Communicating complexity and informing decision-makers: challenges in the data and computation of environmental benefits of renewable energy

by

Tarek Rached

B.S., Mechanical Engineering, Carnegie Mellon University
M.S., Systems and Information Engineering, University of Virginia

Submitted to the Engineering Systems Division
in partial fulfillment of the requirements for the degree of

Master of Science in Technology and Policy

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2008

Author . . . .
Engineering Systems Division
May 18, 2008

Certified by . . . .
Stephen R. Connors
Director, Analysis Group for Regional Energy Alternatives
Thesis Supervisor

Accepted by . . . .
Dava Newman
Professor of Aeronautics and Astronomics and Engineering Systems
Director of the Technology and Policy Program

Communicating complexity and informing decision-makers: challenges in the data and computation of environmental benefits of renewable energy

by

Tarek Rached

Submitted to the Engineering Systems Division
on May 18, 2008, in partial fulfillment of the
requirements for the degree of
Master of Science in Technology and Policy

Abstract

This thesis contrasts the quantification of avoided emission benefits of renewable generation as determined by a marginal emissions analysis and the methodology specified by the Massachusetts Greenhouse Gas Policy and Protocol. Both methodologies are applied to an offshore wind installation that is currently being proposed by the Town of Hull, Massachusetts. The key finding is that the Massachusetts Greenhouse Gas Policy undercounts the avoided emissions benefits of the proposed installation by a range of 30%-50%, depending on the emission. Finally, the policy implications of this finding is explored and expounded upon.

Thesis Supervisor: Stephen R. Connors
Title: Director, Analysis Group for Regional Energy Alternatives

Acknowledgments

Acknowledgements

I have many people to thank for their help on the research that lead to this thesis. Stephen Connors let me loose on the AGREA database and then gently guided me back on course whenever I found myself lost in a morass. Thank you, JP, for being a great sounding board/sanity checker, and for reassuring me after my periodic announcements that I'd found a grave flaw in one of our foundational assumptions. The whole team at the Renewable Energy Research Laboratory at the University of Massachusetts at Amherst were tremendously helpful, both in choosing AGREA to carry out the emissions analysis, and for keeping me well-fed with wind resource data. The folks at the Hull Municipal Light Plant were likewise patient and helpful in helping me track down their data. I'd like to give special thanks to the Town of Hull, for their vision and commitment to renewable energy. Thank you for giving me the opportunity to work on this worthwhile project.

And finally, thank you Liz, for all your support during many days and nights lost in code.

THIS PAGE INTENTIONALLY LEFT BLANK

Executive Summary

Renewable energy is coming online at a tremendous rate around the world. It is not, however, being born into a vacuum—there is a huge existing infrastructure of fossil generation that renewable generation exists alongside. In the absence of large-scale electricity storage, the actual emission reduction benefits of renewable generation thus depends on the *interaction* of the power produced by renewable generation and the power produced by fossil generators.

The Analysis Group for Regional Energy Alternatives (AGREA) has developed an analytical methodology to determine the exact avoided emission benefits of renewable generation that identifies the fossil generators responding to load in every hour, and then matches that up with the amount of renewable generation available in each hour. My research extends this historical analysis into a prospective tool to evaluate the potential avoided emissions benefits for an offshore wind turbine installation of 15 megawatts that the Town of Hull, Massachusetts has proposed. In addition to computing the anticipated avoided emissions, I also explore several graphical methods of communicating both my analysis and my results to the policymakers and residents of Hull.

In the autumn of 2007, the Commonwealth of Massachusetts released the Massachusetts Greenhouse Gas Policy and Protocol that mandates that qualifying projects institute measures to mitigate the emissions of greenhouse gases caused by the development of those projects. The Protocol also specifies a methodology for quantifying the reduction in emissions due to the mitigation measures, and an application of this methodology has been requested by the Secretary for Energy and Environmental Affairs for the Hull Offshore project. Over the course of this thesis, I show that the MA GHG Protocol is inappropriate for the quantification of avoided emission benefits of renewable generation in general, and that in the case of the Hull Offshore Project, in fact under-counts avoided emissions by 33% to 50%. I examine the reasons for this discrepancy, and recommend that a more appropriate policy be developed to accurately quantify the environmental benefits of renewable generation.

My key findings are that if the Town of Hull builds the project out to the full 15 MW of capacity, the project will avoid 29,550 tonnes of CO_2, 52,200 kg of SO_2, and 22,350 kg of NO_x annually in a medium wind year. These numbers increase by 16% for high wind years, and decrease by around 10% for low wind years. These figures are shown in the context of historical data in Figure 3-22. In terms of the Town of Hull's existing footprint, these avoided emissions figures mean that they will avoid 153% of their CO_2, 104% of their SO_2, and 145% of their NO_x. These numbers are shown in Figure 0-1.

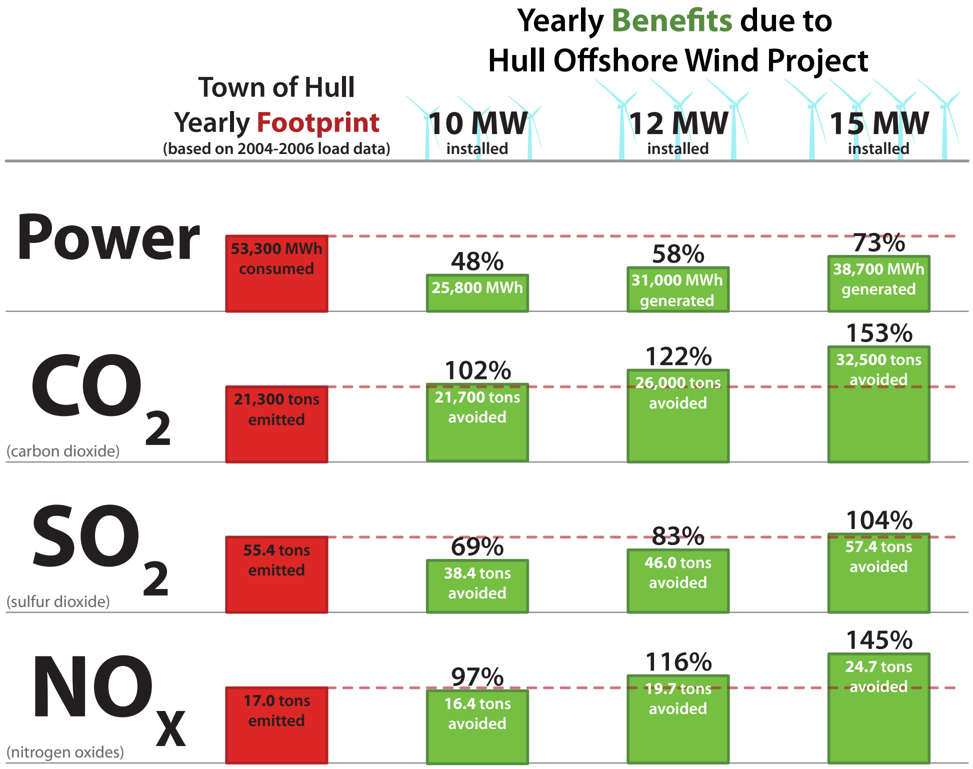

Figure 0-1: Hull Offshore benefits

Contents

THIS PAGE INTENTIONALLY LEFT BLANK

List of Figures

List of Tables

THIS PAGE INTENTIONALLY LEFT BLANK

Chapter 1

Introduction

1.1 Goal

The goals of this thesis are two-fold: first, I will show how to better measure the quantity of air emissions that will be avoided by building episodic, renewable generation or investing in energy efficiency. Then, I will show how to best synthesize and present those results in terms that are relevant to decision-makers and the public. Finally, I aim to demonstrate the superiority of AGREA's situational, marginal emissions methodology to the systemwide average methodology currently specified by the Commonwealth of Massachusetts. These three goals support the larger, overarching goal of accurately assessing the environmental benefits of renewable generation, a goal which takes on increasing importance as society transforms its energy infrastructure to address the challenge of global climate change.

1.2 Context

1.2.1 Quantifying avoided emissions from renewable generation is important

Accurately quantifying avoided emissions from renewable generation is important. Global anthropogenic climate change is *the* issue of our age, and the electricity generation sector accounts for 33% of the United States' greenhouse gas emissions.[1] There are a number of policies underway or under consideration that aim to eliminate the current negative externalities associated with the emission of CO_2 by pricing that emission in some form or another. Amongst these schemes here in New England is the Regional Greenhouse Gas Initiative, which will put a cap-and-trade system on utility generators in the participating states. While the avoided emissions resulting from the operation of renewable generation will not be salable under RGGI, they will reduce the number of permits that a "load serving entity" will need, and the renewable electrons themselves will be comparatively less expensive due to fossil generators' having to purchase carbon permits. It is crucial that the barriers to the construction of renewable generation be lowered so that when RGGI goes into effect, consumers will have renewable capacity available to them.

In addition to greenhouse gas emissions, the other, "traditional" criteria pollutants, SO_2, NO_x, and particulates are still being emitted, and although the various implementations of the Clean Air Act[1] have done much to reduce the rates at which these compounds are emitted, there is still a long way to go, as these pollutants continue to cause adverse environmental and human health effects.

As the United States, and New England in particular, continues to build more electricity generation capacity, it is imperative that we be able to factor in the relative avoided emissions of potential renewable generation projects of different types, sites, and scales in order to accurately prioritize them by their environmental impacts.

1.2.2 Calculating Avoided Emissions is Hard

In order to understand how to accurately measure how much emissions are avoided by adding renewables to the generation mix on the grid, we first need to understand how the grid actually operates.

---

[1] Our marginal emissions analysis is, in fact, based on hourly emissions data collected under the auspices of the Act.

Grid Behavior

The electricity industry in New England was restructured in 1997. Currently, the Independent System Operator (ISO) for New England is responsible for the operations of the grid, and runs the markets that govern the supply and demand of electricity across New England. At a gross level of simplification, here is how those markets operate: Every day, the ISO releases a forecast of the demand for electricity for every hour for the following day. Electricity generators submit (complex, 30-part) bids to supply certain amounts of power under certain constraints. All the bids are stacked in order of ascending price, and the market clears at the marginal bid. This is the **day-ahead market**. The next day, everyone whose bid was accepted produces as much electricity as they promised, but since the actual demand always varies from the forecast demand, there is a secondary, **real-time market** to make up the difference. In the real-time market, generators have standing bids to produce power, and this market is continuously clearing to match supply to demand. Finally, there is an **ancillary services market**, which provides services such as voltage regulation and operating reserves.

Wind power—along with solar and run-of-the-river hydro—usually bids zero dollars into the realtime market, as the nature of the resource is use it or lose it. This means that from a market perspective, wind will theoretically displace whatever the marginal unit bid is in the real-time market. For a low capacity project such as the Hull Offshore project, the amount of generation will most likely not actually affect the dispatch order (e.g. whether fossil units get turned off or on). Furthermore, although the preceding explanation of the electricity markets makes it sound like units' behavior is to either turn on or off when they are dispatched, that complicated 30-part bid actually leads to more complex behavior. One typical behavior is for a generating unit to bid 95% of its capacity into the day-ahead market, and then bid the last 5% of its capacity into either the real-time or ancillary market, where that power can fetch a premium price. The actual response of fossil generation to small amounts of renewable generation is thus not shutting down one or two units; rather many units will throttle down their generation by some marginal fraction. Thus, it is inaccurate to categorize a unit as a "peaking" or a “baseload” unit, when in fact many units share the characteristics of both.

New England Fuel Mix

The electric generation capacity in New England is a mix of natural gas, with significant oil, coal, and nuclear generation as well. As can be seen in the weeklong hourly snapshot of fossil generation in Figure 1-1, coal units operate at high capacities all hours of the day, while the gas and oil units respond to the load at a higher rate. However, even the coal does dips overnight and ramps up in times of high demand. Since the carbon dioxide intensity of coal is higher than that of oil and gas, the average carbon dioxide emission rate varies inversely with the amount of gas and oil generation online. This effect can be seen when the line representing carbon dioxide emissions slopes up as the gas generation goes down in Figure 1-1.

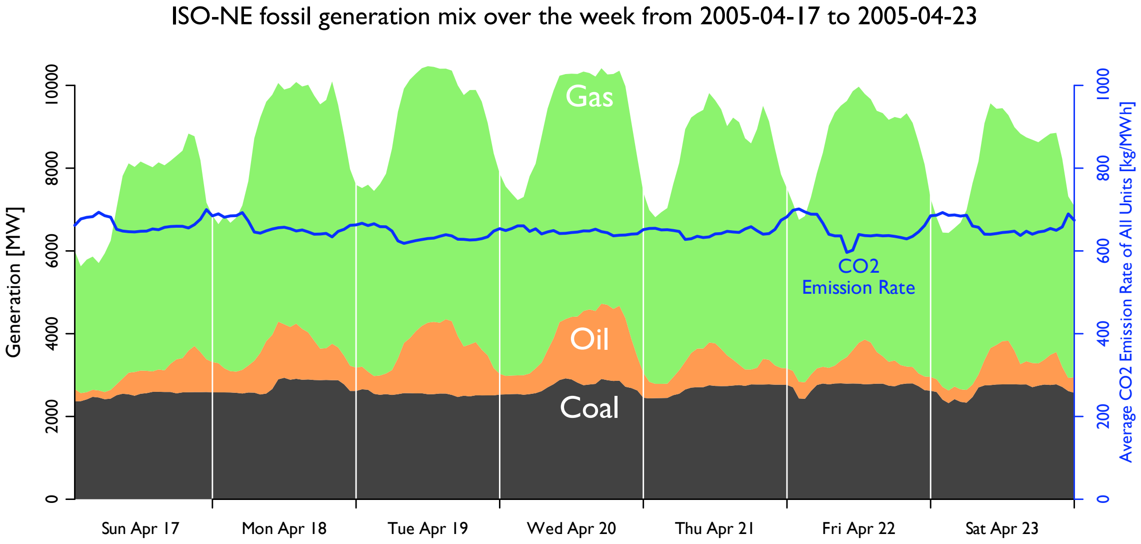

Figure 1-1: Fossil Generation by fuel type over one week, with average carbon dioxide emission rate. (Not shown: nuclear, hydro, and electricity imports & exports)

1.2.3 Proposed Hull Offshore Wind Project

The Town of Hull, Massachusetts, has a progressive municipal electrical utility, the Hull Municipal Light Plant (HMLP), that has already installed two wind turbines onshore in town. The first large turbine on Hull, a 660 kW turbine named Hull Wind I, was installed in 2001 near the Hull High School on Windmill Point. The second turbine, a 1.8 MW model, was erected on the town landfill in 2006. HMLP installed and operates both these units, and is also responsible for purchasing wholesale electricity and selling it to the residents of Hull. Together, these turbines supply an average of 12% of the town's electricity. The town and its utility are now moving forward with the Hull Offshore project, an installation of up to four turbines for a total additional wind capacity of 15 MW. The turbines are to be sited on Harding's Ledge, pictured in Figure 2-1, a little over a mile off of Hull's Natasket Beach.[13]

The Hull Offshore project has been under active development since 2003, when HMLP, with technical support from the Renewable Energy Research Laboratory (RERL) at the University of Massachusetts at Amherst, initiated discussions with the U.S. Army Corps of Engineers and state officials about the permitting process for a single offshore turbine. By 2005, HMLP had expanded their proposal to four turbines in the 3.6 MW range, and in 2006, the Massachusetts Technology Collaborative offered the town of Hull a loan (forgivable in the case that the project is not built) to support detailed technical analysis of the project. The RERL contracted with my research group, the Analysis Group for Regional Energy Alternatives (AGREA) in the Laboratory for Energy and the Environment (LFEE) at the Massachusetts Institute of Technology (MIT) to provide a marginal emissions analysis of the proposed wind installation.

The current status of the project, with regards to state-level environmental permitting is as follows: In December of the past year, 2007, a 16-page Environmental Notification Form (ENF) was submitted to the MEPA Office outlining the scope of the project and its anticipated environmental impacts. This form was accompanied by a more detailed 49-page Narrative that examined each potential impact in depth. The ENF was reviewed by the MEPA Office, and on February 8th, they released their certificate on the ENF, which stipulated that the project "requires the preparation of an Environmental Impact Report (EIR)." The EIR is currently being prepared, and will be submitted in the upcoming months. The ENF was prepared, and the EIR is being prepared, by ESS Group, Inc., of Wellesley, MA.

My analysis is relevant to both the ENF and the EIR, as the avoided air emissions constitute the primary environmental benefits of the project. Preliminary results of my analysis were provided to ESS group for the preparation of the ENF. I will be providing ESS with the final results for the EIR, along with a detailed description of how I arrived at those results. In the certificate of the ENF, Ian A. Bowles, the Massachusetts Secretary of Environmental Affairs, noted that the "displaced emissions" quantified in the ENF lacked an explanation of the methodology used, and directed ESS to use the EEA Greenhouse Gas Emissions Policy and Protocol. In the course of this thesis, I use this Protocol and explore the differences between it and our marginal analysis.

1.3 Other Approaches

The quantification of the avoided emissions that result from renewable generation has been attempted in a number of ways. The myriad approaches can basically be broken down into two categories: modeling the system dispatch, or examining actual emissions data. I take the second approach in this thesis.

