CS205b/CME306

Lecture 8

1 Finite Element Method

This section uses a significant amount of text from Teran et al., Finite Volume Methods for the Simulation of Skeletal Muscle, 2003, with permission of the author.

1.1 Constitutive Modelling (Stress-Strain Relationships)

We have geometric measures of deformation, and we would like to have a formula that expresses the stress of the system in terms of this deformation. Finding these relationships is the heart of issue of modelling in mechanical engineering. Constitutive models may start with one of any number of strain measures and end with any one of the forms of stress. It is important to note that all of the stress formulations are all interchangeable, but the different measures of strain are not interchangeable, and the different constitutive models are not interchangeable. A constitutive model is basically some way of writing down a stress based on measurements of the deformation of a material. In general, a constitutive model may make any measurements of the deformation it needs and perform any computations necessary based on those measurements to yield the stress.

One such constitutive model is obtained by generalizing Hooke's law to higher dimensions, in which case it states that there is a linear relationship between Cauchy stress $\boldsymbol{\sigma}_{ij}$ and Cauchy strain $\boldsymbol{\epsilon}_{km}$ (a linearization of the Green strain that is also linear in the actual deformation). This expresses a linear relationship between deformation and stress and is also referred to linear elasticity. The most general possible linear relationship between stress and strain is given by

$$\boldsymbol{\sigma}_{ij} = \sum_{km} C_{ijkm} \boldsymbol{\epsilon}_{km}$$

where C_{ijkm} contains $3 \times 3 \times 3 \times 3 = 81$ entries. Note that most of these entries are not actually independent. Since $\boldsymbol{\sigma}_{ij}$ is symmetric, the quantities C_{ijkm} and C_{jikm} are not independent. Since $\boldsymbol{\epsilon}_{km}$ is symmetric, the quantities C_{ijkm} and C_{ijmk} are not independent. These observations bring the number of independent entries down to 36.

In the special case of an isotropic material (a material whose properties do not depend on direction), only two independent entries remain. This relationship is often written as

$$\boldsymbol{\sigma} = 2\mu\boldsymbol{\epsilon} + \lambda \mathrm{tr}(\boldsymbol{\epsilon})\mathbf{I}$$

where $\mathrm{tr}(\boldsymbol{\epsilon})$ is the sum of the diagonal elements of $\boldsymbol{\epsilon}$, and μ and λ are Lamé's constants. These can be rewritten in terms of the more useful parameters E (Young's modulus) and ν (Poisson's ratio) as

$$\boldsymbol{\sigma} = \frac{E}{1+\nu}\boldsymbol{\epsilon} + \frac{\nu E}{(1+\nu)(1-2\nu)}\mathrm{tr}(\boldsymbol{\epsilon})\mathbf{I}.$$

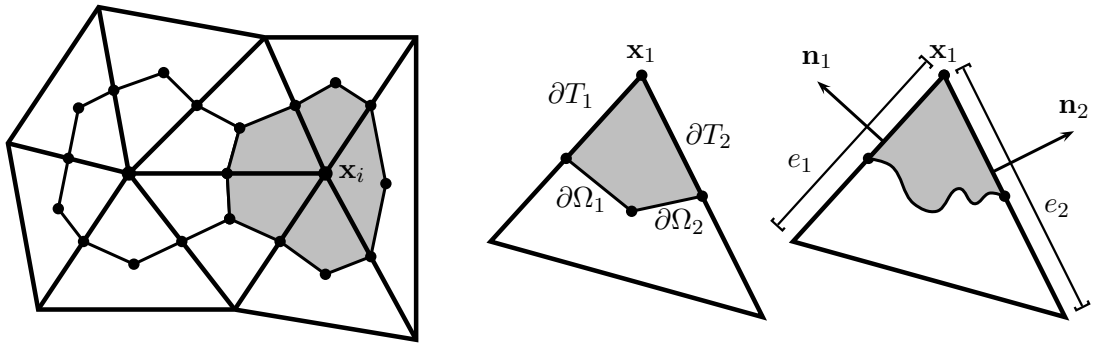

(a) FVM integration regions.

(b) Integration over a triangle.

Figure 1: Region around a node.

If the material were in the shape of a long rod, a stretching force on the rod would cause it to elongate; the ratio of the stress applied and the strain (relative change in length) the rod undergoes is called its Young's modulus. As the rod elongates, it also becomes thinner; the ratio of the relative change in length and the relative change in diameter of the rod is the Poisson's ratio, negated. The Poisson's ratio is a value in the range $(-1, 1/2)$, which can be observed from the denominator of the expression above. The upper limit of $1/2$ corresponding to perfectly incompressible materials (rubber is close). Cork has a Poisson's ratio of roughly zero. Materials with negative values are rare.

