waste Heat	Fans	Humid- ifier	DX Cooling	Total Annual Energy Usage	Peak Electric Demand	Power Usage Effective- ness (PUE)
		lower RH	upper RH									
Baseline	18/22	40%	55%				70	0	98	483	62	1.84
Economizer	18/22	40%	55%				81	0.3	22	419	59	1.60
	18/22	20%	80%				85	0.01	12	413	59	1.57
	18/22	1%	100%				85	0	12	413	59	1.57
	24/29	40%	55%	263	4.4	48	79	1.5	1.3	397	60	1.51
Economizer Plus	24/29	20%	80%				79	0.04	1.4	395	60	1.50
	24/29	1%	100%				79	0	1.5	395	61	1.50
	18/29	1%	100%				33	0	7.8	356	53	1.36

Results present energy consumption for a Baseline scenario (no economizer) and energy efficient scenarios with economizer use under different air temperature setpoints and humidity restrictions. The difference in PUE values indicates a potential savings of 58% in data center non-IT energy use.

Table 5-15. Annual energy modeling results for localized data center energy efficiency in Seattle, WA.

Scenario	Supply/ Return Air Setpoint (°C)	Humidity Restrictions		Server Load	Lights	UPS Waste Heat	Fans	Humid- ifier	DX Cooling	Total Annual Energy Usage	Peak Electric Demand	Power Usage Effective- ness (PUE)																		
		lower RH	upper RH																											
Baseline	18/22	40%	55%	263	4.4	48	70	0	82	467	62	1.78																		
Economizer	18/22	40%	55%										81	0.70	13	410	59	1.56												
	18/22	20%	80%																82	0.02	10	407	59	1.55						
	18/22	1%	100%																						82	0	10	407	59	1.55
	24/29	40%	55%																											
24/29	20%	80%	78										0.11	1.3	394	61	1.50													
24/29	1%	100%																78	0	1.3	394	61	1.50							
Economizer Plus	18/29	1%																						100%	32	0	5.8	353	53	1.34

Results present energy consumption for a Baseline scenario (no economizer) and energy efficient scenarios with economizer use under different air temperature setpoints and humidity restrictions. The difference in PUE values indicates a potential savings of 56% in data center non-IT energy use.

Table 5-16. Annual energy modeling results for localized data center energy efficiency in Chicago, IL.

Scenario	Supply/ Return Air Setpoint (°C)	Humidity Restrictions		Server Load	Lights	UPS Waste Heat	Fans	Humid- ifier	DX Cooling	Total Annual Energy Usage	Peak Electric Demand	Power Usage Effective- ness (PUE)
		lower RH	upper RH									
Baseline	18/22	40%	55%	263	4.4	48	70	0	75	460	62	1.75
Economizer	18/22	40%	55%				77	1.0	34	428	59	1.63
	18/22	20%	80%				80	0.13	27	423	59	1.61
	18/22	1%	100%				81	0	41	438	70	1.67
	24/29	40%	55%				78	2.0	13	408	59	1.55
Economizer Plus	24/29	20%	80%				79	0.33	8.6	404	59	1.54
	24/29	1%	100%				79	0	8.6	403	59	1.53
	18/29	1%	100%				32	0	28	375	64	1.43

Results present energy consumption for a Baseline scenario (no economizer) and energy efficient scenarios with economizer use under different air temperature setpoints and humidity restrictions. The difference in PUE values indicates a potential savings of 43% in data center non-IT energy use.

Table 5-17. Annual energy modeling results for localized data center energy efficiency in Dallas, TX.

Scenario	Supply/ Return Air Setpoint (°C)	Humidity Restrictions		Server Load	Lights	UPS Waste Heat	Fans	Humid- ifier	DX Cooling	Total Annual Energy Usage	Peak Electric Demand kW	Power Usage Effective- ness (PUE)						
		lower RH	upper RH															
Baseline	18/22	40%	55%	263	4.4	48	MWh/year		120	505	63	1.92						
Economizer	18/22	40%	55%				70	0					76	0.46	75	467	59	1.78
	18/22	20%	80%				80	0.03					65	0	104	461	59	1.75
	18/22	1%	100%				83	0					50	1.0	28.0	444	59	1.69
	24/29	40%	55%				78	0.10					28.1	0	79.4	427	59	1.63
Economizer Plus	24/29	20%	80%				84	0					34	0	428	428	67	1.63
	24/29	1%	100%															

Results present energy consumption for a Baseline scenario (no economizer) and energy efficient scenarios with economizer use under different air temperature setpoints and humidity restrictions. The difference in PUE values indicates a potential savings of 32% in data center non-IT energy use.

Table 5-18. Annual energy modeling results for localized data center energy efficiency in Richmond, VA.

Scenario	Supply/ Return Air Setpoint (°C)	Humidity Restrictions		Server Load	Lights	UPS Waste Heat	Fans	Humid- ifier	DX Cooling	Total Annual Energy Usage	Peak Electric Demand	Power Usage Effective- ness (PUE)
		lower RH	upper RH									
Baseline	18/22	40%	55%	263	4.4	48	70	0	94	479	63	1.82
Economizer	18/22	40%	55%				77	0.94	49	443	59	1.68
	18/22	20%	80%				81	0.12	39	436	59	1.66
	18/22	1%	100%				82	0	62	459	79	1.75
	24/29	40%	55%				78	1.8	25	420	59	1.60
Economizer Plus	24/29	20%	80%				81	0.29	15.0	412	59	1.57
	24/29	1%	100%				81	0	16.4	413	65	1.57
	18/29	1%	100%				33	0	46	394	69	1.50

Results present energy consumption for a Baseline scenario (no economizer) and energy efficient scenarios with economizer use under different air temperature setpoints and humidity restrictions. The difference in PUE values indicates a potential savings of 39% in data center non-IT energy use.

Tables 5-19 through 5-23 present PUE modeling results for mid-tier data centers in each climate region. The shift from the DX cooling systems in the server rooms and localized data center to the use of air-cooled chillers results in similar cooling energy savings in the more moderate climates (San Francisco and Seattle) and slightly larger savings in the more humid climates (Chicago, Dallas, and Richmond). Again, the higher temperature setpoints in the Economizer Plus scenario produce larger cooling energy savings. However, unlike the localized data center case, for mid-tier data centers the chiller savings are increased or nearly unchanged as the humidity restrictions are removed under both the higher and lower temperature setpoints, indicating that latent cooling poses less of a threat to achieving economizer savings under this cooling design. The increased chiller savings gained from economizer use in the Economizer scenario of the mid-tier data centers are partially lost to the increase in fan energy caused by the ducted air delivery system. While the air-handlers are more efficient in this larger data center, more exhaust fans must operate during the economizer periods and the greater airflow required to remove the increased heat load is more sensitive to the increase in static pressure associated with the ducts. Because of this high fan energy use, a significant improvement in efficiency is observed in the Economizer Plus scenario when adjusting the return air temperature to 28.9 °C (84 °F) while maintaining the supply air temperature at 17.8 °C (64 °F).

The Economizer Plus PUE values calculated with the increased difference in temperature setpoints and no humidity restrictions represent the greatest potential savings calculated in the chapter that is available to mid-tier data centers. These Economizer Plus PUE values are compared to the Baseline PUE values for mid-tier data center for each

climate region in Figure 5-7. This comparison shows that the shift to chiller use in the cooling system results in less variation of the baseline PUE among the different climate regions. In the Economizer Plus scenario the cooling energy use is nearly eliminated in the milder climate regions.

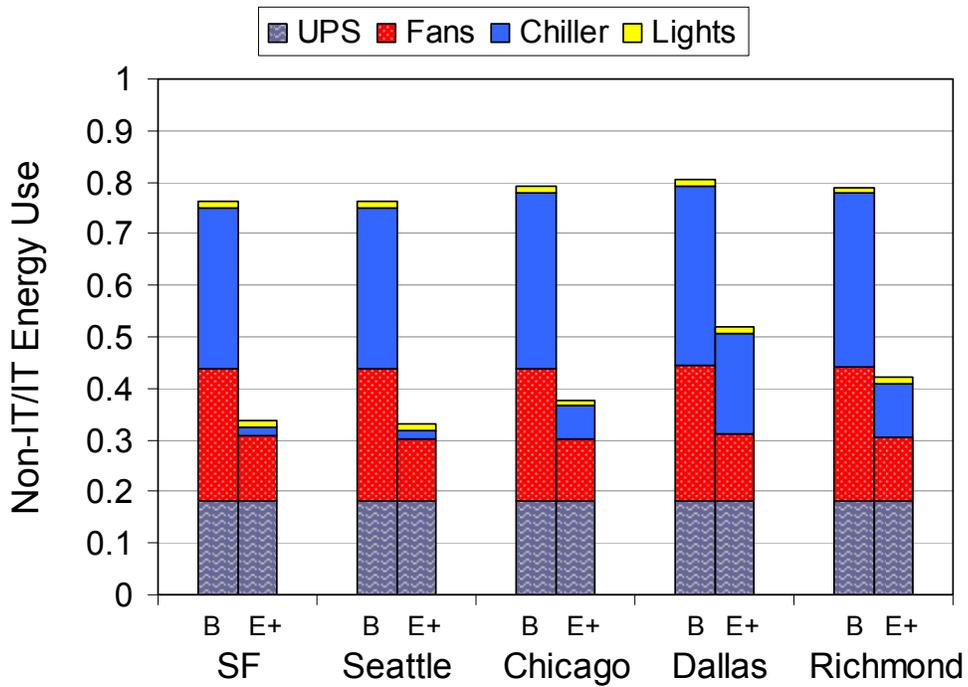


Figure 5-7. Mid-tier data center non-IT component energy use scaled in proportion to the IT energy demand. Each climate region compares energy use under the Baseline (B) and Economizer Plus (E+) scenarios.

Table 5-19. Annual energy modeling results for mid-tier data center energy efficiency in San Francisco, CA.

Scenario	Supply/ Return Air Setpoint (°C)	Humidity Restrictions		Server Load	Lights	UPS Waste Heat	Fans	Humid- ifier	Chiller	Chilled Water Pumps	Total Annual Energy Usage	Peak Electric Demand kW	Power Usage Effective- ness (PUE)
		lower RH	upper RH										
Baseline	18/22	40%	55%				384	3	542	65	3,088	376	1.76
Economizer	18/22	40%	55%				552	2.0	117	12	2,778	389	1.59
	18/22	20%	80%				578	0.06	56	5.9	2,735	389	1.56
	18/22	1%	100%				578	0	56	5.9	2,735	389	1.56
	24/29	40%	55%	1,752	22	321	535	9.2	5.6	0.56	2,645	393	1.51
	24/29	20%	80%				535	0.28	5.9	0.60	2,637	393	1.51
Economizer Plus	24/29	1%	100%				535	0	6.3	0.63	2,637	393	1.51
	18/29	1%	100%				218	0	28	4.0	2,345	340	1.34

Results present energy consumption for a Baseline scenario (no economizer) and energy efficient scenarios with economizer use under different air temperature setpoints and humidity restrictions. The difference in PUE values indicates a potential savings of 56% in data center non-IT energy use.

Table 5-20. Annual energy modeling results for mid-tier data center energy efficiency in Seattle, WA.

Scenario	Supply/ Return Air Setpoint (°C)	Humidity Restrictions		Server Load	Lights	UPS Waste Heat	Fans	Humid- ifier	Chiller	Chilled Water Pumps	Total Annual Energy Usage	Peak Electric Demand	Power Usage Effective- ness (PUE)
		lower RH	upper RH										
Baseline	18/22	40%	55%	1,752	22	321	384	8	539	62	3,089	375	1.76
Economizer	18/22	40%	55%				550	4.4	67	6.7	2,723	389	1.55
	18/22	20%	80%				558	0.12	45	4.6	2,703	389	1.54
Economizer Plus	18/22	1%	100%				558	0	45	4.6	2,703	389	1.54
	24/29	40%	55%				529	13	3.4	0.32	2,641	381	1.51
	24/29	20%	80%				529	0.70	5.6	0.54	2,631	393	1.50
	24/29	1%	100%				529	0	5.6	0.55	2,630	393	1.50
Economizer Plus	18/29	1%	100%				210	0	23	3.1	2,331	339	1.33

Results present energy consumption for a Baseline scenario (no economizer) and energy efficient scenarios with economizer use under different air temperature setpoints and humidity restrictions. The difference in PUE values indicates a potential savings of 57% in data center non-IT energy use.

Table 5-21. Annual energy modeling results for mid-tier data center energy efficiency in Chicago, IL.

Scenario	Supply/ Return Air Setpoint (°C)	Humidity Restrictions		Server Load	Lights	UPS Waste Heat	Fans	Humid- ifier	Chiller	Chilled Water Pumps	Total Annual Energy Usage	Peak Electric Demand kW	Power Usage Effective- ness (PUE)
		lower RH	upper RH										
Baseline	18/22	40%	55%				384	19	577	65	3,141	384	1.79
Economizer	18/22	40%	55%				527	6.4	202	19	2,849	389	1.63
	18/22	20%	80%				549	0.82	156	14	2,815	389	1.61
	18/22	1%	100%				554	0	159	15	2,822	429	1.61
	24/29	40%	55%	1,752	22	321	532	13	83	7.2	2,731	386	1.56
	24/29	20%	80%				541	2.1	58	5.0	2,701	386	1.54
Economizer Plus	24/29	1%	100%				541	0	58	5.0	2,699	386	1.54
	18/29	1%	100%				210	0	99	10	2,414	408	1.38

Results present energy consumption for a Baseline scenario (no economizer) and energy efficient scenarios with economizer use under different air temperature setpoints and humidity restrictions. The difference in PUE values indicates a potential savings of 52% in data center non-IT energy use.

Table 5-22. Annual energy modeling results for mid-tier data center energy efficiency in Dallas, TX.

