The Price of a Yaw: Optimal Deadbands and Filters for Wind-Turbine Yaw Control, Calibrated on SCADA Data

A Random-Walk Wind with Fast Turbulence, the Fourth-Root Deadband, and an Average Measured in Minutes

Working Paper · Miscellaneous

T. Zamrik | 09-NOV-2022

Abstract. A wind turbine must keep its rotor facing a wind whose direction never stops moving, and every yaw that realigns it wears the yaw system. Using a year of 10-minute records from the six turbines of the Kelmarsh wind farm, this paper measures what sets the right policy: below rated power the wind direction wanders like a random walk, at 11.2 degrees per square-root hour on average but at a rate that changes with the weather, with 8.2 degrees of fast turbulence on top, and misalignment costs power as $\cos^k$ with $k = 2.59$. With the turbulence idealised as white noise, the optimal controller separates into a Kalman-Bucy filter whose time constant is $\eta\sqrt{2\tau_u}/s$, 3 to 8 minutes for turbulence correlation times of 10 to 60 seconds, and a deadband of half-width $(6Ks^2/a)^{1/4}$ on the filtered misalignment. An ideal band of 8.9 degrees run on the site's own varying wander reproduces the turbines' bursts of yaws, 40.9% of them back to back against 41.1% in the records and 12.4% under a steady wander, so their controller behaves like that band. It reveals a price of 2.3 kWh per yaw, at which the cost of misalignment and wear together is 1.0% of below-rated energy; a band that widens with the fourth root of the current wander rate lowers it by 18%, and a 30-second average instead of a filter holds about three times the misalignment for the same number of yaws.

1. Introduction

A horizontal-axis wind turbine produces most when its rotor faces the wind squarely. The nacelle that carries the rotor sits on a yaw bearing and is turned by electric yaw drives, and a controller decides when to turn it and how far, from a wind vane and an anemometer mounted on the nacelle behind the rotor [2]. The decision is harder than it sounds. The wind direction is never still: over hours it swings with the weather, over minutes it drifts, and over seconds it flickers with every gust. A controller that followed every change would turn the nacelle all day and wear out the yaw system; one that ignored them would leave the rotor misaligned and lose energy.

Yaw control is one of the older loops in a turbine's controller, and its usual form is simple: average the vane signal over some time, and yaw when the average leaves a deadband [3]. The two numbers that define it, the averaging time and the width of the band, are usually set by experience. This paper asks what they should be, given what the wind at a real site does.

Both sides of the trade-off are real. A misaligned rotor sweeps less of the wind's momentum, and its power falls roughly as a power of the cosine of the misalignment angle; misalignment also changes the loads on the blades and tower, sometimes for the better [4]. The vane that reports misalignment sits in the rotor's wake, and its readings can be biased: turbines fitted with a nacelle-mounted lidar, which measures the incoming wind ahead of the rotor, have been found to run with offsets the vane did not reveal [5]. Deliberate misalignment is even used on purpose, to steer a turbine's wake away from the turbines behind it [6], a different objective from the one here.

On the other side, every yaw costs something. The yaw drives start and stop, the brakes engage, the bearing turns under load, and each event consumes a little of the life of parts that are expensive to replace at the top of a tower. That cost is not a line in any data sheet, and it is the quantity this paper sets out to reveal.

Keeping a drifting quantity close to a target, when every correction carries a fixed cost, is a classical problem. For cash balances it gives the policy of Constantinides [8]: do nothing inside a band, and pay the fee to return to the target when the balance reaches its edge. For a Brownian motion with a fixed cost of intervention it was solved completely by Harrison, Sellke and Taylor [9], and the theory of such impulse control is now textbook material [11], [12]. With a quadratic loss the width of the optimal band grows only as the fourth root of the fixed cost, a law Dixit derived for prices that are changed only occasionally [10].

The yaw problem is an instance of exactly this, but the quantity being controlled is not observed cleanly: the vane sees the slow drift of the wind and its fast turbulence together, and must tell them apart before it acts. The ingredients that decide the answer, how fast the wind direction drifts, how strongly it flickers and how much a degree of misalignment costs, are measurable from the records every modern turbine keeps, and are rarely measured.

The paper makes five contributions. First, it measures the wind direction at a real site and shows that, at the scales where yawing pays, it is a random walk with a varying rate, not a mean-reverting process, with the mean reversion confined to fast turbulence. Second, it states the optimal yaw policy in closed form: a deadband of half-width $(6Ks^2/a)^{1/4}$ with full information (Theorem 1), unchanged in form when the wander rate varies (Theorem 2), and, when turbulence hides the slow direction, a Kalman-Bucy filter with time constant $\eta\sqrt{2\tau_u}/s$ followed by the same band (Theorem 3 and Proposition 2).

Third, it shows from the records that the Kelmarsh turbines behave like an ideal 8.9-degree band on the site's own wander, which reproduces the bursts in their yaws, and reads off the price of a yaw that such a band implies, 2.3 kWh (Corollary 1). Fourth, it shows that a band which widens with the fourth root of the current wander rate cuts the cost of control by 18% at that price (Proposition 1). Fifth, it verifies every formula against second-by-second simulation, and shows what an average over seconds instead of minutes costs.

2. The Wind at Kelmarsh

Kelmarsh wind farm, in Northamptonshire, England, has six Senvion MM92 turbines rated at 2.05 MW, with 92-metre rotors on hubs of 68.5 and 78.5 metres, in operation since April 2016. Its operator released the 10-minute SCADA records of the farm from 2016 to 2021 under an open licence [1]; this paper uses the calendar year 2019, a complete year for all six turbines. Each record holds the 10-minute mean of the wind speed and of the wind direction, the nacelle position and the position of the wind vane relative to the nacelle, the active power, and for most signals the minimum, maximum and standard deviation within the interval.

The analysis keeps the records below rated power, with wind speed between 4 and 11 metres per second and power above 100 kW, where misalignment costs energy. That leaves 37,100 operating hours across the six turbines, 5,912 to 6,462 per turbine, at a mean power of 693 kW. The wind direction in the records is the nacelle position plus the vane reading, to within about 2 degrees root mean square; the vane reading alone is the misalignment the controller sees.

The wind direction wanders like a random walk. *Method.* For lags from 10 minutes to 6 hours, the squared change of the 10-minute mean wind direction, wrapped to ± 180 degrees, is averaged over all pairs of operating records, for each turbine and pooled, and fitted by $V(h) = c_0 + s^2 h + \beta h^2$: noise from the 10-minute averaging, a random walk, and a slow veer. The local rate for each interval is estimated from the squared 10-minute changes over the surrounding hour, less the noise.

Parameters. 36 lags; six turbines; 37,100 operating hours.

Result. The pooled fit gives $s = 11.16$ degrees per square-root hour, $c_0 = 7.9$ degrees squared and a veer of $\sqrt{\beta} = 3.1$ degrees per hour; per turbine s lies between 10.94 to 11.53. The root mean square change is 5.6 degrees after 10 minutes, 11.9 after an hour and 33.2 after six, growing close to the square root of the lag with no sign of levelling off. The local rate is far from constant: its quartiles are 10, 57 and 148 degrees squared per hour, its mean 126, and in 21.8% of intervals it is below what the averaging noise lets one resolve.

Error. The six turbines agree on s to within 5%. A mean-reverting direction with a reversion time of a few hours would make the curve bend over within the six hours shown; it bends upwards instead, the signature of the veer.

The fast turbulence inside ten minutes. *Method.* The within-interval standard deviation of the wind direction, recorded with every 10-minute mean, is averaged over the operating records of each turbine.

Parameters. The same records as the variogram.

Result. The direction's standard deviation within 10 minutes is 8.39 degrees pooled, 7.84 to 9.02 degrees per turbine. A random walk at $s = 11.16$ would contribute only 1.9 degrees of spread inside 10 minutes, so almost all of it is fast turbulence, with standard deviation $\eta = 8.18$ degrees once the random walk's share is removed.

