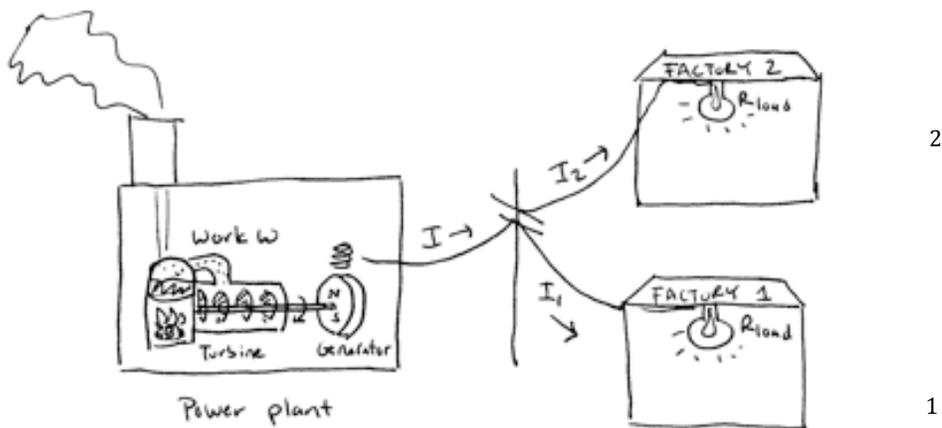


Problem 1: Stability of power generation.

In this problem you are the facility manager of a power plant that provides power to two factories only. You'll try to maintain the quality of your output power as the demand for electricity changes. You have to maintain your voltage near 120 V and your frequency at 60 Hz or lose your job.

You have a coal-fired steam turbine and a single "2-pole" synchronous generator, as in the pictures below. That means your turbo-generator is supposed to be running at 60 Hz, i.e. 3600 revolutions per minute (rpm). It's properly designed so that at this rotation rate it generates alternating voltage at 120 V for its customers. *(In real life, generators operate at voltages of order 10,000 V, and those voltages are first up-converted for transmission and then down-converted before reaching the customers.)*



Both components of your system have limits. The turbine can put out a maximum of 60 MW of work (beyond that it just can't go; no matter how fast you dump coal into your firebox – you can't get steam pressures any higher). The generator will obediently source as much current as you ask for, but the manufacturer says that it will overheat dangerously from resistive heating if you try to generate more than 50 MW of power, because the currents get too large. You haven't bothered to install any automatic shutdown mechanism, so you have to control everything by hand.

At time 0, your power plant is running comfortably. Only Factory #1 is "asking" for power, and it is demanding 20 MW, driving a resistive load. Factory #2 has not plugged in their equipment at this time, so no current is flowing. (In the drawing above, $I_2 = 0$.)

A. How much current is the generator is putting out? What is the resistance of the Factory 1's load? (in Ohms, and check your units).

A short time later, Factory #2 switches on with an identical load: $R_{\text{load}2} = R_{\text{load}1}$. (From this point on in the problem, the two factories will have the same loads).

- B. Just from common sense, how much current will now “want” to flow to the second factory? (If you can, draw the circuit and explain. Note that if MORE current is flowing when the load is larger, the net resistance must be LESS.)**
- C. Again, from common sense, what total power are you now trying to put out?**
- D. What is the effective resistance of the two-factory system?**

Due to a rapid growth in the U.S. economy, both factories simultaneously ramp up production and suddenly “ask” for more power. The extra equipment plugged in means that there are effectively more places for the current to flow, so that total resistance becomes $2/3$ of its former value.

- E. If you could hold your voltage to 120 V, what would your total power output be? What happens? (Describe in words.) (Hint: rather than calculating numbers out, try to just scale from the previous situation).**

Somehow you’ve managed to hold on to your job, the insurance covers everything, and you’ve used the insurance settlement to purchase a new generator that can handle 100 MW of power. But, the economy continues growing, and one day both factories again simultaneously step up production and drop their effective resistance now to $1/2$ their original values in D.

