Physics of the Free-Electron Laser

The free-electron laser can produce intense, coherent light as the result of the interaction between an electron beam and an electric field of light.

For further information contact Ted Scharlemann (415) 422-5795.

A free-electron laser (FEL) directly converts the energy of a beam of electrons into light. The light is emitted as the electron beam bunches into many very thin "pancakes" that radiate coherently, as if the bunches formed an antenna. The bunching occurs when the electrons wiggle in a periodic array of magnets (the wiggler) and interact with the light. As the electrons lose energy, they can continue to interact with the light if the wiggler magnetic-field strength or period decreases with distance along the wiggler. Dramatic increases in the efficiency of conversion between electron-beam power and light are possible when this is done.

The term "free-electron laser" is something of a misnomer because the light in an FEL is created by a mechanism better described as coherent synchrotron radiation than as stimulated emission, the mechanism in a conventional laser. Synchrotron radiation is inevitably produced when a relativistic electron (that is, an electron with velocity very close to the speed of light) is accelerated or decelerated. The electrons in the relativistic electron beam of an FEL are accelerated—in the sense of changing direction, not gaining energy—by the wiggler (see Fig. 1a, described in more detail later).

If nothing else were to occur in an FEL, the device would be a standard synchrotron radiation source, such as the one at the Stanford Synchrotron Radiation Laboratory. There, the radiation power from the electron beam is the sum of the single-particle radiation power from all the electrons in the beam; that is, the synchrotron radiation from individual electrons in the beam adds incoherently.

In an FEL, the electrons physically bunch on the scale of one wavelength of light to form a series of very thin pancakes. The pancakes are 21 μm thick and 1.6 mm in diameter for the examples used in this article. The single-particle radiation from each electron in a pancake has the same phase; the radiation amplitudes, rather than powers, add together. Furthermore, because the pancakes are separated by almost precisely one wavelength of light, the radiation amplitudes from different pancakes also add in phase. The radiation power, which is proportional to the square of the radiation amplitude, then depends on the square of the number of electrons radiating in phase rather than on only the number of electrons. At the single frequency designated for an FEL, the coherence of the radiation (the linear adding of amplitudes) produces an enormous enhancement of radiation power.

The frequency at which an FEL operates is determined by the wiggler period, the electron energy, and the wiggler magnetic field. The optical wavelength is a Lorentz transformation[1] of the wiggler period into the beam frame (a frame of reference moving down the wiggler with the average velocity of the electron beam) followed by a Doppler shift back into the laboratory frame of reference, again with the average velocity of the electron beam down the wiggler. The velocity we used in the Lorentz transformation and the Doppler shift is the longitudinal electron velocity. This velocity is less than the full electron velocity because of the curved path that the electrons take in the wiggler (see Fig. 1b). The difference between longitudinal and full velocity of the electrons depends on the wiggler magnetic field; hence, the signal (optical) wavelength λ_s is the wiggler wavelength λ_w multiplied by $(1 + a_w^2)/2\gamma^2$, where a_w is the dimensionless root-mean-square vector potential of the wiggler, and γ is the Lorentz factor of the electrons.

The light from the FEL could be at any wavelength from microwave to ultraviolet, and eventually, with improvements in available electron beam quality, even at x-ray wavelengths. We will refer to the electromagnetic radiation in the FEL as "light," and ask the reader to remember that the "light" could be at almost any wavelength.

The key to operation of an FEL is the physical mechanism that bunches the electron beam into pancakes and thereby permits the radiation to be coherent. We can best understand the mechanism by examining what the wiggler does to the motion of electrons in a relativistic electron beam and how that motion can couple electrons to the amplified light.

Wiggler Effects on Electrons

Figure 1a is a side view of a linear wiggler in which the periodic magnetic field is in the vertical direction. In this example, the magnets are arranged so that the direction of the magnetic field reverses every 2.75 cm. The magnetic field is shown by the dashed lines from each magnet. Precisely at the midplane of each magnet, the field lines are either up or down, with no sideways component. The electron beam moves horizontally, from left to right, across the magnetic field. Because of the motion across the magnetic field, each electron is affected by a force into or out of the plane of Fig. 1a and at right angles to both the magnetic field and the direction of motion. The force causes the electrons and hence the electron beam to wiggle.