It is interesting to note that the assumption of isotropy is very strong. The most general linear and isotropic function of $\mathbf{A}$ is $\alpha \mathrm{tr}(\mathbf{A}) + \beta\mathbf{A} + \gamma\mathbf{A}^T$ for constants α, β, and γ. The original form for the isotropic case is obtained by using $\boldsymbol{\epsilon} = \boldsymbol{\epsilon}^T$ and then defining $\lambda = \alpha$ and $2\mu = \beta + \gamma$.

Sometimes a different strain measure (e.g., Green strain) is chosen instead of the Cauchy strain, and the stress is assumed to be linear in that strain. Other constitutive models, such as Neo-Hookean, assume nonlinear relationships.

1.2 Finite Element Method (FEM)

Now, we have some measure of what the forces are on a control volume, but we need a numerical discretization to evolve the system forward in time. Going back to our chunk of mass, we need to figure out what the forces are around that point, since we have $\mathbf{x}_i'' = \mathbf{v}_i' = \mathbf{f}_i/m_i$. We do this by integrating the force over the control volume around the point.

FEM provides a simple and geometrically intuitive way of integrating the equations of motion, with an interpretation that rivals the simplicity of mass-spring systems. However, unlike masses and springs, an arbitrary constitutive model can be incorporated into FEM. In the deformed configuration, consider dividing up the continuum into a number of discrete regions each surrounding a particular node. Figure 1(a) depicts two nodes each surrounded by a region. Suppose that we wish to determine the force on the node $\mathbf{x}_i$ surrounded by the region Ω. Ignoring body forces for brevity, the force can be calculated as

$$\mathbf{f} = \frac{D}{Dt}\int_\Omega \rho\mathbf{v}\, d\mathbf{x} = \oint_{\partial\Omega} \mathbf{t}\, dS = \oint_{\partial\Omega} \boldsymbol{\sigma}\mathbf{n}\, dS$$

where ρ is the density, $\mathbf{v}$ is the velocity, and $\mathbf{t}$ is the surface traction on $\partial\Omega$ (the force applied over the area). The last equality comes from the definition of the Cauchy stress on $\mathbf{t} = \boldsymbol{\sigma}\mathbf{n}$.

Evaluation of the boundary integral requires integrating over the two segments interior to each incident triangle. Figure 1(b) depicts one of these incident triangles along with interior segments labeled $\partial\Omega_1$ and $\partial\Omega_2$. Since $\boldsymbol{\sigma}$ is constant in each triangle and the integral of the local unit normal over any closed region is identically zero (from the divergence theorem), we have

$$\oint_{\partial\Omega_1} \boldsymbol{\sigma}\mathbf{n}\, dS + \oint_{\partial\Omega_2} \boldsymbol{\sigma}\mathbf{n}\, dS + \oint_{\partial T_1} \boldsymbol{\sigma}\mathbf{n}\, dS + \oint_{\partial T_2} \boldsymbol{\sigma}\mathbf{n}\, dS = 0$$

where ∂T_1 and ∂T_2 are depicted in the figure. Note that the integral of $\boldsymbol{\sigma}\mathbf{n}$ over $\partial\Omega_1$ and $\partial\Omega_2$ can be replaced by the integral of $-\boldsymbol{\sigma}\mathbf{n}$ over ∂T_1 and ∂T_2. That is, for each triangle, we can integrate over the portions of its edges incident to $\mathbf{x}_i$ instead of the two interior edges $\partial\Omega_1$ and $\partial\Omega_2$. Moreover, even if $\partial\Omega_1$ and $\partial\Omega_2$ are replaced by an arbitrary path inside the triangle, see Figure 1(b), we can replace the integral over this region with the integral over ∂T_1 and ∂T_2.

We choose an arbitrary path inside the triangles that connects the midpoints of the two edges incident on $\mathbf{x}_i$. Then the surface integrals are simply equal to $-\boldsymbol{\sigma}\mathbf{n}_1 e_1/2$ and $-\boldsymbol{\sigma}\mathbf{n}_2 e_2/2$ where e_1 and e_2 are the edge lengths of the triangles. Thus, the force contribution to node $\mathbf{x}_i$ from this triangle is

$$\mathbf{f} = -\frac{1}{2}\boldsymbol{\sigma}(e_1\mathbf{n}_1 + e_2\mathbf{n}_2).$$

In three spatial dimensions, we divide the volume of a tetrahedron among its four corners in such a way that the face areas are divided equally among their three corners. This is analogous to insisting in 2D that edges of the triangle be cut through their midpoints. Extending the above argument from 2D to 3D, we are able to express the force contribution from a tetrahedron to its vertices as

$$\mathbf{f} = -\frac{1}{3}\boldsymbol{\sigma}(a_1\mathbf{n}_1 + a_2\mathbf{n}_2 + a_3\mathbf{n}_3),$$

where a_1, a_2, and a_3 are the areas of three faces surrounding the node $\mathbf{x}_i$ under consideration, and $\mathbf{n}_1$, $\mathbf{n}_2$, and $\mathbf{n}_3$ are the faces' normals.