Error. The records give the size of the turbulence but not its time scale, which would need readings every second. The correlation time τ_u is therefore varied from 10 to 60 seconds throughout, the range the integral scales of lateral turbulence near hub height suggest [7].

The power lost to misalignment. *Method.* An aligned power curve is built from the median power in 0.25 m/s bins of wind speed, using records with a vane reading below 2 degrees and no yaw inside the interval. Each record's power is divided by the curve at its wind speed, and the median ratio is taken in 1-degree bins of the vane reading. The exponent k in $\cos^k$ is fitted by least squares to the logarithms over 0 to 10 degrees, with a bootstrap over records for its interval.

Parameters. 179,187 records with no yaw inside, wind speed 4 to 11 m/s; 200 bootstrap samples.

Result. $k = 2.59$, with a 95% interval of 2.30 to 2.87. At a vane reading of 9.5 degrees the power is 2.9% below the aligned curve. At the mean power of 693 kW this gives $a = 0.2735$ kW per degree squared.

Error. Within 10 degrees the ratios follow $\cos^k$ closely, between $\cos^2$ and $\cos^3$. Beyond about 11 degrees they turn and rise, to 1.17 at 24 degrees, which no real misalignment can cause; the large readings are therefore not large misalignments, and the fit is confined to the range where the vane is credible (Remark 1).

How often the turbines yaw. *Method.* An interval is flagged as containing a yaw when the nacelle position ranges over more than 1 degree within it; separately, steps of more than 5 degrees between consecutive 10-minute means are counted. The misalignment the vane reports is summarised by its root mean square.

Parameters. Operating records of all six turbines, 2019.

Result. The turbines have 28.0 flagged intervals per operating day, 26.6 to 29.7 per turbine, and 24.4 steps of more than 5 degrees. The vane reports a root mean square misalignment of 7.14 degrees, with a mean close to zero on every turbine.

Error. Flagged intervals undercount yaws: two yaws in one interval count once. The count is therefore used only through the calibration of Algorithm 1, which flags simulated intervals in exactly the same way.

Remark 1 (Why the vane's misalignment is not the true misalignment)*.* The vane sits behind the rotor, in air the rotor has already turned, and its reading mixes the true misalignment with errors of its own. Two pieces of evidence in the records show that the errors are large.

The first is the tail of the power ratio in the computations: readings above about 17 degrees come with power above the aligned curve, which a rotor misaligned by that much cannot deliver. The second comes from the model. The band that reproduces the turbines' yaws (Algorithm 1) holds a true slow misalignment of 3.65 degrees root mean square, while the vane reports 7.14; after the averaging noise is removed, the vane's own error is about 5.8 degrees root mean square. Lidar campaigns have found offsets of this kind on other turbines [5]. For this paper the consequence is practical: the cost of misalignment is fitted only where the vane is credible, and no claim rests on the misalignment the vane reports.

3. The Model

Definition 1 (The wind direction: a slow random walk and fast turbulence)**.** The wind direction at hub height, in degrees, is

$$\Theta_t = W_t + U_t,$$

the sum of a slow part W_t and a fast part U_t. The slow part is a random walk, $dW_t = s\, dB_t$, with wander rate s in degrees per square-root hour; from Theorem 2 on, the rate may vary with time, $s = s_t$. The fast part is stationary turbulence with standard deviation η and correlation time τ_u, modelled as an Ornstein-Uhlenbeck process $dU_t = -U_t\, dt/\tau_u + \eta\sqrt{2/\tau_u}\, dB'_t$, with B and B' independent Brownian motions.

The split is a statement about time scales. Over the minutes and hours between yaws the direction drifts and does not come back; over seconds it flickers around its local value and does. The first is what a yaw can follow; the second is what it should ignore. Only the size of the fast part, η, can be measured from 10-minute records; its time τ_u is treated as a parameter between 10 and 60 seconds, the range of integral time scales of lateral turbulence near hub height.

Definition 2 (The nacelle, its yaws and the misalignment)**.** The nacelle heading ϕ_t is constant between yaws and is moved to a new heading at each yaw time $T_1 < T_2 < \dots$, the times and the new headings chosen by the controller from what it has measured. The slow misalignment is $x_t = W_t - \phi_t$, and the misalignment the rotor feels is $x_t + U_t$.

A yaw is treated as instantaneous. At 0.5 degrees per second a typical yaw of about 9 degrees takes some 18 seconds, against tens of minutes between yaws, and the simulations of Algorithm 2 turn the nacelle at its true rate to check that the approximation costs nothing visible.

Definition 3 (The long-run cost of misalignment and of yawing)**.** Below rated wind speed a rotor misaligned by an angle e delivers $\cos^k e$ of its aligned power, with exponent k. For small angles the power lost per unit time is $a\, x_t^2$, with

$$a = \bar{P}\,\frac{k}{2}\Big(\frac{\pi}{180}\Big)^2 \quad \text{kW per degree squared,}$$

where $\bar{P}$ is the mean power, and each yaw costs K, in kilowatt-hours, for the wear it causes. A policy is judged by its long-run cost rate

$$C = \limsup_{T\to\infty} \frac{1}{T}\,\mathbb{E}\Big[\int_0^T a\, x_t^2\, dt + K\, N_T\Big],$$

where N_T counts the yaws up to time T.

The fast part U_t does not appear in the cost. Its effect on power is the same for every controller, because no yaw can follow it, and it is already inside the aligned power curve against which the loss is measured. The price K is the one quantity the records cannot show directly: it stands for the wear on the yaw drives, bearing and brakes, expressed in the energy the operator would give up to avoid one yaw.

Assumption 1 (Standing assumptions)**.** The turbine runs below rated power; misalignments are small enough that $1-\cos^k e \approx \frac{k}{2}e^2$; a yaw takes negligible time compared with the time between yaws; and, until Theorem 2, the wander rate s, the turbulence and the price K do not change.

Above rated power the pitch controller holds the output at its limit, so a small misalignment costs nothing and the problem disappears; the analysis is therefore restricted to wind speeds between 4 and 11 metres per second, where the Kelmarsh turbines spend most of their operating hours. The quadratic loss is accurate to within a few per cent of itself up to about 15 degrees, beyond the misalignments any sensible band allows.

Remark 2 (Why a random walk, and not a mean-reverting wind)*.* A model in which the wind direction always returns to a mean, an Ornstein-Uhlenbeck process for the slow part, would make the variogram level off at twice the process's variance after a few reversion times. Over the six hours measured it does not level off; it grows almost as the square root of time.

The mean reversion is real, but it lives elsewhere: inside each 10-minute interval the direction spreads by 8.4 degrees, where a random walk would spread by 1.9. That is turbulence returning to the local direction within seconds. The model of Definition 1 keeps both: a random walk where the yaw decisions are made, and a fast mean-reverting part that the controller must learn to ignore.

4. The Deadband with Full Information

Lemma 1 (The cost of a band policy)**.** Let the slow misalignment be known and follow a random walk with rate s, and let the controller yaw to the wind, resetting x to zero, whenever $|x_t|$ reaches b. Then the long-run yaw rate and the long-run mean square misalignment are

$$\nu = \frac{s^2}{b^2}, \qquad \overline{x^2} = \frac{b^2}{6},$$

and the long-run cost is $C(b) = a\,b^2/6 + K\,s^2/b^2$.