1.3.1 Examining Historic Emissions Data

In order to determine which units are actually ramping down their output when renewable generation is fed into the grid—and thus which emissions are being avoided—a number of approaches can be taken. The first is to just take the average emissions of all units supplying load at any point in time, e.g. the average hourly system emissions. The next level of detail is to identify the units that are actually responding to changes in the system load, and then calculate *their* average emissions. There are a number of methods to identify units that respond to load, from simply categorizing units according to their fuel (the approach taken by the Massachusetts Greenhouse Gas Policy and Protocol, which I will explore in greater detail in the following section), to examining their actual behavior on an hourly basis. My research team, the Analysis Group for Regional Energy Alternatives, takes this latter approach, and we go one step further by using the marginal emissions that are associated with the response to system load as the basis for our emissions rate. This methodology was developed for the US EPA by Stephen Connors, Mike Adams, Kate Martin, Ed Kern, and Baafour Asiamah-Adjei[9] and was developed further by Mike Berlinski[7]. This approach has been used to asses the avoided emissions potential of solar and wind energy in the continental USA. My contribution will be to extend this historical analysis into an anticipatory tool that can be used to assess potential renewable generation projects.

1.3.2 Massachusetts GHG Policy

In the fall of 2007, the Massachusetts Environmental Policy Act (MEPA) Office promulgated the Massachusetts Greenhouse Gas (MA GHG) Policy and Protocol.[2] Under this policy, certain projects requiring an Environmental Impact Report must "identify measures to avoid, minimize, or mitigate [GHG] emissions." To do so, the project must quantify both the anticipated GHG emissions for a baseline project, and the anticipated GHG savings with mitigation. The project is covered by the policy if it: i. is being undertaken by the Commonwealth; ii. is funded in part or in whole by the Commonwealth; iii. requires an Air Quality Permit; or iv. requires a Vehicular Access Permit. Although the policy is ostensibly concerned with all greenhouse gases, in practice it is currently focused only on CO_2 generated during the Use Phase of the project (the Construction Phase is not considered). Three primary methodologies are prescribed for calculating the CO_2 emissions of a proposed project: for direct emissions from on-site equipment, indirect emissions from the generation of the energy that the project is anticipated to consume, and finally the indirect emissions from transportation caused by the project. It is the second methodology that is of interest to us, as this is the approach that is used to also calculate the anticipated avoided emissions for renewable generation[2] The methodology is comprised of two steps: first, estimate the total electricity that will be consumed by the project over its lifetime; second, multiply this quantity of electricity by

---

[2] It should be noted that this application of the methodology was not considered by any agency—it was requested by Ian Bowles, the Secretary of Energy Environmental Affairs in his response to the Environmental Notification Form submitted by the Hull Municipal Lighting Plant: "Estimates of air quality emissions associated with traditionally produced power should be based on the ISO-New England Marginal Emissions Report which provides emissions factors for a variety of stationary combustion sources."

an “emissions factor” expressed in pounds of CO_2 emitted per megawatt hour generated. The Policy specifies that the “emissions factor” should come from the 2005 ISO-New England Marginal Emissions Report.

1.3.3 ISO-New England Marginal Emissions Report

Every year, the system operator for New England, ISO-NE, releases an analysis of the marginal emissions released by fossil generating units on the grid in the previous year. This analysis is performed in order to quantify the effect that demand side management programs have had on emissions, and to assess the avoided emissions of renewable energy projects. The basic output of this report is a table of “marginal emission rates” for CO_2, SO_2, and NO_x, expressed as the mass of the compound emitted per unit of electricity generated. For each compound, a couple of emission rates are provided: separate on-peak and off-peak rates for CO_2 and SO_2, and every permutation of on- and off- peak and ozone and non-ozone season for NO_x. On-peak is defined as all hours between 8 a.m. and 10 p.m., and ozone season is May through September. Annual averages for all hours are also calculated and presented for each compound.

The algorithm used by ISO-NE to calculate the avoided emission rates is fairly straightforward: take the average emission rates of all marginal units within the time cohort in question. Marginal units are defined as “intermediate fossil units”, which is subsequently defined as all units burning oil and/or gas. The stated reason for excluding coal units from this algorithm is that coal units “typically operate as baseload units and would not be dispatched to higher levels in the event that more load was on the system.”[10] This algorithm is applied to the EPA EGRID data set, and the emission rates are produced. As I will show, the assumption that coal does not respond to load is not born out by a close examination of the EGRID data, and results in an underestimation of the marginal emission rate.

Since the MA GHG Policy does not specify which marginal emission rate (“emission factor,” in the language of the policy) to use, the annual average marginal CO_2 emission rate appears to be the correct number to use, according to the policy. For most projects, this is an entirely appropriate rate, as a representative temporal profile of the electricity that will be used over the use phase of a project is difficult to calculate. However, the use of one emissions rate for all hours is not appropriate for renewable generation projects for which widely accepted and highly accurate estimations of seasonal and diurnal generation patterns exist. The time dynamics, and correlation with the temporal patterns in the emissions rate are lost in such an analysis, and thus the resulting estimates of potential environmental benefits will not be as accurate as they could be. Although the MA GHG Policy results in a coarse analysis, it does so with some degree of self-awareness; in its introduction, the Policy states:

> EEA also recognizes that the GHG quantification required by this Policy will not result in absolutely accurate projections. The intent is not one hundred percent certainty as to the amount of GHG emissions; rather, it is a reasonably accurate quantitative analysis of emissions and potential mitigation that will allow the Project proponent and reviewers to assess the overall impact of the Project as proposed and the reduction in emissions if various techniques are used.[2]

1.4 Summary of Results and Conclusions

In the course of my analysis, I find that the Town of Hull can, if it builds the offshore wind project out to the full 15 megawatts of capacity, avoid all of their emissions, and in fact offset an additional 50% of their CO_2 and NO_x on top of that. I also validate the above hypothesis that the Massachusetts Greenhouse Gas Policy and Protocol is too coarse to accurately capture the dynamics of the avoided emissions characteristics of renewable generation, and my primary recommendation is that the Protocol not be used for such purposes.

1.5 Roadmap

Up next is an exploration of the data and methodology used in the AGREA Load Shape Following (LSF) analysis, along with an identification of assumptions and limitations.

I will share the results of the analysis, and guide you through the salient dynamics of each of the resulting time series. I will subsequently show several graphic explanations of the results for the specific audiences this research is intended for (policymakers and the public), and conclude with a roundup of the results, lessons learned and an identification of future avenues of research.

Chapter 2

Data and Methodology

2.1 Key Methodological Assumptions

This analysis depends on the primary assumption that the amount of electricity produced by the proposed wind installation is so small as to not affect the dispatch order of fossil units. My analysis is targeted at the marginal emissions—those emissions that are produced as a result of generating units ramping up or down their output to match net load, not the emissions resulting from units coming online or being taken offline. Once a proposed wind installation—or group of installations—is large enough that it does affect the dispatch order, and hence violates this assumption, the relevant emissions factor is the average—not marginal—emissions, as this number takes into account the total emissions in any hour. For the proposed 15 MW Hull offshore project, this assumption is a safe one, as the nameplate capacity of 15 MW is smaller than the vast majority of generating units, and actual output of the wind turbines in any given hour will often be lower than that.

A related assumption is that the proposed project requires no additional fast-acting reserve generation capacity to mitigate the intermittency of the renewable resource.[16] This assumption is again based in the size of the project, and the Hull project's 15 MW is sufficiently small for this assumption to hold. The 10 minute hourly reserve requirements in ISO New England are generally in the hundreds of megawatts, so the 15 MW Hull project is not going to have much of an effect. Were the project large enough to start requiring additional reserve generation, the emissions associated with having that additional capacity operating in spinning reserve would have to be factored into the emissions impact of the project, in addition to the actual emissions generated when the reserve would be called up.

Finally, we're assuming that the proposed renewable generation is bidding at zero into the realtime market, acting as a price-taker. This is typical behavior for any type of generation with high capital costs and low to no marginal costs, such as wind, solar, run-of-river hydropower, and nuclear power.

2.2 Data

In my analysis, I draw on four hourly time series of data:

- Emissions measurements from each generating unit in New England (US EPA CEM, EGRID)
- The total electricity consumption in New England (ISO-NE)
- The wind resource at the Hull offshore site (UMASS RERL)
- Electricity sales in Hull and wind generation from Hull 1 & 2 (Hull Municipal Light)

By combining these disparate data, I calculate the emissions that would be avoided by the proposed offshore wind turbine project in Hull, and put those results in the context of Hull's current electricity situation. Before I explain the methodology, it behooves me to introduce you, dear reader, to the data and its structure in some detail, in order to clarify the subsequent explanation of methodology.

2.2.1 Database

The data is all is stored in a single relational database, in order to facilitate its exploration and analysis. Each set of data from the sources listed above takes its own meandering path into the database, but once it is there and normalized, it can all be referenced and utilized in a common format. Since we are dealing with hourly time series, the timestamp of each record is, as I will show, the most crucial piece of connective data. The actual database used is MySQL version 5.0.41,[1] hosted on an Athena Redhat server that enables multiple users to access the database simultaneously. MySQL is a mostly ANSI SQL compliant database, with widespread support from many software vendors, which allows each user to use whichever software package best suits to their analytical needs. Most of my analysis was conducted in the R statistical software environment,[2] and a combination of PERL[3] and PHP[4] scripts were used to get the data into the database.

2.2.2 EPA emissions data

The two datasets that forms the foundation for the entire analysis are the Continuous Emissions Monitoring (CEM) dataset and the Emissions & Generation Resource Integrated Database (eGRID). Both are products of the Environmental Protection Agency. The first dataset, CEM, provides hourly measurements of emissions data and operating parameters for each fossil generating unit in the United States, and the second, eGRID provides metadata about the those units, the generators they power, and the power plants in which they are located.

Continuous Emissions Monitoring

The Continuous Emissions Monitoring is a by-product of the Acid Rain Program, which was set up to achieve the goals of Title IV of the 1990 Clean Air Act. In this program, every electricity generating unit over 25 megawatts capacity, and every unit using high-sulfur fuel, regardless of capacity, must measure and report hourly emissions of CO_2, SO_2, and NO_x. Alongside the emissions data, the units report their operating parameters, such as heat rate and unit output in each hour. This data is collected by the Clean Air Markets Division of the EPA, and made available on its website every year in a compressed EDR format, with one file per plant per quarter.[5][3, 5] To bring this data into the database,

---

[1] http://www.mysql.com/

[2] http://www.r-project.org/

[3] http://www.perl.org/

[4] http://www.php.net/

[5] As the EDR format is currently being phased out, and all future CEM data will be submitted and accessible online in an XML syntax.

I downloaded it and ran a series of scripts developed by my predecessor, Mike Adams, that unpack the raw data and create data files that can be read into the database. Each generating unit is uniquely identified by its Office of the Regulatory Information System PLant number (ORISPL), and its unit identifier. We store one record for each unit for each hour by using the ORISPL, the unit identifier and the timestamp as a unique key. Since there are several million records in each subregion (e.g. power pools like ISO-NE), and our algorithms only utilize one subregion at a time, we create separate tables for each subregion.

The CEM data for the period 1998 through 2002 had already been loaded in the AGREA database to enable the prior LSF analyses.[9, 7] I downloaded and imported the data from 2003 through 2006, although at a late stage in the analysis, I discovered that half of the plants in the New England subregion (NEWE) did not get imported for just the year 2006. Thus, the numbers for this last year should be taken with a grain of salt. I will be remedying this error, as well as adding the emissions for 2007 in an upcoming report that will be released to the public.

eGRID

eGrid is a high-level inventory of all available operational performance and emissions data of electric power systems in the US.[6] The data in eGRID is provided at several levels of aggregation, from individual boiler data all the way on up to power control area, and all data is aggregated by year. This data is provided by the EPA in the form of Excel spreadsheets that are released periodically. We use this dataset to build several tables with information about each fossil boiler, each generator, and each power plant in our dataset. Since eGRID also uses the same ORISPL as the CEM to identify each plant, we can join these two tables together with the appropriate SQL queries.

2.2.3 ISO-NE System Load

The New England Independent System Operator (ISO-NE) publishes a great deal of information about the performance of the New England grid, from market overviews to detailed hourly load and anonymized bid data. We used this latter dataset to determine the total system load within each hour, although the quantity of imports and exports within each hour are not considered.[12]

2.2.4 The Hull Offshore Wind Resource

The wind resource data was obtained from the Renewable Energy Research Laboratory (RERL) at the University of Massachusetts at Amherst. RERL had installed Light Detection and Ranging (LIDAR) equipment capable of measuring the wind speed and direction at multiple heights on Little Brewster Island in 2006. They used the data gathered from this instrument to compute a 10-minute interval time series of wind speeds at 80 meters altitude (the hub height of the proposed turbines) at Harding Ledge, the proposed project location. In addition to this data set, I obtained an hourly time series of wind speeds measured at Logan International Airport from January 1996 through July of 2007. Using an algorithm known as Measure, Correlate, Predict (MCP),[15] I was able to generate an hourly time series of wind speeds for the same time period as the Logan data, based on the correlation between it and the LIDAR data.[6]

I was then able to take this wind speed data and feed it through a power curve for one of the candidate turbines. The power curve basically tells you how much power the turbine will generate at any given wind speed. The turbines under consideration were the GE 3.6s MW, which starts generating power in winds of 4 m/s, reached maximum output at 15 m/s, and a cuts out to prevent damage to the turbine and blades at 25 m/s.

To facilitate a more flexible analysis that wasn't constrained by the total installed capacity, I divided the power data by the capacity of the proposed turbine to calculate the MWh generated per MW of installed capacity, which could then be multiplied by the size of the proposed project to get hourly generation.

---

[6] There are a number of differences between Logan Airport and the Hull Site that could create artifacts in the MCP'ed wind data for Hull. The most obvious of these is the large surface area at Logan that is covered by concrete, which could heat up in the summer, creating a sea-breeze due to the differential between the heating of the runways and the ocean.

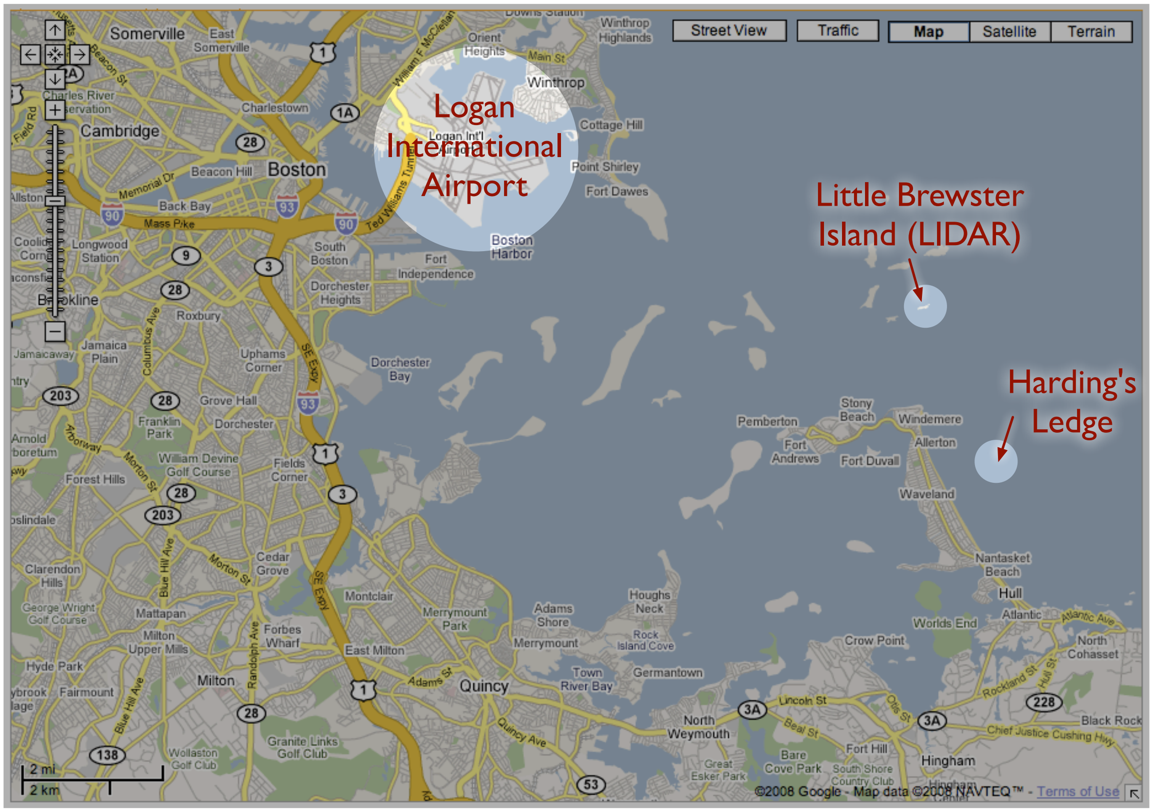

Figure 2-1: Map of sources of wind resource data and Hull Offshore site – Logan International Airport, LIDAR on Little Brewster Island, and the proposed location of the Hull Offshore Turbines on Harding's Ledge (source: Google Maps)

2.2.5 Historical Load and Wind Generation

The final data set I used was the hourly electricity sales and generation from Hull Wind I and Hull Wind II, provided by Mike Lynch at the Massachusetts Municipal Wholesale Electric Company (MMWEC) with the permission of the Hull Municipal Light Plant (HMLP). HMLP reports their hourly sales, along with the output from each of the wind turbines, to MMWEC, from whom I acquired the data. Once I had obtained the hourly data, I stored it alongside the hourly wind speed, power output, and capacity factor data from the previous section in a single table.

