Foreword

These notes cover some things from my relativity lectures that aren't covered in Griffiths. —T. Witten

motivation for relativity

Our work with Maxwell's equations led us to conclude that electric and magnetic waves can excite each other to create an electromagnetic wave propagating through empty space. Moreover, the speed of propagation is fixed by the equations: it is $(\epsilon_0\mu_0)^{-1/2} \simeq 3 \times 10^8$ meters per second. We denote this speed as c. It is paradoxical for an absolute speed to appear in this way, for the equations give us no reference frame with which to measure the speed. In ordinary mechanics such equations would be suspicious. We've learned that such things as speed and kinetic energy have no absolute meaning, but depend on how the observer is moving. And as Newton's Laws tell us, all observers moving at constant velocity with respect to one another observe the same laws of motion. As we developed Ampere's Law and Faraday's law, no mention was made of how the observer had to be moving. Yet when these laws are combined they give the wave equation, with its absolute speed c. By all we know from Newtonian mechanics, the speed of anything, including an electromagnetic wave, should depend on the motion of the observer. If one observer sees a wave moving at speed c to the right, another observer also moving to the right at speed u should measure the same wave moving at speed $c - u$. Yet Maxwell's equations give no way for this expected change of speed to happen. How can this paradox be resolved?

Two possible ways have arisen historically. The first was to follow the logic familiar from previously-known wave phenomena, such as sound waves in a liquid or solid. Here too, the laws of elasticity give rise to a wave equation with a definite speed of propagation. But here the waves represent a *disturbance of a medium.* Thus there is a natural frame of reference: the rest frame of the medium. When the wave equation is derived, it is clear that the motions are relative to the undisturbed medium. If one looks at electromagnetic waves in this context, it is natural to make the same interpretation: E and B must represent disturbances of some *medium.* Indeed, some people interpreted the fields in this way even before electromagnetic waves were conceived of. Accordingly, ingenious experiments were devised to sense the motion of this medium: the "luminiferous æther." But as we now know well, no such motion was detected.

A second possible resolution of the paradox was Einstein's bold re-conceptualization of space and time. Einstein asked whether our notions of space and time could be generalized in such a way that, as in ordinary Newtonian mechanics (1) all observers moving at constant velocity with respect to one another are equally valid and observe the same physical laws. yet there is a particular speed c which does not obey the familiar law of addition of velocities. Rather, (2) anything measured to move at speed c in by one observer is also measured to move at c by any other observer moving uniformly with respect to the first one. These two famous postulates of Einstein appear mutually contradictory. We can bring this contradiction to a head by considering a pair *events*, X and Y. An event denotes something that occupies a specific place and time. (As Griffiths says, a snapped finger is an event; an ocean voyage is not.) For convenience we'll suppose that the spatial separation between X and Y is along the x axis. In general a pair of events have some spatial separation Δx and some temporal separation Δt. The ratio $v \equiv \Delta x/\Delta t$ is the speed required to go from one event to the other. This speed is different when measured by a moving observer. If the observer is moving at a speed u in the x direction, then its origin moves during the time Δt and it measures a separation $\Delta x' = \Delta x - u\Delta t$. Thus it measures a speed $v' \equiv \Delta x'/\Delta t = (\Delta x - u\Delta t)/\Delta t = v - u$. The familiar law of addition of velocities emerges ineluctably, for any velocity including c.

The only way to avoid this outcome is to admit a radical possibility: the two observers must have different time co-ordinates $t' \neq t$ by virtue of their relative motion, just as they have different space co-ordinates, $x' \neq x$. In the sequel we will deduce the transformation law that gives the relation between the co-ordinates of an event for one observer and those for a moving observer. You have seen this "Lorentz transformation" discussed in previous classes. The text by Griffiths also discusses it along similar lines. These discussions generally infer the transformation by considering certain thought experiments involving light flashes. Often they assume part of the answer without discussion. And often the transformation law emerges only as a result of tedious and unmotivated algebra. I shall try to take a complementary approach. My aim is to show how the transformation law emerges inevitably from Einstein's two postulates stated above. I also want to reveal the deep similarity between Lorentz transformations and rotations.

Before doing this, I must lay the groundwork by discussing transformations that preserve length. In due course we shall see that Lorentz transformations are also a type of length-preserving transformation. The ideas below are developed in mathematical physics books. Jackson's electrodynamics book listed as a supplementary text for the course has some information about this subject. A more thorough treatment is given by D. E. Littlewood, *The theory of group characters and matrix representations of groups* 2d ed. (Oxford, Clarendon Press, 1950).