Given an arbitrary stress $\boldsymbol{\sigma}$, regardless of the method in which it was obtained, we obtain a force on the nodes in the following fashion. For each face j, compute the traction as $\boldsymbol{\sigma}\mathbf{n}_i$, and multiply by the face area a_i to obtain the total force for the face $\boldsymbol{\sigma}a_i\mathbf{n}_i$. Then, distribute one third of this force to each of the face's three vertices. Note that the forces applied to the four faces sum to zero, so only three need be computed directly.

If we use the Second Piola-Kirchhoff stress $\mathbf{S}$, we can convert it into a First Piola-Kirchhoff stress $\mathbf{P} = \mathbf{FS}$. From the First Piola-Kirchhoff stress $\mathbf{P}$, we can compute the total force contribution of a tetrahedron to a node from the areas A_i and normals N_i in material coordinates as

$$\mathbf{f} = -\frac{1}{3}\mathbf{P}(A_1\mathbf{N}_1 + A_2\mathbf{N}_2 + A_3\mathbf{N}_3)$$

Since the A_i and $\mathbf{N}_i$ do not change during the computation, we can precompute and store these quantities. Then the force contribution to each node can be computed as $\mathbf{f}_i = \mathbf{P}\mathbf{b}_i$, where the $\mathbf{b}_i$ are precomputed. As before, only three of these need to be computed directly since total force on the four faces sums to zero.

2 Rigid Bodies

At this point, we have a copy of our undeformed object in material coordinates and another copy that is moving and deforming in spatial coordinates. If this objects gets stiffer and doesn't vibrate very much, then in spatial coordinates, the object just translates and rotates. If we examined the body through, for example, its Green strain, nothing would happen, since $\mathbf{F}$ would be a rotation matrix, and $\mathbf{G} = (\mathbf{F}^T\mathbf{F} - \mathbf{I})/2 = (\mathbf{I} - \mathbf{I})/2 = 0$. Thus, for rigid bodies, there is no strain. For a normal object, we have a degree of freedom for every particle, but for a rigid body we can use fewer degrees of freedom. Since the only thing a rigid body is able to do is translate and rotate through space, a rigid body can be characterized by just those six (three in 2D) degrees of freedom.

2.1 Kinematic Description

As with materials that are able to deform and move through space, it is convenient to keep a copy of rigid bodies in their own material space. However, in the case of rigid bodies, this material space defines a coordinate system for the entire rigid body.

At this point, we introduce a bit of new terminology. We typically refer to the coordinate system of the initial configuration (or some reference configuration) as material coordinates. We will often refer to this as object space. Similarly, by world space we will mean the spatial coordinate system. Since the rigid body is fixed in object space, it is convenient to take the rigid body's center of mass to be the origin of the coordinate system. In this way, the position $\mathbf{x}$ assigned to a rigid body corresponds to the position of its center of mass in world space.

2.1.1 Position and Orientation

Because a rigid body does not deform, the position and velocity of any point on the rigid body at any time is completely described by the offset of that point from the center of mass (origin) in object space $\hat{\mathbf{r}}$ and the position $\mathbf{x}$ and orientation $\mathbf{R}$ of the rigid body in world space. To keep things consistent as much as possible, we will refer to quantities that live in object space with a hat $\hat{\mathbf{r}}$ and the corresponding quantities that live in world space without it $\mathbf{r}$. In this case, $\mathbf{r} = \mathbf{R}\hat{\mathbf{r}}$. This is an offset of the location from the center of mass, so the location of this point in world space is

$$\mathbf{x}_r = \mathbf{x} + \mathbf{r} = \mathbf{x} + \mathbf{R}\hat{\mathbf{r}}.$$

Because $\mathbf{R}^{-1} = \mathbf{R}^T$, the *radius vector* $\hat{\mathbf{r}}$ of the point in object space may be expressed as

$$\hat{\mathbf{r}} = \mathbf{R}^T(\mathbf{x}_r - \mathbf{x}).$$

These relations give a mapping between object space and world space. (Note that since the center of mass in object space is the origin, $\hat{\mathbf{r}}$ is the position in object space.)

In 2D, the rotation of the body in world space may be described by an angle θ. It may also be described by the rotation matrix

$$\mathbf{R} = \begin{pmatrix} \cos\theta & -\sin\theta \\ \sin\theta & \cos\theta \end{pmatrix}.$$

The columns of $\mathbf{R}$ are the locations of the object space points $(1, 0)$ and $(0, 1)$ relative to the center of mass of the rigid body.

In 3D, the rotation of the body may be described as a 3×3 rotation matrix $\mathbf{R}$. The columns of $\mathbf{R}$ are the locations of the object space points $(1,0,0)$, $(0,1,0)$ and $(0,0,1)$ relative to the center of mass of the rigid body.

