Generic System Requirements for High Altitude Wind Turbines

Nykolai Bilaniuk, Ph.D., P.Eng.
CEO, LTA Windpower
Ottawa, Canada
v1 - October 2009

ABSTRACT

The basic concept of a high altitude wind turbine (or still more generally, airborne tethered wind generator) appears to have merit based on the fact that abundant energy is present in the wind at heights beyond the reach of current wind turbine towers. However, exploiting that energy is a challenge. Many different technologies in this vein have been proposed but they can't all be equally promising. This paper is intended to assist the prospective investor or purchaser in reasoning through this technology area: it explores the practical requirements that must be met by an airborne tethered wind generator technology to be both technologically feasible and economically viable. The paper concludes that *autonomous survivability* is a largely neglected but critical aspect of performance.

Introduction

Wind power is not usually viable as a sole source of electricity because it is difficult to match supply and demand. However, it can be an attractive supplement to other forms of electrical generation. This includes on-grid use, where windpower isn't economically competitive with fossil fuels or hydro, but brings ecological benefits that governments use to justify subsidies. Wind power makes even more sense in off-grid applications displacing fuel that would otherwise have been consumed by diesel generators.

By way of introduction, it would be useful to define a few windpower related terms, particularly for a non-specialist audience. Some other terminology may be introduced later.
- "Efficiency" is the fraction of the energy in the wind in the area intercepted by the wind system that is usefully captured for conversion to electricity at a given instant. Theory proves that the upper bound on possible efficiency is 16/27ths or slightly over 59%, called the "Betz Limit". The symbol "Cp" ("coefficient of performance) is often used for efficiency.
- "Capacity factor" is the ratio of the long-term average output to the rated output. For a large terrestrial wind turbine in a good location, 25% would be a good result.
- "Nameplate capacity" is the rated maximum output of the wind generator. For commercial units, this ranges from a few hundred watts to a few megawatts.
- Wind energy is captured when the wind pushes against a surface, causing motion of that surface. There are two general possibilities. If the motion of the surface is in the same direction as the wind, as in a cup anemometer for example, then the wind machine is a "drag" design. If the motion of the surface is perpendicular to the motion of the wind, as with a propeller, then the wind machine is an "axial flow" design.

It is worth noting that of all the wind turbine designs invented for terrestrial applications, a large fraction is technically feasible, but just one type dominates the market: the axial flow propeller turbine, accounting for well over 99% of the global installed generating capacity. It would actually have been surprising if two or more technically widely divergent approaches had turned out to offer the exact same net economic benefits, allowing multiple designs to remain competitive. In the early days of the industry, different ideas were tried, but the whole industry eventually converged on what works best. Even achieving a close second place in viability is not good enough to survive in the marketplace. The margins are just too thin.

Airborne wind energy systems do have potential advantages over their terrestrial counterparts, but despite this, airborne systems have yet to be produced in volume at all, although they were pioneered decades ago. The reasons for the commercial failure of airborne wind energy systems - at least so far - must be explored and understood. Thus this paper is structured as a series of exploratory questions that examine the options in airborne wind turbines, seeking to identify core requirements that affect technical feasibility and economic viability.

A shakeout similar to the one among terrestrial wind technologies is inevitable among airborne tethered wind generator technologies as well. The fact that currently every proponent of this basic idea is advocating a completely different technology is a mark of the relative youth of this concept as an industrial sector. This situation cannot last. Either one technology will emerge as a clear winner, or none will prove viable in competition with other approaches to electrical power generation including terrestrial windpower. Which approach to airborne tethered wind power generation offers the best prospects of being competitive? This paper does not answer that fundamental question, but lays out the groundwork allowing the reader to perform informed evaluations and draw reasonable conclusions.

Let us begin by asking the simple question, why even consider airborne systems?

1. Why go airborne to collect wind energy?

Simply put, the fact that winds are stronger and steadier at higher altitudes above terrain provides the justification for reaching higher, and beyond some point tethered systems are bound to be less expensive than ever-higher towers.

That winds are stronger and steadier the further one rises above terrain is generally recognized, so the claim is not controversial. However, it is reassuring to examine some real data. A "wind speed profile" - is a plot of windspeed vs. elevation above terrain averaged over many geographic locations. Fortunately, there are examples of data from the U.S. National Oceanic and Atmospheric Administration (NOAA) being distilled in this way. See for example this plot from the USA at 40 degrees north latitude available here: http://showcase.netins.net/web/wallio/MZW.html and reproduced below:

Figure 1   Winds in the United States at 40 degrees north latitude

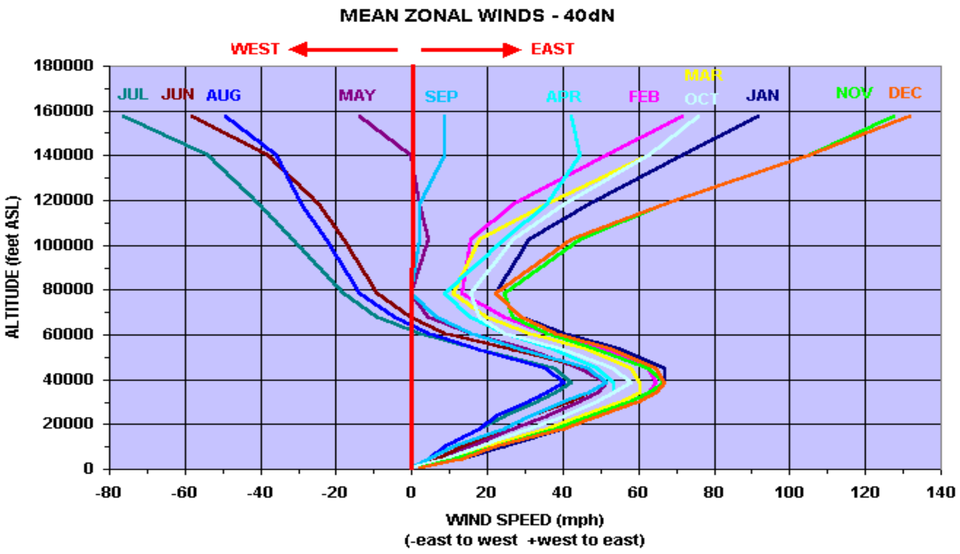

From this data we can see that up to about 40,000 feet (12 km) elevation, wind speeds increase steadily with altitude, and their seasonal variability is relatively small. Since the energy in the wind is proportional to the cube of the windspeed but is linear with air density, the available energy in the wind increases initially with increasing altitude. Above 12 km elevation, wind speeds actually decline until about 80,000 feet (24 km). Also, air density declines throughout with altitude. Compared to the value at sea level, air density declines to about 95% at 0.5 km, 89% at 1 km, 69% at 3 km, and at 12 km reaches just 19% (http://www.engineeringtoolbox.com/air-altitude-pressure-d_462.html) .

Given the wind energy's dependence on air density and wind speed, there cannot be any sense in operating wind energy systems at greater altitudes than 12 km. This still does not mean that 12 km is the ideal operating altitude, since there are other practical limitations that influence the decision, and which make it advisable to operate much lower.

From the above data one cannot tell what happens in the crucial range of elevations within 1 km of terrain. What is intuitive is that there is greater variability in wind strength and steadiness near the surface, where winds are most influenced by terrain roughness and solar thermal effects. The size of surface features, or "surface roughness" plays a leading role in reducing the winds at the surface, as illustrated from the following chart.

Figure 2. Winds close to the surface and their dependence on surface roughness
[The chart is from a Canadian government website, and we wish to duly acknowledge authorship, but regrettably we can no longer locate the original source.]

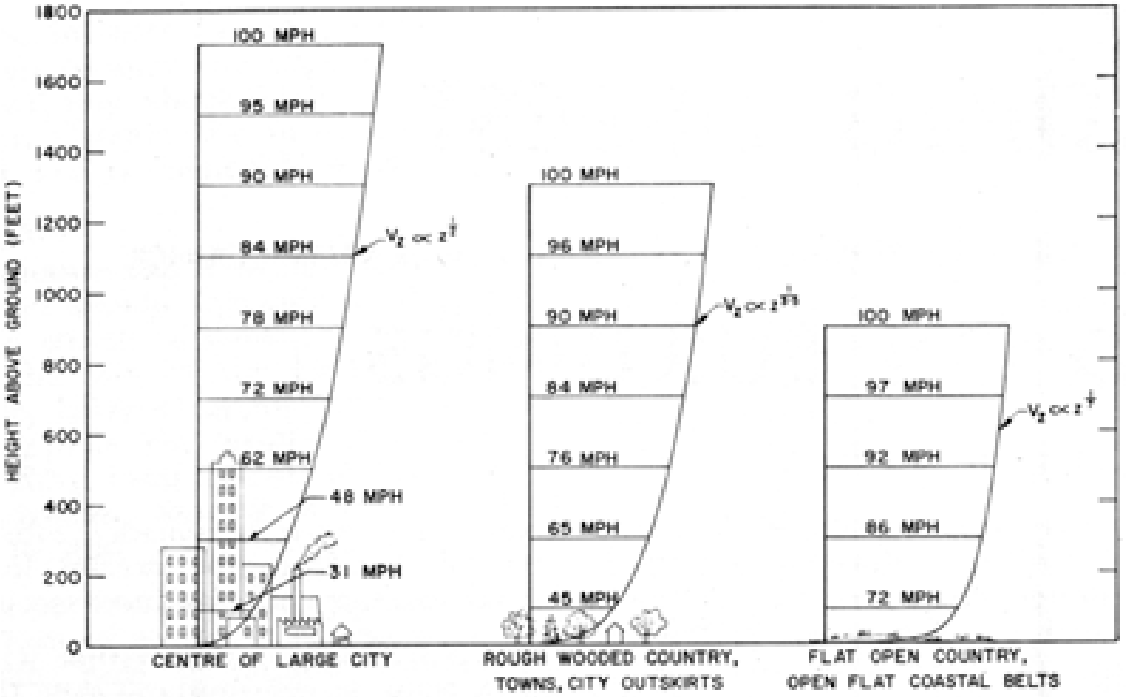

Locations with the best surface winds tend not to coincide with the locations of electricity markets. Thus a compromise is required in site selection for terrestrial turbines. Building a higher tower is a workaround, but the cost of building towers ever higher doesn't scale well: as height is doubled, the cost is much more than doubled. The towers built for today's largest wind turbines are of order 100 metres in height. Meanwhile the cost of making a longer tether for an airborne generator increases linearly with length, and lengths from 100 to beyond 500 metres are feasible. Given that the energy in the wind is proportional to the cube of the wind speed, and both wind speed and wind uniformity increase with altitude at all locations, tethered turbines have a great advantage with regard to capacity factor.

Consider a location with poor surface wind resources, namely Ottawa, Canada:

Environment Canada data shows that mean predicted windspeeds are in the range of 3-4 m/s at 30 m elevation, 4-5 m/s at 50 m, 5-6 m/s at 80 m, 5-6 m/s at 100 m, 7-8 m/s at 270 m, and 9-10 m/s at 550 m elevation. The wind energy at 500 metres is approximately 10 times what it is at ground level.