Transformations that preserve ordinary length

For this digression we will forget about relativity and the speed of light. To prepare for the Lorentz transformation to follow, we consider a simpler, drier subject: ordinary space. Any point in ordinary space can be represented by its co-ordinates x_1, x_2 and x_3 in some established co-ordinate system. In another co-ordinate system, the three co-ordinates of the given point are generally different: *viz.* x'_1, x'_2, x'_3. We use $\vec{\mathbf{x}}$ and $\vec{\mathbf{x}}'$ to denote the set of all three co-ordinates of our given spatial point in the two systems. In general, the primed co-ordinate system might or might have perpendicular axes. Equal intervals of x'_1 might or might not correspond to equal intervals of x_1. The two co-ordinate systems could be wildly different. Still, there has to be some function $\vec{\mathcal{R}}$ that gives the primed co-ordinates in terms of the unprimed ones: $\vec{\mathbf{x}}' = \vec{\mathcal{R}}(\vec{\mathbf{x}})$. This function $\vec{\mathcal{R}}$ is the *transformation* from the unprimed to the primed co-ordinates of points in space. Since all points must still have unique co-ordinates in the primed system, the transformation $\vec{\mathcal{R}}$ has to be one-to-one, and hence invertible. It is often convenient to consider a subclass of co-ordinate systems: those that preserve ordinary length or distance between two points (relative to the original co-ordinates). Any co-ordinate system that preserves lengths in this way should give the same physical laws. Accordingly, we now consider co-ordinate systems such that for every $\vec{\mathbf{x}}$, the corresponding $\vec{\mathbf{x}}'$ has the same length: $|\vec{\mathbf{x}}'|^2 = |\vec{\mathbf{x}}|^2$. That is,

$$|\vec{\mathbf{x}}'|^2 \equiv x_1'^2 + x_2'^2 + x_3'^2 = {x_1}^2 + {x_2}^2 + {x_3}^2 \equiv |\vec{\mathbf{x}}|^2$$

We also want the primed system to preserve distances between *pairs* of points: $|\vec{\mathbf{y}}' - \vec{\mathbf{x}}'|^2 = |\vec{\mathbf{y}} - \vec{\mathbf{x}}|^2$. From this fact it follows readily that any ordinary *dot product* must be the same in the primed and the unprimed co-ordinates:

$$|\vec{\mathbf{y}} - \vec{\mathbf{x}}|^2 = \sum_{i=1}^{3} (y_i - x_i)^2 = \sum_{i=1}^{3} x_i^2 + y_i^2 + 2x_i y_i = |\vec{\mathbf{x}}|^2 + |\vec{\mathbf{y}}|^2 + 2\vec{\mathbf{x}} \cdot \vec{\mathbf{y}}$$

Thus a dot product can be expressed in terms of lengths. Since all the lengths in the primed system are equal to their unprimed counterparts, the dot products must be equal as well:

$$2\vec{\mathbf{x}}' \cdot \vec{\mathbf{y}}' = |\vec{\mathbf{y}}' - \vec{\mathbf{x}}'|^2 - |\vec{\mathbf{x}}'|^2 - |\vec{\mathbf{y}}'|^2 = |\vec{\mathbf{y}} - \vec{\mathbf{x}}|^2 - |\vec{\mathbf{x}}|^2 - |\vec{\mathbf{y}}|^2 = 2\vec{\mathbf{x}} \cdot \vec{\mathbf{y}}$$

This restriction on the primed co-ordinates naturally imposes a corresponding restriction on the transformation $\vec{\mathcal{R}}$. Remarkably enough, length-preserving transformations must be *linear*: $\vec{\mathcal{R}}(\vec{\mathbf{x}} + \vec{\mathbf{y}}) = \vec{\mathcal{R}}(\vec{\mathbf{x}}) + \vec{\mathcal{R}}(\vec{\mathbf{y}})$ for any $\vec{\mathbf{x}}$ and $\vec{\mathbf{y}}$. We can see this by taking the dot product of this equation with an arbitrary other vector $\vec{\mathbf{u}}$. This $\vec{\mathbf{u}}$, whatever it is, must be the image of some other vector $\vec{\mathbf{v}}$: $\vec{\mathbf{u}} = \vec{\mathcal{R}}(\vec{\mathbf{v}})$. Now the equation of linearity can be written

$$\begin{aligned}
&\vec{\mathbf{u}} \cdot (\vec{\mathcal{R}}(\vec{\mathbf{x}} + \vec{\mathbf{y}}) - \vec{\mathcal{R}}(\vec{\mathbf{x}}) - \vec{\mathcal{R}}(\vec{\mathbf{y}})) \\
&\qquad = \vec{\mathcal{R}}(\vec{\mathbf{v}}) \cdot (\vec{\mathcal{R}}(\vec{\mathbf{x}} + \vec{\mathbf{y}}) - \vec{\mathcal{R}}(\vec{\mathbf{x}}) - \vec{\mathcal{R}}(\vec{\mathbf{y}})) \\
&\qquad = \vec{\mathcal{R}}(\vec{\mathbf{v}}) \cdot \vec{\mathcal{R}}(\vec{\mathbf{x}} + \vec{\mathbf{y}}) - \vec{\mathcal{R}}(\vec{\mathbf{v}}) \cdot \vec{\mathcal{R}}(\vec{\mathbf{x}}) - \vec{\mathcal{R}}(\vec{\mathbf{v}}) \cdot \vec{\mathcal{R}}(\vec{\mathbf{y}})) \\
&\qquad = \vec{\mathbf{v}} \cdot (\vec{\mathbf{x}} + \vec{\mathbf{y}}) - \vec{\mathbf{v}} \cdot \vec{\mathbf{x}} - \vec{\mathbf{v}} \cdot \vec{\mathbf{y}} \\
&\qquad = \vec{\mathbf{v}} \cdot (\vec{\mathbf{x}} + \vec{\mathbf{y}} - \vec{\mathbf{x}} - \vec{\mathbf{y}}) = 0.
\end{aligned}$$

Here we have used the fact that dot products are unchanged upon transformation. Basically the linearity of the transformation follows from the linearity of the dot product.

