CHAPTER ONE

SPECIAL RELATIVITY

To understand black holes, we have to learn some relativity. The theory of relativity is split into two parts: special and general. Albert Einstein came up with the special theory of relativity in 1905. It deals with objects moving relative to one another, and with the way an observer's experience of space and time depends on how she is moving. The central ideas of special relativity can be formulated in geometrical terms using a beautiful concept called Minkowski spacetime.

General relativity subsumes special relativity and also includes gravity. General relativity is the theory we need in order to really understand black holes. Einstein developed general relativity over a period of years, culminating in a paper in late 1915 in which he presented the so-called Einstein field equations. These equations describe how gravity distorts Minkowski spacetime into a curved spacetime geometry, for example the Schwarzschild black hole geometry that we will describe in Chapter 3. Special relativity is simpler and easier than general relativity because gravity is neglected—that is, gravity is ignored, or presumed to be too weak an effect to be significant.

Special relativity includes the formula $E = mc^2$, relating energy E, mass m, and the speed of light c. This is one of the most famous equations in all of physics, possibly in all of human understanding. $E = mc^2$ made it possible to foresee the awesome power of nuclear weapons, and it is at the core of our hopes, as yet unrealized, for a clean source of energy from nuclear fusion. $E = mc^2$ is also very relevant to black hole physics. For example, the 3 solar masses' worth of energy ejected from the first observed black hole collision is a prime illustration of the equivalence of mass and energy. To get an idea of just how cataclysmic this collision was, consider that the mass converted into energy in the explosion of a nuclear weapon (assuming a yield of 400 kilotons) is a mere 19 grams.

Special relativity is closely related to James Clerk Maxwell's theory of electromagnetism. Indeed, an early hint of the relativistic view of space and time emerged in the late 1800s in the form of so-called Lorentz transformations, which explain how observers' perceptions of electromagnetic phenomena depend on how the observers are moving. The most familiar electromagnetic phenomenon is light, which is a traveling wave of electric and magnetic fields. A consequence of Maxwell's theory is that light has a definite speed. Relativity is built around the idea that this speed is truly a constant, independent of the motion of the observer.

The motion of observers is described in special relativity in terms of frames of reference. To get an idea of what a frame of reference is, think of a high-speed train. If all the passengers are seated and all the luggage is stowed, then everything on the train indeed is stationary with respect to the train itself. But the train is moving quickly relative to the Earth. Let's assume that the train is moving in a straight line at constant speed. To give a fully precise account of frames of reference, we should stipulate the absence of any significant gravitational field. For example, instead of a train running at constant speed along the Earth's surface, we would need to consider a spaceship coasting at constant speed in otherwise empty space. Earth's gravitational field is weak enough that, for present purposes, we can ignore its effects on the train and work with just the special theory of relativity rather than the general theory.

Without looking out the windows, it's hard to tell how fast the train is moving. In a situation where the train has fantastic suspension and the track is very even, and where the blinds on all the windows are down, it would be impossible to know that the train is moving at all. The train provides a frame of reference—the one that its passengers naturally use to judge whether something inside the train is moving. They can't tell (in the ideal situation just described) whether the entire train is moving. But they certainly know when someone walks up the aisle, because such a person is moving relative to their frame of reference. Furthermore, all physical phenomena, like balls dropping or tops spinning, would behave the same, as observed by an observer on the train, whether the train is actually moving or not. Briefly then, a frame of reference is a way of looking at space and time which is associated with an observer, or a group of observers, in a state of uniform motion. Uniform motion means that the train is not speeding up, or slowing down, or turning. If the train is doing one of these things, then