Figure 1b is a top view of the wiggler in which the electron beam, as before, moves from left to right. However, the magnetic field is now directed into or out of the plane of the figure. The undulating line shows the wiggle motion of the electrons.

The light amplified in the FEL propagates in the same direction as the beam, which in Fig. 1 would be from left to right. The light is polarized such that its electric field is in the horizontal plane, the same

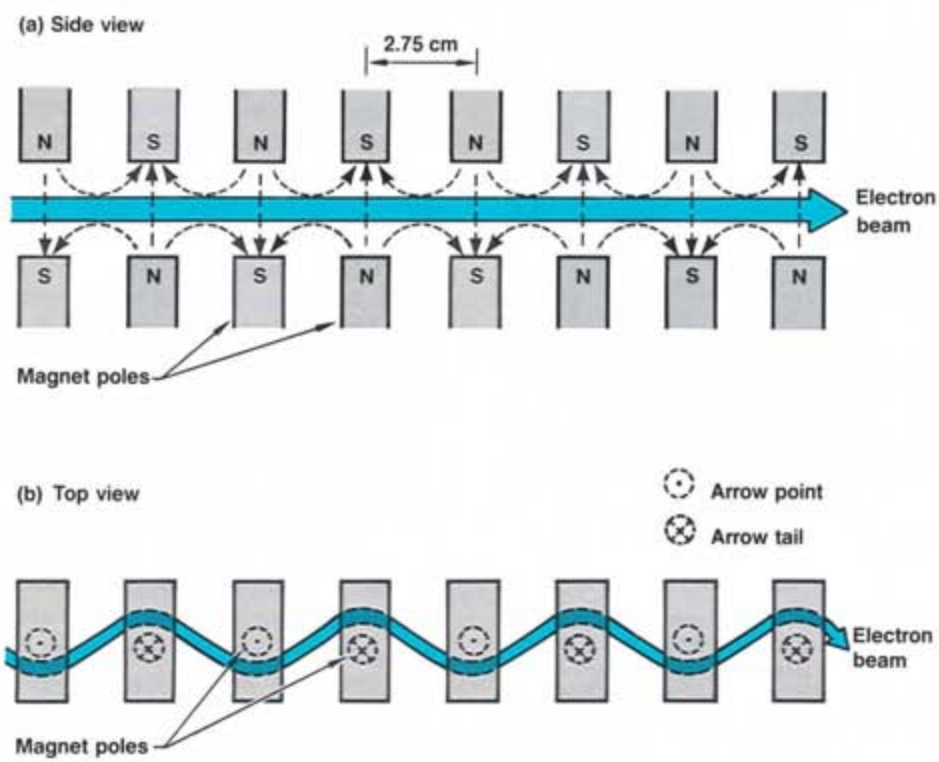

Fig. 1

Diagram of a wiggler in which magnets are arranged linearly every 2.75 cm. (a) Side view of a linear wiggler in which magnets are arranged to provide a magnetic field that periodically reverses direction. On the midplane, the magnetic field lines point either up or down. As the electron beam moves across the magnetic field, each electron wiggles in response to a force at right angles to both the magnetic field and the direction of motion. (b) Top view of the wiggler, with the magnetic field directed into or out of the illustration, as denoted by the tails and tips of arrows, respectively. Only the top set of magnets is shown. The horizontal wiggle motion of the electrons is represented by the undulating blue line.

