"I see the effects of the wind, but I've ***never seen the wind...****"*

-Billy Graham

Seeing the Wind That is Coming

Welcome to the frontline of innovation in wind-sensing technology. We stand on the verge of a dramatic breakthrough, which will change how our man-made systems interact with the nature's dynamic forces. Wind and weather are complex, three-dimensional structures, which are constantly moving and changing. Traditionally, mankind has attempted to understand these forces of nature with overly simplistic macro effects (e.g. weather measurements and predictions) and then respond to them after they have arrived or passed. In fact, up until now, we have only been able to see the *effects* of the wind, not the wind itself. Light Detection and Ranging (LIDAR) technology has been shown to have the ability to sense the motion of microscopic particles in the air at significant distances. The ability to employ this amazing technology in a practical way has, until recently, been out of man's reach. Concurrently we are entering a phase of history when extracting clean energy from nature is of both economic and strategic importance, leading to higher demand for wind power as a renewable energy source. The demand for more efficient, reliable wind power and the ability to sense the wind at a distance have converged to create a fertile environment for advancement and innovation.

Now that we are able to measure air movement, in real time, remotely and accurately, what does that mean for future application? To answer that, this paper intends to start a scientific discussion of the engineering practices and challenges that will be faced in implementing LIDAR wind-sensing technology. It will not address, however, the specific means with which to measure the wind, but rather discuss the impact of that measurement. LIDAR has the potential to give three-dimensional speed and direction of the air mass at high update rates based on the processing of Laser Doppler returns. This measurement can be of an air mass at a single focused range or range-gated measurements, which can provide a wind profile measured as a function of range in a specified direction.

To this point in its development, the wind industry has relied heavily on tax incentives and grants to make wind energy economically viable. Sharply rising oil prices and environmental concerns are making investments in renewable energy net operational profits more competitive. Currently controls on wind turbines are relatively crude and cannot anticipate wind changes. Various studies[1] have concluded that measurement of wind significantly before it arrives would afford a wind turbine operator (or a control system) the ability to make adjustments to the turbine that will reduce its vulnerability to overload or vibration damage. The DOE estimates that such reduction would save approximately 10% of the maintenance costs of a wind turbine. By reduced repair cost and avoiding significant loss of revenue (energy production) due to maintenance down time, the cost-avoidance offered by a look-ahead LIDAR system is very attractive. Other aerodynamic studies[2] indicate that suboptimal alignment and adjustment of the wind turbine often reduces the energy production potential by 30%. All of this argues for the integration of advanced wind sensing and a responsive control system.

One way to focus on the sensing and control responses that are of first order importance to wind turbine profitability is to break down wind conditions into gross effects. As this body of knowledge and experience advances, the second and higher order effects can be studied and captured. A convenient hypothesis is that wind has both superimposed macro- and microstructures and can be thought of in terms of the 3-D vector resultant of both steady and dynamic components in both magnitude and direction. To examine the first order differences among various types of effects, small wind turbulence is not relevant to the control logic, but it could be of interest at the second or higher order effects as the origin of potentially damaging vibration modes of the turbine blades. Linear control theory and design can easily translate sensor input into predetermined mechanical or electrical system adjustments. Of interest is both the mechanical response time to make those adjustments and the "voting logic" which decides when an adjustment will benefit in balance with the energy needed to make the adjustment of large components and the potential for increased wear/decreased service life as a result of more frequent use of

---

[1] National Renewable Energy Laboratory, *Technical Report NREL / TP-500-39154*, January 2006.

[2] Leishman, J. Gordon, *The Principles of Helicopter Aerodynamics*, Cambridge Press 2000, pg. 727

the adjustment mechanisms. The concept is simple, but little actual control logic has been developed because of the lack of the timely sensor information. Without the ability to anticipate the movements of the wind, measurement of wind after passing through the rotors results in control lag, putting the WT out of phase with the changes in wind conditions.

More modern WT designs incorporate a ring gear to give the ability to adjust the yaw of the nacelle to align with what is believed to be the direction of the wind (in the horizontal plane). There is no ability, however, to pitch (or elevate) the nacelle to align with a vertical component (or the vector resultant of wind with a vertical component), even though the existence of up and down drafts (caused by thermal effects and terrain) are known to contribute a potentially important percentage of the total energy in the wind. The fact that the vertical component is so variable leads to the desire to directly measure it. The wind industry de facto standard measurement comes from anemometers (whether mechanical or acoustic) that are, by design, two-dimensional (in the horizontal plane).