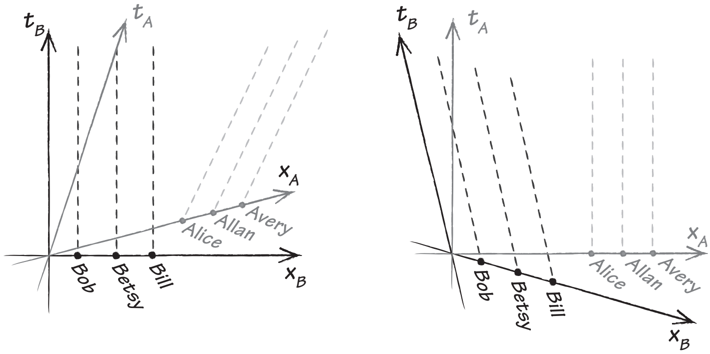

FIGURE 1.1. Left: Minkowski spacetime, showing three observers in the B-frame as stationary and three observers in the A-frame as moving forward. Right: A different perspective on Minkowski spacetime, in which the B-frame observers are now moving backward and the A-frame observers are stationary.

the passengers will notice; for example, rapid acceleration pushes them back in their seats, whereas rapid deceleration throws them forward.

Let's imagine our train passing through a station without stopping or slowing down. The passengers on the train, call them Alice, Allan, and Avery, are observers in a moving frame of reference which we'll call the A-frame. Meanwhile, their friends Bob, Betsy, and Bill stand on the platform, in a stationary frame of reference which we'll call the B-frame. To draw these frames of reference, we put B-frame position on the horizontal axis and B-frame time on the vertical axis, and we map out the trajectories of our various observers through space and time, so that over time, B-frame observers always stay at the same B-frame positions, whereas A-frame observers move forward. The resulting diagram is actually Minkowski spacetime! The word spacetime refers to the fact that we are showing space and time on the same diagram. It's possible to take a different perspective on Minkowski spacetime, such that A-frame observers are shown as stationary while B-frame observers move backward. More on that perspective later.

Special relativity hinges on the assumption that the speed of light is constant. In other words, the speed of light is supposed to be the same when measured by the observers on the train as when measured by the observers on the platform. If that weren't so, then by measuring the speed of light, an observer could tell which of the two frames of reference she was in. But a core tenet in relativity theory is that physics should be the same in any frame of reference, so that you really *can't* tell which frame you're in through any physical measurement. According to this tenet, we cannot pick out a frame and say, "Remaining in this frame is what it means to be stationary. Motion consists of being in a different frame." We can only say, "Any frame is as good as any other. The only idea of motion that we can permit is motion of one observer with respect to another." In other words, states of motion are not absolute; they are relative. Thus it was a misnomer to refer to the A-frame as moving and the B-frame as stationary. All we can really say is that they are moving with respect to one another. (The idea of the B-frame being stationary seemed natural, though, because we were implicitly thinking of motion relative to the Earth.)

The intuition we've explained about relative motion seems like common sense, and we should ask ourselves how we can possibly get any leverage from it on questions relating to the deep nature of space and time. The key ingredient is Maxwell's theory of electromagnetism. What this theory tells us (among other things) is that if Alice pulls out a laser pointer and sends a pulse of light forward, toward the front of the train, and Bob does the same, then the two light pulses travel forward at exactly the same velocity. This seems like another innocuous claim, but it's not! For example, if we arrange for the train to go at 99% of the speed of light (so obviously not an American train), then wouldn't Bob measure a laser pulse shot forward by Alice to be traveling at almost double the speed of light? After all, she is moving forward at 99% of the speed of light relative to Bob, and her light pulse moves forward at the speed of light relative to her, so it seems like Bob should measure her light pulse to be moving forward at 199% of the speed of light. But according to electromagnetism, he doesn't! He measures it to be moving at precisely the same speed of light, relative to him, that Alice would report if she measured its motion relative to her.

How is this possible? The answer is that Alice and Bob measure the passage of time differently, and they also measure length differently. The details of how this happens are encoded in the Lorentz transformation, which is a mathematical expression relating time and length in the A-frame to time and length in the B-frame. A Lorentz transformation is easy to draw using Minkowski spacetime. Before the Lorentz transformation (the left side of Figure 1.1), we can think of the B-frame as stationary and the A-frame as moving forward. After the Lorentz transformation (the right side of Figure 1.1), the A-frame is stationary and the B-frame is moving backward! A Lorentz transformation is just the change of perspective between the account that Bob would offer based on thinking of his frame as stationary, and the one that Alice would offer based on thinking of her frame as stationary.