Since $\vec{\mathcal{R}}$ is linear, it can be readily written down in general. The most general linear transformation can be reduced to a set of coefficients r_{ij} and b_i:

$$x'_i = \sum_{j=1}^{3} r_{ij} x_j + b_i.$$

One subcase of linear transformation is the case where all the coefficients r_{ij} are zero. Then $x'_i = x_i + b_i$. The primed co-ordinate system is simply a translation of the unprimed one. Such translations do preserve length, but they aren't very interesting. Below we will concentrate on transformations that have $b_i = 0$.

The sums in the preceding equation can be simplified using matrix notation. The matrix for the transformation $\vec{\mathcal{R}}$ is called

$$\mathsf{R} \equiv \begin{bmatrix} r_{11} & r_{12} & r_{13} \\ r_{21} & r_{22} & r_{23} \\ r_{31} & r_{32} & r_{33} \end{bmatrix}.$$

Then the linear transformation law can be written in matrix form as

$$\begin{pmatrix} x'_1 \\ x'_2 \\ x'_3 \end{pmatrix} = \begin{bmatrix} r_{11} & r_{12} & r_{13} \\ r_{21} & r_{22} & r_{23} \\ r_{31} & r_{32} & r_{33} \end{bmatrix} \begin{pmatrix} x_1 \\ x_2 \\ x_3 \end{pmatrix}.$$

or equivalently

$$\vec{\mathbf{x}}' = \mathsf{R}\vec{\mathbf{x}}.$$

While any length-preserving transformation $\vec{\mathcal{R}}$ can be written in this matrix form, by no means every matrix R is length-preserving; only very restricted R's are length preserving. There is one straightforward approach to determining the allowable R's. We could simply impose the condition that $|\vec{\mathbf{x}}'|^2 = |\vec{\mathbf{x}}|^2$ and work through the algebra. But there is another approach that is easier and gives additional insight. This approach is to attack a more modest goal: to study transformations which hardly change $\vec{\mathbf{x}}$ at all: *infinitesimal* transformations.

infinitesimal length-preserving transformations If the vectors $\vec{\mathbf{x}}'$ are nearly equal to their $\vec{\mathbf{x}}$ counterparts, it means that R must be nearly the identity matrix

$$\mathsf{1} \equiv \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \quad \text{or} \quad \mathsf{1}_{ij} = \delta_{ij}.$$

Denoting the infinitesimal transformation as R, this amounts to saying

$$\mathsf{R}(\epsilon) \to \mathsf{1} + \epsilon \mathsf{L}$$

where the coefficient ϵ is arbitrarily small. Now finding the transformation amounts to finding the form of the *generator* L. Even for such tiny transformations, the condition of length preservation restricts the generator L. To find out what L has to look like, we consider the dot product of two arbitrary vectors: $\vec{\mathbf{x}}' \cdot \vec{\mathbf{y}}' = \vec{\mathbf{x}} \cdot \vec{\mathbf{y}}$. Writing this equality out in terms of L gives

$$[(\mathsf{1} + \epsilon \mathsf{L})\vec{\mathbf{x}}] \cdot (\mathsf{1} + \epsilon \mathsf{L})\vec{\mathbf{y}} = \vec{\mathbf{x}} \cdot \vec{\mathbf{y}}.$$

Multiplying out and simplifying, we find

$$\epsilon[\mathsf{L}\vec{\mathbf{x}}] \cdot \vec{\mathbf{y}} + \epsilon \vec{\mathbf{x}} \cdot \mathsf{L}\vec{\mathbf{y}} + \mathcal{O}(\epsilon^2) = 0.$$

Since ϵ is arbitrarily small, we may neglect the $\mathcal{O}(\epsilon^2)$ piece. We may simplify the first dot products by reversing its two factors to get

$$\vec{\mathbf{y}} \cdot \mathsf{L}\vec{\mathbf{x}} = -\vec{\mathbf{x}} \cdot \mathsf{L}\vec{\mathbf{y}}.$$

Taking $\vec{\mathbf{y}} = \hat{x}_i$ and $\vec{\mathbf{x}} = \hat{x}_j$, this product simplifies greatly to

$$\mathsf{L}_{ij} = -\mathsf{L}_{ji} :$$

The matrix L must be *antisymmetric*. We have found that any generator L of a length-preservating transformation has to be antisymmetric: that is, L has the restricted form

$$\mathsf{L} = \begin{bmatrix} 0 & L_{12} & L_{13} \\ -L_{12} & 0 & L_{23} \\ -L_{13} & -L_{23} & 0 \end{bmatrix}.$$

There are only three independent parameters, compared to nine for a full matrix. Evidently the antisymmetry condition also works in reverse: a generator preserves length if and only if it is antisymmetric.