The tower construction problem is particularly acute in remote locations, since the cost of bringing building materials and construction machinery to remote sites can be prohibitive, especially to those sites that are off-road and distant from navigable waters. It is also particularly expensive to supply fuel to this class of locations. An example is the arctic. There, the best surface winds are found in coastal areas, while surface winds deep inland are weaker. Distances are vast and roads are few. Construction of both roads and towers can be greatly complicated by permafrost and its thawing. Winter roads suffer from ever-shorter seasons because of climate change. Points of interest such as valuable ore bodies are in inconvenient locations.

In competition with terrestrial wind turbines, airborne systems might not be able to offer increases in nameplate capacity or improvements in efficiency. However, by reaching to higher altitudes they can offer higher capacity factor. Doubling the capacity factor from 10-25% to 20-50% is realistic, which other things being equal lowers the cost per delivered kilowatt-hour by half. The airborne systems may also offer simpler construction, and more freedom in site selection which reduces the cost of transmission lines and their inevitable electrical losses.

Lesson learned:
(1.1) Making a wind generator airborne indeed has its attractions (but they can be realized only if doing so works well enough to be both technically feasible and economically viable).

2. Location of the generator itself

A practical issue that arises in the design of a high-altitude capable windpower collection system is where to put the electrical generator itself. One could either A) collect energy from the wind but keep the generator on the ground, OR B) put the generator in the air.

If one were to opt for solution A), then in lieu of transmitting electricity to the ground, one needs to transmit mechanical power. One way or another, that mechanical power needs to end up in rotary form, since that is how electrical generators expect their input. Also, generators prefer to run at relatively high speed, on the order of a few thousand rpm. For AC generators that are directly connected to the grid, the output frequency and phase have to be exactly matched to the grid. The nominal speed depends on the number of poles as well. If grid synchronization is not necessary, the situation is more flexible.

To understand the topic of mechanical power, it is necessary to consider two simple equations. Let P = power, w = angular frequency, T = torque, r = radius of action, and F = force. The equations describing the power in rotary motion are:

$$P = wT$$
$$\mathbf{T} = \mathbf{r} \times \mathbf{F}$$

The equations tell us that in order to get high power output, either a high frequency of rotation, or high torque are required. High speed / low torque is better. If the raw power source instead provides high torque / low speed, producing substantial power involves high forces on at least some components, thus the parts must be strong and expensive. In particular, in order to get from a low-frequency high-torque regime to the high-frequency low-torque generator input, an expensive high step-up ratio transmission is needed. High step-up ratio also means high "cut-in windspeed" (the speed at which the wind generator begins to work at all), so some of the available power in the wind will be wasted.

How can mechanical power derived from the wind be transmitted from the air to the ground? Several possibilities suggest themselves.

a) Undulations in the tension of the tether

This has been proposed by various researchers. In this approach, kites or gliders are manipulated in their flight paths to vary the force they exert on their tether, which is purely a tensioning member. Both projects envision using multiple kites, but they are otherwise different. One conceptual problem with this idea is that a high angular frequency is difficult to achieve. If a flywheel is used, it may be expected to turn at a fraction of a revolution per minute, with the attendant consequences of high input torque and high cost. A variation on this approach would be to have a cable wrapped around a pulley or drum, and at least two kites working in tandem and pulling the cable back and forth in spurts such that each to-and-fro cycle involves many turns of the pulley wheel. This has the benefit of producing an output with higher angular frequency, but makes the output intermittent. Intermittent output will generally require the system to use a DC power storage scheme, such as batteries or capacitors, which hurts overall efficiency and scalability while increasing cost. In another variation, kites are steered from the ground by multiple tethers, and these are induced to alternately maximize and minimize the tether load. During the tether load minima, the tether is partially retracted using a motor, and spooled out during the maxima while driving a generator.

b) Ganged tethers

This idea can be found in an early patent to Miles Loyd (see U.S. Patent No. 4,251,040). This method involves generating power using propellers mounted on kites, with a tether consisting of multiple reciprocating cables ganged using an arrangement similar to the crankshaft of an internal combustion engine. This approach does allow power to be gathered by kites flying across the wind, and can deliver mechanical power to the ground at frequencies much higher than the time it takes for the kite to fly a complete loop. It i s unknown whether this design has ever been built successfully.

c) Torsion in the tether

Another approach may be to use torsion rather than tension in the tether, for example, in the Selsam patent (U.S. Patent No. 6,616,402). This is challenging from a strength-of-materials standpoint, particularly for large sized systems.

d) Compressed Air

What about using the captured windpower to run a flying air compressor, and transmitting the power to ground via a compressed air hose? The air can then do work on the ground, such as crank a generator via a turbine. This also makes some sense, and has been discussed in engineering circles. However, the efficiency of this process is almost certainly poorer than transmitting electricity.

Some of the above ideas should work. The relevant question is the cost per kilowatt-hour delivered to the customer. Mechanical power transmission to ground involves either increased cost in mechanical components, increased complexity on the ground (or both), or deceased efficiency versus using a tether with an electrical conductor. This may be why, when we transmit hydro power over long distances, we do so in electrical rather than mechanical form.

The balance of costs seems to favour the electrically conductive tether. However, this has the disadvantage of exposing the ground equipment to increased risk of damage during electrical storms, a topic that will be revisited. A definite answer to the question of which method of transmitting power is best is lacking, because the extensive experimental data that would be required to prove the point is not yet available. Therefore one may wish to keep options open but flying the generator is probably more promising.

Lesson learned:
(2.1) Lean towards making the generator airborne, but there is insufficient experimental data available to make a firm determination on this issue.

3. Practical limits to tether length

This question affects both electricity-carrying and purely mechanical tethers, but electrical tethers face more constraints.

Electrical wire is made of copper or aluminum. Since an airborne tether will be flexed a lot, this discourages the use of aluminum because of that metal's poor fatigue properties. Thus stranded copper wire is called for, which is more expensive and heavier than aluminum, but is a better conductor.