Scenario	Supply/ Return Air Setpoint (°C)	Humidity Restrictions		Server Load	Lights	UPS Waste Heat	Fans	Humid- ifier	Chiller	Chilled Water Pumps	Total Annual Energy Usage	Peak Electric Demand kW	Power Usage Effective- ness (PUE)
		lower RH	upper RH										
Baseline	18/22	40%	55%				384	3	607	73	3,162	399	1.80
Economizer	18/22	40%	55%				523	3.0	409	36	3,066	401	1.75
	18/22	20%	80%				548	0.18	349	31	3,024	401	1.73
	18/22	1%	100%				566	0	368	33	3,063	530	1.75
	24/29	40%	55%	1,752	22	321	530	6.7	273	23	2,928	400	1.67
	24/29	20%	80%				572	0.66	141	12	2,820	400	1.61
Economizer Plus	24/29	1%	100%				572	0	141	12	2,820	400	1.61
	18/29	1%	100%				222	0	312	31	2,660	445	1.52

Results present energy consumption for a Baseline scenario (no economizer) and energy efficient scenarios with economizer use under different air temperature setpoints and humidity restrictions. The difference in PUE values indicates a potential savings of 36% in data center non-IT energy use.

Table 5-23. Annual energy modeling results for mid-tier data center energy efficiency in Richmond, VA.

Scenario	Supply/ Return Air Setpoint (°C)	Humidity Restrictions		Server Load	Lights	UPS Waste Heat	Fans	Humid- ifier	Chiller	Chilled Water Pumps	Total Annual Energy Usage	Peak Electric Demand kW	Power Usage Effective- ness (PUE)
		lower RH	upper RH										
Baseline	18/22	40%	55%				384	11	578	68	3,137	391	1.79
Economizer	18/22	40%	55%				525	6.0	270	25	2,921	393	1.67
	18/22	20%	80%				551	0.75	208	19	2,874	393	1.64
	18/22	1%	100%				561	0	218	20	2,894	531	1.65
	24/29	40%	55%	1,752	22	321	533	11	134	12	2,785	393	1.59
	24/29	20%	80%				551	1.8	77	6.5	2,731	393	1.56
Economizer Plus	24/29	1%	100%				551	0	77	6.5	2,729	419	1.56
	18/29	1%	100%				215	0	163	17	2,490	459	1.42

Results present energy consumption for a Baseline scenario (no economizer) and energy efficient scenarios with economizer use under different air temperature setpoints and humidity restrictions. The difference in PUE values indicates a potential savings of 47% in data center non-IT energy use.

Tables 5-24 through 5-28 present PUE modeling results for enterprise data centers in each climate region. The use of water-cooled chillers for the cooling system in these large data centers provides the most efficient PUE values in all the standard operation and energy efficient scenarios. In Dallas and Richmond, the Baseline PUE values for the enterprise data center are lower than the respective server room PUE values for the Economizer scenario. Under both temperature setpoints in the energy efficient scenarios for the enterprise data center, relaxing the humidity restrictions to 20-80% RH results in greater energy savings, but further removing the humidity restriction altogether provides no increased benefit or results in a slight latent cooling penalty depending on the climate. The high airflow required for this large data center causes significant fan energy demand that is further increased by the addition of multiple exhaust fans in the Economizer scenario. Without increasing the difference in temperature setpoints, this high fan energy demand suppresses the large energy savings gained from reduced chiller operation during economizer use. The increased chiller activity when the return air temperature is increased to 28.9 °C (84 °F) while maintaining the supply air temperature at 17.8 °C (64 °F) in the Economizer Plus scenario is minimal relative to the fan energy savings due to the efficiency of the water-cooled chiller based cooling system. Figure 5-8 compares the enterprise data center fans and chiller under the Economizer Plus scenario with the Baseline scenario for each climate region. The Baseline PUE values in Figure 5-8 are nearly identical across all climate regions and the improved PUE values in the Economizer Plus scenario are the most efficient of all the space types evaluated in this chapter.

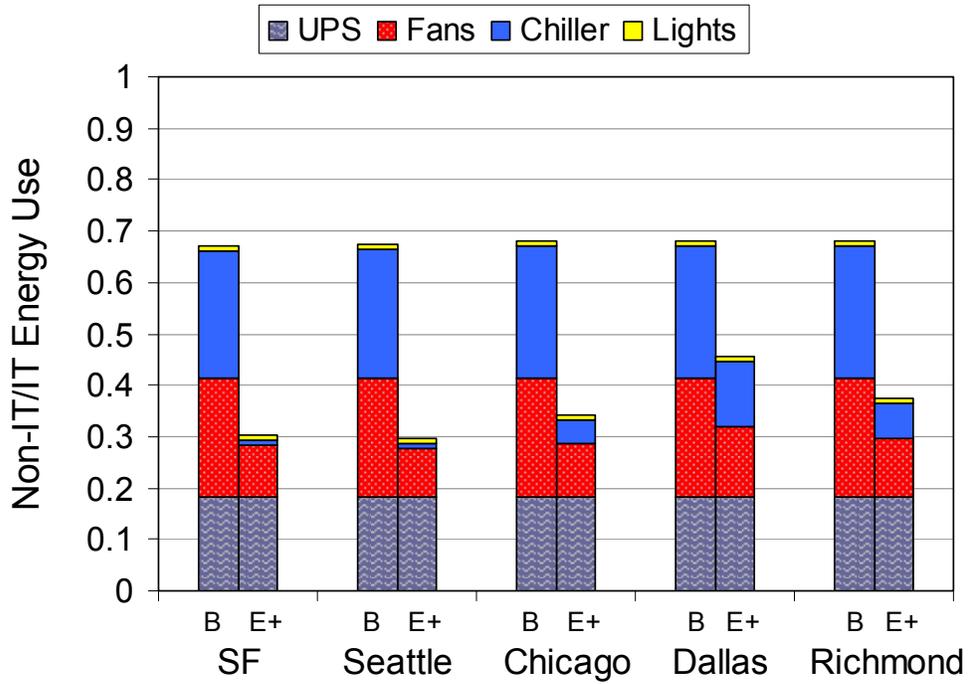


Figure 5-8: Enterprise data center non-IT component energy use scaled in proportion to the IT energy demand. Each climate region compares energy use under the Baseline (B) and Economizer Plus (E+) scenarios.

Table 5-24. Annual energy modeling results for enterprise data center energy efficiency located in San Francisco, CA.

Scenario	Supply/ Return Air Setpoint (°C)	Humidity Restrictions		Server Load	Lights	UPS	Waste Heat	Air Handler		Chilled Water Plant				Total Annual Energy Usage	Peak Electric Demand	Power Usage Effective- ness (PUE)
		lower RH (%)	upper RH (%)					Fans	Humid- ifier	Chiller	Cooling Tower	CHW Pumps	CW Pumps			
Baseline	18/22	40	55	4,380	44	803		756	5	1,071	11	198	54	7,323	845	1.67
Economizer	18/22	40	55					868	4.8	209	2.4	49	13.8	6,374	855	1.46
	18/22	20	80					918	0.14	102	1.3	29	10.4	6,289	862	1.44
	18/22	1	100					918	0	103	1.3	29	10.4	6,289	862	1.44
	24/29	40	55					831	22	9.4	0.12	2.9	1.2	6,094	849	1.39
	24/29	20	80					832	0.67	10.1	0.14	3.3	1.4	6,074	852	1.39
	24/29	1	100					832	0	10.7	0.14	3.4	1.4	6,074	859	1.39
Economizer Plus	18/29	1	100					400	0	50	0.54	22	10.7	5,711	821	1.30

Results present energy consumption for a Baseline scenario (no economizer) and energy efficient scenarios with economizer use under different air temperature setpoints and humidity restrictions. The difference in PUE values indicates a potential savings of 55% in data center non-IT energy use.

Table 5-25. Annual energy modeling results for enterprise data center energy efficiency located in Seattle, WA.

Scenario	Supply/ Return Air Setpoint (°C)	Humidity Restrictions		Server Load	Lights	UPS	Waste Heat	Air Handler		Chilled Water Plant				Total Annual Energy Usage	Peak Electric Demand	Power Usage Effective- ness (PUE)
		lower RH (%)	upper RH (%)					Fans	Humid- ifier	Chiller	Cooling Tower	CHW Pumps	CW Pumps			
MWh/year																
Baseline	18/22	40	55	4,380	44	803		756	17	1,070	11	198	54	7,333	838	1.67
Economizer	18/22	40	55					861	10	119	1.4	30	9.0	6,258	862	1.43
	18/22	20	80					877	0.29	81	1.0	23	8.0	6,217	863	1.42
	18/22	1	100					877	0	81	1.0	23	8.0	6,217	863	1.42
	24/29	40	55					819	32	5.4	0.07	1.6	0.56	6,085	845	1.39
Economizer Plus	24/29	20	80					819	1.7	9.3	0.12	3.1	1.2	6,061	854	1.38
	24/29	1	100					819	0	9.5	0.12	3.1	1.3	6,060	854	1.38
	18/29	1	100					383	0	40	0.47	17	8.0	5,677	783	1.30

Results present energy consumption for a Baseline scenario (no economizer) and energy efficient scenarios with economizer use under different air temperature setpoints and humidity restrictions. The difference in PUE values indicates a potential savings of 56% in data center non-IT energy use.

Table 5-26. Annual energy modeling results for enterprise data center energy efficiency located in Chicago, IL.

Scenario	Supply/ Return Air Setpoint (°C)	Humidity Restrictions		Server Load	Lights	UPS	Waste Heat	Air Handler		Chilled Water Plant				Total Annual Energy Usage	Peak Electric Demand	Power Usage Effective- ness (PUE)
		lower RH (%)	upper RH (%)					Fans	Humid- ifier	Chiller	Cooling Tower	CHW Pumps	CW Pumps			
MW/h/year																
Baseline	18/22	40	55	4,380	44	803		756	37	1,073	17	198	54	7,363	850	1.68
Economizer	18/22	40	55					811	15	334	5.1	79	21	6,492	862	1.48
	18/22	20	80					852	2.0	255	4.2	64	19	6,423	911	1.47
	18/22	1	100					861	0	268	4.7	68	20	6,448	987	1.47
	24/29	40	55					823	30	128	2.6	31	8.8	6,251	848	1.43
	24/29	20	80					840	5.0	88	1.59	24	8.0	6,194	855	1.41
	24/29	1	100					840	0	89	1.59	24	8.0	6,190	855	1.41
Economizer Plus	18/29	1	100					380	0	194	4.2	56	20	5,881	923	1.34

Results present energy consumption for a Baseline scenario (no economizer) and energy efficient scenarios with economizer use under different air temperature setpoints and humidity restrictions. The difference in PUE values indicates a potential savings of 50% in data center non-IT energy use.

Table 5-27. Annual energy modeling results for enterprise data center energy efficiency located in Dallas, TX.

Scenario	Supply/ Return Air Setpoint (°C)	Humidity Restrictions		Server Load	Lights	UPS	Waste Heat	Air Handler		Chilled Water Plant				Total Annual Energy Usage	Peak Electric Demand	Power Usage Effective- ness (PUE)				
		lower RH (%)	upper RH (%)					Fans	Humid- ifier	Chiller	Cooling Tower	CHW Pumps	CW Pumps							
MW/year																				
Baseline	18/22	40	55	4,380	44	803		756	6	1,081	40	200	54	7,365	850	1.68				
Economizer	18/22	40	55					801	7	652	15	154	39	6,894	862	1.57				
	18/22	20	80					850	0.42	551	14	135	36	6,814	862	1.56				
	18/22	1	100					880	0	608	16	149	41	6,921	1127	1.58				
Economizer Plus	24/29	40	55					817	16	416	12	98	25	6,611	848	1.51				
	24/29	20	80					900	1.6	210	5.6	58	19	6,421	883	1.47				
	24/29	1	100					900	0	211	5.6	58	19	6,421	927	1.47				
Economizer Plus		18/29	1					100				405	0	541	17	142	45	6,377	994	1.46

Results present energy consumption for a Baseline scenario (no economizer) and energy efficient scenarios with economizer use under different air temperature setpoints and humidity restrictions. The difference in PUE values indicates a potential savings of 33% in data center non-IT energy use.

Table 5-28. Annual energy modeling results for enterprise data center energy efficiency located in Richmond, VA.

Scenario	Supply/ Return Air Setpoint (°C)	Humidity Restrictions		Server Load	Lights	UPS	Waste Heat	Air Handler		Chilled Water Plant				Total Annual Energy Usage	Peak Electric Demand	Power Usage Effective- ness (PUE)
		lower RH (%)	upper RH (%)					Fans	Humid- ifier	Chiller	Cooling Tower	CHW Pumps	CW Pumps			
Baseline	18/22	40	55	4,380	44	803		756	22	1,076	24	199	54	7,358	850	1.68
Economizer	18/22	40	55					808	14	443	8.5	104	27	6,631	862	1.51
	18/22	20	80					858	1.8	337	7.4	84	24	6,539	893	1.49
	18/22	1	100					874	0	366	8.5	92	26	6,594	1130	1.51
	24/29	40	55					824	27	206	5.7	49	13.2	6,352	845	1.45
	24/29	20	80					860	4.4	116	2.9	32	10.3	6,252	887	1.43
	24/29	1	100					860	0	116	2.9	32	10.4	6,248	995	1.43
Economizer Plus	18/29	1	100					391	0	290	8.3	81	27.3	6,024	1032	1.38

Results present energy consumption for a Baseline scenario (no economizer) and energy efficient scenarios with economizer use under different air temperature setpoints and humidity restrictions. The difference in PUE values indicates a potential savings of 45% in data center non-IT energy use.

Table 5-29 presents a summary of the Baseline PUE values calculated for each data center space type and climate zone. The adjusted Baseline PUE values used to estimate the Current Practices PUE values are presented in Table 5-30, where the weighted average of the Current Practices PUE values is 2.0. A PUE value of 2.0 matches the best available data on current data center performance (Greenberg et al., 2006; Brown et al., 2007; Belady and Malone, 2007; Sullivan, 2009). The PUE values used to estimate data center energy use in the Economizer scenario are presented in Table 5-31, with the closet data centers PUE assumed to remain at 2.0 while the other PUE values reflect the economizer use while maintaining the humidity and temperature settings equal to the Baseline PUE calculations. Table 5-32 presents the Economizer Plus PUE values, which represent economizer operation with 17.8 °C (64 °F) / 28.9 °C (84 °F) supply/return temperature setpoints and no humidity restrictions for the localized, mid-tier, and enterprise data centers. Economizer Plus PUE values remain the same as the Economizer PUE values for the server room and closet data centers since expanded temperature setpoints are not a viable option for these space types.