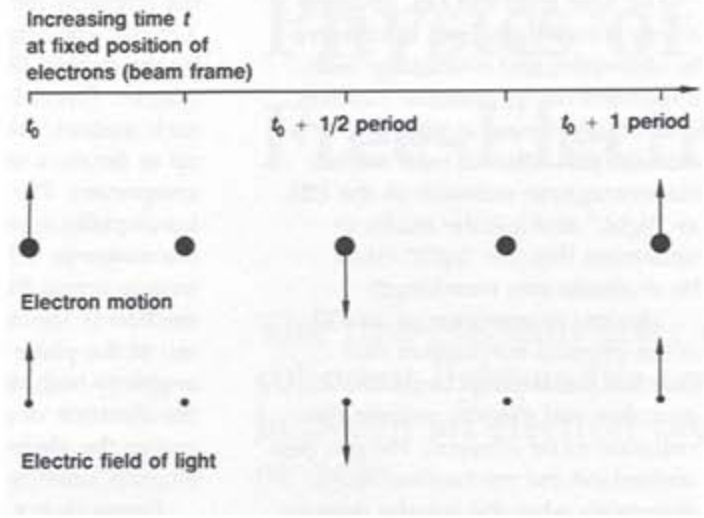

Fig. 2

In an FEL, electrons oscillate in time t at a given position in the beam frame. At each electron position, the electric field of the light also oscillates periodically. The periodic motion of the electron illustrated here remains in the same direction as the electric field of the light. The electric force on the electron is therefore steadily opposite to its direction of motion. The electron consequently loses energy that is gained by the light.

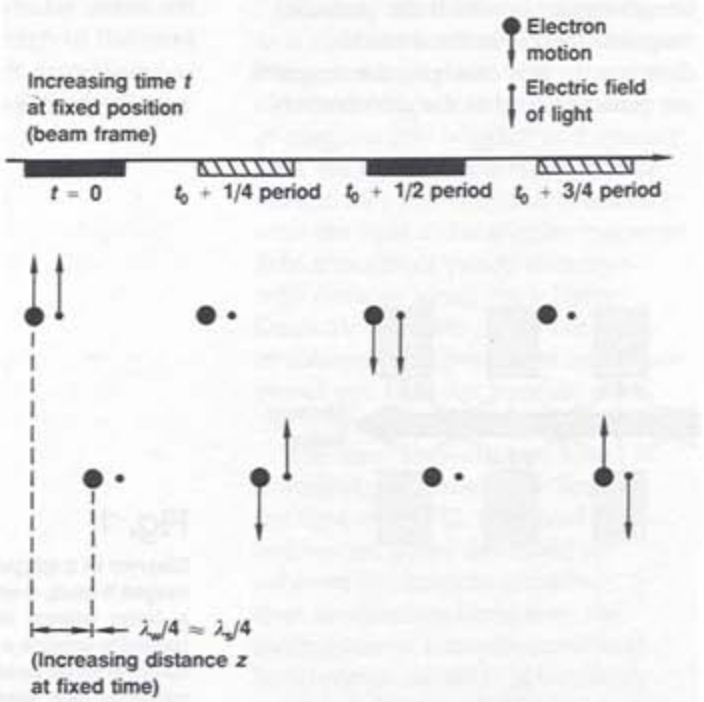

Fig. 3

An electron shifted to the right by one-quarter wavelength of light with respect to the electron of Fig. 2 sees a force from the light that is steadily in the same direction as its motion. That electron, therefore, gains energy from the light. As electrons either gain or lose energy, they move toward the point at which the wiggle motion and electric field of the light are 90° out of phase.

plane as the wiggle motion of the electrons. The interaction of the electric field of the light with the electron wiggle motion couples the electron beam to the light.

Imagine that we could move with the average speed of an electron in Fig. 1b. Then we would see each electron moving back and forth horizontally, as Fig. 2 shows. We would also see the light propagating to the right. At the position of each electron, the electric field of the light would also oscillate, pointing back and forth periodically. In an FEL, we can arrange for the periodic electron motion always to be in the same direction as the periodic direction of the electric field of light, at least for a single electron. With such a configuration, the force on the electron by the electric field of the light is steadily opposite to the direction of motion of the electron. A force consistently opposed to the motion of an electron causes it to lose energy. The energy lost by the electron is gained by the light.

