

Foundations

by Greg Egan

3: Black Holes

Copyright © Greg Egan, 1999. All rights reserved.

The previous article in this series began building the framework of ideas needed for general relativity by describing the geometry of *manifolds* — mathematical spaces without any notion of distance or angle — and then showing how it was possible to add a *metric* that defined these things in a very general way. The idea of *parallel transport* of a vector was introduced: moving along any path, you can carry a kind of “reference copy” of a vector from your starting point with you. A path is called a *geodesic* if it continues to follow the parallel-transported copy of its initial direction, never swerving away from its original bearing. Parallel transport of a vector around a closed loop can produce a reference copy back at the starting point that fails to match the original vector, and this effect is used to quantify the curvature of space (or spacetime), via the *Riemann curvature tensor*.

Einstein's equation links the curvature of spacetime with the presence of matter and energy. We haven't quite said all that we need to about curvature, but this article will begin by attacking the other side of the equation. This will give us some insight into why the equation takes the form it does, before we reach the final goal: examining one solution of the equation, the Schwarzschild solution, which describes a black hole.

Mass

If we want to quantify the amount of matter and energy in a region of spacetime, a good place to start is the idea of **mass**. According to Newtonian physics, when we weigh an object we're measuring the gravitational force that the Earth exerts upon it, and this force is taken to be proportional to the object's mass. Mass is usually defined quite differently, though, through the property of **inertia**: in the absence of complications like friction, when you apply a certain force to an object its rate of acceleration will be inversely proportional to its mass. Imagine pushing two items of furniture on frictionless pallets across a level surface; even though you're not opposing gravity, the same push will accelerate a 100-kilogram sofa half as much as a 50-kilogram bookcase.

Are these two ways of measuring mass — gravitational and inertial — necessarily equivalent? If they are, then neglecting the effects of atmospheric drag, a truck and a pebble should both fall off a cliff with the same acceleration all the way: however much harder it is to accelerate the truck, the gravitational force on it is proportionately greater. In a vacuum, all objects should fall to Earth at exactly the same rate, whatever their mass, and whatever they're made of. Centuries of experiments have confirmed that they do, so

this is no surprise to anyone at this point in history, but from a Newtonian perspective it's quite baffling that there are no known exceptions to this rule. No other force works like this. The electrostatic force between two objects depends on their electric charges; a proton and a positron have identical positive charges, but very different masses, so although they'll experience the same electrostatic force — the same push — if placed in the same electric field, they won't accelerate identically like the truck and the pebble. What's so special about the gravitational force that it's always perfectly matched to an object's inertial mass?

Einstein's answer is that gravity isn't a force at all. Rather, in the absence of forces, any object — whatever its mass and composition — simply follows a geodesic in spacetime: it takes the straightest possible world line in the direction it happens to be heading. In the curved spacetime near the Earth, the geodesic of an object that started out stationary would carry it straight to the centre of the planet if nothing got in its way. The only reason a pebble and a truck sitting motionless on the edge of a cliff aren't following such paths is because the cliff pushes up on them, with an electrostatic force between the electrons of the atoms at the surfaces making contact. The different forces required by the pebble and the truck to keep them from falling aren't really opposing two different “gravitational forces.” If you define an object's acceleration in curved spacetime as *the degree to which its world line fails to be a geodesic* — by analogy with the case in flat spacetime, where having a constant velocity means having a perfectly straight world line — the cliff is simply applying different forces to produce the same acceleration in two different masses.

If the idea that a motionless object can be accelerating strikes you as bizarre, imagine swinging a weight on the end of a rope: once it's swinging in a fixed circle, you still need to apply a constant force to accelerate it towards you, just to keep it from getting further away. What you're doing is *curving* a path that would otherwise be straight: cut the rope and the weight will fly off in a straight line. Letting the rope hang vertically is similar: the force you're applying to keep the weight motionless is still keeping its world line from being the straightest possible path through spacetime, a path that would carry it towards the Earth. Being “motionless in space” (relative to some massive object like the Earth) generally doesn't produce the straightest possible world line in curved spacetime. Compare this to a ship travelling east at a fixed latitude, say 45° S. The ship is “motionless” in the dimension of latitude — it's not drawing closer to either the south pole or the equator — but it can only do this if its engines are constantly applying a south-directed force to keep it from heading north along a great circle, the geodesic it would otherwise naturally follow if merely propelled forward.

So, your inertial mass tells you how much force must be provided (by the ground, or the floor, or the chair you're sitting in) to accelerate you sufficiently to keep you motionless with respect to the surface of the Earth, in exactly the same way as it tells

you how much force must be provided to accelerate you into motion. The idea of a “gravitational mass” that determines your response to a gravitational field is illusory. There is only one kind of mass: inertial mass.

However, as we'll see shortly, matter isn't the only thing with inertia.

Velocity and Acceleration

To provide a full description of matter and energy as the source of spacetime curvature, we need to introduce the relativistic versions of some simple ideas from classical physics. The ordinary velocity vector, \mathbf{v} , of an object in three dimensions tells you how fast the object is travelling in each of three directions — the velocity's coordinates v^x , v^y and v^z describe how fast the object's x , y and z coordinates are changing with time — and the length of \mathbf{v} is the speed of the object, how fast it's moving overall.

This tells you everything you need to know about an object's motion, but there's a way of “re-packaging” the same information that's more useful in relativity. People using different coordinate systems might disagree about every aspect of the three-dimensional vector \mathbf{v} : not just its individual coordinates, but even its overall length, the speed of the object. But what happens if we extend the vector into four dimensions? Let's define a vector \mathbf{u} called the **4-velocity** of the object, with coordinates u^x , u^y , u^z and u^t that describe how all four spacetime coordinates are changing for the object with time. Whose time? We want the 4-velocity to be independent of any coordinate system, so we define \mathbf{u} as the rate of change with respect to the time shown by a clock carried along with the object itself: this is known as **proper time**, and it's usually referred to by the Greek letter tau, τ . We're defining the 4-velocity \mathbf{u} as being ∂_τ , the rate of change of things with respect to τ . For example, $u^x = \partial_\tau x$: the x coordinate of \mathbf{u} is just the rate of change of the object's x coordinate, with respect to a clock moving alongside the object.




Fig 1: Ship's 4-velocity in Earth & Ship Coords

Consider a spaceship moving past the Earth with a constant speed of v , a situation where we only need to worry about one space coordinate, plus time. Call coordinates in which the Earth is stationary x and t , and coordinates in which the ship is stationary λ and τ . It's easy to describe the ship's 4-velocity \mathbf{u} in its own coordinates, because we've defined \mathbf{u} as ∂_τ . So $u^\lambda = \partial_\tau \lambda = 0$ (the ship is motionless in its own coordinates) and $u^\tau = \partial_\tau \tau = 1$ (the ship's clock keeps perfect time with respect to itself). Assuming that we've chosen coordinates for the ship in which the metric \mathbf{g} is just the Minkowskian metric, we then have:

$$\begin{aligned}
 \mathbf{g}(\mathbf{u}, \mathbf{u}) &= (u^\lambda)^2 - (u^\tau)^2 \\
 &= 0^2 - 1^2 \\
 &= -1
 \end{aligned} \tag{1}$$

The negative sign for $\mathbf{g}(\mathbf{u}, \mathbf{u})$ tells us that \mathbf{u} is a *timelike* vector, as you'd expect for the direction of an object's world line, and its length is the square root of $-\mathbf{g}(\mathbf{u}, \mathbf{u})$, which is just 1. To describe \mathbf{u} in Earth coordinates, we use the Lorentz transformation that we derived in the article on special relativity, rewritten slightly to apply to coordinate vectors rather than coordinates themselves:

$$\partial_\lambda = (\partial_x + v\partial_t) / \sqrt{1 - v^2} \tag{2a}$$

$$\partial_\tau = (v\partial_x + \partial_t) / \sqrt{1 - v^2} \tag{2b}$$

As in previous articles, we're making life simple by using units where the speed of light is equal to 1. Since $\mathbf{u} = \partial_\tau$, this immediately tells us:

$$\mathbf{u} = (v\partial_x + \partial_t) / \sqrt{1 - v^2} \quad (3a)$$

$$u^x = v / \sqrt{1 - v^2} \quad (3b)$$

$$u^t = 1 / \sqrt{1 - v^2} \quad (3c)$$

If the ship's speed v increases, both of the individual coordinates of \mathbf{u} grow larger, but due to the nature of the spacetime metric the effects on the overall length of \mathbf{u} cancel each other out. If we compute this with the Minkowskian metric in Earth coordinates:

$$\begin{aligned} \mathbf{g}(\mathbf{u}, \mathbf{u}) &= (u^x)^2 - (u^t)^2 \\ &= v^2/(1-v^2) - 1/(1-v^2) \\ &= -1 \end{aligned} \quad (4)$$

The agreement with Equation (1) should come as no surprise: the length of a spacetime vector is completely independent of the coordinates used. And since we can pick Minkowskian coordinates like λ and τ that are stationary with respect to *any* object — even in curved spacetime this is possible over a small region around the object at a given moment, just as we can always pick Euclidean coordinates over a small region of the Earth's curved surface — it's always going to be true that $\mathbf{g}(\mathbf{u}, \mathbf{u}) = -1$. The 4-velocity is always a unit timelike vector, a vector with a length of 1 that points along an object's world line. You can recover the object's ordinary velocity \mathbf{v} in a given coordinate system by taking the space coordinates of \mathbf{u} and dividing them by the time coordinate, e.g. for the example we've just given, in Earth coordinates, $v^x = u^x/u^t = v$.