2.1.2 Mass and Center of Mass

The mass m is the total mass of the particles that compose the rigid body, and the center of mass $\mathbf{x}_c$ is a mass-weighted centroid:

$$m = \sum_i m_i \qquad \mathbf{x}_c = \frac{1}{m} \sum_i m_i \mathbf{x}_i.$$

Note that the center of mass may be shifted to the origin by $\mathbf{x}_i \leftarrow \mathbf{x}_i - \mathbf{x}_c$, since $\frac{1}{m}\sum_i m_i(\mathbf{x}_i - \mathbf{x}_c) = \frac{1}{m}\sum_i m_i \mathbf{x}_i - \frac{1}{m}\mathbf{x}_c \sum_i m_i = \mathbf{x}_c - \mathbf{x}_c = 0$.

Placing the center of mass of the rigid body in object space at the origin gives us some useful identities. Since $\mathbf{x}$ is the center of mass,

$$\mathbf{x} = \frac{1}{m}\sum_i m_i \mathbf{x}_i = \frac{1}{m}\sum_i m_i(\mathbf{x} + \mathbf{r}_i) = \frac{1}{m}\mathbf{x}\sum_i m_i + \frac{1}{m}\sum_i m_i \mathbf{r}_i = \mathbf{x} + \frac{1}{m}\sum_i m_i \mathbf{r}_i.$$

It follows that $\sum_i m_i \mathbf{r}_i = 0$. Moreover, since $\mathbf{R}\hat{\mathbf{r}}_i = \mathbf{r}_i$, we also have the same identity in object space $\sum_i m_i \hat{\mathbf{r}}_i = 0$.

2.1.3 Velocity and Angular Velocity

The velocity of a rigid body may be defined as it was for particles as $\mathbf{v} = \mathbf{x}'$, the velocity of the center of mass. In 2D, the angular velocity is also straightforward to describe. We may define $\omega = \theta'$, where θ is the angle of rotation.

In 3D, the angular velocity is somewhat more interesting. The position of a point on a rigid body $\mathbf{x}_r = \mathbf{x} + \mathbf{r}$. To get an idea of how a rigid body moves, consider the derivative of this.

$$\mathbf{v}_r = \mathbf{x}_r' = \mathbf{x}' + \mathbf{r}' = \mathbf{v} + (\mathbf{R}\hat{\mathbf{r}})' = \mathbf{v} + \mathbf{R}'\hat{\mathbf{r}} = \mathbf{v} + \mathbf{R}'\mathbf{R}^T\mathbf{r}$$

Here, $\mathbf{v}$ describes the motion due to translation, and $\mathbf{R}'\mathbf{R}^T\mathbf{r}$ describes the motion due to rotation. The quantity $\mathbf{R}'\mathbf{R}^T$ is important. Since $\mathbf{R}$ is a rotation, $\mathbf{R}\mathbf{R}^T = \boldsymbol{\delta}$. (We use $\boldsymbol{\delta}$ to denote the identity matrix when dealing with rigid bodies, since it is typical to use $\mathbf{I}$ to represent the inertia.) Taking the derivative of this yields $0 = \mathbf{R}'\mathbf{R}^T + \mathbf{R}\mathbf{R}'^T = \mathbf{R}'\mathbf{R}^T + (\mathbf{R}'\mathbf{R}^T)^T$, so that $\mathbf{R}'\mathbf{R}^T$ is skew-symmetric and can be written in the form

$$\mathbf{R}'\mathbf{R}^T = \begin{pmatrix} 0 & -\omega_3 & \omega_2 \\ \omega_3 & 0 & -\omega_1 \\ -\omega_2 & \omega_1 & 0 \end{pmatrix} = \boldsymbol{\omega}^*$$

where $\boldsymbol{\omega}^*$ is the matrix such that $\boldsymbol{\omega}^*\mathbf{u} = \boldsymbol{\omega} \times \mathbf{u}$ for any vector $\mathbf{u}$. We call the quantity $\boldsymbol{\omega}$ the angular velocity of the rigid body. Substituting this in for the velocity equation above yields

$$\mathbf{v}_r = \mathbf{v} + \boldsymbol{\omega}^*\mathbf{r} = \mathbf{v} + \boldsymbol{\omega} \times \mathbf{r}.$$

It turns out that $\boldsymbol{\omega}$ has a simple and useful physical interpretation. It describes an angular velocity about the axis $\boldsymbol{\omega}/\|\boldsymbol{\omega}\|$ with speed $\|\boldsymbol{\omega}\|$ in radians per second.

