Environmental Predictions and the Energy Sector: A Canadian Perspective

Case Study Report

24 December 2007
Final Version

Prepared for

Environment Canada

Contract Number: K3A40-06-0028

Prepared by

TURQUOISE
TECHNOLOGY SOLUTIONS INC.

DYDEX
DYDEX RESEARCH AND CAPITAL LTD.

This page intentionally left blank.

Setting the Stage

The purpose of this case study report is to build on the findings and observations of the Literature Review Report by providing concrete examples on the economic benefits of applying Environmental Prediction (EP) for the planning and management of the energy system.

The four case studies presented in this report are selected from different sectors of the energy industry, ranging from upstream oil and gas to renewable energy, and they focus on applications from Canada as well as abroad. For each case study, a valuation framework has been developed in order to estimate the associated economic benefits with a particular EP application.

This report is composed of four self contained case studies. In order to facilitate the navigation through the case studies, we recommend the reader to follow the order as they are presented in this report. Some common pieces of information, such as how wind speed can be used to calculate power output of a wind turbine, are presented once and then cross-referenced in subsequent case studies.

The first case study examines the benefits of EP in upstream oil and gas operations in Canada, in particular the impact of sea ice on operational and strategic decisions, and how EP can be used to make better decisions. This case study provides background information on sea ice, offshore oil and gas operations in Canada, and the threat of sea ice to these operations.

The second case study investigates the benefits of EP for the supply side of hydroelectricity by analyzing, at a conceptual level, the value which can be extracted from EP corresponding to the operational, tactical, and strategic time scales. This case study describes many different ways in which hydrological EP can add value, and follows up on four of them for a detailed quantitative analysis in a Canadian setting.

The third case study highlights the importance of EP when wind turbines are used in tandem with hydroelectricity assets. This synergy has the potential of addressing some of the concerns about wind energy variability, and EP is slated to play a very important role in managing the multiple sources of uncertainty which have an impact on the strategic and operational decisions of wind and hydro generators.

The final case study aims to document applications of EP for various types of renewable energy systems, including off-shore wind and solar energy. All of these applications incorporate satellite data to complement other data sources. As part of this case study, a valuation framework is presented which demonstrates the economic benefit of using satellite data for off-shore wind energy resource assessment.

List of Interviewees

(In alphabetical order by surname)

- Dr. Lindsay Anderson, Environmental Engineering, Cornell University
- Mr. Philippe Beaucage, Doctoral Candidate, INRS, Varennes, Québec
- Mr. Malcolm Dewhurst, Meteorologist, Oceans Inc St John's Newfoundland
- Dr. Hans Duivenvoorden, Manager, Wind Energy, EcoFys Netherlands BV
- Mr. Barry Kirby, Mariner, Newfoundland
- Dr. Adam Kucera, Manager, Quantitative Analysis, Integral Energy, Australia
- Ms. Terry Kwas, Stakeholder Relations Manager, TransAlta Wind
- Mr. Pierre Lang, independent consultant, Montreal
- Mr. John Langan, Water resources consultant, Stantec Engineering, London Ont.
- Ms. Lin Li, (M.Sc. Hydrologist), Systems Analyst, TransAlta
- Dr. Shawn Marshall, Canada Research Chair in Glaciology, University of Calgary
- Mr. Desmond Power, P.Eng, also with C-CORE (interviewed with Mr. Ralph)
- Mr. Freeman Ralph, P.Eng,, Manager, Ice Engineering Centre for Cold Ocean Research & Engineering (C-CORE), St John's Nfld
- Mr. Gaëtan Roberge, Bureau de Recherche, Hydro Québec
- Dr. Vernon Singhroy, Senior Research Scientist, Canada Center for Remote Sensing Natural Resources Canada
- Dr. Hans Tuenter, Risk Management, Ontario Power Generation
- Dr. John Yackel, Sea Ice Expert, University of Calgary
- Dr. Jun Yang, Director General, National Satellite Meteorological Center, China Meteorological Administration
- Dr. Luis Fernando Zarzalejo, CIEMAT, Spain

Summary Table of Economic Benefits from Applications of Environmental Predictions

Description	Benefit	Type of Forecast	Estimated Value
Sea Ice	▪ Making better decisions to manage sea ice threats and managing the overall risk of an off-shore platform	▪ Operational	▪ Expected annual savings of approximately $2.4 million for a 100,000 bpd well if EP is used along with an escort vessel.
		▪ Strategic	▪ EP involved in the decision to deploy $1 billion gravity base structure (Hibernia) instead of deploying cheaper FPSO vessel.s
Supply Side of Hydro	▪ Various benefits ranging from optimal storage to minimizing spinning reserves.	▪ Operational	▪ Reducing the need for spinning reserve results in (very limited) benefits of $55,000 for the Alberta Market.
			▪ Better operation of microhydro facilities (with limited storage) through EP results in approximately $750,000/year value.
		▪ Tactical	▪ Combining snowpack measurements with El Nino/Southern Oscillation climatology yields approximately $3.6 million for the Alberta market.
		▪ Strategic	▪ EP as part of the decision process in embarking on a $1 billion improvement to the hydro generation from Niagara Falls.
Wind/Water Integration	▪ Better operational decisions if a wind energy generator uses EP and has access to a pump storage in order to manage wind variability and maximize revenues	▪ Operational/ Tactical	▪ 1% reduction in the mean absolute percentage error (MAPE) of load forecasts decreases variable generation costs by 0.1% - 0.3%.
		▪ Strategic	▪ EP as a means of reducing the non-dispatchability of wind power, allowing for higher green power penetrations.
Earth Observation for Renewables	▪ More accurate long-term revenue forecasts for off-shore wind farms	▪ Strategic	▪ Only relying on meteorological mast measurements for a year can result in significantly lower revenue estimates from an off-shore wind farm (approximately $5.75 million/year less revenue for the base case). Satellite data driven EP can create value by avoiding wrong decisions based on poor estimates.
			▪ Just 1% error in long-term mean wind speed estimate results in a revenue reduction of $350,000/year.

Table of Contents

Table of Figures

List of Tables

List of Acronyms

ENSO	El Nino Southern Oscillation
EP	Environmental Predictions
ESA	European Space Agency
GPS	Geographical Positioning System
HDD	Heating Degree Days
IEA	International Energy Agency
IESO	Independent Electricity System Operator
kWh	Kilowatt hour
MWh	Megawatt hour
NASA	National Aeronautics and Space Administration
NOAA	National Oceanic and Atmospheric Administration
SAR	Synthetic Aperture Radar

Note: Unless otherwise noted, all currency figures in this report are in Canadian dollars.

Environmental Predictions and the Energy Sector: A Canadian Perspective

Case Study 1: Impact of Sea Ice on Upstream Oil and Gas Operations

Prepared for

Environment Canada

Contract Number: K3A40-06-0028

Prepared by

TURQUOISE TECHNOLOGY SOLUTIONS INC.

DYDEX
DYDEX RESEARCH AND CAPITAL LTD.

1. Introduction

The purpose of this research project is to gather information on the economic benefits of applying Environmental Prediction (EP) for the planning and management of the energy system. In particular, the economic benefit of EP across the whole energy value chain in Canada is investigated.

The first in a series, this case study aims to investigate the benefits of EP in the upstream oil and gas operations in Canada, in particular the impact of sea ice on operational and strategic decisions, and how EP can be used to make better decisions.

The case study provides background information on sea ice, offshore oil and gas operations in Canada, and the threat of sea ice to these operations in Sections 2, 3, and 4. The role of EP and valuation of benefits are discussed in Section 5.

2. The Science of Sea Ice

Sea ice is simply frozen sea water. Icebergs, though posing similar dangers to economic activity, are not considered as sea ice, since they form on land and then fall into the ocean. Nonetheless, this case study considers the risks posed to offshore oil and gas extraction by sea ice and icebergs together, as both pose similar dangers and may be analyzed in similar ways. Research on how sea ice is formed, destroyed, and transported is currently of great importance, particularly to climate change models (Marshall, 2007).

Sea ice plays an important role in such "big picture" models, based on its two main characteristics: its reflectivity and its impact on the salinity of surrounding water. Ice reflects much more solar radiation than water does, and it has a direct effect on the albedo of the planet. As far as the impact on salinity is concerned, when sea ice is formed, it increases the salt concentration and hence the density of the water below it. This causes the body of water to sink and thereby contribute to the "global conveyor belt" which drives ocean circulation. Both of these characteristics make it imperative to use large scale modelling techniques in the global circulation models (GCMs) which are used for climate change research.

This case study concentrates on the energy sector impact of sea ice at relatively short time scales. Indeed, as will be shown below, formation and destruction of sea ice constitutes a very interesting case study for Environmental Prediction (EP), as it requires the integration of many forecasts including temperature, insolation, wind, and ocean currents in a framework requiring nontrivial modelling of thermodynamics, ice mechanics, and even the Coriolis force.

2.1. Formation of Sea Ice

It is interesting to note that the solid form of water, ice, has a lower density than its liquid form. However, sea water is thermodynamically quite different from fresh water, making the formation of sea ice more complicated than the formation of lake ice (as distinct from river ice, which presents its own modelling challenges)[1]. First, due to its salinity, salt water has a lower freezing point (and higher boiling point) than fresh water. For typical sea-water salt concentrations, the freezing point of water is -1.8 degrees Celsius. A less well known difference between the thermodynamics of fresh and salt water is that very cold fresh water, like ice, becomes less dense as it reaches the freezing point. This makes it easy for lake ice to form – cold water floats upward until a thin sheet of ice forms, thickening as the cold water continues to float upwards. Thus both convective and conductive heat transfer assist in the formation of lake ice, allowing it to form quite rapidly.

Salt water is different: the salt in sea water causes the density to increase as it reaches the freezing point, so that very cold sea water sinks. The ice itself still floats and the water still freezes top down, but convection works against the formation of sea ice by continually removing the coldest water from the surface region. Generally, the top 100-150 meters of water must be cooled to the freezing temperature for ice to form. As a result of these differences, sea ice forms more slowly compared to freshwater ice.

2.2. Steps in the Formation of Sea Ice

The following six steps outline the formation of sea ice.

(i) First frazil crystals, small needle-like ice crystals, form. Frazil crystals are nearly pure fresh water.

(ii) The accumulation and bonding of floating frazil crystals form sheets of sea ice. In calm waters this process forms a smooth thin form of ice called “grease ice”, which looks a bit like an oil slick. Grease ice develops into a continuous thin sheet of ice called “nilas”. Currents or light winds can push nilas around so that they slide over each other in a process called rafting. This then thickens into a more stable sheet with a smooth bottom surface, known as congealation ice. Frazil crystals can’t form under congealation ice which also insulates the water from the air a bit slowing down the process. Therefore, the ice crystals in congealation ice are long and vertical. On the other hand, in rough waters, frazil crystals form slushy circular disks called pancake ice, which form raised edges or ridges on their perimeter from bumping into one another. These can also raft together and form ridges, which have keels. In the Arctic, ridges up to 20m thick can form when thick ice deforms. Unlike ice formed in calm waters, sheet ice formed from the agglomeration of pancake ice has a rough bottom surface. Note the importance that the wind plays in this process.

(iii) Once sheet ice is formed it continues to grow through the winter. If it does not reach a sufficient thickness, it will completely melt during the summer. However, it can also be thick enough that it never melts all the way through, allowing the formation of multi-year ice. Eventually, sea ice gets so thick that it completely insulates the water below from the cold air – at this thickness (typically 3m in Arctic waters) a thermodynamic steady state is reached.

---

[1] The freezing of river ice brings its own complications, which are outside the scope of this case study.

(iv) The large scale dynamics of sea ice sheets allows sheets to fold over one another in much the same way as geological processes fold layers of sedimentary rock. For this reason, sea ice much thicker than 3m can be found.

(v) There is a great deal of interesting science around the question of what happens to the salt expelled from the frazils – briefly, this salt remains in first year ice but is much less prevalent in multi-year ice. This means that first year and multi-year ice have different electromagnetic properties that can be detected by satellite sensors.

(vi) When water freezes it expels latent heat which must be absorbed elsewhere in the system. This latent heat can result in the formation of so-called polynas, or holes in the sea ice.

Sea ice is neither continuous nor uniformly smooth, but rather a complex surface that varies a great deal over short distances. This is due to wind and ocean currents, which push sea ice around to form ridges and keels. Snow can also have an impact by pushing ice below the surface of the water creating flooded sea ice. However, this is more common in Antarctic than in Arctic waters.

2.3. Melting of Sea ice

Melting is, if anything, more complicated than the formation of sea ice, in part due to the extreme inhomegeneity of the sea ice. In the melting process, radiative heat transfer between incoming solar radiation and the ice plays a key role, in addition to the conductive and convective heat transfer mechanisms already identified in the formation of sea ice. This solar radiation is responsible for the hard-to-model positive feedback effect in which the areas first to melt, "melt ponds", reflect less sunlight than the surrounding ice, thereby melting more rapidly than their surroundings. Finally, snow cover on sea ice both insulates it from warm air and increases its local albedo (reflecting more incoming radiation), protecting it from melting.

2.4. Gaps in Sea Ice

Of great importance for navigation in ice filled polar waters are gaps in sea ice. These may be divided into two categories: leads and polynas.

Leads are narrow linear cracks in the ice caused by solid mechanics and fracture of ice sheets, similar to those observed in plate tectonics. Leads vary from a few meters to over a kilometer wide. As with melt ponds, leads have lower albedo than the surrounding ice and so melt faster. Leads can in turn influence local weather patterns as the warm ocean water is exposed to the cold winter atmosphere, leading to evaporation and the formation of low-level clouds downwind. Leads are clearly important for navigation, because even when they freeze, they tend to contain thinner and weaker ice more easily traversed by icebreakers.

Polynas are oval shaped areas of water surrounded by sea ice. Polynas may be further subdivided into two categories: Sensible Heat (Open Ocean) Polynas, in which the hole is created by the upwelling of a warm water ocean current, and Latent Heat (Coastal)

Polynas, in which water, warmed by the released latent heat of freezing, is forced to the surface when it meets the shoreline.

2.5. Dynamics of Sea Ice

As mentioned above, sea ice has an equilibrium thickness of about 3m in the Arctic (it tends to be thinner in the Antarctic). However, dynamic processes can make it much thicker, on the order of 10m. In decreasing order of importance, the five dynamic forces acting on sea ice are: wind, ocean currents, Coriolis force, internal ice stress, and sea surface tilt.

As a rule of thumb, sea ice moves at 2% of the wind speed which drags it. However, ocean currents typically act in the opposite direction of the wind speed. The Coriolis force, which is caused by the spinning of the planet, is weak but consistent, and is responsible for the large scale, long duration motion of large bodies of sea ice. Internal ice stresses represent the structural forces acting on sea ice, while sea surface tilt is caused by the fact that different parts of the ocean are at a higher gravitational potential than others at a given time.

Both drifting sea ice (also called pack ice) and icebergs can be a threat to offshore oil platforms, as described later in this case study. As a result, predicting the trajectory of the two types of floating ice represents an important challenge. Both types of floating ice are primarily influenced by the wind and secondarily influenced by currents. However, the effect of wind is even more pronounced on icebergs, which emerge a fair distance from the ocean's surface, than on sea ice. Accurate prediction of floating ice trajectories requires three interacting elements: an accurate wind prediction model, an accurate ocean current model, and a detailed understanding of how wind and current forces impact a given piece of floating ice. Each of these three components represents a major challenge to Environmental Prediction.

According to Malcolm Dewhurst, his employer Oceans Ltd., based in St. John's, Newfoundland, provides customized weather forecasts to the offshore oil and gas industry, and the company can provide spot wind speed forecasts on a 3-hour frequency with a 90-hour forecasting horizon. For the first 24 hours, these forecasts are 80-90% accurate, with a slight decrease in accuracy for the remaining time interval.

Predicting ocean currents is a very difficult task. Off-the-shelf software such as the PALM system (Dewhurst 2007; C-CORE staff 2007) incorporates general principles but must be adapted to local bathymetry. Work is only now beginning on detailed integrated wind-current models for the East coast of Canada.

Even with detailed wind forecasts, the impact of wind and ocean currents on sea ice is a complicated dynamics problem, especially given the irregular and idiosyncratic shape of floating ice. As a result of this complexity, accurate long range floating ice trajectory forecasts remain an emerging technology and require extensive future work.

2.6. Big Picture Characteristics of Sea Ice

At winter's end, sea ice typically covers 14-16 million km^2 in the Arctic and 17-20 million km2 in the Antarctic. By the late summer only 3-4 million km^2 remain in the Antarctic while 7-9 million km^2 remain in the Arctic.

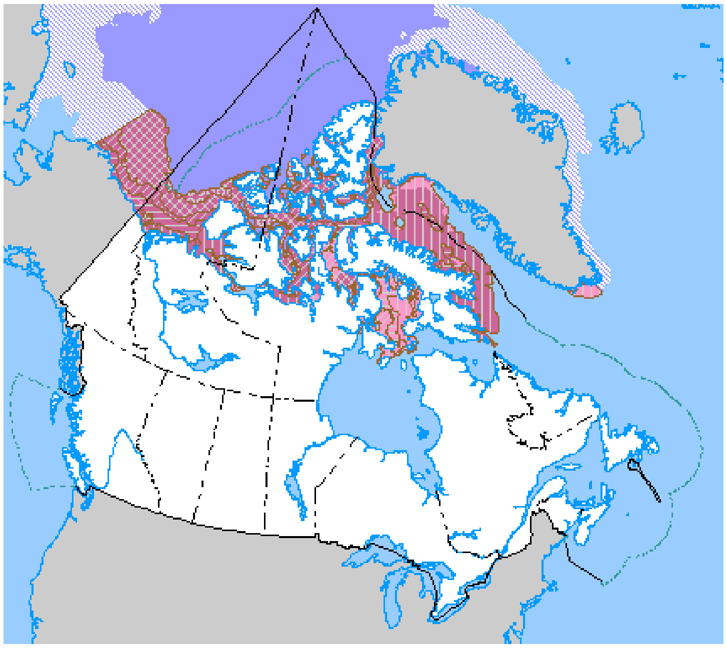

Figure 1:Sea Ice - Late Summer extent
http://atlas.nrcan.gc.ca/site/english/maps/environment/seaice/summertypes

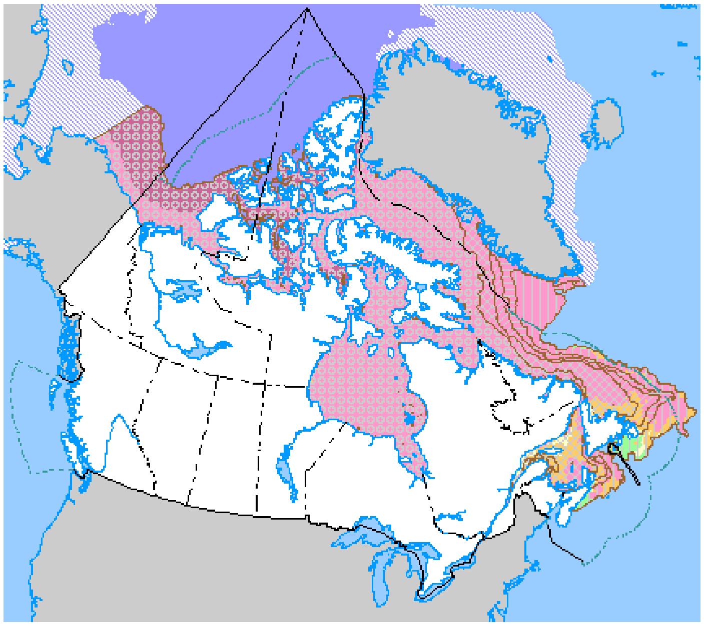

Figure 2: Sea Ice- Late Winter extent
http://atlas.nrcan.gc.ca/site/english/maps/environment/seaice/wintertypes

Like many other quantitative measurements of sea ice extent, these figures may only be obtained using satellite imagery, which has been available since 1972. Analysis of this imagery reveals substantial month-to-month fluctuation in sea ice levels. Additionally, there are well known climate variations on decadal time scales, the North Atlantic Oscillation and the Arctic Oscillation, further complicating the picture.

Nonetheless, analysis of satellite-based sea ice date shows that, over the last 35 years, Arctic sea ice cover has been decreasing at an average rate of about 3% per decade while Antarctic ice has been increasing at about 0.8% per decade. It is also worth noting that between 2002 and 2004, the activity level of sea ice was particularly low.

However, a year of low overall sea ice extent does not necessarily translate into a year of low danger for the energy industry. Pack ice, sea ice that has separated from the main region of sea ice, can drift into the oil producing regions of the North Atlantic, posing an iceberg-like threat to oil platforms. Indeed, in 2003 sea ice would have caused problems for the White Rose drilling platform, had the platform been in position at that time (personal communication Freeman Ralph, June 2007).

2.7. Satellites and Sea Ice Data

As mentioned above, Earth Observation satellites play a very important role in the observation of sea ice. The Canadian Ice Service, a division of Environment Canada, is particularly reliant on two satellite platforms: Radarsat and Envisat. Radarsat is a satellite designed and manufactured in Canada under the supervision of the Canadian Space Agency. Envisat is a European built satellite financed and operated by the European Space Agency. These satellites work by observing the radiation emitted by the Earth at various points in the spectrum, and by using these observations to make inferences about the presence or absence of sea ice and, indeed, its type and thickness.

Satellite sensors observing radiation in the infrared spectrum can characterize the temperature of a region. This allows relatively cold sea ice to be distinguished from the relatively warm surrounding ocean. But it is difficult to differentiate melting sea ice from the surrounding ocean, and cloud cover can limit the ability of sensors to detect infrared signals.

For these reasons observations in the microwave spectrum (so-called "passive microwave") are also very valuable – clouds do not emit much microwave radiation and so do not obscure this picture. Also, the physical properties (atomic composition, crystalline structure) of an object determine the amount and characteristics of the microwave radiation it emits, allowing crystalline sea ice to be distinguished from liquid sea water. However, the energy emitted in the microwave part of the spectrum is quite low, forcing it to be collected over a large spatial area. This makes identifying smaller structures, such as leads and polynas, difficult using this modality.

Satellites are not limited to simply observing the radiation emitted from the Earth – they can also beam microwaves at a location on the Earth's surface and see what is reflected, using radar principles. Active microwave radars such as synthetic aperture radar can be used to measure the thickness of sea ice and to determine (by porosity and salt content), new ice from multi-year ice. The Canadian Space Agency's RADARSAT is a synthetic aperture radar tool.

Significant sea ice data is available at the National Snow and Ice Data Centre (US) http://nsidc.org/.

3. Canadian Upstream Oil and Gas Operations Impacted by Sea Ice

The depletion of oil and gas reserves from accessible and politically stable parts of the world has directed the exploration efforts to more challenging environments such as the Arctic and the North Atlantic. Exploration for oil and gas in the Arctic will be aided by a newly developed technique of sea ice surface seismic analysis (Speece et al. 2006) in which seismic trucks are used on the surface of thick sea ice to probe the Earth's subsurface for hydrocarbon rich zones. While this is a somewhat new development, oil extraction from high latitude ocean regions at risk from drifting sea ice is not. Therefore, this case study focuses on the impact of sea ice on the oil and gas operations off the East coast of Canada. However, there are interesting emerging applications of sea ice EP relevant to oil operations in the Prudhoe Bay region of Alaska and Canada's Mackenzie Delta/Beaufort Sea region.

In particular, a problem of current interest to oil companies is determining how long into the fall a drilling vessel can stay in the Beaufort before the formation of sea ice precludes egress from the area (interview with Dr. John Yackel 2007). According to Dr. Yackel, sea ice formation can be predicted up to about a week in advance, but not with a very high level of skill. For a risk averse operator, the accuracy of such forecasts needs to be greatly improved before the resulting predictive tools are of great practical utility.

The Jeanne D'Arc basin is an oil rich geological structure located below the Grand Banks, East of Newfoundland. Oil is currently being produced from three locations, in a region roughly 350km east of St John's (Nfld). These locations, 30-50km distant from one another are the Hibernia Oil Platform, the Terra Nova and White Rose sites.

Figure 3: The Hibernia GBS being towed from its construction site at Bull Arm, Nfld.
Source: http://www.offshore-technology.com/projects/hibernia/

These three sites produce approximately 314,000, 150,000 and 100,000 barrels per day of light crude oil, respectively. All three locations are at risk not only from icebergs but also from floating sea ice at thicknesses of 1-2m. It is interesting to note that different oil platform technology has been selected for these sites. The Hibernia Oil Platform is a fixed "gravity based structure" (GBS) which is designed to be nearly invulnerable to the encroachment of sea ice or icebergs. In contrast, the extraction at the Terra Nova and White Rose sites is accomplished by means of a floating production, storage, and offloading vessel (FPSO).

Figure 4:The floating production, storage, and offloading vessel for Terra Nova.
Source: http://www.offshore-technology.com/projects/terra_nova/

A wealth of engineering and geological detail about these fields is presented in Appendix 1.

4. Ice Threats

Offshore oil and gas production facilities like the ones described above are susceptible to a variety of environmental dangers including high winds, high seas, and the threat of ice. Icebergs represent a particularly serious threat, since the Grand Banks locations described above are in relatively shallow (approx. 80m) water and a large iceberg can scour the ocean floor in the vicinity. However, floating sea ice also represents a significant threat to these facilities.

The two design approaches, reinforced gravity based structure and FPSO, are very interesting in their approach to ice threats. Hibernia, which cannot be moved, is designed to be invulnerable to ice threats. According to Freeman Ralph (personal communication, June 2007), this invulnerability comes at the heavy financial cost of $1 billion for the additional ice reinforcement. In contrast, the FPSO approach of Terra Nova and White Rose allows the vessel to be disengaged in response to a severe ice threat. In the latter case, the timing of the decision to disengage is critical. If the vessel is moved in response to every threat, production will be delayed at a high cost. On the other hand, if the vessel never leaves, it will be at risk to damage from ice.

Coupled with this set of decision processes is a complicated threat management environment which allows sea ice and icebergs to be diverted from the structure using icebreakers and even water cannons. Environmental predictions of the trajectory of floating sea ice and icebergs are an integral part not only of the threat management system but also of the evacuation system. In the next section, this threat analysis system is described, paying particular attention to the importance of EP.

5. Ice Threat Analysis for Oil Platforms: An EP Perspective

In order to get some insight into the challenges of incorporating EP into sea ice/iceberg management for offshore operations, we conducted interviews with six people or groups. Following initial discussions with Dr. Shawn Marshall (Canada Research Chair in glaciology at the University of Calgary) and Dr. Andy Foster (a professor of Applied Mathematics at Memorial University in Newfoundland), additional experts were contacted. Useful “grassroots” information was obtained from Mr. Barry Kirby, a mariner on the Newfoundland Greenland run. More extensive interviews were conducted with Dr. John Yackel, a sea ice expert at the University of Calgary; with iceberg-expert staff at the Centre for Cold Oceans Research and Environment (C-CORE) in St John’s and with Mr. Malcolm Dewhurst, a meteorologist at St John’s – based Oceans Ltd. The valuation framework proposed in this case study was validated through discussions with the C-CORE staff (however, all responsibility for the implementation rests with the authors).

A closer look at the ice management protocol used at Hibernia can give a good idea about the current practices of risk management and mitigation in offshore operations.

For more information, please see Hibernia’s website:
http://www.hibernia.ca/html/about_hibernia/index.html

As discussed in Section 3, the Hibernia platform is designed to survive the impact of sea ice and icebergs. It is estimated that it can withstand the impact of a one-million tonne iceberg with no damage. It is further estimated that it can withstand contact with a six million tonne iceberg (largest that can drift into that water depth) with repairable damage, an occurrence only expected once in 10,000 years.

Given that the Hibernia platform is located in relatively shallow water (80 metres deep), the odds of a large iceberg ever hitting the platform are extremely low. Nevertheless, an Ice Management Strategy is in place to actively manage this risk.

The operators of Hibernia obtain information about approaching icebergs through a variety of means, including the following:

- airborne surveillance briefings of the Canadian Ice Service (Environment Canada) and the International Ice Patrol (US Coast Guard)
- satellite and radar technology, including Hibernia's own state-of-the-art platform radar system, which can identify approaching icebergs up to 18 nautical miles away
- helicopters which pinpoint an iceberg’s position using radio signals
- onboard technology on support vessels that collect ocean current information as they steam toward the iceberg and transmit it back to St. John's via satellite
- side scan sonar, from vessels that travel alongside the iceberg and record a detailed profile to measure its draught.

Hibernia operators also use the services of Provincial Airlines of St. John's, which provides services such as weather forecasting, ice management and physical and environmental data management. Data collected using different means is entered into Provincial Airlines’ ice management trajectory modelling program. This program includes certain EP characteristics: using complex mathematical modelling systems, and combining wind and wave elements of the weather forecast with the physical and environmental information gathered in the field, it can predict iceberg movement and identify those which may drift close to the production area.

5.1. Risk Mitigation

Hibernia's approach to ice management is two-fold. The first line of defence comprises active countermeasures deployed to prevent iceberg collision. If these fail, both the platform and the Offshore Loading System (OLS) are designed to withstand such a collision.

Hibernia operators report that icebergs representing a threat are tackled proactively, while they are still 20 km or more away from the platform. The platform support vessels encircle the iceberg with a long cable or rope - much like a giant lasso - and tow the iceberg into a different trajectory. It is not necessary to tow the iceberg very far, as even a slight nudge to an iceberg at that distance will change its course considerably over a 20 km drift. Water cannon or propeller wash of the support vessel can be used to deflect smaller pieces of ice.

Close encounters with icebergs could, however, force the platform to stop production and actual contact may require repairs afterward. As well, any bottom-scouring iceberg could potentially cause damage to the platform's OLS, a network of oil transmission pipelines on the ocean floor. For this reason, the OLS pipeline has been encased in concrete for additional protection. A redundant OLS system is in place to serve as an auxiliary, in the unlikely event that the other system is damaged.

5.2. Modelling the Decision Problem: Operational/Tactical

One possible way to frame the decision problem is discussed below. The operational/tactical decisions associated with threat management and evacuation are summarized in the decision tree presented in Appendix 2.