Just as the acceleration of an object is defined in classical physics as the rate of change of its velocity with time, its **4-acceleration** vector, \mathbf{a} , is defined in relativistic physics as the rate of change of its 4-velocity with proper time. How can \mathbf{u} change, if its length must always be 1? Only by the object's world line changing direction in spacetime, which is what it means to change your ordinary velocity.

But how should we judge a “change of direction” when spacetime is curved? The physical evidence that your 4-velocity isn't “changing direction” is simply that you're weightless, because no force needs to act on you in order for you to follow your world line. If you're in a spaceship that's (a) orbiting the Earth, (b) falling straight towards a planet (without atmospheric drag), or (c) cruising through interstellar space, in all three cases you'll be weightless. In all three cases, you're following a geodesic. So acceleration means moving along a world line that is *not a geodesic*. This is true in either flat or curved spacetime, but to compute acceleration in curved spacetime you need to work out the change in an object's 4-velocity from moment to moment by using parallel transport to carry its earlier 4-velocity forward along its world line for comparison with the later value. This is known as taking the **covariant derivative** of

the vector \mathbf{u} , in the direction \mathbf{u} , which we write as $\nabla_{\mathbf{u}}\mathbf{u}$. So $\mathbf{a}=\nabla_{\mathbf{u}}\mathbf{u}$, and for a geodesic $\nabla_{\mathbf{u}}\mathbf{u}=0$.

In the previous article, we used the symbol ∇ to write the changes in coordinate vectors relative to their parallel-transported versions, e.g. on the surface of the Earth, using longitude and latitude as x and y coordinates, $\nabla_x\partial_x=(\sin y \cos y) \partial_y$. This means that as you travel east (take a covariant derivative in the x -direction, ∇_x) in the northern hemisphere (where $\sin y \cos y$ is positive), the local direction east (∂_x) “veers north” (in the direction of ∂_y) relative to a gyroscope bearing or a great-circle geodesic that was pointing east when you first set out. But you can take covariant derivatives in any direction, not just coordinate directions, and you can take covariant derivatives of any vector, not just the coordinate vectors. All you have to do is ask how the vector changes relative to a parallel-transported copy of its initial value, as you travel in the specified direction.

Energy and Momentum

Another powerful concept from classical physics is the **momentum** vector for an object, which is just its velocity vector multiplied by its mass: $\mathbf{p}=\mathbf{mv}$. This quantifies the intuitive notion that a 1-gram bullet travelling at 1 kilometre per second, and a 1-kilogram bowling ball travelling at 1 metre per second, have something in common. Since force is defined as mass times acceleration, and acceleration is the rate of change of velocity, force can equally well be defined as *the rate of change of momentum*. This tells us just what it is that the bullet and the bowling ball have in common: to bring them to a halt in one second, to reduce their momentum to zero, you'd need to apply exactly the same force, 1 Newton, in either case.

Momentum turns out to be *conserved*: for a collection of objects — maybe interacting among themselves, but subject to no external force — the total momentum never changes. Why not? When the objects aren't interacting, they're subject to no forces at all, so they'll simply keep moving with whatever constant velocities they happened to possess. When two of the objects do interact, they'll exert equal and opposite forces on each other, and whatever change in momentum one of them experiences as a result, the other will experience an equal and opposite change. The total momentum vector remains constant.

A closely related idea is that of **kinetic energy**, K , which is a number rather than a vector. Energy in general can be defined as the capacity to “do work,” in the technical sense of moving a load some distance against a resisting force — it's no coincidence that this idea developed most rapidly in the age of steam engines. Suppose you extract energy from a moving object of mass m and speed v by making it drive a piston that resists its motion with a constant force, bringing it to rest in a time t . The

object's average velocity over that period will be $v/2$, so it will travel a distance of $vt/2$. Its deceleration will be v/t , and the force needed to produce this will be mv/t . So the "work done" by the object will be the force applied, mv/t , times the distance moved, $vt/2$, which comes to $mv^2/2$. This is the classical formula for kinetic energy: $K=mv^2/2$. Although the bullet and the bowling ball mentioned earlier have the same momentum, the kinetic energy of the bullet is a thousand times greater: both can be stopped in 1 second by a force of 1 Newton, but the bullet will travel 500 metres in that time (averaging half its initial speed of 1 km/sec, as the force gradually decelerates it), the bowling ball a mere half a metre.

Energy in general turns out to be conserved, like momentum, but kinetic energy is often converted into other forms when objects interact. Sometimes these forms are really just kinetic energy "in disguise": the frictional heating or sound produced by most objects colliding is mainly just a transfer of kinetic energy from the colliding objects to individual molecules. But kinetic energy can also be converted into various kinds of **potential energy**: when you release a plucked guitar string, its energy cycles back and forth between the kinetic energy of motion and the potential energy stored by the material of the string when it's stretched — though of course it all eventually leaks away as sound, and a tiny amount of heat. Like kinetic energy, changes in potential energy can sometimes be "disguised" because they're happening down at the level of individual molecules. When a meteor hits the Earth, most of its kinetic energy ends up as heat, some of which goes to drive chemical reactions in the surrounding rock — rearrangements of atoms which change their overall electrostatic potential energy.

Because the momentum vector $m\mathbf{v}$ and the kinetic energy $mv^2/2$ depend on the ordinary velocity of objects, they depend on the coordinate system you're using. In Newtonian physics that's not a problem: if people are playing pool on a train, a Newtonian analysis in either pool-table coordinates *or* coordinates fixed to a point on the ground beside the track will show energy and momentum being conserved (if absolutely everything, from chemical energy in the players' muscles to the sound of every collision is taken into account). In ground-based coordinates everything on the train will be moving much faster, but that will be equally true before and after each shot.

However, since the classical law of conservation of momentum is phrased in terms of vectors in space, not vectors in *spacetime*, it shouldn't really come as a surprise that it needs to be modified in order to work in relativistic physics.




Fig 2: Spring-loaded objects in x,t frame

If we examine a simple case where the old formulation goes wrong, it's not hard to see what changes need to be made. Figure 2 is a spacetime diagram showing two objects of equal mass, m , pushed apart by coiled springs. One ends up travelling left with a speed of v , and the other travelling right with the same speed. The initial momentum of the system, which we'll call $\mathbf{p}_{\text{before}}$, is obviously zero. The ordinary velocities of the objects after the springs push them apart are $\mathbf{v}_1 = -v\partial_x$ and $\mathbf{v}_2 = v\partial_x$, so their momenta are $\mathbf{p}_1 = -mv\partial_x$ and $\mathbf{p}_2 = mv\partial_x$. The total momentum of the system, $\mathbf{p}_{\text{after}} = \mathbf{p}_1 + \mathbf{p}_2$, is still zero.




Fig 3: Spring-loaded objects in λ,τ frame

Now we re-analyse these events in coordinates which are moving to the right with

a speed of v , relative to our previous choice — coordinates which follow one of the moving objects after the springs have uncoiled. We'll call these coordinates λ and τ ; Figure 3 is a spacetime diagram in which ∂_λ and ∂_τ are drawn as perpendicular. In classical physics, we could convert all the ordinary velocities to the new coordinates just by subtracting the vector $v\partial_x$ from them; that's known as a **Gallilean** coordinate transformation, and it would be appropriate for comparing the two perspectives on our rail-car pool game. But if we're assuming that v is a significant fraction of lightspeed, we need to treat the shift in coordinates as a rotation in spacetime.

To transform all the ordinary velocities, we first need to write the objects' 4-velocities in the original coordinates. Making use of Equation (3a), these are:

$$\mathbf{u}_0 = \partial_t \quad (5a)$$

$$\mathbf{u}_1 = (-v\partial_x + \partial_t) / \sqrt{1 - v^2} \quad (5b)$$

$$\mathbf{u}_2 = (v\partial_x + \partial_t) / \sqrt{1 - v^2} \quad (5c)$$

Having done this, we can apply a Lorentz transformation, which converts the coordinate vectors to the new system:

$$\partial_x = (\partial_\lambda - v\partial_\tau) / \sqrt{1 - v^2} \quad (6a)$$

$$\partial_t = (-v\partial_\lambda + \partial_\tau) / \sqrt{1 - v^2} \quad (6b)$$

Substituting these expressions into Equations (5) gives:

$$\mathbf{u}_0 = (-v\partial_\lambda + \partial_\tau) / \sqrt{1 - v^2} \quad (7a)$$

$$\mathbf{u}_1 = (-2v\partial_\lambda + (1+v^2)\partial_\tau) / (1 - v^2) \quad (7b)$$

$$\mathbf{u}_2 = \partial_\tau \quad (7c)$$

We can now compute the ordinary velocities in the new coordinates, by dividing u^λ by u^τ in each case:

$$\mathbf{v}_0 = -v\partial_\lambda \quad (8a)$$

$$\mathbf{v}_1 = (-2v/(1+v^2))\partial_\lambda \quad (8b)$$

$$\mathbf{v}_2 = 0 \quad (8c)$$

Equation (8b) illustrates an important phenomenon: the relativistic addition of ordinary velocities. You might have been wondering how to reconcile the fact that speeds greater than lightspeed are impossible, with the idea of two spaceships heading away from Earth at, say, 75% of lightspeed in opposite directions. Wouldn't each ship think of the other as moving at 150% of lightspeed? Equation (8b) shows that the speed they'd actually

measure for each other would be $(-1.5/(1+.75^2))$, which is 96% of lightspeed. Compare this with the following situation: you're walking due north, a trail on your left runs slightly north of north-west (4 metres north for every 3 metres west, to be precise), and a trail on your right runs slightly north of north-east (4 metres north for every 3 metres east). Both these trails are moving 3 metres further away from you, sideways, for every 4 metres you advance northwards. Now, suppose you were walking on the left-hand trail. Would you expect the right-hand trail to grow precisely 6 metres further away to your right, for every 4 metres you advanced in the direction you're now walking? Of course not: the trails will separate “sideways” much faster than that, because your idea of “sideways” slices through them very differently now. In the case of the ships, because we're dealing with spacetime geometry, their world lines will separate more *slowly* in the direction one of them would consider to be “space” than you'd predict by adding up two velocities based on Earth's idea of the direction of “space.”