Table 5-29. Baseline PUE values for the standard operation scenario.

	SF	Seattle	Chicago	Dallas	Richmond	Average	Weighted average
Server closet	n/a	n/a	n/a	n/a	n/a	n/a	1.80
Server room	1.79	1.75	1.76	1.85	1.79	1.79	
Localized DC	1.84	1.78	1.75	1.92	1.82	1.82	
Mid-tier DC	1.76	1.76	1.79	1.80	1.79	1.78	
Enterprise-class DC	1.67	1.67	1.68	1.68	1.68	1.68	

PUE values assume humidity restrictions at the ASHRAE recommended range of 40-55% RH and a supply air temperature of 18 °C (64 °F) with the return air temperature increasing to 22 °C (72 °F).

Table 5-30. Current Practices PUE values for the standard operation scenario.

	SF	Seattle	Chicago	Dallas	Richmond	Average	Weighted average
Server closet	2	2	2	2	2	2	2.00
Server room	1.99	2.00	2.01	2.10	2.04	2.03	
Localized DC	2.09	2.03	2.00	2.17	2.07	2.07	
Mid-tier DC	2.01	2.01	2.04	2.05	2.04	2.03	
Enterprise-class DC	1.92	1.92	1.93	1.93	1.93	1.93	

Adjusted baseline PUE values used to estimate data center energy use under current practices. Current Practices PUE values correspond to an overall average PUE value of 2.0 when weighted by the space type distribution presented in Table 5-3.

Table 5-31. Economizer PUE values for the energy efficiency scenario.

	SF	Seattle	Chicago	Dallas	Richmond	Average	Weighted average
Server closet	2	2	2	2	2	2	1.66
Server room	1.63	1.60	1.66	1.80	1.71	1.68	
Localized DC	1.60	1.56	1.63	1.78	1.68	1.65	
Mid-tier DC	1.59	1.55	1.63	1.75	1.67	1.64	
Enterprise-class DC	1.46	1.43	1.48	1.57	1.51	1.49	

PUE values (except server closets) assume 40-55% humidity restrictions and a supply air temperature of 18 °C (64 °F) with the return air temperature increasing to 24 °C (72 °F). No humidity restrictions are placed on server rooms.

Table 5-32. Economizer Plus PUE values for the energy efficiency scenario.

	SF	Seattle	Chicago	Dallas	Richmond	Average	Weighted average
Server closet	2	2	2	2	2	2	1.55
Server room	1.63	1.60	1.66	1.80	1.71	1.68	
Localized DC	1.36	1.34	1.43	1.63	1.50	1.45	
Mid-tier DC	1.34	1.33	1.38	1.52	1.42	1.40	
Enterprise-class DC	1.30	1.30	1.34	1.46	1.38	1.35	

PUE values (except server closets) assume no humidity restrictions and a supply air temperature of 18 °C (64 °F) with the return air temperature increasing to 29 °C (84 °F). Return air temperature setpoint remains at 24 °C (72 °F) for server rooms.

5.3.3. Carbon intensity

Along with climate differences, site location can also affect the mix of primary energy that is used to generate the electricity supplied to a data center. Table 5-33 shows how the electricity mix varies among cities, with electricity generation in some regions more reliant on fossil fuels while hydro- or nuclear- generated electricity provide the majority of electricity for other regions. This variation in electricity source affects the fossil carbon and air pollutant emissions associated with data center operation. Table 5-34 includes previously reported carbon equivalent greenhouse gas emission estimates for the generation of electricity from forms other than fossil fuels. These emission rates represent the entire lifecycle of electricity production. Specifically, nuclear and solar values include the gases emitted during the extraction, processing, and disposal of associated materials (Fthenakis and Kim, 2007). The hydro, and wind values include emissions from infrastructure construction, flooded biomass decay in the reservoir, loss of net ecosystem production, and land use (Pacca and Horvath, 2002). Table 5-34 also presents direct CO₂ emissions from coal and natural gas combustion, based on electricity generation (EIA, 2009a) and emission inventories (EIA, 2009b) for the United State in 2008. While lifecycle emissions are not included for coal and natural gas, the large direct emissions indicate that emissions from other aspects of generation are a relatively minor contribution. The emission estimates in Table 5-34 indicate amounts of greenhouse gases released from the different primary energy sources that significantly contribute to the electricity mixes in Table 5-33. Hydro, nuclear, and renewables provide electricity at approximately 20 g CO₂(e)/kWh. Emissions increase by more than an order of magnitude to 407 g CO₂(e)/kWh when natural gas is the primary energy source. Coal provides electricity with the highest emission level at more than 977 g CO₂(e)/kWh.

Table 5-33. Electricity resource mix for each of the modeled data center geographical regions.

Region	Provider	year	Electricity Resource Mix (%)										
			Coal	Oil	Gas	Other fossil	Bio-mass	Hydro	Nuclear	Wind	Solar	Geo-thermal	Other
San Francisco, CA	Pacific Gas & Electric	2008	8	0	39	0	4	19	22	2	<1	4.8	<1
Seattle, WA	Seattle City Light	2007	<1	0	<1	0	0	91	5	3	0	0	<1
Chicago, IL	Deregulation (state average)	2005	48	<1	4	<1	<1	<1	48	<1	0	0	0
Dallas, TX	Deregulation (state average)	2005	37	1	49	1	<1	<1	10	1	0	0	<1
Richmond, VA	Dominion Virginia Power	2007	46	0	8	0	0	44	44	0	0	0	1

Differences in electricity produced from coal, natural gas, and hydro generated power indicate potential variation in carbon emissions by data center location. Data compiled from regional providers and national data (EPA, 2008). Pacific Gas & Electric and Seattle City Lights data represents the electricity mix supplied to the consumer. National data does not include out-of-state purchases. Dominion Virginia Power data does not include purchased power.

Table 5-34. Carbon intensity values associated with electricity generation in the United States compiled from national inventory data (EIA, 2009a; EIA, 2009b), Fthenakis and Kim (2007), and Pacca and Horvath (2002).

Carbon Intensity	
Fuel Source	g CO₂ (e) /kWh
Coal	976
Natural Gas	407
Nuclear	24
Hydro	20
Wind	5
Solar	22

The carbon intensity values for coal and natural gas only represent direct CO₂ emissions while values for the other fuel sources represent CO₂(e) emitted during the entire life cycle of electricity generation.

Table 5-35. Greenhouse gas intensity (CO₂(e)/kWh) emissions associated with regionally specific electricity generation sources.

	Coal 976 CO ₂ /kWh	Natural Gas 407 CO ₂ /kWh	Nuclear and Renewables ~20 CO ₂ (e)/kWh	Regional Average CO ₂ (e)/kWh
San Francisco, CA	8%	41%	51%	258
Seattle, WA	1%	0%	99%	30
Chicago, IL	48%	4%	48%	493
Dallas, TX	39%	49%	11%	587
Richmond, VA	48%	8%	44%	510

The distribution of electricity generation is based on regional electricity resource mix data in Table 5-32 and the values of CO₂(e) emissions are based on data in Table 5-33.

Table 5-35 categorizes the regional electricity resource mix data from Table 5-33 by these three tiers of the CO₂(e) emissions, which are used to estimate a regional average CO₂(e)/kWh. There is more than an order of magnitude difference among the estimated regional average CO₂(e)/kWh values, with Seattle producing the lowest emissions with 30 g CO₂(e)/kWh and Dallas providing the greatest with 587 g CO₂(e)/kWh. Applying the electricity demand estimates presented in Table 5-12 with the regional CO₂(e)/kWh estimates in Table 5-33 predicts a CO₂(e) emissions rate associated data center operation of approximately 26 Mt CO₂(e)/y and Mt 8 CO₂(e)/y for the Current Practice and Economizer Plus scenarios, respectively. Figure 5-9 shows annual CO₂(e) emission estimates for the Current Practice and Economizer Plus scenarios, separated by region.

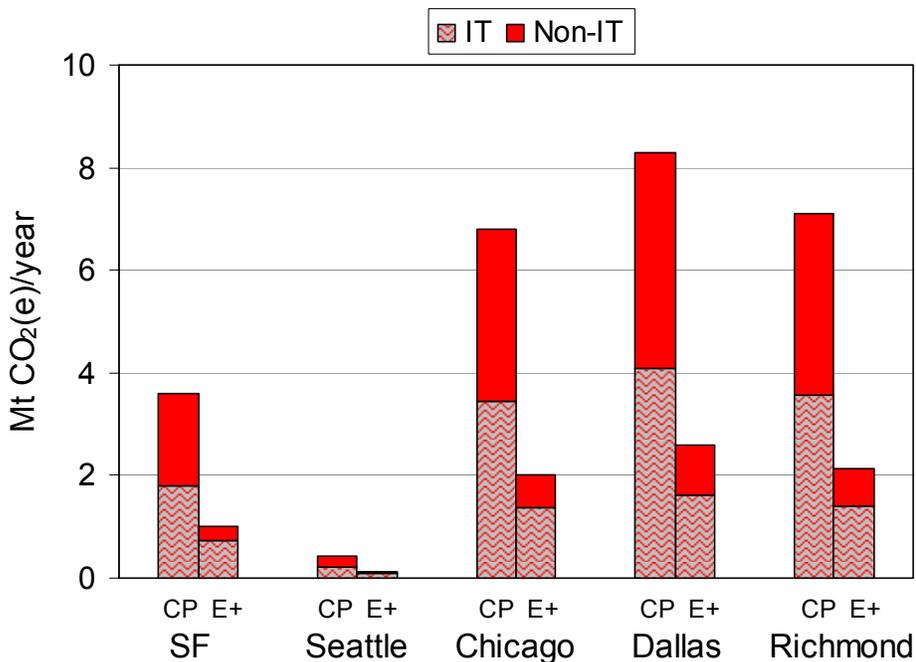


Figure 5-9. Total data center energy use, separated as IT and non-IT components, for each climate region. Each climate region compares energy use under the Current Practices (CP) and Economizer Plus (E+) scenarios. Each region is assumed to equally represent 20% of all data center activity in the United States

Actual CO₂(e) emissions from data center operation depend on the electricity mix available at the location of each building and how those electricity mixes vary throughout the year. Along with regional differences in electricity generation, the primary energy mix used to provide electricity for data center operation may vary by season and time-of-day. Furthermore, the PUE values calculated from modeling results for each region represent annual averages of the mechanical system efficiency, which can vary with meteorological conditions throughout the year. For data centers with economizer use, the average mechanical efficiency is a combination of the periods when the economizer is active and the chiller is not used, periods when the economizer is shut off and the data center operates similarly to a data center without an economizer, and periods of partial economizer use when the temperature of the entering outdoor air is between the supply and return air temperature setpoints allowing for a reduced level of chiller operation. Table 5-36 presents the modeled annual distribution of hours for full, partial, and no economizer activity under the Economizer Plus scenario (17.8 °C (64 °F) / 28.9 °C (84 °F) supply/return temperature setpoints and no humidity restrictions).

The variation in distribution of economizer hours highlights the potential need to evaluate each data center location separately to account for how the electricity mix of that region correlates with data center PUE efficiency. For example, Table 5-37 presents modeled economizer hours for San Francisco, separated by months of the year, and the corresponding PUE value for enterprise data centers in this region under the Economizer Plus scenario. Since hydro-generated electricity represent a significant portion of the annual San Francisco electricity mix, and the amount of hydro-generation varies by month, a more accurate evaluation of the greenhouse-gas emission from the San

Francisco data centers would match the monthly electricity mix averages with the monthly PUE averages presented in Table 5-37. Figure 5-10 shows energy use for enterprise data centers in San Francisco under the Economizer Plus scenario, calculated using the annual average PUE presented in Table 5-24 (PUE=1.30) and the monthly PUE values presented in Table 5-36. The relative increase in energy demand during the warmer months when applying the monthly PUE values represents a shift of 12 million kWh demand, or 3% of the annual energy use, from the cooler months to the warmer months. This shift in energy demand through the year is specific to regional climate and may be greater in areas that experience greater seasonal changes in weather. The effect of unequal energy demand throughout the year on annual CO₂(e) emission estimates depends on the annual variation in CO₂(e)/kWh from electricity generation, though a large variation CO₂(e)/kWh values would be needed to significantly affect CO₂(e) emission estimates. For example, even if the CO₂(e)/kWh values in San Francisco were twice as large during the six months with the highest monthly PUE values in Table 5-37 (344 CO₂(e)/kWh for half of the year and 172 CO₂(e)/kWh for the other half), this change would result in only a 3% increase in annual CO₂(e) estimates.

Table 5-36. Hours of economizer activity for each climate region under the Economizer Plus scenario.

Region	Hours of Economizing		
	Full	Partial	None
San Francisco, CA	7,031	1,694	35
Seattle, WA	7,464	1,268	28
Chicago, IL	5,743	2,434	583
Dallas, TX	2,869	4,555	1,336
Richmond, VA	4,919	3,126	715

Table 5-37. Hours of economizer activity for each month in San Francisco under the Economizer Plus scenario.