Assuming that the platform is surrounded by three concentric circles, which denote various threat levels, one can break down the decision process as follows. Each successive breach by an iceberg reduces the time remaining to the decision maker and forces an intermediate shutdown so that, if necessary, there is enough time to evacuate.

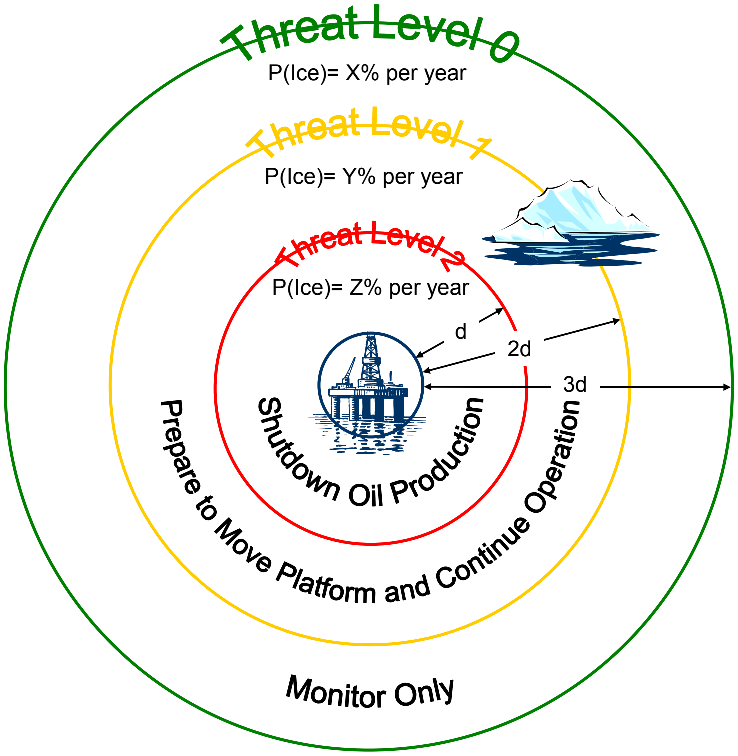

Figure 5: Threat Levels

Once the sea ice or iceberg starts to pose a serious threat, the operator can resort to active countermeasures. Support vessels can try to tow away an iceberg, or even deflect it with high pressure water hoses. If it is approaching pack ice they can break it up with an icebreaker. In both of these cases, EP plays a critical role: the operator wants to make sure that the actions actually move the threat from a trajectory ending in a hit to a trajectory ending in a miss. Moving it from a trajectory ending in a miss to another trajectory ending in a miss is a waste of time and money; moving it from a trajectory ending in a miss to one ending in a hit is a cruel twist of fate. The same reasoning applies to the pack ice threat: the operator wants to break up the part of it that will actually impact the structure.

5.2.1. Tactical Analysis of Oil Rig – in the absence of any EP

The setting here is that a mobile oil platform (FPSO) sits at the centre of a nested set of concentric circles. We assume that the FPSO services an oil well (roughly modelled on the White Rose well in the Jeanne D'Arc Basin operating at maximum capacity) with a daily production of 100,000 barrels of light crude, and a production lifetime of 10 years at this production level. We assume that the production is sold on the international spot oil market. The circles are defined so that an ice threat (sea ice or iceberg) arriving at the outer perimeter of a given concentric circle will take a given time to hit the oil platform, if it hits it at all.

Since floating ice rarely provides a threat to human life (Kirby 2007; C-CORE 2007) we feel justified in assuming an expected value framework for this set of decisions.

In the absence of any EP information, we assume that the probability of ice breaching a given area is proportional to that area. Additional information, about the expected direction from which ice threats come for instance, will change the scaling relationship here (if, for example, ice always comes from due north, the probability of ice breaching a given circular region is proportional to its diameter).

In this study, we further assume that the outer circle has a radius of three times the inner circle, and the middle circle has a radius twice that of the inner circle, which means that the probability of ice breaching a given circle must follow in the progression 9:4:1 (i.e. the outer circle is 9 times more likely to be breached than the inner most circle and the middle circle is 4 times more likely to be breached). At this time, we do not yet have any ability to quantify these probabilities beyond this. However, for the purpose of this analysis, we are working with the assumption that, during a 30-day ice season, the outer circle is breached by ice 9% of days (therefore, the inner two circles are breached 4% and 1% of these days respectively). Countermeasures involving serious economic loss need only be initiated as the second circle is breached, so this corresponds to an assumption of approximately once per year of second circle events and approximately once every 3 years of extremely serious, inner circle breach events.

A rigid threat protocol designed to safeguard both the oil platform workers and the natural environment is instituted so that when the outer circle is breached by ice, the steps of a shutdown and evacuation sequence are initiated. The initial steps of this sequence are easily reversible and involve neither any significant costs nor any disruption to the flow of oil. As the ice breaches the successive circles, the threat level increases until eventually oil production is shut down in preparation for moving away from the threat.

To be more specific, we consider a three-level threat model. When ice is outside the outer concentric circle, a threat level of zero is considered to exist and no action other than the normal business of the oil platform need be initiated. When ice breaches the outermost circle, the FPSO makes ready for sea, incurring a modest cost in fuel burned. If the ice leaves the outermost circle again, this step can be reversed.

If, on the other hand, ice continues to breach the second circle, the FPSO must shut down oil production. This comes at a much heavier cost. Even if we ignore any petroleum engineering consequences of cycling production in this way, the result is still to defer production of oil from the present, adding it back on at the end of the reservoir's life.

This incurs a known cost of 100,000 x today's oil price, with a countervailing benefit of 100,000 x the oil price 10 years from now, discounted to the present day at an appropriate interest rate (spread over the London Interbank Offered Rate, LIBOR). If we assume that today's oil price is US$70 per barrel, discounted at 6%, and that the value of oil grows simply with a 2% inflation rate, the cost of deferring a day's production is about $2.25 million.

It should be noted that this result is very sensitive to the underlying assumptions regarding the oil prices. If the prices declined to late 1990s levels, the valuations would be greatly reduced. The readers of this report are encouraged to use their own assumptions in the companion spreadsheet. We should also note that, if oil prices continue increasing, deferring production 10 years might actually lead to a huge profit. However, many firms cannot afford to wait since they actively need to manage their cash flows as they have to pay debts and salaries to remain in business. Finally, government royalties, calculated on a per-barrel basis, are also exposed to risk and a lengthy outage in oil production might, in a tightly balanced oil market, lead to higher consumer prices at the gas pumps.

We assume that the ice leaves the middle circle after a single day (either to breach the inner circle or to retreat to the outer circle), and that if the ice retreats to the outer circle, production can be restarted immediately.

If the ice breaches the innermost circle, the FPSO needs to disconnect from the wellhead and steam away. The cost of this is estimated to be an entire week of lost production – to cover the time required to steam to safety, return, and perform the delicate operation of re-connecting to the wellhead (note that the wellhead is installed in a pit in the sea floor, the "glory hole", to safeguard it against the effect of iceberg bottom scour).

At each stage countermeasures can be adopted. These include a mobile escort vessel, with tug and at least some icebreaking capabilities, and possibly armed with high-pressure water cannons. According to Freeman Ralph (personal communication, June 2007), such vessels cost $20,000 per day plus fuel burned (we assume a total of

\$25,000 per day)[2]. We can assume that the vessel is employed to either redirect icebergs out of the threat area, using either the water cannons or, for sufficiently small icebergs, by towing, or to break up floating sea ice so that it can pass by the oil platform without damaging it.

We assume that employing such a vessel (without using EP) can clear the threat a fraction E of the time at every stage. Thus the probability that ice penetrates the outer circle is (1-E), the possibility that the ice breaches the middle circle is (1-E)*(1-E), and the probability that ice breaches the inner circle is $(1\text{-}E)^3$. Employing such a vessel will be worthwhile on a cost basis if its use for a single day reduces the probability of a middle circle breach lasting one day by even 1% (1% x 2.25 million = \$30,000 >\$ 25,000). However, the vessel cannot come on a day-by-day basis. It must be engaged for some minimum period of say 1 month, with a cost then mounting to \$750,000. However, there are three oil platforms (Hibernia, White Rose, Terra Nova) within a few hours steaming of one another, so perhaps the three platforms could share such a vessel. Nonetheless, the cost of engaging it may not be worthwhile if the probability of an ice-related shutdown is rather small.

5.2.2. Tactical Analysis of Oil Rig – introducing EP

This section illustrates the benefits of EP. The first thing that EP can do for us is to alter the threat regions. For instance, suppose that we have a reliable forecast that strong winds from the East will persist for the next 3 hours. Since the primary determinant of sea ice motion is the wind, this allows us to simply monitor pack ice which enters the extreme west of our outer circle without even starting the first level threat response of making ready for sea. A similar argument could be used to delay the stoppage of production when sea ice breaches the west end of the middle circle. In essence, good EP allows the threat regions to be shrunk, reducing the probability of having to engage in costly countermeasures. The better the EP, the more comfortable we can be with nearby ice. We should also note that EP will have a dynamically varying effect on the size of threat regions as weather conditions change, however we abstract that aspect of the problem away in this analysis.

For the purposes of the case study, we assume that so called "passive" EP decreases the area of each threat circle, and hence the probability of ice entering the circles, by the same factor q, $0 < q < 1$. Thus, rather than an 9%/4%/1% per day set of probabilities, good use of EP allows a tighter set of 9q%/4q%/q% probabilities.

Good EP can also increase the probability of successful countermeasure operations. With a good prediction of where natural winds and currents will take a given ice threat, a small "nudge" might suffice to place the threat on an oil platform avoiding trajectory. We assume that EP increases the ability of the escort vessel to avert the ice threat from a fraction E to a fraction QE, where $Q > 1$ but QE remains < 1.

---

[2] Such services are sometimes contracted on a standby basis resulting in lower costs than presented here. However, even with the vessel cost figures used here, countermeasures are still worthwhile.

5.2.3. Quantitative Results

Now we present some quantitative results. To replicate the calculations made here, and for some additional commentary, please see the spreadsheet Sea_Ice_Worksheet.xls which accompanies this case study.

Case 1a (Base Case): This base case computes the expected annual cost of ice to an offshore platform similar to the White Rose without any use of countermeasures such as escort vessels nor any use of Environmental Prediction. This provides a baseline against which the other strategies can be judged.

4% of the time the middle circle is breached. A middle circle breach requires a one day production shutdown at a cost of one day's deferred production, estimated to be $2.25 million. There are 30 ice days in a year so the total cost of this middle circle breach is 30*0.04*2.25 = 2.7 million. Conditioned on the middle circle being breached, ¼ of the time the inner circle is breached next. An inner circle breach requires a shutdown for 7 days (after which it is assumed that the ice has left the outermost threat circle) at a cost of 7 days of lost production. Again, there are 30 ice days in a year so the cost of an inner circle breach is 30*0.01*7*2.25 = $4.725 million. So the total expected annual cost of ice, with no countermeasures and no use of Environmental Prediction, is 2.7 + 4.725 = $7.4 million.

Case 1b: In this case, no Environmental Prediction is used, although an escort vessel is employed for the 30 days. The cost of the escort vessel is shared between 3 platforms at a cost of 1/3*(30 days)*$25,000 per day = $0.25 million. We assume that 10% of the time, the escort vessel is able to stop a floating ice threat from moving from one threat level to the next higher level.

When the escort vessel is present, the middle circle is breached only 4%*$(0.9)^2$ = 3.24% = of the time, and the inner circle is breached just $(0.9)^3$x1% = 0.729% of the time. So the expected loss due to ice decreases to 30*0.0324*$2.25 + 30*0.00729*7*2.25 = $5.6 million. But the escort vessel costs $250,000, so the total expected loss with escort vessel is $5.85 million. This represents a savings of 7.4 – 5.85= $1.55 million, suggesting that with the admittedly very crude assumptions employed here, an escort vessel is well worth the money.

Now we add passive EP, and assume that the q factor is 0.9. In other words, passive EP allows the threat areas to be shrunk by 10%.

Case 2a: EP, no escort vessel.

In this case, the middle circle is breached just (0.9)*4% = 3.6% of the time and the inner circle is breached (0.9)*1% = 0.9% of the time. This simply scales the cost of ice events down by 10% for a total estimated savings of 0.1*$7.4 million = $740,000.

Case 2b: EP and escort vessel.

With active EP, the ability of the escort vessel to avert threats is scaled up by an assumed factor of Q =2. At the same time, the passive EP decreases the threat regions by q = 0.9. In this case, an escort vessel must still be retained at a cost of $250,000. The net effect of these assumptions is to save $3.98 million, a substantial improvement from the $1.55 million saved by the escort vessel alone or of the $740,000 saved by passive EP alone. We can therefore attribute a significant savings of $3.98 million - $1.55 million = $2.4 million to the incremental benefit of EP in this context.

While every effort was made to use indicative values for these parameters, these parameters represent estimates rather than hard-edged numbers. Their value lies primarily in their ability to stimulate modelling and critical thought. For example, if q is very small the escort vessel stops being worthwhile, even if it does a really great job, because the EP is already so good at assessing threats that their expected cost is low.

5.3. Modelling the Decision Problem: Strategic

The decision of whether to adopt the ice reinforced gravity base structure or the FPSO is a strategic one. Indeed, a third choice, to adopt a non ice-reinforced structure, might also be appropriate in warmer seas, or perhaps even in cold seas with some future level of ability to incorporate advanced EP with advanced ice control countermeasures.

The strategic decisions are summarized in the decision tree in Appendix 3.

Here, the decision maker has three main design choices: oil platform minimally adapted to high latitudes, fixed oil platform (reinforced) and FPSO

(i) Oil platform only minimally adapted to high latitudes, this is a fixed and not particularly strong structure, the kind used on the North Sea or the Gulf of Mexico. This design would be relatively cost-efficient but could be seriously damaged by an iceberg or by pack ice. If the decision maker had a complete mastery of EP relating to ice, she might opt for this and use detailed knowledge of the system to implement the necessary countermeasures.

(ii) Fixed Oil platform heavily reinforced to be able to take a direct hit from an iceberg and from sea ice (Hibernia type design as discussed above). The extra cost is substantial (on the order of an extra billion dollars according to C-CORE experts), but the operator basically doesn't need to consider EP since the structure is very robust to withstand impact.

(iii) FPSO, a system with a lot of mobility (so that it can "get the hull out of there" when necessary). This is presumably expensive compared to the first design type, however since the vessel can be re-used after the oilfield is depleted there can be a “resale” value. Operating this design also requires a very good use of EP, basically giving the operator the confidence to sit put when ice breaches the outer ring. Using this design requires being fairly confident in the evacuation decisions, if the risk tolerance is too low, the vessel has to leave too frequently (and waste production), if it is too high, the vessel can end up stationary when it actually needs to move and suffer substantial damages.

It is interesting to note that if there are no EP capabilities, the second choice seems to be the only option. Although, even then there might be a legitimate choice between 2 and 3 depending on how far south the operations are. If you are located such that significant ice rarely breaches the outer ring it might be cheaper to go with 3 and move it at the slightest hint of danger than it would be to go with 2. As your location gets icier the balance switches from 3 to 2, but as your knowledge of EP improves, the balance switches back from 2 to 3. As your EP capabilities and countermeasure skills improve (perhaps beyond current state of the art), you might be able to push the region for which the first design is an option into icier and icier regions.

Another interesting note is that the optimal path through the decision tree will depend on two parameters: ice danger (itself requiring some EP about the impact of global warming on ice) and operational EP/threat response capability.

5.3.1. First steps to Strategic EP Valuation

A top-down approach can be used to obtain some figures related to the strategic value of EP in this setting. The cost of ice engineering the Hibernia platform was approximately $1 billion. The corresponding FPSO solution can be obtained for about $200 million. However, Hibernia services a 300,000 barrel per day well about three times the size of the wells serviced by FPSO vessels. Therefore we can suppose that Hibernia has an additional cost of about $400 million. This additional cost brings a corresponding benefit of not having to worry about ice any longer.

From this step, we can impute a value of ice risk reduction[3]. Assuming that this capital cost is amortized over a 20 year lifetime, this represents an additional capital cost of about $30 million per year more than the cost of the White Rose structure (or about $10 million per year for a 100,000 barrel per day well). The above operational/tactical analysis suggests that the ice risk for a 100,000 barrel per day well is about $5-8 million, depending on the sophistication of the countermeasures adopted. Considering that the Hibernia solution guards against other risks in addition to the ice risk, these two numbers seem to be broadly consistent.

Note that, using the FPSO would not be an option if a threat response system like the one described above were not available. Therefore, this threat analysis depends on the availability of EP, and implies a multi-million dollar per year benefit attributable to EP.

6. References

Speece M., Betterly, S., Levy, R., Harwood, D., Pekar, S., Winter, D., Lutz, M., and J. Doren , "A new over sea ice active-source seismic method to support ANDRILL", Geophysical Research Abstracts 8, 013713, 2006

http://www.offshore-technology.com/projects/white_rose/

http://www.offshore-technology.com/projects/terra_nova/

---

[3] Presumably, the robust construction of Hibernia also supplies its operator with assurances against other forms of economic and safety risks.

http://www.offshoreoilandgas.gov.bc.ca/world-offshore-oil-and-gas/

http://www.offshore-technology.com/projects/hibernia/

http://www.oceans.nf.net/main.html

http://www.pr-ac.ca/index.html (note the recent RFP for study of moving large components through the northwest passage to the Mackenzie River delta/Beaufort region)

http://www.ainc-inac.gc.ca/oil/Pdf/beauesize.pdf

Various personal communications as noted in the text.

Appendix 1: Additional Information on Offshore Platforms

For more information regarding the information contained in Appendix 1, please see http://www.offshore-technology.com/

Appendix 1.1 Hibernia, Jeanne d'Arc Basin

Hibernia is located in the Jeanne d'Arc Basin, 315km east of St John's, Newfoundland, in water 80m deep. The field consists principally of two early Cretaceous reservoirs - Hibernia and Avalon - located at average depths of 3700m and 2400m, respectively. Hibernia oil is a light sweet crude, with a density of about 32° to 34° API and a sulphur content, by weight, of 0.4-0.6%. The field contains approximately three billion barrels of oil in-place, and recoverable reserves are estimated to be at around 615 million bbl.

It was decided that the Hibernia field would be developed using a special gravity-base structure, strong enough to withstand a collision with a one-million-tonne iceberg (expected to occur once every 500 years) and a direct hit from a six-million-tonne iceberg (expected just once every 10,000 years).

In September 1990, HMDC awarded the gravity base structure (GBS) contract design to Newfoundland Offshore Development Constructors (NODECO). The detailed design was subcontracted to Doris Development Canada (DDC).

The Hibernia's novel 450,000t gravity base structure design consists of a 105.5m concrete caisson, constructed using high-strength concrete reinforced with steel rods and pre-stressed tendons. The caisson is surrounded by an icewall, which consists of 16 concrete teeth. Structurally, the 1.4m-thick icewall is supported by a system of X and V walls, which transmit the loads to the interior tiewall. The X and V walls have a thickness varying from 0.7m to 0.9m and the tiewall has a thickness of 0.9m. Put together, these walls form the icebelt. The caisson is closed at the bottom and top by horizontal slabs and the base slab has a diameter of 108m. The upper top-surface slab is about 5m above sea level.

Inside the gravity structure are storage tanks for 1.3 million bbl of crude oil.
Four shafts run through the GBS from the base slab to support the topsides facilities: namely the utility shaft, the riser shaft and two drill shafts. Each of the shafts are 17m in diameter and extend to a total height of 111m.

The utility shaft houses the mechanical outfitting required to operate the GBS system. It includes pipework, heating and air-conditioning, and electrical controls. The two drill shafts each house 32 drill slots to accommodate the wells, which will reach depths of more than 3700m below sea level, down into the oil reservoirs.

The topsides have a design capacity of 23,900m^3/d (150,000 b/d), based on the 98 million m^3 (615 million barrel) estimate. The topside facilities consist of five super-modules (processing, wellhead, mud, utilities, and accommodation for 185 people) as well as seven topside mounted structures (helideck, flareboom, piperack, main and auxiliary lifeboat stations, and two drilling modules).

The wellhead module for Hibernia was fabricated at Bull Arm, while the remaining components were made in construction sites located around the world - two in Italy and the remaining two in South Korea. Four of the topside mounted structures (flareboom, helideck, main and auxiliary lifeboat stations) were also fabricated at Bull Arm. The other three topside mounted structures (components of the two drilling rigs and the piperack) were fabricated in Newfoundland and New Brunswick, with some of the components being built in Alberta.

The 37,000t integrated topsides facility was transported by barges to the Hibernia deepwater site and positioned above the partially submerged GBS shafts to form the completed 600,000t production platform. This was then towed to its final site and 450,000t of solid ballast was added to secure it in place.

Oil stored in the GBS will be exported by means of an offshore loading system (OLS) consisting of short subsea pipelines, a subsurface buoy and flexible loading hoses, feeding a purpose-built shuttle tanker.

Appendix 1.2 Terra Nova, Jeanne d'Arc Basin

The Terra Nova field, located 350km ESE of St John's Newfoundland and 35km SE of Hibernia, was discovered in 1984 by Petro-Canada. Field reserves have been estimated at 406 million barrels (Mbbl).

Water depths are shallow - between 90m and 100m. The mean annual wind speed is 35kmph, with the strongest recorded wind speed being 145kmph and the largest recorded wave height being 25m.
The area is characterised by the seasonal presence of floating sea ice, ranging in thickness from 0.5m to 1.5m, produced by the freezing of the ocean's surface layer and icebergs.

Terra Nova is subdivided into three major structural blocks: the Graben, the East Flank and the Far East. The field is estimated to contain over one billion barrels of oil in place, of which about 400Mbbl of oil are recoverable. (The Far East block, which is not yet drilled, is expected to add at least 100Mbbl of reserves to the 300Mbbl that have already been estimated within the Graben and East Flank). The estimated peak production rate is 125,000b/d from the Graben and East Flank portions alone. A total of 32 wells are planned for the Graben and East Flank blocks, including 20 production wells, ten water-injection wells and two gas injection wells. For the Far East, a total of 12 wells are planned, including six production wells and six injection wells. Field life is expected to be 18 years.
Petro-Canada selected the Grand Banks Alliance (SBR Offshore, Doris Conpro, PCL Industrial Constructors, Coflexip Stena, Halliburton Canada and FMC Canada) to carry out engineering, procurement, construction, installation, commissioning and possibly pre-development drilling activities up to the production of first oil. The project partners

and Grand Banks Alliance consequently established a single alliance: the Terra Nova Alliance, with each company participating on a risk-and-reward basis.

The subsea layout will consist of a production well feeding into a template, which, in turn, will be connected by flexible flowlines to a riser-base manifold (RBM). In order to protect the subsea wells from iceberg scour, they will be set in glory holes - large holes drilled in the seabed in which equipment can be installed.

Flexible risers will connect the RBM to an FPSO (floating production, storage, and offloading) vessel. The vessel will have a length of 280m and a width of 45m. The combination of low air and water temperatures with wind and wave action, makes superstructure icing a consideration during the winter months. This means that an allowance of several hundred tonnes of superstructure ice accumulation must be factored into weight and stability calculations. There must also be procedures for monitoring and mechanisms for controlling ice build-up on the structure and substructures of the offshore facilities.

Low water temperatures require that fluids such as hydraulic control fluids be heated or treated to lower their freezing point. Similarly, low temperatures combined with the waxy nature of the crudes require that the flowlines and riser are insulated to reduce wax deposition. The FPSO is designed to operate in moderate sea ice, up to a limit of five-tenths coverage and to disconnect, as required, to avoid heavy pack ice and potential collisions with icebergs.

The 9000t topsides facilities will be installed approximately 4.5m above the main deck. They will contain the necessary equipment to produce 150,000b/d oil, and inject 250,000bbl of seawater/day and 125 MMcf/d of gas. The FPSO hull will have an integrated storage capacity of 900,000bbl.

The export system will be a tandem offloading system for the transfer of crude oil from the storage tanks of the FPSO to ice-strengthened shuttle tankers, ranging in weight from 80,000t to 120,000t. The offloading system will be designed for connection to tankers in 5m significant-wave-height conditions.

Appendix 1.3 White Rose Oil and Gas Field, Jeanne d’Arc Basin

The White Rose field is located 350km east of Newfoundland, approximately 50km from both the Terra Nova and Hibernia fields. It is operated by Husky Oil (72.5%) on behalf of Petro-Canada (27.5%).

The area consists of several oil and gas pools in the Avalon formation sandstones, which were deposited during the early Cretaceous along a north-south trending shoreline roughly paralleling the eastern margin of the Jeanne d’Arc Basin. The White Rose oilfield development will involve recovering an estimated 36 million cubic metres (230 million barrels) of recoverable oil from an area of approximately 40km².

The first three wells - N-22, J-49 and L-61 - were drilled between 1984 and 1986. White Rose E-09 well was drilled in 1987-1988 into the South Avalon oil pool and encountered over 90m of net oil pay. In 1999 and 2000, three additional delineation wells were drilled into the South Avalon, White Rose L-08, A-17 and H-20. A fourth well, N-30, was also drilled in 1999 into the North Avalon pool, downdip from the N-22 well.

The field will be developed from three or four drill centres on the seafloor, with production and water and gas injection wells located at each centre. These drill centres will be located in excavated glory holes that lie below the seabed to protect the wells from iceberg scour.

Current plans envisage up to 10 to 14 production wells. The production from the combined wells is estimated between 12,000m³/day to 18,000m³/day (75,000b/d and 110,000b/d). An additional eight to eleven gas and water injection wells will be drilled for resource conservation and to maintain reservoir pressure. The wells will be drilled in phases over a four to six-year period. Up to four to six production wells, one to three water injection wells and one gas injection well will be required for first oil production.

The drill centres will be connected to a ship-shaped floating production, storage and offloading (FPSO) facility with flexible flowlines and risers. The FPSO's turret will be designed to allow the facility to disconnect from the subsea drill centres and move in the event of an emergency.

This FPSO will be able to store between 111,000m³ and 135,000m³ (700,000 and 850,000 barrels) of oil (approximately eight to ten days of oil production) and will contain topside processing units, accommodation and a turret.

In April 2002 Husky awarded Samsung Heavy Industries of South Korea the contract to build the FPSO hull. The design is derived from the proven purpose-built Grand Banks shuttle tanker design and features an ice-strengthened double hull. It will have dual conventional propulsion systems and two high efficiency rudders.

SBM IMODCO was given the contract for the turret and mooring system. The contract scope includes the engineering, procurement, and construction of a disconnectable turret. The mooring system connects the turret to the seabed and allows the FPSO to weathervane around the turret while connected.

The contract for the engineering, procurement, construction and installation of topsides has been awarded to Aker Maritime Kiewit Contractors (AMKC), a joint venture of Peter Kiewit Sons Co. Ltd. and Aker Oil and Gas Technology Ltd. The topsides are designed to produce oil at a quality suitable for shipment by shuttle tankers to market.

Husky time chartered two newbuild shuttle tankers from Knutsen OAS to transport oil from the White Rose offshore project. Delivery of the Suezmax-size vessels, each with a one million-barrel capacity, is planned for the second quarter of 2005. The vessels will be constructed by Samsung Heavy Industries in South Korea.

Appendix 2: Decision Tree (Operational/Tactical)

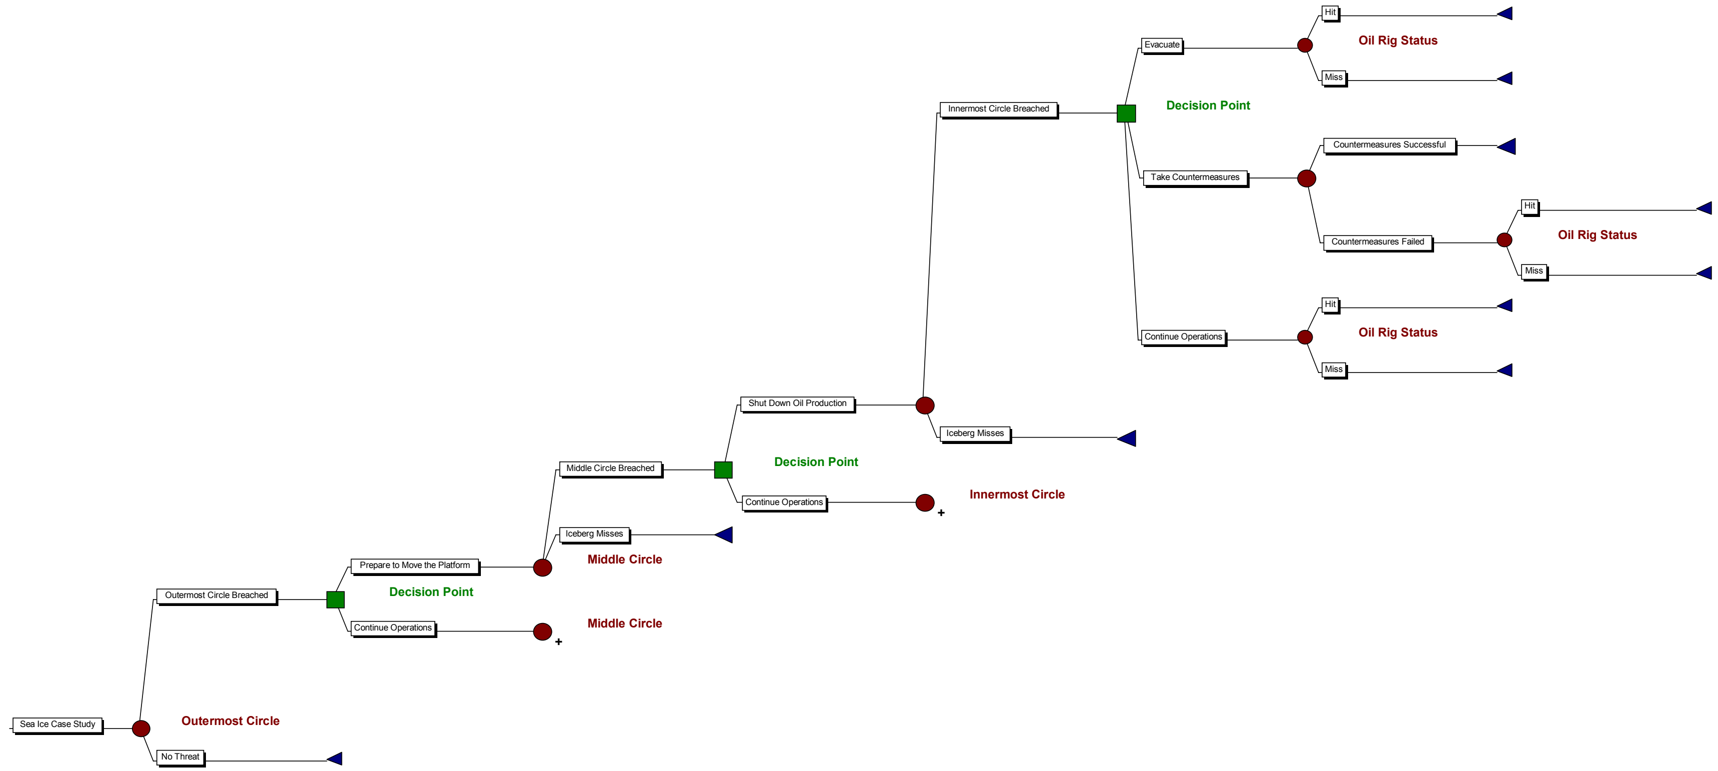

Note: Circles represent chance events, squares represent decision points, and the triangles represent the payoff.