If we use Equations (8) to compute the total momentum of the system before and after the springs uncoil, $\mathbf{p}_{\text{before}}=2m\mathbf{v}_0=-2mv\partial_\lambda$, since the combined objects have mass $2m$, and $\mathbf{p}_{\text{after}}=m\mathbf{v}_1=(-2mv/(1+v^2))\partial_\lambda$, since the second object is stationary and contributes no momentum. These are obviously not the same! Under a Galilean transformation of velocities, \mathbf{v}_1 would just be $-2v\partial_\lambda$ and the two values would agree, but the Lorentz transformation “spoils” everything.

What we've shown is that different observers won't even agree as to whether or not the classically-defined momentum vector has been conserved! Fortunately, there's a closely related spacetime vector that *is* conserved — and since it's a *spacetime* vector, this is a claim that has nothing to do with any particular observer or coordinate system.

The **4-momentum** vector \mathbf{P} is defined as the 4-velocity \mathbf{u} multiplied by the **rest mass** of the object: $\mathbf{P}=m\mathbf{u}$. The “rest mass” of an object is just the inertial mass as we've already defined it, with the proviso that you measure it at a nice low velocity, much smaller than the speed of light; we'll soon see why this is important. Since every object's 4-velocity in its own coordinates is just $\mathbf{u}=\partial_\tau$, every object's 4-momentum in the same coordinates is $\mathbf{P}=m\partial_\tau$. In coordinates x and t in which the object has a speed of v , Equations (3) yield:

$$\mathbf{P} = m(v\partial_x + \partial_t) / \sqrt{(1 - v^2)} \quad (9a)$$

$$P^x = mv / \sqrt{(1 - v^2)} \quad (9b)$$

$$P^t = m / \sqrt{(1 - v^2)} \quad (9c)$$

Just as every object's 4-velocity has a length of 1, every object's 4-momentum has a length of m , its rest mass. This is obvious when we write $\mathbf{P}=m\partial_\tau$, and though it's a little harder to see when we look at a description in someone else's coordinates, the fact remains that everyone will agree on the length of a spacetime vector, so everyone will

agree on an object's rest mass.

Examining Equation (9b), we see that the component of the 4-momentum in the spatial direction looks like the ordinary momentum of an object with mass $m/\sqrt{1-v^2}$, moving with a speed of v . This means, for example, that P^x for any object moving at 80% of lightspeed will be $(1/\sqrt{1-.8^2})=1.67$ times greater than the ordinary momentum p^x for an object with the same mass and speed. What are we to make of the “extra” momentum? This effect is sometimes described by saying that moving objects “gain mass” — though like the idea that moving clocks “run slow,” it isn't really describing any change in the object itself, just a change in your relationship with it. If you apply a force to a particle moving through your laboratory at 80% of lightspeed, and a clock on the wall tells you that the interaction lasted for a nanosecond, a clock moving alongside the particle would only record $\sqrt{1-.8^2}=.6$ nanoseconds of proper time. If you overestimate how long you've applied the force, you'll expect more acceleration than you actually get, and blame the difference on increased mass. It's the rate of change of 4-velocity with proper time that measures an object's true acceleration, and if you stick rigorously to that spacetime view, you never need use any other mass than the rest mass.

Still, objects moving at relativistic speeds *are* effectively harder to push around than their rest mass and velocity alone would suggest — and we already have a name from classical physics for what they've gained: kinetic energy. Taking that point of view, what Equation (9b) is telling us is that kinetic energy, just like matter, possesses *inertia*.

This can be made clearer if we subtract out the rest mass of the object and see what remains. It can be shown that $1/\sqrt{1-v^2} - 1$ is approximately equal to $v^2/2$ for values of v much smaller than 1. (It would be too much of a detour to explain the mathematics behind this claim, but if you doubt it just grab a calculator and work out the two expressions for $v=.001, .002, .003$ and see how close they are in all cases.) The extra mass that a moving object seems to possess, $m/\sqrt{1-v^2} - m$, is then approximately equal to $mv^2/2$, which is the classical formula for kinetic energy. The exact, relativistic formula for kinetic energy is $K=m/\sqrt{1-v^2} - m$.

Equation (9c) shows that the time coordinate of the 4-momentum is equal to $m/\sqrt{1-v^2}$, the kinetic energy plus the rest mass, or **total energy**, E , of the object. You can think of an object's total energy as that part of its momentum that's pointing in the time direction (for some particular observer's definition of time), rather than any spatial direction, making it “momentum standing still” (also known as “inertia”) — or you can think of spatial momentum as energy that looks as if it's moving through space, because the observer is moving relative to the object. From the point of view of the object itself, all its momentum is just rest mass moving through time, $\mathbf{P}=m\partial_\tau$. But however you look at it, the 4-momentum encapsulates both the ordinary momentum and the energy of an object — for this reason it's sometimes referred to as the **energy-momentum** vector —

and there's no need for a separate law of conservation of energy: conservation of 4-momentum does it all.

But if kinetic energy is handled automatically by the spacetime geometry of the 4-momentum vector, how do we account for potential energy? Going back to our spring-loaded projectiles of Figure 2, it turns out that the only way to make the time coordinates of the before and after 4-momenta match up is by realising that the mass of the combined objects *with coiled springs* has to be greater by a factor of $1/\sqrt{1-v^2}$, due to the potential energy in the springs, than it would be if the springs were slack and the objects were at rest. Potential energy must have inertia too.

Measuring the extra mass of a compressed spring is probably a lost cause, but the same effect shows up very starkly in nuclear physics. The different arrangements of protons and neutrons that form atomic nuclei have different potential energy, and if you compare the mass of a given nucleus with the mass of an equal number of separated protons and neutrons, there's a significant difference, known as the **mass defect**. Both nuclear fission and nuclear fusion rearrange nuclei into new combinations with less potential energy than the starting ingredients, extracting the difference as kinetic energy.

What's more, just as kinetic and potential energy can be converted into each other, it's now well known that matter itself can be converted into energy, and vice versa. A particle of matter and a particle of antimatter can combine and annihilate each other; the immediate result is usually two photons, which are particles with zero rest mass — *all* their energy is kinetic energy. How much mass translates into how much energy? In units where c is equal to 1, energy is measured in exactly the same units as mass, so Einstein's famous " $E=mc^2$ " hasn't appeared in any of our calculations. If we'd been using more conventional (but less convenient) units, " mc^2 " would have popped up all over the place instead of " m ."

The usual definition of 4-momentum, $\mathbf{P}=\mathbf{mu}$, doesn't apply to particles with zero rest mass. Rather, a photon's 4-momentum is the null spacetime vector (a vector with an overall length of zero, also known, appropriately, as a *lightlike* vector) whose time coordinate for a given observer is equal to the energy, E , that the observer considers the photon to possess, and whose spatial component points in the direction of the photon's motion. There's a simple relationship between the 4-momentum, based on energy, and the *propagation vector*, based on wavelength (which we used in the article on special relativity when deriving the Doppler shift). In units where $c=1$, a photon's 4-momentum is equal to Planck's constant times the propagation vector.

The Stress-Energy Tensor

All forms of energy have inertia, and everything with inertia must contribute to the curvature of spacetime. So the 4-momentum vector, which keeps track of all forms of

energy, must play a crucial role in describing the source of the curvature of spacetime.

Something's missing, though. The Earth has a certain 4-momentum, which reflects its rest mass and its path through spacetime. If we crushed the Earth down to the size of a boulder, that super-dense, Earth-mass boulder would have exactly the same 4-momentum as the Earth itself. But the boulder would only have the same effect on spacetime as the Earth up to the point where the surface of the planet had once been: satellites would still orbit an Earth-mass boulder in exactly the same way (give or take some tiny deviations caused by the planet's actual lumpiness), but the gravitational field near the centre of the boulder would be very different from the field near the centre of the Earth.

What's missing from the 4-momentum is any notion of **density**. Ordinarily, we think of density as mass per unit volume, say kilograms per cubic metre, and there's no reason why the inertial mass due to various forms of energy can't be included in this — or to put it another way, why we can't look at the total **energy density** in spacetime, counting rest mass as a form of energy, along with kinetic and potential energy.

The total energy of an object is equal to the time coordinate of its 4-momentum, so it depends on whose idea of “time” you're using. The volume of the object also depends on a choice of direction for time, since this determines precisely which directions in spacetime count as “space.” There's a phenomenon similar to time dilation, known as “length contraction”: if a spaceship flew past the Earth at 80% of lightspeed, we'd measure the distance between the world lines for its frontmost and hindmost points along a different direction in spacetime than the people on board, and conclude that the ship was only 60% the length they considered it to be. Like proper time and rest mass, the astronauts' own measurement of their ship's length would be more sensible than ours, but that doesn't change the fact that the ship's energy density would seem greater to us, both from the kinetic energy added to its rest mass, and the way the total energy seemed to be packed into 60% the volume.

So the notion of energy density is very much observer-dependent, but there's still a way to keep our description of it nice and universal. What we need is a vector or tensor that can be used to calculate the energy density of a system according to any observer, just as the 4-momentum vector \mathbf{P} can be used to calculate the total energy.