Key consequences of the Lorentz transformation include time dilation and length contraction. We're going to explain time dilation first because it's easier to describe. Suppose that at noon on Friday you get on a train at Princeton Junction. For convenience, we're going to say that this time and place correspond to the origin of Minkowski space, where the t and x axes cross. Now, there are fast trains and slow trains that go through Princeton Junction; some go north toward New York, and some go south toward Philadelphia; and you can decide which one you want. What you're going to do is ride the train for exactly one hour by your watch, and then get off and mark where you end up. Obviously, if you take a fast train, you get farther. But beware of the assumption that you get exactly twice as far riding a train that goes twice as fast. The tricky part is that you're riding the train for exactly one hour as measured by your own watch. The speed of a train is something observers who are stationary relative to the ground would measure, and their watches run a bit differently from yours because they're in a different frame of reference.

So where do you end up? More generally, if you and a bunch of friends all take different trains (all departing Princeton Junction at the same time), where do you all wind up? The answer is that you all wind up somewhere on a hyperbola in Minkowski spacetime (see figure 1.2). In other words, the hyperbola is the set of all possible final locations that you can reach after precisely one hour of your own travel time. One possible final location is Princeton Junction itself, at precisely 1 p.m. Princeton Junction time. The way you wind up there after an hour is if you are silly enough to spend an hour on a train that doesn't move at all. In that situation, of course it's 1 p.m. Princeton Junction time when you "arrive," because your frame of reference is the same as the

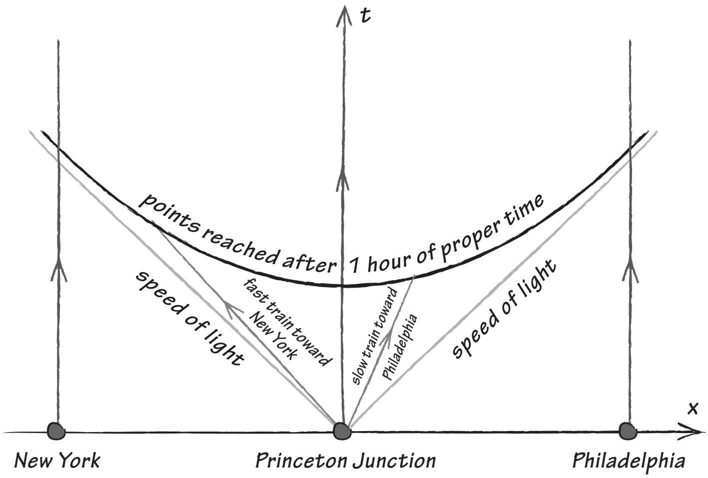

FIGURE 1.2. Trains from Princeton Junction. The curve showing points reached after 1 hour of proper time is a hyperbola.

station's, so that your watch exactly keeps pace with station time. If instead you get on a train that actually goes somewhere, your watch runs slower than station time, so when you get off after an hour of perceived travel time, station time is actually later than you think it ought to be. This later-than-you-think effect, known as time dilation, is captured in Minkowski spacetime by the way the hyperbola curves upward in the time direction as you go to locations farther and farther from your starting point.[1] Minkowski spacetime is sometimes called hyperbolic geometry, in reference to precisely the type of hyperbola we have been discussing.

1 The later-than-you-think effect for a normal train ride from Princeton to New York City amounts to approximately a hundred billionth of a second. So time dilation isn't going to make you late for work.

In Minkowski spacetime, we visualize the constant speed of light by drawing light rays at precisely a 45° angle relative to the vertical time axis. You'll notice that the hyperbola of possible endpoints for a one-hour train ride is wholly within the region of spacetime between two light rays emanating from the origin. This is the way Minkowski spacetime encodes the statement that none of our trains can go faster than light.