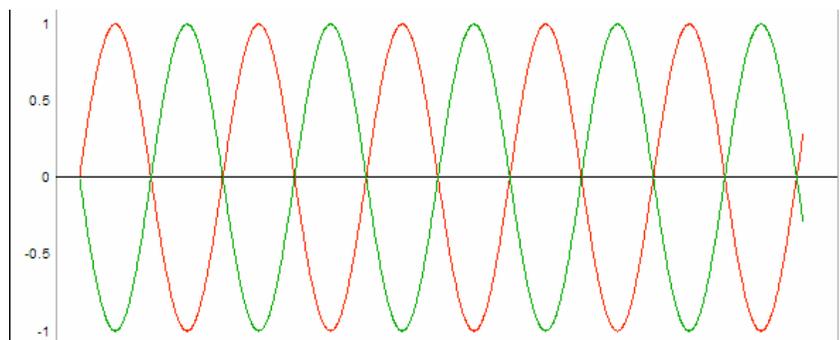
- F. If you could hold your voltage to 120 V, what would your total power output be? What actually happens? (Describe in words.)**
- G. What voltage is the generator running at?**
- H. What speed is the rotor turning?**
- I. What frequency are you putting out?**
- J. Do you keep your job now?**

Problem 2: Electricity distribution – a very local field trip.

In the questions below, answer what you can from your local electrical system. If you can't answer a few, don't worry. These are look-see questions, not time-consuming. For background, read the Hayes chapter on the grid posted on the website.

The electrical power in your house is at a relatively low 110-120 V. You know from class discussions that resistive losses would be severe if you tried to transmit power at this low voltage over any significant distance. (Losses are $I^2 R_{\text{wire}}$ or, since $P=IV$, also $(P_{\text{total}}/V)^2 * R_{\text{wire}}$.) You also know that that transmission and distribution systems minimize resistive losses by transmitting electrical power at much higher voltages (i.e. lower currents). That means there must be a transformer quite close to your house. How close? Edison transmitted DC power at 110 V and couldn't extend his lines longer than 2 km, so you are sure to find a transformer within 2 km. In modern systems it is much closer than that. In the U.S., transformers are usually located on power poles.

In the U.S., the transformer not only steps down voltage but it also duplicates and inverts each phase that is fed into it. A single conductor that goes into the transformer is matched with three wires leading out: a "hot" wire that carries the original signal, only now stepped down in voltage to 120V, a "hot" wire that carries that same 120V signal but inverted (180 degree phase shift, for those comfortable with that language), and a return wire, because the system is no longer balanced and certainly there will be net current flow.



Why the inverting? Because 120V was chosen for household wiring to be safe – the lower the voltage, the less dangerous it is if you accidentally span that voltage and shock yourself. But some household appliances actually need higher voltage – things like clothes dryers and electric stoves that pull a lot of power. (Since $P = I * V$, too low a voltage would require them to draw too high a current, and require wires too thick to be practicable to reduce the $I^2 R$ heating). If your building carries two 120V lines whose signals are inverted relative to each other, you can get a normal 120V sine wave by connecting either line to ground (0 volts) or you can get a 240V sine wave by connecting the two hot lines to each other. This system allows your house to have the best of both worlds (120 V and 240 V). Most of the wiring in your house can be at a safe low voltage, but a few high-power devices can still be wired to get the higher voltage they need.