In addition to the electron with a properly phased electric field and wiggle motion just described, there are other electrons in the electron beam. For example, an electron shifted to the right by one-quarter wavelength of light is affected by a force from the light that is always in the same direction as its motion (Fig. 3). This electron will gain energy and, hence, decrease the energy of the light.

Electron Bunching

When the electron beam first enters the wiggler, it is unbunched, and an equal number of electrons gain and lose energy. Therefore, no net transfer of energy initially takes place between the beam and light. But consider what happens as individual electrons either gain or lose energy. As observed in the laboratory, an electron that gains energy speeds up; one that loses energy slows down. As these two electrons move with the beam, they approach one another at a point midway between their original

positions. The point toward which they move is where the wiggle motion and the force from the light are exactly 90° out of phase. Here, an electron neither loses nor gains energy, and it is toward this point that all electrons in the beam move. The result is a bunching of the electron beam so that the bunches are the size of a wavelength of light and wiggle at a phase (with respect to the electric field) for which no exchange of energy with light takes place.

So far, there has been no net transfer of energy between the electron beam and light. However, the bunched electrons are accelerating (wiggling); therefore, they must necessarily radiate coherent synchrotron radiation. The apparent inconsistency of added radiation without gain is resolved by recognizing that the added radiation is 90° out of phase with the radiation that is present; the effect of the added radiation is not gain, but a phase shift of the light. Instead of changing in amplitude, the electric field at the position of the electron bunches changes in phase.

Now, at last, gain can occur. The shift in phase of the electric field of light is precisely equivalent to a shift of position of the bunch. The sign of the phase shift turns out to be the correct sign to shift the electron bunch to a position at which it loses energy. The coherent loss of energy of the electrons in the bunches becomes a gain in energy of the light, and thus the electrons convert their energy into coherent light energy.

Bunching of the electron beam is essential to the operation of an FEL. Also, the bunches must be in precisely the correct position within a single wavelength of the light. It is counterintuitive, but nonetheless true, that the bunch formation and proper positioning all occur naturally, without external intervention; e.g., it is not necessary to worry about the precise phase of incoming light because the bunches will form at the correct phase regardless of the phase of the input signal.

The electrons in the bunch slow down as they lose energy. Thus, they move away from the phase at which they transfer energy to light and toward the phase at which they regain energy from light. If nothing else were done in an FEL, there would be a small maximum to the fraction of energy the electrons could transfer to light before they would begin to absorb energy. The maximum would be on the order of 10% or less, depending on the device. However, it is possible to prevent the electron bunches from reabsorbing light energy by tapering the wiggler.

Tapering the Wiggler

Tapering involves changing either the magnetic field strength or wavelength of the magnetic field periodicity with position along the wiggler. Figure 4 shows an increasing separation of the magnets as the electron beam moves to the right to represent a weakening magnetic field. The effect on electron motion can be understood by referring to Fig. 1b. To move a given distance to the right, an electron must follow a curved path, which takes longer than movement along a straight line. The result is an average velocity that is less than the velocity of the electron. The decrease in average velocity depends on the magnetic field strength and the electron energy. The larger the field or the smaller the energy, the slower an electron moves to the right.

The velocity of an electron to the right determines how the phase of the electron's wiggle motion changes with respect to the electric field of light and whether it loses or gains energy. As electrons in the bunch lose energy, they slow down by decreasing their average speed to the right unless the magnetic field is decreased to compensate for the loss of energy. It is possible to decrease the magnetic field of the wiggler in precisely the manner required to keep bunches moving at a constant velocity. The electron bunches then see an electric field that is steadily maintained in the direction of their wiggle motion even as they lose energy.

A dramatic increase in the fraction of electron energy that can be converted to light has been experimentally demonstrated at our Electron Laser Facility (ELF), where 7% conversion without tapering was increased to approximately 34% with wiggler tapering.