"The acceptance of this gross data for operational decision-making is an indication of the lack of sophistication in the wind power industry today."

In addition to the lack of three-dimensional wind data, there are further limitations to implementing any advanced controls imposed by the current sensors (anemometers). Anemometers are frequently placed on the rear of the WT nacelle; measuring wind after it has passed through the blades. Not only is this time-late for use in a reactive control, but also the rotors significantly disturb the airflow, meaning wind data recorded at the back of the nacelle is not representative of the wind in front of the blade sweep. Time averaging is often used to smooth through the wake-induced variations in wind velocity. The industry typically takes the anemometer data (which is being determined at about a 1 hz rate) and then time averages. Variations (RMS) are recorded as indications of turbulent content. Clearly, in many if not most actual wind conditions there will be numerous changes in the wind over the averaging period, but no advantage is taken. The acceptance of this gross data for operational decision-making is an indication of the lack of sophistication in the wind power industry today. While the degree of variability in the wind is known to be somewhat a function of the wind farm location as well as seasonal weather, many of the sites with highly regular and predictable Class V wind have already been exploited. Growth in the wind industry is going to result in having to accept and make a profit at less optimal (consistent wind) locations (especially ashore) to build future wind farms. This will require more sophisticated sensing and controls as discussed herein.

What kinds of information does a turbine control system need to anticipate and react to wind conditions? The discussion below will break down simplified models of wind changes. It is recognized up front that while these may be useful for gaining an understanding of the problem, they are likely too simplistic for a basis for designing a system of integrated sensor and controls.

Wind Direction

The first obvious wind change is a step change in direction. To understand the parameters of interest, a starting condition of steady horizontal wind is assumed. A mass of air with a different direction at the same magnitude is detected some distance (and time) in advance of the wind turbine. At the leading edge there is no instantaneous change, but rather a mixing and shift over time. The simple model, however, assumes the arrival of a new wind direction, which will result in the wind turbine being out of yaw with the wind. The advanced detection by the LIDAR tells the control system what new direction the incoming wind will have (Fig. 1). A range-gated LIDAR could also tract the new

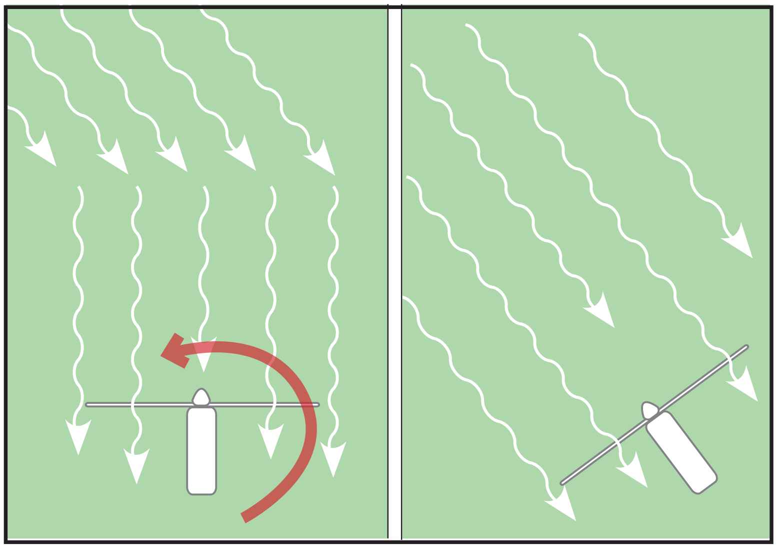

Figure 1 | Change in Yaw Anticipates Change in Wind Direction mass of air as it approaches to determine not only the time of arrival but also the length of time (or the depth) the new direction can be anticipated at the WT. Voting logic can determine when sufficient energy gain (or damage avoidance) can be expected and whether it warrants the rotation in yaw of the nacelle to the new angle. The sensor can tell the control system whether and when the wind may shift back to support that decision process. With the input from the wind sensor, the control system can determine when it should have the nacelle in place at the new angle, what the response time window is, and even the best (most efficient) yaw rate to make the adjustment.