finite transformations Using the infinitesimal transformations as leverage, we may now find out what finite length-preserving transformations are like. This will lead us to a physical interpretation of them. Suppose we replace ϵ in the above equations by a finite number θ. For concreteness, we'll also take a specific value for the generator L: $\mathsf{L} = \begin{bmatrix} 0 & 1 & 0 \\ -1 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}$. With this simple choice of L the transformation acts only on the 1 and 2 co-ordinates; it leaves the 3 co-ordinate alone. Though $\mathsf{1} + \theta \mathsf{L}$ is not length preserving, it is easy to create from it a length-preserving transformation via a limiting process. The trick is to break the transformation into n small pieces.

$$\mathsf{R}(\theta) = \lim_{n \to \infty} \left(\mathsf{1} + \frac{\theta}{n}\mathsf{L}\right)^n$$

We will take n to be large, so that each factor has a small coefficient θ/n. Thus the coefficient approaches zero and each factor becomes an infinitesimal transformation. To compensate for the smallness of each factor, we increase the number of factors in proportion to n. The form of this $\mathsf{R}(\theta)$ can be simplified greatly. First, it is clear that R leaves the 3 axis alone, since each of the infinitesimal factors do. We're thus free to focus on the $1-2$ plane. This product $(\mathsf{1}+(\theta/n)\mathsf{L})^n$ can be simplified by multiplying it out using the binomial theorem. This theorem applies to any objects that can be multiplied and added, including matrices:

$$\left(\mathsf{1} + \frac{\theta}{n}\mathsf{L}\right)^n = \sum_{k=0}^{n} \binom{n}{k} \left(\frac{\theta}{n}\right)^k \mathsf{L}^k.$$

Here $\binom{n}{k}$ means the binomial coefficient $n!/(k!(n-k)!)$. The only case we need for our limit is the case where $n \gg k$. In this limit, the Binomial theorem simplifies:

$$\lim_{n \to \infty} \left(\mathsf{1} + \frac{\theta}{n}\mathsf{L}\right)^n = \sum_{k=0}^{\infty} \frac{1}{k!} \theta^k \mathsf{L}^k.$$

This series has a familiar look. Except for the L^k factors, it is just the power series for $\exp\theta$. To discover the effect of the L factors we use a powerful property of all finite matrices: Any d-dimensional matrix raised to the d power can be expressed as a sum of lower powers of the matrix. Our L matrix is two dimensional. Thus L^2 has to be expressible in terms of L and $\mathsf{1}$. To find this simplification for our L we have but to multiply LL. We find the simple result $\mathsf{LL} = -\mathsf{1}$! Using this identity we may replace every factor of L^2 by $\mathsf{1}$ in our series above. This simplifies the sum greatly. Whenever k is even, the L^{2k} is simply $(-1)^k\mathsf{1}$. Whenever k is odd, we may collapse all but one factor of L to get $\mathsf{L}^{2k+1} = (-1)^k\mathsf{L}$. Now the series can be simplified greatly:

$$\mathsf{R}(\theta) = \sum_k \frac{1}{(2k)!}(-1)^k\theta^{2k}\mathsf{1} + \sum_k \frac{1}{(2k+1)!}(-1)^k\theta^{2k+1}\mathsf{L}.$$

The first sum contains all the even-power terms from the series before; the second sum contains the odd terms. Now the power series are purely numerical. They are merely some function of θ. Indeed, the two series are familiar functions: the first is $\cos\theta$; the second is $\sin\theta$. We have thus arrived at the explicit form of R. Re-introducing the third dimension, we have

$$\mathsf{R}(\theta) = \begin{bmatrix} \cos\theta & \sin\theta & 0 \\ -\sin\theta & \cos\theta & 0 \\ 0 & 0 & 1 \end{bmatrix}$$

Now the meaning of $\mathsf{R}(\theta)$ is clear and familiar. R is simply a *rotation* about the 3 axis by an angle θ. We would have obtained similar results for the other simple generators. It is convenient to denote this generator that had only a 1-2 and 2-1 unit elements as L_3. In similar fashion, we can define the two other similar generators L_1 and L_2. Not surprisingly, the L_1 generates rotations about the 1 axis, and L_2 generates rotations about the 2 axis. Taking these three generators together, we can form the vector matrix $\vec{\mathsf{L}}$ as the vector having the elements $\mathsf{L}_1, \mathsf{L}_2$, and L_3. Any generator can of course be expressed as a linear combination of these three. That is, we can always find a vector $\vec{\omega}$ such that the desired generator is $\vec{\omega}\cdot\vec{\mathsf{L}}$. It can be shown that this general generator produces a rotation along the $\hat{\omega}$ axis with angle $|\vec{\omega}|$. Our search for the nature of length-preserving transformations has led to a simple answer. We've found that all rotations are length-preserving. These can all be generated as a succession of identical infinitesimal transformations, as shown above.