Proof. Between yaws x is a Brownian motion with variance rate s^2 started at zero, and the yaws are renewals. The expected time $T(x)$ to leave $(-b, b)$ solves $\frac{1}{2}s^2T'' = -1$ with $T(\pm b) = 0$, so $T(x) = (b^2-x^2)/s^2$ and a cycle lasts b^2/s^2 on average. The expected integral $Q(x)$ of x^2 until exit solves $\frac{1}{2}s^2Q'' = -x^2$ with $Q(\pm b) = 0$, so $Q(x) = (b^4 - x^4)/(6s^2)$ and a cycle accumulates $b^4/(6s^2)$. By the renewal-reward theorem the long-run yaw rate is s^2/b^2 and the long-run mean of x^2 is $(b^4/6s^2)/(b^2/s^2) = b^2/6$. □

The two halves of the cost pull in opposite directions: a wide band holds more misalignment, as the square of its width, and a narrow one yaws more often, as the inverse square. Neither depends on anything but s, b and the two prices; the renewal argument is the standard one for Brownian control problems [11].

Theorem 1 (The fourth-root deadband)**.** Under Assumption 1, with the slow misalignment known, the policy that minimises the long-run cost among all yaw policies is the band policy of Lemma 1 with half-width

$$b^* = \Big(\frac{6\,K\,s^2}{a}\Big)^{1/4},$$

which yaws at the rate $\nu^* = s\sqrt{a/(6K)}$ and costs

$$C^* = \sqrt{\frac{2\,a\,K\,s^2}{3}},$$

half of it lost power and half of it wear.

Proof. Among band policies, $C(b) = ab^2/6 + Ks^2/b^2$ from Lemma 1 has a single minimum where $ab/3 = 2Ks^2/b^3$, which is b^*; substituting gives C^* and shows the two terms equal. That no other policy does better is shown in the derivation below. ☐

The fourth root is the signature of a fixed cost against a quadratic loss, and it makes the band remarkably robust: a price of a yaw sixteen times higher only doubles the band. The same law appears in the economics of menu costs [10], in the impulse control of Brownian motion [9], and in cash holdings [8], where a price, a stock or a balance drifts and every correction costs a fixed fee.

The long-run problem is solved by a relative value function h and a constant C that satisfy, inside the band,

$$\tfrac{1}{2}s^2h''(x) + a\,x^2 = C, \qquad |x| < b, \tag{4.1}$$

together with the condition that yawing is never cheaper than waiting, $h(x) \le K + \min_y h(y)$, with equality at the edge of the band. By symmetry h is even and its minimum is at zero, so a yaw resets the misalignment to zero. Integrating twice,

$$h(x) = h(0) + \frac{C\,x^2 - a\,x^4/6}{s^2}. \tag{4.2}$$

Two conditions fix b and C. Value matching, $h(b) - h(0) = K$, says that at the edge waiting and yawing cost the same; smooth fit, $h'(b) = 0$, says that the edge is chosen optimally. The second gives $C = ab^2/3$, and the first then gives

$$\frac{ab^4/3 - ab^4/6}{s^2} = K, \qquad b^4 = \frac{6Ks^2}{a}, \tag{4.3}$$

which is Theorem 1 again, with $C = ab^2/3 = \sqrt{2aKs^2/3}$.

It remains to check that no policy does better. Inside the band h is smooth and satisfies the equation with equality; the function $h(x) - h(0) = (Cx^2 - ax^4/6)/s^2$ increases from 0 to K on $[0, b]$, so a yaw from inside the band would cost more than waiting; and a policy that let $|x|$ exceed b would face $\frac{1}{2}s^2h'' + ax^2 > C$ there, a higher running cost. Applying Itô's formula to $h(x_t)$ between yaws and summing the jumps at the yaws shows that every admissible policy has long-run cost at least C, the standard verification argument for impulse control of a Brownian motion [9], [12]. The band policy attains it.

Corollary 1 (The price of a yaw, revealed by the yaw rate)**.** If a turbine's controller behaves like a band of half-width b on a wind that wanders at rate s, it yaws at the rate $\nu = s^2/b^2$, and that band is the optimal one of Theorem 1 exactly when the price of a yaw is

$$K = \frac{a\,b^4}{6\,s^2} = \frac{a\,s^2}{6\,\nu^2}.$$

Proof. Solve Theorem 1 for K and use $\nu = s^2/b^2$ from Lemma 1. ☐

The corollary turns the problem round. The price of a yaw is hard to know in advance, since it bundles wear, fatigue and maintenance, but a controller that has been tuned by experience embodies one. Measuring the wander rate and counting the yaws reads it off, in the same unit as the energy it trades against.

Corollary 2 (A cost per degree as well as per yaw)**.** If each yaw also costs c per degree turned, the long-run cost of a band of half-width b is

$$C(b) = \frac{a\,b^2}{6} + \frac{(K + c\,b)\,s^2}{b^2},$$

and the best band solves $a\,b^4 - 3\,c\,s^2\,b - 6\,K\,s^2 = 0$. With no fixed cost it is $b = (3cs^2/a)^{1/3}$.

Proof. A yaw from the edge turns the nacelle by b, so the cost per yaw is $K + cb$; Lemma 1 gives the rest, and setting $C'(b) = 0$ gives the quartic. With $K = 0$ it reduces to $ab^3 = 3cs^2$. ☐

A cost proportional to the angle, from the energy the yaw drives draw and the production lost while turning, gives a cube-root law instead of a fourth-root one, the classical result for proportional transaction costs [8], [11]. For a turbine the fixed part dominates: the drives use a few kilowatts for a few seconds, a small fraction of a kilowatt-hour.

Theorem 2 (A wander that varies: the band runs on the variance clock)**.** Let the wander rate s_t vary with time, independently of the random walk it scales, with long-run mean square $\overline{s^2}$. Under a band of fixed half-width b the long-run yaw rate is

$$\nu = \frac{\overline{s^2}}{b^2},$$

the long-run mean square misalignment is still $b^2/6$, and the band that is optimal among fixed bands is Theorem 1's with s^2 replaced by $\overline{s^2}$.

Proof. By a time change onto the variance clock [14]; see the derivation that follows. ☐

The theorem says that a varying wind does not change what a fixed band costs on average; it changes when the costs fall. A fixed band yaws in bursts when the wind wanders fast and sits still when it is steady, and the bursts in the Kelmarsh records are exactly this, as the computations show.

Measure time by the variance the random walk has accumulated, $A_t = \int_0^t s_u^2\,du$. By the theorem of Dambis, Dubins and Schwarz [14] the slow direction run on this clock, $\widetilde{W}_\theta = W_t$ with $\theta = A_t$, is a standard Brownian motion. A band policy acts on the path, not on the clock: it yaws when $|x|$ reaches b, whatever the clock reads. So in clock time the misalignment is exactly the process of Lemma 1 with $s = 1$, which yaws at the rate $1/b^2$ per unit of clock and has mean square $b^2/6$ in clock time.

Converting back to real time, the number of yaws up to time t is the number up to clock time A_t, and

$$\frac{N_t}{t} = \frac{N_{A_t}}{A_t} \cdot \frac{A_t}{t} \longrightarrow \frac{1}{b^2} \cdot \overline{s^2}. \tag{4.4}$$

For the mean square, real time spends $1/s_t^2$ units per unit of clock where the rate is s_t. Because s is independent of the Brownian motion, where the misalignment happens to be within the band is independent of how fast the clock runs there, and

$$\frac{1}{t}\int_0^t x_u^2\,du = \frac{\int x^2\,dA/s^2}{\int dA/s^2} \longrightarrow \frac{b^2}{6}. \tag{4.5}$$

The cost of a fixed band is therefore $ab^2/6 + K\,\overline{s^2}/b^2$, minimised by Theorem 1 with $\overline{s^2}$ in place of s^2.