Wind Magnitude

To understand how precision wind magnitude measurement in advance can improve WT operations and profitability, one must separate the effects into the two primary regions of wind turbine performance. At the low end of the WT performance curve, the output increases with wind speed from zero to approaching the rated max power. In this region, the WT can take advantage of additional efficiency at any given wind speed. When the control system knows exactly what wind velocity is coming ahead of time and when it will arrive, the controls can make precise blade pitch adjustments to achieve the highest possible lift, resulting in greater torque. In the absence of the LIDAR data, nominal (overly conservative) pitch settings can be made to prevent vibration and stall at gust speeds, but then will be suboptimal for all the periods of lower wind speed before and after the gust and will not achieve optimum energy production performance. In light and variable wind conditions, when the differences and periods between gusts and non-gust conditions are sufficient to warrant blade pitch adjustments, an active control system can improve outputs in the lower wind speed region. The optimum pitch would be determined by the magnitude of the anticipated wind assuming that the system has yaw alignment with the wind direction. With the blade rotating, the pitch for best lift/drag performance (and to avoid stall) will be determined based on the net vector resultant velocity - a combination of blade rotational (sweep) speed and wind speed. Net velocity, therefore, varies span-wise along the blade as a function of both the wind shear (variance with height above ground) and the local blade velocity (a function of angular rotation and radius). Blade designs attempt to account for nominal span wise velocity and bending moment variation as a function of sweep radius. With current semi-rigid blades, the entire blade angle of attack changes with the pitch changes at the root (hub). With the measurement of incoming wind changes and their duration ahead of the WT (anticipated in advance), a "voting logic" must be created that can determine if the change in magnitude is significant enough and far enough away to make the mechanical adjustments worthwhile. This can be aided by sensing wind magnitude at a number of distances ahead of the WT using range-gated LIDAR mounted on (or in) the nacelle.

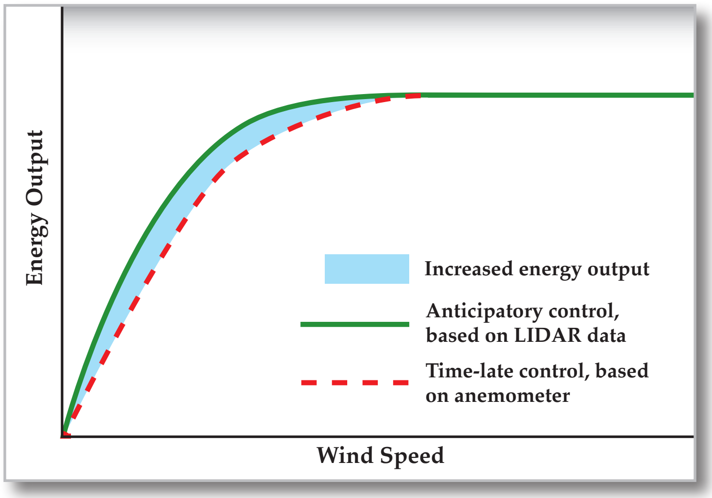

Figure 2 | Look-ahead LIDAR Performance

As shown in Figure 2, performance of a WT with active controls using look-ahead wind sensors for direction (yaw) and magnitude (pitch) optimization is compared with the kind of real world performance a WT would have relative to the theoretical design performance. A substantial percentage of the operation are at sub-optimal yaw and/or blade pitch because the current control system cannot react to wind transients. The higher wind speed region of WT performance is typified by having more than sufficient wind to achieve the maximum rating on the turbine. The control system must react to these conditions, smoothly preventing overloads to the blades and/or structure while maintaining the highest possible mechanical torque transfer to the generator. The system must also have a preset upper limit for maximum wind speed and stress on the WT and will shut the turbine down if that limit is reached. Because of the controllability in response to precise reading of the incoming wind in time to make adjustments, the LIDAR-based system should be able to operate at higher wind speeds with lower safety factors, which translates to more output. At the extreme upper end of wind speed, the anticipating control system can be linked to the braking system; not only determining better when to apply a brake and potentially less violent braking, but also to confirm that it is safe to release it as soon as possible to return to making electricity.