Suppose the system we're trying to describe is simply an object of rest mass m , with a 4-velocity of \mathbf{u} , and hence a 4-momentum of $\mathbf{P} = m\mathbf{u}$. Let's call the 4-velocity of the observer \mathbf{w} , to distinguish it from that of the object. Then the total energy of the system is just the time coordinate, in the observer's frame, of the object's 4-momentum: $E = -\mathbf{g}(\mathbf{P}, \mathbf{w}) = -m \mathbf{g}(\mathbf{u}, \mathbf{w})$, where we're using the spacetime metric, \mathbf{g} , to “project out” the component of \mathbf{P} in the direction of the unit timelike vector \mathbf{w} .

To find the volume that the observer would measure the system as having, we take the **proper volume** V — the volume we'd measure if we were at rest with respect

to the system — and divide it by $-\mathbf{g}(\mathbf{u}, \mathbf{w})$. Why? The “length contraction factor” needed to adjust the volume of, say, the spaceship mentioned earlier, comes from comparing Earth-based and ship-based spacelike vectors that run along the axis of the ship. But the angle in spacetime between those two vectors is exactly the same as the angle between \mathbf{u} and \mathbf{w} — just like the identical angles between ∂_t and ∂_τ and between ∂_x and ∂_λ in Figure 1 — so we can obtain this factor by applying the metric to \mathbf{u} and \mathbf{w} , rather than going to the trouble of calculating the spacelike vectors themselves. All we need is a minus sign to correct for the fact that we've used timelike vectors, not spacelike ones.

Combining these results, we find that the energy density an observer with 4-velocity \mathbf{w} will measure for the system is:

$$\begin{aligned} \text{energy density} &= -m \mathbf{g}(\mathbf{u}, \mathbf{w}) / (V / -\mathbf{g}(\mathbf{u}, \mathbf{w})) \\ &= \rho \mathbf{g}(\mathbf{u}, \mathbf{w}) \mathbf{g}(\mathbf{u}, \mathbf{w}) \end{aligned} \quad (10)$$

where we've introduced the symbol ρ (the Greek letter rho, which is traditionally used for density) for m/V , the rest mass of the system divided by its proper volume, or the **proper density** of the system.

To go any further, we need to introduce some new terminology. If you're working on a manifold with a metric, \mathbf{g} , you can uniquely identify a 1-form \mathbf{f} with any vector \mathbf{v} , and vice versa, by imposing the requirement that $\mathbf{g}(\mathbf{v}, \mathbf{w}) = \langle \mathbf{f}, \mathbf{w} \rangle$ for any other vector \mathbf{w} . How are \mathbf{f} and \mathbf{v} related geometrically? The contours of \mathbf{f} must be perpendicular to \mathbf{v} , so that if \mathbf{w} is also perpendicular to \mathbf{v} , i.e. $\mathbf{g}(\mathbf{v}, \mathbf{w}) = 0$, motion in the \mathbf{w} direction won't cross the contours of \mathbf{f} at all, yielding $\langle \mathbf{f}, \mathbf{w} \rangle = 0$. The coordinates of \mathbf{f} are easy to find: for example, $f_x = \langle \mathbf{f}, \partial_x \rangle = \mathbf{g}(\mathbf{v}, \partial_x)$.

Nice as it would be if the coordinate 1-forms (such as dx) and the coordinate vectors (such as ∂_x) were equivalent in this sense, that's only true when the coordinate vectors are all mutually perpendicular spacelike unit vectors. This holds for rectangular coordinates in space, but not for Minkowskian spacetime coordinates, the one hitch being that it's $-dt$, not dt , that's equivalent to ∂_t , because $\mathbf{g}(\partial_t, \partial_t) = -1$, whereas $\langle dt, \partial_t \rangle = 1$.




Fig 4: Vectors & Equivalent 1-forms

Given this ability to use the metric to convert back and forth between vectors and 1-forms, we can convert any tensor of rank (r,s) into another tensor of the same total rank $r+s$, but which acts on a different combination of vectors and 1-forms, e.g. a tensor of rank $(r-1,s+1)$ or $(r+1,s-1)$. This process is known as **raising and lowering indices**, because the coordinates of tensors are written with subscripts for any 1-forms in the tensor product and superscripts for any vectors.

For example, suppose we define a tensor \mathbf{T} of rank $(2,0)$ with the equation:

$$\mathbf{T} = \rho \mathbf{u} \otimes \mathbf{u} \quad (11)$$

We can “lower both indices” of this tensor to produce another “version” of \mathbf{T} , of rank $(0,2)$, by replacing \mathbf{u} with an equivalent 1-form \mathbf{f} for which $\langle \mathbf{f}, \mathbf{w} \rangle = \mathbf{g}(\mathbf{u}, \mathbf{w})$ for any vector \mathbf{w} . We could give this new version another name, but it's common practice to use a single name for all the versions of a tensor, because it's really just another way of describing the same thing.

$$\mathbf{T} = \rho \mathbf{f} \otimes \mathbf{f} \quad (12)$$

We can now use this tensor to describe the energy density we calculated in Equation (10):

$$\begin{aligned} \text{energy density} &= \rho \mathbf{g}(\mathbf{u}, \mathbf{w}) \mathbf{g}(\mathbf{u}, \mathbf{w}) \\ &= \rho \langle \mathbf{f}, \mathbf{w} \rangle \langle \mathbf{f}, \mathbf{w} \rangle \\ &= \mathbf{T}(\mathbf{w}, \mathbf{w}) \end{aligned}$$

The tensor \mathbf{T} defined by Equation (11) is known as the **stress-energy tensor** for the system. The values of \mathbf{T} throughout a region of spacetime can be thought of as describing a “current” of 4-momentum, \mathbf{P} , giving both the density of \mathbf{P} and the direction in which it's moving. For a particle, the 4-momentum “flows” in the same direction as it points: along the particle's world line. But in more complicated systems, such as those with “shear stress” which we'll describe shortly, momentum can be transported in a direction other than that in which it points.

To see how the stress-energy tensor works, let's check that we can recover the object's energy density in its own coordinates.

$$\begin{aligned}\mathbf{T}(\mathbf{u},\mathbf{u}) &= \rho \mathbf{g}(\mathbf{u},\mathbf{u}) \mathbf{g}(\mathbf{u},\mathbf{u}) \\ &= \rho (-1)(-1) \\ &= \rho\end{aligned}$$

In general, we can *define* the stress-energy tensor, \mathbf{T} , of any system, by the requirement that $\mathbf{T}(\mathbf{u},\mathbf{u})$ is the total energy density of the system according to an observer with 4-velocity \mathbf{u} .

Equation (11) tells us how to construct \mathbf{T} for a single object, such as an asteroid, from its proper density and 4-velocity. \mathbf{T} has different values from point to point in spacetime: in the vacuum around the asteroid \mathbf{T} is zero, whereas inside the asteroid $\mathbf{T}=\rho \mathbf{u}\otimes\mathbf{u}$, and if the density ρ varies from point to point because of the presence of different minerals, \mathbf{T} will reflect that variation.

For more complicated systems, it takes more work to construct the stress-energy tensor. Using Minkowskian x , y , z and t coordinates for our observer, it turns out that the requirement we've used to define \mathbf{T} — that $\mathbf{T}(\partial_t,\partial_t)$ is the total energy density — also demands that $\mathbf{T}(\partial_x,\partial_t)$, $\mathbf{T}(\partial_x,\partial_x)$, $\mathbf{T}(\partial_y,\partial_x)$ and so on, tell us something analogous. In effect, if \mathbf{T} is to work for absolutely *any* observer, the geometry has to make sense even when we substitute unit spacelike vectors in place of the observer's 4-velocity.

Actually, the completely general case is easier to describe if we talk about the (2,0) version of \mathbf{T} , which accepts two unit 1-forms, say \mathbf{i} and \mathbf{j} , rather than two vectors. In that case, $\mathbf{T}(\mathbf{i},\mathbf{j})$ is the density of the \mathbf{i} coordinate of the 4-momentum, in a spacetime region that lies in the contours of \mathbf{j} . If \mathbf{j} is dt for some observer, the contours of \mathbf{j} will lie in what that observer considers to be “space,” and if \mathbf{i} is also dt , the density of the t coordinate of the 4-momentum is the energy density according to that observer, the result we've already described. But if \mathbf{i} is a spacelike 1-form instead, say dx , then $\mathbf{T}(dx,dt)$ will be the density of the x coordinate of momentum.




Fig 5: Density of 4–momentum in two regions

The situation is slightly trickier to interpret if \mathbf{j} is spacelike, e.g. dx . First, suppose that \mathbf{i} is spacelike too. The region in question has two of its dimensions in space and one in time, since the vectors ∂_y , ∂_z and ∂_t all lie in the contours of dx . (In Figure 5, we've left out the z direction, because we can only draw three dimensions at once, so the three-dimensional “ dx ” region is drawn here as a square.) How do we interpret the “density of momentum” in a region like that: a two-dimensional area in the yz plane, swept through an interval of time?

Density is usually the measure of something “per volume,” which is “per length, per length, per length” for each of the dimensions defining that volume. What we have here is a density that is “per length, per length, per time,” or “per area, per time.” In fact, we have a density that is “momentum per area, per time,” or equally well, “momentum per time, per area.” As Figure 5 illustrates, particles that contribute to the density of momentum in this region of spacetime *cross* the two-dimensional area of space during the time under consideration, and so they contribute to the rate of transfer of momentum from one side to the other. The rate of change of momentum per time is force, and force per area is pressure. When \mathbf{i} and \mathbf{j} are spacelike, $\mathbf{T}(\mathbf{i},\mathbf{j})$ measures pressure!