Proposition 1 (A band that widens with the fourth root of the wander)**.** If the wander rate s_t changes slowly compared with the time between yaws, the band $b_t = (6Ks_t^2/a)^{1/4}$, chosen afresh for the current rate, has long-run cost

$$C_{\text{ad}} = \sqrt{\frac{2aK}{3}}\,\bar{s}, \qquad \frac{C_{\text{ad}}}{C_{\text{fixed}}} = \frac{\bar{s}}{\sqrt{\overline{s^2}}} \le 1,$$

where $\bar{s}$ is the long-run mean of s_t and $C_{\text{fixed}} = \sqrt{2aK\overline{s^2}/3}$ is the cost of the best fixed band.

Proof. While the rate is s, Theorem 1 applies with that rate and costs $\sqrt{2aKs^2/3}$ per unit time; averaging over time gives C_{ad}, and the best fixed band costs C_{fixed} by Theorem 2. The ratio is at most one by Jensen's inequality. □

The more unevenly the wind wanders, the more an adaptive band gains: when the direction is steady it can afford a narrow band and hold the rotor close to the wind, and when the direction is restless it widens and saves the yaws a fixed band would spend. The rate itself can be estimated on line, from the squared direction changes over the last half hour or so.

Remark 3 (A veering wind, and why it barely matters). The fitted veer adds a steady drift of about 3.1 degrees per hour to the random walk. With a drift the best band is no longer symmetric, and the best yaw aims slightly ahead of the wind.

For the drift and the price of a yaw at Kelmarsh the best asymmetric band runs from 9.20 degrees on one side to 8.67 on the other, and each yaw overshoots by 0.47 degrees in the direction of the veer. It lowers the cost of control by 0.45%, about 0.005% of the energy, because over the tens of minutes between yaws a drift of a few degrees per hour moves the wind far less than its random wander does. The symmetric band of Theorem 1 loses nothing that matters.

5. The Deadband with Turbulence

Theorem 3 (Separation: a filter and a deadband)**.** Let the controller measure the direction Θ_t but not its parts, and let the turbulence be white, with intensity N: the measurement is $dY_t = W_t\,dt + \sqrt{N}\,dV_t$ for a Brownian motion V. Then the optimal policy estimates the slow direction by the Kalman-Bucy filter

$$d\widehat{W}_t = \frac{s}{\sqrt{N}}\,(dY_t - \widehat{W}_t\,dt),$$

a running average with time constant $\tau_f = \sqrt{N}/s$, and applies the band of Theorem 1 to the estimated misalignment $\widehat{W}_t - \phi_t$, yawing to $\widehat{W}_t$. Its long-run cost is

$$C = \sqrt{\frac{2aKs^2}{3}} + a\,s\sqrt{N},$$

the full-information cost plus an irreducible term from not knowing the slow direction.

Proof. *By the Kalman-Bucy filter [13] and the Gaussian conditional law of the slow direction [15]; see the derivation that follows.* □

The optimal controller does two separate jobs. It estimates where the slow wind is, with a filter that knows nothing about yaws, and it decides when to yaw, with a band that knows nothing about turbulence. Neither job is changed by the other.

The nacelle heading is chosen from the measurements and does not affect the wind, so the conditional law of W_t given the measurements is Gaussian, with a mean $\widehat{W}_t$ and a variance P_t that do not depend on the yaws [13], [15]. The variance solves the Riccati equation of the Kalman-Bucy filter,

$$\frac{dP}{dt} = s^2 - \frac{P^2}{N}, \tag{5.1}$$

whose steady state is $P = s\sqrt{N}$, reached within a few filter times. The gain is $L = P/N = s/\sqrt{N}$, the filter's time constant $1/L = \sqrt{N}/s$.

Conditionally on the measurements, the expected loss is

$$\mathbb{E}\big[a(W_t - \phi_t)^2 \mid \mathcal{F}^Y_t\big] = a(\widehat{W}_t - \phi_t)^2 + aP, \tag{5.2}$$

so the cost splits into a part driven by the estimate and a constant $aP = as\sqrt{N}$ that no policy can remove. The estimate moves only through the innovations $d\nu_t = dY_t - \widehat{W}_t\,dt$, which form a Brownian motion with intensity N, so $d\widehat{W}_t = L\,d\nu_t$ is a random walk with variance rate $L^2N = s^2$: the estimated misalignment wanders exactly as fast as the true one. The remaining problem, to keep $\widehat{W}_t - \phi_t$ small at a fixed price per yaw, is the problem of Theorem 1 with the same s, solved by the same band.

Proposition 2 (An average measured in minutes)**.** If the turbulence is the Ornstein-Uhlenbeck process of Definition 1, its intensity at the time scales of the filter is $N = 2\eta^2\tau_u$, and the optimal averaging time and the irreducible mean square misalignment are

$$\tau_f^* = \frac{\eta\sqrt{2\tau_u}}{s}, \qquad P = s\,\eta\sqrt{2\tau_u}.$$

A running average over τ_a seconds with $\tau_a \ll \tau_f^*$ passes turbulence of variance close to $2\eta^2\tau_u/\tau_a$ into every yaw decision.

Proof. The spectral density of an Ornstein-Uhlenbeck process with variance η^2 and correlation time τ_u is $2\eta^2\tau_u/(1+\omega^2\tau_u^2)$, which equals $2\eta^2\tau_u$ at the low frequencies the filter passes; Theorem 3 then gives τ_f^* and P. An average over $\tau_a \gg \tau_u$ of a process with this spectrum has variance $2\eta^2\tau_u/\tau_a$. □

The formula weighs the two sources of change in the measured direction against each other. Averaging longer removes more turbulence but lets the slow wind move further before the controller notices; the balance falls where the turbulence the average lets through equals the drift it misses. The correlation time τ_u is set by the integral scale of lateral turbulence, which surface-layer spectra put at tens of seconds near hub height [7].

Remark 4 (Yawing part of the way)*.* The filter moves its estimate towards each new reading by only a small fraction of the difference, $L\,dt = dt/\tau_f$, so after a gust it has barely moved, and a yaw goes to the estimate rather than to the latest reading.

A controller that yaws to its latest short average treats every reading as the new truth, and so lands wherever the turbulence happened to be when it fired. The filter's answer is to go only part of the way towards each reading and to let the readings argue among themselves. The intuition of anticipating a wind that comes back survives, but it belongs to the fast turbulence, not to the slow direction.

Remark 5 (Coloured turbulence and the white-noise idealisation)*.* Theorem 3 treats the turbulence as white noise; the turbulence of Definition 1 is an Ornstein-Uhlenbeck process with a correlation time of seconds. The exact filter for that model has two states instead of one, but when τ_u is much shorter than the filter's time constant the difference is small.

The simulations, which use the coloured turbulence, measure how small. At 1.25 yaws per hour the mean square misalignment of the best filtered band is within 5% of Theorem 3's value for every correlation time from 10 to 60 seconds. The best filter time constant in the simulation's grid lies somewhat below the formula's, because the cost is flat near its minimum: a filter a factor of two off costs only a few degrees squared.

Example 1 (The numbers at Kelmarsh)**.** The measured quantities are $s^2 = 126$ degrees squared per hour on average, $a = 0.2735$ kW per degree squared and $\eta = 8.18$ degrees, and the turbines flag 28.0 intervals with a yaw per day.

A steady-wander reading of the yaw count, the band that yaws once per flagged interval, would give a band of 10.3 degrees and a price of 4.2 kWh per yaw. It is wrong, because the wind wanders unevenly and yaws come in bursts that share intervals. The band that reproduces the flagged intervals on the site's own wander is 8.95 degrees; it yaws 1.50 times an hour, 36 times a day, and by Corollary 1 with Theorem 2 it is optimal when a yaw costs $K = 2.32$ kWh.

At that price the best fixed band costs 7.30 kW by Theorem 1, 1.05% of the mean power, half in lost energy and half in wear; the simulation on the site's wander gives 7.12 kW. The adaptive band of Proposition 1 costs 5.81 kW by formula and 5.83 by simulation, 18% less. With turbulence, Proposition 2 puts the best averaging time at 3.3 to 8.0 minutes for correlation times of 10 to 60 seconds, and the irreducible misalignment at 6.8 to 16.7 degrees squared, a loss of 0.27% to 0.66% of the energy that no yaw policy can recover.