It may seem like our discussion of time dilation doesn't have much to do with Lorentz transformations. To see that it really does, let's go back to calling the reference frame of the train the A-frame, while the reference frame of the Earth is the B-frame. Suppose Alice spends an hour in the A-frame on her way from Princeton Junction to New York. Meanwhile, Bob and his friends remain stationary with respect to the Earth. How should they figure out the time of Alice's arrival? It's not very useful for her to call them when she arrives, because the signal she would use could only travel at the speed of light, and Bob and friends would have to do a calculation based on the time they received her call, the speed of the signal, and the distance to New York City to figure out when Alice arrived. That all sounds too tricky. So Bob figures out a better way. He synchronizes his watch with one of his friends, let's say Bill, and Bob and Bill take up positions at the Princeton Junction and New York City train stations, respectively. Bob measures when Alice leaves, and Bill measures when she arrives. No telephony required. It might seem tricky to synchronize watches in a reliable way between distant observers, but one good strategy would be for Bob and Bill both to start out halfway between Princeton Junction and New York City, synchronize their watches while standing next to one another, and then walk at identical speeds to their respective stations, all well before Alice boards her train.

In this whole narrative of Alice's train ride, the A-frame is clearly privileged, because Alice doesn't need any friends to figure out the duration of her train ride, whereas Bob and Bill must cooperate to make their measurement of the time. The time interval that Alice measures is called proper time because she measures it while remaining at a fixed location in her own frame of reference (the A-frame). The time interval that Bob and Bill measure is dilated time, which always must be greater than proper time. Dilated time is part of how the A-frame and B-frame perspectives on spacetime are related. The Lorentz transformation between the A-frame and the B-frame contains time dilation, and more.

A similar discussion can be used to describe length contraction. Instead of a train ride, let's imagine that Bob, Bill, and Alice go to the Olympics, where Alice hopes to set a record in pole-vaulting. Her secret is that she can run really fast, at 87% of the speed of light. (For some reason, she leaves the 100 meter dash to Usain Bolt, even though she figures she could post a time of under 0.4 microseconds.) Alice chooses a 6 meter pole, which is longer than most vaulters want, but after all she is pretty exceptional. Bob and Bill don't believe that Alice is using a pole that long, so they resolve to measure it as Alice charges down the track, holding her pole perfectly horizontal as she goes. Clearly, they've got a tough job. How can they actually make the measurement? Here is what they come up with. First, they synchronize their watches. Then they stand somewhat less than 6 meters apart, and they agree that at precisely the same time, they're going to glance up at Alice and record which part of her pole they see. After many attempts, they manage to arrange themselves so that Bob sees the tail end of Alice's pole, while Bill sees the front tip. Then they measure the distance between themselves. The answer is that they are only 3 meters apart. They reasonably conclude that Alice's pole is 3 meters long. They approach Alice and explain what they found. Alice protests that they can't have gotten it right. She enlists the help of her two friends, Allan and Avery, who run with her (apparently they're equally good sprinters) and measure her pole in her frame. The answer they find is that her pole is 6 meters long.

Once again, the A-frame is privileged in this discussion, because it's the frame in which Alice's pole is stationary. Its length as measured in the A-frame is called the proper length. Its length as measured in the B-frame is always shorter, and it is termed the contracted length. Time dilation and length contraction are closely linked, as we can appreciate in this example by considering what Alice would say about her experience running down the track toward the bar. As measured in her frame, it takes her half as long to get to the bar as Bob and Bill would have measured by the protocol we discussed above in reference to Alice's train ride to New York City. Time dilation, then, involves a factor of two for Alice's record-busting sprint at 87% of the speed of light. Length contraction also involves a factor of two: A-frame observers say her pole is 6 meters long, and B-frame observers say it is 3 meters long. In general, time dilation and length contraction always involve the same factor, sometimes called the Lorentz factor.

There seems to be a disconnect between our discussion of special relativity, which focuses on spacetime geometry, and the famous equation $E = mc^2$. Let's try to bridge this gap by considering a partial derivation of $E = mc^2$, where the most important steps can be illustrated geometrically. Our argument is only a partial derivation because it will involve some approximations and a couple of other formulas which we don't fully justify or derive.