At this point, it would be useful to point out that electricity-carrying tethers must be capable of carrying both electric current and mechanical tension, but not necessarily within one jacket. We might define a "unified" tether as one that does have all the components in one jacket, and a "non-unified" tether as one that doesn't. A non-unified tether may, for example, have the electrical wires suspended at intervals from the tension member by means of clips or standoffs of some kind. A unified tether may be used with a winch, whereas a non-unified tether cannot; it would be damaged if rolled into a winch.

Since a wire for a tethered airborne wind turbine will be flexed continuously in normal use, it would probably be considered a "flexible cable" for the purposes of the electrical code, although this precedent has yet to be firmly established. The appropriate cable parameters are discussed in Article 400 of the US National Electrical Code. The maximum voltage for any recognized type of flexible cable is 2000V, for types G, G-GC, PPE, and W. The numerous other cable types are limited to 600V or less. The maximum permissible current depends on the number of conductors in the cable, and is higher for 1-wire cables than for 3-wire cables. However, for a unified tether, 3 wires would be the minimum appropriate.

Figure 3   Properties of flexible cables as function of wire gauge (AWG) size

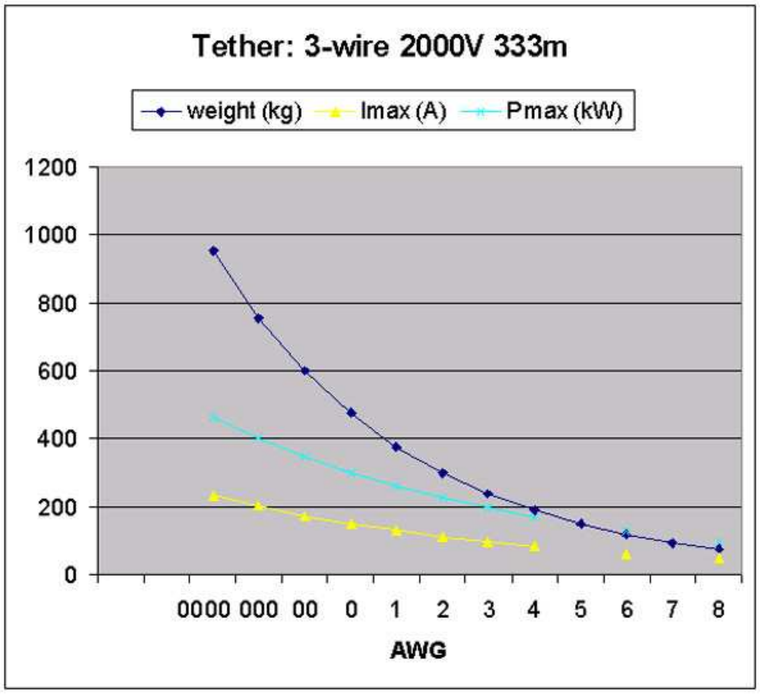

The plot (Figure 3) shows data for a unified tether under an optimistic scenario: a 3-wire 2000V system that is a relatively short 333 metres in length. The yellow trace shows the maximum permissible current allowed per conductor by the electrical code, in amperes. The resulting maximum possible single phase input power to the line (when the maximum current is driven at the full 2000V into a purely resistive load using two wires), expressed in kilowatts, is in light blue. The dark blue trace is the mass of the copper in the wire alone, not counting the insulation and tension members. In practical terms, the optimal operating condition of 2000V transmission is difficult to achieve, and at any lesser voltage or non-unity power factor the maximum power throughput declines and the percent loss rises.

Consider a realistic example of what one might try to build, a 100 kW wind generator with a tether length of 1500 feet (circa 450 metres) running 3-phase output at one of 380, 480, or 600 VAC. The corresponding phase currents at full power and full voltage into a resistive load are 88.6, 70.2, and 56.1 amperes respectively. In turn, the minimum required cable gauges are 3, 4, and 6 AWG respectively, and greater if the reactance of the load cannot be guaranteed to stay low. Even with a three-wire delta with no additional grounding lead, the weight of the copper alone is 325, 257, or 162 kg for the 1500 ft run. The ohmic losses in the tether at full current are 7 kW, 5.5 kW, and 5.6 kW respectively, or 5.5 to 7 percent of total production over 1500 ft. There is a tradeoff: one could trade up to heavier wire, reducing ohmic loss but increasing lofted mass in a regime where reducing both is precious. The three conductors bundled together must also spool onto a winch and the current must pass through a set of slip rings before the power is finally available on the ground. The drum diameter must be very large to accomodate the large minimum bending radius of such unified tethers. Avoiding excess curvature to the cable also means continually re-orienting the winch with changing winds, or developing some other means to achieve the same objective.

If a system were to operate at high power with a substantial portion of the tether spooled onto the winch, heat dissipation at the winch may itself become a problem. A system on the scale of the above example may be close to the practical limit of what is possible with a winched, unified tether using today's approved cable voltages. Even if an order of magnitude improvement is possible, it seems pretty clear that the suggestion made by some promoters of airborne tethered wind generators of operating megawatt scale systems at altitudes of 10's of thousands of feet (on the order of 10 km) is not realistic. Wire mass and wire losses are sufficient reasons to warrant focusing on operating altitudes of order 1 km or less.

According to the electrical code, a non-unified tether may operate at slightly higher currents. At some point, it may also be possible to build non-unified tethers with higher voltage than 2kV, but this will probably require specific regulatory approval. Non-unified tethers are less expensive than their unified counterparts but can't be winched.

Some people have suggested the use of superconducting cables for the tether, but this is completely outside the realm of possibility given the state-of-the-art in superconductivity. The highest temperature superconducting wires still require liquid nitrogen cooling, and a tether cannot be cooled to such low temperatures because of the prohibitive cost and bulk of the cryogenic equipment. In addition, these so-called high-temperature superconductors are too brittle to be used in a tether.