6. Computation

Algorithm 1 — From 10-minute records to the deadband and the price of a yaw

Inputs: for each turbine, 10-minute means of wind direction, nacelle position, vane position, wind speed and power, and the minimum and maximum of the nacelle position within each interval. *Output:* the wander rate s, its local values s_t, the fast turbulence η, the exponent k and slope a of the power loss, the yaw rate, the best band and the price of a yaw.

1. Keep the intervals below rated: wind speed 4 to 11 m/s and power above 100 kW.
2. For lags of 10 minutes to 6 hours, average the squared change of the wind direction (wrapped to ± 180 degrees) and fit $V(h) = c_0 + s^2 h + \beta h^2$: c_0 is the averaging noise, s the wander rate, $\sqrt{\beta}$ the veer.
3. Estimate the local rate for each interval from the squared 10-minute changes over the surrounding hour, less c_0.
4. Take η from the within-interval standard deviation of the direction, less the random walk's share.
5. Build the power curve from intervals with $|\text{vane}| < 2$ degrees and no yaw inside; fit k to the median power ratio in 1-degree bins up to 10 degrees; set $a = \bar{P}\,(k/2)(\pi/180)^2$.
6. Flag an interval as containing a yaw when the nacelle position ranges over more than 1 degree within it.
7. Find the fixed band that, run second by second on the site's own local rates, flags the same share of intervals; read the price of a yaw from Corollary 1 with $\overline{s^2}$, as Theorem 2 allows.

Cost. One pass over the records per step, seconds for a year of six turbines; step 7 is a simulation of 27,000 hours of wind, about a minute on many cores. The step that matters is step 7: counting flagged intervals undercounts yaws that come in bursts, and calibrating the band on the same flags removes that bias instead of ignoring it.

Algorithm 2 — Simulating a yaw controller second by second

Inputs: the wander rate s, the turbulence η and τ_u, the yaw rate of the drive (0.5 degrees per second) and a controller: either a running average over τ_a seconds with a deadband d, or an exponential filter with time constant τ_f and a band b. *Output:* the yaw rate and the mean square slow misalignment.

1. Start 1000 independent turbines with the nacelle aligned.
2. Each second, advance the slow direction by a Gaussian step of variance $s^2/3600$ and the turbulence by its exact Ornstein-Uhlenbeck step; the measured direction is their sum.
3. Update the controller's view: the running average of the measured misalignment, or the filter $\widehat{W} \leftarrow \widehat{W} + (\Theta - \widehat{W})/\tau_f$.
4. If the nacelle is not already turning and the view is outside the band, start a yaw towards the average or the estimate.
5. Turn a turning nacelle by at most 0.5 degrees towards its target.
6. After a one-hour burn-in, accumulate the yaws and the square of the slow misalignment $W - \phi$.
7. Repeat for 24 simulated hours and every combination of controller settings.

Cost. $1000 \times 86{,}400$ steps per setting, a few seconds each; 557 settings in parallel take under a minute. With the turbulence switched off and the filter set to follow the measurement, the same code checks Lemma 1.

The band formulas against simulation. *Method.* Algorithm 2 with the turbulence switched off and the controller following the true direction, for bands of 4 to 18 degrees, with the nacelle turning at its true rate of 0.5 degrees per second.

Parameters. $s = 11.16$ degrees per square-root hour; 1000 turbines for 24 hours each per band, after a one-hour burn-in.

Result. For a band of 10 degrees the mean square misalignment is 17.21 against $b^2/6 = 16.67$, and the yaw rate 1.214 per hour against $s^2/b^2 = 1.245$. Across all seven bands the simulated mean square exceeds $b^2/6$ by 2% to 8%, and the yaw rate falls short of s^2/b^2 by 1% to 5%.

Error. Both differences come from the yaw taking time: while the nacelle turns, the wind keeps moving, so each yaw starts a little later and ends a little off. They are largest for the narrowest band, where a yaw's duration is the largest share of the time between yaws, and shrink to about 2% at 18 degrees. Lemma 1 is exact for instantaneous yaws.

The best averaging time against the formula. *Method.* Algorithm 2 with the exponential filter, for time constants from 30 seconds to 15 minutes and bands from 4 to 22 degrees, for each of five turbulence correlation times; for each time constant the mean square misalignment at 1.25 yaws per hour is interpolated along the band.

Parameters. $s = 11.16$, $\eta = 8.18$; $\tau_u = 10$, 20, 30, 45 and 60 seconds; 1000 turbines for 24 hours per setting.

Result. Proposition 2 puts the best time constant at 197, 278, 341, 417 and 482 seconds; the best of the simulated grid is 120, 240, 240, 360 and 360 seconds. The cost rises steeply for shorter filters, by half or more once the time constant is a minute or less, and gently for longer ones.

Error. The grid's best lies 14 to 39% below the formula's value, within a flat minimum where neighbouring grid values differ by only a few per cent. The formula is exact for white turbulence; with the coloured turbulence simulated, the optimum is slightly shorter.

The cost decomposition against simulation. *Method.* The mean square misalignment of the best filtered band at 1.25 yaws per hour, from the simulations of the previous item, against Theorem 3's prediction $s^2/(6\nu) + s\eta\sqrt{2\tau_u}$, the band's share plus the filter's irreducible error.

Parameters. As before; the band's share at 1.25 yaws per hour is 16.6 degrees squared.

Result. For $\tau_u = 10$, 20, 30, 45 and 60 seconds the prediction is 23.4, 26.2, 28.4, 31.0 and 33.3 degrees squared, and the simulation gives 24.5, 26.2, 28.3, 30.8 and 32.3.

Error. The largest difference is 5%, at 10 seconds, where the yaw's own duration adds most. The decomposition is confirmed: the filter's error adds a constant to the misalignment and leaves the band's trade-off unchanged.

An average over thirty seconds. *Method.* Algorithm 2 with a running average over 30 seconds and deadbands from 4 to 24 degrees, against the best filtered band and the best running average of any length, at equal yaw rates.

Parameters. As before, for each of the five correlation times.

Result. With a 30-second average no deadband up to 24 degrees brings the yaw rate below 1.5 to 2.6 per hour, depending on the turbulence, because gusts keep crossing the band. At 3 yaws per hour it holds 47 to 68 degrees squared of misalignment, against 14.6 to 22.5 for the filtered band: 3.0 to 3.6 times as much. The best running average is always 4 to 8 minutes long and comes within 0.8 to 2.8 degrees squared of the filter at 1.25 yaws per hour.

Error. Each value is interpolated along the deadband from simulations of 24,000 turbine-hours. The conclusion does not depend on the filter's shape: an average of the right length does almost as well as the exponential filter; an average of the wrong length cannot be rescued by any deadband.

The bursts of yaws, reproduced by the site's own wander. *Method.* In the records, the gaps between consecutive intervals flagged with a yaw are counted in 10-minute steps, within stretches of continuous operation. The same count is made on two simulations: an ideal band run second by second on the site's own local wander rates, with its width chosen to flag the same share of intervals, and the same band idea on a steady wander at the average rate.

Parameters. 31,292 gaps in the records; stretches of at least 6 hours; band 8.95 degrees on the site's wander, 10.38 on the steady one.

Result. In the records 41.1% of gaps are a single step, yaws in consecutive intervals; the band on the site's wander gives 40.9%, the band on a steady wander 12.4%. The next shares are 16.4%, 9.3%, 6.6% in the records against 18.1%, 10.1%, 6.0% on the site's wander and 18.1%, 15.1%, 11.8% on a steady one. The mean gaps are 4.39, 4.59 and 5.08 steps.