Appendix 3: Decision Tree (Strategic)

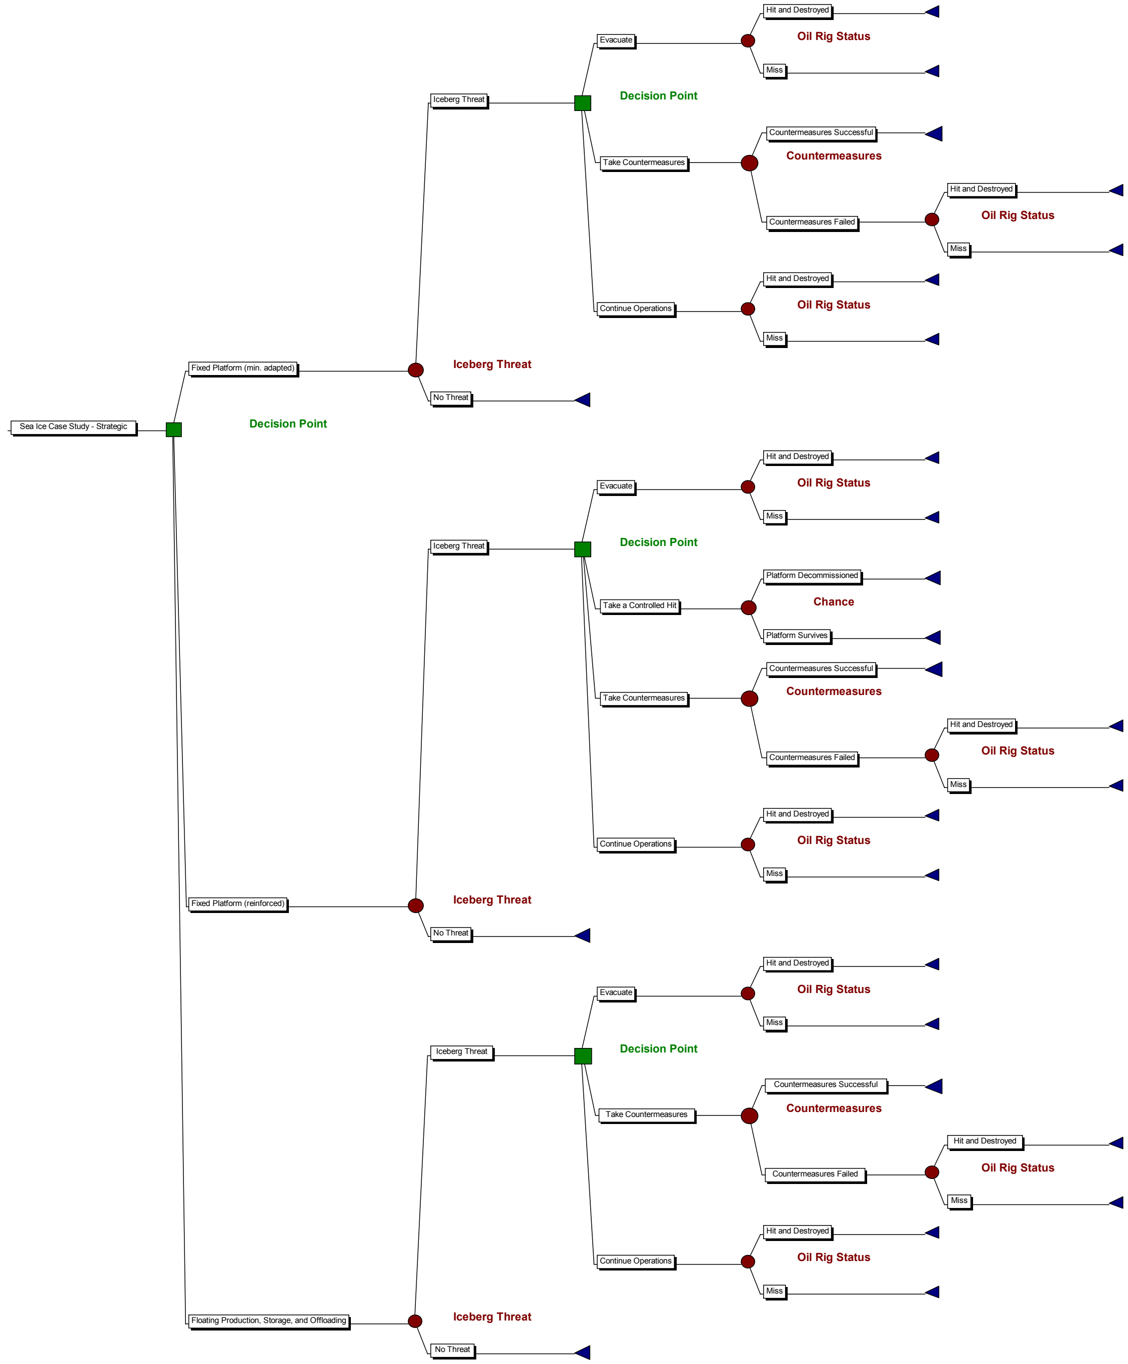

Environmental Predictions and the Energy Sector: A Canadian Perspective

Case Study 2:
Hydrological Environmental Prediction

Prepared for

Environment Canada

Contract Number: K3A40-06-0028

Prepared by

TURQUOISE
TECHNOLOGY SOLUTIONS INC.

DYDEX
DYDEX RESEARCH AND CAPITAL LTD.

1. Introduction

Worldwide, approximately 2890 TWh of hydroelectricity is generated each year, accounting for about 16.5% of the world's total electricity production (IEA, 2006). In Canada, the share of hydroelectricity is much higher, accounting for approximately 60% of the national total. Some Canadian provinces, most notably Québec, Manitoba, British Columbia and Newfoundland and Labrador have nearly all of their electricity needs provided by hydropower.

Hydroelectric plants provide electricity by converting the potential energy of water stored at a high elevation into kinetic energy of moving water. This kinetic energy is then used to move the blades of a turbine, thereby generating electrical energy in a very efficient way. The amount of energy generated is proportional to the amount of water flowing through the turbine and the height the water falls. The height the water falls, or the "head", is determined from the relative level of the water upstream of the dam and the turbine.

Both flow rates and water levels vary over time, and this variability has a significant impact on various strategic and operational decisions. Indeed, in the annual report of Hydro Québec, the largest risk factor is identified as the hydrological inflow risk, with other risk factors such as interest rate risk and foreign exchange risk having a smaller impact. For a very large utility trading extensively in the U.S. markets, this insight is particularly significant (Hydro Québec, 2006).

However, while significant sensitivity to an environmental variable is necessary for a valuable Environmental Prediction (EP), such a sensitivity is by no means sufficient. It is also necessary that information about future environmental events be actionable. In this case study, we attempt to quantify the potential value of sophisticated hydrological EP to hydroelectricity companies. We underline "potential" as a recurring theme throughout the interviews we conducted and the literature we reviewed as part of this study, since it signifies the curious absence of EP in current industry practices.

The structure of this case study is as follows. In Section 2, a brief review of the science of hydrology is given, focusing on hydrological predictability over a number of time scales corresponding to each of the operational, tactical, and strategic time scales introduced in the Literature Review Report[4]. In Section 3, a valuation framework is presented in order to estimate the benefits which can be extracted from EP, corresponding to the operational, tactical, and strategic time scales. This section describes various ways in which hydrological EP is valuable, and focuses on four of these in a detailed quantitative analysis within a Canadian setting. It should be noted that all benefit estimates for EP are based on certain assumptions, which are outlined throughout the report.

---

[4] In the hydrological prediction field, the general practice is to consider a few weeks as medium term and six months (or more) as long term. In this case, short term refers to a period of few days.

On the operational time scale (hours to days), we identify two different areas in which forecasts can add value: buffering fluctuations and the role of storage facilities. On the short end of the operational time scale (1-24 hours), accurate hydrological inflow predictions can buffer fluctuations in power demand (also called "load"). Such buffering ability has a high value as it allows minimization of so called "spinning reserve" with associated savings in wear and tear, fuel consumption, and emissions. We quantify this value using a framework originally developed by Hobbs et al. (1999). On the longer end of the operational time scale (24 to 72 hours), we present brand new research (Zhao and Davison 2007) quantifying the value of accurate hydrological inflows to the operator of a hydroelectric generation facility with water storage capacity. We note that there is a synergy between hydroelectric facilities with storage capacity and wind turbine power generation, a topic examined in detail in the Wind/Water Case Study.

At the cusp between the tactical and strategic time scales lies the month to year level predictability afforded by accurate snow-pack modelling and a proper understanding of El Nino/Southern Oscillation events. We follow the lead of Hamlet, Huppert and Lettenmaier (2002) in quantifying the potential value of such predictions to hydroelectric operators possessing the ability to store large quantities of water.

Section 3 closes with a discussion of the value of extremely long time scale hydrological predictions when very long term decisions need to be made, such as building or upgrading capital intensive hydroelectric projects with very long life cycles.

Section 4 focuses on the gap between the potential value of hydrological EP versus the limited use of EP in today's hydroelectricity industry and possible explanations for this gap, including economic, scientific, and cultural factors. The case study conclusions are reported in Section 5, along with some suggestions for follow up work.

2. Review of Hydrological Environmental Prediction

2.1. Hydrological Forecasting as Interdisciplinary EP

Hydrological forecasting is the science of predicting future water levels and flow rates within a given watershed. For instance, one might be interested in predicting the water levels in each of the five Great Lakes as well as the flow rates in the St. Mary's river, the Niagara River, and the St. Lawrence River. This type of forecasting application is perhaps the prototypical problem of holistic interdisciplinary prediction, uniting the need for understanding the processes within the atmosphere and the hydrosphere, as well as other natural and human-induced processes on the Earth. It is also an application in which multiple time scales are relevant.

Hydroelectricity is intimately related to the global water cycle, where outflow is in the form of stream flow and evapotranspiration (a combination of evaporation from open bodies of water, evaporation from soil surfaces, and transpiration from the soil by plants). This water vapour joins the general circulation in the atmosphere and, subsequently, falls back on land through precipitation (rain and snow). Increased precipitation leads to increased overland flow. This water can be intercepted and channelled through turbines to generate power (Hill, O'Keefe and Snape, 1995).

Understanding the principles of the water cycle is the first step to forecasting flows and water levels in a given watershed, where a watershed is defined as "a combination of both the surface drainage area and the parcel of subsurface soils and geologic formations that underlie it" (Freeze and Cherry, 1979).

A hydrological forecast for a given region begins with a model for forecasting precipitation over the watershed in question. Precipitation is delivered to streams both as overland flow to tributary channels and by subsurface flow rates as groundwater (Freeze and Cherry, 1979). Determining where and how the water travels, as overland flow, on the terrain requires an understanding of hydraulics. A detailed model also requires accurate topographic information (which can be obtained from various data sources, such as digital elevation models and then entered into Geographical Information Systems – GIS).

In principle, once this hydrologic information is integrated, a computationally intensive process can be used to make predictions of water levels (referenced to some datum such as sea level) and flow rates. Water level information can then be used to calculate the volume stored in every body of water within a GIS framework. Given the inherent difficulty of making accurate point forecasts of rainfall, this approach is very challenging to implement, yet it is conceptually rather straightforward.

The other missing piece of the puzzle is measuring the amount of time it takes for a raindrop to travel from where it fell to another point on Earth's surface. This travel time is called the "characteristic time" and, with this measurement in hand, the flow and level at this point can be obtained by "looking backwards" and determining the rainfall at that time (either by forecasting or by looking at observed data of past rainfalls). It should be noted that even this "backcasting" step, if applicable, is not trivial as detailed fine grained rain measurements are rarely made throughout a watershed, let alone recorded.

The forecasting step is even more challenging, because the grid- or mesh- size on a General Circulation Model (GCM) is much larger than the area over which rainfall forecasts are needed[5]. There is also some fractal self-similarity at play here between watersheds of different sizes (please see Appendix 1 for a more detailed discussion).

2.2. Hydrological Prediction on Short Time Scales

A significant amount of hydrological forecasting efforts is oriented to the prediction of floods. The basic prediction idea (Castelli 1995) is that the watershed has a characteristic response time, which, for a small or medium sized watershed, may only be a few hours (Siccardi and Adom 1993). This response time represents the time required for an input of rainfall to be collected and arrive at some point on the river, where it manifests itself in a storm surge. As a warning of a flood with only a few hours lead time is of little value, the quest for accurate flood forecasts requires accurate rainfall forecasts.

---

[5] This is the case for GCM-based temperature measurements as well, but since the spatial distribution of temperature is more uniform, the forecast at a central grid point can be used throughout the entire cell. But rainfall is extremely variable on spatial and temporal scales, with scaling reported to be fractal (Breslin and Belward, 1999).

However, the prediction of rainfall is very difficult. Rainfall is extremely variable, both in space and in time. It is created by a variety of interacting physical processes which occur on scales varying from micrometers to hundreds of kilometers (Balaji and Tarboton 1993). As weather models have a grid size of just a few kilometres, physical processes, including the important ones relating to cloud formation occuring on smaller length scales must be approximated, for instance by the "dynamical downsizing" of these grids (Cluckie, Yuan, and Wang 2006). Most operational rain forecasts are for fewer than seven days into the future.

The attention paid to the prediction of extreme hydrological events such as floods is understandable given their severe economic impact. However, it is a simple statistical fact that prediction, or even characterization, of extreme events tends to be much more challenging than predicting low order moments such as the mean and the variance of the same process[6]. For hydropower producers, these low order moments (such as expected flow and variance of that flow) are of significant economic value.

It should be noted that, because of the challenges in precipitation forecasting, combined with the difficulty of modelling watersheds, hydrologists were early proponents of probabilistic forecasting (Krzysztofowicz et. al. 1993).

2.3. Hydrological Prediction on Longer (six month) Time Scales

In cold climates, determining the amount of water stored in the "snow pack" and how the snow pack will melt in the coming season can provide significant insight into the behaviour of future water flows. Modelling of snow pack melt is a challenging subject with many ties to glaciology (Marshall, 2007). Snow pack thickness can be measured by a variety of ways, including field studies to obtain in-situ measurements and remote sensing (as discussed in the Earth Observation for Renewables Case Study). This data is used along with statistical analysis to estimate the total snow pack (Marshall, 2007).

On annual time scales, some insight into next year's climate variables can be obtained via El-Nino Southern Oscillation climatology. Even on very large basins like the Columbia River Basin, water inflows vary a great deal from year to year. See Figure 6 for details.

---

[6] Paradoxically, this might not be so much the case for hydrology as forecasting the flood might only require us to understand the precipitation and the simple surface geometric paths taken by the water, without requiring us to understand the subsurface processes.

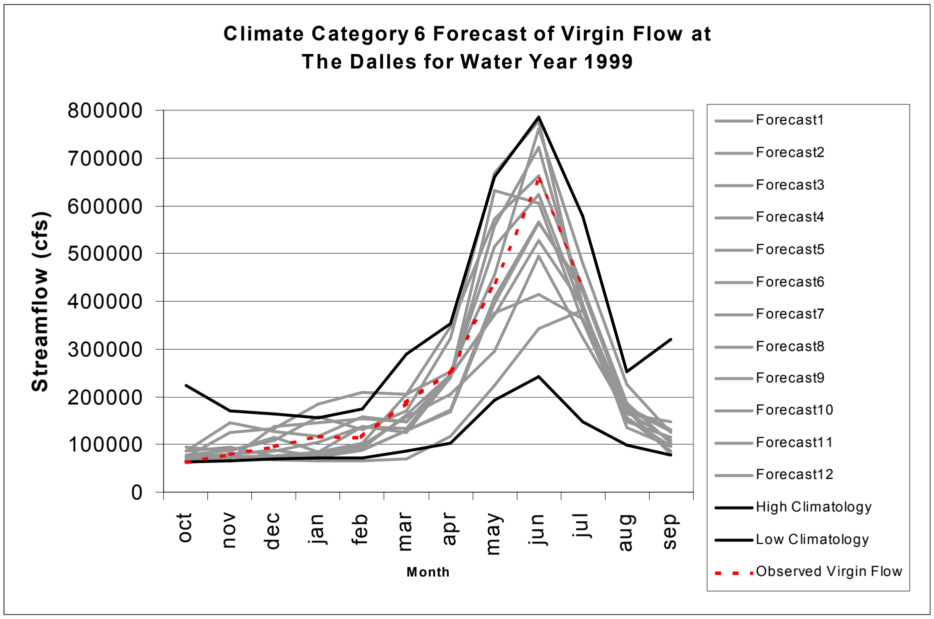

Figure 6: Long-range ensemble streamflow forecast for water year 1999 (PDO cold / ENSO cold)

Source: Hamlet Huppert and Lettenmaier (2002). Light gray lines are forecast ensemble members, heavy black lines are the observed minimum and maximum streamflow for each month for 1948-1988.

The historical approach to forecasting water levels on the Columbia river, is summarized by Hamlet Huppert and Lettenamier (2002), as follows:

> "The reservoir operating system for the Columbia River Basin has evolved to make use of streamflow forecasting techniques that become available on January 1 based on observed snowpack and statistical relationships between snowpack and spring and summer streamflow. These forecasts are updated monthly until July. As the snow accumulation season progresses, the forecasts of spring runoff become increasingly accurate until about April or May as knowledge of the total snow water equivalent that will contribute to runoff during the spring melt improves (Koch and Buller, 1993). As the snow melts through the summer, the subsequent streamflow becomes progressively less dependent on the accumulated snow from the previous winter season."

However, Hamlet Huppert and Lettenmaier (2002) have been able to improve the quality of this forecast using a macro-scale model which incorporates, at a coarse level, the hydrology described above.

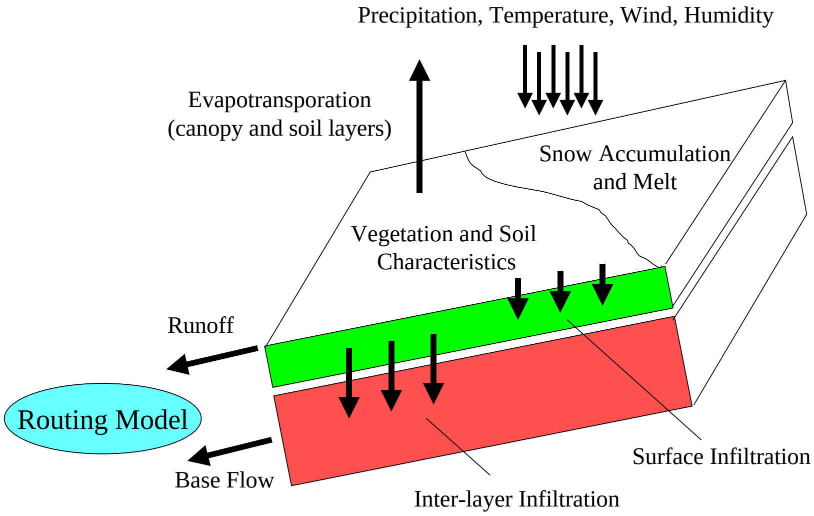

Figure 7: Schematic of grid-based macro-scale hydrology model
Source: Hamlet Huppert and Lettenmaier 2002.

This model enables a significant improvement in flow forecast over the long-range historical base implicit in the statistical relationship-driven snowpack model described in Koch and Buller (1993). In Section 3.2 we will describe the way in which Hamlet, Huppert & Lettenmaier (2002) describe how this improved predictive power can be monetized in a way which improves the ability of a hydro dam operator, similar to those on the snow pack dominated Columbia River Basin, to use their storage capacity.

2.4. Long term Hydrological Prediction

Of course in a larger watershed, or in one which is dominated by the storage of frozen precipitation in the form of snow, much longer term predictions (which are not so dependent on the day-to-day details of precipitation) is possible. For example, in the North American Great Lakes basin, the Canadian Department of Fisheries and Oceans (DFO 2007) prepares monthly probabilistic water level forecasts six months in advance. Note that for the Great Lakes, the difference in time scales between available data and the range over which useful forecasts can be made is incredible: six minute interval data can be downloaded from NOAA (2007).

Some researchers suggest that the water levels of the Great Lakes might be predicted, or at least the behaviour of their time series usefully characterized, on much longer time scales. For instance, a recent study (International Lake Ontario St. Lawrence River Study, 2007) included a simulation model which characterizes Great Lakes water levels and St. Lawrence River flows by providing 500 possible 100 year scenarios (at ¼ month time steps for Lake Ontario and monthly time steps for the upper Great Lakes). This model was created by Dr. Laura Fagherazzi of Hydro Québec in consultation with Prof. J. Salas of Colorado State University. This simulation effort preserves the time correlation observed in historical levels of the Great Lakes as well as the correlation from lake to lake. It should be noted that, like many hydrological systems, the Great Lakes is a managed one, with both Lake Superior and Lake Ontario outflows being controlled under the management of the bi-national International Joint Commission.

Finally, the impact of climate change is likely to have an impact on Great Lakes' water levels (Lam and Schertzer 1999). For instance, only half of the water lost each year from Lake Superior leaves via the St. Mary's river – the other half evaporates. Ice cover in the winter can slow this evaporation but the warming trend delays or even prevents the formation of this ice cover. Langan (2007) confirms that the ice cover on Lake Superior suppresses evaporation in the winter, and climate change can be an important factor impacting the amount of evaporation in the long-run.

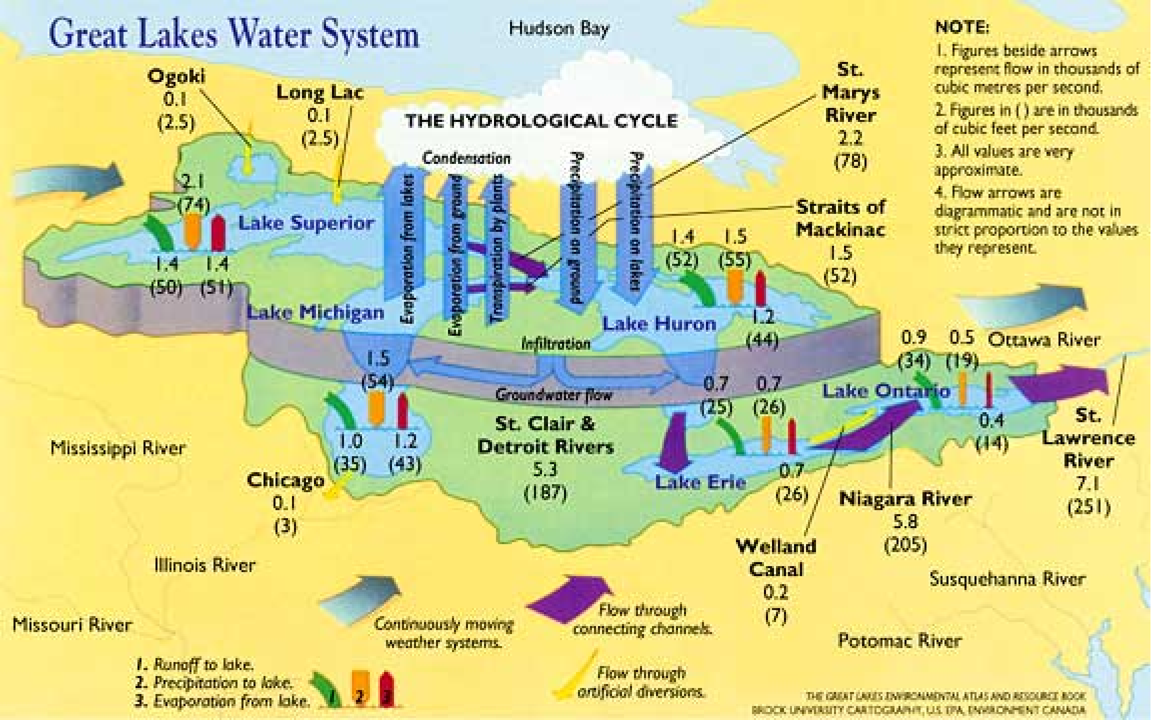

Figure 8: A depiction of a lumped system model for the Great Lakes watershed

Source: The Great Lakes Environmental Atlas and Resource Book.

3. Characteristics of the Decision Problem

3.1. Modelling the Decision Problem: Operational

A relatively recent trend in renewable energy projects is the designation of "microhyro" systems, small intermittent watersheds with a limited amount of storage. For such watersheds, with short characteristic times, prediction of water levels and flow rates is largely dependent on rainfall forecasts integrated over the (rather small) watershed area.

The small watershed size implies that the rainfall does not have a long lead time to reach the watershed. However, the small area also means that the type of rainfall forecast required is close to a point forecast, the most difficult type. Thus, reasonable time horizons for hydrological forecasts in this setting are measured in hours to a few days. The value of these hydrological predictions lies in their contribution to short term “unit commitment” or plant scheduling. Unit commitment has two main values: the first to allow the buffering of supply/demand imbalances in the system (as described in Section 3.1.1), and the second to optimally utilize the water inflows provided by these watersheds (as decribed in Section 3.1.2).

Since the marginal cost of producing hydro power is very close to zero, the decision of whether or not to dispatch the power produced by the microhydro facility is a non-issue: if we assume that the amount of water that can be stored behind the turbine is small, the water will certainly be used, and the power dispatched, within the day. Therefore, the value of an accurate hydrological forecast lies in the impact this dispatch of power has on the remaining units within the system[7].

Hydroelectric power plays a very important role in any modern electrical power system by providing an invaluable means of buffering the inevitable fluctuations on short time scales. Since it is impossible to keep an “inventory” of electricity on hand, there must always be enough electricity being generated in the system to meet the fluctuating demand. So called “quick-response hydro” is available in most Canadian jurisdictions, in which the amount of water sent through the turbines can be changed on a very short notice (Tuenter, 2007). In jurisdictions with very little hydro, the equivalent balancing services must be provided by the more expensive “spinning reserve” in which coal- or gas- fired plants are warmed and ready to be ramped up to generate power. It is clear that the more accurate the short-term prediction of water inflows, the more prominent the role of hydro will be for balancing services. Minimizing the spinning reserve of fossil fuel systems can yield multiple benefits, including the reduction of GHG emissions.

An appropriate modelling approach to capture the value of accurate short term forecasts of water inflows is to treat the power generated by a turbine as “negative load” (a negative source of demand, which reduces the amount of power required from the other generators in the system). This approach is detailed in the next section.

3.1.1. The Value of Hydrological EP to Operational Valuation: Reducing the need for Spinning Reserve

If we consider hydroelectric power as "negative load", valuing short term hydrological forecasts fits well within the theoretical framework of Hobbs et al. (1999). This framework, based both on an extensive survey of U.S. electrical utilities and on mathematical modelling, allows the value of a decrease in mean absolute percentage error (MAPE) of a load forecast to be assessed. It was found that a 1% reduction in MAPE decreases variable generation costs by 0.1%-0.3% when MAPE is in the range of 3-5%[8].

For completeness, it should also be noted that the primary determinant of load forecasts is weather forecasts, and the value of these forecasts within the framework established by Hobbs et al. (1999) is analyzed by Teisberg, Weiher, and Khotanzad (2004).

---

[7] See Yamin (2004), for a more detailed discussion of this “unit commitment problem”.

[8] An earlier work by Ranaweera, Karady and Farmer (1997) and the references cited therein suggest that MAPE tends to lie in the range of 2-5%.

Our valuation framework is based on modelling hydro power as "negative load". We ignore any diversification effect inherent in combining inflow and load uncertainty and consider only the savings attributable to EP by requiring less backup for uncertain hydro inflows.

In Alberta, hydroelectric generation accounts for about 8% of an 11,500 MW market for a notional capacity of about 1,000 MW. However, hydroelectric generation in this province is quite variable and very seasonal (Li, 2007) as discussed further in Section 3.2.1. During peak flow "spring snowmelt" season, when power output is highest, the variability in this output will be negligible. Likewise, at the end of the fall season, when natural water flows are lowest, variability is negligible because any flow across dams is the result of planned reservoir withdrawals.

We assume that there is appreciable variability in water inflows during about 3 months of the year, and that during these times the capacity factor[9] for hydro is about 25%[10]. During this variable time of the year, hydro power will require some backup buffering. Even though a hydro dam might be able to provide its own buffering through release of stored water, this release will primarily be due to load balancing reasons and not economic ones. Therefore the opportunity cost modelled by Hobbs and his co-workers is present in this case.

During this 25% capacity season, we assume that the uncertainty in hydrological inflows at a given dam is about 10% over the 2-day time window. As TransAlta alone has 13 dams (albeit some of which are on the same river system) we can suppose that the actual variability over the entire system is scaled down, by diversification, to about 3%. A further 33% relative improvement in operational scale river inflow forecasting would, therefore, result in a further 1% improvement of hydro system MAPE. Please consult the MAPE sheet of the Hydrology_Worksheet.xls spreadsheet for the (simple) implementation of this analysis.