- A. Go outside and find the transformer that serves your house /dorm / building, photograph the high(er) voltage distribution lines coming in to the transformer, the transformer itself, and the lines coming into your building. Attach the images to your problem set solutions.** (If you absolutely can't find a digital camera, make hand sketches.)
- B. Label the picture/s of the distribution lines with all interesting /relevant features.** Refer to the Hayes chapter for more information... Be sure to note whether your local distribution lines have a fourth "neutral" or "return" wire, or whether they carry the three phases without a neutral, as do transmission lines. (Both are possibilities – the high-voltage transmission lines always skip the return, but sometimes it's used for local distribution).
- C. Label the picture/s of wires entering the transformer, and again comment on all interesting features.** How many of the primary phase conductors (the 3 "hot wires" of the distribution system) feed the transformer? If only one conductor feeds it, the transformer can only produce single-phase power.
- D. How many buildings are served by your transformer?**
- E. Label the picture/s of lines leaving the transformer and entering your building and again comment on features.** From the "low side" of the transformer (the low-voltage power leaving the transformer), how many wires leave the transformer? What are those wires carrying? (See explanation above). Look for where those lines go – how many phases are sent to each building served by the transformer? Remember that since each primary phase conductor would normally become three lines (one phase and its inverse plus a neutral return wire), having 3 wires coming in doesn't mean you're getting 3-phase power. For three-phase you'd need four wires, 3 for the three phases plus a neutral.
- F. Comment on the thickness of the wires leading to your building relative to the distribution lines.** Are they thicker or thinner than the main distribution lines? Explain why. Would they carry more or less power?
- G. If you can see a max power rating written on the transformer, write it down.** This will be usually written in "kVA", which is essentially kW (as you know if you remember that 1 Volt at 1 Amp = 1 Watt). **Comment on whether that total power is reasonable for the expected number of the buildings/people the transformer is serving.** That is, estimate the max electrical power per person or house provided.
- H. (Optional) Find a substation, photograph it, and discuss.** Some will be hidden:
<https://www.wbez.org/shows/curious-city/hidden-in-plain-sight-inside-downtown-chicagos-windowless-doorless-buildings/6488570c-525b-47df-942d-bcfa59ab9ff>, Or
<http://www.messynessychic.com/2013/12/13/lights-on-but-nobody-home-behind-the-fake-buildings-that-power-chicago/>. The Hyde Park substation is hidden away on Harper Ave. between 57th and 56th. There is an exposed substation in the South Loop near IIT at 823 S. Jefferson.

Problem 3 – Energy densities of flows

In class we discussed the energy extractable from natural flows (wind, water), which is best understood with Bernoulli's equation. For incompressible fluids that is:

$$\text{Energy/mass} = \varepsilon = \frac{1}{2} v^2 + g \cdot h + p/\rho = \text{ct.}$$

(In a compressible fluid like air, the pressure term can trade off with a thermal energy term, so would be written differently. We need not worry about that for wind, though, because in free flows like wind or free-running rivers we extract the kinetic energy component only.) In dam hydro, we extract the pressure component from water at the dam base. That term is equivalent to the potential energy $g \cdot h$ at the top of the dam. The energy content $g \cdot h + p/\rho$ is constant for all water behind the dam; the relative importance of the terms just varies.

- A. Prove that the mass energy density $g \cdot h$ for water at the top of the dam is the same as the mass energy density p/ρ for water at the bottom of the dam. Consider water at the bottom of the dam: the pressure on (force/area) it is created by the mass of the column of water above. Write down the force/area exerted by a column of water, then divide by density ρ to get the $+ p/\rho$ term.

From A you should see that the height or "head" of the dam determines the energy density of its water. For comparing energy densities different fluids, engineers sometimes convert to units of "effective head": the height of a dam behind which water would have the same *volume* energy density as that in the fluid. They use volume density (energy/volume) rather than mass energy density (energy/mass) because that informs you how much power you could get from a turbine of a given physical size. In the problems below, if you're considering air, you need to account for the fact that air is 1000 x less dense than water. ($\rho_{\text{air}} \sim 1 \text{ kg/m}^3$ but $\rho_{\text{water}} \sim 1000 \text{ kg/m}^3$.) To get the effective head for a flow of air, find the energy/mass ε (which is equivalent to $g \cdot h$), then multiply by the density ratio $\rho_{\text{air}}/\rho_{\text{water}} \sim 1/1000$, then divide by g to get a height.

In the problems below, calculate both the **mass energy density** (in J/kg) and **volume energy density** (as effective head, in m) of several fluids from which we might want to extract energy.