It is also convenient that the $\boldsymbol{\omega}$ defined for 3D and the ω defined for 2D are closely related. If we consider the 2D body as living in the xy plane, we can extend the usual 2D definition of rotation to

$$\mathbf{R} = \begin{pmatrix} \cos\theta & -\sin\theta & 0 \\ \sin\theta & \cos\theta & 0 \\ 0 & 0 & 1 \end{pmatrix}.$$

Then, we would define $\mathbf{R}'\mathbf{R}^T = \boldsymbol{\omega}^*$, which leads to

$$\mathbf{R}'\mathbf{R}^T = \begin{pmatrix} -\theta'\sin\theta & -\theta'\cos\theta & 0 \\ \theta'\cos\theta & -\theta'\sin\theta & 0 \\ 0 & 0 & 0 \end{pmatrix} \begin{pmatrix} \cos\theta & \sin\theta & 0 \\ -\sin\theta & \cos\theta & 0 \\ 0 & 0 & 1 \end{pmatrix} = \begin{pmatrix} 0 & -\theta' & 0 \\ \theta' & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}.$$

Using the original 2D definition $\omega = \theta'$ we have $\boldsymbol{\omega} = (0, 0, \omega)$.

Note that the equation $\mathbf{R}'\mathbf{R}^T = \boldsymbol{\omega}^*$ can also be written $\mathbf{R}' = \boldsymbol{\omega}^*\mathbf{R}$, which is a first order ODE that must be solved (in the 3D case) to evolve rotations.

2.1.4 Momentum, Angular Momentum, and Inertia

The momentum of a system is the sum of the momenta of the individual particles

$$\mathbf{p} = \sum_i m_i\mathbf{v}_i = \sum_i m_i(\mathbf{v} + \boldsymbol{\omega} \times \mathbf{r}_i) = m\mathbf{v} + \boldsymbol{\omega} \times \sum_i m_i\mathbf{r}_i = m\mathbf{v}.$$

From this we see for the first time that the center of mass is *preferred*. That is, if we measure the velocity of the rigid body at its center of mass, then the momentum of the rigid body has the form $m\mathbf{v}$, even if the body is rotating.

The angular momentum is defined about a point in space. In the case of a rigid body, we choose the center of mass for this purpose. The angular momentum is the sum of the angular momenta for the individual particles about the rigid body's center of mass,

$$\mathbf{L} = \sum_i m_i\mathbf{r}_i \times \mathbf{v}_i = \sum_i m_i\mathbf{r}_i \times (\mathbf{v} + \boldsymbol{\omega} \times \mathbf{r}_i) = \mathbf{v}\sum_i m_i\mathbf{r}_i + \sum_i m_i\mathbf{r}_i^*\mathbf{r}_i^{*T}\omega = \left(\sum_i m_i\mathbf{r}_i^*\mathbf{r}_i^{*T}\right)\omega$$

where we have made use of the fact that $\boldsymbol{\omega} \times \mathbf{u} = -\mathbf{u} \times \omega = -\mathbf{u}^*\omega = \mathbf{u}^{*T}\omega$. (The use of transposes that seem otherwise unnecessary has two benefits. The first benefit is that with suitable definitions of cross product and star in 2D, the equations expressed for 3D become equally valid in 2D. The second benefit is that it will make an observation about inertia more clear shortly.) This motivates the definition of an angular equivalent for mass, called the inertia, represented by $\mathbf{I}$, and defined as

$$\mathbf{I} = \sum_i m_i\mathbf{r}_i^*\mathbf{r}_i^{*T} \qquad \mathbf{L} = \mathbf{I}\omega.$$

Note from its definition that $\mathbf{I}$ is symmetric positive semidefinite, and for any reasonable body it is symmetric positive definite.

In 3D, the inertia tensor $\mathbf{I}$ has the complication that it depends on the orientation. Thus, it is useful to also define it in object space. Using the identity $\mathbf{u}^*\mathbf{u}^{*T} = \mathbf{u}^T\mathbf{u}\boldsymbol{\delta} - \mathbf{u}\mathbf{u}^T$, one may show that $(\mathbf{R}\mathbf{u})^*(\mathbf{R}\mathbf{u})^{*T} = \mathbf{R}\mathbf{u}^*\mathbf{u}^{*T}\mathbf{R}^T$. Then, we have

$$\mathbf{I} = \sum_i m_i\mathbf{r}_i^*\mathbf{r}_i^{*T} = \sum_i m_i(\mathbf{R}\hat{\mathbf{r}}_i)^*(\mathbf{R}\hat{\mathbf{r}}_i)^{*T} = \sum_i m_i\mathbf{R}\hat{\mathbf{r}}_i^*\hat{\mathbf{r}}_i^{*T}\mathbf{R}^T = \mathbf{R}\left(\sum_i m_i\hat{\mathbf{r}}_i^*\hat{\mathbf{r}}_i^{*T}\right)\mathbf{R}^T.$$

Then, we can make an analygous definition for the inertia tensor in object space

$$\hat{\mathbf{I}} = \sum_i m_i \hat{\mathbf{r}}_i^* \hat{\mathbf{r}}_i^{*T} \qquad \mathbf{I} = \mathbf{R}\hat{\mathbf{I}}\mathbf{R}^T.$$

The important reason for introducing the inertia tensor in object space is that $\hat{\mathbf{I}}$ does not change during the simulation, so it may be precomputed and used to compute the inertia $\mathbf{I}$ during the simulation.