2.3 Methodology

2.3.1 Historical Avoided Emissions Analysis

Overview

The AGREA avoided emissions methodology combines a number of disparate data sources—for clarity's sake, here is an overview of how they all fit together:

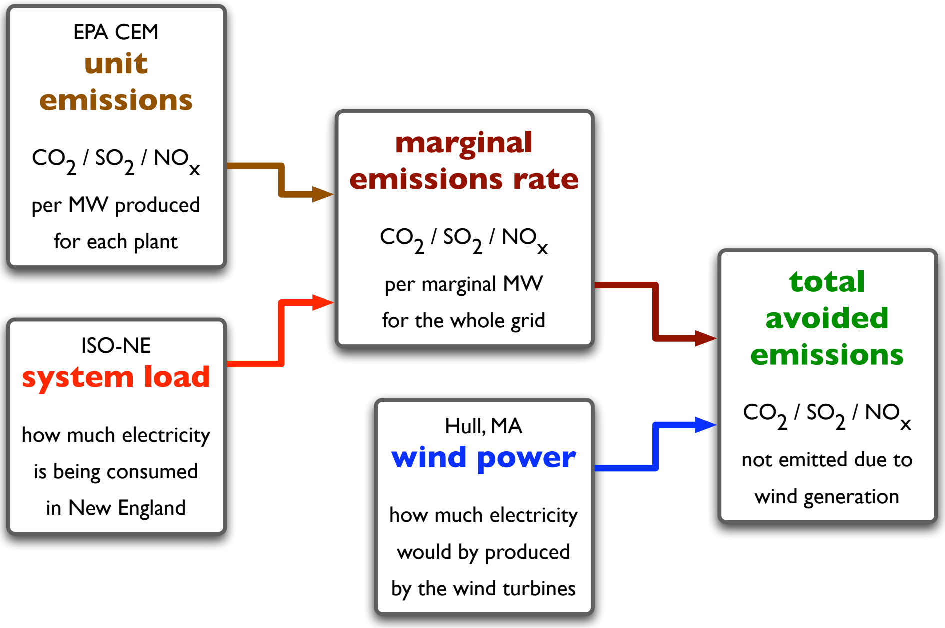

Figure 2-2: Overview of the data flows in the AGREA avoided emissions methodology

First of all, we load the data from the sources described in the data section into the database. As we load the EGRID fossil generating unit data, we also compute the following for each unit in each hour: the change in unit load, the unit's load state, and whether or not the unit's is flagged as "Load Shape Following" (LSF). The basic algorithm is as follows: the unit's change in load is simply the difference between the output in each hour and the output in the hour preceding it.

To identify the load state, we first determine the unit's maximum output in the current year and the two years preceding it. If the current year is the first year of the dataset, we look at the current year and one year ahead of it. We take this maximum output to be the unit's operational capacity—the nameplate capacity (analogous to the capacity of a computer hard drive), is not an accurate enough indicator of the actual operational capacity of the unit. In addition to this capacity, we also calculate yearly summer capacities, based only on the data from May through September, as the operational capacity of thermal power plants is reduced when air temperatures are higher. So now that we have a maximum capacity, we calculate the unit's output during each hour as a percentage of that maximum capacity, and use that percentage to bin the unit in that hour into one of four "load states", as shown in Table 2.1.

Load Condition in any hour	Load State in that hour
output $<$ 5% capacity	Turning on or off
5% $\leq$ output $<$ load 55%	Standby
55% $\leq$ output $<$ 90%	Spinning Reserve
90% $\leq$ output	Full Load

Table 2.1: Unit Load States and Conditions

Finally, we determine if the unit is following the load. First of all, any units in spinning reserve are considered to follow load, and so they are automatically flagged as such. Units in all other load states are also flagged as load shape following in each hour if one of the following is true: either the unit's load is changing in the same direction as the system's load, or the unit was flagged as load shape following in the previous hour and both the total system load and the unit's load have changed less than 2.5% over the past hour.

Thus, for each unit, for each hour, we have three new pieces of data: the change in unit load, the unit's load state, and whether or not it is following the system load (Load Shape Following, or LSF). We write these data back into the same table in which the hourly unit emissions stored, and also store the hourly change in total system load into the total system load table. With this new data, we are now ready to calculate the ISO's hourly marginal emission rates.

To calculate the marginal emission rates in each hour, we take a weighted average of the emissions from units following the system load. It's basically the marginal emissions of the marginal generation. We do this as follows: in each hour, we take the rate of emissions (amount of emission in that hour divided by the unit's electricity output) and then weight it by the change in load of that unit with respect to the total change in load of all units that are following the load in that hour. We perform this calculation for each of the three emissions in the data set: CO_2, SO_2, and NO_x. This gives us three hourly time series of marginal emissions rates, expressed as mass of compound emitted per unit of electricity produced.

$$\sum_{lsf} \frac{CO_2}{\text{load}} \frac{\Delta\text{unit load}}{\sum_{lsf} \Delta\text{unit load}} = \text{hourly marginal emission rate} \tag{2.1}$$

Exclusion of unit-hours displaying anomalous emission rates due to measurement lag

In calculating the marginal emissions, we exclude all emissions from qualifying units in the hours that those units' average emissions exceed either 5 tons of CO_2 per megawatt-hour, or 100 pounds of NO_x per megawatt hour. This serves to exclude a number of outlying hours whose emission rates are physically impossible. We theorize that these anomalous data points are the result of the emissions monitoring equipment being mounted downstream of the point at which the operating parameters of the unit are measured, resulting in a lag between the measurements. Thus the emissions caused by burning fuel near the end of an hour may be measured and recorded in the next hour. This effect is especially pronounced when the unit's output is ramping up or down—in steady state, the emissions and power output of the unit will not vary from hour to hour.

Now that we've calculated the marginal emissions rates, it is fairly straightforward to determine emissions avoided by the renewable generation. Since we assume that the electricity produced by the renewable generation can't be stored, and thus must be used in the same hour in which it was produced, emissions will be avoided at the marginal rate that we just computed. Thus, we simply multiply the hourly renewable electricity production time series (expressed in MWh produced in each hour) by the marginal emissions rate (expressed in kilograms of compound emitted per MWh produced by fossil generation in that hour) to give us the avoided emissions in that hour (expressed in mass of compound that was not emitted as a result of the renewable electricity production). We actually calculate this using a nominal 1 MW installed capacity so that the results can be easily scaled to any size of installation under consideration. For example, to compute the emissions avoided by a 12 MW installation of four 3 MW turbines, one simply multiplies the avoided emissions by 12. We conclude this historical analysis by recording the time series of marginal emissions rates and total avoided emissions rates into the database so that they can be quickly accessed for further analysis.

2.4 Forward Looking Analysis

All of the above numerical gymnastics yields a strictly historical analysis. To project the analysis into the future–to anticipate how many emissions might be avoided by a renewable generation project that is yet to be built–we proceed in two steps, separating out insights on the emissions profile over the course of the day and year from the wind generation profile, and then recombining them to give us an anticipated range of avoided emissions. First, we find the mean emissions rate for each season and time of day cohort, which gives us a new matrix of emissions rates for winter day, winter evening, and winter nights, spring day, spring evening, spring nights, etc. Thus the first cell in the matrix contains the mean emissions rate for all winter daytime hours over all years in the dataset. The other cells are similarly filled. This matrix is shown here:

$$\begin{pmatrix} \text{winter day} & \text{spring day} & \text{summerday} & \text{autumnday} \\ \text{winter evening} & \text{spring evening} & \text{summerevening} & \text{autumnevening} \\ \text{winter night} & \text{spring night} & \text{summernight} & \text{autumnnight} \end{pmatrix}$$

With the mean LSF emissions rate for each seasonal and time of day cohort in hand, we turn to the wind resource to develop a range of typical wind years, subdivided into seasons and times of day. We start by determining the mean wind power for each season in year. We then identify the minimum, median, and maximum of each season, so we know which winter had the highest wind, which winter had the lowest wind, and which winter had the median wind. Once we identify those representative seasons, we record the mean wind power generated in each time of day cohort within those seasons. Thus, we build three matrices representing the range (minimum, median, and maximum) of wind speeds that we anticipate in future years; each matrix contains 12 elements, one for each season and time of day cohort.

Now that we have the typical emission rates, and a set of wind generation profiles that covers the anticipated range of wind generation, we merely multiply the the latter by the former, multiply each cell of the resulting matrix by the number of hours in each cohort in a year, and sum the resulting avoided emissions. Thus, we calculate the anticipated avoided emissions for a low wind year, for a medium wind year, and a high wind year. This analysis is admittedly coarse, and relies to a great deal on the emissions profile of the grid remaining the same. This, along with other limitations is discussed in greater detail in Section 2.6.

2.5 Representation of the Results of the Analysis to Policymakers and to the Public

All of the above analysis yields a great deal of data, but the question now before us is: how do we distill those numbers into meaningful insights, and how do we communicate those insights to decision-makers and the public? I answer this question by first determining what the analysis can add to the two groups' understanding of the proposed wind installation, and then I relate the avoided emissions insights to their existing knowledge and interactions with the world.

2.5.1 Decision-makers

The first group I examine is the ostensible target audience of this analysis: those to whom the Environmental Impact Report must be circulated as mandated by section 11.16 of the MEPA regulations. This group includes: the Massachusetts Department of Environmental Protection, the Massachusetts Coastal Zone Management office, the Division of Marine Fisheries, and the Department of Public Health.[14] To these readers, I am communicating the following:

- The range of emissions that we anticipate that the project will offset
- How that anticipated range compares with the historical range of emissions
- How and why our projections differ from the anticipated avoided emissions calculated by the methodology prescribed by the Massachusetts Greenhouse Gas policy
- Why our methodology is more appropriate than the Massachusetts GHG policy's for the purpose of evaluating wind generation

To accomplish these disparate goals, I construct two figures, each communicating two of the above messages. To display the range of anticipated emissions in comparison with the historical range, I construct a coarse time series that that shows some of the seasonal dynamics, but is focused on the total annual emissions. Since the hourly time series of avoided emissions is not of interest, and may in fact provide too much detail, I bin the emissions up by season. I then stack the seasons for each year, and display the consecutive years in order. After the most recent year in our analysis, I show the anticipated range of avoided emissions, similarly binned by season. I demarcate the boundary between the historical emissions and the anticipatory with a strong line, and label both regions appropriately. I then mark the historical range, and extend its boundaries down into the anticipatory region to show that the anticipated range fully contains the historical range. Thus, we have a relatively clear chart that shows our projections, and puts them in the context of the historical range of avoided emissions had the wind turbines been in operation. All emissions on this chart are expressed in terms of quantity of CO_2 avoided per megawatt of installed capacity, so to calculate the total avoided emissions for a given configuration of wind turbines, the decision-maker simply multiplies the emission numbers on the chart by the number of megawatts under consideration. This enables the comparison of several different configurations, but at the expense of some computational complexity, so the anticipated range is calculated and displayed for the primary candidate configuration.

2.5.2 Public

The residents of Hull, on the other hand, are not, for the most part, going to be nearly as interested in the exact amount of CO_2, SO_2, and NO_x that will be avoided by this project. Nor are they going to be terribly interested in the temporal dynamics of the analysis. Instead, what they want to know is, "If these turbines get built, what does that mean in terms of my (and my town's) electricity usage? Plus, I've heard that a couple of different size turbines are being considered, so what's the trade-off for the big turbines compared to the small ones?" To answer these questions, I construct a table of plots. Across the top of the table, I have several candidate turbine configurations: the first, no new turbines, just the existing Hull 1 & 2 output; 3 × 3MW turbines; 3 × 3.6MW turbines, 4 × 3MW turbines, 4 × 3.6 MW turbines. Each of these configurations is represented by a diagram of the turbines, highlighting the relative turbine size and the total capacity. Each row beneath that header row represents a different measurement, expressed in terms of the town's current energy and emissions footprint. Thus, the first row is simply the annual generation, the second avoided CO_2, and so on. The resulting chart, Figure 3-23 shows the residents of Hull how much of their emissions will be avoided for each of the candidate configurations, and puts that next to the currently avoided emissions for the existing turbines.

2.6 Limitations

The AGREA LSF avoided emissions approach has a number of limitations that are primarily driven by uncertainty about the future, lack of data, or restrictions on access to data. Some of these limitations can be addressed through further analysis. I have identified the following primary limitations:

- Lack of access to hourly non-fossil generation data
- Lack of identification of the marginal units in the day-ahead, real-time, and ancillary services markets
- Uncertainty as to the fuel mix and emissions profile of the future grid

2.6.1 Lack of access to hourly non-fossil data

The EGRID data set is, by definition, comprised of emissions and generation data from only those units that emit CO_2, SO_2, or NO_x. This means that the operating parameters of all nuclear, hydropower, and wind plants are not represented, nor are electricity imports or exports. Since our analysis weights the marginal emissions produced by each unit by the total marginal fossil response to the system demand, and this response is assumed to be the total response, there may be a component of responsive generation that is not captured.

Nuclear power plants will generally be generating at full capacity whenever they can, so the lack of access to data on their generation doesn't generally affect our analysis. However, nuclear plants do get taken offline periodically for routine maintenance, and this loss in generation capacity must be made up for by fossil units, which may create artifacts in our emission rates when the fossil units respond to the reduced nuclear capacity instead of the system demand.

Hydroelectric plants comes in two basic flavors: run-of-the-river and reservoir. Run-of-the-river hydroelectricity is very similar to wind and solar power, in that electricity is generated whenever the renewable resource (in this case, running river water) is available. Reservoir hydroelectric, on the other hand, can and is scheduled for generation at the most profitable times, as long as the reservoir isn't full. Since reservoir hydropower can come online very quickly, it is often used in the ancillary services market to balance demand and supply of electric generation. Not having access to the hourly generation profile of reservoir hydroelectric generation does adversely affect our analysis, as this generation does respond to load. However, this effect is mitigated by the relatively small (<5%) proportion of reservoir hydroelectric generation capacity in New England.[11]

Run-of-the-river hydro can by considered alongside wind and solar—these renewable sources of electricity have no marginal costs, and are highly temporally variable, so they generate whenever they can. All three of these generation types do, however, have known seasonal and diurnal patterns, so further analysis can and should be undertaken to mitigate their impact on our marginal analysis. Furthermore, they represent an even smaller proportion of the generative capacity in New England than does reservoir hydropower, so the extent of their impact is limited.

Imports and exports of electricity across the grid boundary are an issue, as they are included in the total system load data from ISO-NE, but are not captured in the EPA CEM dataset. This means that we don't have a complete picture of the source of the electricity that is serving the load at any point in time. Unfortunately, this limitation is more difficult to address than the missing hydro and nuclear generation, as ISO-NE doesn't have nearly as much information about which units are producing the electricity it imports.

Finally, there are a number of fossil plants that are exempt from reporting into the Continuous Emissions Monitoring program. These are units that burn low sulfur (¡ 0.05% by weight), non-coal fuels for a total nameplate capacity of less than 25 megawatts.[4] I do not anticipate that the lack of data on these units will cause a significant adverse impact on my analysis, given their low generation capacity.

2.6.2 Lack of identification of the marginal units in the day-ahead, real-time, and ancillary services markets

Our analytical approach, and in particular our method of identifying the units that are responding to load in any particular hour, is based on the only data we have available to us: the actual operating parameters of each unit in each hour. We know only how each unit is behaving; we don't know why. Specifically, we don't know what sorts of bids the unit operators submitted to the day-ahead, realtime, and ancillary services markets, and which of those bids was accepted. As Connolly shows in his 2008 thesis,[8] plants across the fuel spectrum bid into each of these markets. Since renewable electricity is produced at the rate at which the resource is available, it will likely be bidding into the realtime market, and will thus avoid the emissions associated with the marginal bid in that market.

If we knew which unit the marginal bid was placed for in each hour, our analysis would be more representative of the the actual behavior of the market. Unfortunately for us, while ISO-NE does provide a great deal of information about the bids submitted, it does not identify the units that the bids are submitted for to prevent anticompetitive gaming of the market. While this is good for the economic performance of the market, it does impair our ability to accurately quantify the avoided emissions from renewable generation. However, since ISO New England does have access to this data, it could, as a part of preparing its annual Marginal Emissions Report, perform this analysis while still sufficiently obfuscating the original bid data to prevent collusion.

2.6.3 Uncertainty as to the fuel mix and emissions profile of the future grid

The extension of our historical, marginal LSF analysis into the future raises a number of questions, chief among them: what fuels will generators be burning in the coming years, and how much will they be emitting? While our historical analysis can give us some idea of the answers to these questions in the past, there are several exogenous factors that may affect both the fuel mix and the emissions profile of the future grid.

Of the generating units currently online, the subset of those that are in operation (and their level of output) at any time is determined primarily by the results of the the markets managed by ISO-New England. Since fossil units marginal costs (and thus their bids into the various markets) are heavily dependent on fuel costs, a shift in the price of one fossil fuel relative to another can have immediate effects in the bidding behavior of various plants, and the fuel mix of the set of dispatched plants could be affected as a result.

