

Particle Systems

**Brian Curless
CSE 457
Spring 2013**

Reading

Required:

- ◆ Witkin, *Particle System Dynamics*, SIGGRAPH '01 course notes on Physically Based Modeling.
- ◆ Witkin and Baraff, *Differential Equation Basics*, SIGGRAPH '01 course notes on Physically Based Modeling.

Optional

- ◆ Hockney and Eastwood. *Computer simulation using particles*. Adam Hilger, New York, 1988.
- ◆ Gavin Miller. "The motion dynamics of snakes and worms." *Computer Graphics* 22:169-178, 1988.

What are particle systems?

A **particle system** is a collection of point masses that obeys some physical laws (e.g, gravity, heat convection, spring behaviors, ...).

Particle systems can be used to simulate all sorts of physical phenomena:

Fire

Fluids

Snow

Objects

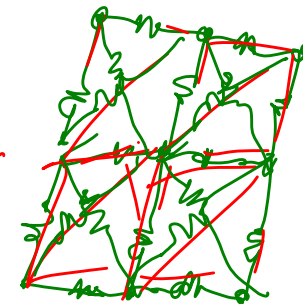
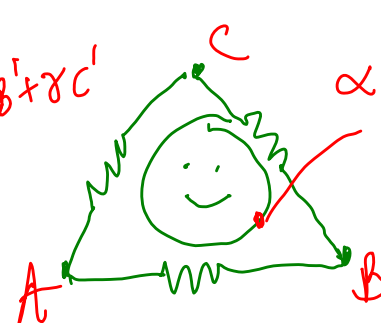
Smoke

Strand .. rubber band, hair

Projectiles

Cloth

Swarms



Particle in a flow field

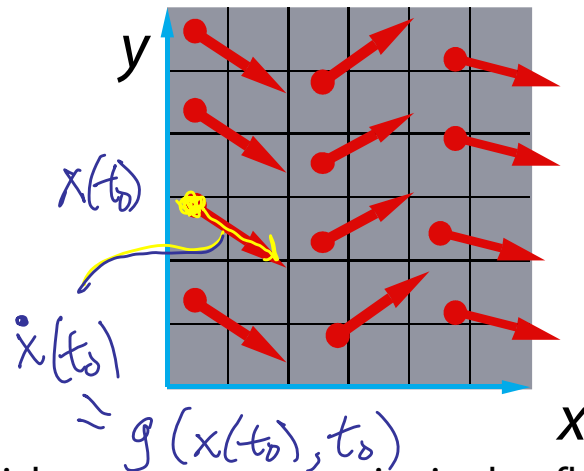
We begin with a single particle with:

♦ Position, $\mathbf{x} = \begin{bmatrix} x \\ y \end{bmatrix}$

♦ Velocity, $\mathbf{v} \equiv \dot{\mathbf{x}} = \frac{d\mathbf{x}}{dt} = \begin{bmatrix} dx/dt \\ dy/dt \end{bmatrix}$

Suppose the velocity is actually dictated by a driving function, a vector flow field, \mathbf{g} :

$$\dot{\mathbf{x}} = \mathbf{g}(\mathbf{x}, t)$$



If a particle starts at some point in that flow field, how should it move?

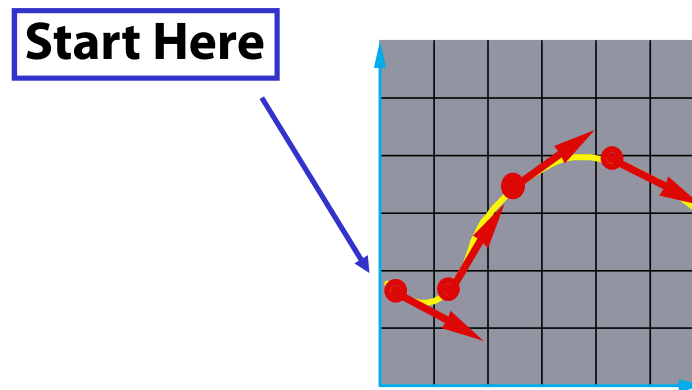
Diff eqs and integral curves

The equation

$$\dot{\mathbf{x}} = \mathbf{g}(\mathbf{x}, t)$$

is actually a **first order differential equation**.

We can solve for \mathbf{x} through time by starting at an initial point and stepping along the vector field:



This is called an **initial value problem** and the solution is called an **integral curve**.