We follow Hobbs et al. (1999) and suppose that this improvement saves about 0.2% of the variable generation cost during that same time period. If we assume the variable generation cost in the Alberta market is largely natural gas fueled, for a total fuel cost of $50/MWh, then saving 0.2% of this cost would amount to saving $0.10/MWh.

Over three months, a 1,000 MW facility working at 25% capacity factor generates 546,000 MWh of electricity[11]. Therefore, saving variable costs of $0.10/MWh leads to overall cost savings of about $55,000 per year for the entire Alberta market.

---

[9] A capacity factor of 100% means that the system can work at its maximum design capacity all the time.

[10] This way we have 100% capacity ¼ of the year, 25% capacity ¼ of the year, and essentially 0% capacity for the rest of the year. This assumption leads to correct order-of-magnitude annual capacity factors.

[11] 1,000 MW* 0.25 capacity factor*168 hours/week *13 weeks = 546,000 MWh

Obviously, this estimate is based on fairly strong assumptions. However, even ten times higher savings would not be particularly material in the context of the Alberta electricity market. This result indicates that this form of EP is not as valuable for industry in comparison to other applications investigated in this report.

3.1.2. The Value of Hydrological EP to Operational Valuation: Better Operation of Microhydro with Limited Storage

The operator of a hydroelectric facility with some limited ability to store water can profit from fluctuations in the price for electrical power. During periods of low power prices, water can be held back and the amount of power generated can be decreased. The amount of water saved can then be channelled through the turbine during periods of high power prices. However, implementing this strategy requires overcoming two challenges: first, the future price is not perfectly predictable, and second, future water inflows may not be perfectly predictable either.

In regulated markets with time-of-day pricing (for example, the seasonal power pricing present in Ontario, albeit over a longer time scale) the price will be predictable by definition. Even in deregulated markets, Davison et al. (2003) show that the price is largely determined by the balance between the load and the supply of power. Over short time horizons, the power supply is largely predictable, barring unplanned outages which are not very likely to occur within the next few days. The demand for power is, over short time horizons, again almost entirely determined by foreseeable cultural factors (such as day-of-week effects) and the humidity-adjusted temperature (Zhou and Davison, 2007; Kucera, personal communication 2007) which can also be predicted accurately over short (24-72 hour) time scales.

In order to isolate the value of improved hydrological EP in this decision problem, Zhou and Davison consider a simple model of a pump assisted hydroelectric facility operating in a market with time-varying but deterministic power prices and hourly water inflows following a simple random model. The engineering details of the facility are described by a model containing several parameters.

The next step is to determine the optimal operating strategy for this facility based on the available information. Once this strategy is determined, it is relatively easy to compute the expected value of this facility using a dynamic programming approach. By modelling the problem with a 48-hour time horizon, Zhou and Davison assume that, while the form of the stochastic model governing the future inflows is known, the actual inflows are unknown. To build insight, hourly inflows are modelled as either high or low, both having equal and independent probability. Given this assumption, stochastic dynamic programming is used to calculate the optimal control and the associated value of the pump storage facility.

Next, they assume that the future inflows, although generated from sampling within the same stochastic model, are known and use deterministic dynamic programming to obtain the optimal controls and associated values for each set of inflows. Finally, using Monte Carlo techniques they compute the average of these two values. The first optimal control assumes random inflows, the second assumes deterministic inflows (which are drawn from the same distribution as the random inflows, after the fact). In this case, they are able to conclude that, irrespective of the initial fill level of the reservoir, an accurate 48-hour flow forecast increases the value of the pump storage facility by about 2%. This result does not change fundamentally even if the pump storage facility is replaced by a facility which can merely store water for later release without pumping ability.

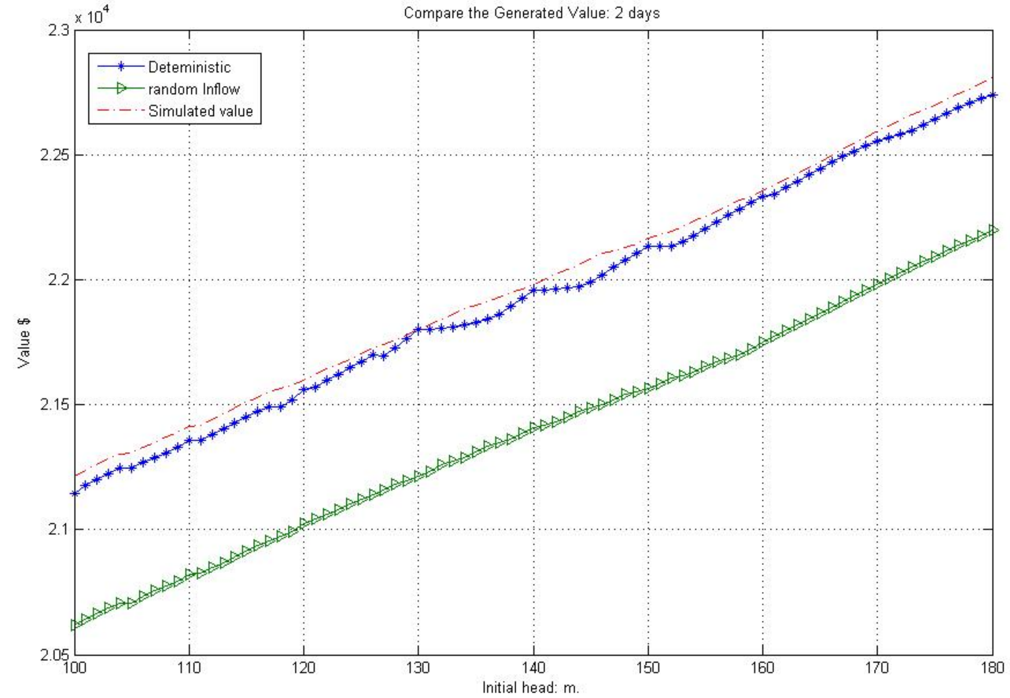

Figure 9: Comparing the value of a pump storage plant with and without accurate flow rate prediction
Source: Zhou and Davison (2007)

In order to determine the value of this prediction capability for the Alberta setting, we use the same base case scenario as in Section 3.1.1., but this time with storage as an additional model feature. If we take the 2% savings at face value, then the annual value of perfect 48-hour load forecasts to Alberta will be approximately \$3 million per year[12]. As discussed in Section 3.1.1, if during the time of the year when ¾ of the annual hydropower generation takes place, water forecasts are already such that no improvement in control is possible (because the best strategy available is "use it or lose it"), then a smaller value of \$750,000 would be our estimate for the value of EP. Although this number is also based on many assumptions, using the powerful mathematical and computational tools outlined here, a more detailed analysis can be conducted to gain more insight as part of a future study.

---

[12] 1,000 MW x 8,760 hours/year x \$70/MWh x 2% x 25%. The \$70 is the total average price (thus including fixed costs and profit) to purchase electricity in Alberta over the course of the year.

3.2. Modelling the Decision Problem: Tactical

Several interviewees (Anderson, Tuenter) as well as personal communication with Kucera (2007) indicated that longer range hydro forecasts could be used to schedule planned maintenance for other generation units. The reasoning behind this is quite simple: if in two weeks' time a large influx of water is predicted, this flow can be used to generate large amounts of power thereby freeing some capacity from other parts of the generation system, which can instead be diverted for planned maintenance.

In a larger watershed, longer time scale hydrological forecasts are possible. In particular, this is the case when there is a large amount of frozen precipitation stored every winter in the form of snowpack[13].

In such systems, which are prevalent in the western part of North America, both in British Columbia and in the hydroelectric assets controlled by the Bonneville Power Administration, water levels in reservoirs can be used to manage seasonal fluctuations in demand. They can also be used as a means to profit from seasonal price fluctuations, by means of offering hydro production into neighbouring markets during times of high prices and in conserving water in the reservoirs for future use, even perhaps by purchase of power from the neighbouring jurisdictions when prices are low.

The ability of such "water trading" entities to generate profits from time varying electricity prices in neighbouring markets is limited by two main factors: the elasticity of prices to additional supply or demand of the neighbouring market[14] and the ability to store water. Even the largest reservoirs cannot prudently be filled above an upper threshold, nor drawn down below a lower threshold. If future water inflows are uncertain, these "safety limits" must be computed by taking into account certain risk factors. If, on the other hand, future water inflows can be well characterized, water supplies can be managed much more aggressively. This approach has significant potential economic value as described by Hamlet, Huppert, and Lettenmaier (2002).

It is important to note that most watersheds are used by more than one group of stakeholders. Rivers provide not only hydroelectric generation but also important societal benefits (for instance through recreational use). In many cases, they also provide unique habitats for fish and other flora and fauna. These other uses provide additional constraints governing the use of the rivers. As a case in point, the amount of water allowed to flow over Niagara Falls is constrained by the need to display the majesty of the falls during the daytime. This and other constraints on the Great Lakes watershed are specified in the Niagara River Treaty. Moreover, regulators tend to be

---

[13] Even in a hydro system with limited or no water storage possibilities, month-to-year time scale water forecasts are useful and actionable since they allow pre-buying of appropriate amounts of coal (in markets in which coal fired power plants make a contribution), and also since they allow routine maintenance of fossil fuel and nuclear plants to be scheduled for a time during which hydro electric contributions to the supply stack are at their peak. While these benefits are concrete, they are difficult to quantify and are also minor in magnitude compared to the value of longer time horizon hydrological forecasts in hydro systems which allow storing large amounts of water.

[14] The amount the prices in the neighbouring market change in response to large influxes or withdrawals of power from their market.

very risk averse and prefer erring on the more conservative side when it comes to managing water levels. For instance, policy makers charged with the regulation of water levels at key points on the Columbia River watershed have consistently used extreme observations in records, ranging from the lowest water level year ever recorded to the third-lowest water level year (in about a century of measurements). Such conservative protocols, while ensuring adequate water levels at all times and for all purposes, also have the negative effect of leaving a lot of water "on the table" in most years – water which could have been used to generate valuable hydroelectric power. If a better forecast of extreme levels is possible, and if this forecast is accepted by regulating bodies, it can actually provide significant economic value.

3.2.1. First steps to Tactical EP Valuation

The information provided by Lin Li (2007), a hydrologist with TransAlta Corp (the operator of many Alberta hydroelectric facilities), together with data on variable Alberta electricity market prices, was used to obtain a rough estimate of the value of annual EP projections for the Alberta market.

Two important parameters for a hydroelectric system are the capacity factor of the system and the storage capacity (the ratio of the amount of water that can be stored in the reservoir to the average annual flow rate of the rivers).

Table 1 includes the data for two Alberta river systems, the Bow River system and the North Saskatchewan River system (Li, 2007). TransAlta's reservoirs are drawn down to their minimum capacity each spring to prepare them to accept as much of the spring runoff as possible.

River System	Capacity Factor	Storage Capacity
Bow River	31%	24%
North Saskatchewan River	21%	44%

Table 1: Data for Alberta river systems

Note: Capacity Factor denotes the ratio of average annual energy production from the river to the amount of energy which would be produced if the generating stations on the river were working at full output all the time. Storage Capacity denotes the ratio of average annual runoff to the reservoir volume.

The Alberta Electricity Systems Operator reports that the capacity of the hydroelectric plants in Alberta is about 8% of the 11,500 MW generation capacity of Alberta, or about 1,000 MW when working at full capacity. Public power price and load data for Alberta, for 2005-2006, as archived by Davison, reveal the following pattern of prices (Table 2)[15].

It is clear from Table 2 that in a system like Alberta's, where high flow rate times correspond to low price times (May-July), the ability to save water for use in August-November is very valuable. Indeed, such storage could allow power to be sold for $100/MWh rather than $50/MWh.

---

[15] On–peak defined as 7AM to 11PM.

The experience of Hamlet Huppert and Lettenmaier (2002) on the Columbia River watershed suggests that the difference between a dry- and a wet-year is substantial – this is also the case in the intermittent Alberta market. Assume that an EP application for annual Alberta flow rate allowed for an additional 10% of the high capacity spring runoff water to be stored for release in the late summer or winter high price season[16]. In this case, the value of EPs for the Alberta market is very significant: about $3.6 million[17] (accounting for the difference between the roughly $100/MWh available during the high demand, low flow months and the roughly $50/MWh available during the spring months).

	Avg price ($/MWh)	Avg load (MW)	On Pk Price ($/MWh)	On pk Load (MW)
Jan	72.12	8174	85.82	8458
Feb	54.07	8160	63.33	8431
March	44.08	8006	51.34	8260
April	42.87	7549	51.40	7790
May	56.26	7453	69.74	7717
June	51.16	7349	63.92	7620
July	37.75	7422	44.69	7645
August	88.33	7452	102.95	7709
Sept	74.30	7360	93.57	7632
Oct	121.92	7530	144.30	7798
Nov	124.80	7859	155.21	8157
Dec	103.03	8205	127.60	8491

Table 2: Alberta price and load data (from Davison's data archive)

3.3. Modelling the Decision Problem: Strategic

Hydro power assets are very expensive to construct but, once commissioned, they are operational for a long time. Consider the following two examples which illustrate the expected operational lifetime of dams. The Hoover Dam, which created Lake Mead in Nevada, was constructed in the 1930s at a cost of approximately $670 million in today's dollars and is still functioning. Much of the work constructing Ontario Power Generation's Adam Beck power plant at Niagara Falls was done over a century ago (in terms of tunnel construction, preparation of site, etc.) and is still being used today. Turbines may occasionally be refurbished and replaced, but investments in hydroelectric facilities have an extraordinarily long life.

Of course, predicting water inflows on such a long time scale, even for a watershed as huge as that of the Great Lakes is very difficult. As far as the scope of this case study is concerned, climate change effects may need to be taken into account (Robinson 1997, Visconti et al. 2003). Even without considering this major confounding factor, it is true that long term forecasting of water levels is a great challenge, as discussed in Section 2.

---

[16] This assumption is relatively conservative given Columbia River experience, and given the fact that the storage capacity of the two river systems allows for the storage of roughly an entire month of max power production flow rate, as seen in Table 1.

[17] 10% x 1,000 MW x 720 hours x ($100 - $50)

According to Tuenter (2007), Ontario Power Generation's approach to this calculation was to use the insights of their in-house hydrology experts coupled with the plethora of data assembled on Great Lakes' water levels and inflows over the last century. However, the forecast methodology used was more about "pattern recognition" than any other method.

3.3.1. First steps to Strategic EP Valuation

An investment decision to construct a project with a very long service lifetime is a difficult problem, as there are many uncertainties which are very hard to characterize in advance. Consider the risk management process described by Tuenter (2007) that factored into the decision to build a third tunnel (Ontario Power Generation 2007) to channel water into the Adam Beck plant at Niagara Falls. This tunnel, when completed, will be able to divert more water from the Niagara River through Adam Beck's turbines (currently the flow rate in the Niagara river exceeds the ability of Adam Beck to utilize it about 65% of the time; after the construction of the third tunnel this figure will drop to 15%). For more details of this analysis, see the ThirdTunnel worksheet of the Hydrology_Worksheet.xls spreadsheet.

If the current water flow patterns of the Niagara River remain more or less unchanged, this investment will be able to generate an additional 1.6 million MWh of electricity per year. At current Ontario wholesale power rates of about $50 per MWh, the value of this electricity will be between $80 and $100 million dollars per year. On the other hand, the cost of the third tunnel project is estimated to be $985 million[18]. Construction began in September 2005 and is slated for completion in 2010.

This investment will not be profitable until 2033, assuming the construction project takes the scheduled 5 years; an equal cost is incurred in each of these years; OPG pays about 6% on its debt; the hydrological forecasts yielding an annual additional power production of 1.6 million MWh are correct; the average annual wholesale value of power to OPG is $50 per MWh, increasing after 2010 by an annual growth rate of 2% per year, and the average annual costs to maintain the project are $5 million, also increasing with the 2% annual growth rate. Clearly there are many financial sources of uncertainty to this project but a major input is the mean water flow on the Niagara river, especially considering that for this project to pay off, the additional flow over and above that already used by Beck and the corresponding U.S. side plants must be as advertised. In other words a small percentage decrease in the annual flow will have a disproportionately larger impact on the flow available to service the third tunnel project.

Of course, for a project as big and as long-lived as this, many factors will go into making the final go/no-go decision and Environmental Predictions are just one input. Also important will be long-range predictions for electricity prices, long term predictions for interest rate levels, and political and organizational considerations. Therefore, it is difficult to tease apart the value directly attributable to EP in this case study, but it is nonetheless interesting to see that the EP is important in this context.

---

[18] All figures are from official Ontario Power Generation sources as summarized in Ontario Power Generation 2007.

4. Why is Hydrological EP not used more frequently?

4.1. EP Strategy of Transalta

At the present time, TransAlta is not using any hydrological EP in their decisions (although they are now in preliminary discussions with a German vendor who supplies Manitoba Hydro (Li 2007)). One reason for this lack of EP lies in the extreme complexity of making hydrological forecasts in the highly variable streams stemming from the Rocky Mountains. The near intermittency of many of these rivers and streams also means that the inputs of other stakeholders take on a large importance.

4.2. EP Strategy of OPG

While OPG employs some hydrologists and meteorologists (Tuenter 2007), they are limited in using the information these experts generate. Most of Ontario's hydroelectric resource is run-of-the-river, with little ability to store water. This inability to store water eliminates many of the ways that EP can be monetized. Moreover, the large hydro units in Ontario, like the Adam Beck unit at Niagara Falls, are considered by the government to be "legacy" assets that are only allowed to earn a small price per MWh. Therefore, sophisticated bidding strategies which capitalize on EP are not available for these large hydro units. At the other extreme, microhydro units in Ontario count as "green power" and are guaranteed a large ($0.11 per kWh) payoff regardless of the time the power is dispatched. While this last limitation works to the advantage of the microhydro operator, it is still a limitation nevertheless. Unless extreme power prices are experienced, the microhydro producer cannot do better than taking the fixed "green" rate.

OPG did have to make some long term projections about Great Lakes' water levels and flow rates for the construction of the new "Tunnel 3" at Adam Beck (as discussed above). However, the EP approach they adopted for this very difficult prediction problem was more based on pattern matching with the shape of existing water level and inflow curves than by any sophisticated hydrological modelling.

4.3. EP Strategy on the Columbia River Watershed

The U.S. Army Corps of Engineers restricts outflows on the Columbia River watershed using very conservative "observed worst year" or only slightly less conservative "observed third worst year" data. As indicated in the literature review, worst case or "minimax" decision makers are often in the poorest position to make use of probabilistic EP, since the reduction of a remote, though possible, event to an even more remote, yet still possible, event does not change worst case scenarios at all.

5. Conclusions and Future Work

In this case study, we were able to show four concrete cases in which Environmental Prediction could create economic value for hydroelectric operators.

We showed that at operational time scales, better EP could reduce the need for "spinning reserve", albeit at a rather modest savings of $55,000 for the entire Province of Alberta.

A new model for the improvement in water storage protocols at operational time horizons seemed, again in our rough approximations, to have more potential value of about $750,000 per year for Alberta's hydroelectric sector. Next, at a tactical level, we were able to characterize the value of combining snowpack measurements with El Nino/Southern Oscillation climatology to the Albertan hydro sector as being worth as much as $3.6 million per year. However, all of these numbers are based on relatively strong assumptions and, while they may represent profit opportunities, they do not necessarily represent the status quo value of EP to this sector.

Our final case study was based on a real-life case, this time in Ontario. Long range Great Lakes levels Environmental Prediction has been used in Ontario as one input in deciding to proceed with the so-called "third tunnel" project in which a new tunnel is being bored to generate electrical energy from currently unused water passing over Niagara falls. This project has the potential to generate significant revenue for Ontario Power Generation and Environmental Prediction was one of the tools which factored into the decision.

6. References

Anderson, C. L.,. "Personal Communication": See notes of interview held Sept 29 2007 (Lindsay Anderson is a professor in the department of Environmental Engineering at Cornell University and an expert in electricity markets and green energy)

Breslin, M.C. and J.A. Belward.,"Fractal Dimensions for rainfall time series", Mathematics and Computers in Simulation, 48, nos. 4-6, pp437-446, 1999.

Castelli, F., "Atmosphere Modelling and Hydrologic Prediction Uncertainty", US-Italy Research Workshop on the Hydrometeorology Impacts and Management of Extreme Floods, Perugia Italy Nov 1995

Cluckie, I.D. , Yuan Y. and Y. Wang, Uncertainty Analysis of hydrological ensemble forecasts in a distributed model utilizing short-range rainfall prediction. *Hydrology Earth Systems Science Discussions* 3, 3211-3237, 2006

Department of Fisheries and Oceans (Canadian Hydrographic Service), Government of Canada (2007) http://www.waterlevels.gc.ca/C&A/glfcst_e.html, accessed Oct 2 2007, archived as CdnHydrogSvceGrtLksFcst.doc

Doeschl, A., Ashmore P. and Matt Davison "Scale issues in natural and modelled braided rivers". Pp 177-198 of Sambrook Smith G. H., Best, J. L., Bristow, C., 2006

Freeze, R. A. and Cherry, J. A., *Groundwater*, Prentice-Hall, New Jersey, 1979

Gurnell, A. M. and Petts, G. E. (eds.) *Braided River: Processes, Deposits, Ecology, and Management*. Special Publication of the IAS, Blackwell Science, Oxford.

Hamlet, A.F., Daniel Huppert D. and D. P. Lettenmaier (2002). "Economic Value of Long-Lead Streamflow Forecasts for Columbia River Hydropower", *Journal of Water Resources Planning and Management*, March/April 2002, 91-101.

Hill, R., O'Keefe P. and C. Snape, The Future of Energy Use, Earthscan Publications Limited, London, 1995

Hobbs, B.F., Jitprapaikulsarn, S., Konda S., Chankong, V., Loparo, K and Dominic J Maratukulam, "Analysis of the Value for Unit Commitment of Improved Load Forecasts", IEEE Transactions on Power Systems 14(4), 1342- 1348. November 1999

Hwang, Y, Clark, M.P., and Rajagopalan, B., "Analysis and Use of Input Uncertainties in Watershed Modelling" Under review by *Water Resources Research*, 2005

Hydro Québec, "Annual Report", 2006

IEA, International Energy Agency, 2006 Electricity Information, OECD/IEA, Paris

International Energy Agency, Paris France, June 2005.

International Lake Ontario St. Lawrence River Study (2003) http://www.losl.org/twg/hhmodelling-e.html#1, accessed Oct 2 2007, archived as IJCHydrologyStudyGrp.doc, document written by Syed M. A. Moin

Koch, R.W. and, Buller, D., *Forecasting Seasonal Streamflow: Columbia River at The Dalles*, Report to Bonneville Power Administration, 1993

Konnie, L., "On the Levels", downloaded from www.lakesuperior.com October 3 2007.

Krzysztofowicz, R, Drzal, W.J., Drake,T.R., Weyman, J.C. and L.A. Giordano, Probabilistic Quantitative Precipitation Forecasts for River Basins. *Weather Forecasting* **8**, 424-439, 1993

Krzysztofowicz, R. and Herr, H.D., "Hydrologic uncertainty processor for probabilistic river stage forecasting: precipitation-dependent model**"** *Journal of Hydrology,***249***,* 46-68, 2001

Kucera, A. (2007), he holds a Ph.D. in mathematics and works for Integral Energy, a large Australian power utility, where he is the head of the risk management department, personal communication, April 2007

Lam, D.C.L. and William M. Schertzer, eds. *Potential Climate Change Effects on Great Lakes Hydrodynamics and Water Quality*. Published by the American Society of Civil Engineers, 1999

Langan, J., ,consulting hydrologist for Stantec limited, personal communication, see notes of interview conducted October 13 2007

Lee, D.H. and Anne H. Clites, "Assessing Risk in Operational Decisions Using Great Lakes Probabilistic Water Level Forecasts", *Environmental Management* **21**(1) pp 43-58, 1997

Lin, L.L., "Personal Communication": See email exchange October 2007 (Lin Li holds a Master's degree in hydrology and is employed by Transalta inc), 2007

Marshall, S., Canada Research Chair in Glaciology at the University of Calgary, personal communication, 2007

National Oceanographic and Atmospheric Administration (2007) Tides and Currents database: Available at http://tidesandcurrents.noaa.gov. Accessed Oct 2 2007.

Ontario Power Generation, "The New Niagara Tunnel Project". Website http://www.opg.com/power/hydro/new_projects/ntp/ntp_faq.asp (stored as "Why is OPG building a third tunnel.doc"), 2007

Rajagopalan, B. and David G. Tarboton "Understanding Complexity in the structure of rainfall" *Fractals* 1(3) 606-616., 1993

Ranaweera, D.K., Karady, G.K. and Richard G. Farmer, "Economic Impact Analysis of Load Forecasting", *IEEE Transactions on Power Systems* **12**(3), 1388-1392., 1997

Robinson, P.J., (1997). "Climate Change and Hydropower Generation", *International Journal of Climatology* **17**(9) 983-996, 1997

Roulin, E. and S. Vannitsem, "Skill of medium-range hydrological ensemble prediction", *J. Hydrometeorology* **6** 729-744., 2005

Roulin, E.,"Skill and relative economic value of medium-range hydrological ensemble prediction", *Hydrological Earth Systems Science* **11**, 725-737. (Section 3.2 on page 728 quantifies economic value), 2007

Roulston, M.S., Kaplan, D.T., Hardenberg, J. and L.A. Smith :Using medium-range weather forecasts to improve the value of wind energy production". *Renewable Energy* 28, 585-602, 2003

Siccardi, F. and D. Adom, "A non-structural policy for the mitigation of flood effects: The Arno Project". pp. 3-12 of Prediction and Perception of Natural Hazards, J. Nemec, J.M. Nigg and F. Siccardi, eds, Kluwer Dordrecht, 1993

Teisberg, T.J., Weiher, R.F. and A. Khotanzad "The Value of National Weather Service Forecasts in Scheduling Electricity Generation", Manuscript, 2004

Tuenter, S., "Personal Communication": See notes of Interview held July 2007.

Visconti, G., Beniston, M., Iannorelli, E. D. and Diego Barba, editors. *The Effects of Global Warming on Mountain Regions: A summary of the 1995 report of the Intergovernmental Panel on Climate Change*. Springer Netherlands. , 2003

Yamin, H.Y., "Review on methods of generation scheduling in electric power systems", *Electrical Power Systems Research* **69** 227-248, 2004

Zhao, G. and Davison, M., "Estimating the value of accurate hydrological Forecasts for a Pump Storage Facility", Working Paper, 2007

Zhao, Z. and Davison, M. "A principal component/Fourier series analysis of Ontario Electricity Load Shapes and their relation to temperature and humidex" Working Paper, 2007

Appendix 1: More on Hydrological Forecasting

If for a large watershed with area A_1 the characteristic time scale is t_1, another smaller scale watershed with area A_2 might have characteristic time scale $t_1(A_2/A_1)^f$. Although such relationships do allow suitably scaled measurements of sand-table type lab experiments to be extrapolated to natural scale rivers (see Doeschl, Ashmore & Davison 2006), the continuum of results they imply makes it very difficult to make categorical statements about for how long hydrological variables can be predicted. Each watershed is different, not only in size but also in many other ways including subsurface soil geologic formations and vegetation. Therefore, for each watershed, different studies must be employed based on the same general principles.

To illustrate these general principles we discuss so called "impulse-response" hydrological models, in which the aggregate rainfall across a given watershed, or portion thereof, is predicted, estimated, and recorded. These general principles may be illustrated by considering a simplified situation, in which processes below the surface of the Earth are ignored (i.e., assuming that a cement surface exists over the whole watershed).

In these impulse-response models, a single characteristic time is also estimated (in order to establish the time required for a given supply of rainfall to reach a given point). With these variables in hand the flow at location x and time t is the sum off all the rainfall that happened at time $t\text{-}s_1$ which travels along paths that cause it to arrive at point x s_1 time units after being precipitated, plus all the rainfall that happened at time $t - s_2$ travelling along paths causing it to arrive at point x s_2 time units after being precipitated, and so on. These figures must also be corrected for evaporation.

The next step is to remove the simplifying assumption of the impervious shell and consider permeability. Earth is made of rock and soil, each with its own degree of permeability, which may or may not be covered by flow-altering vegetation. Rain falling at a given point now has an option of flowing along the surface of a stream network to the next measuring point (as discussed above), it but can also soak into the Earth (at which point its flow rate is defined by more complicated "flow through porous media" equations). The geometry of the flow is by necessity largely unknown, since water can sojourn in a wetland, or be taken up by a plant and transpired back into the atmosphere, etc. Modelling these processes requires much more data, most of which is not measured, as well as understanding yet more complicated physical processes to an already challenging suite of problems[19]. Moreover, the geographical inputs or "boundary conditions" are rarely fixed – changes in land use can dramatically change watershed performance (e.g., paving over farmland leads to dramatic increase in storm runoff).

---

[19] Note that in sedimentary rock environments where geological action has resulted in the tilting of sedimentary layers, the porosity and permeability of outcropping surface rock can range over three orders of magnitude within a kilometer.

Faced with this many layers of complexity, the course often taken is to consider the hydrological forecast to be a random or stochastic model, in which random rainfall inputs are considered to be passed through a network of random time scales (corresponding to the different paths a raindrop could take through the system).

$$F(x,t) = \int_{-\infty}^{t} K(\tau,t)R(\tau)d\tau$$

Where $K(\tau,t)$ is the so-called system response kernel and $R(\tau)$ is the (random) rainfall forcing function.

Such integration of random impulse can lead to so-called "fractional Brownian Motion" models – probabilistic, lumped system, hydrological models which have become popular in recent years for identifying possible long-range watershed scenarios.

Environmental Predictions and the Energy Sector: A Canadian Perspective

Case Study 3:
Environmental Prediction for Wind/Water Integration

Prepared for

Environment Canada

Contract Number: K3A40-06-0028

Prepared by

TURQUOISE
TECHNOLOGY SOLUTIONS INC.

DYDEX
DYDEX RESEARCH AND CAPITAL LTD.