The first step is to say in an equation what mass actually is. The best equation to use is $p = mv$, where p is the momentum and v is the velocity of a slowly moving massive body whose mass is m. The relation $p = mv$ comes directly from Newtonian mechanics, and it's OK for us to use it provided we make v much less than the speed of light. The next step is to relate energy to something. Here we are going to have to take yet another result from electromagnetism on faith: The momentum p of a light pulse is related to its energy E by the equation $p = E/c$. As we have learned, light pulses are peculiar in that they move at a fixed velocity in any frame. That's very unlike the way massive objects behave. In a given frame of reference, massive objects can either stand still, or they can move with some velocity v, which—according to special relativity—must always be less than the speed of light.

We now know the momentum of a massive object ($p = mv$) and the momentum of a light pulse ($p = E/c$). It would be wrong to set these two momenta equal, because massive objects are different from light pulses! What we must do instead is to figure out how to build a massive object out of light pulses. Then we will be able to use our momentum equations to derive $E = mc^2$.

Here is the crucial idea. Let's set up two perfectly reflecting mirrors exactly facing one another, and arrange for two identical light pulses to be going back and forth between the mirrors in such a way that they are always going in opposite directions. We claim that this setup is, effectively, a massive body. Assume that we can make the mirrors very light—so

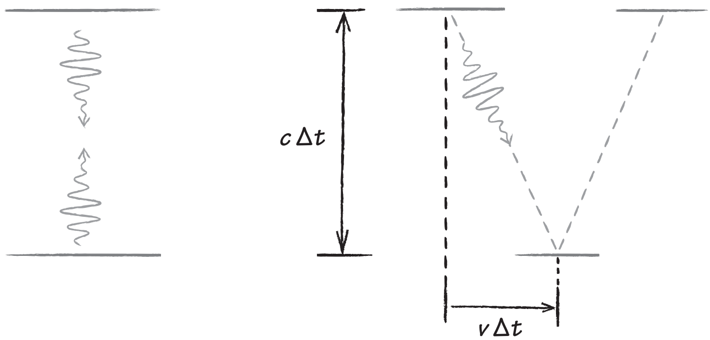

FIGURE 1.3. Left: Two identical light pulses travel back and forth between two mirrors. Right: The mirrors move to the right with a velocity v. In the time Δt that it takes one of the light pulses to get from one mirror to the other, the light pulse travels a distance of approximately $c\Delta t$ upward or downward, and a distance $v\Delta t$ sideways.

light that we can ignore them in calculations of the mass as well as the energy. Then the energy of our massive body is twice the energy of one of the light pulses. Its momentum is exactly zero because one light pulse has upward momentum at the moment when the other has downward momentum, and these upward and downward momenta cancel. They cancel because the body as a whole has no upward or downward motion; only its parts are moving.

To arrive at the right argument leading to $E = mc^2$, what we need to do is coax our whole light-and-mirrors contraption into motion. We're going to simplify the discussion by tracking the behavior of only one light pulse. If we tracked them both, we'd just get double the energy and double the mass. It's also going to simplify our discussion for the motion of our contraption to be sideways relative to the original up-and-down motion of the light pulse that we're tracking. Once this motion is in progress, the light pulse isn't just going up and down anymore. It's going a little sideways too. This is where geometry starts to come in. The light pulse's sideways motion is at a speed v, while its up-and-down motion is at a speed c. (Actually, its up-and-down motion is just a little slower than c because the *total* velocity of the light pulse is c. At the accuracy we need, it's OK to ignore this detail.) Another way to put it is that a fraction v/c of the motion of the light pulse is sideways. So it seems reasonable to assert that the sideways momentum p_{sideways} of the photon is v/c times its total momentum $p = E/c$. That is, $p_{\text{sideways}} = Ev/c^2$. We now assert that $p_{\text{sideways}} = mv$, which makes sense because p_{sideways} is the sideways momentum of the contraption as a whole (tracking only one of the two light pulses), and we're thinking of the contraption as a massive body. If we now combine our two ways of writing p_{sideways}, we get $\frac{Ev}{c^2} = mv$. Simplifying this equation, we get . . . drum roll . . . $E = mc^2$!