If we return to the consideration of 2D as living in 3D, we note that $\mathbf{L} = \sum_i m_i \mathbf{r}_i \times \mathbf{v}_i$ is the sum of terms that are cross products of vectors in the xy plane, so that we can write $\mathbf{L} = (0, 0, L)$. Along with $\boldsymbol{\omega} = (0, 0, \omega)$ and $\mathbf{L} = \mathbf{I}\boldsymbol{\omega}$, we can replace the use of $\mathbf{I}$ with $I = \mathbf{I}_{33} = \sum_i m_i \mathbf{r}_i \cdot \mathbf{r}_i$. Thus, in 2D, we have $L = I\omega$, where all three are scalars.

2.1.5 Force and Torque

For individual particles, there is a single vector quantity, force, that describes the outside influence on the evolution of the particle. For rigid bodies, there is a linear component and an angular component. As in the particle case, we define force according to Newton's second law

$$\mathbf{F} = \mathbf{p}' = m\mathbf{v}' = m\mathbf{a} = \frac{d}{dt}\sum_i m_i \mathbf{v}_i = \sum_i \mathbf{F}_i.$$

In particular, we may compute the acceleration on a rigid body by adding up the forces applied to each particle in the body and then dividing off the total mass of the rigid body. Then, we can evolve the body using $\mathbf{a} = \mathbf{v}' = \mathbf{x}''$.

We also have an angular equivalent of force, referred to as torque

$$\boldsymbol{\tau} = \mathbf{L}' = \frac{d}{dt}\sum_i m_i \mathbf{r}_i \times \mathbf{v}_i = \sum_i m_i \mathbf{v}_i \times \mathbf{v}_i + \sum_i m_i \mathbf{r}_i \times \mathbf{a}_i = \sum_i \mathbf{r}_i \times \mathbf{f}_i = \sum_i \boldsymbol{\tau}_i.$$

This means we can compute the torque on the rigid body by adding the torque contribution $\boldsymbol{\tau}_i = \mathbf{r}_i \times \mathbf{f}_i$ at each particle.

In 2D as the xy plane, $\mathbf{L} = (0, 0, L)$ and $\boldsymbol{\tau} = \mathbf{L}'$, so that that $\boldsymbol{\tau} = (0, 0, \tau)$ and $\tau = L'$. Letting $\boldsymbol{\tau}_i = (0, 0, \tau_i) = \mathbf{r}_i \times \mathbf{f}_i$, we have $\tau = \sum_i \tau_i$. We can evolve rotations by $\tau = L' = I\omega' = I\theta''$, where I is the scalar moment of inertia. Thus, we find that the task of evolving rotations in 2D is not much more difficult than evolving positions.

The story in 3D is quite different. The main complication in 3D comes from the dependence of $\boldsymbol{\omega}$ on $\mathbf{R}$. Because the angular momentum L is a conserved quantity and $\mathbf{L} = \mathbf{I}\boldsymbol{\omega} = \mathbf{R}\hat{\mathbf{I}}\mathbf{R}^T\boldsymbol{\omega}$, we can write $\boldsymbol{\omega} = \mathbf{R}\hat{\mathbf{I}}^{-1}\mathbf{R}^T\mathbf{L}$. In practice, this may be updated by advancing $\mathbf{L}$ using $\boldsymbol{\tau} = \mathbf{L}'$, then computing $\boldsymbol{\omega} = \mathbf{I}^{-1}\mathbf{L}$ and finally evolving rotations by evolving the ODE $\mathbf{R}' = \boldsymbol{\omega}^*\mathbf{R}$.

2.2 Collisions

When an object modeled as a mass spring system collides with another object, it is sufficient to model the collision response locally. However, in the case of rigid bodies, it is necessary to consider the effects on position and rotation of the entire body. In this case, the response to a simple collision can cause the entire body to move in interesting ways. When rigid bodies become involved in more than one collision, things become complicated very quickly. With mass spring systems, we have enough degrees of freedom that we can deform the mesh locally.

2.2.1 Impulse Application

When applying impulse between two particles, the particles apply equal and opposite impulses, and this is sufficient to guarantee conservation of linear momentum. For rigid bodies, it is necessary to conserve angular momentum as well, making it necessary to apply the equal and opposite forces *at the same location*. Rigid bodies may also exchange torque. Another important aspect of rigid body simulation is the application of impulses and angular impulses. This is important since there are some aspects of rigid body simulation that cannot be treated properly with forces (e.g., collisions or contact with friction), and impulses are easier to apply than forces because the orientation does not change during their application. For the remainder of this section, we assume that the orientations of the rigid bodies do not change, so that the inertia tensor $\mathbf{I}$ also does not change.