Does *any* length-preserving transformation amount to a rotation, of the form $\lim_{n\to\infty}(1 + \vec{(\theta/n)}\mathsf{L})^n$ for some antisymmetric L? In some cases it clearly does. For example suppose we make two successive transformations using the same generator L, but with different angles θ_1 and θ_2. The result can be written

$$\mathsf{R}(\theta_2)\mathsf{R}(\theta_1) = \lim_{n\to\infty}\left(\mathsf{1} + \frac{\theta_2}{n}\mathsf{L}\right)^n\left(\mathsf{1} + \frac{\theta_1}{n}\mathsf{L}\right)^n.$$

Since the L's are all identical, we may re-arrange the factors to form

$$\mathsf{R}(\theta_2)\mathsf{R}(\theta_1) = \lim_{n\to\infty}\left[\left(\mathsf{1} + \frac{\theta_2}{n}\mathsf{L}\right)\left(\mathsf{1} + \frac{\theta_1}{n}\mathsf{L}\right)\right]^n.$$

We may multiply these two infinitesimal factors to get

$$\mathsf{R}(\theta_2)\mathsf{R}(\theta_1) = \lim_{n\to\infty}\left[\mathsf{1} + \frac{\theta_2 + \theta_1}{n}\mathsf{L}\right]^n.$$

Here we have neglected the term $\theta_1\theta_2/n^2$, because it becomes negligible in the large-n limit. The quantity on the right is simply $\mathsf{R}(\theta_1 + \theta_2)$. Thus we have shown that when two successive transformations are done using the same generator, the result is again a transformation using that generator, and the coefficient or "magnitude" of the combined transformation is the sum of the magnitudes of the constituent transformations.

But not all length preserving transformations have this form. A major counterexample is the parity transformation, which inverts all the axes. There is no way to build this transformation out of infinitessimal pieces.

Let us then consider length transformation that *can* be made from infinitessimal pieces, but not identical ones. We may call these continuous transformations. An example would be

$$\mathsf{R}_1(\alpha)\mathsf{R}_2(\beta) = \lim_{n\to\infty}(1 + (\alpha/n)\mathsf{L}_1)^n(1 + (\beta/n)\mathsf{L}_2)^n$$

Can *this* product can be built from a single generator $\vec{\omega}\cdot\vec{\mathsf{L}}$? If so, how does $\vec{\omega}$ depend on the given α and β? From our intuitive knowledge of rotations such a $\vec{\omega}$ can be found: these two rotations amount to a single rotation about some axis or other. Indeed any product of continuous, length preserving transformations can be built from a single generator. It can be done in spaces of arbitrary dimension. But to find the right combination of generators, *i.e.,* $\vec{\omega}$ is not an easy matter. For example, $\vec{\omega}$ changes if the two rotations are done in the reverse order. All this is the subject of books on the rotation group.

Co-ordinates of equivalent observers

The above ideas about length-preserving transformations provide a way of deducing how the co-ordinates of two mutually moving observers have to be related in order to satisfy Einstein's two postulates. The reason is that the invariance of the speed of light can be expressed as a kind of distance-preserving condition. Using this condition, we can deduce the nature of the Lorentz transformation by following the reasoning above for ordinary length-preserving transformations.

invariance of the speed c Before launching into our reasoning, we set up a convenient co-ordinate system. Since we wish to describe events, we must define time and space co-ordinates in each co-ordinate system. It's convenient to define a co-ordinate $x_0 \equiv ct$ in addition to the ordinary x_1, x_2, x_3. (To give you a feeling for the scale, if t is a nanosecond, x_0 is about a foot.) We define the set of all four co-ordinates x_0 ... x_3 of an event by a four-dimensional vector denoted $\underline{X}$. If two events are separated by spatial co-ordinates $\Delta\vec{\mathbf{x}}$ and Δx_0, a speed of $v = |\Delta\vec{\mathbf{x}}|/|\Delta x_0|$ is needed to go from one to the other. The speed v' as viewed in the primed co-ordinates has the same form, with Δx replaced by $\Delta x'$ everywhere. The postulate tells us that whenever $v = 1$, v' must equal 1 as well. This condition can be expressed in terms of a generalized "length", called the Lorentz length, or the Lorentz distance. We denote the Lorentz length of an event $\underline{X}$ by $||\underline{X}||$:

$$||\underline{X}||^2 \equiv -x_0^2 + |\vec{\mathbf{x}}|^2.$$

This Lorentz length-squared resembles ordinary length-squared in some ways, but not in others. Like ordinary length-squared the Lorentz length is a quadratic expression in the components that gives a scalar number for each vector in the space considered. Unlike ordinary length-squared, it can be negative as well as positive. Further, a vector can have a vanishing Lorentz length and still not be the zero vector. Indeed, these "null vectors" are precisely the ones that have $|\vec{\mathbf{x}}|^2/x_0^2 = 1$: there are lots of null vectors but only one zero vector. To reach such points from the origin requires a speed c. Thus whenever two events are separated by a speed c, the Lorentz distance between them is zero, and vice versa. Since the speed is c in all co-ordinate systems if it is c in any, we conclude that the Lorentz distance between two events is 0 in all if it is zero in any. These zero lengths are invariant under different co-ordinate systems.