Concerning tethers for systems that bring the power to ground mechanically, not electrically, the constaints are more relaxed. E.g. the possible lengths are greater, linear density (weight per unit length) is lower, allowed radii of curvature against a winch drum may be tighter, and ohmic losses are irrelevant. Still significant limitations do arise. For one thing, winches can only be so large. Tethers do have mass and have some amount of spring to them. Thus they can store mechanical energy in gravitational potential and using elasticity instead of passing it down into the ground-based hardware. The longer the tether, the more this will be a problem. Non-electrical tethers are also not immune to the effects of electrical storms. A cable made of an insulator such as nylon or kevlar, but wetted on the surface, is a more attractive path for discharging lightning than the surrounding air. With these factors considered, while it *is possible to imagine* tethers several kilometres long, it *is not realistic to build* such tethers, even if they are transferring power to ground by mechanical rather than electrical means.

Lessons learned:
In general:
(3.1) Very high altitudes of order 10 km are not realistic no matter where the generator is located.
For electricity-bearing tethers:
(3.2) Use stranded copper wire.
(3.3) Focus on designs that fly the generator less than 1 km above terrain. Be careful with heat concentration around winches.
(3.4) For scalability to large sizes or higher altitudes, avoid winches to preserve at least the possibility of using a non-unified tether with separate distantly spaced power wires. This is also less costly.

4. Sources of lift

There are essentially two approaches to getting lift, A) wind-derived lift such as Bernoulli lift, and B) buoyant lift.

A) Wind-derived lift

Wind-derived lift refers to any lift mechanism driven by the wind itself. The most commonly used form of wind-derived lift is lift due to Bernoulli's principle. This principle describes how aircraft wings operate: as air flows around a wing or similar structure, a higher density, lower speed is obtained below the wing compared to above. The resulting difference in pressure is the lift. Bernoulli lift is used in a number of proposed airborne tethered wind generators. Another wind-derived lift mechanism is Magnus effect, which results from the rotation of an object with an irregular surface within an airflow.

Reliance purely on wind-derived lift has a serious deficiency. What if there is a calm in the wind? Advocates of going to very high altitudes claim that the winds are so steady there that a calm is not a serious threat. Even if that is so (and that is far from certain), we have already established that the prospects for very long tethers are not great, especially if they have to carry electrical power. For that matter, what if the system needs to be brought down for a landing because of maintanance, intolerable environmental conditions, or some other reason? This will surely be true at least occasionally. Then that very same airborne tethered wind turbine which was designed for the strong winds of very high altitudes must survive in the weak and rapidly varying winds close to ground.

The threat of wind speed variability when taking off and landing is very real. Consider the following data from a location in Canada that has a reputation for some of the strongest surface wind conditions the country has to offer - Sable Island, Nova Scotia. This island is home to a famous herd of wild horses, and hosts a research station manned by some 20 or so persons, but is otherwise uninhabited. There is no pier and no natural harbour on this sandbar, thus no sizable ship can dock. The island's electrical power was long derived from three main diesel generators and one last-resort backup, but two years ago five 7.5 kW wind turbines were added. The fuel for the diesel gensets has to be brought in by helicopter, and before the wind farm was installed, the island's annual fuel consumption was on the order of 20,000 litres.

As stated, Sable Island is considered an excellent location for windpower. It is a fairly flat island far from any other land mass. The island's characteristics are instructive, and can be discerned from the table in Figure 4.

Figure 4 Wind Statistics for Sable Island, Nova Scotia, Canada

Source: Principal Station Data for Sable Island. 1984. Canadian Climate Program, Atmospheric Environment Service, Environment Canada (PSD/DSP-92), found at: http://www.greenhorsesociety.com/Wind-Energy/Wind-1k.htm

	Jan	Feb	Mar	Apr	May	June	July	Aug	Sep	Oct	Nov	Dec
Prevailing direction	W	W	W	W	SW	SW	SW	SW	SW	W	W	W
avg. windspeed	32.0	31.0	30.0	26.2	23.2	20.3	18.5	18.9	21.6	25.8	29.4	31.3
% obs calm	0.8	0.7	0.7	0.7	0.6	1.0	1.2	1.3	0.8	1.0	0.7	0.9
% obs 1-10 km/h	5.8	6.1	7.0	8.7	11.1	14.2	16.0	17.1	13.1	9.0	6.7	6.2
% obs 11-20 km/h	16.3	19.2	20.7	26.2	34.0	40.8	45.7	43.5	36.5	26.3	21.0	17.9
% obs >21 km/h	77.1	74.0	71.6	64.4	54.3	44.0	37.1	38.1	49.6	63.7	71.6	75.0
peak wind km/h	126	114	140	114	100	109	98	97	124	158	174	130

Notes:
The values in the table are based on 30 years of observations ending in 1984. The anemometer was said to be at 10 metres elevation, but it isn't clear whether that meant 10 metres above terrain, or 10 metres above mean sea level.
The "avg. windspeed" is the average monthly value calculated from daily values which are averages of hourly readings.
The "% obs" means the proportion of hourly observations in which windspeeds within a particular range were recorded.

Although the average windspeeds at 10 metres elevation are a robust 18 to 32 km/h year round, calm nonetheless happens 1% of the time and low wind another 10% of the time or so. It stands to reason that at least some of the calms must begin suddenly, which would cause a low-flying Bernoulli-lifted device to stall and crash, possibly causing catastrophic failure and unacceptable expenses in downtime and repair (unless the device can be motored and have the power level change on extremely short notice). Airplanes do not recover from a stall in a very short time, and wind generators are not free to fly in any direction required to recover from a stall. They also start from a very low altitude to begin with. To boot, low altitude winds are less predictable, making it impossible to bring the system down in a controlled manner when you want or need to do so.