Error. Only the band's width was fitted, to the overall rate; the shape of the distribution was not. That the site's own wander reproduces it to within about two points, where a steady wander misses by a factor of three, is independent evidence that the turbines behave like a band on a random walk whose rate varies, and that the bursts are the wind's.

An adaptive band under the site's wander. *Method.* The simulation of the previous item, with the band chosen afresh in each 10-minute interval as $(6Ks_t^2/a)^{1/4}$ from the local rate, against the fixed band that the turbines' yaws reveal, both at the revealed price.

Parameters. $K = 2.32$ kWh, $a = 0.2735$ kW per degree squared; the site's local rates for 37,100 hours.

Result. The fixed band yaws 1.50 times an hour and holds 13.34 degrees squared, a cost of 7.12 kW. The adaptive band yaws 1.25 times an hour and holds 10.70 degrees squared, a cost of 5.83 kW, 18% less; Proposition 1 predicts 5.81 kW. The fixed band's mean square misalignment of 13.34 matches Theorem 2's $b^2/6 = 13.34$.

Error. Proposition 1 assumes the rate changes slowly compared with the time between yaws; the simulation does not, and still agrees with it to within 0.4%. The ratio of the mean rate to its root mean square at Kelmarsh is 0.796, the whole of the gain.

Source. `compute_numbers.py`: the simulations of Algorithm 2 at the Kelmarsh wander and turbulence, with a turbulence correlation time of 30 seconds; each entry is the least mean square slow misalignment, in degrees squared, reachable at the given yaw rate.

Columns. Yaws per hour; the perfectly informed band of Lemma 1, $s^2/(6\nu)$; the filtered band of Theorem 3; the best running average of any length, with its length; and a running average over 30 seconds.

yaws per hour	perfect band	filter + band	best running average	30-second average
0.50	41.5	55.2	55.3 (8 min)	unreachable
0.75	27.7	39.6	40.9 (8 min)	unreachable
1.00	20.7	32.6	33.8 (8 min)	unreachable
1.25	16.6	28.3	29.5 (8 min)	unreachable
1.50	13.8	25.5	26.9 (8 min)	unreachable
2.00	10.4	21.7	22.8 (8 min)	unreachable
3.00	6.9	18.6	19.4 (8 min)	63.2

The first two columns differ by the filter's irreducible error, 11.8 degrees squared, at every rate, as Theorem 3 says. The best running average is always long, and close to the filter. The 30-second average cannot yaw less than 2.6 times an hour, and when it can be compared it holds about three times the misalignment.

Script. `fig_wander.py`, which draws the variogram of the 10-minute wind direction for each turbine and pooled, with the fitted noise, random walk and veer, and the distribution of the local wander rate over the year.

Shows. The slow direction is a random walk at the scales where yawing pays, and its rate varies widely.

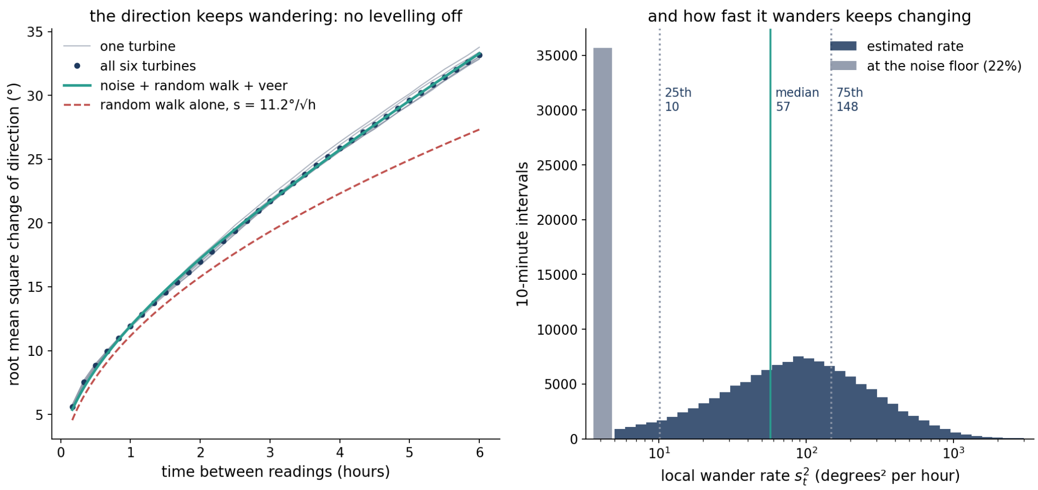

Figure 1: Left: the root mean square change of the wind direction at Kelmarsh against the time between readings, from 10 minutes to 6 hours, for each of the six turbines and pooled, with the fitted noise, random walk and veer and a random walk alone for comparison; right: the distribution of the local wander rate over the year on a logarithmic scale, with the intervals at the noise floor shown apart.

The pooled points grow almost as the square root of the lag and bend slightly upwards, the veer, rather than levelling off as a mean-reverting direction would. The six turbines lie on top of one another. On the right the local rate spreads over two orders of magnitude, with quartiles of 10, 57 and 148 degrees squared per hour; 21.8% of intervals are too calm to resolve and sit at the floor of the estimate. This spread is what Theorem 2 and Proposition 1 are about.

Script. `fig_power.py`, which draws the median power ratio in 1-degree bins of the vane reading, from records with no yaw inside, with the fitted $\cos^k$ and the number of records in each bin.

Shows. Where the vane can be trusted, and what a degree of misalignment costs.

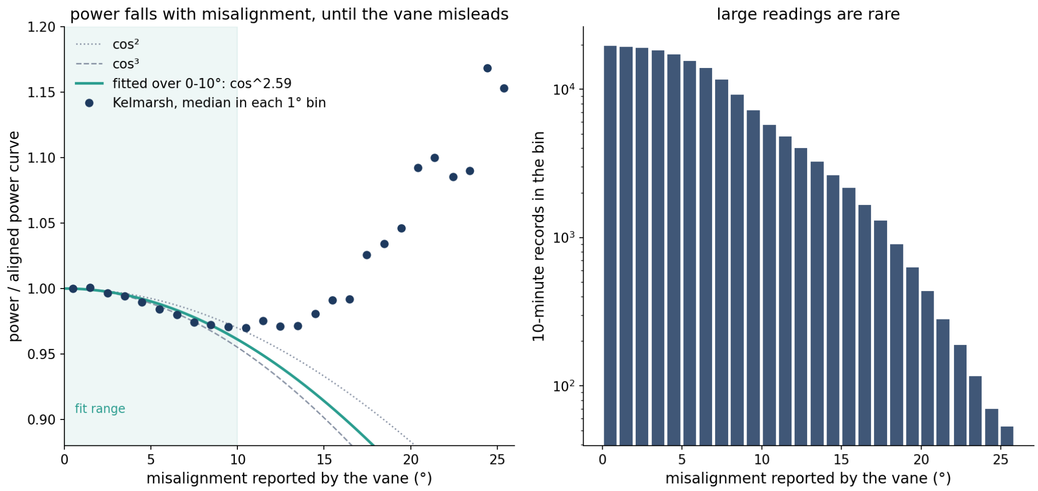

Figure 2: Left: the power of the Kelmarsh turbines relative to their aligned power curve against the misalignment the vane reports, with medians in 1-degree bins, the fitted cosine power over 0 to 10 degrees and cosine squared and cubed for comparison; right: the number of 10-minute records in each bin, on a logarithmic scale.

Up to 10 degrees the medians fall along $\cos^{2.59}$, between the cosine squared and cubed. Beyond about 11 degrees they turn upwards and exceed the aligned curve, which a misaligned rotor cannot do: in those records the vane is not reporting the true misalignment (Remark 1). They are also few, as the right panel shows, and the fit ignores them.