Lorentz length is fixed for equivalent observers Remarkably, the Lorentz distance is invariant not only for null vectors, but for any pairs of events, null or not. To see this requires some further reasoning. (I'm grateful to Bob Geroch for providing the main ideas.) To help with this reasoning, we note that there is a natural counterpart of a dot product associated with the Lorentz length. It appears when we write out the Lorentz length of the sum of two events $\underline{X}$ and $\underline{Y}$:

$$||\underline{X} + \underline{Y}||^2 = ||\underline{X}||^2 + ||\underline{Y}||^2 + 2(-x_0 y_0 + 2\vec{\mathbf{x}} \cdot \vec{\mathbf{y}}).$$

The term in parentheses is the counterpart of a dot product. We may call it the Lorentz dot product. We can express it in four-vector notation by defining a special matrix

$$\mathsf{g} \equiv \begin{bmatrix} -1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}.$$

Then the equation above becomes

$$||\underline{X} + \underline{Y}||^2 = ||\underline{X}||^2 + ||\underline{Y}||^2 + 2(\underline{X} \cdot \mathsf{g}\underline{Y}).$$

(Since g is diagonal, $\underline{X} \cdot \mathsf{g}\underline{Y} = \underline{Y} \cdot \mathsf{g}\underline{X}$. The $\cdot$ has the usual definition: $\underline{X} \cdot \underline{Y} \equiv \sum_0^4 x_i y_i$.) We want to show that the Lorentz length of any event vector $\underline{X}$ is the same for all observers. Clearly we must do this by relating the arbitrary vector $\underline{X}$ to some null vector.

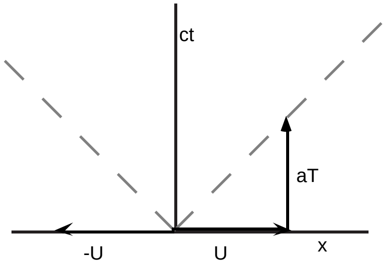

For this purpose it is useful to define the unit vector along the x_0 direction, $\underline{T}$. First we consider an event with only spatial components. We call it $\underline{U}$. For any $\underline{U}$ one can always make a null vector by adding some multiple of the pure-time vector $\underline{T}$, as shown in the figure, with $t_1 = t_2 = t_3 = 0$. That means we can find a coefficient a such that $||a\underline{T} + \underline{U}|| = 0$. We can expand out this Lorentz length to relate it to the lengths of $a\underline{T}$ and $\underline{U}$:

$$0 = ||a\underline{T} + \underline{U}||^2 = ||a\underline{T}||^2 + ||\underline{U}||^2 + 2a\underline{T} \cdot \mathsf{g}\underline{U} = a^2\,||\underline{T}||^2 + ||\underline{U}||^2$$

The cross term vanished because $\underline{T}$ has only a 0 componet while $\underline{U}$ was constructed to have none. By the same reasoning we see that $a\underline{T} - \underline{U}$ must also be null.

Now we consider the transformed four-vectors $(a\underline{T})'$ and $\underline{U}'$. There is no knowing how the primed co-ordinates have been transformed. $\underline{U}'$ may have a nonzero time component, and $(a\underline{T})'$ may have spatial ones. Still, the sum $(a\underline{T})' + \underline{U}'$ must be null, since it was null in the unprimed co-ordinates. We can consider a to be an integer, by just taking the unit of time to be very small. Then $a\underline{T}$ is the sum of a unit vectors. The image of the vector $a\underline{T} + \underline{U}$ in the primed frame is obtained by transforming $\underline{T}$ and $\underline{U}$ and adding the results. The resulting vector $a\underline{T}' + \underline{U}'$ must be null, since it was in the unprimed frame.

$$0 = \left|\left|a\underline{T}' + \underline{U}'\right|\right|^2 = a^2\left|\left|\underline{T}'\right|\right|^2 + \left|\left|\underline{U}'\right|\right|^2 + 2a\underline{T}' \cdot \mathsf{g}\underline{U}' \qquad (*)$$

We may obtain an analogous result starting from our other null vector $a\underline{T} - \underline{U}$:

$$0 = \left|\left|a\underline{T}' - \underline{U}'\right|\right|^2 = a^2\left|\left|\underline{T}'\right|\right|^2 + \left|\left|\underline{U}'\right|\right|^2 - 2a\underline{T}' \cdot \mathsf{g}\underline{U}'$$

Adding these two equations yields

$$0 = a^2\left|\left|\underline{T}'\right|\right|^2 + \left|\left|\underline{U}'\right|\right|^2 \qquad (**)$$

This gives us information about the relative magnitudes of $\left|\left|\underline{T}'\right|\right|$ and $\left|\left|\underline{U}'\right|\right|$. Recognizing that a^2 be written as $-\left|\left|\underline{U}\right|\right|^2 / \left|\left|\underline{T}\right|\right|^2$, we find

$$\frac{\left|\left|\underline{T}'\right|\right|^2}{\left|\left|\underline{T}\right|\right|^2} = \frac{\left|\left|\underline{U}'\right|\right|^2}{\left|\left|\underline{U}\right|\right|^2}$$

The ratio of $\underline{T}$ on the left is independent of the U chosen. Thus we see that the Lorentz squared-length of all $\underline{U}$ with no time component grows by the same factor when we transform to the primed co-ordinates. We call this factor λ.

Another important fact emerges from these equations: the transformation preserves Lorentz-orthogonality: $a\underline{T}' \cdot \underline{U}' = 0$. This follows from comparing Eqs. * and **.