Wind Shear

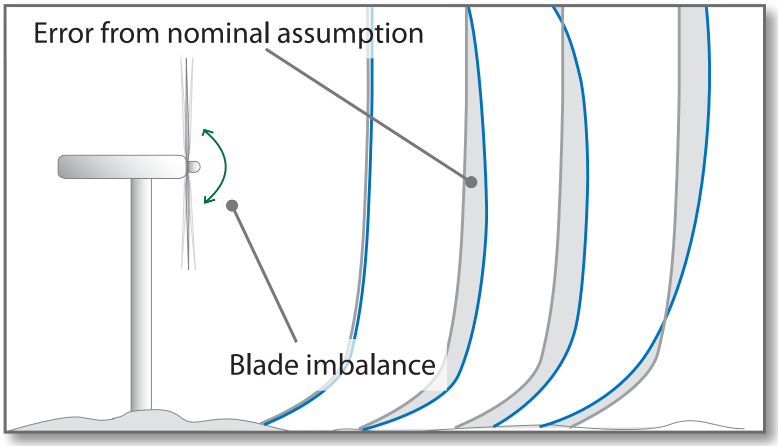

Figure 3 | Wind Shear and Blade Imbalance

In many situations the wind will be significantly different at various heights above the ground or water. This shear creates an imbalance in forces across the blade spans and an undesirable couple at the WT hub as the blades sweep across the velocity profile. As the wind changes in magnitude and direction, the wind profile can be expected to change as well (Fig. 3). A number of damage and wear modes result in premature failures when the WT doesn't adjust for the wind shear. Reacting to the wind shear also affords the opportunity to make more energy than when the blades are adjusted uniformly to some nominal setting or are slow to react to changes in wind conditions. Given input on the wind profile from the LIDAR sensor, the control system may be able to change blade pitch as a function of sweep position if the blades can be controlled independently. If not, it may at least select a more conservative operating condition, giving up some energy performance for reduced risk of damage or wear. The important thing is knowing the wind shear ahead of the WT with enough warning to make one of these choices, provided the necessary control system and capable WT. Current practice of measuring wind profiles with anemometer towers located in the vicinity (but not precisely upstream of each WT) gives some indication of shear, but clearly not as accurately nor as dynamically as can be determined for each WT by a co-axial LIDAR.

Combinations of Direction and Magnitude Change

The dynamic nature of wind is composed of combinations of the two simplistic change inputs above plus the wind shear and smaller turbulent and three-dimensional instability. The sensor will be able to present the three dimensional wind field data with high resolution and update rate. The control system has to sort through that information and determine appropriate actions with combinations of nacelle yaw and blade pitch (potentially as a function of angular location). The WT mechanical system is very large and the response time for adjustments (at practical force levels) is significant. This has to be accounted for in the control response.

Control System Response

The simple cases above show that the control system has to be able to optimize the orientation and aerodynamics of the blades with the incoming wind to either produce the most power or to minimize stresses, damage, and unscheduled maintenance. To do this efficiently, the system will not only have to decide which adjustments are appropriate to the changing conditions, but also have to provide the logic (see Fig. 5 on pg. 7) to decide when the structure and duration of the approaching wind justifies making mechanical adjustments and timing such changes to be in phase with the wind conditions.

The figure on the following page (Figure 4) shows measurements and displays for a wind magnitude increase and the resulting mechanical stress. The assumption in these examples is a sensed incoming change in wind magnitude without a direction change. The first graph shows a snapshot of the incoming wind as a function of range (range-gated) from the LIDAR sensor. The outer range-gate outputs from the forward-looking LIDAR show the tracks that would result from two typical conditions. Firstly (yellow), an incoming gust of rather short duration, and secondly (green), the leading edge of a longer-term change in magnitude.

The next two figures plot the nominal blade pitch and the notional stress on the system as a function of time rather than range (which are linked by the wind speed). The control system voting logic looks at the wind profile as it approaches and determines that, in the first case, it is not worth making a blade pitch change because of the magnitude and duration of the gust. The system over-stresses somewhat and briefly, but within safe margins, and rides through the gust, returning to normal. In the second case, a longer-term wind magnitude increase triggers a pitch change enough ahead of the arrival of the higher velocity air to accommodate the control and mechanical lag time. The control response can be designed to ensure that (since wind changes are not instantaneous), after a brief