- B. Hoover Dam, Colorado River, Water behind the Hoover Dam on the Colorado River, the dam that forms Lake Mead, the second highest in the U.S. at 726.4 feet (221.4 m). When constructed in 1935 this was the largest concrete structure in the world.
- C. Water behind Croton Dam on the Muskegon River in Michigan, still a working powerplant, at 12 m height.
- D. Water flowing in the Mississippi River, near New Orleans at its end, at a speed of 3 m/s (the fastest rate along the river – it's 1.2 m/s higher up). Discuss.



Hoover Dam, Colorado River, AZ/NV border
Photo: Wikimedia Commons



Croton Dam, Muskegon River, MI
Photo: Gray Drake Lodge website

- E. Wind at a very good wind speed for wind turbine operation, 10 m/s. Here remember that the effective head is the dam height for **water** that would have the same volume energy density. (*Remember: you don't need to worry about the pressure term for wind so the fact that air is a compressible fluid doesn't affect this problem*).
- F. What is the ratio of the head of Hoover Dam to the effective wind head in E? That ratio tells you how much more power you can get from a big hydro turbine than from a wind turbine of similar size. Discuss.
- G. (*Optional*) Consider rates of fluid consumption. Your car burns gasoline to turn chemical energy into kinetic energy to turn the drivetrain. In highway driving you consume about 2 gallons of gasoline every hour. Now imagine you powered your car via hydro instead. Imagine you have an electric car and charge your battery with energy generated from a hydro turbine. In order to drive for an hour, how many gallons of water per hour would you have to let flow through the hydroelectric plant, if you're getting power from Hoover Dam? From Croton Dam?

Problem 4: Size of wind turbines

You saw above that the energy density of wind is very small. This means that to get appreciable power, wind turbines must be very large. A wind turbine is enormous (over 100 meter diameter) compared to even a big hydro turbine, but produces less than 1% of the same power. The average wind turbine installed in the U.S. today is a bit over 2 MW rated power. (Rated power is the maximum power that can be put out with optimal wind.) You saw in class that large hydro turbines are well over 500 MW.

The size of wind turbines is however problematic – it means that they can't rotate too fast or they'll shake themselves apart. Why? Consider the speed experienced by different parts of the blade. The linear velocity will differ along turbine blade will differ: speeds are slower near the hub and maximum at the tip of the blade. In one revolution of the turbine, the tip of the blade travels around a circle of distance $2\pi R$ or πD . (If this is confusing, see the more extended discussion posted with the readings.) The blades cannot survive if the tip speed breaks the sound barrier (which is 340 m/s).

- A. Visualize a wind turbine in your mind. Estimate its rotation rate, in rotations per second (Hz). Then check your imagination by finding some video of a rotating turbine, and estimate from the video.
- B. For a turbine the size of those at the Kelly Creek wind farm (blade length 49 m, so that total rotor diameter is 98 m), what is the maximum rotation rate that would still keep the tip speed below the speed of sound? Compare to your estimates in A. Is the real turbine safe?
- C. You might have found in B that wind turbines seem even slower than is required for safety. It turns out that wind turbine designers have an additional constraint, that the turbines obtain optimal efficiency only if their tip speed ratio “tsr” – the ratio $v_{\text{tip}}/v_{\text{wind}}$ between speed of the blade tip and the speed of the wind – is in a narrow range of 5-6. If the average wind at Kelly Creek is around 7 m/s (typical for Illinois sites), what should the tip speed be for optimum efficiency? What should the turbine rotation rate be? Compare to your estimates.

Note that the tip speed ratio constraint means that the ideal rotation rate for the turbine would vary slightly with wind speed. All the big generators we discussed before - hydro and gas/steam – are synchronized to the electrical grids so their rotation matches 60 Hz power. Wind turbines used to be that way as well, with their slow rotation fixed and matched to the grid (producing 60 Hz AC via gearboxes and multiple magnetic poles). Modern power electronics now allow “variable speed wind turbines”, which get slightly higher ϵ because they can always rotate at their optimal rate: faster if the wind is blowing faster. The resulting power is then frequency-shifted to match the grid, via power electronics housed inside the tower.