It might be objected that our light-and-mirrors contraption is very unlike the massive objects of everyday experience. That's not quite true. Protons and neutrons comprise most of the mass of everyday materials, and they can be approximately understood as little regions of spacetime inside which three nearly massless quarks bounce around at close to the speed of light. If this were the whole story, then the mass of the proton would come entirely from the motion of its constituent quarks, just as the mass of the light-and-mirrors contraption comes from the light pulses. In fact, there is more to the story: Quarks interact strongly with one another, and these interactions also contribute significantly to the total energy of the proton and hence its total mass. Nevertheless, the essential origin of most of the mass of everyday matter actually has more to do with our light-and-mirrors analysis than any intrinsic mass of fundamental constituents of matter.

The further we go with special relativity, the clearer it is that Maxwell's theory of electromagnetism is a key precursor to it. Better yet, it is in many ways a precursor to general relativity! Let's end this chapter with a tour of the highlights of Maxwell's amazing theory.

Before electromagnetism was properly developed, people understood the attraction between positive and negative charges in much the same way that Newton understood the gravitational pull between the Earth and the Sun. Briefly, they didn't really understand either one. Newton knew he didn't understand. He wrote of his quest to understand the origin of gravitational pull, "I have not as yet been able to discover empirically the reason for these properties of gravity, and I frame [or feign] no hypothesis." (This is an approximate translation from Newton's original Latin.) Of course, Newton had a highly useful quantitative law describing the strength of the gravitational pull. In particular, he knew that the pull weakens as the inverse square of the separation between the gravitating bodies. The attractive force between positive and negative charges follows a similar inverse square pattern. But it bothered him and his many successors that there could be *any* force acting over a distance. In other words, it's strange that a force on one object can be caused by the existence of another object which is far away. Michael Faraday championed the modern resolution of this puzzle. According to his ideas, a charged object both creates and responds to electric fields, which spread out in space according to four equations whose final form Maxwell discovered.

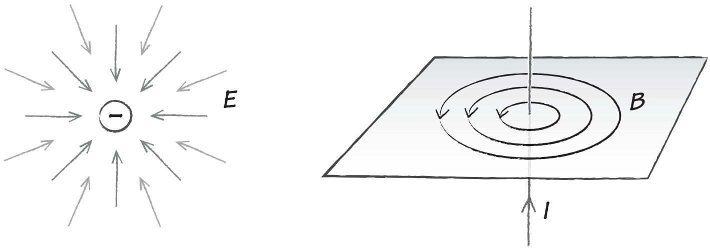

FIGURE 1.4. Left: The electric field E near a negative charge points everywhere inward. Right: A wire carrying a current I creates a magnetic field B that circulates around it.

In Faraday's picture, negative charges do *not* directly attract positive charges. Instead, a negative charge orients the nearby electric field so that it points directly toward the negative charge. The electric field in turn pulls on a positive charge at some distance away from the negative charge, and the net result is that the positive charge is drawn toward the negative charge. We could equally well say that the positive charge orients the nearby electric field to point directly away from it, and the electric field in turn exerts a pull on the negative charge. Both effects are happening at once. If all we watch is the charges, we conclude (correctly) that they feel equal and opposite forces drawing them together. Faraday's point is that these forces arise only through the action of the electric field, which has an existence independent of whatever charges might have produced it.

A similar story can be presented for magnetic forces and magnetic fields. Without entering into details, moving electric charges both create and respond to magnetic fields, which spread out in space in a fashion dictated by Maxwell's equations. A particularly important example is that magnetic fields form around a wire carrying an electric current. Electric current is the motion of microscopic charges inside the wire, so this is just a special case of the more general rule that moving charges produce magnetic fields.