Figure 4 | Wind Speed Change and Pitch Control

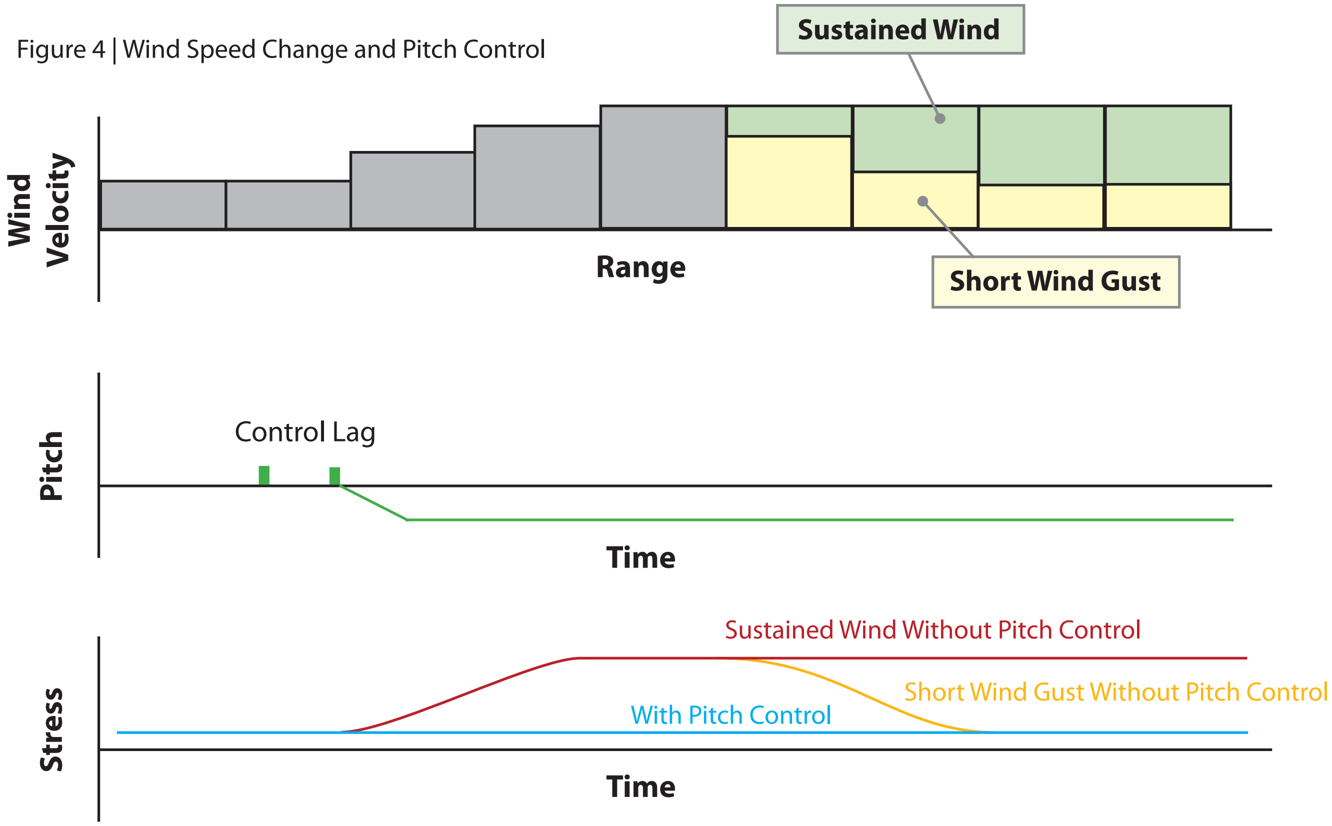

transition, the stress level on the system can be kept at either a normal or an acceptably higher level during the wind event. Reducing the blade pitch will control stress/strain on the blades and ensure that stall conditions along the blades are avoided. It is assumed that wear, mechanical damage, and reliability are progressive and result in cumulative effects over time and cycles up until some catastrophic damage occurs. Therefore, it is expected that maintenance requirements track with the net accumulated stress conditions under which the WT has been operated. The red line on the stress curve represents the result if control action is not taken to the second longer-term wind event such as would occur awaiting actions if time-late averaged data was being used as a basis for adjustments. While not shown, clearly the look-ahead sensor gives the system the opportunity to completely feather the blades and/or apply a brake before wind that would cause failures or blade strikes reaches the turbine.

Voting Logic

An example of the kind of logic for the controls that could be developed is given in the simplified flow diagram (Figure 5) that follows. Once a look-ahead LIDAR senses an increase in wind speed, this information passes through a series of decision steps starting with whether the magnitude of the approaching wind exceeds stress/strain limits on the WT. Then it must be decided whether the increase goes beyond safe operational limits on the WT. If it could cause damage, then emergency actions to feather the blades and/or stop the rotation can be made with sufficient warning time before any damage can occur. If the level is not expected to require emergency stoppage, then the decision process assesses whether the increase is of sufficient duration to justify the control changes as discussed above. A mechanical vibration or strain sensor could operate in conjunction with the look-ahead wind measurement LIDAR to over-ride decisions if actual stress experienced at the WT is higher than predicted.