Wind turbine “control cabinet” (Ingeteam)



Lowering the tower over the control cabinet (Grasmere Wind Farm)

Problem 5: Areal energy density: can wind power the world?

In previous problem sets you identified a target for finding a renewable power source that could run the whole world in the future and let everyone live like Americans. The criterion we picked in class was that the areal energy flux had to be 10 W/m^2 or better. That corresponded to 5% conversion efficiency of sunlight, assuming an average 200 W/m^2 reaching the ground. Can wind get us there? In this problem you'll consider limitations on individual turbines, design a wind farm, and calculate the areal energy flux it can extract from the wind.

- A. In class we wrote down an expression for the power carried by any moving fluid: $P_{\text{kin}} = \epsilon A \rho v$, where ϵ is the energy density (energy/mass). Plug in the energy density of moving wind to get an expression for the extractable power in a flow of air. What is the dependence on velocity? If wind at one site is 2x as fast as at another, how different are the power flows? Discuss in the context of wind farm siting. As you'll see on Tuesday, wind farm developers are extremely picky about their locations, generally refusing to build on sites below $\sim 7 \text{ m/s}$ wind. Why?

Betz's law limits: Of course, you cannot in practice extract ALL the kinetic energy from wind. Extracting energy via a wind turbine slows the wind down, but since the flow of air can't be completely stopped, not all kinetic energy can be removed from the wind. The *extractable* power is then $P_{\text{ext}} = e P_{\text{kin}}$, where e is the efficiency of extracting power. It is straightforward (see optional problem) to derive that the limiting case for the ideal turbine is $e = 0.59$. Commercial wind turbines are not ideal but they are very close, and it's reasonable to assume that they operate at $e \sim 0.5$.

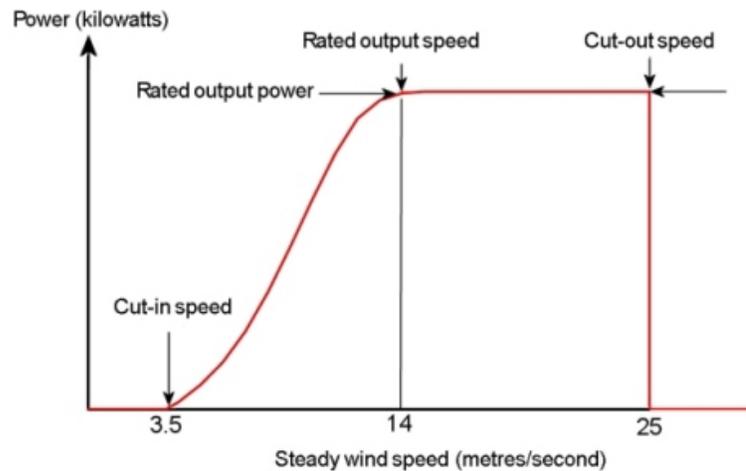
Wind variability: Even when a wind turbine operates with 0.5 extraction efficiency at all times, you face the additional constraint that wind speeds are not constant, and a given turbine will experience a range of wind velocities v .

- B. Explain why your expression in A above means that you cannot compute the mean power produced by a wind turbine by simply plugging in the mean wind speed. Instead you must consider the real-world distribution of speeds.

Wind turbine power curve: It turns out that operational constraints introduce still more complications. You can think of a wind turbine as experiencing four different regimes. In three of these regimes (#1,3,4), extractable wind power is further limited.

1. **The wind is blowing too slowly to be worth letting the turbine turn.** There are always some frictional losses, and if the turbine can't produce enough electricity to overcome those, it would act as a motor instead of a generator: you'd have to draw electricity from the grid to keep the turbines rotating. Running the turbine would consume rather than produce electricity. In these conditions you'd set a brake, stopping the turbine.