Actually, the term “pressure” is usually reserved for the case where the force is perpendicular to the area involved, as in most gases or liquids: if you're deep in the ocean, the water pushes directly against every exposed surface with exactly the same pressure, and there's no significant sideways force. $\mathbf{T}(dx,dx)$ gives you the pressure of such a fluid, and it will be the same as $\mathbf{T}(dy,dy)$ and $\mathbf{T}(dz,dz)$. Only in viscous fluids, or solids (such as the Earth's mantle and crust) is it possible to have “shear stresses,” sideways forces that are trying to deform the material rather than just compress it. These show up in the stress-energy tensor as values for $\mathbf{T}(dx,dy)$, $\mathbf{T}(dx,dz)$, etc.

In the case of the Earth, the effect on the gravitational field of pressure and shear stresses is infinitesimal, and so long as you get the density of rest mass right, you'll be able to calculate spacetime curvature in and around the planet with great precision. However, if you're an astronomer studying white dwarves or neutron stars, the pressure in the interiors of such highly compressed objects can be great enough to have a significant effect on their gravitational field.

The last combination to consider is a spacelike \mathbf{j} with a timelike \mathbf{i} . Again, the contours of \mathbf{j} define an area followed over an interval of time, so the “density” we get is again a rate of change with time, per unit area — in this case, of the timelike \mathbf{i} -coordinate of momentum, i.e. of energy. For example, $\mathbf{T}(dt,dx)$ measures the **energy flux** across the yz plane. This might record something like the flow of energy from sunlight — though of course rest mass counts as energy too, so a flying brick would be an equally good example.

Conservation of 4-momentum in Curved Spacetime

In relativistic physics, the 4-momentum \mathbf{P} takes over the role of classical energy and momentum as the quantity that is conserved for any isolated system: so long as no external forces are applied, the total 4-momentum of the system won't change. For a lone object cruising through space along a geodesic, we can write conservation of 4-momentum in a very straightforward way: $\nabla_{\mathbf{u}}\mathbf{P}=\nabla_{\mathbf{u}}(m\mathbf{u})=m\nabla_{\mathbf{u}}\mathbf{u}=0$. This is both a “global law” where we can make comparisons between times that are far apart — it's meaningful to talk about \mathbf{P} at time $t=0$ being equal to \mathbf{P} at time $t=1000$, since the object's world line provides an obvious path to use to parallel-transport the earlier 4-momentum forward for comparison — and a “local law” that applies from instant to instant to dictate the shape of the world line: conservation of 4-momentum, $\nabla_{\mathbf{u}}\mathbf{P}=0$, is easily seen to be equivalent to the statement that the world line is a geodesic, $\nabla_{\mathbf{u}}\mathbf{u}=0$.

However, for a complicated system spread over a large volume of space, there might not be any obvious way to add up all the \mathbf{P} vectors at two different times, and then compare them. In general, it's easier to concentrate on a local statement of the conservation law, and it turns out that there's a way to do this that applies absolutely everywhere. *In any small region of spacetime, the 4-momentum that flows in must equal the 4-momentum that flows out.* This is true even when the region isn't isolated from external forces, because we can take account of those forces by treating them as a flow of momentum across the boundaries of the region, just as we did when considering the role of pressure in the stress-energy tensor.

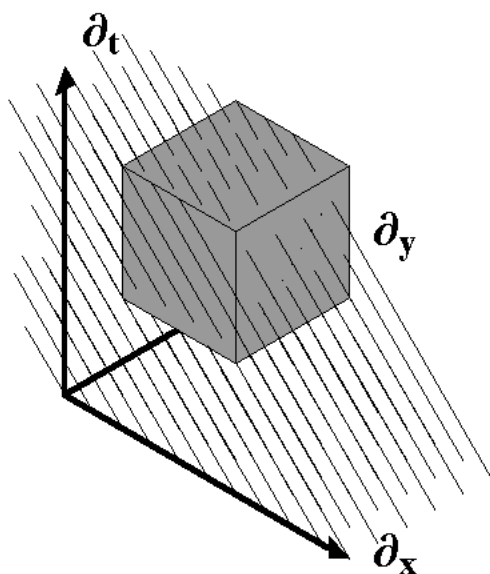


Fig 6: Conservation of 4-momentum

Suppose you decide to observe the conservation of 4-momentum in a region of spacetime that is a certain cubic metre of your back yard, over a time of one minute, from noon until 12:01. 4-momentum can “flow into” the region in either of two ways: in the time direction — just by being in the right place already, like the rocks and ants that were in the chosen space at noon — or in a spatial direction, like the insects that crawl or fly in during the chosen period. Similarly, 4-momentum can “flow out” either by still being there at 12:01, like the rocks and some of the insects, or by sneaking out earlier. Everything that was there initially, plus everything that entered, minus everything that exited, minus everything that was still there at the end ... leaves you with nothing. What we want is a mathematical version of this statement, a measure of the combined inflow and outflow of 4-momentum that we know will be equal to zero.

The only trouble with analysing a cubic metre of garden is that the density of 4-momentum varies enormously from place to place (there's much more energy density in rock than in air, for example). So let's consider instead a region of spacetime so small that the stress-energy tensor, \mathbf{T} , is almost constant, and its *rate of change* in any direction can be considered constant.

What exactly do we mean by the rate of change of the stress-energy tensor in a given direction? It should be clear by now that in curved spacetime, the *only* standard against which things can be judged to have changed is parallel transport, so what we need is a definition of parallel transport for a tensor. This turns out to be especially easy for a tensor of the form $\mathbf{a} \otimes \mathbf{b}$, where \mathbf{a} and \mathbf{b} are vectors: you just parallel-transport the vectors separately, then take the tensor product. For example, since parallel transport of the 4-velocity \mathbf{u} of a free-falling object along that object's geodesic world line always produces a reference copy of \mathbf{u} that exactly matches the actual 4-velocity at each point, a free-falling

object whose proper density is unchanging will also have a stress-energy tensor, as defined by Equation (11), in agreement everywhere with a parallel-transported reference copy. The covariant derivative of the stress-energy tensor along the world line — the rate of change between the tensor itself and a reference copy of an earlier version — will thus be zero: $\nabla_{\mathbf{u}}\mathbf{T}=0$.

Returning to our tiny spacetime region, assume for the sake of simplicity that we've chosen units such that the dimensions of the region in both space and time are all equal to one. Focus on the coordinate of the 4-momentum in some direction \mathbf{i} . The amount of \mathbf{i} -coordinate present initially in the region is $\mathbf{T}(\mathbf{i},dt)$ evaluated at $t=0$, and the amount present finally is $\mathbf{T}(\mathbf{i},dt)$ evaluated at $t=1$. So the net outflow from the spacetime region in the time direction is equal to *the rate of change* of \mathbf{T} in the time direction, $\nabla_t\mathbf{T}(\mathbf{i},dt)$, multiplied by the length of time being considered, 1.

Similarly, the amount of \mathbf{i} -coordinate flowing in through the $x=0$ side of the cube is $\mathbf{T}(\mathbf{i},dx)$ evaluated at $x=0$, and the amount flowing out through the $x=1$ side of the cube is $\mathbf{T}(\mathbf{i},dx)$ evaluated at $x=1$. So the net outflow is $\nabla_x\mathbf{T}(\mathbf{i},dx)$. Identical results hold for the y and z sides of the cube. So in order for the net outflow in all directions to come to zero, we must have:

$$0 = \nabla_x\mathbf{T}(\mathbf{i},dx)+\nabla_y\mathbf{T}(\mathbf{i},dy)+\nabla_z\mathbf{T}(\mathbf{i},dz)+\nabla_t\mathbf{T}(\mathbf{i},dt) \quad (13)$$

An expression like this, where the rate of change is taken in the same direction as one of the coordinate 1-forms fed into a tensor, and the results added up for all possible coordinate directions, is known as the **divergence** of the tensor, $\text{div } \mathbf{T}$. Since there's one "slot" into which we can still feed any 1-form, \mathbf{i} , $\text{div } \mathbf{T}$ here is defining a rank (1,0) tensor — which is really just a vector. So our local law of conservation of 4-momentum can be written as:

$$\text{div } \mathbf{T} = 0 \quad (14)$$

and interpreted as saying that the amount of 4-momentum being conjured up out of thin air in every unit 4-volume of spacetime is zero. A tensor that has a divergence of zero is described as being **divergence free**.

There's one form of energy from classical physics that we've deliberately left out of the stress-energy tensor: "gravitational potential energy." The reason we've left it out, and the reason we're putting it in quotes, is because, like "gravitational force," there's no need for such a thing in general relativity. According to Newtonian physics, when you toss a ball into the air, its kinetic energy is converted into gravitational potential energy as it rises above the ground. In general relativity, once the ball leaves your hand it simply follows a geodesic, and there's no need to worry about potential energy — the curved

geometry of spacetime accounts for everything. By using the covariant derivative in Equation (13) and, implicitly, Equation (14), measuring all changes against the standard of parallel transport and geodesics, we're putting the burden that used to be carried by "gravitational potential energy" entirely on the geometry, where it belongs.

The Einstein Tensor

The stress-energy tensor \mathbf{T} is all we need to describe the presence of matter and energy, but there are still two problems standing in the way of equating \mathbf{T} with spacetime curvature. The first is that the Riemann curvature tensor \mathbf{R} is a tensor of rank (1,3): you can feed it a 1-form and three vectors to get a number, or feed it three vectors and leave the first slot "unfed" to get a vector, but however you look at it, it's something quite different from \mathbf{T} , which we've defined as having rank (2,0) or (0,2). Raising and lowering indices won't help: \mathbf{R} has a total rank of four, and \mathbf{T} has a total rank of two.