Because the fossil fuels differ in their carbon intensity (unit of CO_2 emitted per unit of electricity generated), changes in the fuel mix of the grid will be accompanied by changes in the emissions rate of the grid. As more electricity is produced from coal relative to natural gas, more CO_2 is emitted. The relative fuel prices can be affected by any number of factors: their inherent volatility, supply constraints, seasonal variations in demand patterns, etc. Fossil fuel prices, especially those of oil and natural gas, are volatile in and of themselves, and although their variability is somewhat correlated over the long term, in the short term the price of each fuel can jump without the other responding.

The Regional Greenhouse Gas Initiative, New England's first try at a CO_2 cap-and-trade system, will also affect operating costs of plants burning different fuels at different rates, due to the variation in carbon intensities previously discussed. This will have the effect of raising the cost of burning all fossil fuels, and will cost coal plants approximately 1.2 times as much as oil plants, and 1.7 times as much as gas plants per kWh produced. Thus, one of RGGI's potential effects will be to move coal units higher up in the dispatch queue. However, since cap-and-trade systems don't actually set a price for carbon credits, the actual magnitude of RGGI's effect on the economic dispatch of fossil units is uncertain.

Finally, there may be legislative mandates that, given the existence of enabling technologies, restrict emissions in and of themselves. An example of this type of an effect is the precipitous drop in SO_2 and NO_x emissions due to the Clean Air Act. This drop can be clearly seen in our historical analysis of these two emissions: emissions fell every year from 2000 to 2003, and stabilized thereafter. This drop was caused by the mandated installation of SO_2 and NO_x scrubbers onto the stacks of all fossil burning plants. Should

Congress decide to tighten the emission limits, we would expect to see a similar drop in the future, if the technologies are capable of it. Unfortunately, there currently exist no economically viable technologies to scrub CO_2 out of the stacks, so until such technologies are developed, we can safely assume that plants will emit whatever CO_2 they produce.

Chapter 3

Results

3.1 Marginal Emission Rates

Utilizing the methodology described in the previous sections, I generated an hourly time series for the marginal emission rates in the New England grid for each of the following compounds: CO_2, SO_2, and NO_x. The full data set can be seen in the figures on the following pages, in the form of a series of 8760 diagrams, so named because they display each of the 8,760 hours in a year—the 24 hours in a day running across the horizontal axis, and the 365 days in a year descending the vertical axis. Thus, time-of-day patterns can be seen as vertical stripes, whereas seasonal patterns show up as horizontal bands. As can be seen in the diagrams, there is a tremendous amount of variability in and between the emission rates. Let's examine the marginal emission rates for each compound, in turn, to explore some of these episodic patterns.

3.1.1 Marginal CO_2 Emission Rates

The marginal emission rates for CO_2 exhibit two patterns that are of interest to us. First, there is a bimodal diurnal pattern, with a streak of high emission rates in the early morning hours and late evening hours. This can be seen especially clearly in the 2005 8760, which has a yellow streak running down the left side, throughout the year, at around three or four in the morning, and a red streak running down the right side at around 9

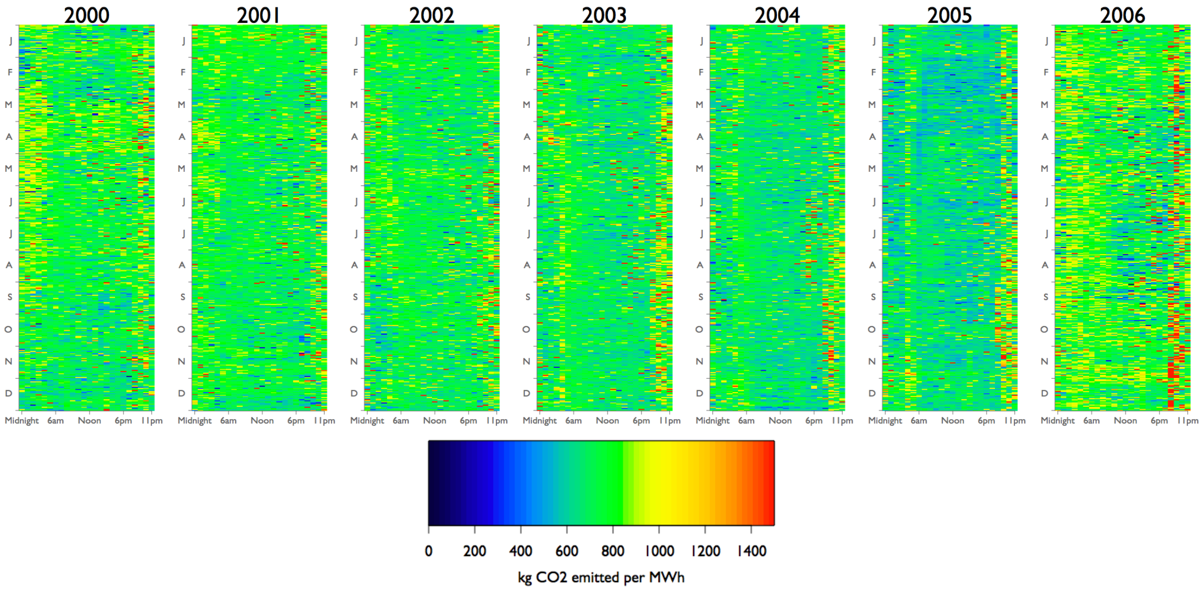

Figure 3-1: Hourly Marginal CO_2 Emission Rates $\left[\frac{\text{kg } CO_2}{\text{MWhr}}\right]$

or 10 o'clock in the evening. This diurnal pattern can also be seen in the hourly average chart in Figure 3-2; note the spike at 9pm. The other pattern is actually interesting in its absence: there is a no discernible seasonal pattern—the summers look much like the winters. This makes the marginal emission rate of CO_2, as I will demonstrate shortly, the exception amongst the three compounds measured here.

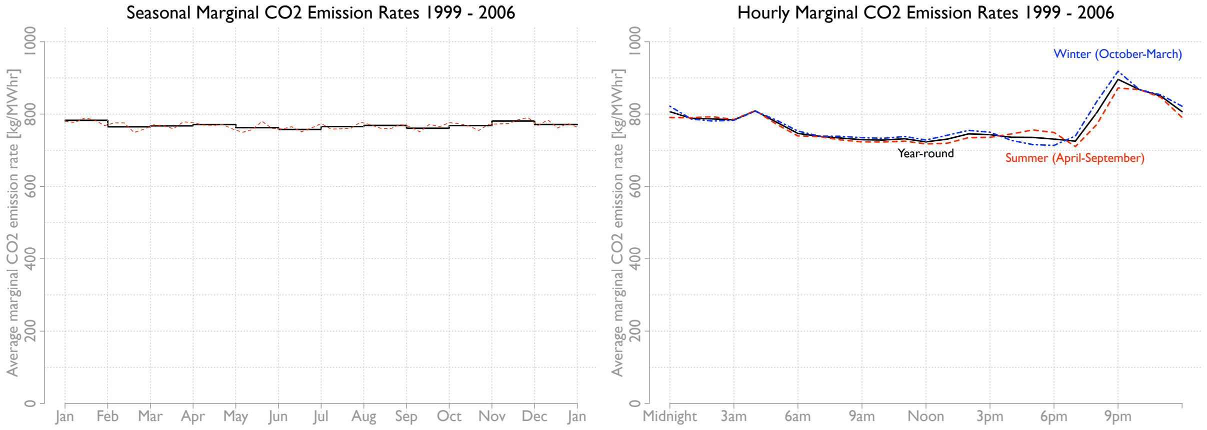

Figure 3-2: Hourly Marginal CO_2 Emissions – Seasonal and Diurnal Trends

3.1.2 Marginal SO_2 Emission Rates

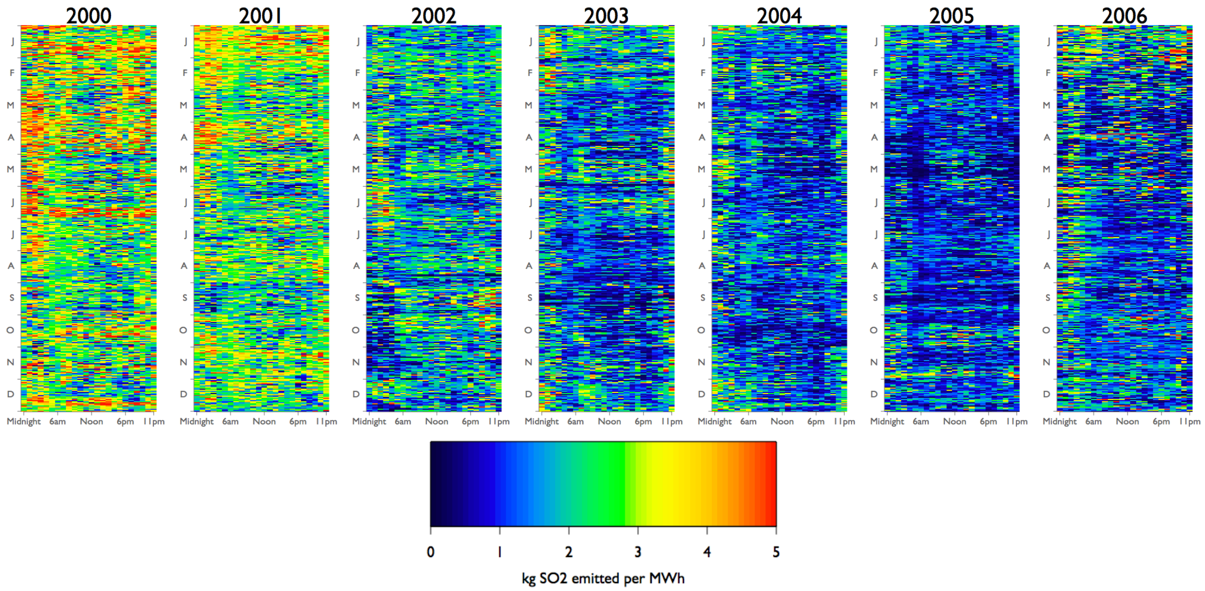

Figure 3-3: Hourly Marginal SO_2 Emission Rates $\left[\frac{\text{kg } SO_2}{\text{MWhr}}\right]$

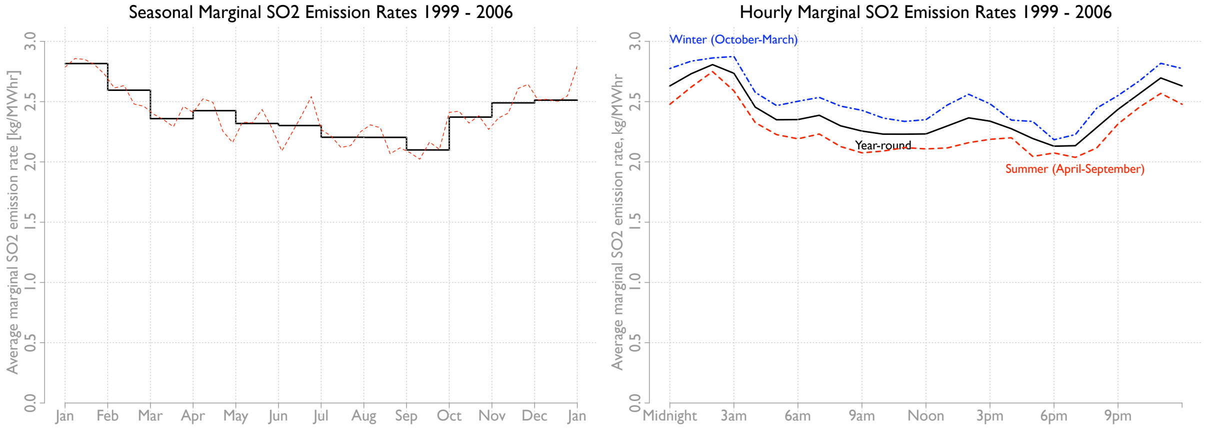

Figure 3-4: Hourly Marginal SO_2 Emissions – Seasonal and Diurnal Trends

In contrast to CO_2, the marginal SO_2 emission rates show far more than simply a diurnal pattern. The first thing that jumps out is the drastic decrease in emission rates from 2000 through 2003. Besides that inter-annual trend, we also see a seasonal pattern that peaks at around 2.8 $\frac{\text{kg } SO_2}{\text{MWhr}}$ a peak in December and drops to a low of around 2.1 $\frac{\text{kg } SO_2}{\text{MWhr}}$ in September. There is also a diurnal pattern with similar range, peaking at 2am and hitting its low at around 6pm. The diurnal pattern's peak and zenith don't vary that much over the course of the year, but there is a decent gap between the summer and winter average in the intervening hours.

3.1.3 Marginal $\mathrm{NO_X}$ Emission Rates

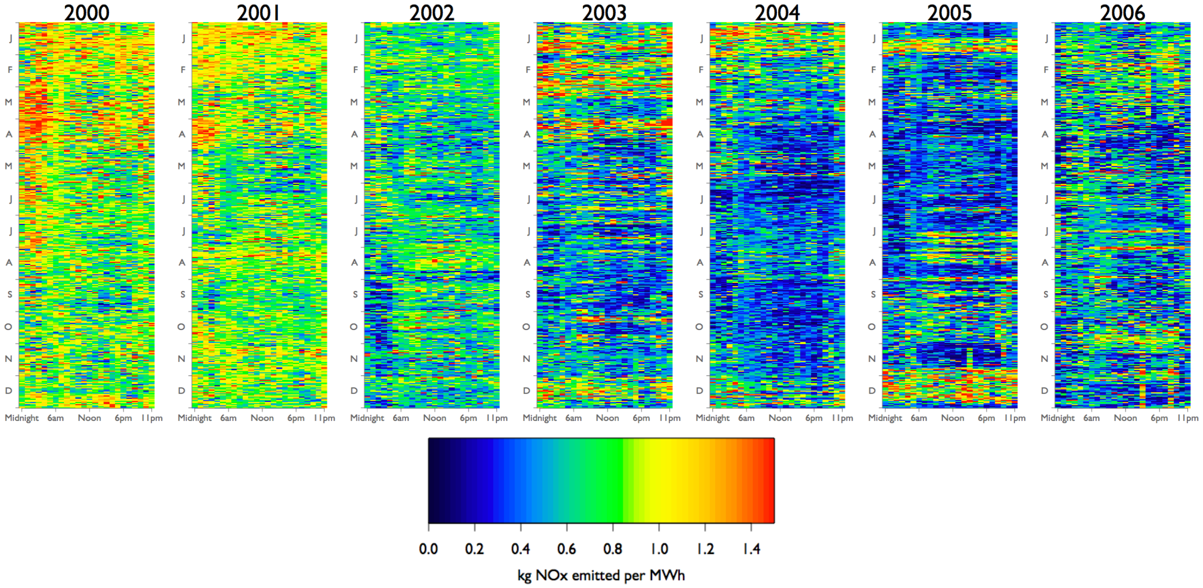

Figure 3-5: Hourly Marginal $\mathrm{NO_X}$ Emission Rates $\left[\frac{\mathrm{kg\ NO_x}}{\mathrm{MWhr}}\right]$

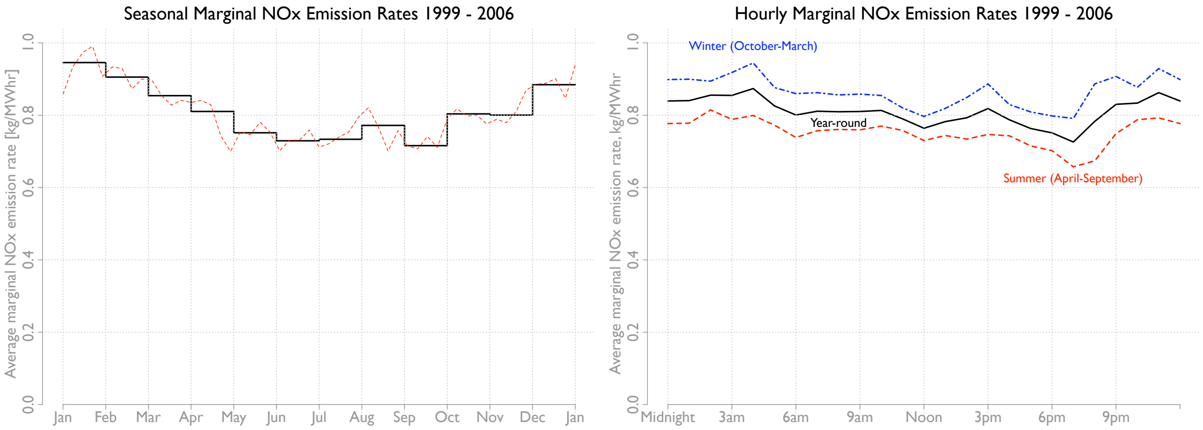

Figure 3-6: Hourly Marginal $\mathrm{NO_x}$ Emissions – Seasonal and Diurnal Trends

The marginal $\mathrm{NO_x}$ emissions rates exhibit some of the same characteristics as the $\mathrm{SO_2}$

rates: the precipitous drop in emissions from 2000 through 2003—which we will explore in greater detail shortly—and seasonal and diurnal patterns. The seasonal pattern is again high marginal emission rates in the winter—peaking at approximately 0.95 $\frac{\text{kg NO}_\text{x}}{\text{MWhr}}$ in December—and low marginal emission rates during the summer—from 0.7 to 0.8 $\frac{\text{kg NO}_\text{x}}{\text{MWhr}}$ from May through September. The diurnal pattern is less pronounced than that of SO_2, but interestingly, the seasonal variation of the diurnal pattern is reversed: the difference between the summer and winter averages of both the peak and zenith emission rates is greater than that of the intervening hours.