2. **The wind is lower than ideal, but high enough to be worth running.** In this case you as the turbine operator would try to produce as much power as possible. You'd set the turbine blades optimally and let it run. The resulting power would then be a function of wind speed. In the diagram below, this regime is between the "cut-in speed" and the "rated output speed".



Typical wind turbine power output with steady wind speed.

3. **The wind is optimal or a bit higher, and manageable.** One you hit the "rated" wind speed, the turbine is producing as much power as it is safe to produce. As you know from Problem 1, there's an upper limit on the power you can get from any turbogenerator, set by the basic turbine and generator designs. If you try to source too much power from a generator it can melt. And if you try to push a turbine too hard – if you apply too much torque to it – you can snap off the blades. That means that if the wind is blowing too hard, you as the operator would adjust the turbine blade angle to a "worse" position, making the turbine less efficient at extracting power from the wind (reducing e to below 0.5). The stronger the wind, the more inefficient you try to make the turbine, to keep the power output at exactly its rated maximum. In the diagram above, this regime is between the "rated speed" and the "cut-out speed".
4. **The wind is so strong that it's not manageable.** If winds are high enough, the risk of overpowering the turbine becomes too great to let it keep spinning. At this point you'd again sets the brake on the turbine and stops it. (You'd also usually point the blades into the wind to minimize force on them, just as if you were sailing in a storm, you'd point your boat directly into the wind for safety.)

To figure out the actual average power the turbine can produce, you convolve the full distribution of real-world wind speeds with the turbine power curve above. Sometimes you'll run at rated power, sometimes in regime 2 where the power output depends on the wind speed, sometimes the brake will be set. The result of that calculation gives you what's known as the **capacity factor** of the wind turbine: the average power produced divided by the rated power of the turbine. A capacity factor of 0 would mean that the

turbine never produces power: there is always not enough wind, or too much wind. A capacity factor of 1 would mean that the turbine always produces its rated power, i.e. that the wind was always ideal for that turbine design. But there is no site on Earth so consistent to produce this scenario. If you install a 2 MW turbine even on a good wind site, it will usually be producing less than 2 MW. The capacity factor therefore serves as a kind of metric of the site's quality, describing how close to its rated power a turbine will actually get. A normal situation in the U.S. would be to install a 2 MW turbine and find that in practice it averages out at 0.7 MW, so that its capacity factor is $0.7/2 = 0.35 = 35\%$. (*But, see note at end of problem for the complexities of defining capacity factors.*)

Wind farm geometry. Wind turbines are always spaced some distance apart. Because each turbine disrupts the velocity field of the wind, they can't be placed right next to each other, or right behind each other - the velocity field downstream of a wind turbine takes a long time to recover. The rule of thumb is to place turbines at least 3 rotor diameters apart along the direction facing the wind, i.e. to leave room for two other wind turbines in between any two that you build, and to leave 10 rotor diameters behind each turbine.

- C. Based on the above discussion, pick reasonable values for an average wind speed v , an efficiency e of extracting energy from the wind flow, a capacity factor c , and a turbine rotor diameter, and write them down. Then, compute the power/area carried by the average wind, where here area is the x-sectional area of the turbine, the "swept rotor area". Then multiply that value by $e*c$ to get the average power/swept rotor area extracted by a wind turbine.
- D. Draw a diagram of a wind farm layout and wind turbine spacing in terms of the diameter D of the rotors. Then derive the resulting power per area in W/m^2 , where the area is now the area of the ground where the wind farm is built, not the swept rotor area.
- E. Does your answer depend on the size of the turbines? (By "size" I mean here the rotor diameter.) That is, by using physically larger turbines, can you extract more power per area? Think carefully about this.
- F. Does windpower meet your criteria of $10 W/m^2$?

Notes

On capacity factor. The capacity factor is often explained in either of two ways, both of which are wrong. In one perspective it is a description of site quality. In another it is a measure of the "efficiency" of the turbine. Both of those views are misleading. First, all commercial wind turbines operate close to Betz' limit; they can't get more efficient. Second, the capacity factor is only defined for a particular turbine type, and what it really tells you is something about how that chosen turbine is matched to a particular site. What the capacity factor literally tells you is how close to its rated power a

particular turbine would perform, on average. The capacity factor need not inform you about wind speed at all, or even power.