Script. `fig_band.py`, which evaluates Lemma 1 and Theorem 1 at the Kelmarsh values of a and s^2 and the revealed price of a yaw.

Shows. The trade-off behind the deadband, and how weakly its width depends on the price.

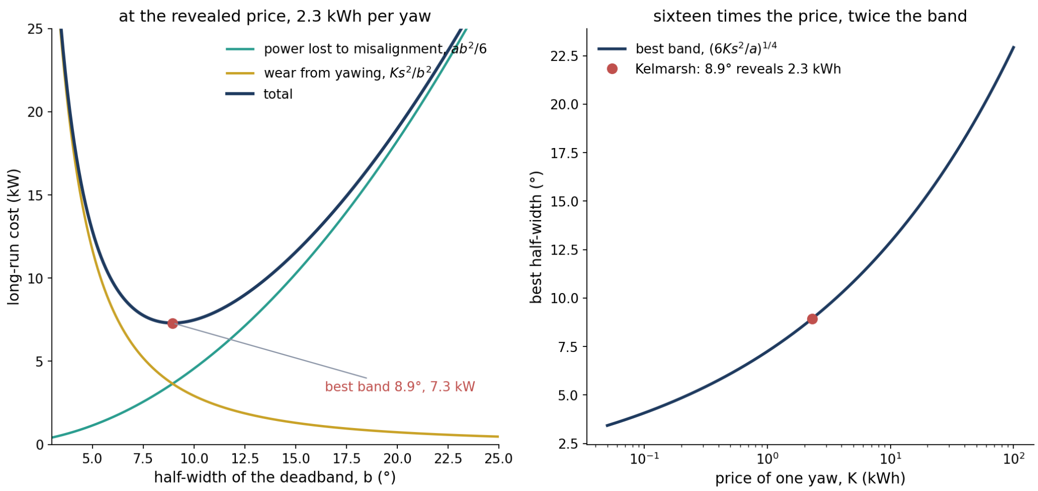

Figure 3: Left: the long-run cost of a deadband against its half-width at the revealed price of a yaw, split into the power lost to misalignment and the wear from yawing, with the optimum marked; right: the best half-width against the price of a yaw on a logarithmic axis, the fourth-root law, with the Kelmarsh band marked.

On the left the lost power rises with the square of the band and the wear falls with its inverse square; they cross at the optimum, 8.9 degrees and 7.3 kW, where each is half the total. On the right the best band grows only as the fourth root of the price: from about 3.4 degrees at 0.05 kWh to 23 degrees at 100 kWh. A price known only to within a factor of four would still fix the band to within a factor of $\sqrt{2}$.

Script. `fig_frontier.py`, which draws, from the simulations of Algorithm 2, the least mean square misalignment reachable at each yaw rate by four controllers, for turbulence correlation times of 10 and 60 seconds.

Shows. What each controller buys with a given number of yaws.

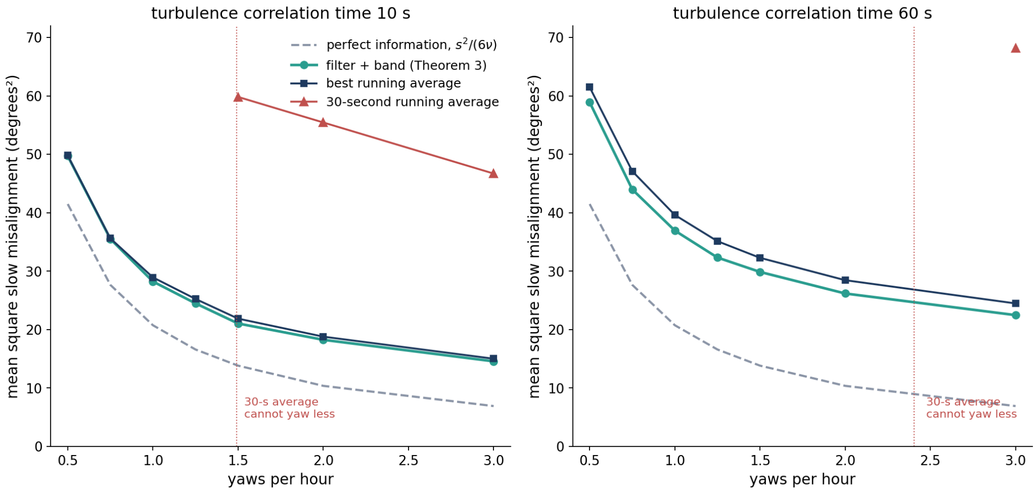

Figure 4: The least mean square slow misalignment reachable at each yaw rate by a perfectly informed band, the filtered band of Theorem 3, the best running average and a 30-second running average, from second-by-second simulations of the Kelmarsh wind, with turbulence correlation times of 10 seconds on the left and 60 seconds on the right.

Every curve falls as yaws are allowed to become more frequent. The filtered band runs parallel to the perfectly informed one, above it by the filter's irreducible error, which is larger when the turbulence is slower. The best running average almost coincides with the filter. The 30-second average cannot reach low yaw rates at all, the dotted line marking its limit, and where it can it sits far above the others.

Script. `fig_window.py`, which interpolates the simulations of Algorithm 2 to a common yaw rate and draws the misalignment against the filter's time constant, and Proposition 2's optimal time against the turbulence correlation time.

Shows. Why the right average is measured in minutes.

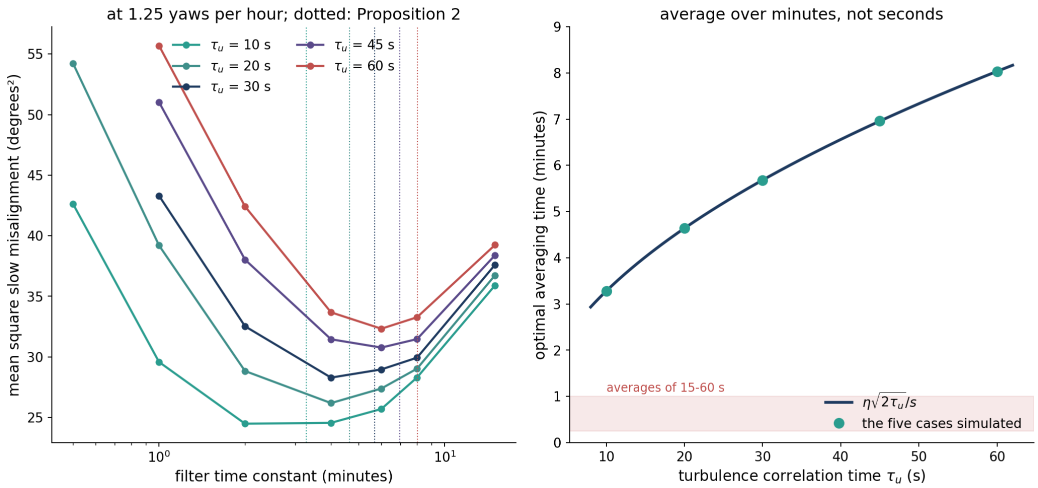

Figure 5: Left: the mean square slow misalignment at 1.25 yaws per hour against the time constant of the filter, on a logarithmic axis, for turbulence correlation times of 10 to 60 seconds, with the optimal times of Proposition 2 dotted; right: the optimal averaging time against the turbulence correlation time, with the five simulated cases and the range of averages of 15 to 60 seconds shaded.

On the left every curve has a broad minimum between 2 and 8 minutes and rises steeply towards 30 seconds, where the filter passes the gusts into its decisions. On the right the optimal time grows as the square root of the correlation time, from 3.3 minutes at 10 seconds to 8.0 at 60, well above averages of a few tens of seconds.

Script. `fig_gaps.py`, which draws the shares of the gaps between yaw intervals in the records and in the two simulations of the bursts.