The other problem is that spacetime can be curved even in a vacuum, where $\mathbf{T}=0$. The reason the Earth orbits the sun is because of spacetime curvature *due to the sun*, whereas the only thing contributing to \mathbf{T} at the Earth is the Earth itself. The Earth's own density has nothing to do with the orbit it's following; a piece of styrofoam placed the same distance from the sun, with the same velocity, would follow the same orbit.

Fortunately, each of these problems sheds light on the other. We can't set \mathbf{R} equal to \mathbf{T} , because the tensors are the wrong rank — but that would be a bad idea anyway, because it would imply that spacetime was flat wherever there was a vacuum. So \mathbf{T} must be equated with some aspect, or "part," of spacetime curvature that we've yet to identify, something that can be zero in a vacuum *without* making \mathbf{R} itself zero.

How can we find the appropriate aspect of curvature? Newtonian gravity comes to the rescue: it turns out that there's a very simple classical calculation we can do, relating *the density of matter to the coming together of objects in free fall*, which points to the need for a similar relationship in general relativity. Suppose the Earth suddenly gave way beneath our feet and began to collapse under its own gravity — all the forces within the rock below that prop it up having magically vanished. The instant that happened, the surface of the Earth would still be stationary, so if you asked "how fast is the Earth shrinking?" the answer would be "not at all, right now." However, it wouldn't be stationary for long, so you could ask instead "at what rate is the Earth's volume 'accelerating' towards a smaller value?"

In Newtonian physics, the acceleration due to gravity at a distance r from a mass of m is given by $a=\kappa M/r^2$, where κ is the "universal gravitational constant." (You're probably used to seeing this written as G , not κ . Annoyingly, in general relativity G has come to be used for the Einstein tensor, which we'll describe shortly, so the gravitational constant is written as κ instead.) The surface area of a sphere is $4\pi r^2$, and multiplying

this by the acceleration downwards shows that the volume of the Earth will be “accelerating” at a rate of $-4\pi\kappa M$. As a proportion of the total volume of the Earth, V , this is just $-4\pi\kappa(M/V) = -4\pi\kappa\rho$, where ρ is the average density of the Earth.

What we've been calling the “acceleration” of the volume is the rate of change (with time) of the rate of change (with time) of volume, so we can write this result as:

$$(\partial_t \partial_t V)/V = -4\pi\kappa\rho \quad (15)$$

We've only shown this for one particular situation, but it turns out that any small collection of particles in free fall through a region where the density is ρ will have a volume V that changes according to Equation (15). In a vacuum, where $\rho=0$, a volume that starts out unchanging will never change. Imagine a small cloud of space junk, initially motionless with respect to the Earth, high above the atmosphere. If this junk then falls straight down, the shape of the cloud will change: it will grow narrower in all horizontal directions, as individual particles fall straight towards the centre of the Earth, while growing longer vertically, as particles that were initially closer to the Earth experience a slightly greater gravitational acceleration (in the Newtonian view) than particles that were higher up, and so increase their head start even more. But these two changes cancel out exactly, and the overall volume of the cloud won't change.

In general relativity, $\mathbf{T}(\partial_t, \partial_t)$ measures density, so Equation (15) suggests that we should look for a tensor, let's call it \mathbf{C} , such that $\mathbf{C}(\partial_t, \partial_t)$ is the second rate of change with time of a unit volume bounded by geodesics, since geodesics are the world lines of particles in free fall. We could then try to relate \mathbf{C} to \mathbf{T} in an analogous relativistic equation.

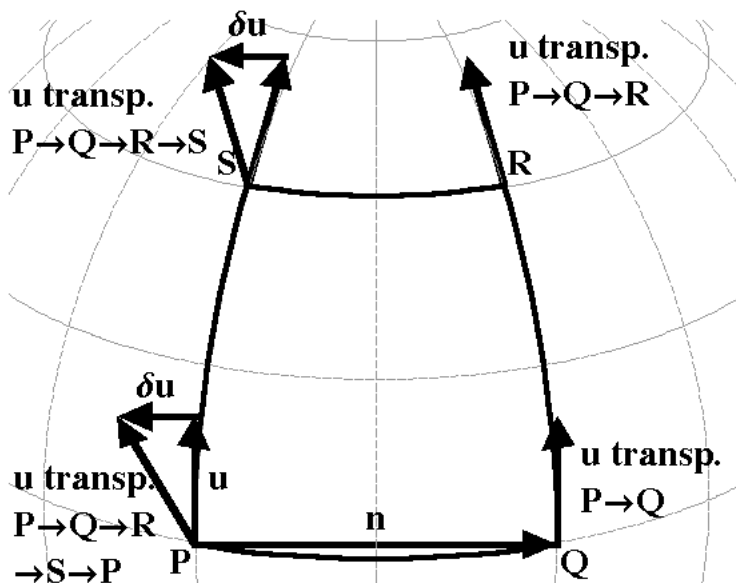


Fig 7: Geodesic Deviation

It's not hard to find the second rate of change of the separation between individual geodesics; this is known as **geodesic deviation**. Figure 7 shows two nearby geodesics, PS and QR, that both start out pointing in the direction \mathbf{u} , and are separated initially by a unit vector \mathbf{n} . (We're dealing with a small enough region of spacetime that it's meaningful to compare vectors at different points, and to describe the separation between points with a vector.) If we parallel-transport \mathbf{u} from one geodesic to another (P to Q), forward a unit distance along the second geodesic (Q to R), back to the first geodesic (R to S), and finally back to its starting point (S to P), then it will return with a small change, $\delta\mathbf{u}$, which we can compute with the Riemann curvature tensor. Since the plane of the loop we've moved \mathbf{u} around is defined by the vectors \mathbf{u} and \mathbf{n} , and the vector we're transporting is \mathbf{u} , we have:

$$\delta\mathbf{u} = -\mathbf{R}(\mathbf{u}, \mathbf{n}, \mathbf{u})$$

But \mathbf{u} doesn't change relative to the geodesics as it's parallel-transported along them, between Q and R and between S and P — that's the definition of geodesics — so we can attribute this entire discrepancy, $\delta\mathbf{u}$, to the difference in direction of the geodesics at S and R. Since the two geodesics start out parallel, the first rate of change of their separation \mathbf{n} is zero. But since they nonetheless manage to acquire a relative “tilt” of $\delta\mathbf{u}$, after we follow them a unit distance in the \mathbf{u} direction, the *second* rate of change of their separation is $\delta\mathbf{u}$, which is $-\mathbf{R}(\mathbf{u}, \mathbf{n}, \mathbf{u})$. In other words:

$$\nabla_{\mathbf{u}}\nabla_{\mathbf{u}}\mathbf{n} = -\mathbf{R}(\mathbf{u}, \mathbf{n}, \mathbf{u}) \quad (16)$$

To compute the second rate of change in the *volume* between the geodesics of a whole cluster of particles (which we'll assume for simplicity to have an initial volume of 1), we need to take the second rate of change of the distance between them in each of the three dimensions perpendicular to \mathbf{u} , and add up the results. But we might just as well do this over all four coordinate directions instead, because any contribution parallel to \mathbf{u} will always be zero. We can write this most succinctly by defining a new tensor, known as the **Ricci tensor**, extracting the second rate of change of distance in each of the coordinate directions by feeding a coordinate 1-form into the very first slot of \mathbf{R} (the one that we usually leave “unfed” in order to get a vector, rather than a number, as the final result) while setting \mathbf{n} , the initial separation between geodesics, to the corresponding coordinate vector.

$$\begin{aligned}\mathbf{Ricci}(\mathbf{v}, \mathbf{w}) &= \mathbf{R}(dx, \mathbf{v}, \partial_x, \mathbf{w}) + \mathbf{R}(dy, \mathbf{v}, \partial_y, \mathbf{w}) + \\ &\quad \mathbf{R}(dz, \mathbf{v}, \partial_z, \mathbf{w}) + \mathbf{R}(dt, \mathbf{v}, \partial_t, \mathbf{w}) \\ (\partial_{\mathbf{u}} \partial_{\mathbf{u}} V)/V &= -\mathbf{Ricci}(\mathbf{u}, \mathbf{u})\end{aligned}\tag{17}$$

A tensor defined this way — by slotting coordinate 1-forms and vectors into another tensor and adding up over all the coordinate directions — is called a **contraction** of the original tensor. We say that the Ricci tensor is “the contraction of the Riemann tensor on its first and third slots.” You can form a contraction over any two slots of a tensor, but if they both take vectors or both take 1-forms, you must lower or raise one index first, so you can feed coordinate vectors to one, and coordinate 1-forms to the other. If you don't, the result isn't coordinate independent.

The negative of the Ricci tensor gives the proportional second rate of change of the volume between geodesics, which we'd like to relate somehow to the stress-energy tensor \mathbf{T} . In analogy to Equation (15), a reasonable first guess would be:

$$\mathbf{Ricci} = 4\pi\kappa \mathbf{T} \text{ (maybe?)}$$

There's a problem, though: if you calculate $\text{div } \mathbf{Ricci}$, the divergence of the Ricci tensor, it's not zero. This means the equation we've just written is incompatible with $\text{div } \mathbf{T} = 0$, the conservation of 4-momentum!