3.1.4 Impacts of Clean Air Act on SO_2 and NO_x Emission Rates

The precipitous drop in both NO_x and SO_2 marginal emission rates from 2000 through 2003 can be directly attributed to the implementation of the Clean Air Act, and the emissions trading programs that were instituted as a result. This is a highly encouraging trend to see, as it confirms that well crafted policies have the capability to positively affect the behavior of fossil plants emissions. It should also be noted that it is these very emissions trading programs that necessitated the collection of this level of detailed emissions data, so I owe my ability to perform this analysis upon the Clean Air Act.

3.2 Wind Resource at Hull Offshore Site

3.2.1 Wind Speed

The time series of wind speeds at the Hull Offshore site that I developed by applying the Measure, Correlate, and Predict algorithm to the wind data from the LIDAR on Little Brewster Island and Logan International Airport is depicted here. Two patterns of note can be discerned: first, the seasonal pattern, with a peak from October through March, and a low from July through September. Second, the diurnal pattern, with wind speeds peaking in the afternoon, and a relative lull from midnight to 6am in the morning. The diurnal pattern is more pronounced in the summer, but the seasonal variation of the diurnal pattern is mostly in the overnight lull—in the winter it is more consistently

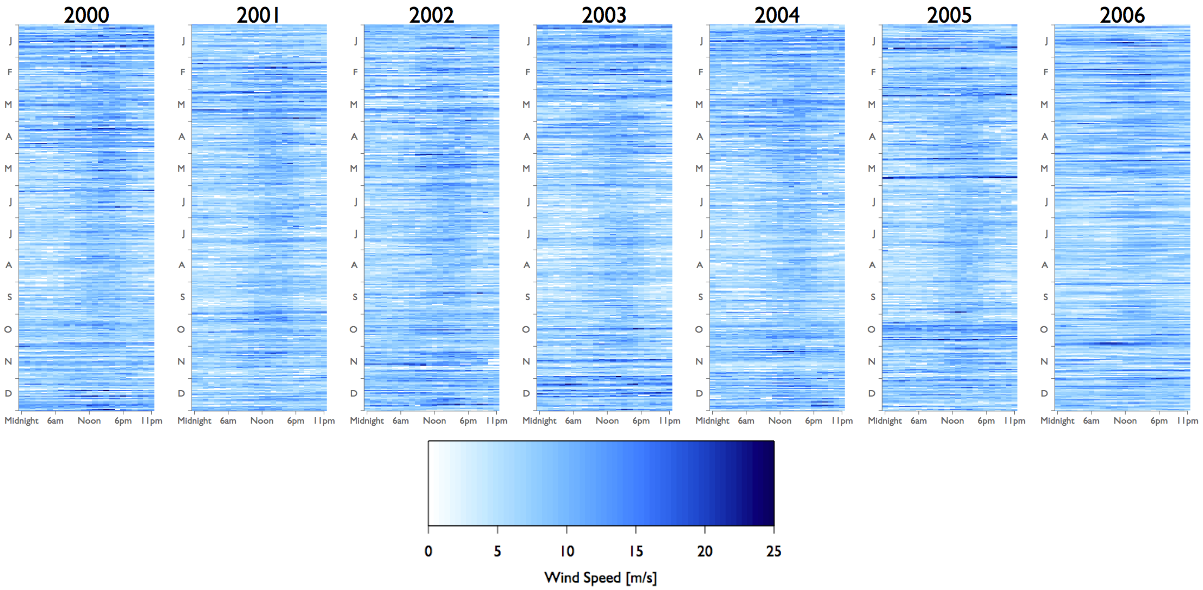

Figure 3-7: Hourly Wind Speed at Hull Offshore Site (80m hub height) $\left[\frac{\mathrm{m}}{\mathrm{s}}\right]$

windy day-round, whereas the summer has calm nights, and the mid-afternoon peak is of similar magnitude year-round.

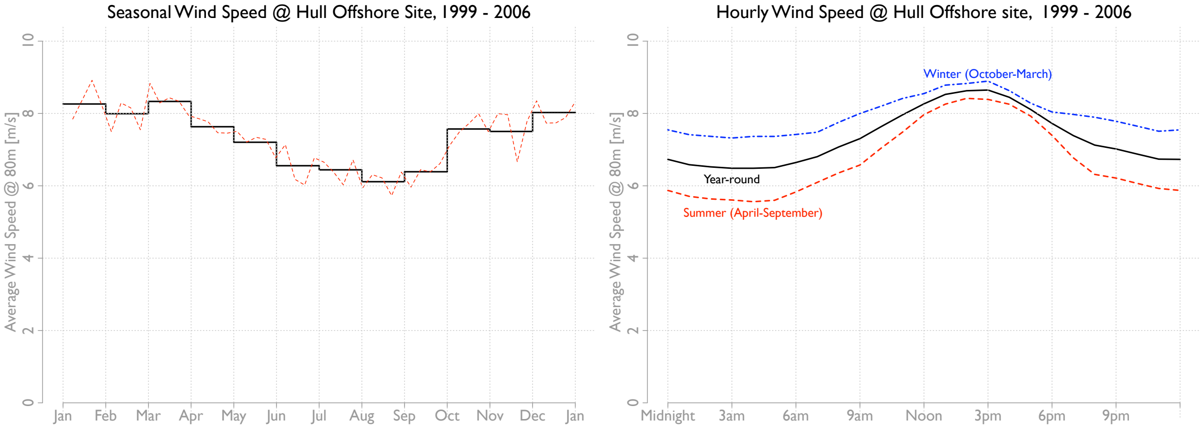

Figure 3-8: Hourly Wind Speed at Hull Offshore Site (80m hub height) – Seasonal and Diurnal Trends

3.2.2 Wind Energy Production

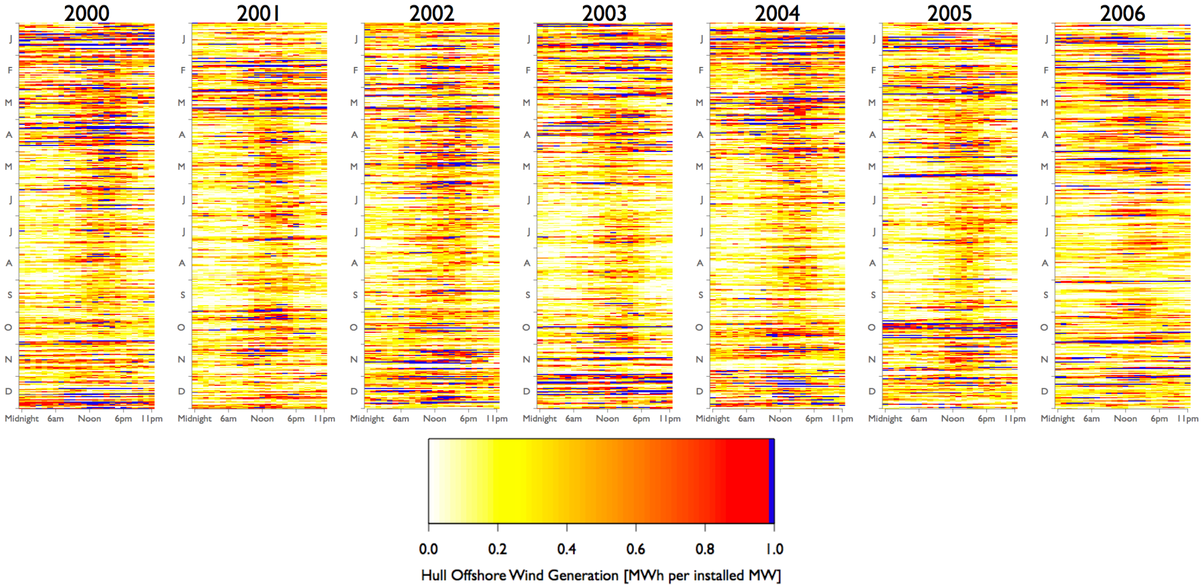

Figure 3-9: Hull Offshore Hourly Wind Energy Production $\left[\frac{\text{MWh generated}}{\text{MW installed}}\right]$

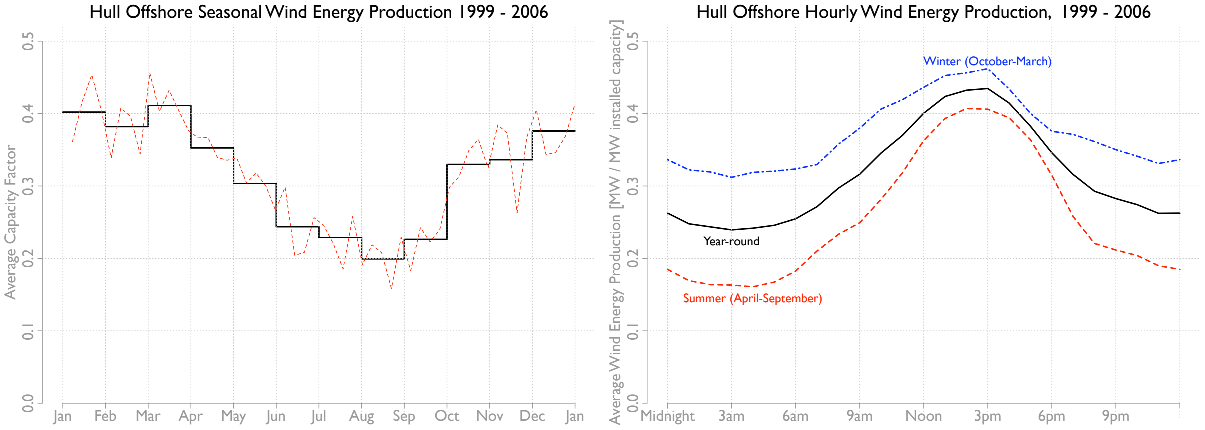

Figure 3-10: Hull Offshore Wind Energy Production – Seasonal and Diurnal Trends

When I mapped the time series of wind speeds to the power curve for the candidate wind turbine (the GE 3.6s), the temporal characteristics of the variation of the time series was preserved, but the range of the variation was reduced and the variation within that range was increased. This was due to a couple of factors: first, the turbine produces no power in wind speeds less than 4 $\frac{m}{s}$, and produces its maximum capacity at all speeds above 15 $\frac{\mathrm{m}}{\mathrm{s}}$. Second, from 4 to 15 $\frac{\mathrm{m}}{\mathrm{s}}$, the power output increases drastically for every additional bit of wind speed available—typical behavior for turbines, whose output increases as the cube of the wind speed. Thus, we have the same seasonal and diurnal patterns as the wind speed, but with larger variation, as can be seen in the trend plots in Figure 3-10.

3.2.3 Net Generation

So now that I've shown the dynamics of the wind generation, how does that match up against the Town of Hull's patterns of electricity demand? Using the hourly load data supplied by the Hull Municipal Light Plant, I generated an hourly times series of the Town of Hull's net generation, assuming 15 MW of additional offshore wind capacity. The results can be seen in Figure 3-11: green areas represent times of surplus generation—during which time the HMLP can sell into the grid—and the red areas are when the Hull demand outstrips the wind power generated in that hour.

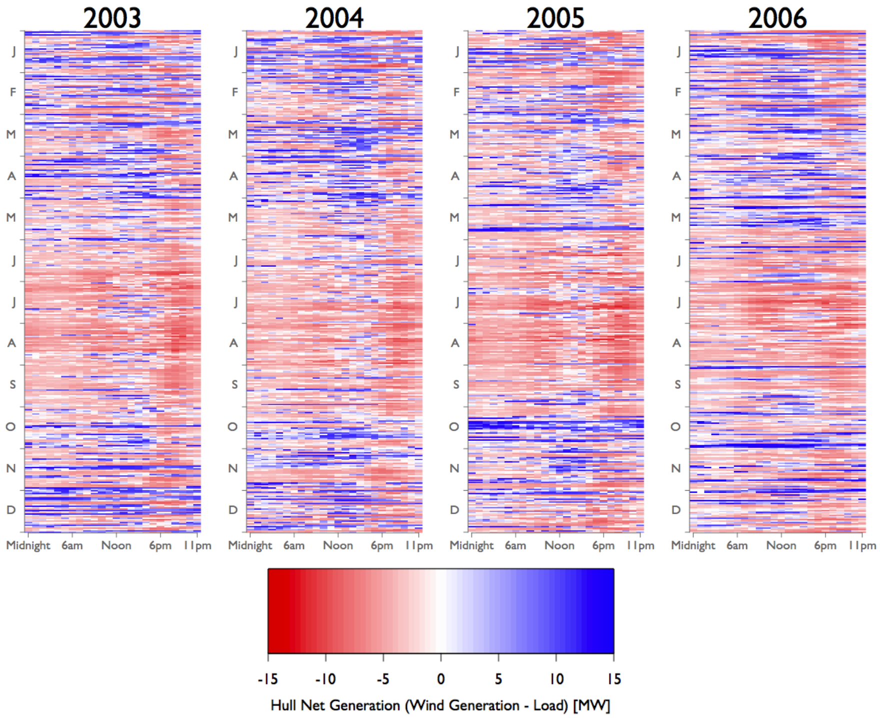

Figure 3-11: Town of Hull Net Load with 15 MW of Offshore Wind installed plus Hull Wind I and II

What is striking is that even though there are clearly times of the year—winter, primarily—when there is generally more wind than there is demand for electricity, there are net-positive generation hours throughout the year, which is testament to the tremendous variability and episocity of the wind resource. Although the Town of Hull will still need to purchase electricity throughout the year, during the winter months, it becomes a net exporter, on average, as can be seen in Figure 3-12. The diurnal pattern shows an early afternoon peak as wind energy production matches against the mid-afternoon lull in load, followed by a precipitous drop in net generation as everyone comes home and turns on their televisions just as the wind dies down at the end of the day. Keep in mind that these results, in particular, are specific to the Town of Hull and its usage patterns. Hull does not have a great deal of commercial or industrial activity, which results in a lower daytime peak load.

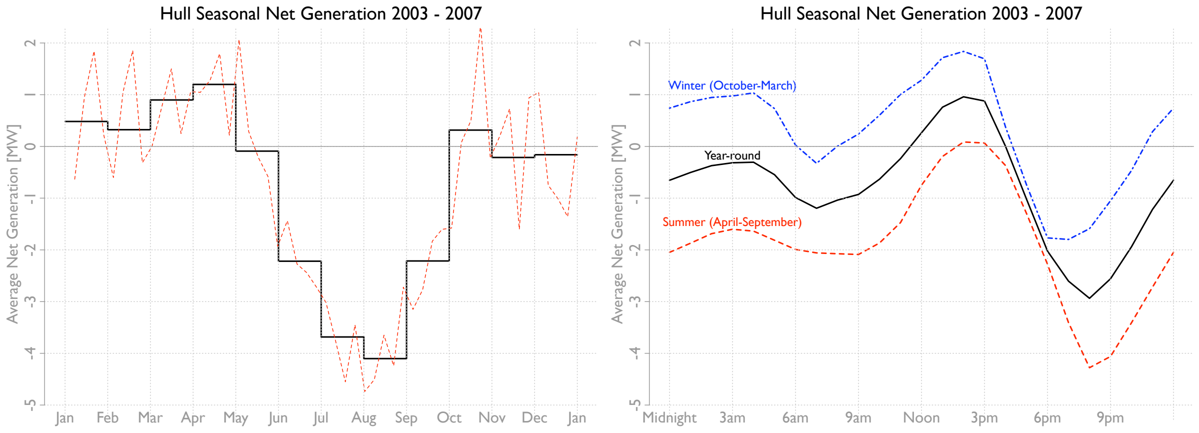

Figure 3-12: Hull Offshore Wind Energy Production – Seasonal and Diurnal Trends

3.3 Hull's Historical Avoided Emissions

Now that we've got both the marginal emission rates for the New England grid and the hourly capacity factor at the Hull Offshore site, we multiply them together to get the hourly avoided emissions. In doing so, the seasonal and diurnal pattens in the marginal emission rates and the wind resource will interact—in some cases magnifying each others effects, in other cases canceling them out—yielding a new set of trends. Let's take a look at what happens.