Month	Hours of Economizing			Average PUE
	Full	Partial	None	
January	737	7	0	1.28
February	648	24	0	1.28
March	692	52	0	1.29
April	625	95	0	1.29
May	575	161	8	1.31
June	492	225	3	1.32
July	461	280	3	1.33
August	471	272	1	1.33
September	408	301	11	1.34
October	517	218	9	1.32
November	664	56	0	1.29
December	741	3	0	1.28

Figure 5-11 indicates that evaluating greenhouse gas emissions associated with data center operation could also be affected by diurnal changes in the electricity mix. Since PUE efficiency is a function of outdoor temperature in data centers with economizers, reduced efficiency (increased PUE) can be expected during the warmest hours of the year. Figure 5-11 shows how the modeled PUE value for enterprise data centers changes throughout the day during different times of the year in San Francisco under the Economizer Plus scenario (17.8 °C (64 °F) / 28.9 °C (84 °F) supply/return temperature setpoints and no humidity restrictions). During the winter period, the PUE is constant, since the outdoor temperature never increases above the supply temperature setpoint throughout all 24 hours of the day, allowing the HVAC system to constantly remain in full economizer mode. During the spring and fall periods, the HVAC system remains in full economizer mode during the evening, night, and early morning hours of the day. Increases in outdoor temperature during the mid-day hours of the spring and fall periods cause the HVAC system to switch to partial economizer operation, which

increases the PUE. During the summer period, the HVAC system operates in full economizer mode during the nighttime hours, partice economizer mode during the morning and evening hours, and with no economizer during the mid-day hours. These three modes of economizer operation during the summer cause the PUE to increase in the morning hours and to further increase during the mid-day before beginning to decrease again in the evening. The PUE values in Figure 5-11 are used to estimate the hourly energy use values presented in Figure 5-12 for enterprise data centers in San Francisco for the Economizer Plus scenario. Calculations using an hourly PUE, rather than a daily average PUE, show a shift in energy demand from night and early evening hours to mid-day hours. This shift represents 5%, 16%, and 7% of the total daily energy demand for these spring, summer, and fall days, respectively. Since the winter PUE value is constant (always operating in full economizer mode) there is no change in winter between the calculated HVAC energy demand using hourly PUE values or daily PUE averages.

The effect of any daily variation in PUE on CO₂(e) emission estimates depends on the variation in CO₂(e)/kWh between electricity during on-peak and off-peak periods. For example, if the electricity provided during the peak hours of noon-4:00PM are primarily generated with natural gas (~400 CO₂(e)/kWh) and the remaining 20 hours correspond to a lower base-load carbon intensity of about 220 CO₂(e)/kWh, the average carbon intensity will remain the same (258 CO₂(e)/kWh). The CO₂(e) emission estimates in this scenario would increase by 2%, 6%, and 3% for these spring, summer, and fall days, respectively, when applying hourly PUE values in Figure 5-11 rather than a daily average PUE value. The hourly changes in PUE value are unique to the climate of each data center location, as indicated by the regional variation in the distribution of hours of

economizer activity in Table 5-36. Greater variation in PUE may be observed in climates less mild than San Francisco or in areas where less efficient peak electricity is employed. Given these potential variations in PUE value when economizers are used in data center cooling systems, evaluation of the annual greenhouse gas emissions associated with data center operation would benefit from determining how the hourly electricity mix for each data center location varies throughout the year. Significant variation in the electricity mix could then be matched to hourly PUE values to better estimate the annual CO₂(e) emissions for data center operation.

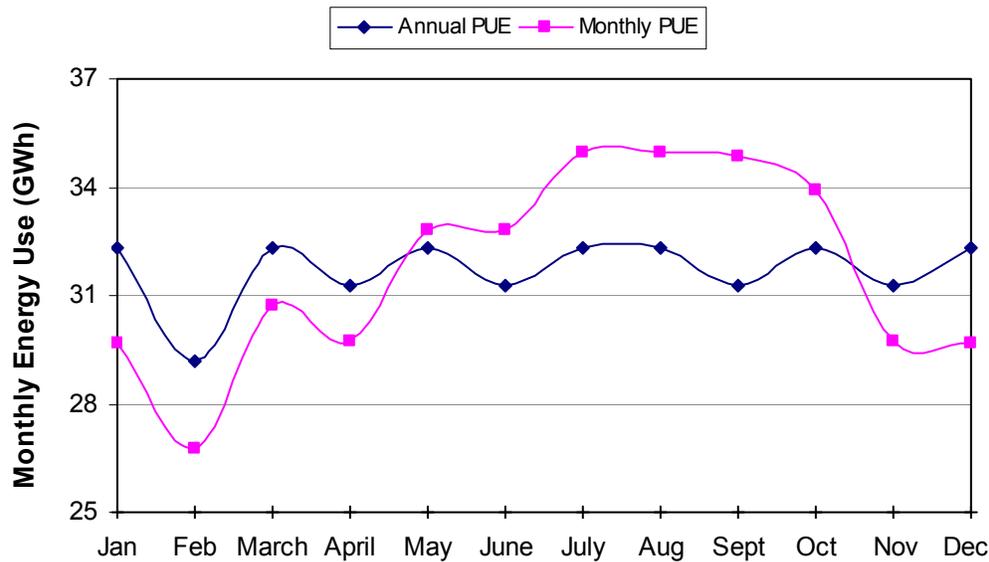


Figure 5-10. Monthly energy use for enterprise data centers in San Francisco under the Economizer Plus scenario. The annual PUE data points represent monthly energy use based on the annual average PUE presented in Table 5-24 (PUE=1.30). The monthly PUE data points represent monthly energy use based on month-specific PUE values, where the PUE increases during warmer months when the economizer operates less hours of the day. The increase in PUE during the warmer months represents a shift of 12 million kWh demand, or 3% of the annual energy use, from the cooler months to the warmer months. The variation in monthly energy use for the annual PUE data points is due to the different number of days in each month.

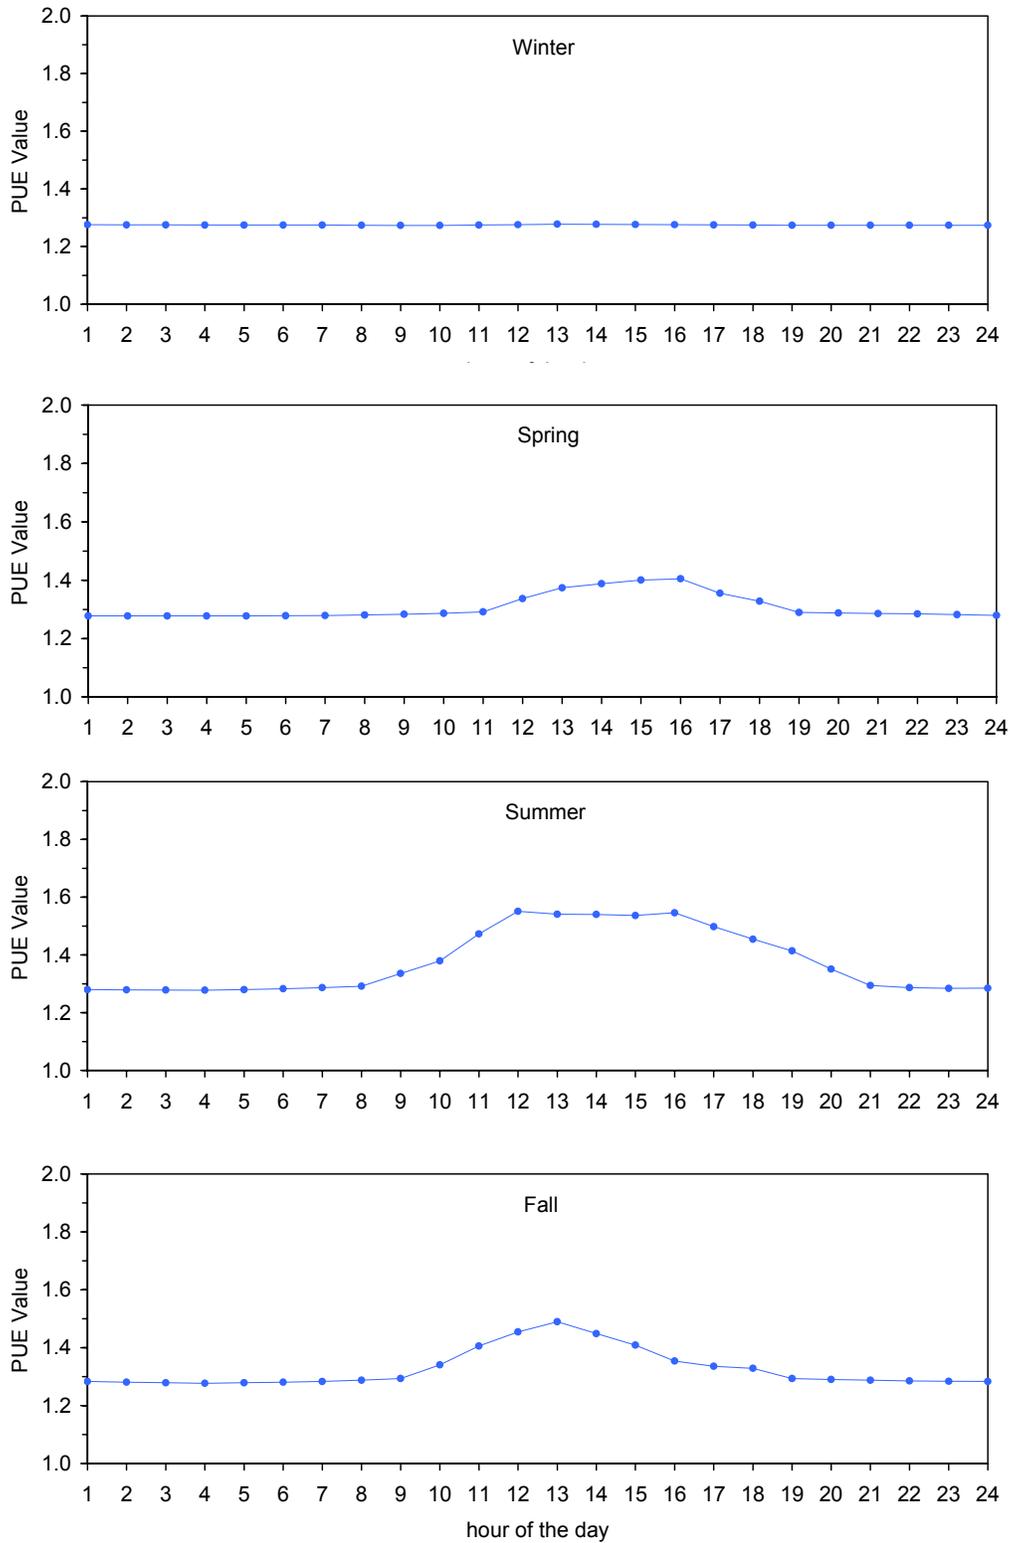


Figure 5-11: Daily variation in PUE value of an enterprise data center during a winter (Jan. 6th), spring (April 10th), summer (July 18th), and fall (Oct. 7th) day in San Francisco under the Economizer Plus scenario (18 °C (64 °F) / 29 °C (84 °F) supply/return temperature setpoints and no humidity restrictions).

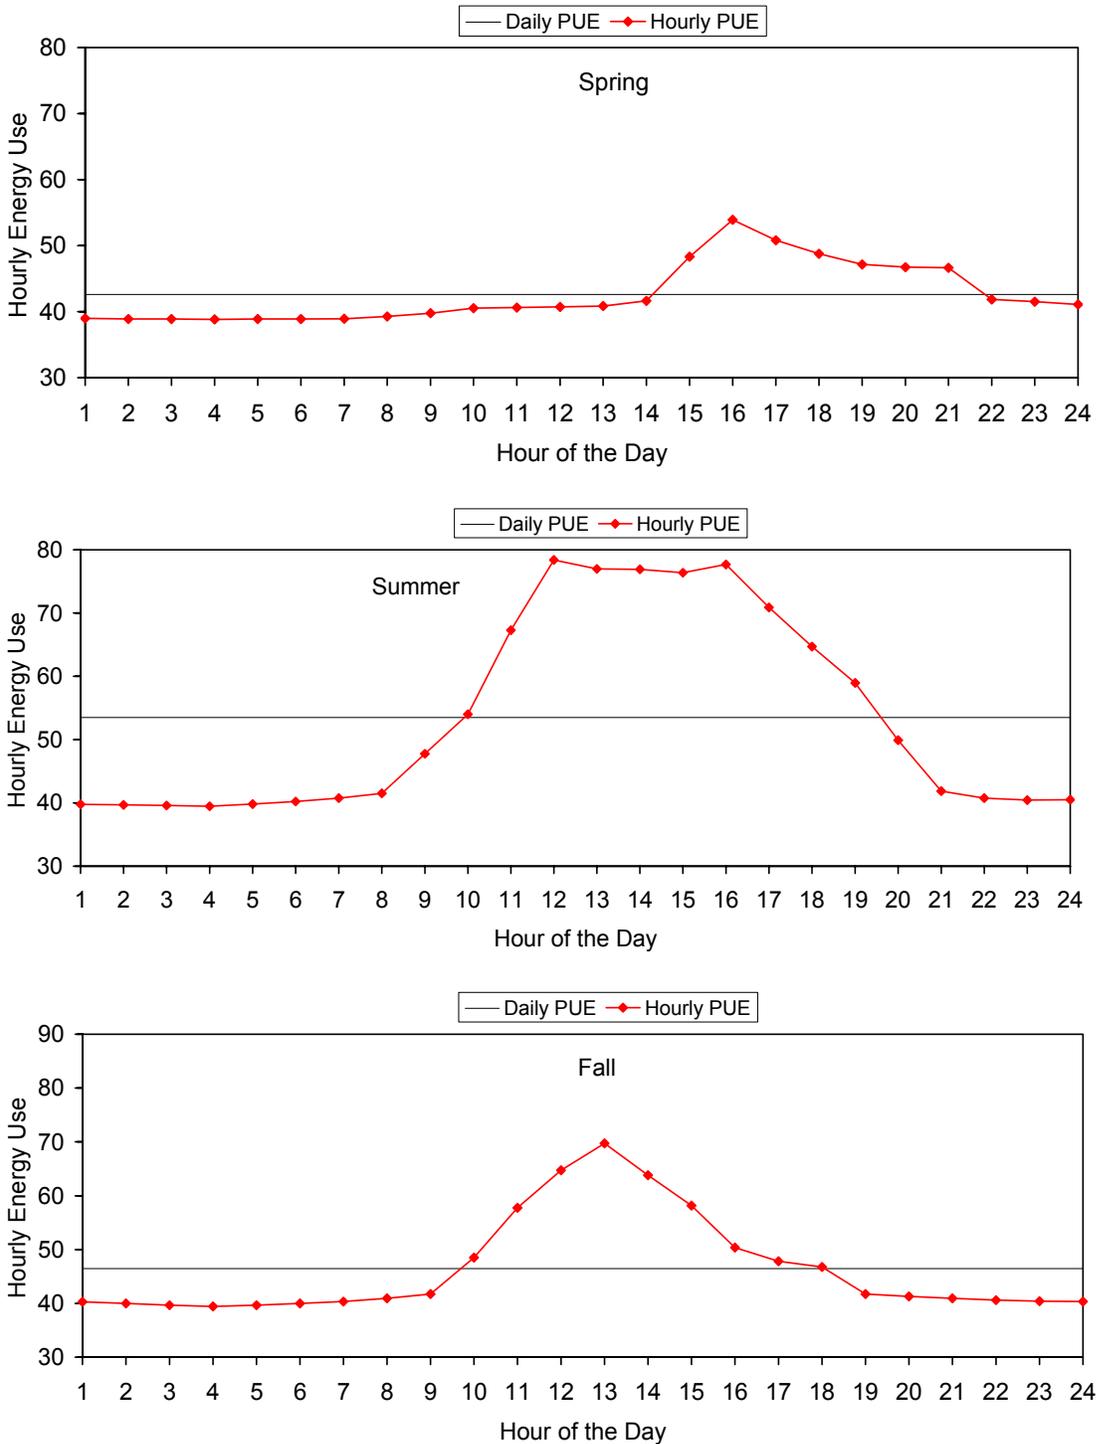


Figure 5-12: Hourly data center energy use (MWh) using the hourly PUE values presented in Figure 5-11. Calculations using an hourly PUE, rather than a daily average PUE, show a shift in energy demand from night and early evening hours to mid-day hours. This shift represents 5%, 16%, and 7% of daily energy demand for the spring, summer, and fall days presented in Figure 5-9, respectively