Meanwhile, the monthly peak winds on Sable Island are far in excess of the average wind, indeed, in excess of 100 km/h most months. This indicates that any windpower system must be very robust, and if it is an airborne system, the operators must be prepared to land it on short notice for its own protection unless it can withstand winds as high as 175 km/h.

For various reasons, Sable Island is not a good candidate for airborne tethered wind generators, most obviously because of the extreme conditions encountered, but also because surface winds are good enough, and because of the risk of injury to the horses who are free to roam as they please.

The salient points from the Sable Island data are that regardless of location, reliance on Bernoulli lift alone is not feasible for airborne tethered wind generators because of the danger from calm. Occasional wind calms are certain to occur. Also, any system regardless of lift mechanism must be prepared to deal with brief and not necessarily predictable episodes of excess wind.

B) Buoyant lift from lighter-than-air gas

A system using buoyant lift can survive in the air even when the wind drops to a near-calm or full calm, because there is no worry about stall speed. This is good. However, if there is a transient wind gust then the system gets blown towards the ground because of the increased wind loading. As it gets blown, momentum will carry it forward like a pendulum, and it will overshoot its new equilibrium position, so there is still a risk of a crash. For a start, it would be helpful to know where an airborne tethered wind turbine does sit under equilibrium conditions.

Figure 5   Forces acting on an airborne tethered body

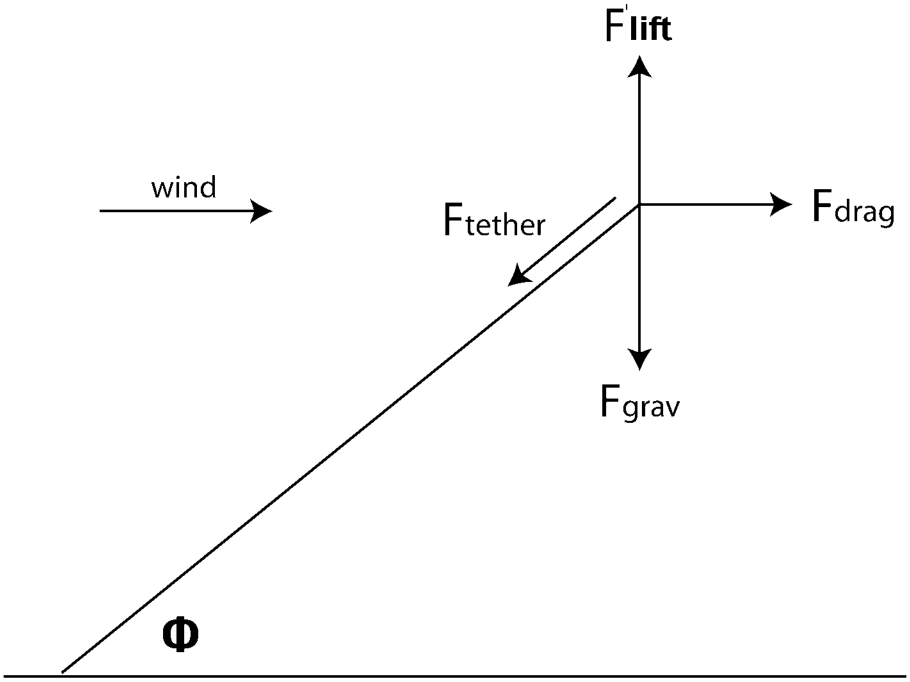

This diagram shows the aggregate forces acting on the airborne tethered wind generator. For simplicity, it is assumed here that the forces all act at one point, and that the tether itself is massless. In order for the generator to maintain position, the sum of all forces acting on it must be zero. Note that the only force opposing wind drag is the horizontal portion of the tether force, and it is from a portion of this drag that electrical power is generated. This has an interesting corollary: If the wind is horizontal and the tether is exactly vertical, then the generator cannot possibly produce any power, no matter what the detailed design of the system. Also, it isn't necessarily desirable to have the system sit at a high angle when producing a lot of power, because that means that the vertical portion of the tether tension must be enormous.

In more detail: As the wind increases, the drag increases and more power can be extracted. The increased drag will cause the system to lean downwind, and the angle **Φ** will get smaller unless the lift increases such that the sum of the vertical forces (lift+gravity) increases by the same proportion as the drag. In the event of a sudden wind gust, the system moves downwind but its momentum carries it beyond its equilibrum position before settling. The amount of overshoot depends on the speed of onset of the gust, the details of the design especially its lofted mass, and its dynamic behaviour. For a system reliant on Bernoulli's principle, the lift ought to increase with increasing windspeed, at least partially offsetting increasing drag and minimizing overshoot. For a system dependent solely or mostly on buoyant lift, it is not possible to offset increasing drag, and such a system will suffer from greater transient overshoot.

How much equilibrium downlind lean will occur can be quantified. First we resolve the tether force Ftether into its horizontal and vertical components,

$$\text{Ftether\_h} = \text{Ftether} \cos \mathbf{\Phi}$$

$$\text{Ftether\_v} = \text{Ftether} \sin \mathbf{\Phi}$$

Under equilibrium conditions,

$$\text{Flift} + \text{Fgrav} + \text{Ftether\_v} = 0$$

$$\text{Fdrag} + \text{Ftether\_h} = 0$$

$$\tan \mathbf{\Phi} = (\text{Flift} + \text{Fgrav})/\text{Fdrag}$$

The power extracted from the wind is equal to

$$P = Cp\ 0.5\ p\ A\ V^3$$

where Cp is the extraction efficiency, p is the density of the air, A is the area intercepted by the airborne tethered wind turbine, and V is the wind velocity. Note that Cp and A must agree as to whether they include the "passive" areas of the wind turbine that do not contribute to power production.