3.3.1 Avoided CO_2 Emissions

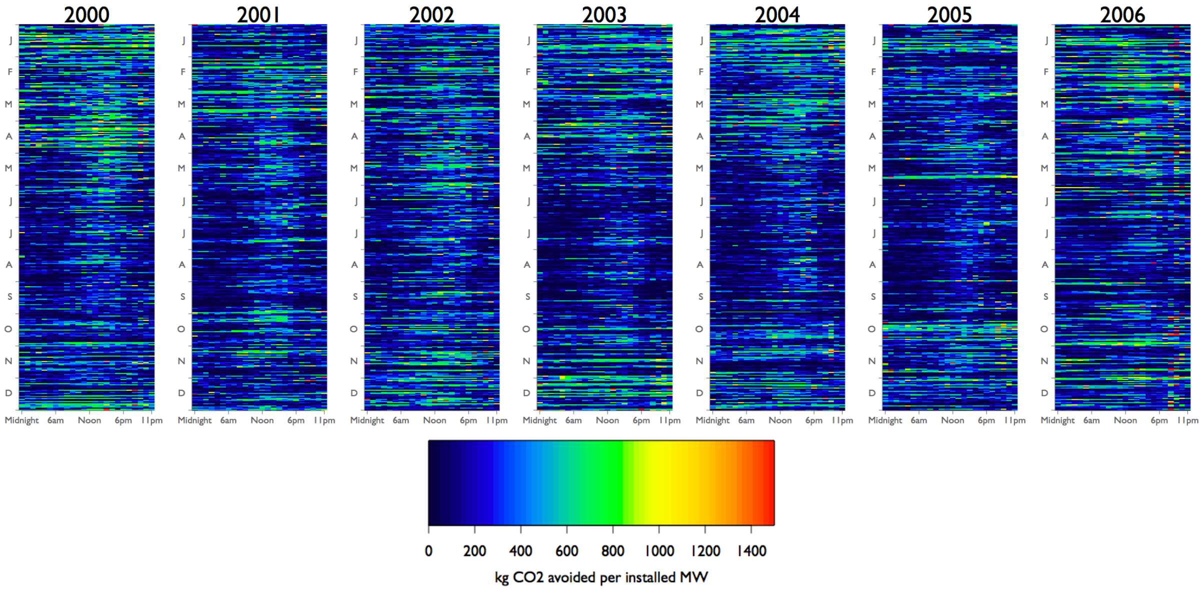

Figure 3-13: Hull Offshore Avoided CO_2 Emissions $\left[\frac{\text{kg } CO_2 \text{ avoided}}{\text{MW installed}}\right]$

If you recall, the marginal CO_2 emission rate exhibited little seasonal variation, and only a small diurnal pattern. Thus, the patterns in the hourly avoided CO_2 are almost solely the result of the temporal patterns of the wind resource: high in winter, low in summer, high afternoons, and low overnight. The two features that do bubble through from the marginal emission rates are the spike in avoided emissions at 9pm, and a small bump at around 4am, as can be seen in Figure 3-14

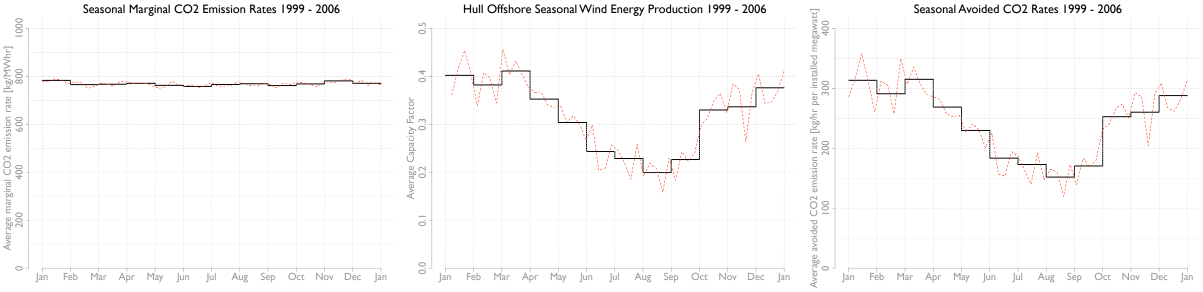

Figure 3-14: Hull Offshore Avoided CO_2 Emissions – Seasonal Trends

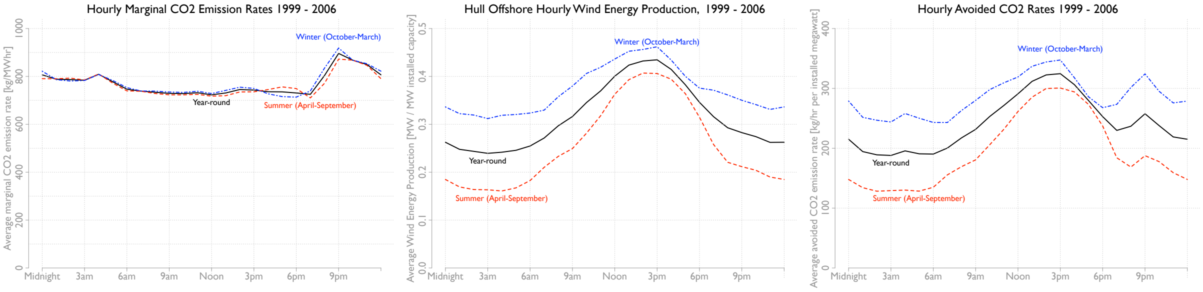

Figure 3-15: Hull Offshore Avoided CO_2 Emissions – Diurnal Trends

3.3.2 Avoided SO_2 Emissions

The avoided SO_2 emission rates, on the other hand, show more seasonal variation, and less diurnal variation than the marginal emission rates and the capacity factors from which they are calculated. The seasonal pattern is amplified, resulting in a higher ratio of winter avoided emission rates to the summer rates. The diurnal pattern is muted, as the peaks and valleys of the marginal emissions rates and the capacity factors do not coincide.

3.3.3 Avoided NO_X Emissions

Finally, we come to the avoided NO_x emission rates. Once again, this mirrors the SO_2 patterns, amplifying the seasonal pattern and dampening the diurnal variation. Although the diurnal range is reduced, the inter-seasonal variation of the diurnal patten is increased, mostly due to a lower summer peak in the afternoon.

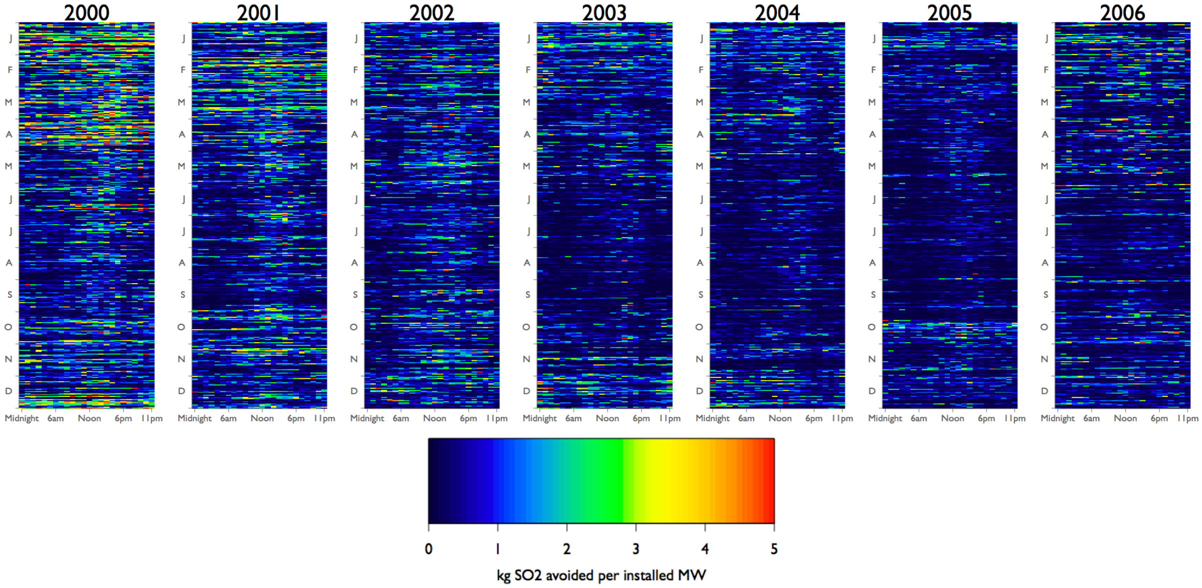

Figure 3-16: Hull Offshore Avoided SO_2 Emissions $\left[\frac{\text{kg SO}_2\text{ avoided}}{\text{MW installed}}\right]$

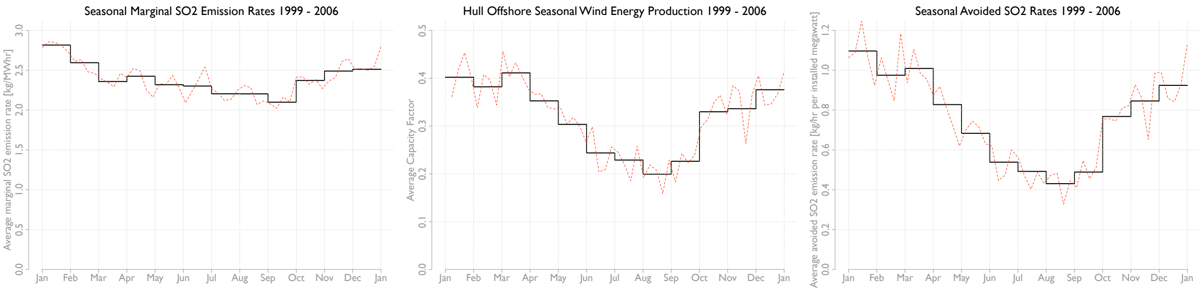

Figure 3-17: Hull Offshore Avoided SO_2 Emissions – Seasonal Trends

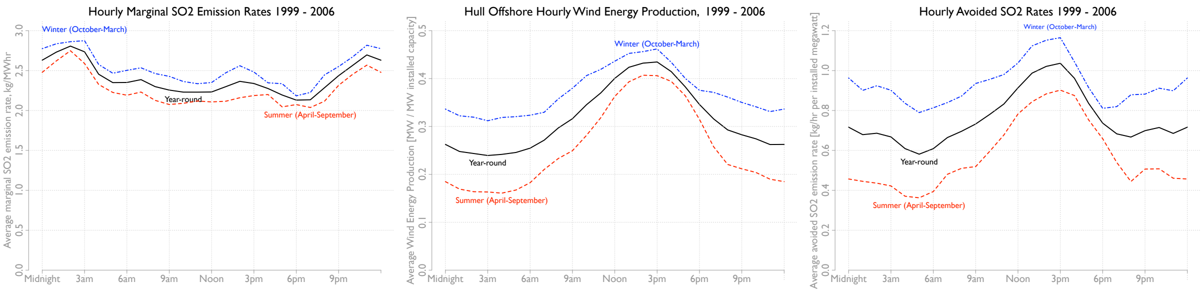

Figure 3-18: Hull Offshore Avoided SO_2 Emissions – Diurnal Trends

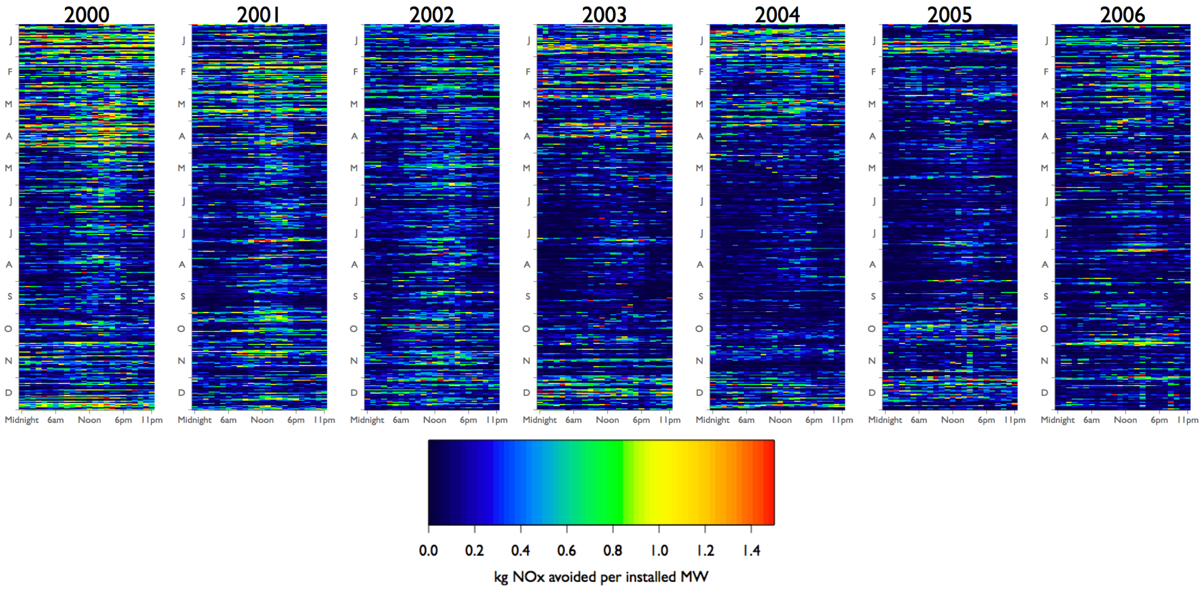

Figure 3-19: Hull Offshore Avoided NO_x Emissions $\left[\frac{\text{kg NO}_\text{x}\text{ avoided}}{\text{MW installed}}\right]$

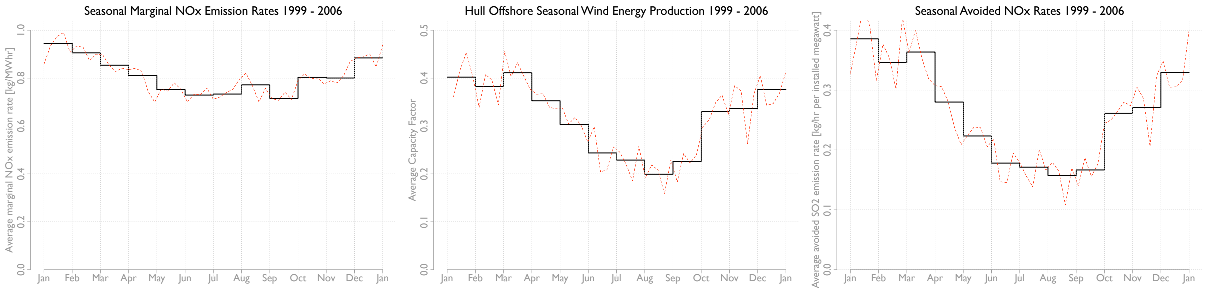

Figure 3-20: Hull Offshore Avoided NO_x Emissions – Seasonal Trends

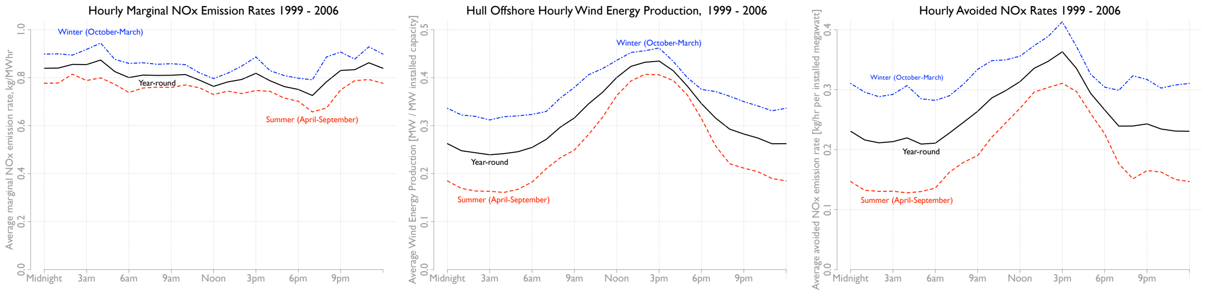

Figure 3-21: Hull Offshore Avoided NO_x Emissions – Diurnal Trends

3.4 Anticipated Avoided Emissions

Now that we have a complete understanding what emissions would have been avoided had the wind turbines been installed off the coast of Hull from 2000 through 2006, we turn our gaze to the future. To develop our anticipated range of avoided emissions, I build up a characteristic set of marginal emission rates and a range of capacity factors, broken down into seasonal and time-of-day cohorts. I use both of those sets of numbers to produce a range of anticipated avoided emissions for each compound.

3.4.1 Marginal Emission Rate Cohort Means

The first step, as previously described in the Section 2.3, is to calculate the mean marginal emission rates in each season and time-of-day cohort in the historical dataset. For CO_2, I use the full range of years, but for the other two compounds, I only take the mean of the last 3 years, from 2004–2006, as I anticipate that the emission reductions resulting from the Clean Air Act will be enduring. The mean rates for each compound in each cohort are shown in Tables 3.1, 3.2, and 3.3.