6. Conclusions

The rapid growth of data center services and the resulting increase in electricity to provide those services highlight the need to pursue energy efficiency opportunities in this sector of the economy. The analysis presented here builds on previous IT equipment energy modeling efforts developed by Brown et al. (2007) and included updated IT equipment stock data. Building space types identified in Bailey et al. (2007) and Brown et al. (2007) are evaluated to account for differences in mechanical equipment and operations. The analysis results indicate about a 65-70% potential reduction of energy use associated with current data center operation, equivalent to an annual energy efficiency resource of nearly 40-50 billion kWh available at the United States level.

The energy estimates in this chapter are based on the best available data at the time of this analysis, but the accuracy of the resulting estimates limited owing to the inherent uncertainties associated with the data, assumptions, and modeling techniques. Estimates of the installed server base quantity and distribution, both of which significantly influence overall energy estimates, are limited to servers that are designated as server products when sold. Server estimates in this chapter do not include custom built servers used by large internet companies (e.g., Google) because no public data about them are currently available. Custom servers were estimated to represent only a small fraction of the total number of U.S. servers as of 2006 (Kooimey, 2007), but the contribution of custom servers to the total energy use in U.S. data centers may increase significantly in the future from the growth of companies such as Google. Also, the estimates of servers installed in the U.S. by space type developed by Brown et al (2007) can only be viewed as general approximations since the spatially disaggregated data is limited to 2005

conditions reported in one study (IDC, 2007) and have not be verified using other data sources.

The modeled PUE values used to estimate the energy contribution from non-IT data center components provide the most comprehensive evaluation to date of the mechanical cooling energy demand from data centers in the United States. However, the modeling results are based on assumptions of equipment efficiency and operation that may differ in practice owing to the many different mechanical designs that could be used in data centers. Furthermore, potential non-IT energy savings associated with closet data centers, while limited, are not included in this analysis because of the challenge of disaggregating this energy use from overall office building energy. Estimates of CO₂ equivalent greenhouse gas emissions from data centers provide an understanding of how changes in this industry can contribute to climate-change mitigation. These modeling efforts, however, would be improved with electricity power mix data specific to each data center location that also identifies diurnal changes in the power mix. Full understanding of the CO₂ equivalent emissions associated with data centers would also require accounting for the embodied energy of the data center buildings, building equipment, and IT equipment.

While the estimates in this chapter are limited by data availability, the results provide insight into data center energy use and highlight areas where potential energy savings can be realized. The efficiency measures applied to the volume servers significantly reduce overall data center energy use, which is consistent with results from Brown et al. 2007) and indicates that the nascent use of server virtualization has the potential to significantly reduce the growth of data center energy use. The modeled PUE

values significantly improve when the temperature setpoints are increased to allow for more hours of economizing in the Economizer Plus scenario for the larger data centers. This finding illustrates the value of data centers operating at the highest temperature allowable without compromising equipment reliability. This measure was not applied to the closet and server room data centers, however, which constrained the total energy savings potential.

Increasing the design temperature difference between the supply and return air is also important to incorporate with economizer implementation since the additional exhaust fans needed with economizer use result in an increase in fan energy, which can be minimized with reduced airflow rates. The modeling results indicate that data centers should explore different temperature setpoints to optimize the balance between chiller and fan energy. While the results presented in Chapter 2 show severe energy penalties when air-side economizer use was paired with ASHRAE recommended humidity restrictions, adding active humidity controls and replacing traditional electric humidifiers with more efficient adiabatic humidifiers results in nearly the same energy demand between the no humidity restrictions and the 20-80% operational range, and only a minor energy penalty is observed with the 40-55% range. The energy use reductions calculated for the energy efficient scenarios also highlight the increased contribution of the data center components where no efficiency measures were assumed, such as the UPS systems and storage equipment. Efficiency advancements in these categories could further improve overall data center energy efficiency.

Overall, PUE results show that larger data centers are typically more efficient than smaller data centers, and this disparity increases in the energy efficient scenarios.

Consequently, the trend toward increased data center consolidation (Carr, 2005) portends potential efficiency improvements in overall data center energy use. The absolute regional difference of total data center energy between the Baseline and Economizer Plus scenarios is minimal even though regional differences in the PUE values significantly increase when economizers are used. This finding highlights the value of the volume server efficiency measures, which have energy saving benefits that carry over to the infrastructure energy demand as well as directly reducing IT loads. However, this result indicates that as data center services increase and more efficiency measures are incorporated, locations of data centers will have a greater effect on their overall energy demand. Future data center development will need to consider site location, along with IT and non-IT efficiency measures, when attempting to minimize the environmental impact attributable to this increasingly prominent economic sector.

Chapter 6: Conclusion

This chapter reflects on the results and discussions from the previous chapters and proposes future research based on a holistic assessment of the conclusions reached during the research process. The dissertation closes with final thoughts for moving forward.

6.1. Economizer implementation and energy savings potential

Throughout all building sectors, improving building energy efficiency has the potential to reduce global energy use and curb the greenhouse gas emissions associated with building operation. The rapid growth and high power density of data centers highlight the importance of energy efficiency in the operation of these buildings. This dissertation estimates data center energy use and explores energy efficiency strategies that can significantly reduce data center energy demand. Much of this dissertation focuses the use of economizers to cool IT equipment with large volume of outside air. Economizers can potentially change the composition of indoor air in contact with the IT equipment in the data center, which expands the evaluation of economizer use to include the impact on IAQ. This dissertation contributes to understanding the relationship between energy efficiency and IAQ in data centers by evaluating how indoor particle concentrations and building energy demand change under different design strategies. The buildings that support the growing IT economy are the focus of this dissertation, but the framework presented to evaluate economizer use and particulate matter in data centers could contribute to addressing the IAQ and energy efficiency of other building types.

With IAQ concerns previously identified as a barrier to economizer implementation (Tschudi et al., 2004), Chapter 2 begins addressing this issue by first identifying particulate matter as a prominent IAQ concern and then quantifying the

impact of economizer use on data center indoor particle concentrations. Evaluating the potential for changes in indoor particle concentrations to noticeably affect equipment reliability proved to be an insurmountable challenge. Documented failure data of IT equipment are not publicly available owing to the private nature of the data center business structure. IAQ concerns, and the energy decisions they influence, were identified to be based on a combination of anecdotal and theoretical evidence, with the most specific articulated concern being current leakage owing to the potential deliquescence of deposited hygroscopic particles between isolated conductors in IT equipment (Weschler, 1991). Theoretically, such current leakage would require humidity levels to be above the deliquescence point specific to the hygroscopic species. It would also require enough particles to deliquesce so as to create a current bridge between the isolated conductors. However, the specific environmental parameters required for IT equipment damage to occur are not clear because of the lack of documented equipment failure. Identified particle concentration limits specified in guidelines ranged by an order of magnitude, from $15 \mu\text{g}/\text{m}^3$ to $150 \mu\text{g}/\text{m}^3$, with no indication that these guidelines are based on any systematic of failure analysis (ASHRAE, 2009). The particle measurements reported in Chapter 2 represent an important early step in addressing this unknown potential damage by characterizing particle concentration levels in data centers and quantifying how economizer use affects indoor particulate matter concentrations. Identified changes in particle concentrations owing to economizer use could then be addressed either by determining if these differences have a distinguishable impact on IT equipment reliability or by engineering methods to remove these differences altogether. Size and time resolved OPC measured concentrations of particles $0.3\text{-}5.0 \mu\text{m}$ in diameter

were used as a proxy to represent changes in particle concentration and to provide an estimate of hygroscopic particle concentrations to which IT equipment is exposed in data center environments. Results from multiple data centers revealed that particle concentrations without economizers were consistently less than $1 \mu\text{g}/\text{m}^3$. This concentration is significantly lower than typical indoor or outdoor particle concentrations and is the result of the abnormally high levels of air recirculation in these buildings, even though relatively low-efficiency air filtration is typically used. When economizers were used, the measured particle concentration increased to about $10 \mu\text{g}/\text{m}^3$. This level is still lower than typical indoor or outdoor concentrations and below even the strictest concentration guidelines identified; however, this particle concentration level increase was an order of magnitude higher than the concentrations measured in the non-economizer data centers. Material balance modeling advanced the understanding of particle sources and sinks in data centers, indicating that there were no significant indoor particle sources and confirming that the particles measured in the data centers were primarily of outdoor origin. The modeling effort also showed that indoor/outdoor ratio of specific hygroscopic particles of concern, such as ammonium sulfate, would likely increase relative to the values measured for particle mass concentration because of differences in the particle size distributions.

Chapter 3 broadened the understanding of particle concentrations in data center that was established in the previous chapter and explored a strategy to overcome one of the barriers to economizer implementation. Along with measuring particles $0.3\text{-}5.0 \mu\text{m}$ in diameter, additional monitoring equipment was used to measure chemical constituents and other attributes of particles to determine if the trends seen with economizer use in

Chapter 2 extend to particles with different physical and chemical characteristics. Collected data showed that economizer-induced changes in concentration for many components of particulate matter could be eliminated with the use of enhanced (MERV 14) filtration. Metered electricity data showed that enhanced filtration had minimal effect on the overall energy savings, allaying the concern that enhanced filtration would negate the energy benefits of economizer use. Theoretical calculations indicated an approximately 10% increase in fan power, from 64 kW to 70 kW, due to the increased pressure drop associated with enhanced filtration, a difference that is small relative to the chiller savings of approximately 100 kW from economizer use. Since the relationship between particle concentration and equipment failure in data centers is poorly understood, any economizer-induced increase above the levels measured in conventionally operated data centers can be construed to be potentially damaging and accordingly hinder deployment. While the actual effect of the economizer-induced particle concentration increase was not evaluated, the results of Chapter 3 showed that pairing economizer use with enhanced filtration could avoid a particle concentration increase, suggesting that elevated particle contamination risk associated with economizer use does not justify avoiding the use of this technology in data centers.

With empirical data supporting implementation of a method to avoid the increases in particle concentration from economizer use without significantly affecting energy efficiency, the focus of the dissertation shifted in Chapter 4 to evaluating the energy savings available from using economizers. Energy models were used to quantify data center energy use in different California climates and the PUE was introduced as a metric to compare the energy efficiency performance of different data center mechanical

designs. PUE values were calculated for a baseline mechanical design (no economizer), a typical economizer system that increases the ventilation air exchange rate (air-side economizer), and an alternative economizer design that maintains the baseline ventilation rate (water-side economizer). The results provide insight into the energy efficiency and relative value of each design. The calculated baseline mechanical design PUE values were considerably better than average measured PUE estimates (Greenberg et al., 2006, Sullivan, 2009), highlighting a disparity between modeled and measured data center building performance. The modeled baseline PUE values represent large data centers built to meet current building design standards, so part of this disparity could have been due to smaller data centers with older, less efficient mechanical equipment, which was not accounted for in the scope of this chapter. Poorly operating equipment and inefficient airflow design could also contribute to this modeled-measurement disparity, a finding that emphasizes the value of commissioning mechanical equipment and properly managing airflow in currently operating data centers. The water-side economizer system, which could be an appealing energy efficiency measure that avoids the need to address the ramification of potential changes in IAQ, was less efficient than the air-side economizer in the five California climate zones modeled. This finding highlights a potential missed energy savings opportunity when water-side economizers are specified rather than air-side economizers owing to IAQ concerns.

Modeling the energy demand of data centers equipped with traditional economizers (air-side economizers) revealed that, relative to the baseline (non-economizer design), energy demand is significantly dependent on other operational parameters, such as temperature and humidity restrictions. Modeling each mechanical

design under varying levels of humidity restrictions showed that, for the five California climate zones modeled, the energy demand remained fairly constant for the baseline and water-side economizer designs, while increasing the humidity restrictions under the air-side economizer design drastically increased energy demand. In fact, under the narrow ASHRAE-recommended humidity range of 40-55% RH (ASHRAE, 2005), the economizer design was less efficient than the baseline design in most of the California climate zones modeled because of that extra energy required for latent cooling and humidification. The temperature setpoints for all modeled mechanical designs were held fixed to maintain consistency between the simulations, but the modeling results showed that these temperature setpoints limited the energy savings potential of economizer use. Increasing the temperature setpoint will reduce chiller demand in any mechanical system, but the economizer design would receive the added benefit of increasing the number of hours throughout the year that data center cooling can be met without operation of the chiller. The insight gained from these modeling results stressed the need for conventional humidity and temperature settings to be evaluated and possibly changed when data centers incorporate economizers into the mechanical design.