We begin by assuming that an impulse $\mathbf{j}$ and angular impulse $\mathbf{j}_{\boldsymbol{\tau}}$ are applied at location $\mathbf{x}_r = \mathbf{x} + \mathbf{r}$. The total change in linear momentum is just $\mathbf{j}$. The rigid body's angular momentum is changed by the linear impulse in addition to the angular impulse. The total change in angular momentum is $\mathbf{r} \times \mathbf{j} + \mathbf{j}_{\boldsymbol{\tau}}$.

$$\begin{pmatrix} \Delta\mathbf{p} \\ \Delta\mathbf{L} \end{pmatrix} = \begin{pmatrix} \mathbf{j} \\ \mathbf{r} \times \mathbf{j} + \mathbf{j}_{\boldsymbol{\tau}} \end{pmatrix} = \begin{pmatrix} \boldsymbol{\delta} & 0 \\ \mathbf{r}^* & \boldsymbol{\delta} \end{pmatrix} \begin{pmatrix} \mathbf{j} \\ \mathbf{j}_{\boldsymbol{\tau}} \end{pmatrix}$$

From this we can compute the change in velocity and angular velocity, taking advantage of the fact that inertia tensor is not changing

$$\begin{pmatrix} \Delta\mathbf{v} \\ \Delta\boldsymbol{\omega} \end{pmatrix} = \begin{pmatrix} m^{-1}\Delta\mathbf{p} \\ \mathbf{I}^{-1}\Delta\mathbf{L} \end{pmatrix} = \begin{pmatrix} m^{-1}\boldsymbol{\delta} & 0 \\ 0 & \mathbf{I}^{-1} \end{pmatrix} \begin{pmatrix} \boldsymbol{\delta} & 0 \\ \mathbf{r}^* & \boldsymbol{\delta} \end{pmatrix} \begin{pmatrix} \mathbf{j} \\ \mathbf{j}_{\boldsymbol{\tau}} \end{pmatrix}$$

We can now use the change in the body's linear and angular velocity to express the change in linear and angular velocity $\mathbf{v}^r$ and $\boldsymbol{\omega}_r$ at the point $\mathbf{x}_r$ where the impulse was applied. We do this using $\mathbf{v}_r = \mathbf{v} + \boldsymbol{\omega} \times \mathbf{r} = \mathbf{v} + \mathbf{r}^{*T}\boldsymbol{\omega}$. Then, we have

$$\begin{pmatrix} \Delta\mathbf{v}_r \\ \Delta\boldsymbol{\omega}_r \end{pmatrix} = \begin{pmatrix} \Delta\mathbf{v} + \mathbf{r}^{*T}\Delta\boldsymbol{\omega} \\ \Delta\boldsymbol{\omega} \end{pmatrix} = \begin{pmatrix} \boldsymbol{\delta} & 0 \\ \mathbf{r}^* & \boldsymbol{\delta} \end{pmatrix}^T \begin{pmatrix} m^{-1}\boldsymbol{\delta} & 0 \\ 0 & \mathbf{I}^{-1} \end{pmatrix} \begin{pmatrix} \boldsymbol{\delta} & 0 \\ \mathbf{r}^* & \boldsymbol{\delta} \end{pmatrix} \begin{pmatrix} \mathbf{j} \\ \mathbf{j}_{\boldsymbol{\tau}} \end{pmatrix}$$

Observe that the relationship is a symmetric positive definite matrix, which we will denote $\mathbf{K}$.

$$\mathbf{K} = \begin{pmatrix} \boldsymbol{\delta} & 0 \\ \mathbf{r}^* & \boldsymbol{\delta} \end{pmatrix}^T \begin{pmatrix} m^{-1}\boldsymbol{\delta} & 0 \\ 0 & \mathbf{I}^{-1} \end{pmatrix} \begin{pmatrix} \boldsymbol{\delta} & 0 \\ \mathbf{r}^* & \boldsymbol{\delta} \end{pmatrix} = \begin{pmatrix} m^{-1}\boldsymbol{\delta} + \mathbf{r}^{*T}\mathbf{I}^{-1}\mathbf{r}^* & \mathbf{r}^{*T}\mathbf{I}^{-1} \\ \mathbf{I}^{-1}\mathbf{r}^* & \mathbf{I}^{-1} \end{pmatrix}$$

The matrix $\mathbf{K}^{-1}$ has an interesting interpretation. When angular and linear velocity and momentum are measured from the center of mass of the body, the momentum is related to the velocity through the mass, and the angular momentum is related to the angular velocity through the inertia tensor. If these quantities are measured about some point other than the center of mass, the momentum will in general depend on both velocity and angular velocity. Thus, it is necessary to consider the relationship of linear and angular velocity to linear and angular momentum at the same time. The "coupled inertia tensor" that relates these quantities is $\mathbf{K}^{-1}$. Note that when $\mathbf{r} = 0$, it decays back into mass and the inertia tensor.