Now it is not surprizing that the squared Lorentz length of *any* vector $\underline{X}$ must grow by the same factor λ when transformed into the primed co-ordinates. This is because any $\underline{X}$ can be expressed as a combination of our $\underline{T}$ and some $\underline{U}$ with no 0 component. For some coefficient b $\underline{X} = b\underline{T} + \underline{U}$. To find its squared length in the primed frame, we use

$$\left|\left|\underline{X}'\right|\right|^2 = \left|\left|b\underline{T}' + \underline{U}'\right|\right|^2 = b^2\left|\left|\underline{T}'\right|\right|^2 + \left|\left|\underline{U}'\right|\right|^2 + 2b\underline{T}' \cdot \mathsf{g}\underline{U}'$$

The dot product vanishes, as we've just noted. Meanwhile, the $\left|\left|\underline{T}'\right|\right|^2$ and $\left|\left|\underline{U}'\right|\right|^2$ are just λ times their unprimed counterparts. Inserting these facts and recombining, we find

$$\left|\left|\underline{X}'\right|\right|^2 = \lambda\left|\left|b\underline{T} + \underline{U}\right|\right|^2 = \lambda\left|\left|\underline{X}\right|\right|^2$$

In general this factor λ may depend on the relative motion of the two co-ordinate systems: It can depend on their relative speed v. But it does not depend on the vector chosen. Now, the primed co-ordinates are as valid as the unprimed ones. Thus, we could have started from the primed frame and worked through all the above reasoning, concluding that

$$\left|\left|\underline{X}\right|\right|^2 = \lambda\left|\left|\underline{X}'\right|\right|^2$$

The λ for this reverse transformation cannot be different from what it was in the forward direction, since there is no objective basis for saying, *e.g.* , which λ should be bigger. The only value of λ that can work in both directions is $\lambda = 1$. (We discard $\lambda = -1$, since λ must be 1 in the limit where the relative speed $v \to 0$. This argument applies to vectors $\underline{U}$ that have no time component. But it is readily generalized to all four vectors. This shows that the Lorentz distance between all any pair of points is the same in all equivalent co-ordinate systems, moving at any relative speed.

linearity Now that we have shown that the the Lorentz length is preserved for all equivalent co-ordinate systems, we can copy our reasoning for ordinary length-preserving transformations and obtain the analogous results: i) Since it preserves the Lorentz length of a given vector, the transformation must also preserve the Lorentz dot product.

$$2\underline{X}' \cdot \mathsf{g}\underline{Y}' = \left|\left|\underline{Y}' - \underline{X}'\right|\right|^2 - \left|\left|\underline{X}'\right|\right|^2 - \left|\left|\underline{Y}'\right|\right|^2 = \left|\left|\underline{Y} - \underline{X}\right|\right|^2 - \left|\left|\underline{X}\right|\right|^2 - \left|\left|\underline{Y}\right|\right|^2 = 2\underline{X} \cdot \mathsf{g}\underline{Y}$$

Further, ii) the transformation $\underline{X}' = \mathcal{R}(\underline{X})$ has to be linear. $\mathcal{R}(\underline{X} + \underline{Y}) = \mathcal{R}(\underline{X}) + \mathcal{R}(\underline{Y})$ for any $\underline{X}$ and $\underline{Y}$. We can see this by taking the Lorentz dot product of this equation with an arbitrary other vector $\underline{U}$. This $\underline{U}$, whatever it is, must be the image of some other vector $\underline{V}$: $\underline{U} = \mathcal{R}(\underline{V})$. Now the equation of linearity can be written

$$\begin{aligned}
&\underline{U} \cdot \mathsf{g}(\mathcal{R}(\underline{X} + \underline{Y}) - \mathcal{R}(\underline{X}) - \mathcal{R}(\underline{Y})) \\
&\qquad = \mathcal{R}(\underline{V}) \cdot \mathsf{g}(\mathcal{R}(\underline{X} + \underline{Y}) - \mathcal{R}(\underline{X}) - \mathcal{R}(\underline{Y})) \\
&\qquad = \mathcal{R}(\underline{V}) \cdot \mathsf{g}\mathcal{R}(\underline{X} + \underline{Y}) - \mathcal{R}(\underline{V}) \cdot \mathsf{g}\mathcal{R}(\underline{X}) - \mathcal{R}(\underline{V}) \cdot \mathsf{g}\mathcal{R}(\underline{Y})) \\
&\qquad = \underline{V} \cdot \mathsf{g}(\underline{X} + \underline{Y}) - \underline{V} \cdot \mathsf{g}\underline{X} - \underline{V} \cdot \mathsf{g}\underline{Y} \\
&\qquad = \underline{V} \cdot \mathsf{g}(\underline{X} + \underline{Y} - \underline{X} - \underline{Y}) = 0.
\end{aligned}$$

so that except for a possible translation, the transformation can be expressed as a matrix R.