Shows. That the turbines' bursts of yaws come from the wind.

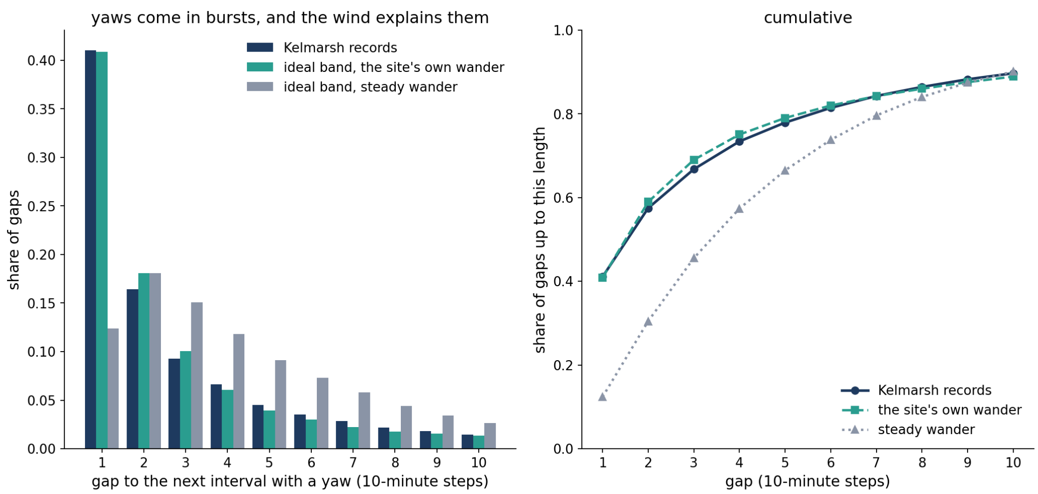

Figure 6: The gaps between 10-minute intervals that contain a yaw, as shares of all gaps: the Kelmarsh records, an ideal band run on the site's own varying wander, and an ideal band on a steady wander of the same average rate, with the cumulative shares on the right.

The records and the band on the site's own wander are almost indistinguishable: 41.1% and 40.9% of gaps are a single step. The band on a steady wander, with the same average rate, has only 12.4% and a flatter shape. The bursts are the turbines responding to spells of restless wind, not a controller chasing its own tail.

7. Discussion

Remark 6 (What an operator can change). Three settings follow from the results, each computable from the turbine's own records: the averaging time, the width of the band, and whether the width should adapt.

The averaging time should be minutes: Proposition 2 puts it at $\eta\sqrt{2\tau_u}/s$, and a filter a factor of two away from the optimum costs little, while one measured in seconds cannot be rescued by any band. The width should follow the fourth-root law with the site's wander rate and the operator's view of the price of a yaw; because of the fourth root, that price need only be known roughly. And a band that widens in restless weather and narrows in steady weather saves 18% of the cost of control at Kelmarsh, from nothing more than an estimate of the recent wander rate. Corollary 1 offers a check on any existing controller: the price of a yaw that its behaviour implies, which an operator can compare with what a yaw actually costs in maintenance.

Remark 7 (What ten-minute data cannot show). The records resolve the wind's wander and the size of its turbulence, but not the turbulence's time scale, the individual yaws, or the true misalignment.

The turbulence correlation time is therefore a parameter here, and the averaging time it implies is given as a range. Yaws are counted through flagged intervals, and the calibration corrects for it rather than measuring each one. The true misalignment is inferred from the model; the vane's own error, which lidar campaigns have measured directly on other turbines [5], is estimated only indirectly. The cost of a yaw is revealed, not measured, and the loads that misalignment itself changes [4] are left out of the cost. Finally, the analysis covers one site and runs below rated power, where misalignment costs energy; readings every second, at more sites, would pin down the turbulence and turn the ranges here into single numbers.

8. Conclusion

Yaw control is an old problem with a clean answer once the wind is measured. At Kelmarsh the wind direction is a random walk at the scales where yawing pays, at a rate that changes with the weather, with fast turbulence on top. The best controller separates into two parts: a filter that estimates the slow direction by averaging over a few minutes, $\eta\sqrt{2\tau_u}/s$, and a deadband on the filtered misalignment whose width grows only as the fourth root of the price of a yaw.

The records show that the turbines already behave like such a band, about 9 degrees wide, because the bursts in their yaws are exactly those a band produces on the site's own wander. Read backwards, their behaviour prices a yaw at 2.3 kWh, and at that price keeping the rotor aligned costs about 1% of the energy produced below rated power, half in lost output and half in wear. Letting the band widen with the fourth root of the current wander rate cuts that by 18%, and averaging over minutes rather than seconds avoids paying three times the misalignment for the same number of yaws.

References

[1] Plumley, C. (2022). *Kelmarsh wind farm data* [Dataset, open access]. Zenodo, version 0.0.3, licence CC BY 4.0. https://doi.org/10.5281/zenodo.5841834

[2] Burton, T., Jenkins, N., Sharpe, D., & Bossanyi, E. (2011). *Wind Energy Handbook*. Wiley. https://doi.org/10.1002/9781119992714

[3] Bossanyi, E. A. (2000). The design of closed loop controllers for wind turbines. *Wind Energy*. https://doi.org/10.1002/we.34

[4] Kragh, K. A., & Hansen, M. H. (2014). Load alleviation of wind turbines by yaw misalignment. *Wind Energy*. https://doi.org/10.1002/we.1612

[5] Fleming, P. A., Scholbrock, A. K., Jehu, A., Davoust, S., Osler, E., Wright, A. D., & Clifton, A. (2014). Field-test results using a nacelle-mounted lidar for improving wind turbine power capture by reducing yaw misalignment. *Journal of Physics Conference Series*. https://doi.org/10.1088/1742-6596/524/1/012002

[6] Fleming, P. A., Gebraad, P. M. O., Lee, S., van Wingerden, J.-W., Johnson, K., Churchfield, M., Michalakes, J., Spalart, P., & Moriarty, P. (2014). Evaluating techniques for redirecting turbine wakes using SOWFA. *Renewable Energy*. https://doi.org/10.1016/j.renene.2014.02.015

[7] Kaimal, J. C., Wyngaard, J. C., Izumi, Y., & Coté, O. R. (1972). Spectral characteristics of surface-layer turbulence. *Quarterly Journal of the Royal Meteorological Society*. https://doi.org/10.1002/qj.49709841707

[8] Constantinides, G. M. (1976). Stochastic cash management with fixed and proportional transaction costs. *Management Science*. https://doi.org/10.1287/mnsc.22.12.1320

[9] Harrison, J. M., Sellke, T. M., & Taylor, A. J. (1983). Impulse control of Brownian motion. *Mathematics of Operations Research*. https://doi.org/10.1287/moor.8.3.454

[10] Dixit, A. (1991). A simplified treatment of the theory of optimal regulation of Brownian motion. *Journal of Economic Dynamics and Control*. https://doi.org/10.1016/0165-1889(91)90037-2

[11] Harrison, J. M. (2013). *Brownian Models of Performance and Control*. Cambridge University Press. https://doi.org/10.1017/CBO9781139087698

[12] Øksendal, B., & Sulem, A. (2019). *Applied Stochastic Control of Jump Diffusions*. Springer. https://doi.org/10.1007/978-3-030-02781-0

[13] Kalman, R. E., & Bucy, R. S. (1961). New results in linear filtering and prediction theory. *Journal of Basic Engineering*. https://doi.org/10.1115/1.3658902

[14] Revuz, D., & Yor, M. (1999). *Continuous Martingales and Brownian Motion*. Springer. https://doi.org/10.1007/978-3-662-06400-9

[15] Liptser, R. S., & Shiryaev, A. N. (2001). *Statistics of Random Processes*. Springer. https://doi.org/10.1007/978-3-662-13043-8