Mean Marginal CO_2 Emission Rates 2000-2006 $\left[\frac{\text{kg CO}_2}{\text{MWhr}}\right]$

	Winter	Spring	Summer	Fall	All Year
Day	740	735	733	725	733
Evening	848	829	835	901	854
Night	791	790	782	779	785
All Day	773	767	764	770	769
Δ	Winter	Spring	Summer	Fall	All Year
Day	-3.7%	-4.3%	-4.7%	-5.7%	-4.6%
Evening	10%	7.9%	8.6%	17%	11%
Night	3.0%	2.8%	1.7%	1.4%	2.2%
All Day	0.6%	-0.2%	-0.6%	0.2%	–

Table 3.1: Mean marginal CO_2 emission rates by seasonal and time-of-day cohorts

Mean Marginal SO_2 Emission Rates $\left[\frac{\text{kg } SO_2}{\text{MWhr}}\right]$
2004-2006 (post-Clean Air Act)

	Winter	Spring	Summer	Fall	All Year
Day	1.49	1.20	1.11	1.20	1.25
Evening	1.71	1.03	1.26	1.31	1.33
Night	1.87	1.35	1.48	1.44	1.54
All Day	1.64	1.21	1.24	1.29	1.34
Δ	Winter	Spring	Summer	Fall	All Year
Day	11%	-11%	-18%	-10%	-7.2%
Evening	27%	-23%	-6.4%	-2.9%	-1.4%
Night	39%	0.6%	10%	7.0%	14%
All Day	22%	-9.7%	-7.6%	-4.1%	–

Table 3.2: Mean marginal SO_2 emission rates by seasonal and time-of-day cohorts

Mean Marginal NO_x Emission Rates $\left[\frac{\text{kg } NO_x}{\text{MWhr}}\right]$
2004-2006 (post-Clean Air Act)

	Winter	Spring	Summer	Fall	All Year
Day	0.692	0.507	0.500	0.506	0.551
Evening	0.766	0.492	0.505	0.605	0.591
Night	0.794	0.533	0.484	0.541	0.587
All Day	0.734	0.512	0.496	0.533	0.568
Δ	Winter	Spring	Summer	Fall	All Year
Day	22%	-11%	-12%	-11%	-3.1%
Evening	35%	-13%	-11%	6.4%	4.0%
Night	40%	-6.1%	-15%	-4.8%	3.4%
All Day	29%	-9.8%	-13%	-6.2%	–

Table 3.3: Mean marginal SO_2 emission rates by seasonal and time-of-day cohorts

3.4.2 Range of Capacity Factor

In order to construct a range of anticipated wind power production, I first split the entire hourly time series of capacity factors into yearly, seasonal, and time-of-day cohorts. I then take the minimum, median, and maximum across all years from each of the seasonal and time-of-day cohorts, and then reconstruct three full years, one for each of the minimum, median, and maximum. Thus, I have a low wind year, medium wind year, and high wind year, broken down by seasonal and time-of-day cohorts, as shown in Table 3.4.

Low Wind

	Winter	Spring	Summer	Fall	Year-round
Day	0.363	0.382	0.253	0.302	0.325
Evening	0.341	0.284	0.139	0.204	0.242
Night	0.305	0.263	0.110	0.204	0.220
All-Day	0.343	0.331	0.192	0.257	0.280

Medium Wind

	Winter	Spring	Summer	Fall	Year-round
Day	0.405	0.427	0.275	0.331	0.359
Evening	0.337	0.296	0.181	0.256	0.267
Night	0.348	0.253	0.136	0.240	0.243
All-Day	0.377	0.354	0.219	0.292	0.310

High Wind

	Winter	Spring	Summer	Fall	Year-round
Day	0.466	0.483	0.342	0.377	0.417
Evening	0.468	0.349	0.195	0.271	0.320
Night	0.404	0.269	0.161	0.265	0.274
All-Day	0.448	0.398	0.265	0.327	0.359

Table 3.4: Range of mean capacity factors, by seasonal and time-of-day cohorts

3.4.3 Anticipated Avoided Emission Results

It is then a simple operation to multiply the mean marginal emissions in each cohort for each compound by the corresponding cohort in each of the anticipated capacity factor years. This produces a range of anticipated avoided emission rates for each compound, still broken down by seasonal and time-of-day cohort. I obtain the final annual avoided emission totals by multiplying each cohort by the number of hours in that cohort in a year, and summing across cohorts. This produced the following ranges of avoided emission for each compound:

CO_2	Avoided Emissions $\left[\frac{\text{tonnes } CO_2 \text{ avoided}}{\text{installed MW} \cdot \text{year}}\right]$	Δ
High Wind	2,280	15.70 %
Med. Wind	1,970	-
Low Wind	1,780	-9.62 %

SO_2	Avoided Emissions $\left[\frac{\text{kg } SO_2 \text{ avoided}}{\text{installed MW} \cdot \text{year}}\right]$	Δ
High Wind	4,030	15.90 %
Med. Wind	3,480	-
Low Wind	3,140	-9.66 %

NO_x	Avoided Emissions $\left[\frac{\text{kg } NO_x \text{ avoided}}{\text{installed MW} \cdot \text{year}}\right]$	Δ
High Wind	1,730	16.00 %
Med. Wind	1,490	-
Low Wind	1,350	-9.57 %

Table 3.5: Range of anticipated avoided emissions, [tonnes per installed megawatt per year]

I have performed all of the preceding contortions to obtain the these numbers: in a medium wind year, I anticipated that the Hull Offshore project would avoid, per installed megawatt, the emission of 1,970 tonnes of CO_2, 3,480 kg of SO_2, and 1,490 kg of NO_x, assuming that the grid mix and dispatch behavior remains the same. In a high wind year, I anticipate that these quanities increase by about 16%, and in a low wind year, I anticipate avoiding approximately 10% less than in the medium wind year. This means that if the project is built out to 15MW, the town of Hull can expect to avoid the emission of 29,500 tonnes of CO_2, 52.1 tonnes of SO_2, and 22.4 tonnes of NO_x in a typical year.

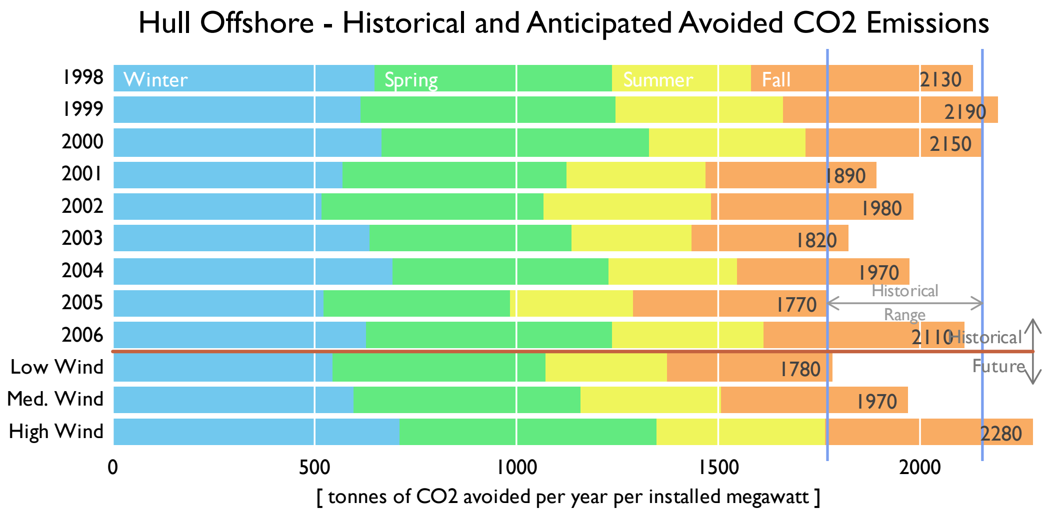

Hull Offshore - Historical and Anticipated Avoided CO2 Emissions

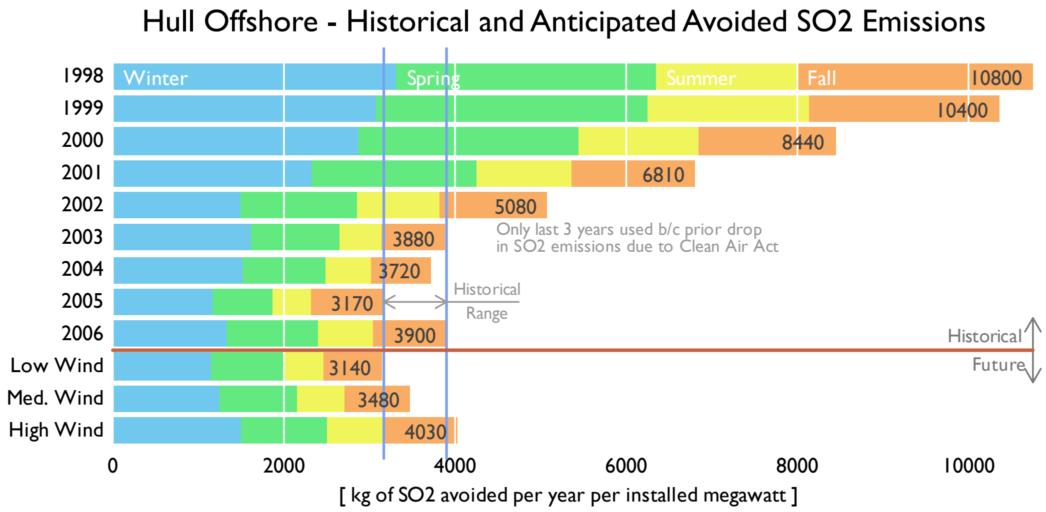

Hull Offshore - Historical and Anticipated Avoided SO2 Emissions

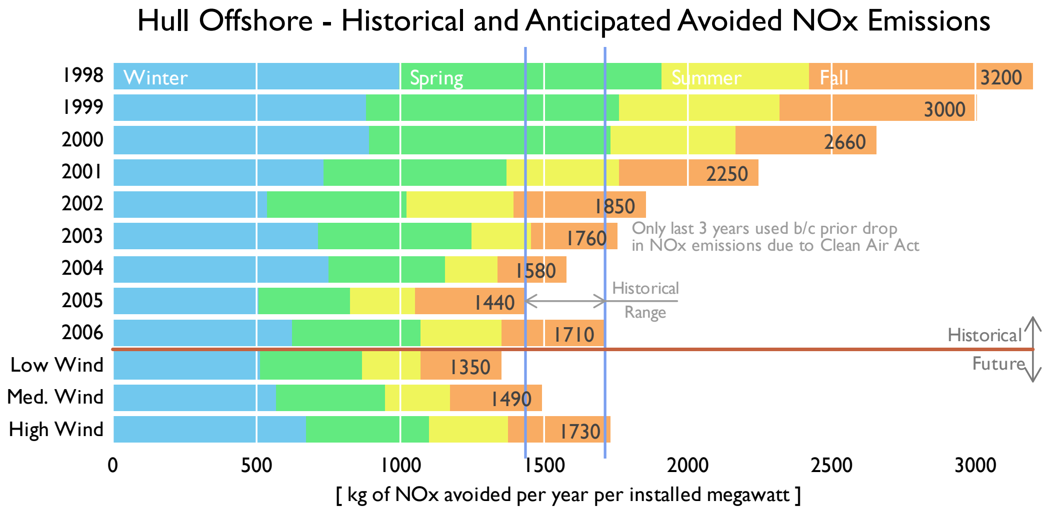

Hull Offshore - Historical and Anticipated Avoided NOx Emissions

Figure 3-22: Avoided Emissions

In Relation to Historical Electricity Usage in Hull

I have shown you the anticipated avoided emissions for the Hull offshore wind project, but what does that mean to the people of Hull? What does 29,500 tonnes even mean? To answer that question, I put the anticipated avoided emissions in terms of Hull's emissions footprint, and to do that, I will use the average emissions rate[1] of the grid and the hourly time series of Hull's electricity usage.

		Potential Impact of Offshore Wind in a medium wind year		
		10 MW	12 MW	15 MW
MWh	Power	25,800	31,000	38,700
kg	CO_2	19,700,000	23,600,000	29,500,000
kg	SO_2	34,800	41,700	52,100
kg	NO_x	14,900	17,900	22,400

Table 3.6: Anticipated emissions benefits of Hull Offshore project

Hull's Current Footprint		Potential Impact of Offshore Wind		
	Yearly Average from 2003–2006	10 MW	12 MW	15 MW
Power	53,300 MWh	48%	58%	73%
CO_2	19,300,000 kg	102%	122%	153%
SO_2	50,300 kg	69%	83%	104%
NO_x	15,400 kg	97%	116%	145%

Table 3.7: Anticipated benefits of Hull Offshore project relative to Hull's current footprint

These may be somewhat surprising results—we anticipate that the Offshore Wind Project will offset 102% of Hull's CO_2 emissions by supplying 48% of the town's electricity. Why the apparent discrepancy? There are two primary reasons: first, the average emission rates used to determine Hull's footprint are low because they are diluted by non-fossil, non-emitting generation. Second, the marginal emission rates are generally higher than the average fossil emission rate, and therefore much higher than the total emission rates when we include non-fossil generation. In essence, the turbines do more

---

[1] I have defined the "average emissions rate" as the total emissions from all fossil units in an hour divided by the total system load, which includes non-emitting generation such as hydro and nuclear.

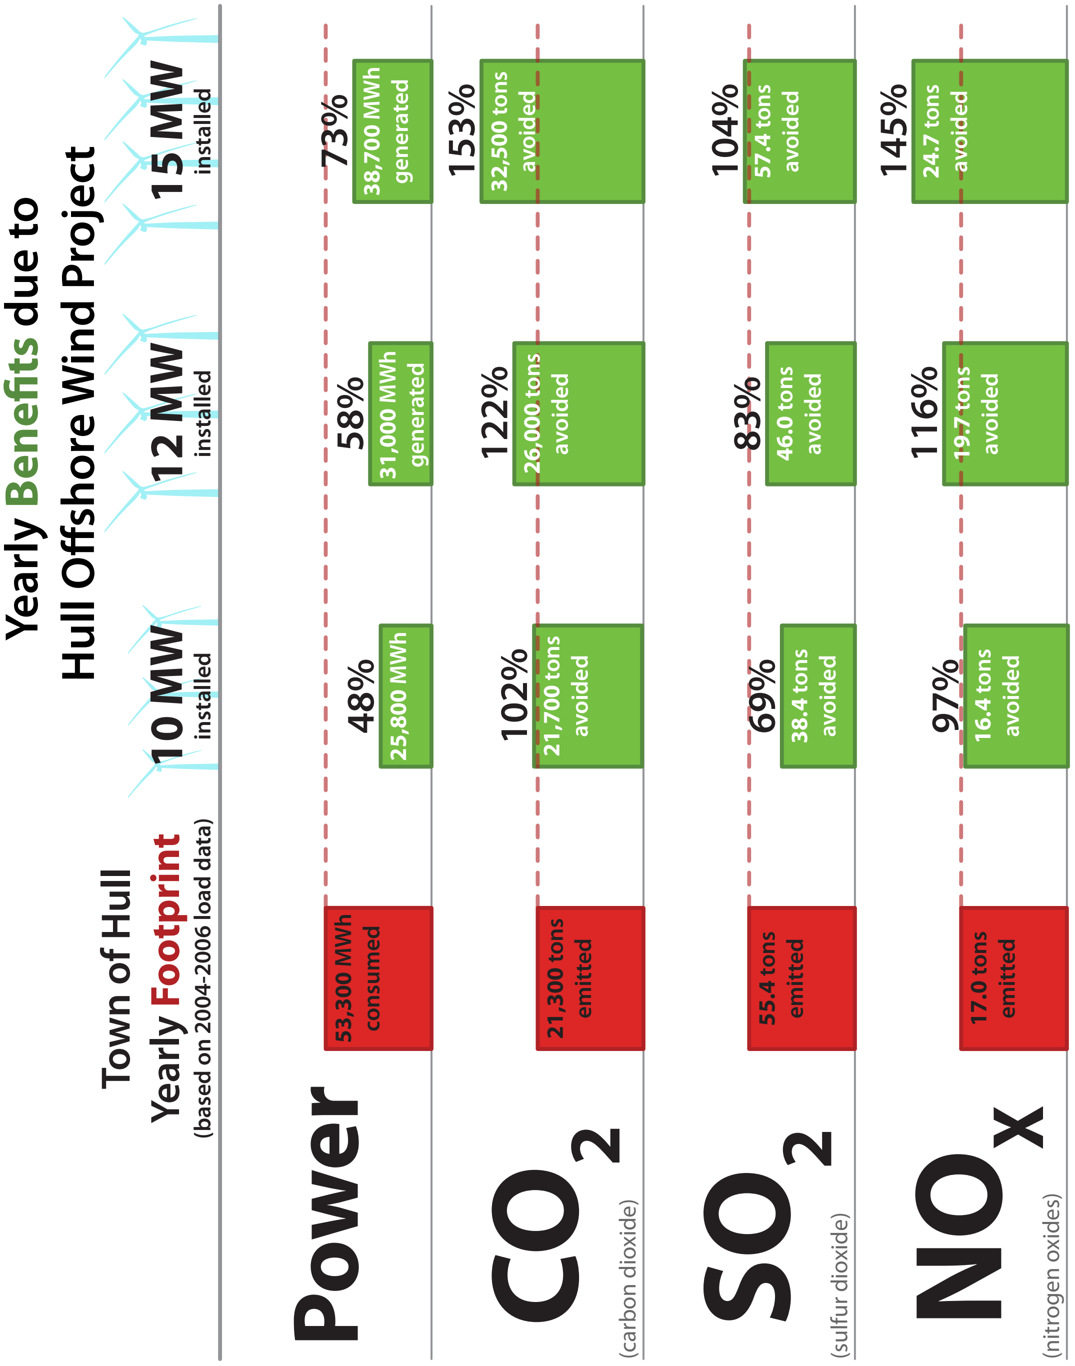

Figure 3-23: Hull Offshore benefits

good (in terms of emissions benefits) than the Town of Hull did bad through its electricity usage. So there we have it: building out the Hull Offshore project to 15 megawatts of capacity will supply the town of Hull with 73% of its electricity, and will offset all of its SO_2, and over 140% of its CO_2 and NO_x. This same data is also shown in Figure 3-23, with the emissions converted to imperial units for easy consumption by the residents of Hull.