Luckily, it turns out that we can use the Ricci tensor to construct another tensor that *is* divergence free. First, define a contraction known as the **Ricci scalar**, which is normally written as R (not in bold face, since it's a number, not a tensor). Because the Ricci tensor as we initially defined it had rank (2,0), we have to perform the contraction on a version which has had one index lowered, to become rank (1,1).

$$R = \mathbf{Ricci}(dx, \partial_x) + \mathbf{Ricci}(dy, \partial_y) +$$

$$\mathbf{Ricci}(dz, \partial_z) + \mathbf{Ricci}(dt, \partial_t) \quad (18)$$

There's a certain combination of the Ricci tensor, the metric \mathbf{g} , and the Ricci scalar that's divergence free. This is known as the **Einstein tensor**, and it's always written as \mathbf{G} .

$$\mathbf{G} = \mathbf{Ricci} - (R/2)\mathbf{g} \quad (19)$$

In the next section we'll say a bit about why this tensor is divergence free, but before doing that let's write the equation connecting \mathbf{G} to the stress-energy tensor. First, note that in Minkowskian coordinates:

$$\begin{aligned} \mathbf{G}(\partial_t, \partial_t) &= \mathbf{Ricci}(\partial_t, \partial_t) - (R/2)\mathbf{g}(\partial_t, \partial_t) \\ &= -(\partial_t \partial_t V)/V + (R/2) \end{aligned}$$

using Equation (17), and the fact that the Minkowskian metric gives $\mathbf{g}(\partial_t, \partial_t) = -1$. Now, in spacetime that isn't very strongly curved, the Ricci scalar, R , turns out to be “dominated” by the last term in Equation (18), $\mathbf{Ricci}(dt, \partial_t)$. Because we're using Minkowskian coordinates, the equivalent expression for the (0,2) tensor is $-\mathbf{Ricci}(\partial_t, \partial_t)$, which in turn is equal to $(\partial_t \partial_t V)/V$. So $\mathbf{G}(\partial_t, \partial_t)$ is approximately equal to $(-\partial_t \partial_t V)/2V$ — half the value we'd get from the Ricci tensor — and to be compatible with Equation (15), we must have:

$$\mathbf{G} = 8\pi\kappa \mathbf{T} \quad (20)$$

This, at last, is the Einstein equation, linking spacetime curvature with the density of matter and energy!

This equation is not unique in meeting the requirement that $\text{div } \mathbf{T} = 0$. Because of the compatibility of the metric with parallel transport, all covariant derivatives of the metric are zero, and hence the divergence of any constant multiple of the metric is also zero. So there's no fundamental reason why the true equation for spacetime curvature might not be:

$$\mathbf{G} + \Lambda \mathbf{g} = 8\pi\kappa \mathbf{T} \quad (21)$$

The symbol Λ (this is a Greek letter, the capital lambda) stands for a number called the **cosmological constant**, and its value is still very much a matter of debate. A negative Λ would cause empty spacetime to be curved as if it contained energy; a positive Λ would cause it to be curved as if it contained “negative energy,” in the sense that it would cause geodesics to move apart rather than come together. When Einstein first developed

general relativity, he chose a small positive value for Λ that would balance the curvature caused by the overall density of matter in the universe, keeping everything static, because at the time there was little observational evidence to support what is now common knowledge: the universe is expanding. When Einstein learnt of this, he declared the cosmological constant to be the greatest mistake of his life, and decided that the true value was exactly zero. However, recent astronomical observations suggest a positive value, sufficient not only to overcome the mutual attraction of matter, but to cause the universe to expand ever more rapidly in the future. Whether or not this is the final verdict, there's still plenty of scope for quantum mechanical treatments of the vacuum, and of gravity itself, to shed more light on the issue of *why* Λ takes whatever value it actually has.

Although Λ is immensely important in cosmology, on any “small” scale — at least up to the size of clusters of galaxies! — it's definitely insignificant, and for the remainder of this article we'll simply assume that $\Lambda=0$, and use Equation (20).

The Bianchi Identity

Figure 8 shows a path that leads from a point, P, around a small cube whose edges are all one unit long, and point in the directions \mathbf{u} , \mathbf{v} and \mathbf{w} . This path traverses every face of the cube exactly once, but it traverses every edge an even number of times, backwards as many times as forwards.




Fig 8: Path to Show Bianchi Identity

If you parallel-transport a vector \mathbf{b} around this path, it will come back unchanged, because every step you travel along an edge in one direction, you eventually travel again in reverse, undoing the effect. However, we can write this *overall* lack of change as a

sum of the changes we get from parallel transport around six simple loops: in each of three planes defined by pairs of the three vectors (e.g. \mathbf{u} and \mathbf{v}), we do one loop for the face of the cube that's closest to P, and another for the opposite face, which is displaced one unit in the direction of the remaining vector (e.g. \mathbf{w}). For the loop around the opposite face we have to get there and back from P along an edge of the cube, but since we use the same edge for both trips, the effect of that part of the path cancels out.

We move around opposite faces in opposite directions; for example, as we travel around the closest face to P in the \mathbf{u} - \mathbf{v} plane, the change in \mathbf{b} is $\delta\mathbf{b} = -\mathbf{R}(\mathbf{b}, \mathbf{u}, \mathbf{v})$, but for the opposite face it's $\delta\mathbf{b} = -\mathbf{R}(\mathbf{b}, \mathbf{v}, \mathbf{u}) = \mathbf{R}(\mathbf{b}, \mathbf{u}, \mathbf{v})$. However, these two terms might not cancel each other out, because \mathbf{R} can be different on the two faces. Different by how much? By the length of the distance between the faces, which is one unit, times the rate of change of \mathbf{R} in the direction \mathbf{w} , which is $\nabla_{\mathbf{w}}\mathbf{R}$. So the change in \mathbf{b} due to these two loops is $\nabla_{\mathbf{w}}\mathbf{R}(\mathbf{b}, \mathbf{u}, \mathbf{v})$. Combining this for all three planes, and equating it to the overall result of zero change that we know we must get, yields:

$$0 = \nabla_{\mathbf{w}}\mathbf{R}(\mathbf{b}, \mathbf{u}, \mathbf{v}) + \nabla_{\mathbf{u}}\mathbf{R}(\mathbf{b}, \mathbf{v}, \mathbf{w}) + \nabla_{\mathbf{v}}\mathbf{R}(\mathbf{b}, \mathbf{w}, \mathbf{u}) \quad (22)$$

This equation is known as the **Bianchi identity**, and it's the reason that \mathbf{G} is divergence free. We won't go through the proof that $\text{div } \mathbf{G} = 0$, but basically it consists of a bit of algebraic rearrangement of Equation (22). So you can ultimately trace the fact that $\text{div } \mathbf{G} = 0$ back to Figure 8, and what it says about the way changes in curvature must fit together over any volume of spacetime.

There are two ways to interpret this. One is to take $\text{div } \mathbf{G} = 0$ as merely a handy clue that \mathbf{G} is the correct choice of tensor to equate with \mathbf{T} , since we already know that $\text{div } \mathbf{T} = 0$. Another is to consider Einstein's equation as *explaining* conservation of 4-momentum. Given Einstein's equation, 4-momentum must be conserved, because $\text{div } \mathbf{G} = 0$ isn't an additional, physical hypothesis that might or might not hold, it's a geometrical tautology: the undeniable fact that every edge in the cube in Figure 8 is traversed in opposite directions an equal number of times.

The Schwarzschild Solution

In empty space, $\mathbf{T}=0$, so Einstein's equation becomes $\mathbf{G}=0$, and since most of the universe is near enough to vacuum, metrics whose curvature satisfies the “vacuum Einstein equation” are enormously important. One obvious vacuum solution is flat Minkowskian spacetime: if the Riemann curvature tensor \mathbf{R} is zero, **Ricci** and \mathbf{G} are also zero. This is a pretty good description of small regions of interstellar and intergalactic space — though not of the galaxy, or the universe, as a whole.

A more interesting vacuum solution is that which allows the moon to orbit the

Earth, and planets to orbit the sun. To analyse the spacetime geometry around a star or a planet, we'll assume that the geometry is **spherically symmetrical**. It turns out that there's only one possible "class" of solutions that meet this criterion, all with the same general shape. The sole freedom left is to plug in a number that lets you set the scale — and by comparison with Newtonian gravity it's easy to identify that number with the mass of the star or planet that lies at the centre of the vacuum geometry.

This class of solutions is known collectively as the **Schwarzschild solution**, and the metric is given by Equation (23). M here stands for the mass of the star, and we've chosen units where not only is the speed of light, c , equal to 1, but the gravitational constant κ is also 1. This makes all the algebra much simpler, and though it's a pain to convert to and from conventional units, the less cluttered equations in between are generally worth it. In **geometric units**, as this system is called, everything is measured in distances — we'll use metres. Time is measured in metres (the time it takes light to travel 1 metre, 3.3 nanoseconds), and mass is measured in metres (the mass that Newtonian gravity predicts would cause an acceleration, at a distance of 1 metre, of 1 metre per metre squared; this is 1.35×10^{27} kilograms, making the mass of the sun, 2×10^{30} kilograms, equivalent to about 1480 metres).

$$\mathbf{g} = -(1-2M/r) dt \otimes dt + 1/(1-2M/r) dr \otimes dr + r^2(\cos \theta)^2 d\phi \otimes d\phi + r^2 d\theta \otimes d\theta \quad (23)$$

The spacetime coordinates used for the Schwarzschild metric are called r , ϕ , θ and t . If you picture a sphere centred on the star, ϕ can be thought of as the longitude and θ the latitude of any point on the surface of that sphere. (It doesn't matter where you put the "equatorial plane" and which hemisphere you call "north," because the geometry is spherically symmetrical.) If you compare the part of the metric involving ϕ and θ with the metric we derived in the previous article for the surface of the Earth, you'll see that it's identical; we've just changed the names of the coordinates from x and y , and the radius of the sphere from E to r .