To understand the complications this introduces, consider a recent trend in wind farm development. Designers have begun using smaller generators relative to the blade size. Since the big rotor area catches more wind power, these turbines reach their maximum limit at smaller wind speeds, so spend more of their operational time at rated power. By definition that means their capacity factors are closer to 1. Reported capacity factors of wind farms have therefore increased, even though the sites are actually getting worse, and the turbines aren't getting more physically efficient. Instead, it's just that turbines are scooping up more wind and applying that power to a smaller generator.

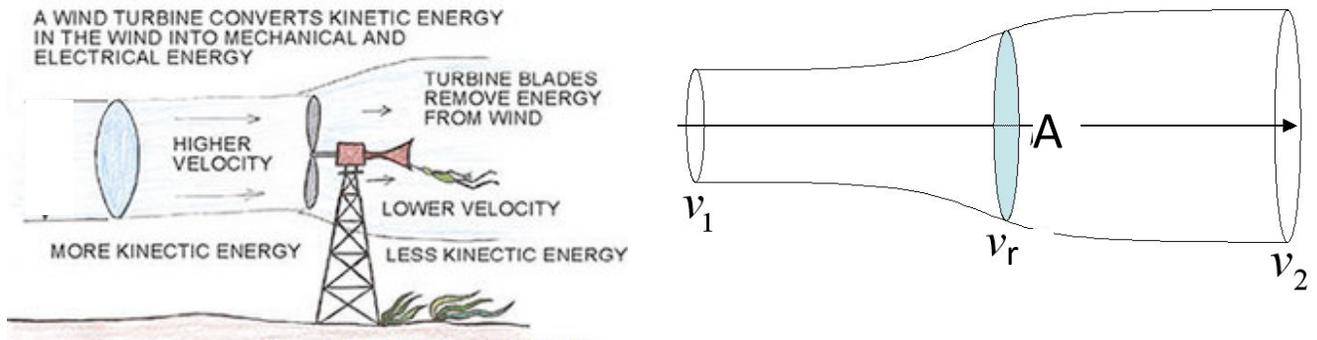
Is going to big blades and undersized generators a good thing at all? Let's think through it using the example stated in the problem, of a 2 MW turbine installed at site where that turbine produces, on average, 0.7 MW, and let's assume reasonably that power is low because wind on this site is frequently weak. The trendy developer would try to compensate for that weak wind by building a higher tower and putting longer blades on the same 2 MW turbine. The turbine then maxes out in lighter wind, producing more power at low wind speeds. But, is that really a good thing? You could actually consider that this new turbine is in a way wasteful of wind. With that big rotor area you should be able to make huge power when the wind is blowing hard, but instead you're artificially limited to 2 MW by having a puny generator. The limit also means the areal energy density of your wind farm must be lower than previously, because you'd have to space the turbines more widely, appropriately for their large rotors, but the turbines wouldn't really act like big wind turbines. Instead they'd act like maxed-out small ones. The choice is purely driven by economics: wind farms could make more power with bigger generators, but bigger generators also cost more.

On wind turbine speed. What happens to a wind turbine when the wind speed increases? Based on problem 1, you might assume that the wind turbine would have to speed up. But Problem 1 concerned a single generator driving a single load. In real life, the real electric grid is driven by many generators, all rotating synchronously. An old-style "fixed speed" wind turbine that is directly connected to the grid is forced by that connection to rotate synchronously with the 60 Hz grid, regardless of the wind speed. If the wind increases, the turbine makes more power, but rather than speed up it effectively pushes a bigger load. This is good in some ways – the turbine has to produce power that is of appropriate frequency – but if turbine speed is constant regardless of the wind then the tip speed ratio (rotor speed/wind speed) is not, meaning the wind turbine efficiency is not optimal at all times.