3.5 Compared to Massachusetts GHG Policy

We now know what emissions we anticipate avoiding according to our marginal analysis, but unfortunately, our methodology is not the one mandated by the Massachusetts Greenhouse Gas Policy. Therefore I will now compare the potential avoided emissions using their methodology with the avoided emissions using our methodology, so that we can see how the approaches differ. So, let's now to the MA GHG Policy for direction: the nut of the technique is to calculate the total electricity the project will consume, and then "multiply [that quantity of electricity] by an emissions factor that calculates the CO2 emitted through the generation of electricity."[2] Now, although the Policy specifies that the emissions factor (i.e. rate) to be used should come from the ISO New England 2005 Marginal Emission Rates Analysis, it says nothing about which factor to use, as the Analysis has both on-peak and off-peak factors, in additional to an annual average. [2] Since the Policy is ambiguous as to the proper factor to use, I chose the more detailed approach of applying the on- and off-peak factors to the appropriate hours in which the turbines would be generating electricity, and used the proportion of Hull on-peak and off-peak generation. I thus calculated that for every megawatt of installed capacity, this project would avoid 1,320 tonnes of CO_2, 2,080 kg of SO_2, and 701 kg of NO_x every year, according to the Massachusetts GHG Policy. This is well below the quanitities of emissions that I anticipate avoiding using the marginal analysis, as can be seen in table Table 3.8.

---

[2] The Policy indicates the ISO New England 2005 Marginal Emission Rates Analysis contains CO_2 emission rates "for a variety of stationary combustion sources"; however, the Analysis contains only three CO_2 two rates, which are aggregated over all intermediate (oil and gas) power plants.

	AGREA Marginal	MEPA GHG Policy	Δ
CO_2 [kg]	1,970,000	1,320,000	-33%
SO_2 [kg]	3,480	2,050	-41%
NO_x [kg]	1,490	690	-54%

Table 3.8: Comparison of AGREA marginal avoided emissions analysis vs. MEPA Greenhouse Gas Policy [quantities per installed megawatt per year]

So why the huge (33%-54%) discrepancy? There are a number of reasons that can be summed up as follows: first, the use of just two emissions factors (on-peak and off-peak) ignores the tremendous episodicity of both wind power and emission rates; and second, the definition of "intermediate units" as any unit that burns oil or gas is inaccurate. Fundamentally, however, the problem is one of an unfamiliarity with the evaluation of renewable generation. The Massachusetts GHG policy is intended to ensure that developers—e.g. energy consumers—take steps to mitigate their greenhouse gas *emissions*,[3] and thus the protocol for evaluating the mitigation measures is necessarily simple. It is, however, completely inappropriate for a renewable generation project such as this one, for which reduced greenhouse gas emissions are a primary motivating factor. The entire project is a mitigation measure, and as such deserves a far more nuanced analysis than simply applying a year-round emissions factor to the total anticipated electricity generation.

The next three sets of plots (Figures 3-24, 3-25, and 3-26) explore some of the relative dynamics of the AGREA LSF emissions rate directly against those of the MA GHG Protocol. Plotted on each of the figures is the average AGREA LSF emission rate for the appropriate time period, the relevant marginal emission factors from the 2005 ISO-NE Marginal Emissions Report (mandated by the MA GHG Protocol), and the average emission rate of all oil and/or gas units, the rate upon which the ISO-NE's marginal emissions are putatively based. I'm only looking at the emissions rates for 2005, which accounts for the increased variation in the seasonal and hourly rates, but this is necessary to compare

---

[3] As a matter of fact, the MA GHG Policy states that the only electricity to be measured is that which is consumed, not generated.

the AGREA approach with that of the Massachusetts Greenhouse Gas Protocol.

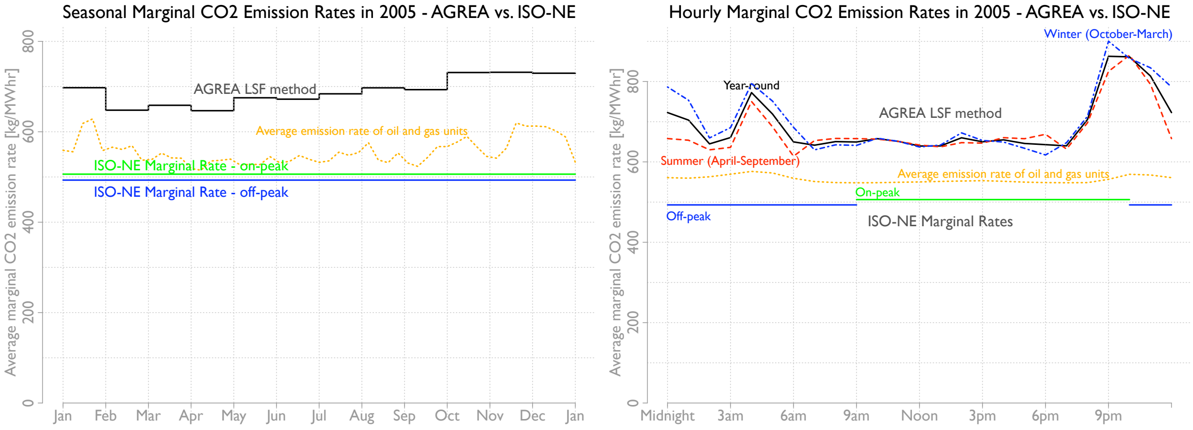

Figure 3-24: Seasonal and diurnal comparison between AGREA LSF and MA GHG CO_2 emission rates

The CO_2 emission rates, shown in Figure 3-24, lack seasonal variation, which makes for a rather staid story—the ISO-NE marginal rates are fairly close to one another, and they are well beneath the AGREA LSF rate. What *is* interesting to note is that the ISO-NE rates are also below the average emissions of oil and gas units, which is supposed to be what the marginal emissions are derived from. As I will show you presently, this effect is not seen in the SO_2 or the NO_x rates, so this may warrant further investigation. On the diurnal side of things, two points of note: first, the ISO-NE Marginal on and off peak rates do not reflect the average emissions in those hours - the average emissions actually dip slightly during the daytime and ramp up a bit at night. The major morning and evening spikes in the AGREA LSF rate are, of course, not captured.

Turning to Figure 3-25, the SO_2 emission rates, on the other hand, behave largely as we would expect. The average oil and gas emissions display similar gross seasonal dynamics as the AGREA LSF emission rates, with the exception of a deeper trough from April though July. Unsurprisingly, the static ISO-NE rate doesn't pick up any of this, but does approximate the average magnitude better than it did for CO_2. When we turn to the diurnal trends, the ISO-NE marginal rates match the oil and gas averages quite accurately, although both of these miss the overnight spike seen in the AGREA LSF rates. The AGREA LSF rates do show increased emission rates during the on-peak

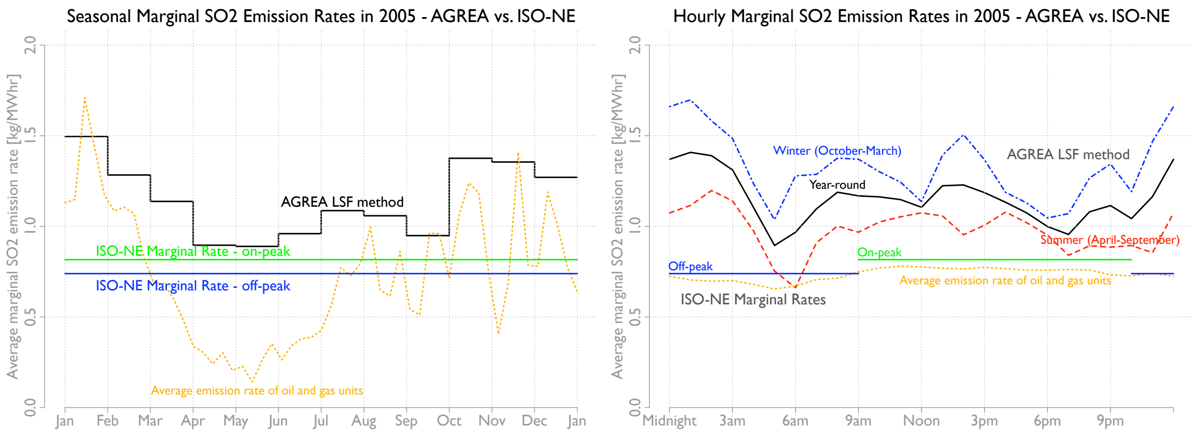

Figure 3-25: Seasonal and diurnal comparison between AGREA LSF and MA GHG SO_2 emission rates

hours, albeit shifted earlier in the day by about two hours.

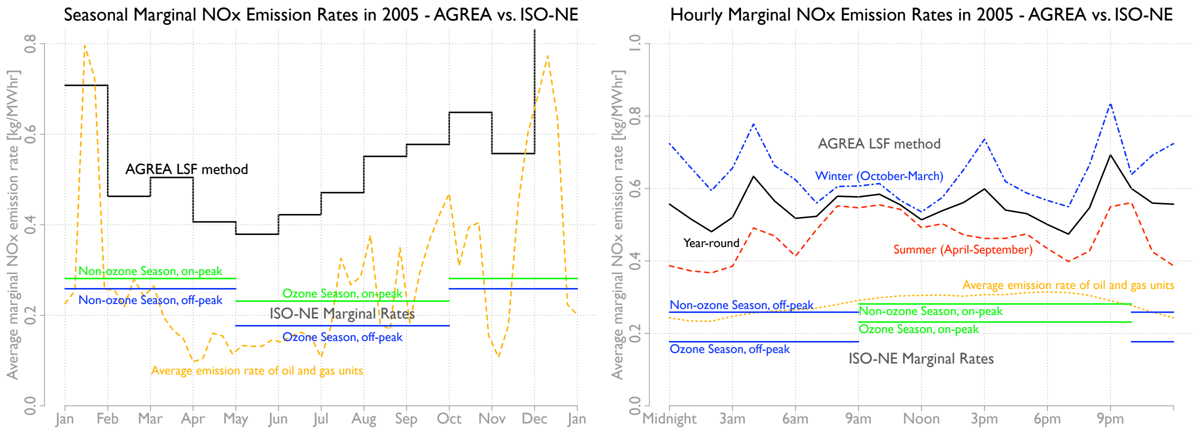

Figure 3-26: Seasonal and diurnal comparison between AGREA LSF and MA GHG NO_x emission rates

Finally, Figure 3-26 shows the NO_x emission rates, which exhibit significant seasonal as well as diurnal variation. Indeed, the 2005 ISO-NE Marginal Emissions Report recognizes the seasonal variation by separate calculating on- and off- peak rates for both ozone and non-ozone seasons. Unfortunately, as can be seen from the plot, the ozone season does not actually correspond to the period of low NO_x emission rates—it appears that the actual dip occurs from April through July, whereas ozone season is defined as May through September. This offset may be a result of 2005 being an anomalous year with respect to emission rate patterns. This period of lower emission rates *is* reflected in the AGREA LSF rate. Once again, we see that both the ISO-NE Marginal Rates and the average oil and gas emissions are significantly less than the AGREA LSF marginal rates. On the diurnal front, the average oil and gas emission rates match up well with the hourly pattern of on- and off-peak ISO-NE marginal rates.

THIS PAGE INTENTIONALLY LEFT BLANK

Chapter 4

Conclusions

So what do these results mean for the evaluation of the environmental benefits of renewable generation, and for Massachusetts nascent Greenhouse Gas Policy? Fundamentally, what I have shown is that you *need* to be operating at this level of detail—hourly time-series—to accurately quantify the avoided emissions of renewable generation. The coarse, yearly average emission factors mandated by the Massachusetts Greenhouse Gas Policy hide too much detail. The rationale behind using these coarse numbers is ease of computation, and this is entirely appropriate for most greenhouse gas mitigation measures, which are the focus of the Policy. However, the developers of renewable generation build—as a matter of course—detailed temporal forecasts of their anticipated generation, and it is really a trivial matter to cross-multiply these forecasts with the appropriate marginal emission rate time series.

The Massachusetts Greenhouse Gas Policy is an admirable first effort towards encouraging GHG mitigation measures in construction projects, but the fact that the Protocol was requested for the Hull Offshore Wind project is an indication that the Policy's purpose is not fully understood by the Executive Office of Energy and Environmental Affairs. The Policy states, quite clearly, that its intent is to encourage the *mitigation* of greenhouse gas emissions, to the extent that the Policy doesn't so much as mention clean generation of any sort.

The Massachusetts Executive Office of Energy and Environmental Affairs is to be commended, for recognizing the importance of reducing greenhouse gas emissions, but it needs the right tools for the right job. For construction projects, mitigation measures are appropriate, but as more renewable energy is proposed and developed, higher resolution tools such as AGREA's hourly marginal emissions algorithm are more appropriate, and do not incur a large additional computational burden on renewable developers, who will have already acquired detailed forecasts of their renewable generation.

Furthermore, the myriad stakeholders in renewable generation development—consumers, renewable suppliers, distribution utilities, the system operator, and state environmental regulators—all have different analytical requirements and require even the same data and information to be presented in different manners to help them understand their piece of the puzzle. The "one size fits all" approach of the MA GHG Policy is not appropriate for such a diversity of needs, so what is needed is a suite of tools that can operate off of a common, agreed-upon dataset.

I spent approximately a third of my time on this thesis just looking for data and normalizing it into the forms necessary for the analysis, which leads me to my next point: The potential role of federal and state agencies and ISO-NE in collecting and publishing long-term time series of the data that will support a complete suite of analytical tools. The US EPA has led in this effort with their publication of the hourly emissions measurement from their continuous emissions monitoring program, and they aren't standing still—they are currently revising the format in which they publish the CEM data to the eXtensible Markup Language, which will facilitate automated parsing and analysis of the data without the need for the convoluted import process that is currently a feature of the AGREA method.

The Town of Hull deserves recognition for its bold vision of wind power for its residents, and the analysis in this thesis validates their decision. However, the question before us is now how do we make this data set and methodology accessible to other communities in the Commonwealth? The answer to that question has two parts. First, the public report that will come out of this thesis will contain full details on the methodology, but more importantly the AGREA database will be populated with marginal emission rates through 2007 in the next month. Greater access to this database may well enable communities across the Commonwealth to consider the full range of benefits of renewable generation. To make these analyses feasible, the communities will also need detailed wind resource data for their locale, so support on that front is also necessary.

4.1 Future Work

Since this is the first prospective application of the AGREA marginal emissions analysis, a follow-up study would be advised to determine how closely our forecasts match the actual generation patterns of both the Hull Offshore project, should it be built, and the changing mix of generation resources available on the New England grid.

ISO-NE's Role in Calculating Marginal Emissions

The Independent System Operator in New England has a unique role that it can play in this analysis. The AGREA LSF methodology is based on determining the units that are responding to load according to their actual behavior. We do this for a number of reasons, but the fundamental one is that we don't know why a unit changes its output because ISO-NE only releases anonymized bid data to prevent anticompetitive behavior. The ISO *does* have access to the full bidding data, in addition to the complete information on dispatch. They know who the marginal units are at any point in time, and with this information, they could, and should, calculate the actual marginal emission rates on an hourly basis. ISO-NE could then publish this time series, just as it currently does with the bidding and generation data, and the level of aggregation would be such that anti-competitive safeguards would be retained. The publication of such a dataset would go a long way to helping prioritize renewable generation in the ISO, and would be in the ISO's best interests.

THIS PAGE INTENTIONALLY LEFT BLANK

Bibliography

[1] US Energy Information Administration. Emissions of greenhouse gases in the united states 2006, November 2006.

[2] Massachusetts Environmental Protection Agency. Mepa greenhouse gas emissions policy and protocol, 2007.

[3] US Environmental Protection Agency. Clean air markets. Accessed at http://www.epa.gov/airmarkt/index.html.

[4] US Environmental Protection Agency. Code of federal regulations, title 40 § 72.7.

[5] US Environmental Protection Agency. Continuous emissions monitoring data. Accessed at ftp://ftp.epa.gov/dmdnload.

[6] US Environmental Protection Agency. The emissions & generation resource integrated database (egrid). Accessed at http://www.epa.gov/cleanenergy/energy-resources/egrid/index.html.

[7] Michael Berlinski and Stephen Connors. Economic and environmental performance of potential northeast offshore wind energy resources. Technical report, MIT Laboratory for Energy and the Environment, January 2006.

[8] J. P. Connolly. Effect of real-time electricity pricing on renewable generators and system emissions. Master's thesis, Massachusetts Institute of Technology, 2008.

[9] Stephen Connors, Katherine Martin, Michael Adams, Edward Kern, and Baafour Asiamah-Adjei. Emissions reductions from solar photovoltaic (pv) systems. Tech-

nical Report 2004-003, Laboratory for Energy and the Environment, Massachusetts Institute of Technology, August 2005.

[10] ISO New England. 2005 new england marginal emission rate analysis, July 2007.

[11] ISO New England. 2006 annual markets report, June 2007.

[12] ISO New England. Hourly historical data post-market, 2008. Accessed at http://www.iso-ne.com/markets/hrly_data/index.html.

[13] James F. Manwell, Jon G. McGowan, Christopher N. Elkinton, and John A. MacLeod. Status report on the hull offshore wind project. In *Windpower 2007 Conference*. American Wind Energy Association, 2007.

[14] Code of Massachusetts Regulations. 301 cmr § 11.16.

[15] Anthony L. Rogers, John W. Rogers, and James .F. Manwell. Comparison of the performance of four measure-correlate-predict algorithms. *Journal of Wind Engineering and Industrial Aerodynamics*, 93(3):243–264, March 2005.

[16] Hugh Sharman. Why wind power works for denmark. In *Proceedings of ICE - Civil Engineering*, volume 158, pages 66–72, May 2005.