Another important concept is the drag coefficient Cd. This can be defined in different ways, but roughly speaking, it is equal to

$$Cd = \text{Fdrag} / (0.5\ p\ A\ V^2)$$

with the variables having the same meanings as before. Drag coefficients were traditionally determined empirically for a given shape by means of wind tunnel tests, but computational predictions are also possible.

For an airborne tethered wind generator, the relationship between power output and drag can be expressed as

$$\text{Fdrag} = (Cd\ \text{Pout}) / (Cp\ V)$$

Since it is undesirable to have the system assume a position too low to the ground, it is important to minimize the amount of drag per unit power output. This means striving for high efficiency is very important to airborne tethered wind turbines. Looking at it another way, there is a need to ensure that what drag there is, is "useful" drag connected with generating electricity.

Note that there aren't many turbine designs that are efficient. E.g. the Danish Wind Industry Association (http://www.windpower.org/en/tour/wtrb/persian.htm) rules out drag based systems as a class: "*Drag based windmills are very inefficient, and really only of interest for educational or hobby purposes. (With all due respect, think before you act - there are thousands of worthless wind turbine patents around. Contrary to popular belief patenting guarantees neither efficiency, usefulness nor workability)."* This conclusion by DWIA may be too drastic, in that a very inefficient system may still be able to produce power economically if it is proportionately cheaper to build and install, but we have established that inefficiency does increase the risk of crashes and does increase the tether loading, which are not desirable properties of themselves, and which do not make the total cost of ownership any lower. However, it is a telling point that most of the proposed airborne tethered wind generators being advocated today are drag devices such as kites.

In sum, considering the main sources of lift that could be utilized in airborne tethered wind turbines, it would appear that the use of Bernoulli lift alone is very risky because of the uncontrollability of the machine during calms, and during takeoff and landing operations. Meanwhile using buoyant lift alone seems technically feasible, but there is a cascade of economically undesirable consequences:

1) Using buoyant lift alone requires large surplus buoyancy to maintain clearance above the ground, and this adds to costs in lifting gas, envelope size, tether strength, and other expenses.
2) Surplus buoyancy means a winch is required to pull the system down when landing.
3) Reliance on a winch means a unified tether is needed (if the generator is lofted), and a turntable or cable-forming funnel to manage the direction changes of the wind, and a power source for the winch itself.
4) Should the heavy winch run on a combustion engine or electricity? If there is plenty of either

combustion motive power or electricity on location, why bother installing the airborne tethered wind generator? The winch is thus a major economic deterrent for remote locations.

Lessons learned:
(4.1) Using wind-derived (Bernoulli) lift alone does not appear to be technically feasible, for risk of crashes.
(4.2) Using buoyant lift alone does not appear to be economically viable.
(4.3) The winch should be optional if at all possible, rather than a necessity.

5. Environmental hazards

In this context, what is meant is not what hazards airborne tethered wind generators pose to the environment, but what hazards the environment poses to the generators. There are several potential stumbling blocks: A) lightning; B) excess wind and/or sudden changes in wind; and C) precipitation (rain, snow, and especially freezing rain that leads to icing).

A) Lightning

Lightning is a burst of electrical power that occurs when an electrical charge substantially different from neutral accumulates in the sky, and is suddenly equalized by discharge to the earth's surface. In temperate latitudes, electrical storms are more prevalent in the summer than in the winter. The energy pulse from lightning is short but very intense: it is a very high current driven by a very high voltage. Although there are devices to protect electrical equipment from voltage spikes, a direct lightning strike is almost certain to cause serious damage to a tethered airborne wind generator, particularly but not exclusively to one that carries electrical conductors its tether.

Note this critical problem does not cripple terrestrial wind turbines. Why not? Their nacelles and towers have ample steel in their structures and/or other very heavy gauge metal paths to ground, and that mass of metal can take both the current and the consequent heating from lightning without serious damage.

An airborne tethered wind generator has only the wires going down its tether. Even if there is a separate grounding conductor meant for lightning, this will not have the desired effect: the rapidly time-varying current will be induced into the other wires, and high voltage may easily exceed 2kV and overcome the insulation in a tether, or even melt the insulation. In the event of a direct strike, the first electronic "brick" downstream of the tether - likely an inverter - will almost certainly be destroyed, and perhaps more gear will go with it.

The incidence of lightning varies by region, but it is a non-negligible effect. In southern Ontario, Canada, lightning strikes approximately 2 times per square kilometre per year on average. However, prominent features are struck often. Toronto's CN tower is said to average almost one direct strike per week.

One cannot afford to have a power generating system absorb serious damage anywhere near that often. If lightning cannot be safely handled by an airborne tethered wind generator's tether, then that generator must avoid getting struck in the first place. That means bringing it down to the ground when electrical storms are anticipated. Fortunately, the incidence of such storms is reasonably predictable based on weather data in most populated areas, and in remote areas a storm front may be observed approaching, usually with at least some minutes' notice.

B) Excess wind and sudden changes in wind

This issue has been substantially covered in connection with the different possible sources of lift, i.e. item 4 above. To review: an unexpected calm is fatal to systems that rely on Bernoulli lift alone, while systems that rely on buoyant lift alone will have a tendency to overshoot their equilibrium position as they drift downwind in response to a gust. A system that has lower mass or balances both sources of lift, will be relatively more stable as winds change, and will therefore not overshoot its equilibrium position by as much.