A similar series of control reaction figures could be shown to illustrate how an approaching reduction in wind magnitude can be accommodated by increasing pitch to maintain lift. Under low wind conditions the voting logic would decide to make blade pitch adjustments when the energy gained offsets the effort to execute the control actions. Look-ahead wind sensing provides time to react to changing wind conditions, catching the wind energy as it approaches rather than waiting for the wind to pass, due to the limitations of the current time-lagged sense-control systems. Failure to make the timely accommodation for reduced wind conditions results in period of reduced capacity factor (normalized energy output).

Figure 5 | Logic Flow Diagram

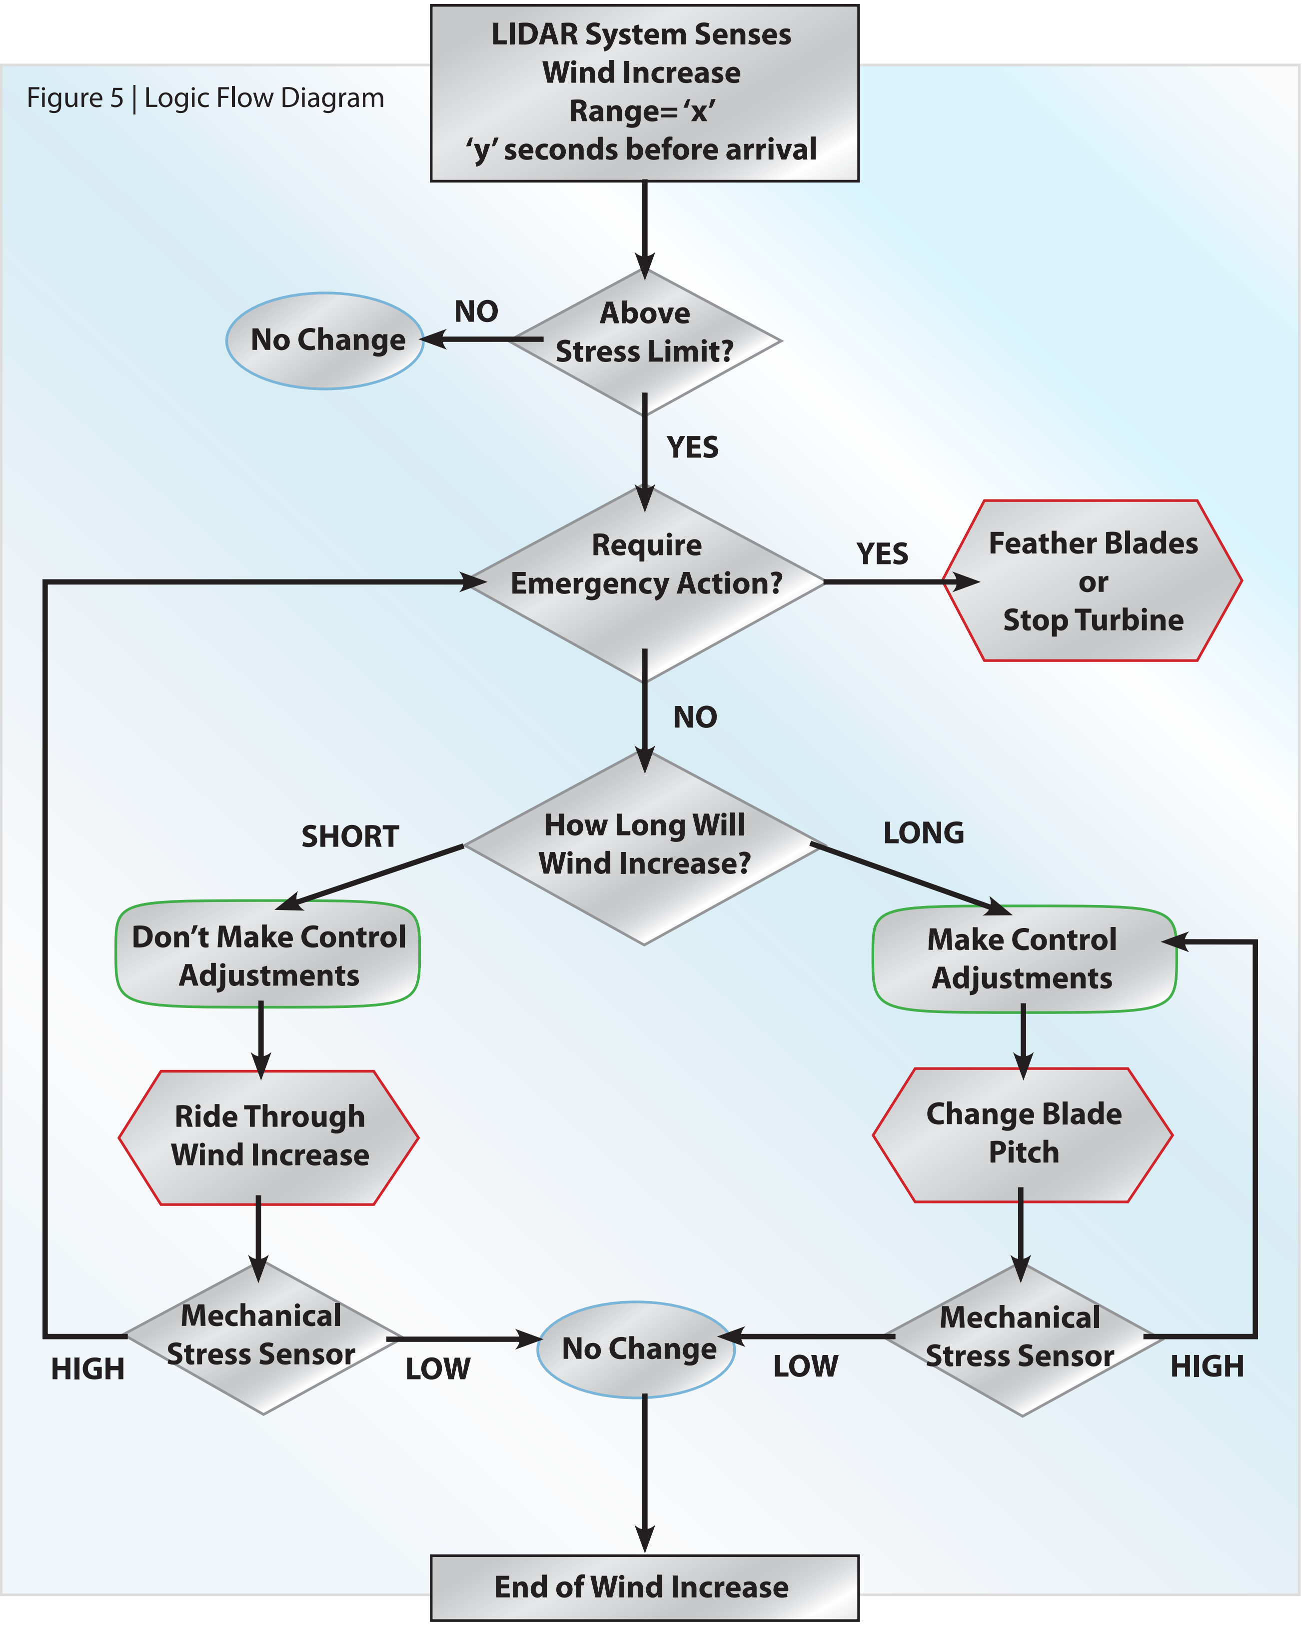

The sense and control response curves for an incoming directional shift in the wind trigger yaw changes to align the WT with the flow. This alignment has a strong effect on the efficiency of the WT in capturing the wind energy and also for limiting vibration occurring from the cross flow and the resulting imbalance in forces at the blades. The graphs below (Fig. 6) illustrate a yaw response to wind direction change sensed ahead by a LIDAR sensor. The top row of dials show the sensed direction of the wind as a function of time at a fixed range (x meters) ahead of the WT. The second row shows the wind direction as sensed at the WT over the same time periods. The nacelle yaw is shown for a LIDAR controlled system (blue) and for a time-late system relying on the measurements at the

WT only (red). As expected, the reduced performance of the WT when out of yaw shows as a reduced output (normalized as capacity factor i.e. % of potential across the swept area). The cumulative gains from avoiding the lags in yaw reaction are expected to result in significant gains for wind farm operation over time.