1. Introduction

Wind energy constitutes one of the most dynamic sectors of the energy industry both in Canada and abroad. The installed capacity rates have been increasing steadily since the early 1990s and the rate of growth seems to be continuing with great strength. Wind energy industry is booming in Canada: the installed capacity more than doubled from 683 MW in 2005 to 1,459 MW at the end of 2006. There are indications that this boom is very likely to continue, since provincial governments are targeting to have a minimum of 10,000 MW of installed wind energy capacity in place by 2015 (GWEC, 2007).

Environmental Predictions are already having a big impact on the way wind energy projects are built and operated. This case study highlights the importance of EP in the decision-making process, both at the individual wind farm level as well as the systems-level. Valuation models are provided indicating the economic value of EP in strategic and operational decision-making.

A wind turbine generates power by converting the kinetic energy of wind into electricity. The conversion starts with the rotation of the blades caused by wind. The blades rotate a shaft inside the turbine, which moves a magnetic field in the generator, thus creating electricity.

Wind turbines come in different sizes and generation capabilities. At the utility scale, an individual wind turbine can generate up to 5 MW of power. Wind turbines can be installed in clusters called "wind farms", and generate large amounts of electricity. At the other end of the spectrum, small wind turbines, (as defined by less than 100 kW of power), are designed for residential and small-scale commercial use to serve as a backup electricity source, or to reduce electricity bills.

There are a number of factors which make wind power appealing from the perspective of individual developers as well as policy makers. In contrast to most other generation types (with the exception of hydro and other renewables), the "fuel" of wind energy is free. In addition, wind energy has a very limited environmental footprint, and except for some greenhouse gas emissions during the manufacturing process, it does not cause any harmful emissions[20]. Of course a complete environmental accounting of the costs and benefits of wind power would need to incorporate the impact of wind turbines on birds and bats, as well as the additional cost, both financial and environmental, of building additional transmission line and access road infrastructure to wind sites.

However, as will be discussed in this case study, there are also some disadvantages of wind power. The free fuel comes with a caveat: there is no guarantee of stable flow, and wind speed can change significantly over time. This characteristic is shared by many other types of renewables which also rely on natural processes with varying degrees of intermittency. This renders wind energy "non-dispatchable", as it cannot be summoned at will and the production cannot be ramped up. Therefore, "wind needs a dance

---

[20] In fact, wind energy is environmentally friendly, even compared with hydroelectricity as the environmental impact is lower. Notable impacts of wind energy include risks to birds, bats and insects, and limited visual and sound pollution.

partner" (Frost, 2007) and an investment in wind power must continue to be "backed up" by traditional, often highly polluting, sources of power such as coal fired plants. This duplication of investment is wasteful, but will remain necessary until wind intermittency can be fully addressed, either through storage (either directly or indirectly) or through more sophisticated methods of load balancing.

One very suitable dance partner for wind is hydroelectricity. Electricity can be stored, at least indirectly, by storing potential energy by means of retaining water in an elevated reservoir or even by pumping it there. One example with considerable potential is the case in which a dam is built to form a large reservoir with plenty of additional water storage capacity.

This case study is structured as follows: In Section 2, the current state-of-the-art in Environmental Prediction of wind is presented; Section 3 focuses on the potential of wind-coupled storage facilities and the synergies of wind and hydroelectricity; Section 4 presents various valuation approaches to capture the economic benefits of such systems.

2. Environmental Prediction of the Wind

EP can make a significant contribution to wind energy projects throughout their lifecycle. These contributions can be grouped under two main categories: resource assessment and forecasting. Resource assessment is essential to identify ideal locations for constructing wind farms and, as will be discussed below, there are many factors in addition to long-term wind speed averages which need to be taken into account. Forecasting is emerging as a very useful tool for better managing existing wind farms and ensuring their successful integration into electricity grids.

2.1. Resource Assessment

Almost all renewable energy types, with the exception of geothermal, rely on the Sun as their energy source. The wind is a result of the pressure gradient due to the uneven heating of Earth's surface by the Sun. The air above the equator is heated up by the Sun while the air around the poles is much cooler due to the angle of solar radiation reaching these regions. Since the density of air decreases with increasing temperature, the lighter air from the equator rises, causing a pressure drop around this region. This pressure drop attracts cooler air from the poles towards the equators (Mathew, 2006).

The solar effect creates weather patterns which vary at different time scales (e.g., diurnal and seasonal patterns in wind). Furthermore, local conditions, such as proximity to mountains, valleys and coastlines, land-use and complexity of the landscape also have an impact on wind. For instance, due to the fact that the relatively flat sea surface results in less friction, wind speeds tend to be consistently higher around coastlines (when the wind is blowing from the sea), and on icy surfaces.

All of these factors complicate resource assessment studies for prospective wind farm sites. Since wind speed and direction can show significant temporal and spatial variability understanding these patterns is critical before an investment decision is made.

Moreover, the focus of wind resource assessment is not directly on the surface of a location, but at around 60-100 metres above the surface where the modern-day wind turbines are located.

Wind speed is one of the critical inputs of wind resource assessment, since the energy output of a wind turbine is a function of wind speed. There is a cubic relationship between wind speed and the maximum theoretical power output. However, Betz' law stipulates that the maximum engineering efficiency deriving from this physical limit results in a proportionality constant of 0.58. Therefore, in the wind energy industry, a "power curve" is used to describe the impact of wind speed on power output. The power curves in Figure 10 are (interpolated) experimental results from one of the most commonly used wind turbines, GE's 1.5 MW series wind turbine..

Note that when the wind speed is less than 2 m/s, there is no power output. Above this level, even a small increase in wind speed can increase the power output significantly. Moreover, the output levels off after about 12 m/s, and the turbine is shut off when the wind speed reaches 20-25 m/s (depending on the model) in order to protect the turbine.

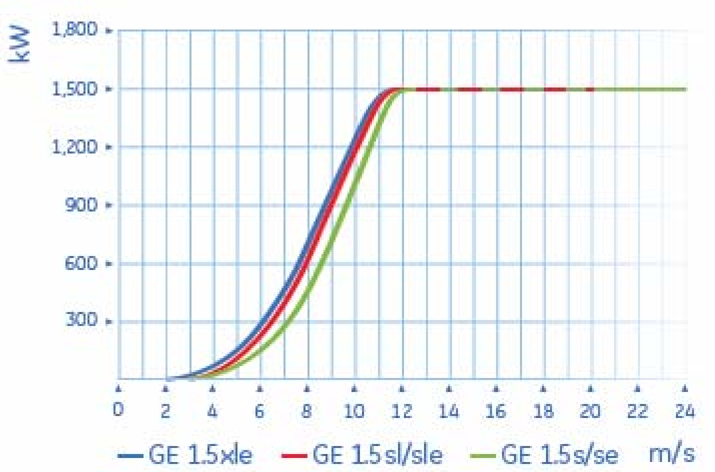

Figure 10: A sample wind curve (source: GE, 2007)

At the resource assessment stage, wind developers are interested in finding optimal locations for building wind farms. These locations should harbour long-term resource availability even though wind speed variability in shorter time scales is inevitable. Moreover, it is important to characterise the variability pattern of wind in candidate locations to determine if the power output is "in sync" with the load characteristics in the target electricity market. For instance, if the power output peaks in the afternoons when the demand for electricity is also at its peak, the revenue which can be generated from a wind farm can increase significantly. In other words, rather than being indifferent between two separate locations with long-term average wind speeds of 9 m/s each, one will prefer the location which can better match the load characteristics of the target market.

There are two main approaches to resource assessment: measurements or modelling. These two approaches are not mutually exclusive and they can be used in combination.

The industry standard for measurement is at least one year of data obtained from a meteorological mast installed in the candidate location, which measures wind speed using an anemometer and wind direction using a wind vane. This time scale is recommended to reveal the wind pattern throughout the year. However, in order to decrease any measurement errors, longer time series are always recommended. Unfortunately, although this approach provides the highest measurement accuracy, it is costly and very time-intensive. Most developers are not in a position to wait for years before making a go/no go decision.

Therefore, EP-based modelling efforts provide a viable, complementary approach when some direct measurements are available. If there are no ground-based measurements which are available, then modelling remains the only option. Some ground-based measurements can be obtained from existing meteorological masts (and other instruments) close to the location being studied. However, due to micro variations around these locations (mainly due to landscape), interpolating the measurements obtained from various sites can give misleading results. One major advantage of tapping into existing databases from nearby instruments is ability to increase the accuracy of resource assessments by time-series dating back as many as 40 years.

AnemoScope, a simulation toolkit developed by Environment Canada, uses two advanced meteorological models (Environment Canada's MC2 and MS-Micro) to calculate and predict wind flow patterns around a given location (Environment Canada, 2007). Other models used in the industry incorporate North American Regional Reanalysis data provided by the National Centers for Environmental Prediction (NCEP) in the U.S.

In order to give an indication of the cost associated with installing meteorological masts, consider that a 60-metre tower structure costs $13,000. With the addition of sensors, data loggers, telecommunication systems and one year of operations, the cost of installing and a single mast can run in the $20,000-$40,000 range.

2.2. Forecasting

An emerging EP application in the wind energy industry is forecasting. Forecasting the expected energy output from a wind farm is becoming increasingly important as the amount of installed wind power increases in a grid. Although there can be significant regional differences, the rule of thumb is a 10% threshold for installed wind power capacity in the overall electricity generation mix. Beyond this level, as the weight of wind energy generation increases, there are a number of issues that individual wind farm operators as well as system operators have to deal with.

One significant risk for the whole system is rapid loss of electricity which would normally be generated by wind farms. Although the probability of all wind farms in a given region losing electricity generation capability is very low, it is nonetheless a non-zero probability. In the case of well-connected grids with multiple connections to other generators, the sudden loss of electricity generation can be compensated for by buying electricity from outside the regional grid. However, if the regional grid is isolated (as is the case in most Canadian provinces), then balancing becomes even more challenging (this is one of the main reasons why wind energy needs to be complemented by other generation types).

Another rationale for forecasting is derived from the imbalance penalties imposed on the wind energy generators by certain system operators. In this case, the wind energy generator is responsible for the shortfall of energy promised, and is forced to buy the difference in energy in the spot market, hence paying a penalty for overestimating the electricity output from its wind turbines.

Accurate wind forecasting is also important to reassure the system operators regarding the stability of grids as more and more wind energy is installed. The Canadian wind energy industry recently passed a very significant milestone, when Alberta Electric System Operator lifted a 900MW cap which was imposed on the wind developers in this province. The cap was a result of AESO's desire to better understand wind variability issues before giving the green light for more wind energy projects. After a $1 million study on wind forecasting, AESO was confident in the capability of EP-powered tools to manage wind energy production in the province and lifted the cap, opening the way for more wind turbines (Blackwell, 2007; Kwas, 2007).

Another wind forecasting initiative is being led by Hydro Québec and Environment Canada. This system, called SPEO (Système de prévision éolienne), was developed to provide wind forecasts at various temporal and spatial resolutions. Currently, the system is being tested by Hydro Québec for 24-hour and 48-hour forecasts at the 1.5km, 2.5km and 200m resolutions. More work is being done to extend the system to 72-hour forecasts as well as very short term forecasts with significantly more accuracy (less than 6 hours) (Roberge, 2007).

3. Wind-Water coupled Pump Storage facilities

Hydroelectric plants based on reservoirs can be used to buffer demand variability in the following way: water is released through a turbine when electricity is required. When electricity is not required, the water is stored. Of course every reservoir has a finite, and possibly time dependent, capacity (due to factors such as cottage lake levels, visual beauty of Niagara falls, etc.), so water cannot be stored indefinitely. Nonetheless, hydroelectric generating facilities of this type represent an ideal way to (indirectly) store electricity. This characteristic enables hydroelectric plants to compensate for the intermittency of wind and increase overall system reliability when hydro and wind energy assets within the same grid system can be managed in an integrated fashion.

However, topographical and hydrological considerations do not allow for such facilities everywhere. For instance, in New Brunswick, hydroelectricity generation increases significantly in springtime; due to spring run-offs, this peak cannot be maintained and energy cannot be stored (Brown, 2007).

In order to address this problem and create an energy storage option, the pump storage facility was developed. The invention of pumped storage dates back to the 1890s in Italy and Switzerland. Later, during the 1930s, reversible hydroelectric turbines were developed which can operate as both turbine-generators and in reverse as electric motor driven pumps (Wikipedia, 2007). In a pump storage facility, surplus (or cheap) electricity can be used to power a pump which elevates water to a high reservoir, where it is stored until electrical power is needed (or expensive), at which time the water is released to drive an electricity generating turbine (Figure 11). This process involves energy losses but is still efficient enough to have been implemented in several places in the world including Australia, Wales, Switzerland, U.S., and Canada (Niagara Falls).

Water storage facilities, pumped or not, were originally designed to smooth fluctuations in electricity demand (or load). They can also be used to profit from the large price variability observed in modern deregulated electricity markets – water is stored (or pumped) when power prices are cheap, and released when power prices are expensive. Naturally, the net effect of such a price-driven operation will still be to smooth load. It is an interesting problem in optimal control to determine the optimal storage/pump and release protocols and in this way quantify the value of coupling wind and water.

Other innovative ways to store wind energy are also being investigated. One such approach is to use compressed air. An initiative by a group of Iowa cities in the U.S. aims to use wind-generated electricity to drive compressed air into underground aquifers. The air can then be released to generate electricity when needed (MSNBC, 2006).

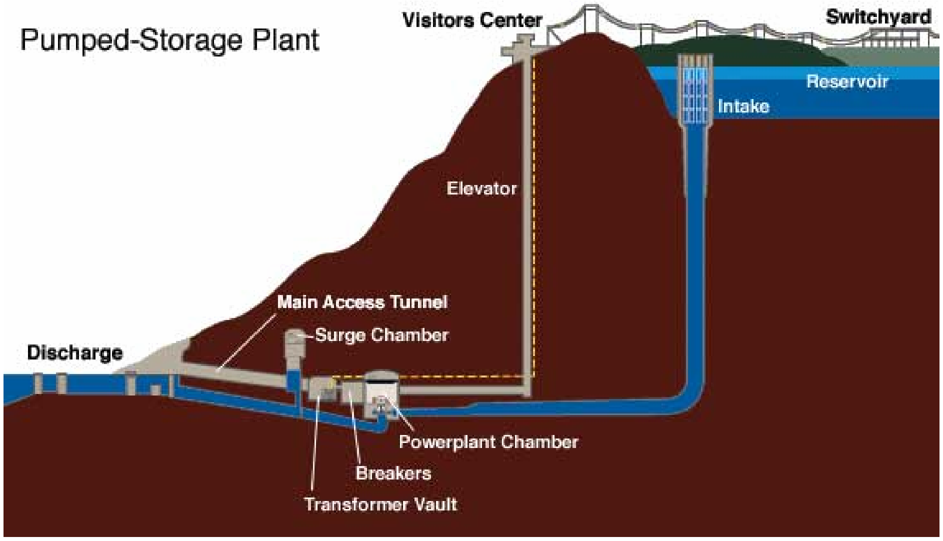

Figure 11: Diagram of a Pumped-Storage Plant (Tennessee Valley Authority)
Source http://www.tva.com/power/pumpstorart.htm

As it will be discussed in the valuation section, the presence of a pumped storage facility or a large reservoir can make an immense difference in the way wind variability is managed.

4. Valuation

The valuation approach covers decisions both at the operational/tactical as well as strategic levels. Operational decisions are required to effectively use the pump storage as a means of smoothing out short (daily, hourly) fluctuations in wind speed. Strategic decisions are based on a good understanding of the tactical value, so that proper investment decisions to build pumped storage facilities can be made.

4.1 Modelling the Decision Problem: Operational/Tactical

4.1.1. Value of Forecasts- the UK Experience

This section will discuss the value of an accurate wind forecast to the wind energy sector in the United Kingdom, in the absence of a pump storage. In a deregulated market system with both a "Day Ahead market" and an "Actual Real Time Market" structure (terms will vary), the value of both a hydro and a wind forecast can be quantified. Note that "Day Ahead" really denotes a forward market very close to delivery – the delivery might be only a few hours ahead. The system described below is similar to the UK "New Electricity Trading Arrangements" (NETA) system[21].

Here's how this "two market system" works. At time t-T a generator offers and amount of power A for time t delivery into the market. These offers are stacked and matched against the corresponding bids for power to determine at what price the market clears (the stacking process was described in some detail in the Literature Review Report). If the generator has offered power into the market below this market clearing price P, the offer is accepted and the generator is guaranteed the market clearing price for its supplied power. In the UK market, which clears every half hour, T was 3.5 hours until recent reforms reduced it to one hour.

At time t, the generator actually delivers power into the market. If the amount of power delivered is exactly A, the generator receives A times P for their power. On the other hand, if the generator delivers either more or less power (A*) into the market, it needs to buy (A-A*) units of power in the real time market (if A > A*) or sell (A*-A) units of power on the real time market (if A* > A). Let the real time market price be Pr. The generator then receives a cash payment of AP + (A* - A) Pr for the actual power generated.

In a situation where, for most generators, A and A* are usually identical or at least close, the real time market price will not be too different from the day ahead market price, on average. Still, let's consider St = (Pr – P)_t, the spread at time t between the real time and day ahead prices for time t power. It is fairly clear that St, like power prices themselves, will be skewed to positive numbers. The reason for this is that in the case of a large market outage, a great deal of power must be attracted into the market at the last minute, with a corresponding high price.

It should also be noted that there is some positive correlation between a unit under-delivering on its promised load and the spread. If the unit under-delivers, other units must make up the slack and they will require a payment for this. For a single small unit under-delivering a small amount, this correlation is probably not that important but, as we shall see, for wind turbines, the correlation will become more important.

---

[21] See IEA (2005) for a detailed discussion of this system.

For a wind turbine generator, which cannot be all that certain of the power it can deliver, this two market system encourages the wind turbine generator to underpromise. This reduces its downside risk of having to source power on the expensive real time market during a dead calm, at the cost of increasing its potential for "regret" – having to sell excess power into a depressed real time market.

The optimum balance between downside risk and upside regret will be found by a generator as a function of the market behaviour and of the meteorological conditions of its generation site, and of the amount of trust it feels it can place on its wind forecast. For instance, a Danish wind turbine operator might feel safe in bidding most of its expected power. The Danish electricity market is rather placid, linked as it is to the hydro dominated Norwegian market. At the same time, the wind in Denmark is fairly reliable. On the other hand, an Alberta wind generator might be more cautious, as the natural gas dominated Alberta electricity market has volatile pricing.

In this setting, the value of an improved forecast can be estimated by computing the additional expected risk-adjusted value to a generator of a bidding algorithm incorporating a better forecast. In order to build insight, suppose that a wind forecast could be obtained that was perfect on the time horizon T (which separated the day ahead market from the real time market). In this case, a generator could always bid the perfect amount of power into the day ahead market, as a result sidestepping the real time market entirely. It would be more complicated to determine the value of improved, though still imperfect, forecast information, as such a calculation would require optimal bidding rules given different endowments of forecast information.

If one was doing this calculation in an actual UK type market, with all data provided, the calculation would be fairly routine. This was done in an IEA publication (2005) which found values of £2.5/MWh to £5.7/Mwh (in British Pounds) for an improved wind forecast.

Strbac et al (2002) completed a detailed study titled 'Quantifying the System Costs of Additional Renewables' for the UK Department of Trade and Industry. This study assumed that 20% of the UK's generation capacity would come from renewables, including both wind power and biomass. This study recognized the fact that non-dispatchable wind and other renewables imposed additional costs on the system, the most important of which were the cost of providing the additional backup generation required to cope with long periods of calm winds and the cost in operational reserves (including what we term 'spinning reserve'). The other important cost they considered was due to the fact that many UK wind turbines are located far from population centres and therefore require additional transmission capacity which experiences line losses when used.

Building on this study, Auer et al. 2004, as referenced in (IEA 2005) considered the cost of having a 20% market share for wind power in the UK market. They did not distinguish between operational and capacity reserve costs, and were able to determine that these costs were between 4.5 and 6 euros/MWh, while transmission and distribution upgrade costs lay between 2.5 and 3 euros/MWh.

This 20% level of wind penetration is quite extreme by current Canadian standards. A third study, by Mott MacDonald (2004) estimated all costs of balancing wind for a variety of penetration levels. In British pounds/MWh, these costs were as follows:

Penetration (%)	Cost (£/MWh)
5.3%	Between 0.9 and 1.47 £/MWh
7.6%	Between 1.12 and 2.03 £/MWh
10%	Between 1.25 and 2.38 £/MWh
14.2%	Between 1.47 and 2.59 £/MWh

Table 3: Cost of wind integration for a variety of penetration levels

These figures can be used to build some insight into the possible range of wind integration costs in Canada. For instance, since the Alberta market is about 10,000 MW, using the exchange rate of C$2 = 1GBP and 10% penetration level, 1000 MW of wind in the province of Alberta would impose additional annual system costs of about $35 million[22]. This large cost, which comes largely from the requirement for backup costs, is definitely a motivation for the introduction of pump storage facilities.

4.1.2. Value of Forecasts- BPA Experience

This section will discuss the upper bound for the value of EP for wind energy when the generator can either store water, or team up with a hydro partner. In the U.S. Pacific Northwest, the ability to store water behind hydro dams is very large (Hamlet Huppert and Lettenmaier (2002), as cited in the Hydrology case study).

If wind turbine variability is buffered with water storage, the problem of under-delivering and getting caught short (as in the UK example) can disappear. Instead, the wind operator will simply replace the lost wind generation with an equivalent hydro generation from stored water. Later on, wind power can be used to substitute hydro power to "repay" the loan. The only losses will be draining water at low prices when higher prices are paid to pump that water. For this reason, the Bonneville Power Administration (BPA) has a scheme where wind power operators can "deposit" wind power behind their reservoirs (by replacing hydro generation with wind generation), for later withdrawal of hydro power instead of wind power for a cost of US$4.50/MWh (IEA 2005).

Such a system can provide nearly infinite storage capacity depending on the geographic environment, in which case the value of a wind forecast is limited. For instance, in the US Pacific Northwest where the BPA "cleans" wind power for $4.50/MWh (IEA, 2005), the value of accurate wind forecasts can be no more than this amount.

4.1.3. Valuing The Role of Pump Storage in Buffering Wind Non-dispatchability

Finally, as discussed in the Hydrological EP Case Study, Hobbs et al. (1999) conclude that a 1% reduction in the mean absolute percentage error (MAPE) of load forecasts decreases variable generation costs by 0.1% - 0.3% when MAPE is in the range of 3-5%. The reasons for these savings include better fuel allocation, reduced need for

---

[22] $4/MWh*1000 MW* 8760 hours/year ~ $35 million per year

spinning reserve, reduced chances of load interruption and better short-term hydropower scheduling.

The idea here is to consider the wind forecast error to be lumped in with the load forecast error in a "net systems availability" setting. Clearly, buffering the wind error with hydro would dramatically decrease the MAPE of such a systems availability forecast. Indeed, reservoir rich hydro based systems don't really need spinning reserve and rarely have system outages. However, it is the buffering with hydro pump storage that has value in this case, not the capability of accurate short term wind forecasts.

4.1.4. Some public policy implications

It is important to note that the value of EP for wind energy will largely depend on which energy player is using it in which jurisdiction. As part of renewable energy incentives, many jurisdictions guarantee a single (and very high) price for every MWh of wind energy generated (such as the Standard Offer program in Ontario[23]). However, the unintended negative consequence of such programs is that it discourages all attempts to smooth wind power output, since at a constant high price, the wind energy generators are better off generating more power, even at a less appropriate time, than less power at a more appropriate time.

In other parts of the globe, there are other examples: for instance, in Australia, the incentive is based on a "market price + x" model, and motivates the wind energy generator to be more proactive (Kucera 2007). In certain other Canadian provinces, such as Alberta, wind power is treated just like any other generation type without any special incentives, which also increases the need for wind forecasting by the individual generators.

In a two-market system (day-ahead vs. real-time), all of the forecast risk is placed on the wind turbine operator, with no residual forecast risk remaining with the Independent Electrical System Operator (IESO). In the Ontario system described above, the tactical value of a wind forecast to a wind energy generator is nil. In this regime, the wind generators bear no forecast risk, but this forecast risk remains, in this case with the IESO. Therefore, even in this regime the wind forecast still has tactical value, in this case to the IESO.

4.2. Buffering Wind Turbine Variability with Pump Storage: A Strategic Case Study

As discussed in Section 3, a pump storage facility works by converting the potential energy in elevated water to kinetic energy. For a detailed review of pump storage facilities, see Thompson, Davison and Rasmussen (2004) and Zhao and Davison (2007).

---

[23] OPA states that "small generators (under 10MW) using selected renewable resources like solar, wind, small hydro and some biomass, will have the opportunity to sign a 20-year contract to sell power to the OPA and receive a guaranteed price per kilowatt hour for the energy delivered into the Grid, over the life of the contract."

A remarkably large pump storage plant is required to store a relatively small amount of energy. As discussed in Appendix 1, the common unit of electricity used in the electrical power industry (and by electricity traders) is the curious hybrid unit of Megawatt-hours (MWh). Since a Watt is a Joule per second, one MW-h is 3.6 gJ (billion Joules). In contrast, the potential energy involved in raising 1 kg of mass 1 meter, at the surface of the Earth, is a puny 9.8 Joules (rounded to 10 in these calculations). Thus, in order to store even just a single MWh of electricity in a pump storage facility for which the upper reservoir is 50 meters above the lower reservoir (consider that Niagara falls has a drop of about 50m), we need a water storage capacity of 7.2 million L of water (the mass of 1 L of water is 1kg). As there are 1,000 L in a cubic metre, the storage capacity is 7,200 cubic meters, much like a shallow 2m pan of water 60 metres by 60 metres.

The round trip efficiency of pump units is usually between 70% and 80%. Corresponding to the best-case 80% round trip efficiency is a one way efficiency of about 90% (0.9x0.9 = 0.81) so this pump storage facility would actually allow about 1.1 MWh of electricity to be exchanged for 0.9 MWh of electricity at some later time. As a rough order of magnitude, a MWh of electricity has a wholesale value of between $30 and $100 (sometimes even $1000) in the Alberta market, so if this facility could be used each day to purchase electricity at $30 and sell it at $100, the annual income from the facility would be $20,805, not a very promising amount especially considering the limitations imposed by our assumptions [24].

It can therefore be seen that pump storage, while mechanically rather efficient, is economically less so, as the cost of constructing such a facility will be high. As such, the optimal size of such a facility is of interest. We want to construct a facility which is large enough to store a useful amount of energy, but not so large as to remain empty most of the time.

We begin with an attempt to transform the energy generated by wind turbines into “always on” baseload. We use wind data recorded at the Brockett ADGM in Alberta between June 15, 2005 and June 14, 2006 as our base case. We assume that the power output for the day can be obtained by transforming the average wind speed for the day, adjusted to estimate the 70m elevation wind speed, and by using the power curve of a Siemens 1MW turbine[25]. From this process, we can derive a power measurement in kilowatts which we transform to an energy measurement in MWh by multiplying by the 24 hours in the day. Note that the average capacity factor in this case is 43%.

This gives us a sequence of daily energy outputs as seen in Figure 12:

---

[24] (0.9*100-1.1*30)*365 = $57*365 = $20,805.

[25] Height adjustments for wind were made using a wind shear coefficient of 0.14 in the RETScreen analysis tool.

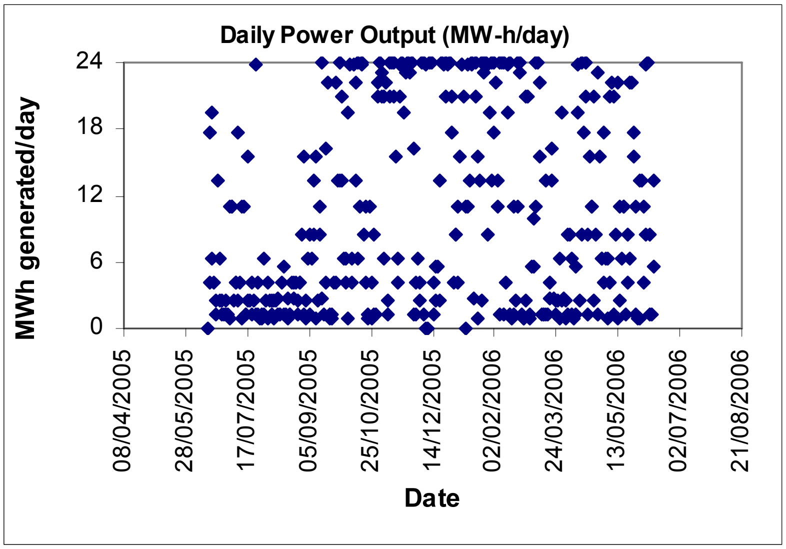

Figure 12: Daily energy generated by a hypothetical 1MW Siemens turbine near Brocket Alberta

It can be seen that the power delivered by this turbine is extremely variable. It should be noted that the data used here somewhat overstates the variability. In fact, wind power does benefit to a certain degree from geographical diversification and from the ability to "wheel" power over long distances including across different markets. Nonetheless, a detailed (but unpublished) study conducted by Dr. Hans Tuenter from OPG suggests that geographical diversification is by no means a panacea for the Ontario wind sector, as wind speeds are fairly strongly correlated over relatively long distances. This qualitative conclusion is supported by other studies, as discussed, for instance in Boland (2007).

As a case study, what if we wanted to convert this turbine into a constant baseload type generator, how much energy would we need to be able to store? To see this, suppose that each day we deposited the amount of energy generated into a bank, and simultaneously withdrew the average annual energy generated from the bank (for this order of magnitude exercise we neglect the fact that these "deposits" and "withdrawals" would, if done via a pump storage facility, come with some heavy transaction costs). Figure 13 illustrates the resulting energy balance.