Like electric fields, magnetic fields are supposed to have some existence independent of any particular configuration of moving charges that might produce them. To understand what we mean by this, let's consider a setup used by Maxwell in his development of the final form of electromagnetism. Put two metal plates parallel to one another without touching, and attach a wire to each one. This setup is known as a capacitor. Let an electrical current flow into one plate and out of the other. This flow results in an increase of positive charge in one plate over time (actually, a growing deficiency of electrons) and an equal increase of negative charge in the other plate (a superabundance of electrons). Because of the growing charge imbalance in the plates, there is a growing electric field between the plates. That electric field runs from the positively charged plate to the negatively charged plate, and as the charges of the plates grow in magnitude, so does the magnitude of the electric field.

We know that a magnetic field forms around a current-carrying wire. In particular, a magnetic field forms around the wires supplying current to the capacitor. But there is no current flowing from one plate to the other, and naively it would therefore seem that no magnetic field should be expected between the plates. Maxwell found this hard to reconcile with his understanding of capacitors, and he proposed an amazing solution: An increasing electric field generates a circulating magnetic field in the same way that a current does. This idea is an important step beyond the original notion that charges produce and respond to fields, because now we see that fields produce fields.

Actually, it was previously understood (by Faraday) that an increasing magnetic field generates a circulating electric field; this is essentially the principle on which electrical generators work. Two of Maxwell's four equations basically formalize these two reciprocal relations between electric and magnetic fields. The other two equations are simpler, saying that magnetic fields have no sources or sinks, whereas the only sources or sinks of electric field are positive and negative electric charges. All of Maxwell's equations are differential equations, which means that they are framed in terms of the rate of change of electric and magnetic fields over time, as well as the way in which these fields vary over space. The differential equations depend on the way fields behave in very small neighborhoods of spacetime. There is no action at a distance in Maxwell's equations. Everything is framed locally in terms of how nearby fields push and pull on one another.

Maxwell's greatest triumph was to show that his equations imply the existence of light. Light, as Maxwell understood it, is a combination of fluctuating electric and magnetic fields, where the spatial variation of the electric field causes the time variation of the magnetic field, and vice versa. The physical constants entering into Maxwell's equations describe the strength of electrostatic and magnetic interactions, but when they are combined in the right way, they give a numerical prediction for the speed of light—a prediction that can be verified experimentally.

Looking ahead, we will want to understand two crucial connections between electromagnetism and general relativity: Both theories involve Faraday's field concept, and both culminate in differential equations for the behavior of fields

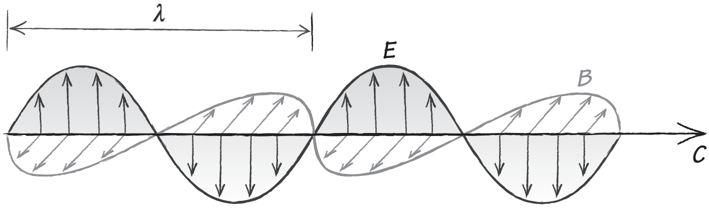

FIGURE 1.5. A light ray is a traveling disturbance of electric fields *E* and magnetic fields *B*, all moving in the same direction at the speed of light *c*. Interpreted as drawn to scale in the print edition of this book, the wavelength λ shown here is several centimeters, which is in the microwave range, a bit shorter than wavelengths found inside a typical microwave oven.

which imply some form of radiation. In electromagnetic radiation, electric fields beget magnetic fields, and vice versa, in a self-sustaining cascade through spacetime described by Maxwell's equations. This cascade has a characteristic wavelength, across which electric fields and magnetic fields vary from zero to their maximum value, back through zero to another maximum, and once again back to zero. Visible light is a special case in which the wavelength is about half a micron. Increasingly longer wavelengths lead to infrared light, microwaves, and radio waves, while shorter wavelengths give rise to ultraviolet light, X-rays, and gamma rays.

Einstein found the gravitational analogs of Maxwell's equations, and they are the main content of the general theory of relativity. The fields in Einstein's equations are stranger than electric and magnetic fields: Surprisingly, they are the curvature of spacetime itself. Another big surprise is that general relativity can describe massive objects in terms of pure geometry. This is quite unlike electromagnetism, in which charges remain fundamental throughout. These purely geometrical massive objects are black holes.

SPECIAL RELATIVITY