Many of the lessons learned during the modeling presented in Chapter 4 were incorporated into the expanded modeling effort presented in Chapter 5. Data center energy models were developed to compare the difference in energy demand between a baseline (no economizer) and an economizer system employing high outside air ventilation rates. For both systems, different data center space types were identified with the mechanical efficiency and operation specified accordingly. This approach allowed the PUE results to better represent the current stock of data centers. However, the

modeled baseline mechanical system PUE values were still below the estimated national average of approximately 2.0 for the current stock of data centers (Brown et al., 2007; Koomey, 2007). This 2.0 PUE estimate is supported by the small amount of empirical data available from Greenberg et al. (2006) and Belady and Malone (2007), the industry consensus established in Brown et al. (2007), as well as recent industry survey data gathered by the U.S. Environmental Protection Agency (Sullivan, 2009). Improper airflow management in data centers, such as hot exhaust air mixing with cold supply air, is one potential source of the modeled-measurement disparity. To account for this disparity, the modeled baseline PUE values were scaled to a weighted average of 2.0. This scaling allowed current data center energy use estimates to account for efficiency differences in data center space type while still representing current national PUE estimates. Energy use for each space type was modeled in different national climate zones, each representing an area identified to have significant data center activity. PUE values improved by approximately 5-25% with the economizer design, with minor variation between climate regions, but significant variation between space types. Greater PUE improvement and absolute efficiency were observed in larger data centers, indicating an energy benefit from consolidating IT equipment in large data centers.

IT equipment data were gathered to update results from a previously developed, bottom-up approach for estimating total national IT energy demand (Brown et al., 2007). IT energy demand itself was estimated for both a baseline and an energy-efficient scenario, which incorporated established but nascent IT efficiency measures identified by Brown et al (2007). The modeled PUE values were applied to the IT estimates, indicating that current total data center energy demand is about 62-70 billion kWh

annually. This total data center energy demand dropped to about 21-23 billion kWh under the energy efficient scenarios, which represents an energy savings potential that includes the use of economizers as well as IT efficiency measures. Similar to results found in Brown et al. (2007), much of this 65-70% potential to decrease energy use in data centers is attributable to energy-efficiency measures for volume servers, which account for a large amount of the overall IT energy demand and, accordingly, a large amount of the heat generated in data centers that must be removed by the mechanical system.

PUE values were calculated for several variations of humidity and temperature operating conditions, and for each data center space type. Given the energy penalty observed in Chapter 4 from humidity restrictions, more efficient adiabatic humidifiers replaced previously used electric humidifiers in this modeling iteration. This humidifier change significantly reduced the energy penalty observed from humidity restrictions and, with the climate regions now expanded to the national level, the ASHRAE “allowable” humidity restrictions (20-80 RH) actually reduced overall energy use in some of the more humid climates. Increasing the temperature of the supply and return air in the economizer scenario significantly increased energy efficiency by reducing the number of hours during the year that require chiller operation. These results emphasize the potential value of increasing the temperature setpoints and the need to consider the energy savings and IT equipment reliability risk associated with including this operational change when using economizers. Greater flexibility in temperature setpoints can also help reduce fan energy, which becomes the prominent component of non-IT energy demand as expanded economizer use reduces the need for chiller operation. Maintaining a lower supply air

temperature, but exposing the IT equipment to higher air temperatures with a lower airflow rate results in reduced fan energy. This operational design shifts some of the energy required for fans to the cooling demand (due to the lower supply air temperature), but since much of the cooling can be met with economizers this design achieves a net energy savings benefit. This finding indicates that future research on data centers should explore different temperature setpoints to optimize the balance between chiller and fan energy.

Chapter 5 concludes by evaluating the energy savings potential available from data centers in relation to greenhouse-gas emissions. Modeling results show that the regional variation in non-IT equipment energy use shifts from being relatively minor (3-5%) to about 30% when the mechanical design includes economizers, highlighting the increased importance of location to data center efficiency when economizers are used. This increased difference, however, is dwarfed by the potential variation in regional carbon intensity associated with electricity production. Even the substantial 65-70% savings potential available through data center efficiency is small relative to the order-of-magnitude difference between the estimated CO₂(e)/kWh emissions for Seattle (30 CO₂(e)/kWh), which relies primarily on hydro-generated electricity, compared to the emissions for Chicago, Dallas, or Richmond (493, 587, 511 CO₂(e)/kWh, respectively), where nearly half of the electricity is generated from coal. In San Francisco, the carbon intensity was estimated in the middle of this range at 258 CO₂(e)/kWh owing to the significant use of natural gas in this region. Monthly and daily variation in carbon intensity by location may also play a role when considering that mechanical efficiency can also vary along these timescales. The flexibility of data center location may be

greater than other building types, which are more restricted by the desired location of the building occupants. The results in Chapter 5 indicate that site location will play a significant role in affecting the CO₂(e) emissions associated with data center operation. Along with the increased mechanical efficiency gains available in cooler climates, the carbon intensity associated with regional electricity generation could also influence future data center locations.

6.2. Opportunities looking forward

The dissertation contributes to the field of understanding and improving building energy efficiency by exploring the relationship between IAQ and operational energy demand in data centers. Specifically, the effect of economizer use on particle concentrations in data centers has been studied and the potential energy savings from this mechanical design have been estimated and compared to estimates of current United States data center energy use. This section builds on the knowledge acquired and described throughout this dissertation to propose and outline other research opportunities where exploration may further improve data center energy efficiency and minimize the environmental impact of this building sector.

6.2.1. Expansion of data center operating conditions

Operating data centers under traditional environmental conditions hinders the energy savings potential of economizers. These conditions include low cooling temperatures, tightly controlled relative humidity, and minimal indoor particle levels. The need for these strict operating conditions has been accepted without much question

within the industry, since meeting these conditions requires relatively little additional energy demand in conventional non-economizer data centers. With economizers being used to curb the increase in data center electricity demand, however, the potential loss in energy efficiency to maintain such strict operating conditions is much higher. The economizer implementation strategies discussed in this dissertation primarily focused on pairing additional technologies with economizer use to maintain strict operating conditions while still achieving substantial energy savings. Chapter 3 showed that operating economizers with enhanced filtration allows indoor particle concentrations to remain at levels measured in non-economizer data centers with only a minor energy penalty. Chapter 5 showed that almost all of the economizer energy savings could be maintained under ASHRAE humidity restrictions (ASHRAE, 2005) when less efficient electric humidifiers are replaced with more efficient adiabatic humidifiers. The need to include additional equipment and to compromise some economizer savings, however, may not be necessary and deserves evaluation. More importantly, Chapter 5 showed that lower data center temperature setpoints significantly reduce economizer energy savings. Lower supply air temperatures in traditional non-economizer data centers result in greater chiller demand, but this demand is exacerbated in economizer-equipped data centers as lower temperatures also reduce the number of hours that the economizer can operate. As economizer use becomes more prevalent, and given the cost of these strict operating conditions, it would be valuable to determine if such operating conditions are necessary to maintain high reliability and, if so, determine more energy efficient ways to achieve similar levels of reliability.

Often cited concerns for maintaining traditional operating conditions include hot spots/overheating, electrostatic discharge (ESD), and conductor bridging. Many data centers operate at or below the lower bound of recommended temperature ranges. Data center managers are hesitant to increase setpoint temperatures, even within operating guidelines, because of concern about “hot spots” being created by poor airflow conditions in high-density areas (Miller, 2008a). The potential for hot spots to cause some servers to overheat and malfunction may motivate data center managers to operate at temperatures lower than necessary. Moisture levels in data centers are regulated to prevent both high and low relative humidity extremes. High moisture levels are thought to cause condensation on electronic components in data centers (Miller, 2007), resulting in numerous equipment problems (ASHRAE, 2009). Low moisture levels can potentially increase electrostatic charge generation and accumulation (Swenson and Kinnear, 2009). As a result of these concerns, humidity in data centers is often restricted to a narrow range using humidification controls. Chapter 2 discussed how deliquescent particles can potentially deposit between isolated conductors on electronic circuit boards and, under increased humidity, dissociate to become electrically conductive (Weschler, 1991). This increased conductivity could lead to electronic equipment failure (Litvak et al., 2000), which dissuades any operational changes that may cause indoor particle concentrations to deviate from conventional conditions.

Identified established modes of failure could be empirically induced under the measured temperature, humidity, and IAQ conditions to better understand how these failures manifest and to identify vulnerable equipment components. This empirical process will also help determine the actual benefit (increase in reliability) from adhering

to the traditional environmental parameters. There is an increasing amount of anecdotal evidence indicating that strict operating conditions provide minimal, if any, reliability benefit. For example, Microsoft observed no server failure while operating within a tent in Washington with no climate or air quality controls between November 2007 and June 2008 (Miller, 2008b). During an economizer study, Intel (2008) exposed servers to considerable variation in temperature and humidity, with minimal filtration of particles, and only recorded a small increase in equipment failure that was similar to baseline expectations. Recent research indicates that humidity control may have little effect on the potential for ESD issues in data centers (Swenson and Kinnear, 2009). Current leakage is only expected to occur under high humidity, which is unlikely at the circuit board given the elevated temperatures within the servers.

Once the temperature, relative humidity, and particle levels have been expanded enough to cause unacceptable failure rates, the theoretical basis for these methods of failure should be explored. An improved understanding of the failure mechanism could allow metrics for the environmental thresholds to be developed and adjusted to changes in IT equipment. For example, empirical failure induced through current leakage could be evaluated to determine if failure can be predicted by the physical bridging of isolated regions with particles through deposition and percolation theory, which has been previously proposed (Weschler, 1991). Modeling methods established to match this empirically observed mode of failure could then be applied to IT equipment with different types of isolated conductors or exposed to different particle concentrations. Knowledge gained through the measured and modeled failure studies would help develop strategies to fortify IT equipment against identified failure methods. The goal of this

fortification would be to maximize economizer savings through chillerless data center operation or, given the high fan energy noted in Chapter 5, a combination of low airflow complemented with modest chiller use. The energy and cost associated with increased server fortification or reduced characteristic time for IT replacement should be compared to potential energy savings gained through expanded operational conditions. These results could lead to a commercialization path that would involve working with equipment manufacturers to develop a certification and labeling process for server operation under less tightly controlled environmental conditions and expanded hours of economizer use. The certification may include different tiers of environmental conditions and indicate anticipated equipment lifetimes.

6.2.2. Improvement of metrics

Improved metrics for some of the areas explored in this dissertation could greatly assist the future research of economizer use and data center IAQ, which could ultimately help increase data center energy efficiency.

Air quality is a concern in data centers due to the potential for pollutants to adversely affect the IT equipment and compromise reliability. In Chapter 2, concentrations of particles 0.3-5.0 μm in diameter were used to measure and compare data center IAQ. Chapter 3 attempted to better tailor the IAQ comparison for data center concerns by measuring concentrations of chemically specific particles that have been identified to impact the reliability of electronic equipment. This improvement, however, still does not address what aggregate effects these individually measured concentrations may exhibit or include how other pollutants, such as gases, may contribute to eventual

equipment damage. The purpose of IAQ measurements in data centers is to understand the risk to IT equipment and an IAQ metric is needed to properly assess that risk. Such a metric could be developed through equipment testing, which could identify pollutant sources and types of equipment failure. Each potential method of failure would be a function of the responsible pollutants, allowing the most immediate mode of failure to be isolated for any measured set of pollutants. Such a metric could be expressed as a timescale value, indicating the IT equipment exposure time to such air quality before experiencing increased rates of failure. An effective data center IAQ metric could also be represented through a responsive architecture design that would act as a precautionary system to warn data center operators of critically high pollutant concentrations. The system could be similar to VESDAs (Very Early Smoke Detection Apparatus), which are currently used in many data centers as an early detection fire alarm. Another option is the use of a digitized coupon, which is an apparatus designed to be affected by particle bridging at a faster rate than server equipment, thus able to act as a harbinger to obviate equipment problems. Such responsive architecture could measure data center IAQ in terms of system replacement or resetting frequency.

Improving the metrics for data center energy use could also be beneficial in promoting energy efficient design and operation. PUE (Power Use Efficiency) is the metric currently used to quantify and compare the energy efficiency of data center buildings. In Chapter 4, PUE was defined as the ratio of total building energy use to IT energy use. This metric highlights the efficiency of non-IT equipment, such as the HVAC equipment, UPS losses, and the building lighting system by quantifying this energy as a function of the energy consumed by the IT equipment. Since total building

energy and energy dedicated to the IT equipment are often metered separately, this metric is simple to apply. However, it becomes less useful as the line between IT and non-IT energy use is blurred in more advanced data center design. Figure 4-3, for example, showed that power demand in servers is distributed to many different components and some, such as the internal fan and PSU losses, are not part of the actual IT processing. More efficient IT equipment orientation could remove the internal fan and provide all airflow through the HVAC fans, which are more efficient owing to their larger size. This energy efficient approach, however, would be penalized using the PUE metric as energy previously attributable to IT operation would now be considered non-IT energy. A more advanced PUE metric would designate only power dedicated to digital processing, storage, or networking as IT power. Further improvement of the PUE metric would be to account for the efficiency of the IT processing power itself. Do so would allow data center efficiency to be normalized to actual IT service rather than simply IT energy demand. Developing such a metric is a challenging endeavor, but initial efforts could focus on measuring services by the processing power required to provide the energy intensive activities that dominate IT energy demand, such as the transfer of rich media. Developing metrics for IT service would also provide the opportunity to compare the associated energy use to that needed to provide similar services through alternative means, for example, comparing energy use between music downloads and compact disks (Weber et al., 2009). Improving data center energy-efficiency metrics through such an approach would not only promote more efficient IT services, but would also highlight and promote IT services that displace energy that would otherwise be used through other non-IT avenues.