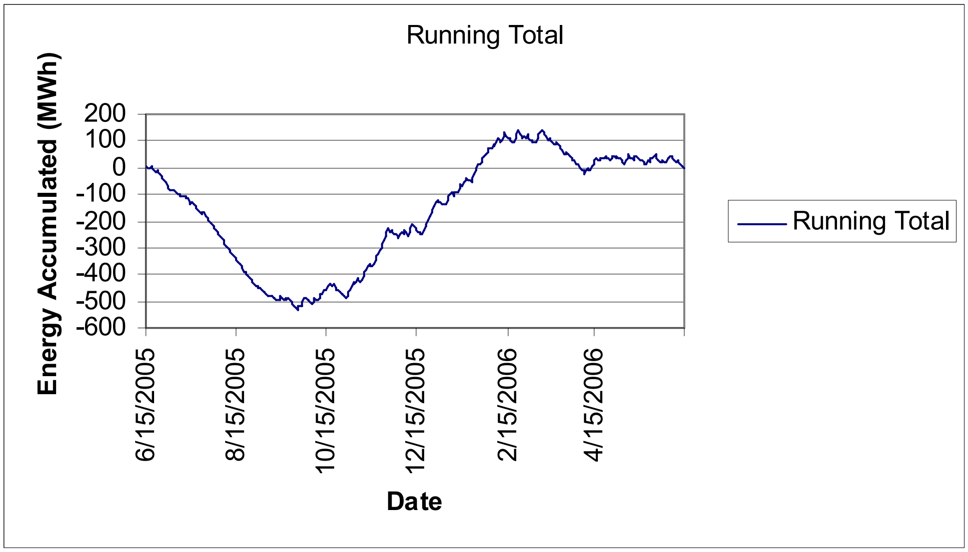

Figure 13: Net accumulation using variable input, constant average output

One measure of how much energy we need to save here is the maximum swing experience in this bank account, here roughly 600MWh. However, saving and releasing is not perfectly efficient in our bank, so even if an efficient 80% round trip wind storage turbine is considered, we would need to have the capacity to store about 750 MWh in our facility to convert the output of a single, highly variable, 1MW nameplate capacity wind turbine into a 430 kW[26] baseload generator.

750MWh is 3.6 million megawatt-seconds, or 3.6 million Megajoules. One cubic metre of water stored at an elevation of 50m has a potential energy of 0.5 megajoules, therefore we need to store 5.4 million cubic metres of water to store 750 MWh equivalent power. If we store this in a pond 10m deep, the area of the pond needs to be 540,000 square meters, which is 0.72 km^2 or 54 hectares. It is important to note that this facility would be big enough to buffer the wind output between June 15 2005 and June 14 2006 for a *single* wind turbine.

It should also be noted that not all geographical locations have the required difference in land elevation to make the construction of pump storage facilities possible, and in extremely dry areas with little precipitation, issues of evaporation may further degrade the potential for pump storage to help.

In order to design this storage facility properly, the performance over many more years of historical data as well as statistical analysis of the existing data would need to be considered. This could be done by bootstrapping from existing data. However, the correlation coefficient between daily average wind speed on successive days is, for this dataset, observed to take the fairly high value of 48%, (see the scatter plot of Figure 14).

---

[26] Average capacity factor x 1MW nameplate capacity

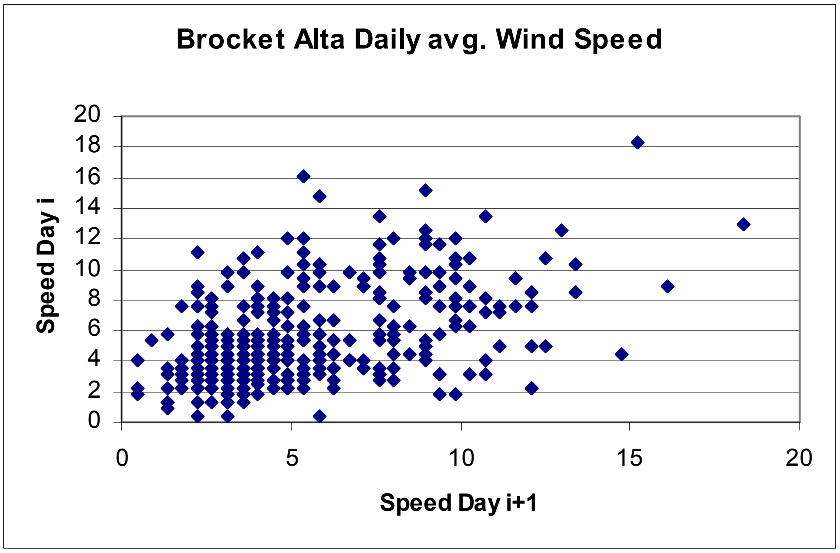

Figure 14: Scatter plot of met mast level wind speed measured on successive days between June 15 2005 and June 14 2006 at Brocket Alberta
Note: The value of the correlation coefficient corresponding to this scatter plot is 48%.

The fair degree of correlation between daily wind levels, together with the massive swing in the fill level data of Figure 13, suggests that this facility may be nearly large enough to buffer most possible years of wind data. Perhaps, following standard engineering practice, we will just build it a little bigger by adding a "safety factor" of 2. So our pump storage facility will have a 100 Hectare (1 km^2) artificial lake[27] as an upper reservoir, which can be filled up to a maximum height of 10m.

4.2.1. Valuing the base load Pump Storage Facility

What would the annual receipts be from this huge storage facility, together with the wind turbine that drives it? Over the same time period of June 15 2005 to June 14 2006 the average power price in Alberta was $72/MWh. As seen in Figure 15, this average annual power price conceals a great deal of variability even in daily average prices – an hourly price plot would be yet more variable. This "buffered" facility would generate 10.4 MWh per day, for 365 days, at $72/MWh for a total annual value of about $272,000. This is sufficient cash flow to finance the construction of such a massive facility which, according to a table published in a BPA report (Bonneville Power Administration, 2006, p. 79) would involve capital costs in the neighbourhood of $1 million.

---

[27] A better solution could be to identify an existing lake to fill this role if the necessary environmental permissions could be obtained.

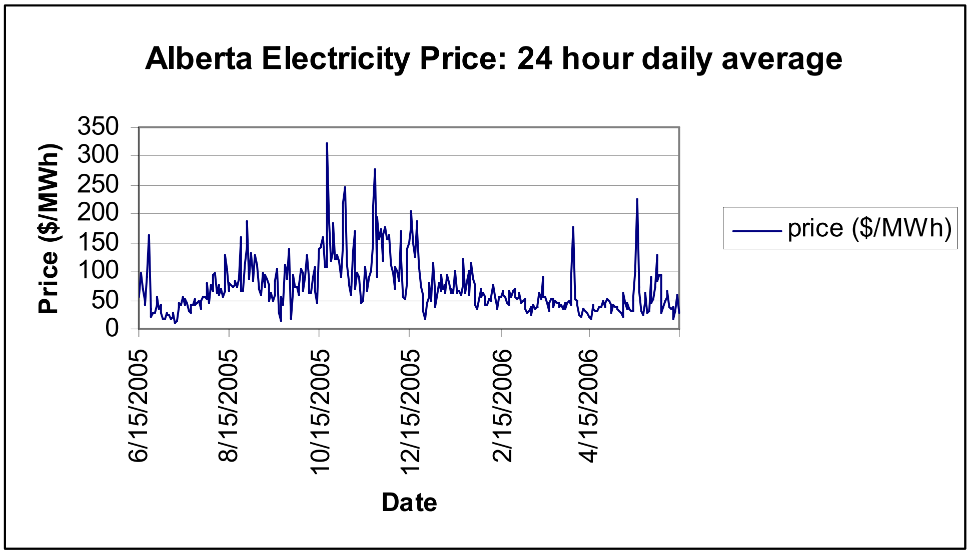

Figure 15: Alberta daily average power price; June 15 2005 – June 14 2006
Source: Data from AESO as stored by Matt Davison

While converting the intermittent production of wind power to a constant baseload is quite a nice trick, it has to be recognized that society's use of electrical power is quite variable in predictable ways – the so-called "daily load shape" as discussed in the Literature Review Report. Market electricity prices reflect this fact, with the average price recorded on every day (this time only between the hours of 7AM and 11PM) for the same period, being $88. This value is also quite variable as shown in Figure 16. Indeed, if the relatively low-load weekends are removed from this average, the resulting average power price over the same window was $95.

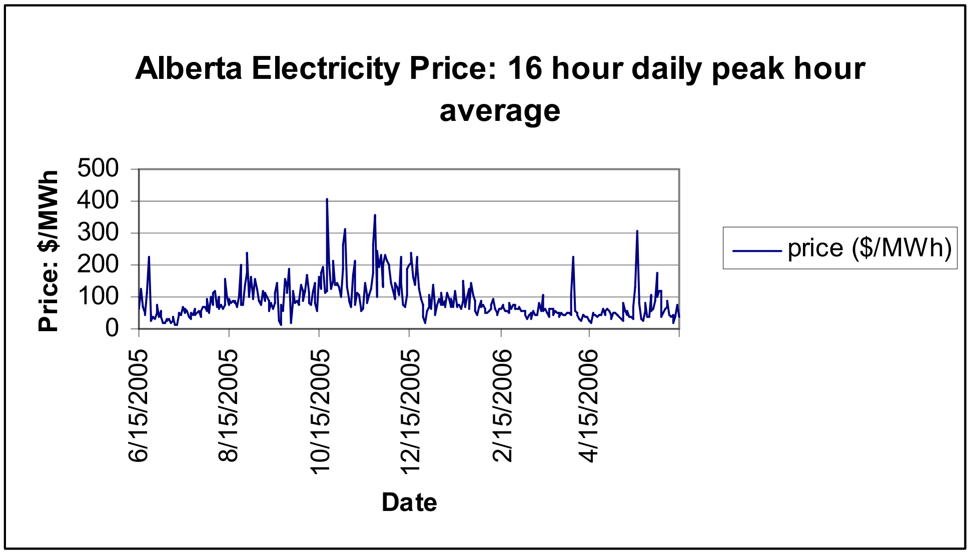

Figure 16: Alberta daily on-peak (7AM – 11PM) average power price; June 15 2005 – June 14 2006
Data source: From AESO as stored by Matt Davison

Based on this idea of "selective availability", we present the design for a pump storage facility, incorporating efficiency losses, for converting wind turbine energy into a generator that reliably gives a given unit of power during the on-peak, weekday, hours of each week (80 hours/week in total). The design process behind this will be reviewed in the next section.

4.2.2. Designing an On Peak Price Facility

For this design, we plot the fill level required in the notional 100 hectare lake located 50m above a reference plane designed above, to convert the June 15 2005 – June 14 2006 energy output of our hypothetical 1MW wind turbine to a steady on-peak load. We consider the exact amount of this on-peak load to be a design variable. If we construct a reservoir with minimum fill level of 0 which was filled to a height of 4m on June 15 2005, and maximum fill level of 6m, and if we withdraw energy at a rate of 12.5MWh/16 hour on peak day, or a 781 kW generator that dependably runs every Monday – Friday, 7AM – 11PM, and if we assume a one-way efficiency of 85% for both the turbine and the pump (or a round-trip storage efficiency of $(0.85)^2$ = 72%, the water level in the pump storage facility during the course of the year is as follows (Figure 17):

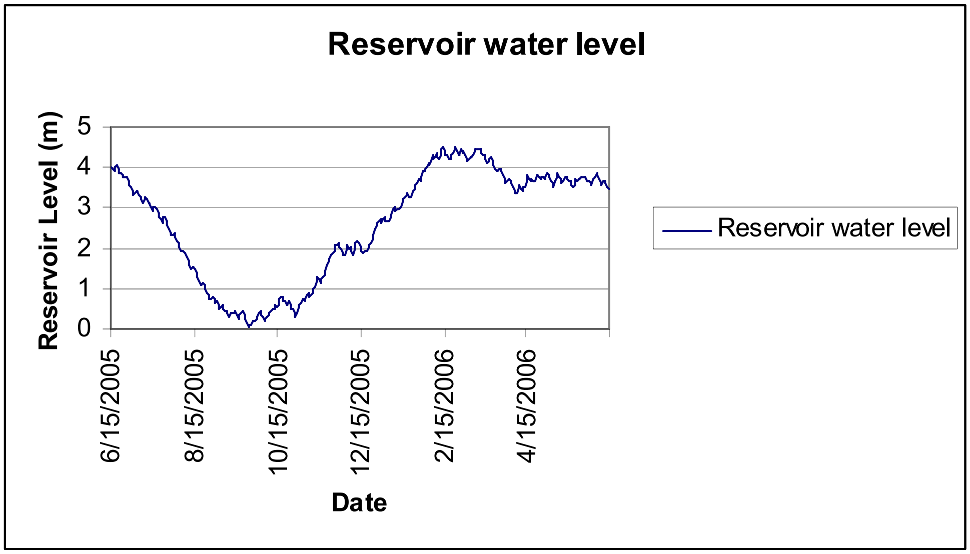

Figure 17: Fill level in 100 Ha reservoir 50m above pump and turbine as a function of time.

Note: Calculations are done assuming that turbine and pump efficiencies are 85%. Energy input is from a hypothetical 1MW Siemens wind turbine located near Brockett Alberta, while energy is generated from the facility at a rate of 781kW during the 80 on-peak hours of each week. Initial fill level of reservoir is 4m and maximum fill level (not binding here) is 6m.

4.2.3. Valuing the On-Peak Pumped Storage Facility

The value of this facility is 260 days x 12.5 MWh-day times average on-peak Alberta price of $95 per day, which corresponds to $308,300. However, the final level of the reservoir is about 60cm less than the initial value, which contributed $7,120 of this amount, but which should really be reckoned as drawing down on capital. So the total value of this facility is $301,170 in the 2005-2006 year considered above. Note that this is more than the $272,000 estimate for the "baseload" facility described above.

4.2.4. Tactical and Operational Extensions of this modelling Exercise

This analysis suggests that, with good statistical knowledge of both wind and price patterns, even more advantageous contracts could be created from this pumped storage facility, leading to the tactical use of wind power EP.

One way to begin creating such a contract was to build from the idea of "swing options" commonly used in energy trading. Swing options work by providing for some initially constant delivery of energy at different time scales (for instance per on-peak hour). They grant their holder a certain number of rights to "swing up" (increase the amount delivered) or to "swing down" (decrease the amount of energy delivered). The flexibility granted by such a contract might be of great value to a wind-water storage operator who is able to predict the future wind even over just a few hours or, probabilistically, days. Sale of such contracts could help finance the pumped storage facility and the associated wind turbines.

At the operational level, if the wind-water system reservoir indicator is close to "empty", wind forecast information could be used to determine whether it is better to default on the promise to deliver power today, when the market price of wind was cheap, in the assurance that there would be no wind tomorrow, or whether to not default on the contract in the certainty that tomorrow's wind would begin to refill the reservoir.

Once the specifications of such a facility is known, we then need to decide the optimal way to operate it, both with and without a forecast. As the optimal operation strategy is itself dependent on the type of uncertainty, ideally this analysis would take the form of a dynamic program with two sources of uncertainty (wind speed and electricity price) having two control variables (amount of wind energy to buffer into the pump storage facility and amount of water to release from the facility). Completing such an analysis is feasible but is beyond the scope of this project.

5. References

Blackwell, R., "Alberta Opens Gate to Wind Power", The Globe and Mail, Sept. 27, 2007.

Boland, J. "Windfarm Output Variability", presentation made at AMSI-MASCOS workshop on Electricity Markets, Surfer's Paradise, Australia, April 2007, available at http://www.amsi.org.au/images/electricity/pdfs/Boland.pdf

Bonneville Power Administration. Renewable Energy Technology Roadmap: Wind, Ocean Wave, In-Stream Tidal and Solar Photovoltaic. Technology Innovation Office, September 2006.

Brown, S., "New Brunswick Wind Energy Integration Overview", presentation made at the 23rd Annual Conference of the Canadian Wind Energy Association, Québec City, 2007

Environment Canada, "AnemoScope Web Site", http://anemoscope.ca/Intro_en.html, accessed Oct 8, 2007

Frost, W., "Market and Operational Framework Update", Alberta Electricity System Operator, presentation made at the 23rd Annual Conference of the Canadian Wind Energy Association, Québec City, 2007

GE Energy, 1.5 MW Series Wind Turbine technical documentation, available at http://www.gepower.com/prod_serv/products/wind_turbines/en/downloads/ge_15_brochure.pdf , 2007

GWEC - Global Wind Energy Council, "Global Wind 2006 Report", Brussels, Belgium, on-line report available at http://www.gwec.net/fileadmin/documents/Publications/gwec-2006_final_01.pdf

Hobbs, B.F., Jitprapaikulsarn, S., Konda S., Chankong, V., Loparo, K and Dominic J Maratukulam, "Analysis of the Value for Unit Commitment of Improved Load Forecasts", IEEE Transactions on Power Systems 14(4), 1342- 1348. November 1999

IEA, "Variability of Wind Power and other Renewables: Management Options and Strategies", International Energy Agency, Paris France, June 2005.

Kucera, A. (2007), holds a Ph.D. in mathematics and works for Integral Energy, a large Australian power utility, where he is the head of the risk management department, personal communication, 2007.

Kwas, T., Stakeholder Relations Manager, TransAlta Wind, personal communication,2007
Mathew, S., Wind Energy Fundamentals, Resource Analysis And Economics, Springer, Berlin Heidelberg, 2006

Mott MacDonald (2004), Intermittency Literature Survey, Annex 4 to the Carbon Trust and DTI Renewables Network Impact Study, available online at www.thecarbontrust.co.uk/carbontrust/about/publications/Annex4.pdf

MSNBC "Store wind power for later use? Cities bet on it: Iowa project would compress air in underground caverns", Jan 4, 2006

Roberge, G., Bureau de Recherche, Hydro Québec, personal communication, 2007

G. Strbac and ILEX Consulting (2002). "Quantifying the System Costs of Additional Renewables in 2020", London: Department of Trade and Industry, available online at www.dti.gov.uk/energy/developep/080scar_report_v2_0.pdf

Thompson, M., Davison, M. and H. Rasmussen, "Valuation and Optimal Control of Electrical Power Plants in Deregulated Markets", Operations Research 52(4) 546-562, 2004

Wikipedia, “Pumped-storage hydroelectricity”, http://en.wikipedia.org/wiki/Pumped-storage_hydroelectricity, site accessed Oct. 6th, 2007

Appendix 1

A Joule (J) is the basic SI unit of energy. It is defined in terms of force, but that definition is not very useful for this case study. From an energy perspective, it is the amount of energy required to raise one gram of liquid water about 0.1 Celsius degree.

In potential energy terms the potential energy possessed by one kg of mass raised 1 meter is a meager 9.8 Joules. However, this is not the unit of energy used in electricity markets. The SI unit of power is a Watt -- one Joule per second. This isn't much power on the scale of a power grid (a single lightbulb draws 40-60 Watts). So instead, the unit of a megawatt, or million watts, is used.

The energy market energy version is a curious hybrid unit -- the megawatt-hour (MWh). This is the amount of energy obtained by running a one MW facilty for one hour (or 3600 seconds). This corresponds to 3.6 giga (billion) joules.

Appendix 2: Decision Tree (Operational/Tactical)

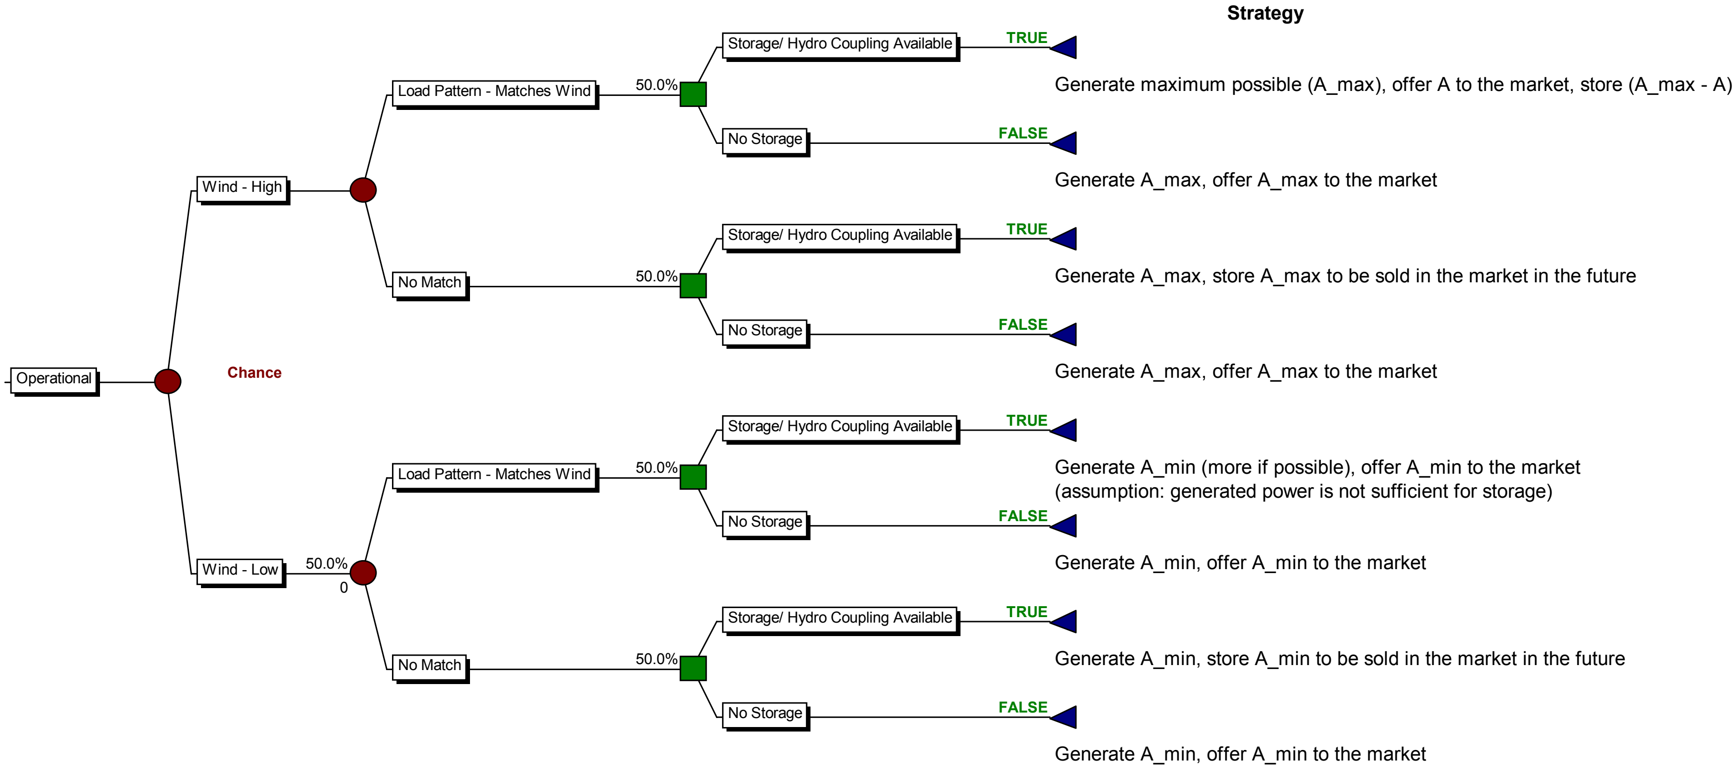

Generate maximum possible (A_max), offer A to the market, store (A_max - A)

Generate A_max, offer A_max to the market

Generate A_max, store A_max to be sold in the market in the future

Generate A_max, offer A_max to the market

Generate A_min (more if possible), offer A_min to the market
(assumption: generated power is not sufficient for storage)

Generate A_min, offer A_min to the market

Generate A_min, store A_min to be sold in the market in the future

Generate A_min, offer A_min to the market

Note: Circles represent chance events, squares represent decision points, and the triangles represent the payoff.

Environmental Predictions and the Energy Sector: A Canadian Perspective

Case Study 4:
Earth Observation for Renewable Energy Projects

Prepared for

Environment Canada

Contract Number: K3A40-06-0028

Prepared by

TURQUOISE
TECHNOLOGY SOLUTIONS INC.

DYDEX
DYDEX RESEARCH AND CAPITAL LTD.

1. Introduction

As discussed in the Literature Review report, renewable energy sources and technologies play a modest role in the global energy mix. With the exception of hydro, renewable energy sources provide only 2% of the world's total energy supply for electricity generation (IEA, 2006). This figure includes almost all types of renewable energy sources, including geothermal, solar, wind, marine (wave and tidal), heat as well as combustible renewables and waste. In Canada, the percentage of electricity generated from renewables (other than hydroelectricity) is less than the global average, and is currently 1.4% of the total (McCarthy, 2007).

European countries have been systematically investing in renewable energy systems in the last decade, and they are beginning to reap some significant rewards. Their investment was not solely in the development of new renewable energy systems, but also in finding innovative solutions for resource assessment, design, grid integration and operational monitoring of these systems using multiple streams of environmental data. These streams include in-situ as well as remotely sensed data from airborne and space-based platforms. Today, the expertise in renewable resource assessment and management is helping European companies export their services and gain a significant market share in the renewable energy market outside Europe.

This case study aims to document applications of EP for various types of renewable energy systems, including off-shore wind and solar energy. Additionally, a valuation framework is presented which demonstrates the economic benefit of using satellite data for wind energy resource assessment.

Widespread adoption of renewable energy systems is held back by a number of factors: cost, reliability and continuity of service. Given that most renewable energy sources are based on natural processes which fluctuate over different time scales, the expected energy yield varies considerably, both daily and seasonally. Therefore, being able to predict these fluctuations and developing the necessary mitigation mechanisms are critical to ensure overall system security and reliability. There are numerous ways to achieve this objective particularly where satellite-powered EP can play a key role, as discussed in the rest of this case study.

1.1. GEOSS

In February 2005, ministers from nearly 60 countries, including Canada, endorsed the 10-year implementation plan for the Global Earth Observation System of Systems (GEOSS). This international cooperation framework is designed to enable sharing timely, high-quality and long-term global information for better decision-making. Moreover, it envisions the integration of data obtained from many different instruments (satellites, aircraft and in-situ). Space-based Earth Observation plays a key role in this mix, since the majority of information is expected to be provided by satellites (Lautenbacher, 2006).

At least three of the nine societal benefit areas identified for GEOSS are related to the Canadian energy industry:

- Improving management of energy resources
- Understanding, assessing, predicting, mitigating, and adapting to climate variability and change
- Improving weather information, forecasting and warning

In order to implement GEOSS for energy applications, an "Energy Community of Practice" (ECP) has also been started[28]. The objective of ECP is to support GEOSS activities related to the application of Earth Observation data for the energy industry. Areas covered under ECP include: siting of power plants and facilities taking into account environmental and sociological issues; optimized design of power systems and facilities; yield estimation and resource monitoring based on historic information; yield forecasts based on near real-time weather forecasting; operation and management of power plants, including automatic failure detection; and trading and monitoring of emissions credits.

1.2. Earth Observation Market Development Programme

The European Space Agency (ESA), initiated the "Earth Observation Market Development Programme" in 2000 in order to stimulate the operational use of Earth Observation in different economic sectors. One of the main thrusts of this initiative is applications in the renewable energy sector. Specifically, the EOMD programme supports demonstration projects which enable the partnership of smaller companies specialized in Earth Observation with larger downstream companies. A number of demonstration projects have targeted the needs of the renewable energy sector with a particular focus on solar, wind and hydroelectricity (Mathieu, 2005).

2. Earth Observation and Renewable Energy

The data for EP applications for renewable energy projects can come from multiple platforms, including in-situ, airborne and space-based sensors. This section will discuss how satellite data can be used to better manage renewable energy projects, and the types of advantages satellite data can offer, when used together with other sources of data for the development of EP-related products and services. The focus will be on the role of EP for solar energy, off-shore wind energy and other applications including hydroelectricity and sea ice monitoring.

2.1. Satellite Data Sources for EP

Satellites are placed into various orbits around the Earth, and depending on their orbital characteristics and the on-board instruments, they can provide different types of EP-related data. Generally speaking, there are three orbital "belts" where satellites are located: Low Earth Orbit (LEO), Middle Earth Orbit (MEO) and Geosynchronous Earth Orbit (GEO). LEO satellites generally have an altitude of less than 1000 km from the Earth's surface, and when placed in polar orbits, their footprint can cover the whole globe in a matter of days. GEO satellites, on the other hand, enjoy the advantage of a

---

[28] http://www.geoss-ecp.org/

"stationary" point in the sky (for observers from the Earth), since the orbital period at the altitude of approximately 36,000 km matches the period of the Earth's rotation.

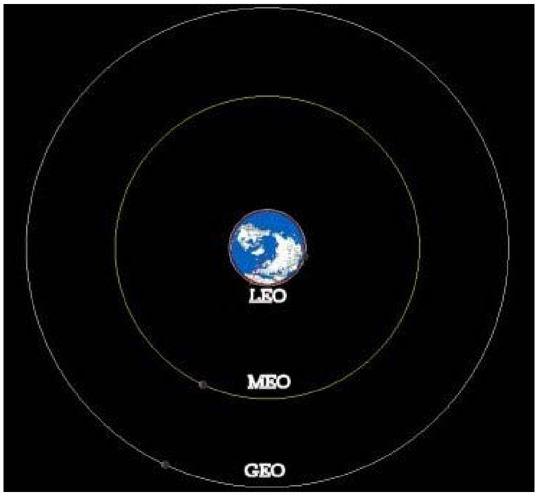

Figure 18: Relative position of LEO, MEO and GEO
Source: Gurtuna and Trepanier (2003)

The characteristics of LEO and GEO make them particularly useful for Earth Observation satellites. At GEO there are many weather satellites, such as Europe's Meteosat series, the United States' GOES series, and China's Feng-Yun series. All of these satellites are carrying instruments that monitor large swaths of the globe. At LEO, instruments on polar orbiting weather satellites have much higher spatial resolution than those on GEO satellites, since they are much closer to the Earth. Some examples of LEO satellites include Canada's Radarsat-1, NOAA series of the US, Europe's ENVISAT and China's Feng-Yun (polar) satellites.

The advantages and disadvantages of LEO and GEO satellites are summarized in Table 4.