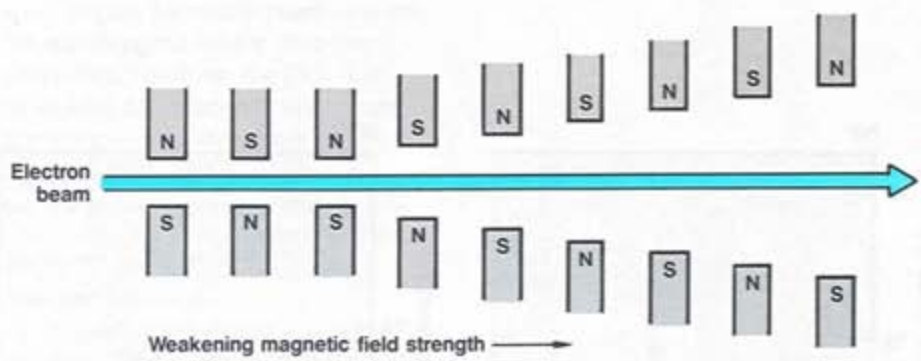

Fig. 4

In a tapered wiggler, the magnetic field strength or wiggler period decreases with distance along the wiggler. The wiggler represented here has a constant period but decreasing field strength, as indicated by the increasing distance between magnets.

Simulated Electron Bunching

Figure 5 shows various stages of gain and tapering (these images were generated by the Laboratory's FEL-simulation code, FRED). The positions of the electron energy γ (in units of the electron rest mass) and relative phase ψ are plotted for three groups of electrons at different positions along the z axis in the wiggler. The phase ψ is defined so that electrons with $\psi > 0$ lose energy and those with $\psi < 0$ gain energy; the position of initial bunching is therefore at $\psi = 0$. An electron at constant ψ is actually moving down the wiggler at nearly the speed of light. Phase ψ is a measure of the electron position in the frame moving with the average velocity of the electron beam.

Figure 5a shows the initial electron distribution on entry to the wiggler. The beam is unbunched (uniform in ψ) and monoenergetic. Figure 5b shows initial bunching, which appears as a concentration of electrons around $\psi = 0$. Figure 5c shows the fully bunched electron beam. However, the phase shift of the light has moved the bunch to $\psi > 0$, where the electron beam loses energy (on average) to light. In Fig. 5d the bunch has lost as much energy as it can in an untapered wiggler.

To extract more energy from the electron beam, the wiggler must be tapered. Figure 5e shows the effect of tapering the magnetic field so that the bunch continues to lose energy. Not all the electrons are in the decelerating bunch. The electrons at the top of the figure are nearly uniformly distributed