6.2.3. Sourcing of electricity and greenhouse-gas emissions

Minimizing the impact of data center operation on climate change would require evaluating the energy demand for these buildings within the context of emitted greenhouse gases. Chapter 5 showed that modeled regional differences in data center energy efficiency vary much less than the potential range of greenhouse gas emissions associated with electricity production in different regions. Furthermore, Chapter 5 showed that as the fluctuation in data center demand increases with the implementation of economizers, which only operate during cooler periods, seasonal and daily variation in electricity source mix could potentially affect total data center greenhouse gas emissions. Cataloging the greenhouse gas emissions associated with providing electricity for different data center locations and periods can help promote data center design measures that reduce the greenhouse-gases emitted as a function of IT energy use or IT service provided. A successful greenhouse gas accounting approach could correlate emissions with electricity prices, for both the location of the electricity service and for time-of-day pricing. Expressing both the costs and the emission benefits could drive mitigation efforts such as peak power savings through the use of thermal storage and on-site generation.

6.2.4. Embodied energy

Evaluating engineering strategies to reduce the energy use and the greenhouse gas emissions associated with providing IT services would benefit from expanding the scope of evaluation to include the embodied energy related to data centers. Along with the

energy required to provide electricity to data centers, IT services also consume energy through the life cycle of the data center building, and of both non-IT and IT equipment. Energy is used during the manufacturing of building materials, such as concrete, as well as during the design, construction, maintenance, and end-of-life phase of the building itself. Non-IT equipment that may contain significant embodied energy includes complex HVAC chillers and pumps. The high turnover rate of IT equipment itself increases the importance of accounting for the embodied energy from this component of data centers. Equipment and materials with potentially high embodied energy could be identified and life-cycle assessments could be performed to quantify energy inputs, which could then be normalized based on estimated lifetimes for buildings (~50 years), mechanical equipment (~15 years), and IT equipment (~5 years). Many of the strategies discussed in this dissertation to reduce operational energy and greenhouse-gas emissions can potentially affect the contribution of embodied energy. Operating IT equipment under broader environmental conditions may reduce operational energy for cooling but could increase the flow of IT equipment through the data center. Chillerless data center cooling would reduce the need for chiller equipment, but Chapter 5 noted improved operational energy savings when lower fan energy was balanced with moderate chiller activity. Considering embodied energy in data center design decisions could increase the opportunities for improving the contribution from data centers in addressing climate change, and ultimately provide a more environmentally benign IT infrastructure.

6.3. An opportunity for sustainability

Data center energy estimates and projections indicating the rapid growth of this economic sector add to an already formidable societal challenge of curbing greenhouse gas emissions. The evolution and expansion of IT services, however, also portend potential opportunities to manage global energy demand as the standard of living increases throughout less industrialized regions of the world. Future IT services, and the data centers that support them, may be able to supplant more energy intensive services currently associated with more advanced economies. For example, simply shifting to an economy focused on the digital, rather than physical, transfer of goods and services creates two opportunities to use energy more efficiently. First, digital transfer may reduce the need for many of the energy-intensive processes involved with manufacturing, packaging, and transportation. Jonathan Koomey (2009) expressed this possibility by stating that, “Moving electrons is always less environmentally damaging than moving atoms.” Initial evaluations have shown potential savings from telecommuting (Atkyns et al., 2002; Kitou and Horvath, 2003), methods of print-product delivery (Toffel and Horvath, 2004), and online retail (Mathews et al., 2001; Hendrickson et al., 2006; Weber et al., 2008). Energy savings through the complete dematerialization of music delivery products (i.e., compact disks) to music downloads (Weber et al., 2009) may also be indicative of future energy savings potential through IT markets. Second, the increased consolidation of data centers combined with the digital transfer of goods and services will result in fewer buildings representing a greater portion of global energy demand. The consolidation of energy demand allows for efficiency efforts that are concentrated on this economic sector to reap significant benefits.

This dissertation contributes to these efficiency efforts by quantifying potential energy savings and identifying design parameters that significantly influence this estimated energy demand. Through the application of air quality engineering and energy analysis, this dissertation provides insight into the energy and IAQ impact of economizer use in data centers. Methods are identified and evaluated to reduce data center energy use while maintaining a low indoor particle concentration to ensure IT equipment reliability. Results provide information about indoor particle concentrations and energy use to help data center designers, operators, and owners make more informed decisions. Results are also applicable for use by policy and decision makers and provide the foundation for future prescriptive guidelines and performance metrics that can be applied to data centers. Current data center design and operation require that nearly half of the building energy demand be dedicated to non-IT equipment (Tschudi et al., 2004; Greenberg et al., 2006; Sullivan, 2009). As this non-IT energy demand is reduced through better design, IT can evolve into a more energy-efficient service with the potential to facilitate a more sustainable expansion of goods and services.

References

- Airguard, 2009. New product cut sheets. www.airguard.com/new.html.
Last accessed January 12, 2009.
- AMD, 2006. Power and Cooling in the Data Center. Advanced Micro Devices. 34246C.
http://enterprise.amd.com/Downloads/34146A_PC_WP_en.pdf.
Last accessed Dec. 10, 2009.
- Andreae, M.O., Gelencser, A., 2006. Black carbon or brown carbon? The nature of light-absorbing carbonaceous aerosols. *Atmospheric Chemistry and Physics* 6, 3131–3148.
- ASHRAE, 1992. Gravimetric and dust spot procedures for testing air-cleaning devices used in general ventilation for removing particulate matter. ANSI/ASHRAE Standard 52.1-1992. American Society of Heating, Refrigerating, and Air-Conditioning Engineers, Inc., Atlanta, GA.
- ASHRAE, 1999. Method of testing general ventilation air-cleaning devices for removal efficiency by particle size. ANSI/ASHRAE Standard 52.2-1999. American Society of Heating, Refrigerating, and Air-Conditioning Engineers, Inc., Atlanta, GA.
- ASHRAE, 2005. ASHRAE Handbook, HVAC Systems and Equipment. American Society of Heating, Refrigerating, and Air Conditioning Engineers, Atlanta, GA.
- ASHRAE, 2009. Particulate and gaseous contamination in datacom environments. American Society of Heating, Refrigerating, and Air-Conditioning Engineers, Inc., Atlanta, GA.
- Atkyns, R., Blazek, M., Roitz, J., 2002. Measurement of environmental impacts of telework adoption amidst change in complex organizations: AT&T survey methodology and results. *Resources, Conservation, and Recycling* 36, 267-285.

- BAAQMD, 2009. Forecasting and data analysis, Bay Area Air Quality Management District. www.baaqmd.gov. *Last accessed Jan. 5, 2009.*
- Bailey, M., Eastwood, M., Grieser, T., Borovick, L., Turner, V., Gray, R.C., 2007. Special Study: Data Center of the Future. New York, NY: IDC. IDC #06C4799.
- Belady C.L., Malone, C.G., 2007 Metrics and an infrastructure model to evaluate data center efficiency. IPACK2007-33338: Proceedings of the ASME InterPACK 2007, Vancouver, BC.
- Bodik, P., Armbrust, M., Canini, K., Fox, A. Jordan, K., Patterson, D., 2006. A Case for Adaptive Datacenters to Conserve Energy and Improve Reliability. Berkeley, CA: University of California, Berkeley Reliable Adaptive Distributed (RAD) Systems Laboratory. <http://radlab.cs.berkeley.edu/>. *Last accessed Dec. 10, 2009.*
- Brown, R., Masanet, E., Nordman, B., Tschudi, B., Shehabi, A., Stanley, J., Koomey, J., Sartor, D., Chan, P., Loper, J., Capana, S., Hedman, B., Duff, R., Haines, E., Sass, D., Fanara, A., 2007. Report to Congress on server and data center energy efficiency: Public Law 109-431. Report LBNL-363E, Lawrence Berkeley National Laboratory, Berkeley, CA.
- Brusse, J., Sampson, M., 2004. Zinc whiskers: hidden cause of equipment failure. IT Professional 6, 43-47.
- Butler, D., 2008. Architects of a low-energy future. Nature 452, 520-523.
- Carr, N.G., 2003. IT Doesn't Matter, Harvard Business Review, May: 41-49.
- Carr, N.G., 2005. The End of Corporate Computing. MIT Sloan Management Review, 46:3, 66-73.
- CB ECS, 2007. Commercial building energy consumption survey. Energy Information Administration. www.eia.doe.gov/emeu/cbecs/contents.html. *Last accessed Dec. 10, 2009.*

- CEC, 2005. Nonresidential Compliance Manual For California's 2005 Energy Efficiency Standards. California Energy Commission, Sacramento, CA.
- Chow, J.C., Watson, J.G., 1998. Guideline on speciated particulate monitoring. Report prepared for US Environmental Protection Agency, Research Triangle Park, NC, by Desert Research Institute, Reno, NV.
- Chow, J.C., Watson, J.G., Lu, Z., Lowenthal, D.H., Frazier, C.A., Solomon, P.A., Thuillier, R.H., Magliano, K., 1996. Descriptive analysis of PM_{2.5} and PM₁₀ at regionally representative locations during SJVAQS/AUSPEX. Atmospheric Environment 30, 2079–2112.
- Christensen, D., 1996. Anionic analysis by ion chromatography. IBM unclassified report 20260.01, International Business Machines Corporation, Rochester, MN.
- Dell, 2007. Data center efficiency in the scalable enterprise. Dell Power Solutions, <http://www.dell.com/downloads/global/power/ps1q07-20070210-CoverStory.pdf>.
Last accessed Dec. 10, 2009.
- EIA, 2009a. Net generation by energy source by type of producer. Energy Information Administration. <http://www.eia.doe.gov/cneaf/electricity/epa/epat1p1.html>.
Last accessed Dec. 18, 2009.
- EIA, 2009b. Emissions of greenhouse gases in the United States 2008. Energy Information Administration.
<ftp://ftp.eia.doe.gov/pub/oiaf/1605/cdrom/pdf/ggrpt/057308.pdf>.
Last accessed Dec. 18, 2009.
- EIA, 2008. Use of energy in the United States. Energy Information Administration. <http://tonto.eia.doe.gov/energyexplained>. *Last accessed Dec. 10, 2009.*
- EPA, 2009. EPA announces Energy Star label for computer servers. United States Environmental Protection Agency, <http://www.energystar.gov/datacenters>.
Last accessed Dec. 10, 2009.

- EPA, 2008. Emissions and Generation Resource Integrated Database 2007 version 1.1 year 2005 summary tables. United States Environmental Protection Agency, [www.epa.gov/cleanenergy/documents/egridzips/eGRID2007V1_1_year05_Summary Tables.pdf](http://www.epa.gov/cleanenergy/documents/egridzips/eGRID2007V1_1_year05_Summary_Tables.pdf).
- Eubank, H., Swisher, J. Burns, C., Seal, J., Emerson, B., 2003. Design Recommendations for High-Performance Data Centers: Report of the Integrated Design Charrette. Snowmass, CO: Rocky Mountain Institute. February 2-5.
- Fthenakis, V.M., Kim, H.C., 2007. Greenhouse-gas emission from solar-electric and nuclear power: A life-cycle study. *Energy Policy*, 35(4), 2549–2557.
- Fan, X., Weber, W., Barroso, L.A., 2007. Power Provisioning for a Warehouse-sized Computer. Proceedings of the 34th International Symposium on Computer Architecture in San Diego, CA. Association for Computing Machinery, ISCA '07, http://labs.google.com/papers/power_provisioning.pdf. *Last accessed Dec. 10, 2009.*
- Fisk, W.J., Faulkner, D., Palonen, J., Seppanen, O., 2002. Performance and costs of particle air filtration technologies. *Indoor Air* 12, 223-234.
- Fisk, W.J., Faulkner, D., Sullivan, D., Mendell, M.J., 2000. Particle concentrations and sizes with normal and high efficiency air filtration in a sealed air-conditioned office building. *Aerosol Science and Technology* 32, 527-544.
- Fontecchio, M., 2007. Data center humidity levels source of debate. SearchDataCenter.com, June 18.
- Fowler, G.A., Worthen, B., 2009. The Internet industry is on a cloud -- whatever that may mean. *The Wall Street Journal*. <http://online.wsj.com/article/SB123802623665542725.html>. *Last accessed Dec. 10, 2009.*
- Greenberg, S., Mills, E., Tschudi, W., Rumsey, P., Myatt, B., 2006. Best practices for data centers: Results from benchmarking 22 data centers. Proceedings of the 2006 ACEEE Summer Study on Energy Efficiency in Buildings. Asilomar, CA.

- Hanley, J.T., Ensor, D.S., Smith, D.D., Sparks, L.E., 1994. Fractional aerosol filtration efficiency of in-duct ventilation air cleaners. *Indoor Air* 4, 169-178.
- Harrison, R.M., Pio, C.A., 1983. Size-differentiated composition of inorganic atmospheric aerosols of both marine and polluted continental origin. *Atmospheric Environment* 17, 1733–1738.
- Hendrickson, C.T.; Lave, L.B.; Matthews, H.S. 2006. *Environmental Life-cycle Assessment of Goods and Services: An Input-Output Approach*, 1st ed.; RFF Press: Washington, DC.
- Hinds, W.C., 1998. *Aerosol Technology: Properties, Behavior, and Measurement of Airborne Particles*. 2nd edition, Wiley, New York.
- Hydeman, M., Tschudi, W. 2009. Energy efficiency for California states' business technology. Presented at Sacramento Municipal Utility District for the Save Energy Now program. Sacramento, CA. September 29.
<http://hightech.lbl.gov/presentations/2009-09-29-smud.pdf>.
Last accessed Dec. 10, 2009.
- IDC, 2007. IDC's Worldwide Installed Base Forecast, 2007-2010. International Data Corporation, Framingham, MA.
- IDC, 2009. IDC's Worldwide Installed Base Estimates, 1996-2008. International Data Corporation, Framingham, MA.
- IMF, 2001. World economic outlook: the information technology revolution. International Monetary Fund. www.imf.org/external/Pubs/FT/weo/2001/02/.
Last accessed Dec. 10, 2009.
- Intel, 2008. Reducing data center cost with an air economizer. Intel Information Technology, White Paper
www.intel.com/it/pdf/Reducing_Data_Center_Cost_with_an_Air_Economizer.pdf.
Last accessed Dec. 10, 2009.