	Advantages	**Disadvantages**
GEO	• Satellite is "stationary" and therefore visible from a third of Earth's surface • Instruments can be used to monitor a given spot on the Earth's surface continuously	• Polar regions cannot be observed • Lower spatial resolution due to the distance from the Earth's surface
LEO (Polar)	• Global coverage • Increased spatial resolution • All weather and night time observation capabilities in the case of satellites with active instruments, such as radar	• Continuous observation of a given spot is not possible • The satellite can observe some spots on Earth only twice per day for a limited amount of time during each pass

Table 4: Advantages and disadvantages of GEO and LEO

For EP applications in Canada, the global coverage of polar orbiting satellites makes it feasible to acquire data from a large number of LEO satellites. However, for GEO weather satellites, geographical coverage is a limiting factor. For instance, Meteosat satellites do not cover most of North America, and therefore, data and tools developed for EP applications using these satellites cannot directly be used in Canada[29].

2.2. Solar Energy

As in the Wind/Water Case Study, EP can be particularly useful for resource assessment and forecasting applications in the solar energy industry. Resource assessment is a very important prerequisite for making good siting decisions for all types of solar energy projects, including solar photovoltaic and solar thermal projects. Understanding the variability of solar energy over time is also possible by using data from historical time series.

Solar energy is following the path of wind energy worldwide and the installed capacity of both solar PV and solar thermal plants is rapidly increasing. Although North America has been lagging behind Europe and Japan, momentum is building, especially in the U.S. In June 2007, a 64 MW solar thermal power plant was commissioned in Nevada, the first plant of this type in North America since the 1980s (Economist, 2007).

In Europe, especially in Germany and Spain, a number of R&D efforts are underway to use satellite data to support solar energy projects. As part of this case study, a series of interviews were conducted with researchers from Spain's Centro de Investigaciones Energéticas Medioambientales y Tecnológicas (CIEMAT), a public research institute. One of the main activities of CIEMAT, through the Renewable Energy Division and Solar Platform of Almeria Division, is to work in cooperation with Spanish industry and universities to solve key scientific and technological problems of the solar industry in Spain.

Specifically, CIEMAT provides various analyses to the solar energy industry in Spain using satellite-derived information. Some research groups look into EP for various applications such as pollution monitoring, and there is a dedicated research group focusing exclusively on solar radiation measurements from satellites.

It is notable that most of the interest from Spanish energy players in this kind of analysis is due to the requirements of the banks who provide financing to solar energy projects. As part of their own loan assessments, banks require studies which incorporate satellite information.

Other sources of demand include large-scale solar energy projects, such as solar power towers in the southern part of Spain (e.g., Solar Tres Power Tower in Andalusia , PS10 near Seville and Andasol 1 in Andalusia). Industrial consortia which work on these plants also work with CIEMAT to obtain solar radiation assessments derived from satellites.

---

[29] However, given the interconnected nature of weather systems, it should be noted that data from other parts of the world, such as the Indian Ocean and Western Pacific, can be very valuable for Canada.

These companies include SENER Grupo de Ingeniería, Grupo Cobra, Solucar Energía (a subsidiary of the Abengoa Group) and Iberdrola (the largest wind energy producer in the world, and one of the top renewable energy generators in the world).

Beyond these two main sources of demand for EP (banks and large-scale solar power system integrators and operators), the rest of the demand is very much policy-driven. In most cases, government regulations and legal enforcement have to be in place to support EP-related services. For example, solar heating is now mandatory for certain buildings in Spain, yet architects and construction companies are designing their solar heating systems based on solar maps which are outdated and provide poor information. However, since updating such maps is not mandatory, not much work is being done for updating the solar radiation measurements which can actually help architects and construction companies to improve the efficiency of their designs. One lesson learned from this observation is that in order to realize the full potential of EP, demonstrating its benefits to various stakeholders and obtaining policy-driven incentives as well as sanctions may be necessary.

2.2.1. Sources of Satellite-derived Information

CIEMAT obtains satellite imagery from the European Organisation for the Exploitation of Meteorological Satellites (EUMETSAT). The main purpose of EUMETSAT is to deliver weather and climate-related satellite data, images and products to the National Meteorological Services of the organisation's 20 Member and 10 Cooperating States in Europe, as well as other users world-wide.

CIEMAT receives the imagery using its own satellite antenna for downlink from various meteorological satellites operated by EUMETSAT. In order to obtain this access, Spain’s National Meteorological Institute pays an annual licensing fee to EUMETSAT. As part of the licensing agreement, the high-resolution imagery obtained by CIEMAT can only be used for its own research and it cannot be distributed to third parties. The pricing of imagery for use by the private sector would be much higher and result in higher costs for EP services. However, as long as the cost-benefit analysis is deemed favorable by industry, there will be demand for satellite-derived solar resource assessments in Spain.

2.2.2. Uses of Satellite-Derived Information

Currently, CIEMAT provides its services through a cooperation and technical assistance model with the private sector. The private sector indicates its analysis needs, and provides a set of inputs to CIEMAT, which conducts solar resource assessment. The location is also an input given by the industry (in other words, the analysis of CIEMAT is not used to make strategic siting decisions, but to analyze a given geographical location). Siting decisions are impacted by a number of factors including political interest, current status of roads, electricity grids, etc.

CIEMAT provides a 1-year assistance service to industry by providing solar radiation information and satellite-derived measurements. Participating industry partners help validate some of the findings by investing in in-situ instrumentation. CIEMAT does not receive any payment from industry for this service, and an annual budget allocation of approximately 60,000 euros is used for the resource assessment of six different locations. These resource assessment services are not priced at their actual market value, and they are provided within a government-industry partnership framework.

Such reports are used mainly for the initial years in the design and construction of a solar plant for dimensioning and other design decisions. During operations, the solar energy generators generally install their own radiometric stations in order to acquire very precise, real-time radiation data. However, such instruments cannot be readily used for forecasts, since they are not forward-looking such as meteorological models based on satellite information.

CIEMAT researchers indicated that to their knowledge, there is currently no private sector provider of such services in Spain, but CIEMAT is currently working on a spin-off company concept together with Universidad Autónoma de Madrid. In Europe, other private sector actors are located mainly in Germany. Meteocontrol GmbH and Solemi are two private sector providers of such services.

Starting from January 1st, 2005, wind farms under the fixed tariff regime in Spain are required by law to provide their wind energy output forecasts no later than 30 hours prior to the start of each day. These forecasts can later be updated up to one hour before predetermined periods during the day. For wind farms participating in the electricity markets, the requirement is different: they have to provide their day-ahead forecast before the day-ahead market closes at 10am. These forecasts can later be updated as well. (GH, 2006)

Solar industry will most likely follow the path of wind energy and such information may be required from solar energy generators within the next 2-3 years in Spain. Therefore, in terms of the forecasting horizon, CIEMAT is mostly interested in 24-, 48- and 72-hour forecasts with 5-day forecasts also seen as a long-term target.

2.2.3. Advantages of Satellite-derived Information

Satellite-derived information has some obvious advantages over other methods for solar radiation assessment in Spain. Currently, there are only 30 radiometric stations in the country which can provide similar information. However, these stations do not provide the necessary spatial resolution required for solar radiation assessments. Archived satellite data can provide a 10km spatial resolution over the last 12 years. This enables advanced statistical analyses (e.g., using neural networks, wavelet methods) and reports which are much more accurate than those which are based on interpolating in-situ measurements.

2.3. Off-shore Wind

Another renewable energy area in which EP has an important role is off-shore wind energy. In recent years, off-shore wind energy projects have been gaining momentum in Europe. A confluence of factors, including the scarcity of land which can be developed as wind farms, some very favourable wind conditions on the oceans and maturing wind farm designs have all contributed to an increase in the number of off-shore wind farms. Following the lead of Denmark in off-shore wind farms, Germany has aggressive plans to increase its off-shore wind energy capacity (Knight, 2007).

Figure 19: Offshore Wind Turbines – Denmark
Image Source: Europa Technologies, Scankort, Google Earth

The energy that can be generated by wind turbines, both off-shore and on land, is highly dependent on the local average wind speed. The specific geographical characteristics of the wind farm sites can be affected by many factors (elevation, vegetation, surface roughness, etc.). However, in general, the regions with the most attractive potential are along the coastline, inland areas with open terrain or on the edge of bodies of water, such as lakes. In addition to these regions, some mountainous areas also have good potential for wind power (CETC, 2004). Off-shore wind farms can offer some of the best sites for wind energy generation, since wind speeds are consistently high on ocean surfaces and there is no vegetation which can cause friction and decrease wind speeds.

For the financial success of a wind energy venture, siting decisions are critical. Given the critical dependence of wind energy output on long-term mean wind speed, a proper wind resource assessment can make a sizeable impact on the financial viability of off-shore wind farm investments. In this process, EP applications using satellite data can be particularly useful.

2.3.1. EO-Wind farm (Environmental Information Services for Wind Farm Management)

EO-Wind farm is a research group funded by the European Space Agency (ESA) in order to identify the present and long-term needs for geo-spatial information related to global offshore (near-coastal) wind mapping. This project aims to estimate wind power potential derived from satellite Synthetic Aperture Radar (SAR), scatterometer and altimeter data, numerical model results and in situ measurements. The priorities of EO-Wind farm are "to develop an action plan to meet the need for wind farm siting and specify the service chain and different products (wind, current, waves etc), customize and test the individual components of the wind farm siting service chain for maximum site selection efficiency and validate the service with key customers involved in the project, and identify improvements required for developing a fully operational service" (EO-Wind farm, 2007).

2.3.2. Satellite Data for Off-Shore Wind Resource Assessment

Satellite data is already being used for onshore wind farms to determine surface roughness, build digital elevation models and classify land-use around prospective wind farm locations. However, to date, no operational method has been developed which can determine wind speeds on land using satellite-based instruments, although there are experimental projects.

For off-shore wind measurements, there are three sets of satellite instruments which provide useful information: passsive microwave instruments (such as the Special Sensor Microwave/Imager - SSM/I), scatterometers (such as NASA's QuikSCAT satellite) and Synthetic Aperture Radar (such as RADARSAT-1, ERS-2 and ENVISAT). Passive microwave instruments provide relatively low spatial resolution but since the satellites have been operational for a long time, they provide data for up to 17 years, the longest global time-series on ocean winds (Hasager, 2006). Scatterometers and SAR are both active (radar) instruments and can provide all weather and night-time coverage capability. Canada has its own SAR satellite, RADARSAT-1, which has been providing data since 1995. Other SAR instruments are onboard European ERS-2 and ENVISAT satellites.

One specific advantage of SAR is the ability to provide detailed spatial information, which can be very valuable for the feasibility analysis of an off-shore wind farm. Although SAR analysis is high-cost (due to the cost of imagery and labour-intensive post-processing and analysis techniques), it can nevertheless be a very useful source of information to complement other data sources. The archived time-series and global data are additional advantages. Disadvantages include limited temporal resolution (due to these satellites being in LEO) and low absolute accuracy when compared to in-situ measurements (Hasager, 2006).

2.4. Other Applications

Earth Observation using satellites has other applications which can be of interest for other sectors of the energy industry. Since it is not feasible to present an exhaustive discussion of these applications, a number of them will be briefly mentioned here.

As discussed in the Sea Ice Case Study, Earth Observation satellites play a very important role in the observation of sea ice. The Canadian Ice Service, a division of Environment Canada, is particularly reliant on two satellite platforms: RADARSAT-1 and ENVISAT. As mentioned above, RADARSAT-1 is located on a LEO polar orbit and it has an active instrument onboard which can provide data in all weather conditions. ENVISAT, a European satellite, has multiple instruments onboard including a SAR instrument. Both of these satellites can provide data regarding the presence or absence of sea ice as well as its type and thickness.

Snow cover monitoring and snow pack measurements are important to forecast the amount of power hydroelectricity plants can generate. An interview with Dr. Jun Yang, Director General of the National Satellite Meteorological Center (a division of China Meteorological Administration) revealed that his office provides assistance to the Chinese hydroelectricity sector by monitoring the snow cover on the Tibetan plateau, an important hydro source for China. Dr. Yang also mentioned services for sea ice monitoring for the coastal areas along the North East of China.

Use of satellite data is of course not limited to renewable energy sources. In order to better understand the benefits of satellites for the Canadian energy industry in general, an interview was conducted with Dr. Vernon Singhroy, senior research scientist at the Canada Center for Remote Sensing, a division of Natural Resources Canada. He also mentioned other types of applications including analyzing the morphology of sediments using remote sensing techniques, using interferometry techniques to help design better oil and gas pipelines (especially in the Arctic regions) and using hyperspectral techniques for identifying mineral and oil deposits.

Demand side applications are also being developed, such as more accurate load forecasting models, which are also partially based on satellite data.

3. Strategic Decision Making: An EP Perspective

3.1. Modelling the Decision Problem: Strategic

Various EP applications in the renewable energy sector constitute a rich set of choices on which to base our valuation framework. It was decided to focus on off-shore wind energy for a number of reasons: there are currently no off-shore wind farms in Canada, and the valuation exercise can help understand the value of EP for this emerging application. Moreover, Canadian researchers are actively involved in developing EP applications for the off-shore wind energy industry using RADARSAT-1 data.

3.2. The Case: Horns Rev Offshore Wind Farm

In order to illustrate the impact EP can have on siting decisions, the following calculations are based on a real-life case: Horns Rev (Devil’s Horn) off-shore wind farm. Horns Rev wind farm, completed in 2002, is one of the world´s largest off-shore wind farms and it is located in the North Sea, 14 km west of Denmark. It has 80 wind turbines capable of generating 2MW of power each, resulting in a maximum production capacity of 160 MW. The wind farm was built at a cost of 270 million euros and the annual expected production is 600,000 MWh, enough to supply the annual demand of 150,000 households (more information can be found in Appendix 1).

As the first step of our analysis, the RETScreen[30] analysis tool provided by Natural Resources Canada was used to determine the power curve for the 2 MW wind turbines used in Horns Dev. The power curve plots the power output (in kW) of a wind turbine as a function of wind speed (m/s). The values provided by RETScreen were used to fit a polynomial curve which was consequently used to generate more power output values using interpolation (Figure 20).

---

[30] http://www.retscreen.net

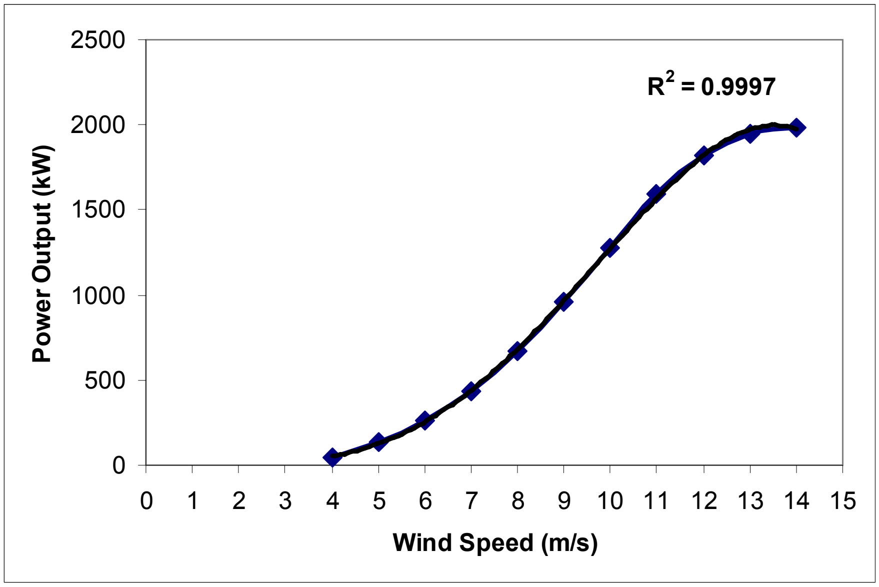

Figure 20: Power curve and the polynomial curve fit used for the wind turbine

The nameplate capacity (the maximum rated output) designated by the manufacturer of this wind farm is 1,401,600 MWh/year (obtained by 160MW times the number of hours in a year, 8760). However, the expected output indicated by the wind farm operator is 600,000 MWh/year, corresponding to a capacity factor of about 43%. It is important to note that this value is based on the long-term mean wind speed, and substantial amounts of variation are to be expected from year to year (due to various reasons discussed in the Wind/Water Case Study).

As part of the feasibility analysis for the Horns Dev wind farm, a comprehensive wind resource assessment study was conducted before construction, resulting in the estimate of a long-term mean wind speed of 9.7 m/s. This value indicates the speed of wind not at the ocean surface, but at 62 m, close to where the turbine blades are located (70m above sea level). Figure 21 shows the distribution of wind speed along with the observed wind directions.

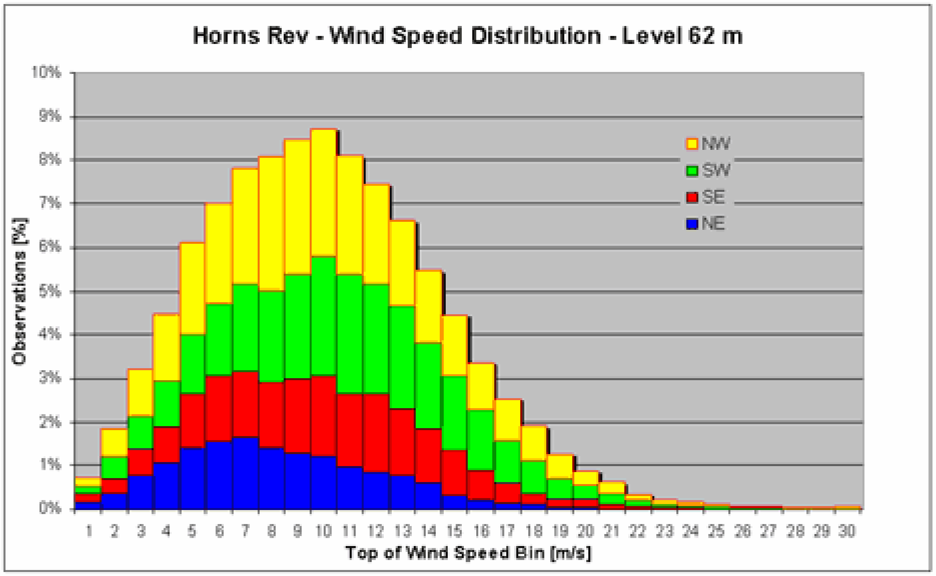

Figure 21: Distribution of wind speed and wind direction at Horns Dev
Source: http://www.hornsrev.dk/Engelsk/default_ie.htm

In order to estimate the total expected output from the wind farm as a function of wind speed, the annual average atmospheric pressure and temperature were estimated using the data from a near-by weather station (Copenhagen). These values, along with the 9.7 m/s wind speed at 62 m height were incorporated into our analysis using the RETScreen tool. Adjusting the resulting figure with some assumptions regarding various energy losses in the wind farm (such as wind turbine wake effects, downtime losses, etc) and using a Rayleigh distribution for the wind speed, gave a value of 603,962 MWh for the expected annual generation. Given that this result is within 1% of the operator's estimate, the same procedure was used to analyze various "what-if" scenarios.

This analysis allows for the estimation of expected revenue from the wind farm at various power output levels by making some additional assumptions. For instance, at 9.7 m/s wind speed and annual average of $50/MWh wholesale electricity prices, this wind farm would generate annual revenues of about $30.2 million. Clearly, the calculations for the economic benefits are very sensitive to the assumptions made regarding the market price, and the benefits can vary from year to year as a function of both the power price and the wind speed.

As discussed in the Wind/Water Case Study, power output is a function of wind speed and even a minor overestimation of long-term mean wind speed can result in actual performance drastically underperforming projected performance, and consequently, lower revenues than originally estimated. In order to build some insight, consider the following graph, which maps the expected annual revenue from this wind farm as a function of the mean wind speed for that year (Figure 22)[31]. Note that the "base case" scenario of 9.7 m/s yields $30.2 million per year (not shown on the graph).

---

[31] Note that the annual revenue is capped just above $37 million (with these price and wind assumptions), because of the "decreasing returns" of higher wind speeds as the turbines get closer to their maximum generation capacities.

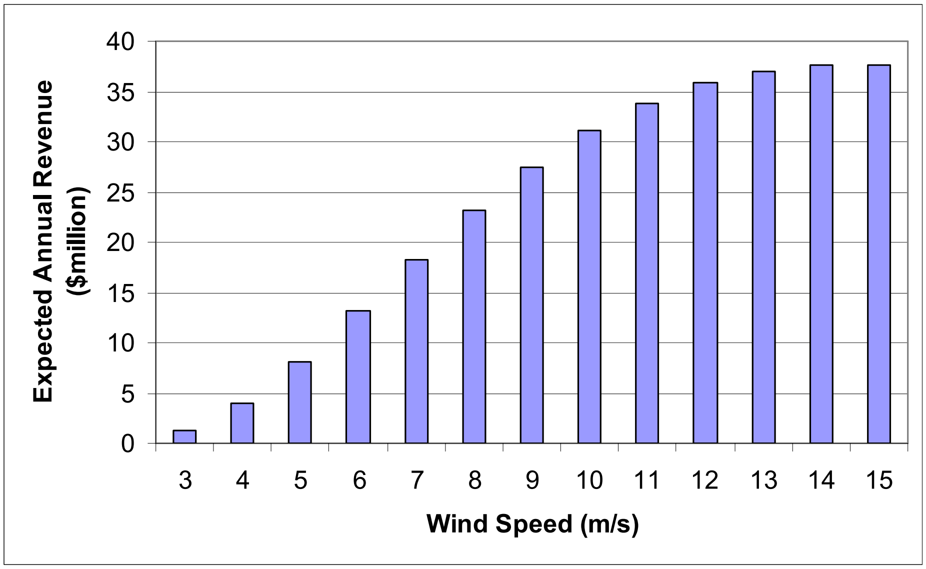

Figure 22: Expected annual revenue as a function of wind speed

In order to demonstrate the impact of wind speed and wholesale electricity prices on annual wind farm revenues, a sensitivity analysis was conducted. This table (in Appendix 2) maps the annual revenue of the wind farm as a function of annual average wind speed and annual average wholesale electricity prices[32].

3.3. First steps to Strategic EP Valuation

As it can be observed on Figure 22, even a slight decrease in wind speed has a big impact on the bottom line of wind farm operations. Therefore, obtaining very accurate long-term mean wind speed estimates as well as understanding the variability characteristics of wind speed and direction are critical to ensure the long-term financial viability of off-shore wind farm investments.

EP plays a very important role in the estimation of wind speed and wind direction. Annual wind speed averages show considerable variation from one year to the next. Consider the case for the North Sea: based on three different data streams from satellites, Figure 23 shows the annual average wind speeds for a location on the North Sea (56.5N, 5.5E) (Hasager et al., 2006).

---

[32] Using an annual average wholesale electricity price for these calculations is justifiable as long as wind speed is uncorrelated with power price. Our initial tests have indicated that there is currently no correlation. This will not be true if there is some systematic anticorrelation between high winds and low prices. Also, if and when wind energy constitutes a much higher portion of total generation capacity, due to supply/demand reasons a correlation may emerge.

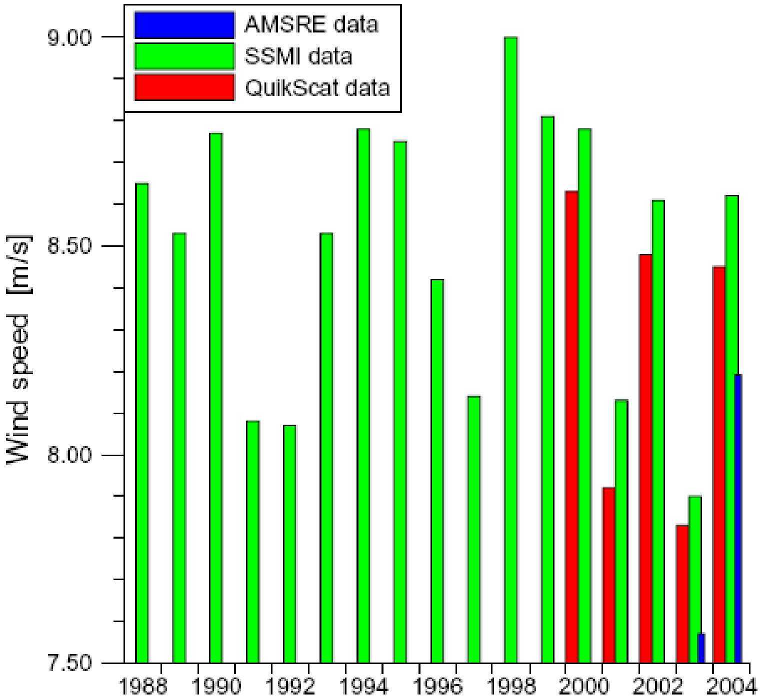

Figure 23: Annual average wind speeds - North Sea
Source: Hasager et al., 2007

The variability of wind speed is evident, and the magnitude of this variation can create a big impact on the revenues of an off-shore wind farm located at this site. Between the "best year" of 1998 and the "worst year" of 2003, there is about 1.1 m/s difference. Note that the wind speed values on Figure 23 are for 10 meters above sea level and these figures correspond to even higher mean speed differences at hub height.

The difference of 1.1 m/s at 10m corresponds to about 1.44 m/s difference at the hub height of 70m[33]. Using the valuation model mentioned above, a difference of 1.1 m/s from the base case scenario of 9.7 m/s corresponds to a wind speed of 8.26 m/s for the wind farm. At this level, the wind farm can generate about 489,303 MWh/year, resulting in annual revenues of $24.47 million (again with the assumption of $50 MWh price average for the year). Therefore a reduction of 1.1 m/s in wind speed (at 10m) corresponds to losses of $30.2 - $24.47 = $5.73 million *in just one year*[34].

3.4. Defining Accuracy

Both for on-shore and off-shore wind energy projects, the most accurate way to calculate actual wind speeds is to obtain direct measurements using anemometers by installing a meteorological mast (metmast) on location. Such masts are generally equipped with various instruments to measure wind speed and direction, and they also

---

[33] This figure was computed using a wind shear exponent of 0.14 in the RETScreen tool 1.1*(70/10)^0.14 = 1.44 m/s.

[34] Another way to look at this relationship is to calculate the loss in revenue from the base case when average wind speed drops by just 1% (i.e.., from 9.7 m/s to 9.603 m/s): about $350,000/year.

have data loggers and data transmission systems (such as satellite uplink/downlink systems, or systems based on mobile phone technology such as CDMA or GSM).

One disadvantage of metmasts is cost: for on-shore masts, the cost of installation is on the order of $20,000 or more depending on the type of instruments installed, the remoteness of the terrain, etc. For off-shore metmasts, due to the complexity of marine operations, the cost of operating a single mast can be on the order of 1 million euros a year (ESA, 2003). Clearly then, an exhaustive approach to install metmasts in every desired location is not a feasible approach.

One other shortcoming of metmasts is the need to obtain at least one year's worth of data before any detailed analysis is conducted for wind resource assessment. Even one year's worth data cannot capture the long-term variability characteristics of wind, as shown in Figure 23. If a metmast had been installed at this location in 1998 (the "best year") and the investment decision was made only based on this measurement, the investors and wind farm operator would have unrealistic electricity output expectations, as evidenced by the wind speeds in 2001 and 2003. In other words, although a metmast is the most accurate way to measure wind speeds once it's operational, what really matters for wind energy projects is to have a well defined statistical model for wind speed properties that is valid over many years. Therefore, in order to capture longer-term variations, climatological adjustments are needed before metmast data can be used in decision-making (this requirement applies to other short-term data sources as well, regardless of their source).

There is also an opportunity cost associated with waiting to obtain an additional year of accurate data. Consider the following example: historical off-shore wind speed measurements for a site ("K13") reported in Coelingh et al. (2001) show that the wind speed average (at 73.8 m above sea level) was 9.8 m/s in 1996 and 9.5 m/s in 1997. The long-term mean wind speed, based on 12 years of historical data (1985-1997), was 9.87 m/s. The two-year average of 1996 and 1997 is 9.65 m/s (0.22 m/s less than the long-term mean). Therefore, if a metmast was installed in January 1996 and very accurate yearly data was collected for two years (at a cost of close to $2.8 million), the wind revenue estimates would have resulted in an underestimation of approximately $790,000/year. In this particular case, obtaining the second year of data not only costs $1.4 million, it also results in a lower revenue estimate. Moreover, one additional year of data collection causes a delay in construction and operations, thereby postponing revenue generation.

This is precisely where modelling and satellite data come into play. Polar orbiting satellites have been collecting wind data for many years now, in some cases (such as the SSM/I data) as far back as 17 years. Therefore, even if a site already has a metmast installed, satellite data analysis can still make a significant contribution by shedding light on longer-term trends in wind speed and direction. In fact, complementing metmast measurements with satellite data can achieve the best of both worlds: enabling the calibration of satellite data with metmast data, thereby increasing the accuracy of estimates obtained from past satellite data.

Wind resource assessment based only on metmast data from one year can differ from the long-term mean wind speed by a significant amount. A study funded by the European Commission, Predicting Offshore Wind Energy Resources, provided insight into variations of offshore wind patterns. For various sites around the North Sea, long-term mean wind speed was calculated using actual observations with a time-series of 12 years (between 1985 and 1997). It was found that annual mean wind speeds can be above the long-term mean for these sites from anywhere between 3% to 12%[35](Coelingh et at., 2001). Not surprisingly, the industry practice is not to rely on any single source of short-term data, but to calculate the climatological average of wind speed (using more than 10 years of data) for a given site. Since the valuation framework presented in this case study is based on the assumption that no climatological average was available, the net effect of this assumption can result in an overestimation of the value of EP in offshore wind resource assessment.

3.5. Some Lessons Learned for Canada

Up to this point, the discussion has centred around research work conducted in Europe. There are currently no offshore wind farms in Canada. However, this situation can change in the future. Researchers at the Institut National de la Recherche Scientifique (INRS) in Québec have been working on using various satellite measurements in order to conduct wind resource assessments, mainly around the Gaspe Peninsula.

For locations without any in-situ measurements, currently the only available data source is the Canadian Wind Energy Atlas, developed by Environment Canada. This atlas is based on the Mesoscale Compressible Community (MC2) model.