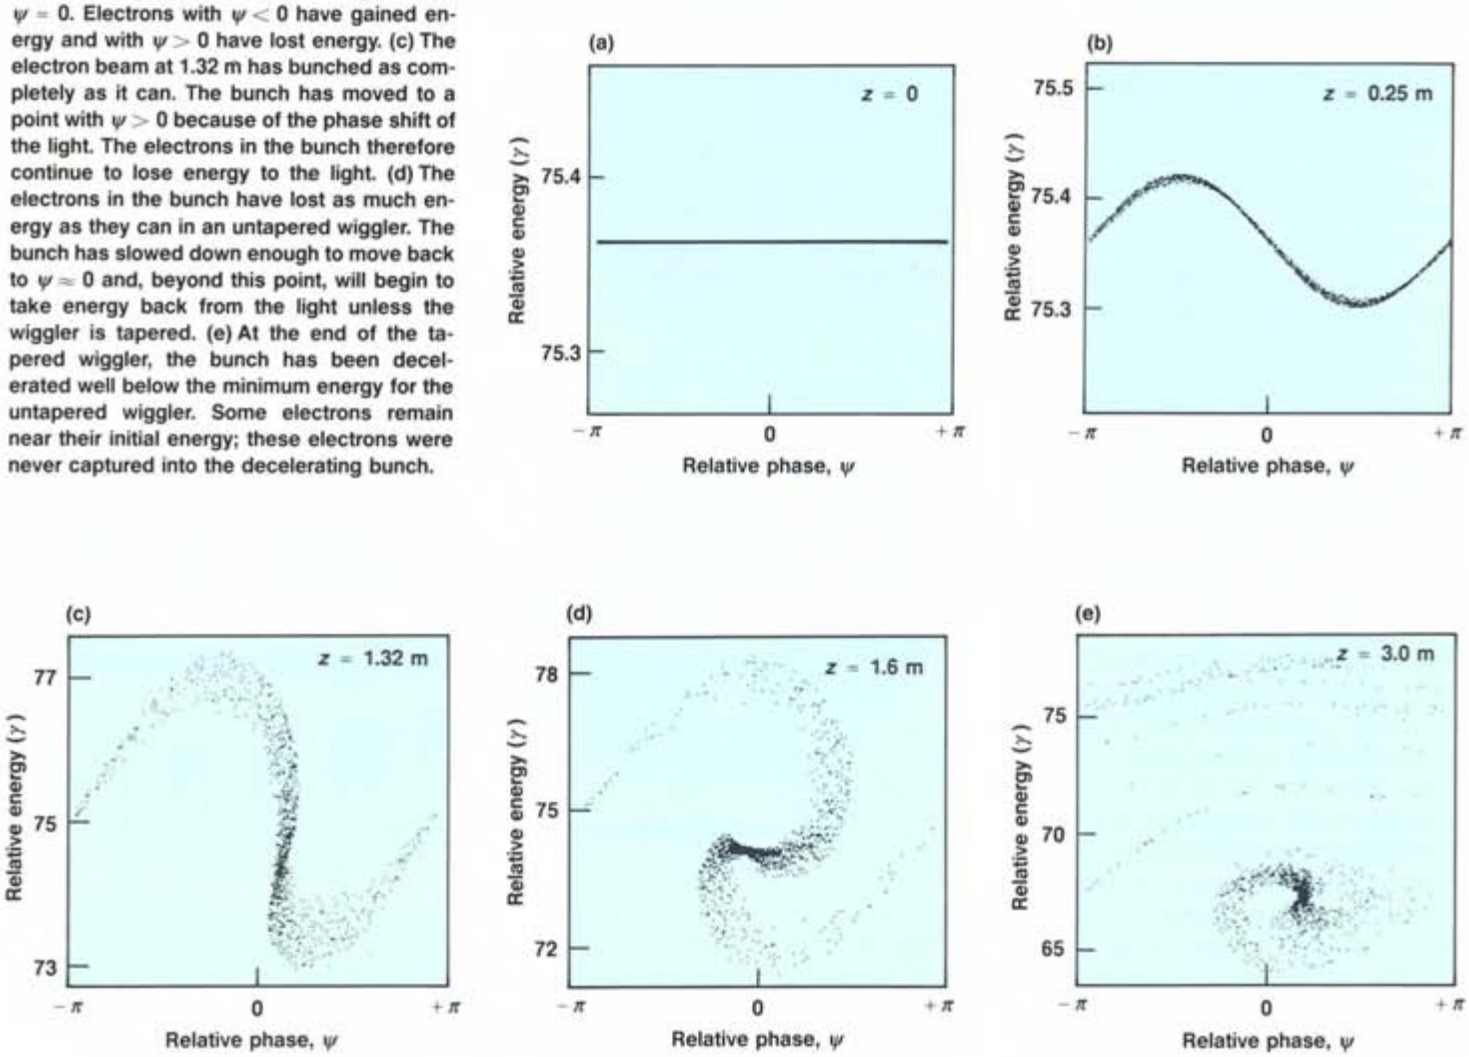

Fig. 5

The effect on the electron beam of the bunching force and wiggler tapering can be illustrated with plots of electron energy γ and relative phase ψ between wiggle motion and electric field direction. (a) The initial electron beam distribution, at entry to the wiggler. The beam is unbunched and all electrons have the same energy. (b) At 25 cm into the wiggler, the electron beam has begun to bunch around the positions at which the wiggle motion and electric field direction are 90° out of phase, or $\psi = 0$. Electrons with $\psi < 0$ have gained energy and with $\psi > 0$ have lost energy. (c) The electron beam at 1.32 m has bunched as completely as it can. The bunch has moved to a point with $\psi > 0$ because of the phase shift of the light. The electrons in the bunch therefore continue to lose energy to the light. (d) The electrons in the bunch have lost as much energy as they can in an untapered wiggler. The bunch has slowed down enough to move back to $\psi \approx 0$ and, beyond this point, will begin to take energy back from the light unless the wiggler is tapered. (e) At the end of the tapered wiggler, the bunch has been decelerated well below the minimum energy for the untapered wiggler. Some electrons remain near their initial energy; these electrons were never captured into the decelerating bunch.

in ψ and, for various reasons, never move into the confined and decelerating bunch.