- Kawamoto, K., Koomey, J.G., Nordman, B., Brown, R.E., Piette, M., Ting, M., Meier, A.K. 2001. Electricity Used by Office Equipment and Network Equipment in the U.S.: Detailed Report and Appendices. Berkeley, CA: Lawrence Berkeley National Laboratory. LBNL-45917. <http://enduse.lbl.gov/Info/LBNL-45917b.pdf>.
Last accessed Dec. 10, 2009.
- Kean, A.J., Harley, R.A., Littlejohn, D., Kendall, G.R., 2000. On-road measurement of ammonia and other motor vehicle exhaust emissions. *Environmental Science and Technology* 34, 3535–3539.
- Kirchstetter, T.W., Novakov, T., 2007. Controlled generation of black carbon particles from a diffusion flame and applications in evaluating black carbon measurement methods. *Atmospheric Environment* 41, 1874–1888.
- Kirchstetter, T.W., Corrigan, C.E., and Novakov, T., 2001. Laboratory and field investigation of the adsorption of gaseous organic compounds onto quartz filters. *Atmospheric Environment* 35, 1663–1671.
- Kitou, E., Horvath, A., 2003. Energy-related emissions from telework. *Environmental Science & Technology* 37, 3467-3475.
- Koomey, J., 2007. Estimating total power consumption by servers in the U.S. and the world. <http://enterprise.amd.com/Downloads/svrpwrusecompletefinal.pdf>.
Last accessed Dec. 10, 2009.
- Koomey, J.G., 2008. Worldwide electricity used in data centers. *Environmental Research Letters*, 3(3) 034008.
- Koomey, J., 2009. The environmental cost of cloud computing: Assessing power use and impacts. Presented at Green:Net, San Francisco, March 24.
<http://events.earth2tech.com/greennet/09/schedule/>. *Last accessed Dec. 10, 2009.*
- LaRosa, L.B., Buckley, T.J., Wallace, L.A., 2002. Real-time indoor and outdoor measurements of black carbon in an occupied house: an examination of sources. *Journal of the Air and Waste Management Association* 52, 41-49.

- Litvak, A., Gadgil, A.J., Fisk, W.J., 2000. Hygroscopic fine mode particle deposition on electronic circuits and resulting degradation of circuit performance: an experimental study. *Indoor Air* 10, 47-56.
- Long, C.M., Suh, H.H., Catalano, P.J., Koutrakis, P., 2001. Using time- and size-resolved particulate data to quantify indoor penetration and deposition behavior. *Environmental Science & Technology* 35, 2089-2099.
- Loper, J., Parr, S. 2007. Energy Efficiency in Data Centers: A New Policy Frontier. Alliance to Save Energy, Washington, DC. <http://ase.org/content/news/detail/4071>. *Last accessed Dec. 10, 2009.*
- Lunden, M.M., Revzan, K.L., Fischer, M.L., Thatcher, T.L., Littlejohn, D., Hering, S.V., Brown, N.J., 2003. The transformation of outdoor ammonium nitrate aerosols in the indoor environment. *Atmospheric Environment* 37, 5633-5644.
- Matthews, H.S., Hendrickson, C.T., Soh, D.L., 2001. Environmental and Economic Effects of E-Commerce: A Case Study of Book Publishing and Retail Logistics. *Transportation Research Record* 1763, 6-12.
- McKinsey & Company, 2007. Reducing U.S. Greenhouse Gas Emissions: How Much at What Cost? www.mckinsey.com/client/service/ccsi/pdf/US_ghg_final_report.pdf. *Last accessed Dec. 10, 2009.*
- McMurry, P.H., Shepherd, M.F., Vickery, J.S., Eds., 2004. Particulate matter science for policy makers: a NARSTO assessment, Cambridge University Press, Cambridge, England.
- Milford, J.B., Davidson, C.I., 1987. The sizes of particulate sulfate and nitrate in the atmosphere – a review. *JAPCA — The International Journal of Air Pollution Control and Hazardous Waste Management* 37, 125-134.
- Miller, R., 2008a. Intel: Servers Do Fine With Outside Air. [datacenterknowledge.com, www.datacenterknowledge.com/archives/2008/09/18/intel-servers-do-fine-with-outside-air/](http://datacenterknowledge.com/archives/2008/09/18/intel-servers-do-fine-with-outside-air/). *Last accessed Dec. 10, 2009.*

- Miller, R., 2008b. New from Microsoft: data centers in tents. Datacenterknowledge.com, www.datacenterknowledge.com/archives/2008/09/22/new-from-microsoft-data-centers-in-tents/. Last accessed Dec. 10, 2009.
- Miller, R., 2007. Data center cooling set points debated. Datacenterknowledge.com, www.datacenterknowledge.com/archives/2007/09/24/data-center-cooling-set-points-debated/. Last accessed Dec. 10, 2009.
- Mitchell-Jackson, J., Koomey, J., Blazek, M., Nordman, B., 2002. National and Regional Implications of Internet Data Center Growth. Resources, Conservation, and Recycling (also LBNL-50534) 36:3, 175-185.
- Morawska, L., Jamriska, M., Guo, H., Jayaratne, E.R., Cao, M., Summerville, S., 2009. Variation in indoor particle number and PM_{2.5} concentrations in a radio station surrounded by busy roads before and after an upgrade of the HVAC system. Building and Environment 44, 76-84.
- Morawska, L., Jayaratne, E.R., Mengersen, K., Jamriska, M., Thomas, S., 2002. Differences in airborne particle and gaseous concentrations in urban air between weekdays and weekends. Atmospheric Environment 36, 4375-4383.
- Nature, 2009. Editorial: overrated ratings. Nature 461, 146.
- Nazaroff, W.W., 2008. Climate change, building energy use, and indoor environmental quality. Indoor Air 18:4, 259-260.
- Nazaroff, W.W., Klepeis, N.E., 2004. Environmental tobacco smoke particles, in Indoor Environment: Airborne Particles and Settled Dust, Morawska, L., and Salthammer, T., Eds., Wiley-VCH, Weinheim, Germany, 245-274.
- Ott, W., Wallace, L., Mage, D., 2000. Predicting particulate (PM₁₀) personal exposure distributions using a random component superposition statistical model. Journal of the Air and Waste Management Association 50, 1390-1406.

- Pacca, S., Horvath, A., 2002. Greenhouse gas emissions from building and operating electric power plants in the upper Colorado River Basin. *Environmental Science & Technology* 36, 3194-3200.
- Patel, C. D., Shah, A. J., 2005. Cost Model for Planning, Development and Operation of a Data Center, Technical Report HPL-2005-107R1, Hewlett Packard Laboratories, PaloAlto, CA.
- Patterson, M.K., Pratt, A., Kumar., P. 2006. From UPS to silicon: an end-to-end evaluation of data center efficiency. *Enterprise Servers and Data Centers: Opportunities for Energy Savings Conference*. Santa Clara, CA. January 31. www.energystar.gov/ia/products/downloads/MPatterson_APratt_Case_Study.pdf.
- Pitz, M., Cyrus, J., Karg, E., Wiedensohler, A., Wichmann, H.E., Heinrich, J., 2003. Variability of apparent particle density of an urban aerosol. *Environmental Science & Technology* 37, 4336-4342.
- Pio, C.A., Harrison, R.M., 1987. The equilibrium of ammonium chloride aerosol with gaseous hydrochloric acid and ammonia under tropospheric conditions. *Atmospheric Environment* 21, 1243-1246.
- Riley, W.J., McKone, T.E., Lai, A.C.K., Nazaroff, W.W., 2002. Indoor particulate matter of outdoor origin: importance of size-dependent removal mechanisms. *Environmental Science & Technology* 36, 200–207.
- Roth, J., 2005. Contamination issues from belts in critical environment cooling. *Data Center Journal*, November 30, 2005. www.datacenterjournal.com.
Last accessed Dec. 10, 2009.
- Roth, K., Goldstein, F., Kleinman, J., 2002. Energy Consumption by Office and Telecommunications Equipment in Commercial Buildings--Volume I: Energy Consumption Baseline. Washington, DC: Prepared by Arthur D. Little for the U.S. Department of Energy. A.D. Little Reference no. 72895-00.
<http://www.eren.doe.gov/buildings/documents>. *Last accessed Dec. 10, 2009.*

- Rumsey Engineers, 2008. Network Appliance: Building 02 Datacenter. Pacific Gas and Electric Non-Residential New Construction Incentive Program. January.
- Rumsey Engineers, 2005. Network Appliance Building 11 Datacenter. Pacific Gas and Electric Industrial Savings by Design Program. June.
- Sarnat, J.A., Long, C.M., Koutrakis, P., Coull, B.A, Schwartz, J., Suh, H.H., 2002. Using sulfur as a tracer of outdoor fine particulate matter. *Environmental Science & Technology* 36, 5305-5314.
- SCAQMD, 1996. 1997 Air Quality Management Plan. South Coast Air Quality Management District. <http://www.aqmd.gov/aqmp/97aqmp/index.html>.
Last accessed Dec. 10, 2009.
- Seinfeld, J.H., Pandis, S.N., 2006. *Atmospheric Chemistry and Physics: From Air Pollution to Climate Change*. 2nd edition. Wiley, New York.
- Shehabi, A., Horvath, A., Tschudi, W., Gadgil, A.J., Nazaroff, W.W, 2008. Particle concentrations in data centers. *Atmospheric Environment* 42, 5978-5990.
- Shields H.C., Fleischer D.M., Weschler C.J., 1996. Comparisons among VOCs measured in three types of US commercial buildings with different occupant densities. *Indoor Air* 6, 2-17.
- Shields, H.C., Weschler, C.J., 1998. Are indoor air pollutants threatening the reliability of your electronic equipment? *Heating/Piping/Air Conditioning Engineering* 70(5), 46-54.
- Sloan, J., 2008. Recovering from rejection. Proceedings from the ASHRAE 2008 Winter Conference , New York, NY.

- Stanley, J. R., Brill, K. G., Koomey, J. 2007. Four Metrics Define Data Center Greenness Enabling Users to Quantify Energy Efficiency for Profit' Initiatives. The Uptime Institute, Santa Fe, NM. <http://www.uptimeinstitute.org>. *Last accessed Dec. 10, 2009.*
- Stern, N., Peters, S., Bakhshi, V., Bowen, A., Cameron, C., Catovsky, S., Crane, D., Cruickshank, S., Dietz, S., Edmonson, N., Garbett, S.L., Hamid, L., Hoffman, G., Ingram, D., Jones, B., Patmore, N., Radcliffe, H., Sathiyarajah, R., Stock, M., Taylor, C., Vernon, T., Wanjie, H., Zenghelis, D., 2006. Stern Review: The Economics of Climate Change, HM Treasury, London.
- Sullivan, A., 2009. Data center efficiency with Energy Star. United States Environmental Protection Agency, Presentation at the 7x24 Exchange 2009 Spring Conference, Boca Raton, FL.
- Swenson, D., Kinnear J., 2009. The role of relative humidity and dew point on electrostatic charge generation and electrostatic discharge (ESD), American Society of Heating, Refrigerating, and Air-Conditioning Engineers, Inc., Conference Presentation, Chicago, IL.
- Syska Hennessy Group, 2007. The Use of Outside Air Economizers In Data Center Environments. White paper 7. Charlotte, NC.
- Toffel, M.W., Horvath, A., 2004. Environmental Implications of Wireless Technologies: News Delivery and Business Meetings. *Environmental Science & Technology* 38, 2961- 2970.
- Tschudi, W., Xu, T., Sartor, D., Nordman, B., Koomey, J., Sezgen, O., 2004. Energy efficient data centers. Report LBNL-54163, Lawrence Berkeley National Laboratory, Berkeley, CA.
- Turpin, B.J., Huntzicker, J.J., Hering, S.V., 1994. Investigation of organic aerosol sampling artifacts in the Los Angeles basin. *Atmospheric Environment* 28, 3061-3071.

- Üerge-Vorsatz, D., Harvey, L.D.D., Mirasgedis, S., Levine, M.D., 2007. Mitigating CO₂ emissions from energy use in the world's buildings. *Building Research & Information* 35, 379-398.
- Uptime Institute, 2000. Heat Density Trends in Data Processing, Computer Systems, and Telecommunications Equipment. Santa Fe, NM.
- U.S. Congress, 2006. An Act to Study and Promote the Use of Energy Efficient Computer Servers in the United States. Federal Register, National Archives and Records Administration. Public Law 109-431, 120 Stat 2920.
- USGBC, 2009. LEED 2009 for New Construction and Major Renovations Rating System. United States Green Building Council. www.usgbc.org.
Last accessed Dec. 10, 2009.
- Walker, I.S., Wray, C.P., Dickerhoff, D.J., Sherman, M.H., 2001. Evaluation of flow hood measurements for residential register flows. Report LBNL-47382. Lawrence Berkeley National Laboratory, Berkeley, CA.
- Wallace, L., 2005. Real-time measurements of black carbon indoors and outdoors: a comparison of the photoelectric aerosol sensor and the aethalometer. *Aerosol Science and Technology* 39, 1015–1025.
- Waring, M.S., Siegel, J.A., 2008. Particle loading rates for HVAC filters, heat exchangers, and ducts. *Indoor Air*, doi:10.1111/j.1600-0668.2008.00518.x.
- Weber, C.L., Koomey, J.G., Matthews, S.H., 2009. The Energy and Climate Change Impacts of Different Music Delivery Methods Final report to Microsoft Corporation and Intel Corporation, August 17,
<http://download.intel.com/pressroom/pdf/CDsvsdownloadsrelease.pdf>.
Last accessed Dec. 10, 2009.

- Weber, C., Hendrickson, C., Jaramillo, P., Matthews, S., Nagengast, A., Nealer, R. 2008. Life cycle comparison of traditional retail and E-commerce logistics for electronic products: a case study of buy.com. Green Design Institute, Carnegie Mellon University, Pittsburgh. http://valcoprams.com/images/static_images/green_study.pdf.
Last accessed Dec. 10, 2009.
- Weschler, C.J., 1991. Predictions of benefits and costs derived from improving indoor air quality in telephone switching offices. *Indoor Air* 1, 65-78.
- Whitby, K.T., 1978. The physical characteristics of sulfur aerosols. *Atmospheric Environment* 12, 135-159.