Researchers from INRS and Environment Canada have estimated the accuracy of measurements from QuikSCAT by comparing the satellite measurements to in-situ buoy measurements and concluded that satellite data can give very accurate estimates for the long-term mean wind speeds (Beucage et al., 2006)[36]. This finding is in-line with another study conducted by Danish researchers who compared QuikSCAT data with in-situ tower measurements at Horns Rev (Hasager et al., 2006).

A recent study by the INRS/Environment Canada research team found that compared to data from QuikSCAT scatterometer, MC2-computed winds tend to overestimate wind speeds by about 4% for the Gaspe region (Beaucage et al., 2007). Based on the valuation framework presented in Section 3.2, with a long-term mean wind speed (at hub height) of 9 m/s, and an average price per MWh of $50, this 4% overestimation can result in $1.44 million of lost revenue per year[37].

---

[35] In many cases, annual mean wind speeds are less than the long-term mean as well. This can also have a negative impact on offshore wind energy investments, by eliminating sites which can actually yield higher energy output in the long-term.

[36] For the period of 2000 to 2004, the root mean square (RMS) error of QuikSCAT measurements is 1.74 m/s with a bias of 0.04 m/s for a total of 1351 co-located comparison points.

[37] For a similar wind farm configuration as the Horns Dev, assuming that the wind farm is built with the expectation of 9.36 m/s mean wind speed at hub height, and the actual mean wind speed of 9 m/s.

Given the complementary nature of satellite data and in-situ measurements, one EP powered strategy could be to conduct pre-feasibility studies based on satellite data, identify a few promising offshore sites, and then install meteorological masts on these selected locations (as illustrated in Appendix 2). This strategy can minimize the risk of placing an expensive meteorological mast in a site not representative of long-term averages, but also help calibrate satellite data for the feasibility analysis stage when actual in-situ measurements for at least 1 year will be available. The cost of satellite data for the pre-feasibility stage can be very reasonable depending on the desired spatial resolution. For instance, QuikSCAT data can be obtained free of charge, whereas SAR data cost can quickly add up for commercial projects[38].

4. Conclusions

Just like any other remote sensing system, satellites cannot provide direct measurements and a series of calibrations and data processing is needed to obtain the results. However, as discussed in this case study, in some situations, satellite data analysis can compensate for low short-term accuracy by achieving accurate results for longer time scales, thanks to the satellite databases containing data from as many as 17 years ago.

New ways of combining data from various satellites as well as in-situ measurements address the individual weakness of each approach, and set a new standard for renewable energy resource assessment. The experience gained from applying EP in solar and wind energy investments can also be leveraged to help various emerging renewable energy sources such as wave and tidal power systems.

Finally, it should be noted that the benefits of Earth Observation quantified in this case study are based on a very specific application. Satellites provide many other benefits through a myriad of EP-related applications, including the monitoring of sea ice, measurement of snow packs for hydrology and other applications. A possible angle for a future research project could be to make a comprehensive evaluation of these benefits, including societal ones.

5. References

Beaucage, P., Glazer, A., Choisnard, J., Yu, W., Bernier, M., Benoit, R and G. Lafrance, “Wind Assessment in a Coastal Environment Using the Synthetic Aperture Radar Satellite Imagery and a Numerical Weather Prediction Model”, INRS and Environment Canada, Research Paper, 2006

Beaucage, P., Lafrance, G., Lafrance, J., Choisnard, J. and Bernier, M., “A New Strategic Sampling Using Synthetic Aperture Radar Satellite Data for Offshore Wind Assessment”, INRS and Environment Canada, Research Paper, 2007

---

[38] Prices quoted by MDA Corporation, commercial provider of Radarsat-1 imagery, are on the order of $2000 per image, and a few hundred SAR images are needed for each location. In comparison, European SAR data are less expensive.

CETC, CanMet Energy Technology Centre Varennes, "Wind Energy Project Analysis", 2004

Coelingh J.P., Folkerts, L., van Zuylen, E.J., and Wiegerinck G.F.M., Predicting Offshore Wind Energy Resources, Final Report, Chapter 10 Comparisons with Measured Data, 2001 available at http://www.eru.rl.ac.uk/POWER_project/POWER_project.htm

Economist, "Solar Power: Trapping sunlight", Sep 13th 2007, The Economist print edition

ESA, "Space Shows way to Europe's Renewable Energy Future", European Space Agency news release, 31 July 2003, available at http://www.esa.int/esaEO

Garrad Hassan Canada Inc., Short-term Wind Energy Forecasting: Technology and Policy, report prepared for the Canadian Wind Energy Association, 2006 available at http://www.canwea.ca/images/uploads/File/Resources/FinalGHReport.pdf

Gurtuna, O and Trepanier, J., On-Orbit Satellite Servicing: A Space-based Vehicle Routing Problem, (co-authored with J. Trepanier), in T.Ciriani (ed.), "Operations Research in Space and Air", Kluwer, Dordrecht, 2003

Hasager, C. B., "Satellites in offshore wind energy: wind resources and wake analysis", Wind Energy Department Risø National Laboratory, Denmark, presentation made at the GEOSS and Wind Energy Workshop, Milan, May 2007

Hasager, C.B.; Astrup, P.; Christiansen, M.B.; Nielsen, M.; Barthelmie, R.J., Wind resources and wind farm wake effects offshore observed from satellite. In: Proceedings (online). 2006 European Wind Energy Conference and Exhibition, Athens (GR), 27 Feb - 2 Mar 2006.

IEA, 2006, "Key World Energy Statistics 2006", International Energy Agency report, Paris, France, 2006.

Knight, S. "Offshore Reality Hitting Home", Wind Power Monthly,Vol. 23, No.9, September 2007-10-13

Lautenbacher, C. C., "The Global Earth Observation System of Systems: Science Serving Society", Space Policy, vol. 22, pp.8-11, 2006

Mathieu, P.P., "Earth Observation from Space in Support of the Renewable Energy Sector", Refocus Volume 6, Issue 4, July-August 2005, Pages 52-54

McCarthy, Shawn, "Can Dirty Old King Coal Manage to Clean Up Its Act?", The Globe and Mail, 19 March 2007.

Websites of interest:
http://www.retscreen.net/ang/home.php
http://www.eumetsat.int/Home/index.htm
http://www.solarpaces.org/News/Projects/Spain.htm

http://www.grupocobra.com/home.html
http://www.sener.es/SENER/index.aspx?&lang=en
http://www.abengoa.com/
http://www.iberdrola.es/wcorp/corporativa/iberdrola?IDPAG=ENINICIO
http://www.meteocontrol.com/
http://www.solemi.de/home.html
http://www.eo-wind farm.org/index.php
http://www.ciemat.es/

Appendix 1: Technical Information – Horns Rev Wind Farm

Figure 24: A 62-metre meteorological mast
(Source: Hasager, 2007)

Wind turbine type: Vestas V80 - 2 MW
Total wind farm output: 160 MW
Expected annual production: 600,000,000 kWh
Rotor diameter: 80 m
Hub height: 70 m
Weight, blade: 6.5 tonnes
Weight, nacelle: 79 tonnes
Weight, tower: 160 tonnes
Weight, foundation: 180-230 tonnes
Total weight per wind turbine: 439-489 tonnes
Cut-in wind speed: 4 m/s
Full power output from: 13 m/s
Cut-out wind speed: 25 m/s
Mean wind speed at 62 metres' height: 9.7 m/s
Water depth: 6-14 m
Distance from land: 14-20 km
Distance between wind turbines: 560 m
Wind farm area: 20 km2
Total project costs: DKK 2 billion/EUR 270 million

Source:
Horns Rev Offshore Wind Farm: Ground-Breaking Wind Power Plant In The North Sea, Elsam Energy
http://www.hornsrev.dk/

Appendix 2: Sensitivity Analysis

Wholesale Electricity Price ($/MWh) (all revenue figures in million dollars)							
Annual Average Wind Speed (m/s)	**20**	**30**	**40**	**50**	**60**	**70**	**80**
3	$ 0.54	$ 0.81	$ 1.07	$ 1.34	$ 1.61	$ 1.88	$ 2.15
4	$ 1.60	$ 2.39	$ 3.19	$ 3.99	$ 4.79	$ 5.59	$ 6.38
5	$ 3.24	$ 4.86	$ 6.47	$ 8.09	$ 9.71	$ 11.33	$ 12.95
6	$ 5.24	$ 7.86	$ 10.48	$ 13.10	$ 15.72	$ 18.34	$ 20.96
7	$ 7.32	$ 10.98	$ 14.65	$ 18.31	$ 21.97	$ 25.63	$ 29.29
8	$ 9.29	$ 13.93	$ 18.57	$ 23.21	$ 27.86	$ 32.50	$ 37.14
9	$ 11.01	$ 16.52	$ 22.03	$ 27.53	$ 33.04	$ 38.55	$ 44.05
10	$ 12.45	$ 18.67	$ 24.89	$ 31.11	$ 37.34	$ 43.56	$ 49.78
11	$ 13.55	$ 20.33	$ 27.11	$ 33.89	$ 40.66	$ 47.44	$ 54.22
12	$ 14.34	$ 21.51	$ 28.68	$ 35.85	$ 43.02	$ 50.19	$ 57.36
13	$ 14.82	$ 22.24	$ 29.65	$ 37.06	$ 44.47	$ 51.89	$ 59.30
14	$ 15.05	$ 22.57	$ 30.09	$ 37.62	$ 45.14	$ 52.66	$ 60.18
15	$ 15.05	$ 22.57	$ 30.10	$ 37.62	$ 45.15	$ 52.67	$ 60.20

Appendix 3: Decision Tree (Strategic)

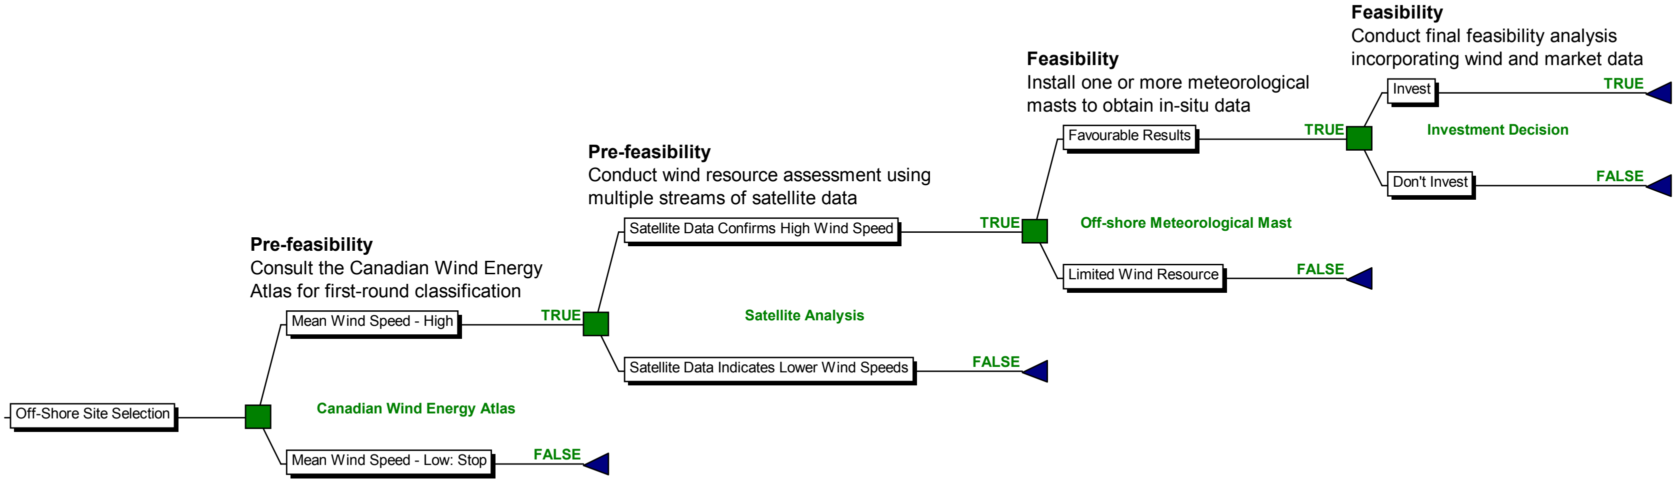

Environmental Predictions and the Energy Sector: A Canadian Perspective

Recommendations for Future Work

Prepared for

Environment Canada

Contract Number: K3A40-06-0028

Prepared by

TURQUOISE
TECHNOLOGY SOLUTIONS INC.

DYDEX
DYDEX RESEARCH AND CAPITAL LTD.

1. Recommendations for Future Work

Both the Literature Review and the Case Study reports were designed to answer some fundamental questions regarding the benefits of Environmental Prediction within the Canadian energy industry context. We hope that some of these answers and the valuation framework itself will add value to the decision making process in public and private sectors. We also feel that these findings motivate many ideas for future work.

The following sections present some future research concepts which were identified as both feasible and high-value research targets.

1.1. Enabling Financial Applications

One of the themes that came up repeatedly in the list of case studies was the fact that, although environmental variables are crucially important to the energy industry, the direct steps that can be taken to utilize environmental predictions are often proportionally quite small. As a result, we believe that there is a large opportunity for financial markets in this sector. Financial markets can be used to trade weather risk between counterparties with varying exposures to and appetite for risk. Such markets for weather derivatives, in particular those on one site temperatures, have already been developed and are discussed in Appendix 1 of this final section.

Other derivatives, on wind speed and precipitation, have also been created – it is possible that even more derivatives on cloud-free days, stream flows, and even sea ice coverage might find their application one day. A goal of such a forward looking study should be to determine what products might be useful (rather than identifying which already exist).

Another related research topic is understanding why weather derivative markets have not really taken off that much internationally. In theory they seem perfect: they allow physical market participants to offset risk while giving financial investors a new asset class uncorrelated with the rest of the market. We suspect that one limiting factor could be the quality and depth of publicly available EP. For instance, taking a counter position against Hydro Québec, which employs a full-time meteorology desk, could be a very risky move if the other party has just two math PhDs at a hedge fund sitting in a room in Greenwich Connecticut.

Therefore, the role of a federal department such as Environment Canada can be to catalyze such markets by working together with financial institutions (e.g. Montreal Exchange), energy industry stakeholders and other government departments (e.g. Industry Canada) by providing leadership and actionable environmental data. Such services do not have to be free but can be provided at a reasonable cost and on a very timely basis. This would encourage more traders to be involved and increase the market depth. It would also pave the way for next generation derivatives on novel indices such as wind speed, sea ice coverage, etc.

1.2. Capturing the Societal Benefits

The mandate of this research project was to understand and quantify the benefit of EP for private sector applications. However, from the very early stages of the research effort, it became apparent societal benefits of EP are also very significant. Based on some of our findings, one very relevant question is how to quantify the broader societal benefits of EP for the energy industry. For instance, EP also has a role in reducing the risk of power outages such as those experienced during the Québec/Eastern Ontario ice storm. Another example is the incorporation of air quality indices into scheduling smog producing "peaking" power units. Another obvious candidate is studying the possible impact of climate change on the energy value chain and the possible roles of EP to manage the associated risks and opportunities.

The value of such applications transcends the merely financial concerns, since they would also enable better safety for energy industry employees, more efficient use of fossil fuels, with concomitant reduction of both health- and happiness- destroying smog and greenhouse gas emissions. Interested readers can refer to Appendix 4 of the Literature Review Report for a discussion on EP as a public good.

1.3. Determining the Value of EP for Different Jurisdictions

Our research identified fundamental differences regarding the way energy is generated and traded in Canadian provinces. Therefore an up-to-date survey of the energy market structure in Canada and the corresponding value of EP can be very valuable for both provincial and federal decision makers. For instance, the users and benefits of wind forecasts can be quite different between Alberta (which has a deregulated energy market and no feed-in tariffs) and Ontario (which also has a deregulated market but offers feed-in tariffs).

1.4. Understanding the Value of Diversification

Another recurrent theme during the research activities was the role of diversification to mitigate the intermittency of renewable energy. This diversification can be geographical (e.g. by installing wind farms in various locations of a province exhibiting vastly different wind regimes), financial (e.g., maximizing electricity revenues in the wholesale market by switching from one asset to another based on generation forecasts) or it can be based on generation type (e.g., smoothing out the variability of renewables by combining wind and solar).

We believe that understanding the potential and limitations of diversification is essential to create synergy with other methods of dealing with intermittency (such as energy storage).

1.5. Enabling Conservation through better EP

One research stream with a lot of potential benefits is understanding the value of EP for demand side energy market participants. This topic ties into Lovin's idea of 'negawatts': why bear the environmental and financial costs of constructing new power plants if you could simply use less power at no cost to comfort? This is particularly relevant since many power plants are constructed to meet peak demand. What if you had a good temperature prediction and some thermal inertia in your office tower? Presumably you could optimize your use of electricity for air conditioning by increasing the amount consumed at lower demand (and hence price) times and decrease the amount used at high demand (and price) times. Such a process could save money for individuals and businesses while providing additional societal benefits due to a flatter load shape.

1.6. Designing Public Policy

Research findings indicate that the value of EP can be drastically different for various decision makers. An expected value decision maker processes and acts on information differently from a minimax decision maker. Generally speaking, the value of EP tends to be higher if the decision maker can afford to make mistakes and survive until the next round. For insight into this, see Appendix 2 of this section, which analyzes the very simple EP-enabled decision making process of using an imperfect rain forecast to decide whether or not to wear a raincoat.

Therefore, for certain applications, maximizing the value of EP would require transforming minimax decision making settings into expected value ones. We believe that one way to achieve this is by replacing hard constraints by soft penalty-based constraints (a classic example is emissions trading schemes, but there are also many other market-based mechanisms).

EP can empower private sector participants to size up their environmental and economic risks and develop alternative ways to manage such risks. On the flipside, new financial markets powered by EP can create new opportunities for other participants who are willing to share some of this risk or present innovative ways of managing it by better understanding the nature.

1.7. Carrying the Message

In addition with this list of new study ideas, an important feature of future work should be disseminating the work already done to a broader industry- and academic audience, through workshops or conference presentations and publications. These publications could include peer reviewed publications and possibly even a book built from the bones of this research project.

Another concrete next step might be the organization of a large workshop with industry, government, and academic participants. As the findings of this study have highlighted some resistance in the industry to actively use EP, the goals of such a workshop could include understanding the concerns and priorities of industry first-hand, encouraging more active use of EP and determining the specific types of EP industry is interested in. Partnering with various industry associations such as Canadian Association of Petroleum Producers, Canadian Wind Energy Association and Canadian Solar Industries Association, can be a very fruitful step in this direction.

Appendix 1: Temperature Case Study Sketch

Both the natural gas and the electricity industry have significant temperature dependence.

Natural gas is chiefly used for heating, which in the winter is nearly completely driven by the temperature as shown in Figure A.1. Electricity demand, particularly in the summer, is driven by the need for air conditioning, which is driven by temperature. In Davison's analysis of temperature vs. electricity demand in Ontario, little increase in predictive power was found if Humidex corrections are included, but Adam Kucera (2007, personal communication) indicates that in Australia the humidity correction is very important.

In terms of strictly physical responses, the strong temperature dependence on demand does not imply the ability for a temperature prediction to avoid costs. However electricity utilities are able to use accurate weather forecasts to appropriately start up the fires in fossil fueled power plants in order that these plants can quickly react to demand swings. A complete analysis of the value of temperature forecasts in this setting has been performed by Teisberg et al. (2005) who found that, in the United States, accurate 24-hour temperature forecasts are worth US$166 million per year in optimizing generation mix. They also characterize the economic benefit in this respect of a single Celsius degree improvement in temperature forecast of US$60 million per year (note that this is a tiny, about 0.012%, improvement on the huge $US50 billion cost of fueling electricity generators in the United States over a year).

To transform this number to a Canadian context, one is tempted to use the common practice of simply dividing by 10. However it is likely that a somewhat larger divisor is needed as easily dispatchable hydroelectric power is much more prevalent in Canada than in the United States, reducing Canadian dependence on fossil fuel based spinning reserve. Our rough estimate of this direct impact of weather forecasts is about C$10 million per year.

It has already been noted that the direct savings emerging from accurate weather forecasts represents a tiny fraction of a huge industry. At the same time, the change in power demand necessitated by even small changes in temperature is immense. As seen in Figure A.1, the difference between a 25 degree day and a 26 degree day can impact the electricity load in the province of Ontario by about 500 MW. Over a 16 hour on-peak day the impact of this is 8000 MW-h which may be worth, during a summer heat wave, $100 per MW-h or even more. So, in a single day, the temperature dependence of a single degree is worth nearly a million dollars just in a single province.

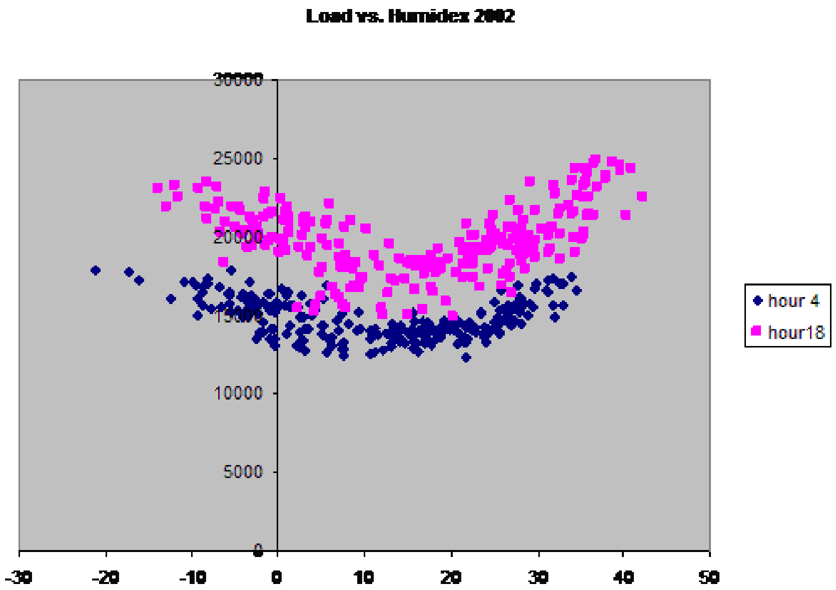

Figure A.1 Data Sources: Ontario IESO as archived by Dydex (load); Environment Canada (2006,2007) for temperature and humidity data. Aggregate Ontario load data is plotted against Toronto (Pearson airport) climate data.

Figure A.1 Load (in MW) vs. Humidex (Ontario)

A financial market has emerged in which even the financial impact of these temperature-induced load swings can be hedged. This is the market for so-called "Weather Derivatives", most of which are temperature based. Temperature derivatives work by constructing a contract whose payoff depends on the accumulated temperature (Pierre Lang, 2007, personal communication). To be more precise, these payoffs are usually reckoned in terms of Heating Degree Days (HDD) or Cooling Degree Days (CDD).

If Figure A.1 is examined, it is apparent that when the temperature is above about 15 degrees Celsius, the load is a roughly linearly increasing function of the temperature (in which case we discuss cooling degree days, since power is being used for air conditioning) while when the temperature falls below the same threshold the load increases with falling temperature (in which case we discuss heating degree days, since power is here being used for heating). The number of cooling degrees on a given day in a given city might be computed by first determining the average temperature for the day (for instance via 0.5(Tmin + Tmax) and then subtracting 15 degrees from this number. If the result of this subtraction is negative, then the number is replaced by zero. So the cooling degree days calculated on a day with a maximum temperature of 27 and a minimum of 19 would be max(0.5(27+19) -15,0) = 8, while the heating degree days calculated on a day with a maximum temperature of 16 and a minimum of 10 would be max(0.5(16+10)-15,0) = 0.
The total number of cooling degree days in a month is determined by summing this expression over each day in the month (or each weekday in the month).

A very simply weather derivative contract might pay $1000 for every heating degree day tallied up for the month. A dizzying array of contract specifications is, however, possible.

By participating in this market, both power consumers and power producers may hedge away the temperature dependence of their portfolios. Investors looking for investment products, the returns of which are completely uncorrelated with financial markets, are also attracted to this market. Finally, speculators who have a view on future temperature prices will provide liquidity to this market.

Clearly, the ability to predict temperatures has value in this market. As in the example shown above, weather derivatives contracts often have long averaging periods, diminishing the value of a strictly short term weather forecast. But in principal, any contract which allows mutually beneficial risk exchange between counterparties can be traded. Thus, for instance wind derivatives, precipitation derivatives or even sea ice derivatives and stream flow derivatives could be traded if there is sufficient interest.

Financial markets have an important role in unlocking the value of EP to the energy industry and a future work building on this study should definitely investigate their power for good.

References:

Teisberg, Thomas J, Rodney F. Weiher and Alireza Khotanzad, "The Economic Value of Temperature Forecasts in Electricity Generation". Bulletin of the American Meteorological Society (2005) pg 1765-1771.

Various personal communications as noted in the text.

Appendix 2: The "Raincoat Problem"

A very simple problem in acting on imperfect prediction information is the decision problem of whether or not to wear a raincoat. Wearing a raincoat always imposes a cost: wearing an unattractive and heavy garment. However, if it rains, this cost is counterbalanced by the benefit of staying dry. Furthermore, if you are caught without a raincoat during a rainstorm, there is a cost associated with getting soaked.

The resulting cost matrix for the decision may therefore be described by:

	Wearing a Raincoat	Not Wearing a Raincoat
It Rains	1	M
It is Fair	1	0

Table A2.1: Raincoat cost matrix.
Note: We assume that there is a cost of 1 to wearing a raincoat but a (generally bigger) cost of M to not wearing a raincoat when it rains.

Merely by examining the cost matrix some special cases of the decision problem already emerge. For instance, we might consider the "teenage boy" cost matrix in which m = 1 and M = 0 (you couldn't care less about getting wet but you would rather die than wear your geeky raincoat to school to the disapprobation of your peers). For an agent with such a cost structure, the optimal strategy will be never to wear a raincoat regardless of climatology or weather forecast. A somewhat less severe case is the "Tarzan" scenario in which M is only a very tiny bit > 1. Intuitively, Tarzan would have to be pretty certain of his forecast before bringing his raincoat to work. Another opposite case is the "Wicked Witch of the West" case (named after a character in Frank Baum's Wizard of Oz). The Wicked Witch of the West disintegrates on contact with water, so she would choose a very large value for M, perhaps an infinite one. For this reason, she would presumably always wear a raincoat regardless of EP information.

Next, we need to establish a climatological base rate. Suppose that with no weather forecast information we may assume that it rains a fraction p of the time (and is fair the remaining 1-p of the time). Thus for Dublin, Ireland, p is nearly 1 while for Las Vegas, Nevada, p is nearly 0. Using only climatological information of this type, agents lying between Tarzan and the Wicked Witch of the West in their aversion to rain contact may evaluate an optimal decision.

Suppose for instance M = 2 and p = 0.3. Then an agent who decided always to wear a raincoat would incur a (certain) daily cost of 1 while an agent who decided never to wear a raincoat would incur an expected daily cost of 0.3*2 + 0.7*0 = 0.6, and so would optimally choose never to wear a raincoat (Note that taking expectations here does not imply a risk neutral decision maker, since any risk aversion could be incorporated into the cost matrix). The value of any weather forecast for raincoat wearers needs to be referenced to the appropriate climatologically optimal decision.

Now we are ready to model rain forecasts. We do this via a "skill" matrix in which the probability of rain given a rain forecast, fair given a rain forecast, etc. is summarized. Thus:

	Rain Forecast	**Fair Forecast**
Rain	**α**	**1-α**
Fair	**1-β**	**β**

Table A2.2: Forecast of rain skill. Here, for example, the probability it is fair given that rain was forecast is 1- β.

A perfect "oracle" rain forecast would have α=1 and β= 1. Note that it must be that both α > ½ and β > ½ or the forecaster would simply reverse her forecasts.

With this modeled, we are now in a position to determine the incremental expected benefit of a rain forecast to an agent with a cost matrix as in Table A2.1. The formulas are obtained merely by a simple application of Bayes' Theorem but are, even in this very simple setting, algebraically quite imposing. The companion spreadsheet RainForecastExample.xls contains an implementation of them.

Intuition can be obtained by investigating some special cases. For instance, consider the case examined above of an agent with M=2, living in a region with climatology such that it rains 30% of the time. As discussed above, her no-forecast strategy is always to wear a raincoat at an expected cost of 0.6. For her, a perfect forecast would enable her to wear a raincoat during exactly the 30% of days for which it rained, for a cost of 0.3*1 = 0.3, so the value of a perfect weather forecast is, in this setting, 0.6 -0.3 = 0.3 per day.

If, however, the skill decreases in this setting to have α= 0.9 and β= 0.9, it turns out the value of a forecast falls to 0.2 (essentially the wearer is now exposed to the cost of wearing a raincoat when it isn't necessary and to the cost of not wearing a raincoat when it is necessary).

If, in this example the risk aversion is turned up so that M=4, the optimal no-forecast solution is to always wear a raincoat (since 1 < 4*0.3 + 0*0.7). In this setting, a perfect forecast becomes quite valuable (0.7), since it allows the raincoat to be left at home 70% of the time, at a savings of 1 per time. However, an imperfect weather forecast here, in particular in which it occasionally rains when fair weather was promised, dramatically decreases the weather forecast value. If α=1, in other words it always rains when rain is forecast, but β= 0.9, in other words it also rains 10% of the time that fair weather is forecast, the value of the weather forecast decreases to 0.47.

For the highly rain averse Wicked Witch of the West, with M = 1000, a perfect rain forecast is still worth 0.7, but even a nearly perfect forecast which has α=1, in other words it always rains when rain is forecast, but β= 0.99, in other words it also rains just 1% of the time that fair weather is forecast, is completely worthless. The thinking here is that if she acts on the forecast she saves 0.7 units (by never wearing a raincoat when it is fair) but loses approximately 0.01*0.3*1000 = 3 by leaving her raincoat at home on 1% of rainy days. For someone with this level of risk aversion, even a very accurate weather forecast is of no value.

The reader is encouraged to look at RainForecastExample.xls for more insightful examples and to play around with the cells to create their own examples.

The main insight here is that risk aversion destroys the value of forecasts. This is, to a certain extent, consistent with the messages of several of our case studies, in which seemingly valuable EP opportunities were ignored by industry.