Returning to the problem of exchanging linear and angular impulses between two bodies a and b, whose attributes are indicated with subscripts as $\mathbf{r}_a$, $\mathbf{r}_b$, m_a, m_b, $\mathbf{I}_a$, $\mathbf{I}_b$, $\mathbf{K}_a$, $\mathbf{K}_b$, etc. Assume

$\mathbf{j}$ and $\mathbf{j}_{\boldsymbol{\tau}}$ are applied to body a, and $-\mathbf{j}$ and $-\mathbf{j}_{\boldsymbol{\tau}}$ are applied to body b. The relative change in angular velocity and linear velocity that this causes is

$$\begin{pmatrix} \Delta\mathbf{v}_{rel} \\ \Delta\boldsymbol{\omega}_{rel} \end{pmatrix} = \begin{pmatrix} \Delta\mathbf{v}_{ra} - \Delta\mathbf{v}_{rb} \\ \Delta\boldsymbol{\omega}_{ra} - \Delta\boldsymbol{\omega}_{rb} \end{pmatrix} = \mathbf{K}_a \begin{pmatrix} \mathbf{j} \\ \mathbf{j}_{\boldsymbol{\tau}} \end{pmatrix} - \mathbf{K}_b \begin{pmatrix} -\mathbf{j} \\ -\mathbf{j}_{\boldsymbol{\tau}} \end{pmatrix} = (\mathbf{K}_a + \mathbf{K}_b) \begin{pmatrix} \mathbf{j} \\ \mathbf{j}_{\boldsymbol{\tau}} \end{pmatrix}$$

Thus, if we know the linear and angular velocity change we would like to achieve at a particular point, we can compute the angular and linear impulse needed to achieve it by solving a 6×6 symmetric positive definite system (3×3 in 2D).

2.2.2 Collision Impulse

Consider the situation where two rigid bodies are colliding. At the collision point, we can identify a plane with normal $\mathbf{n}$ where they touch. The bodies will exchange an impulse of magnitude j in the direction $\mathbf{n}$ so that $\mathbf{j} = j\mathbf{n}$. The bodies do not apply an explicit angular impulse, so that $\mathbf{j}_{\boldsymbol{\tau}} = 0$. We also do not care about the angular velocity at the collision point. This reduces equation that describes the change in relative velocity at the collision point to

$$\Delta\mathbf{v}_{rel} = (m_a^{-1}\boldsymbol{\delta} + \mathbf{r}_a^{*T}\mathbf{I}_a^{-1}\mathbf{r}_a^{*} + m_b^{-1}\boldsymbol{\delta} + \mathbf{r}_b^{*T}\mathbf{I}_b^{-1}\mathbf{r}_b^{*})\mathbf{j}.$$

Letting $\tilde{\mathbf{K}}$ be the top left block of $\mathbf{K}$,

$$\tilde{\mathbf{K}}_a = m_a^{-1}\boldsymbol{\delta} + \mathbf{r}_a^{*T}\mathbf{I}_a^{-1}\mathbf{r}_a^{*} \qquad \tilde{\mathbf{K}}_b = m_b^{-1}\boldsymbol{\delta} + \mathbf{r}_b^{*T}\mathbf{I}_b^{-1}\mathbf{r}_b^{*},$$

we can write an expression the change in relative velocity as

$$\Delta\mathbf{v}_{rel} = (\tilde{\mathbf{K}}_a + \tilde{\mathbf{K}}_b)\mathbf{j}.$$

In this case, we only care about the relative velocity in the normal direction,

$$\Delta\mathbf{v}_{rel,n} = \Delta\mathbf{v}_{rel} \cdot \mathbf{n} = \mathbf{n}^T(\tilde{\mathbf{K}}_a + \tilde{\mathbf{K}}_b)\mathbf{j} = \mathbf{n}^T(\tilde{\mathbf{K}}_a + \tilde{\mathbf{K}}_b)\mathbf{n}j.$$

We now only need to know the desired normal relative velocity to solve for the impulse. This comes from physical property known as the coefficient of restitution ε, which establishes the ratio between the velocities before and after a collision. If we let $\mathbf{v}_{rel} = \mathbf{v}_a - \mathbf{v}_b$ be the relative velocity of the bodies before the collision, then the normal velocity after the collision should be $-\varepsilon\mathbf{v}_{rel} \cdot \mathbf{n}$. Since the normal relative velocity before the collision was $\mathbf{v}_{rel} \cdot \mathbf{n}$, the change required is

$$\Delta\mathbf{v}_{rel,n} = -(1 + \varepsilon)\mathbf{v}_{rel} \cdot \mathbf{n}.$$

We can now solve for the impulse magnitude j as

$$j = \frac{-(1 + \varepsilon)\mathbf{n}^T\mathbf{v}_{rel}}{\mathbf{n}^T(\tilde{\mathbf{K}}_a + \tilde{\mathbf{K}}_b)\mathbf{n}},$$

and the impulse exchanged during the collision is $j\mathbf{n}$.