Euler's method

One simple approach is to choose a time step, Δt , and take linear steps along the flow:

$$\underline{\mathbf{x}(t + \Delta t)} = \underline{\mathbf{x}(t)} + \underline{\Delta \mathbf{x}} = \underline{\mathbf{x}(t)} + \Delta t \cdot \frac{\Delta \mathbf{x}}{\Delta t}$$

$$\approx \mathbf{x}(t) + \Delta t \cdot \dot{\mathbf{x}}(t)$$

$$\approx \mathbf{x}(t) + \Delta t \cdot \underline{\mathbf{g}(\mathbf{x}(t), t)}$$

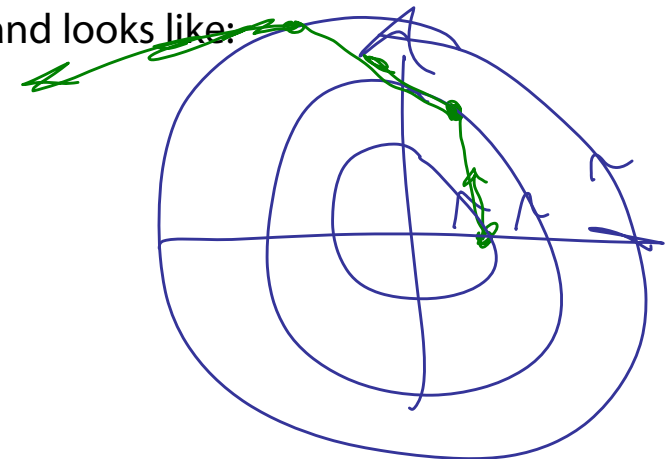
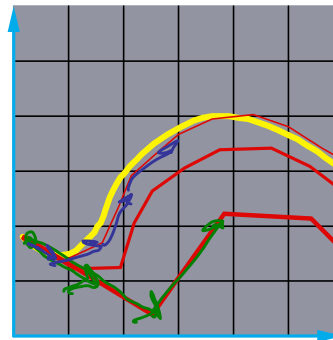
$$\frac{\Delta x}{\Delta t} \approx \dot{x}(t)$$

$$t = i \Delta t$$

Writing as a time iteration:

$$\underline{\mathbf{x}^{i+1}} = \underline{\mathbf{x}^i} + \Delta t \cdot \underline{\mathbf{g}^i} \quad \text{with} \quad \underline{\mathbf{g}^i} \equiv \underline{\mathbf{g}(\mathbf{x}^i, t = i \Delta t)}$$

This approach is called **Euler's method** and looks like:



Properties:

- ◆ Simplest numerical method
- ◆ Bigger steps, bigger errors. Error $\sim O(\Delta t^2)$.

Need to take pretty small steps, so not very efficient. Better (more complicated) methods exist, e.g., adaptive timesteps, Runge-Kutta, and implicit integration.

Particle in a force field

Now consider a particle in a force field \mathbf{f} .

In this case, the particle has:

- ♦ Mass, m
- ♦ Acceleration, $\mathbf{a} \equiv \ddot{\mathbf{x}} = \dot{\mathbf{v}} = \frac{d\mathbf{v}}{dt} = \frac{d^2\mathbf{x}}{dt^2}$

The particle obeys Newton's law:

$$\mathbf{f} = m\mathbf{a} = m\ddot{\mathbf{x}}$$

So, given a force, we can solve for the acceleration:

$$\ddot{\mathbf{x}} = \frac{\mathbf{f}}{m}$$

The force field \mathbf{f} can in general depend on the position and velocity of the particle as well as time.

Thus, with some rearrangement, we end up with:

$$\ddot{\mathbf{x}} = \frac{\mathbf{f}(\mathbf{x}, \dot{\mathbf{x}}, t)}{m}$$

Second order equations

This equation:

$$\ddot{\mathbf{x}} = \frac{\mathbf{f}(\mathbf{x}, \dot{\mathbf{x}}, t)}{m}$$

is a **second order differential equation**.

Our solution method, though, worked on first order differential equations.

We can rewrite the second order equation as:

$$\begin{bmatrix} \dot{\mathbf{x}} = \mathbf{v} \\ \dot{\mathbf{v}} = \frac{\mathbf{f}(\mathbf{x}, \mathbf{v}, t)}{m} \end{bmatrix} \text{ or } \begin{bmatrix} \dot{\mathbf{x}} \\ \dot{\mathbf{v}} \end{bmatrix} = \begin{bmatrix} \mathbf{v} \\ \frac{\mathbf{f}(\mathbf{x}, \mathbf{v}, t)}{m} \end{bmatrix}$$

where we substitute in \mathbf{v} and its derivative to get a pair of **coupled first order equations**.

Phase space

$$\begin{bmatrix} \mathbf{x} \\ \mathbf{v} \end{bmatrix}$$

Concatenate \mathbf{x} and \mathbf{v} to make a 6-vector: position in **phase space**.

$$\begin{bmatrix} \dot{\mathbf{x}} \\ \dot{\mathbf{v}} \end{bmatrix}$$

Taking the time derivative: another 6-vector.

$$\begin{bmatrix} \dot{\mathbf{x}} \\ \dot{\mathbf{v}} \end{bmatrix} = \begin{bmatrix} \mathbf{v} \\ \mathbf{f}/m \end{bmatrix}$$

A vanilla 1st-order differential equation.