So we can imagine the star surrounded by spheres like onion layers, each with a different r coordinate, and each with the same geometry as the surface of a sphere in Euclidean space with a radius of r . The surface area of each onion layer is $4\pi r^2$, and since you can measure this without going any nearer to the star, this offers the simplest way to interpret r . But is the r coordinate actually the distance to the centre of each sphere? No. Distance is defined by the metric, and assuming that you're stationary relative to the star, so that ∂_r is your idea of a purely spatial direction, $|\partial_r| = \sqrt{\mathbf{g}(\partial_r, \partial_r)}$ is equal to $1/\sqrt{1-2M/r}$. For r greater than $2M$, this will be greater than 1, which means that distances measured radially are going to be greater than changes in the r coordinate. There are "more onion layers" packed in here than there would be in Euclidean space.

That tells us a bit about the geometry of space according to stationary observers, but what about the passage of time? It's sometimes said that “clocks run slow” in a strong gravitational field, and there are a number of works of science fiction where the protagonists deliberately travel close to a massive object (such as a black hole, of which we'll have more to say shortly) in order to experience additional time dilation, aging even less compared to Earth-bound people than they would from the effects of travelling through flat spacetime at the same velocity. This effect is certainly real, but the statement about clocks “running slow” needs to be treated as cautiously as the same statement about moving clocks. No clock ever truly *runs slow* unless it's broken — and blaming the “flow of time” is as misleading as blaming the “flow of distance” if you happen to travel from one town to another by a longer route than someone else. Some paths through spacetime from A to B are simply shorter than others, and while curvature complicates this whole business, clocks are no more “slowed down” by gravity than your odometer is “sped up” when you drive over a mountain and register more kilometres from one side to the other than someone who took a road tunnel instead.

It's straightforward in principle to use the metric of Equation (23) to find the proper time along any world line, but the detailed calculations for a complete journey to and from the vicinity of a massive object are a bit too messy to present here. Fortunately, there's a much easier way to quantify gravitational “time dilation” that also tells us something about the view of the stars from near such an object. Suppose you follow the world line of a photon, as it travels from a point in space far from the object and strikes the eye of someone who is stationary relative to the object. That is, someone whose world line is a line of constant r , ϕ and θ , and hence whose 4-velocity will be pointing solely in the direction of ∂_t . Everyone's 4-velocity \mathbf{u} must satisfy $\mathbf{g}(\mathbf{u},\mathbf{u})=-1$, so if $\mathbf{u}=u^t\partial_t$:

$$\begin{aligned}
 -1 &= \mathbf{g}(\mathbf{u},\mathbf{u}) \\
 &= (u^t)^2 \mathbf{g}(\partial_t,\partial_t) \\
 &= -(u^t)^2 (1-2M/r) \\
 u^t &= 1/\sqrt{1-2M/r} \\
 \mathbf{u} &= 1/\sqrt{1-2M/r} \partial_t
 \end{aligned} \tag{24}$$

This tells us, incidentally, that the t coordinate isn't a measure of proper time for our observer, any more than the r coordinate is a measure of proper distance. The proper time that elapses along this observer's world line will be less than any change in the t coordinate, because $\partial_\tau t$ — that is, the rate of change of t with respect to proper time τ — is equal to $\mathbf{u}(t)=1/\sqrt{1-2M/r}$, which is greater than 1.

The t coordinate is useful, though: because it doesn't appear directly in the metric, Equation (23), the geometry of spacetime is independent of the value of t . You

can think of the whole of Schwarzschild spacetime as being made up of lots of slices with different values for t , all piled one on top of the other, with the pile stretching from the past into the future. Unlike the onion layers of different r coordinates, which each have the geometry of a different-sized sphere, all these t -slices are identical. In fact, you can take any shape “drawn” on spacetime and increase the t -coordinate of every point by the same amount, and the new version will be identical to the original.

Ways of moving things that preserve their size and shape in this way are called **isometries** (Greek for “same distance”), and the vectors that produce them, such as ∂_t , are known as **Killing vectors** (after the mathematician Wilhelm Killing). For example, adding thirty degrees to the longitude of every point on the coastline of Africa would just rotate the continent around the Earth, leaving its size and shape unchanged, whereas adding thirty degrees to the latitude of every point would distort the shape enormously. The longitude coordinate vector is a Killing vector, the latitude coordinate vector isn't.

Though we won't prove it, the projection of a Killing vector onto the tangent to a geodesic is the same everywhere along that geodesic. (If you want to test this claim with a simple example, consider the projection of the longitude coordinate vector onto the tangent to a great circle.) In the Schwarzschild geometry, since ∂_t is a Killing vector and the world line of an astronaut in free fall is a geodesic, $\mathbf{g}(\partial_t, \mathbf{w})$ is constant for the astronaut's 4-velocity \mathbf{w} . A photon's world line is also a geodesic, but in that case we have to use the photon's 4-momentum, \mathbf{P} , as the tangent. (The 4-velocity of a photon is a meaningless idea, because the 4-velocity must have a length of 1, but any lightlike vector has a length of zero.) The energy that an observer with 4-velocity \mathbf{u} measures for a photon is $\mathbf{g}(\mathbf{u}, \mathbf{P})$, so using the value of \mathbf{u} from Equation (24) we have:

$$\begin{aligned} E &= \mathbf{g}(\mathbf{u}, \mathbf{P}) \\ &= \mathbf{g}(1/\sqrt{1-2M/r} \partial_t, \mathbf{P}) \\ &= 1/\sqrt{1-2M/r} \mathbf{g}(\partial_t, \mathbf{P}) \end{aligned} \quad (25)$$

Since \mathbf{P} is the tangent to a geodesic, and ∂_t is a Killing vector, $\mathbf{g}(\partial_t, \mathbf{P})$ must be constant along the photon's entire path. Let's call this constant value E_∞ , since for very large values of r , $1/\sqrt{1-2M/r}$ gets so close to 1 that it might as well be 1, and hence the energy someone far away would measure for the photon is just $\mathbf{g}(\partial_t, \mathbf{P})$. This lets us write:

$$E = 1/\sqrt{1-2M/r} E_\infty \quad (26)$$

This equation is known as the **gravitational blue shift**, since it describes how the energy of a photon looks greater — pushing it towards the blue end of the spectrum — to someone deeper in a gravitational field. For example, at a distance of $r=5.55M$,

$E=1.25E_{\infty}$, so an observer would see all the stars in the sky as being 25% “bluer” than someone far away in space.

Because the energy of light is proportional to its frequency — the number of complete oscillations the light wave performs in a second — this immediately tells us something about time as well. By measuring a greater energy for the photon, our observer is also using the light as a signal to compare his or her local clock with a clock far away, and by this method, local time seems to be “running slower” by 25%. This is *not* to say that the frequency of distant stars represents some kind of absolute standard for time. Like the comparison between two clocks in relative motion that we made in the article on special relativity, this is just a way of drawing a connection between two different observers — both of whom are correctly measuring proper time along their respective world lines.

However, it does offer a useful way to get an approximate idea of the effect on relative aging of going near a massive object. If you start out in a mother ship far from the object, descend in a scout ship to a certain r coordinate, and then return, you will have been struck (at some point) by every single wavefront of light from the stars as the people who stayed behind. But at each r value, Equation (26) implies that the time you would have measured between wavefronts was different by a factor of $\sqrt{1-2M/r}$ from that measured on the mother ship. If you spent a large part of your journey hovering at, say, $r=3.125M$, to a good approximation you'll have experienced a total elapsed time only 60% as much as the other travellers.

For values of r smaller than $2M$, $\mathbf{g}(\partial_r, \partial_r)$ is negative, meaning that the coordinate vector ∂_r has switched from being spacelike to being timelike! Similarly, $\mathbf{g}(\partial_t, \partial_t)$ is positive, showing that ∂_t has become spacelike. The distance $2M$ in geometric units is known as the **Schwarzschild radius** for a given mass, and an object that becomes compressed to within its Schwarzschild radius collapses into a **black hole**. You don't need to know the “distance to the centre” of such an object: because r is defined in terms of surface area, any non-rotating spherically symmetric object whose surface area is less than or equal to $16\pi M^2$ is doomed to become a black hole.

Why is this unavoidable? The fact that ∂_r becomes timelike for r less than $2M$ means that “motion” in the r direction becomes the same as motion in any other timelike direction. We have no choice about the fact that our world lines run into the future, so any object that crosses the onion layer at $r=2M$, the **event horizon**, must have a world line that runs in the direction of decreasing r . That's the definition of “the future” within the event horizon.

Couldn't you change the direction of the future by changing your velocity? Yes, but not enough. Figure 9 shows the light cones in the spacetime around a black hole, the cones traced out by all the light rays that could be sent, in every possible direction, from various events. Your world line can't cross the light cones — that would mean travelling

The presence of even a small amount of matter falling into a black hole would alter the geometry inside the horizon — though this would probably only make the whole experience of being there even more violent. There was once considerable speculation about black holes forming various kinds of wormholes connected to other regions of space, but most relativists now consider this impossible. Everything that crosses the horizon will eventually be torn apart — crushed in two directions and stretched in the third, like the falling cloud of space junk we considered earlier — then the remnants will hit the **singularity** at $r=0$. General relativity predicts infinite spacetime curvature there, but the true nature of the singularity will depend on the details of quantum gravity, a discipline still in its infancy.

Further reading: *Spacetime Physics* by E.F. Taylor and J.A. Wheeler (W.H. Freeman, 1966) is an excellent introduction to special relativity. *Gravitation* by C.W. Misner, K.S. Thorne and J.A. Wheeler (W.H. Freeman, 1970) is the Bible of general relativity, with a detailed treatment of almost every aspect of the subject. *Black Holes and Timewarps: Einstein's Outrageous Legacy* by Kip Thorne (Macmillan, 1995) is a non-mathematical account of general relativity, with a wealth of fascinating biographical and historical detail on the subject's development.