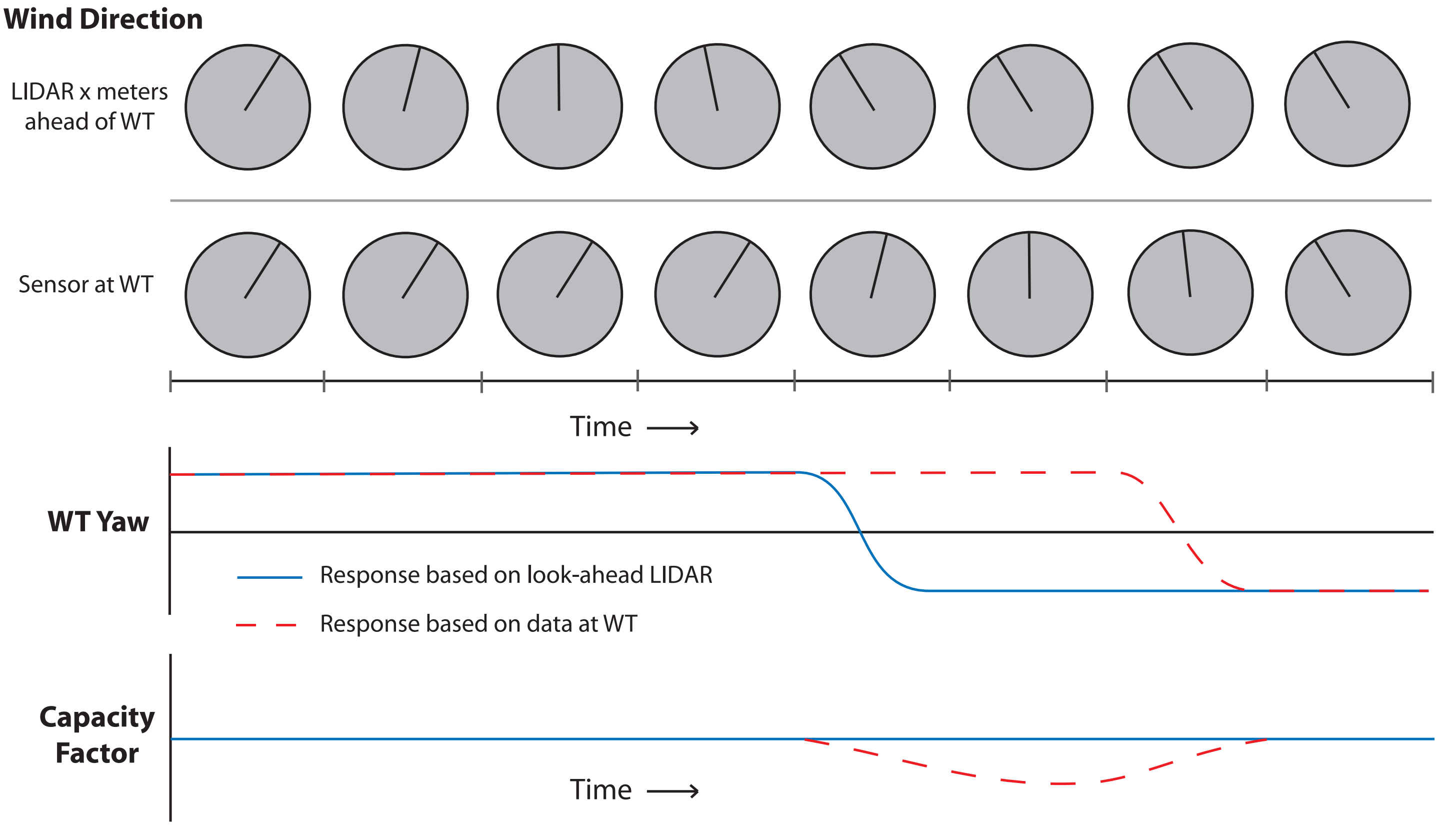

Figure 6 | Look-Ahead LIDAR Response to Wind Change

Quantum Leap

So, by utilizing the ability to remotely sense the wind, a control system will know the speed, direction, and duration of the wind sufficiently ahead of its arrival to make deliberate mechanical adjustments needed for making more energy and controlling stress and strain on the wind turbine components. As a direct result, the long-awaited development of a practical LIDAR for field use enables a quantum leap in profitability for wind turbines; not only by increasing output potential, but also by reducing cost of operation through reduced maintenance and downtime. Ultimately, the introduction of this next generation of sophisticated controls will influence more economical design and more efficient operational guidelines for wind turbines.