infinitesimal Lorentz-length preserving transformations As with ordinary length-preserving transformations, it simplifies life to consider transformations that do almost nothing: infinitesimal transformations of the form

$$\mathsf{R}(\epsilon) = (\mathsf{1} + \epsilon\mathsf{M})$$

We may use the invariance of the Lorentz dot product $\underline{Y} \cdot \mathsf{g}\underline{X}$ to write

$$0 = (\mathsf{M}\underline{Y}) \cdot \mathsf{g}\underline{X} + \underline{Y} \cdot \mathsf{g}\mathsf{M}\underline{X}$$

To simplify this, we use the fact the g is diagonal to change the left side to

$$(\mathsf{g}\mathsf{M}\underline{Y}) \cdot \underline{X} = \underline{X} \cdot \mathsf{g}\mathsf{M}\underline{Y}.$$

We have also interchanged the two dotted factors in the last equation. Combining with the previous equation, we now have

$$\underline{X} \cdot \mathsf{g}\mathsf{M}\underline{Y} = -\underline{Y} \cdot \mathsf{g}\mathsf{M}\underline{X}.$$

This means that the matrix $\mathsf{g}\mathsf{M}$ has to be antisymmetric. This can be verified directly by taking $\underline{X}$ and $\underline{Y}$ to be unit vectors along any co-ordinate axes. From this condition we readily get the form of the generator M. We note that $(\mathsf{g}\mathsf{M})_{ab} = g_{aa}M_{ab}$, since g is diagonal. Our condition $(\mathsf{g}\mathsf{M})_{ij} = -(\mathsf{g}\mathsf{M})_{ji}$ becomes for $a = 0$, $-M_{0j} = -M_{j0}$, or $M_{0j} = M_{j0}$. For the case where both i and j not equal to zero, the $\mathsf{g}_{ii} = 1$, so that $M_{ij} = -M_{ji}$. (We're often led to discuss spatial indices separately from generic four-component indices. For this purpose we'll agree to let the subscripts $_{ijklmn}$ represent spatial indices that run from 1 to 3. If we want to indicate generic components that run from 0 to 3, we'll use use other letters, viz $_{abcdefgh\ opqrstuvwxyz}$.)

Finite Lorentz transformation As before, we may now use the generators to make finite transformations:

$$\mathsf{R}(\zeta) = \lim_{n\to\infty}\left(\mathsf{1} + \frac{\zeta}{n}\mathsf{M}\right)^n$$

The case where M_{0i} $(= M_{i0}) = 0$ generates the now-familiar rotations. But when they are not zero, the M_{0j} generate something new. We consider the simplest case, where all the elements of M are 0 except for $M_{0,1}$ and $M_{1,0}$. Then the transformation R leaves the 2 and 3 co-ordinates alone, and we can thus focus on the 0-1 subspace. Once again, we can write the transformation as

$$\mathsf{R}(\zeta) = \sum_{k=0}^{\infty} \frac{\zeta^k}{k!}\mathsf{M}^k$$

Once again we can simplify by expressing higher powers of M in terms of lower powers. Squaring M explicitly, we find $\mathsf{M}^2 = +\mathsf{1}$. This compares with $\mathsf{L}^2 = -\mathsf{1}$ for the ordinary rotations. This sign makes a big difference in the series expansion. The sines and cosines appearing in rotations now become hyperbolic sines and cosines. Thus the transformation law is

$$\mathsf{R}(\zeta) = \begin{bmatrix} \cosh\zeta & \sinh\zeta & 0 & 0 \\ \sinh\zeta & \cosh\zeta & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

interpretation of ζ Clearly the parameter ζ plays a role analogous to the rotation angle. It is called the rapidity. To find the physical meaning of rapidity, we consider the unprimed image of the primed spatial origin: $x'_1 = x'_2 = x'_3 = 0$. Looking at the equation $\mathsf{R}\underline{X} = \underline{X}'$, we find the condition

$$0 = \cosh\zeta x_1 + \sinh\zeta x_0$$

or $x_1/x_0 = -\tanh\zeta$. This ratio is simply the speed of the primed origin measured in the unprimed co-ordinates, in units of c. Clearly our transformation is a transformation to a moving co-ordinate system moving at speed $c\ \tanh\zeta$ in the $-x_1$ direction. Such a transformation is called a boost. We often call this dimensionless speed $\beta \equiv \tanh\zeta$. Like a rotation angle, the rapidity is an additive quantity. Two successive boosts with the same M and magnitudes ζ_1 and ζ_2 amount to a single boost with magnitude $\zeta_1 + \zeta_2$. The reason is just as explained above for rotations. Note that the relative velocity β is *not* additive in this way. Using trig identities, one can readily express $\cosh\zeta$ in terms of β: $\cosh\zeta = (1-\beta^2)^{-1/2}$. It is conventional to denote this $\cosh\zeta$ as γ This leads us to the familiar form of the Lorentz transformation:

$$\mathsf{R}(\beta) = \begin{bmatrix} \gamma & \gamma\beta & 0 & 0 \\ \gamma\beta & \gamma & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

To make a boost in a general direction, one combines the three time-containing generators: For a boost in the $\hat{n}$ direction, the generator is

$$\mathsf{M} = \begin{bmatrix} 0 & n_1 & n_2 & n_3 \\ n_1 & 0 & 0 & 0 \\ n_2 & 0 & 0 & 0 \\ n_3 & 0 & 0 & 0 \end{bmatrix}$$