New variable-speed wind turbines circumvent this problem by using power electronics to connect to the grid. They are therefore permitted to increase their rotation rate as wind speed increases, until they reach their rated power.

Problem 6 (Optional): The limiting efficiency for extracting energy from the wind

In this problem you'll derive Betz' Law, the limit for the fractional power extracted by an ideal free-stream wind or hydro turbine. (That is, the ratio of power extracted to power carried by the fluid as kinetic energy.) The important thing to realize is that the wind turbine extracts energy by slowing the wind down:



In the pictures above, the speed v_2 is less than v_1 , and the speed at the rotor v_r is somewhere in between. You want to figure out the optimal loss of speed. (This is stated in the slides, but you'll derive it here.) First, you figure out how much power is lost during the slowdown. Then you'll compare that to the power in the wind that would flow past the place where the wind turbine is, if you *hadn't* installed a wind turbine. You take the derivative to determine the slowdown that produces optimal power extraction, and plug that optimum back in to derive the fractional power extracted.

The figure on the right shows the "streamlines" of passing through the turbine. Because mass can't be created or destroyed (the wind turbine is not a nuclear reactor), the mass flow (mass/time) in the diagrams above must be constant at all points shown (point 1, point 2, and at the rotor). When the wind slows down, the area that it occupies therefore necessarily expands. As an analogy: think of trying to move the same number of cars on a road where the speed limit gradually reduces. To keep cars from piling up and making a traffic jam, you'd have to widen the road as you slow down the speed limit, so that cars can still move through at the same rate. That rule means that the flow from point 1 to point 2 occupies a gradually increasing cross-sectional area.

- A. First write an expression for the energy/mass extracted as the air slows down from v_1 to v_2 . The average air molecule passing through the turbine had a certain energy/mass at point 1; it has less at point 2; the difference is the power that the turbine extracts. Write down an expression for that difference.
- B. Write down the mass flow rate (mass/time) of this flow. Calculating the mass flow rate is a little tricky, since you need to know an area and a speed at the same place. But in the diagram above, at the wind turbine rotor you know the area but you don't know the speed, and at the places where you defined your

speeds of v_1 and v_2 , you don't know the area. So, I'll give you a piece of information that is actually not obvious and requires proof: the speed at the rotor is exactly the average of the upstream and downstream velocities v_1 and v_2 . Now you know a speed and an area at the same point, so you can define the mass flow in the moving "sausage" of air shown in the diagrams.

- C. Multiply A by B to get the total power that could be extracted from the flow (energy/mass * mass/time = energy /time). Then simplify your expression by writing it in terms of the ratio v_2/v_1 . That's the logical variable to use, since we don't really care about the absolute wind speeds and instead are thinking about by what fraction we can slow down the wind.
- D. Now ask the question: if you had not installed a wind turbine at all, what would be the power that flowed through the area swept out by the wind turbine? Write down this expression. Then divide your expression in C by this quantity to get an expression for the fraction of this power that you can capture with your wind turbine. That is, you're writing $e = P_{\text{extracted}}/P_{\text{no turbine}}$ as a function of the "slowdown ratio" v_2/v_1 .
- E. Determine the optimal "slowdown ratio" v_2/v_1 that gives you the largest e . If you don't have calculus, or just want to do this by hand: make a graph of the function e vs. v_2/v_1 plotting values until you understand the shape of the expression and can see where its maximum is. Otherwise, take the derivative of e ; the maximum is where the derivative is zero. Your value must be somewhere between 0 (air comes to a complete halt) and 1 (no slowdown at all).
- F. Plug your optimal slowdown ratio in to the expression for e to solve for the optimum efficiency. You should get a simple fraction, whose value will be between 0 and 1. This is an actual physical limit, that applies to any wind turbine, any design, any wind speed, anywhere: the maximum fraction of wind power you can extract.
- G. (*Optional*): try to prove that the velocity at the rotor is the mean of the upstream and downstream velocities.