The effect of winds in excess of the maximum survival windspeed of the system must also be considered. If the wind rises to a level close to the maximum survival windspeed at its operating altitude, it would be prudent to bring the system the ground. Ideally, it should be protected while on the ground, but even if it isn't possible to provide protection in a timely manner, it is still better to land quickly than to not do so. A wind-driven catastrophic failure on the ground is likely to be less costly than a catastrophic failure followed by a fall from a great altitude. The latter event would probably destroy the airborne equipment beyond salvage.

The data from Sable Island shows that brief episodes of very high winds can occur everywhere, and even a short episode is enough to cause catastrophic damage. It is also impossible to predict such gusts with enough accuracy. Conservative deployment ("never send the airborne tethered wind generator into the air unless one is absolutely certain there will be no high winds" is a receipe for loss of most potential production. The other extreme ("damn the torpedoes! keep that thing up there all the time!) is an invitation to frequent destruction. What is needed is an ability for the system to ground itself autonomously in response to excess winds.

C) Precipitation

Experience with both fixed wing and LTA aircraft indicates that rain (with temperatures clearly > 0 celsius) and snow (accompanied by temperatures clearly < 0 celsius) can be designed around readily. Conditions near the freezing point, especially with supercooled rain and its potential for freezing on contact, are a challenging problem. One solution used on modern fixed wing aircraft - the spraying of de-icing fluid - is not viable, as it is too maintenance intensive for this application. Generator flights need to last a month, not just a few hours. In the case of soft shelled lifting bodies on craft using buoyant lift, it is likely that their constant slight deformation will promote shedding and prevent ice buildup on their surfaces. This effect can be enhanced by operating the ballonet(s), if the craft is equipped with such. However, rigid components may be adversely affected by icing, and this has to be remedied. One solution may be a return to a technology already tested in World War II, for example on B-24 "Liberator" bombers. These wre sometimes equipped with a rubber de-icing boot on the wing, which could be inflated to cast off any accumulating ice.

Lessons learned:
(5.1) The only reasonable way of dealing with lightning is by landing before every storm, sitting it out on the ground, and re-launching afterwards. It should be possible to actuate the grounding remotely and on short notice, and it would be even better if the process could be completed autonomously.
(5.2) Managing in variable wind conditions calls for a solution that balances buoyant and Bernoulli lift, and allows rapid but controlled descent to ground in case of excess wind. It should be possible for the system to initiate and execute the temporary grounding autonomously.
(5.3) Icing on rigid surfaces poses a problem that must be resolved, but there are promising approaches to this problem, including some old solutions.

6. On-Location Ground Crew Requirements

Terrestrial wind turbines, even large wind farms, operate with no technicians in attendance whatsoever. It is easy to understand why: Cost. Consider that a technician may earn $25/hr while electricity is worth at most 10 cents/kWh without a generous subsidy. Thus a single technician costs as much per hour as 250 kWh of product. Since the capacity factor of a terrestrial wind turbine in a good location is approximately 25%, this means the technician costs as much as the output from a 1 MW turbine. Extending the argument further, if the gross margin on electricity produced at a wind farm is 10%, then a single full-time technician wipes out the profits from 10 MW of turbine capacity. It is no wonder that in real life technicians do the rounds at a wind farm only occasionally: they do routine inspections, lubrication of key parts, and repairs when required. During routine operation, wind farms must be able to get by with zero staff in order to be economically viable.

In order to be competitive against terrestrial turbines, airborne tethered wind generators cannot rely solely on their higher capacity factor. They must not require a ground crew beyond the occasional visit comparable in frequency and duration to the requirements of terrestrial turbines. Then when a visit is required, a small lightly equipped ground team must suffice. Above it was noted that lightning is a serious threat in warm months, icing in cold months, and brief high wind episodes year round. None of these events can be predicted with specificity a long time in advance, in fact sometimes there will be only minutes of warning, but all instances may require a rapid descent in order to avoid catastrophic damage.

The combination of inescapable environmental hazards and the economic necessity of avoiding keeping a standby crew means that an airborne generator must be able to undertake routine landings and takeoffs completely unattended to be economically viable. The deployment and retrieval technology must be well adapted to this mode of operation. The need for crew, especially a highly skilled crew equipped with heavy vehicles or other machinery, makes a design non-viable.

Less frequent but predictable environmental hazards, such as "once a year" level ice storms or sustained high winds such as hurricanes, would call for sending out a ground crew to secure the system on the ground. This is undesirable and hurts competitiveness, but is not a show stopper.

Lesson learned:
(6.1) Routine unattended takeoffs and landings must be possible, and minimal ground crews and effort should be sufficient to secure the system on the ground when absolutely necessary.

Conclusions

The topic of airborne tethered wind generators was introduced with the observation that the market for terrestrial wind generators has converged completely to one design, the axial flow propeller turbine. It was noted that this convergence should have been expected, because it would have been surprising if two or more approaches that are technically widely divergent had turned out to enjoy the same economics. The thesis was proposed that a similar shakeout among airborne technologies is inevitable, and that exactly one or none will emerge a winner.

From the gathered "lessons learned" about airborne tethered wind generators, one may summarize that the overarching requirement for their simultaneous technical feasibility and economic viability is *efficiency with autonomous survivability* under a large range of conditions, which must include everything except truly extreme conditions.

A certain number of airborne tethered wind power technology proposals are known either from the websites of their promoters, or from the patent literature which by now includes at least tens of patents. Evaluation of these known proposals for airborne tethered wind generators based on the criteria developed here leads to the unfortunate conclusion that *the number of truly promising candidates already described is a null set*. Some fail because they appear technically infeasible, while others may work, but their prospects are dim because they are not economically viable. This may explain why we do not see airborne tethered wind generators competing in the market today, but it does not mean that such a technology cannot succeed if a good enough design is devised. LTA Windpower was formed to exploit the opportunity that thus presented itself. The first challenge was to find whether there was a solution after all that would take the lessons learned into account. We believe we have found an answer.