The bunches shown in Fig. 5a through Fig. 5e are restricted to the range $-\pi < \psi < \pi$. This range approximately corresponds to an optical wavelength of 21 μm. The electron beam actually can be many meters long and is initially continuous and unbunched over its entire length. While the effect of the FEL interaction is to break the electron beam into many microscopic bunches (approximately 10^6), only one of these bunches is shown in Fig. 5.

The power in the light as a function of distance down the wiggler is shown in Fig. 6a. The first knee in the curve at about 1.6 m corresponds to Fig. 5d and is the maximum power obtainable from an untapered wiggler. Note the significant enhancement of power that can be obtained through tapering. Figure 6b shows the variation of the magnetic field along the wiggler. The period of the alternating field is 5.5 cm in this case, but only the amplitude of the peaks of the field is shown.

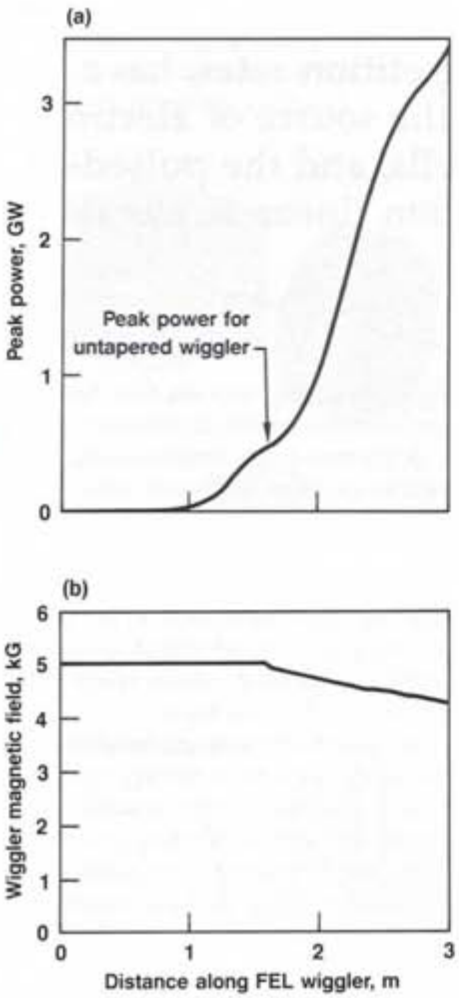

Fig. 6

(a) Tapering the wiggler greatly enhances the maximum power available from an FEL amplifier. (b) Magnetic field strength at the peaks of the periodically oscillating field in a FEL wiggler. In the examples of Figs. 5 and 6a, the wiggler is tapered by decreasing the magnetic field strength, as shown here.

Summary

The wavelengths at which FELs can operate are limited only by the energy and quality of the electron beam. The energy must be high enough that the Lorentz transformation plus Doppler shift produces the desired wavelength, and the intrinsic spread in average velocities of the electrons must be small enough (i.e., the beam quality good enough) to permit the beam to bunch. It is possible, in principle, to operate an FEL at almost arbitrarily short wavelengths, although current technology limits will prevent operation below ultraviolet wavelengths for many years to come. For wavelengths longer than the ultraviolet, however, the FEL is a promising candidate for very many applications requiring high-power coherent light.

Key Words: computer code—FRED; Electron Laser Facility (ELF); free-electron laser (FEL); synchrotron radiation; wiggler.

Notes and References

1. A Lorentz transformation is a mathematical transformation used in the special theory of relativity to relate certain space and time variables.