Differential equation solver

Starting with:

$$\begin{bmatrix} \dot{\mathbf{x}} \\ \dot{\mathbf{v}} \end{bmatrix} = \begin{bmatrix} \mathbf{v} \\ \mathbf{f}/m \end{bmatrix}$$

Applying Euler's method:

$$\begin{aligned} \mathbf{x}(t + \Delta t) &\approx \mathbf{x}(t) + \Delta t \cdot \dot{\mathbf{x}}(t) \\ \dot{\mathbf{x}}(t + \Delta t) &\approx \dot{\mathbf{x}}(t) + \Delta t \cdot \ddot{\mathbf{x}}(t) \end{aligned}$$

And making substitutions:

$$\begin{aligned} \mathbf{x}(t + \Delta t) &\approx \mathbf{x}(t) + \Delta t \cdot \mathbf{v}(t) \\ \mathbf{v}(t + \Delta t) &\approx \mathbf{v}(t) + \Delta t \cdot \frac{\mathbf{f}(\mathbf{x}(t), \mathbf{v}(t), t)}{m} \end{aligned}$$

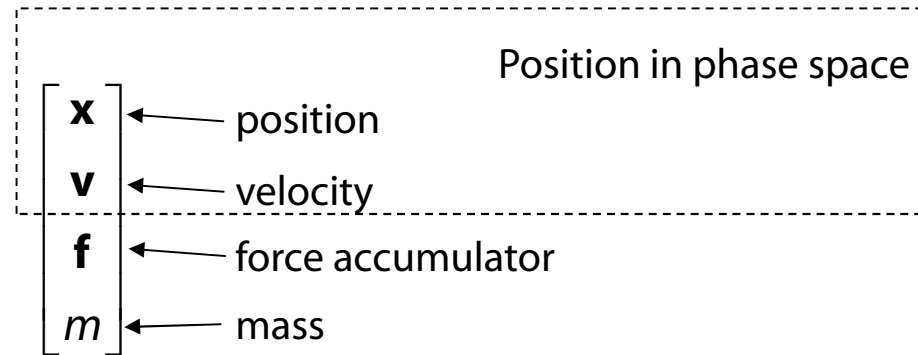
Writing this as an iteration, we have:

$$\begin{aligned} \mathbf{x}^{i+1} &= \mathbf{x}^i + \Delta t \cdot \mathbf{v}^i \\ \mathbf{v}^{i+1} &= \mathbf{v}^i + \Delta t \cdot \frac{\mathbf{f}^i}{m} \end{aligned} \quad \text{with} \quad \mathbf{f}^i \equiv \mathbf{f}(\mathbf{x}^i, \mathbf{v}^i, t)$$

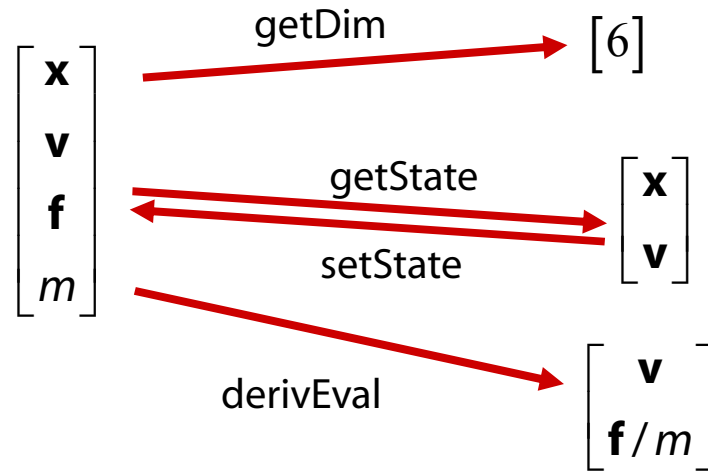
Again, performs poorly for large Δt .

Particle structure

How do we represent a particle?



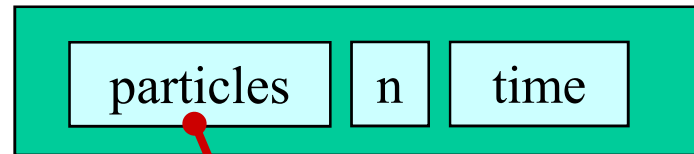
Single particle solver interface



$$S_{old} = \text{getState}()$$
$$S_{new} = (S_{old} + \Delta t \cdot \text{derivEval}())$$
$$\text{setState}(S_{new})$$

Particle systems

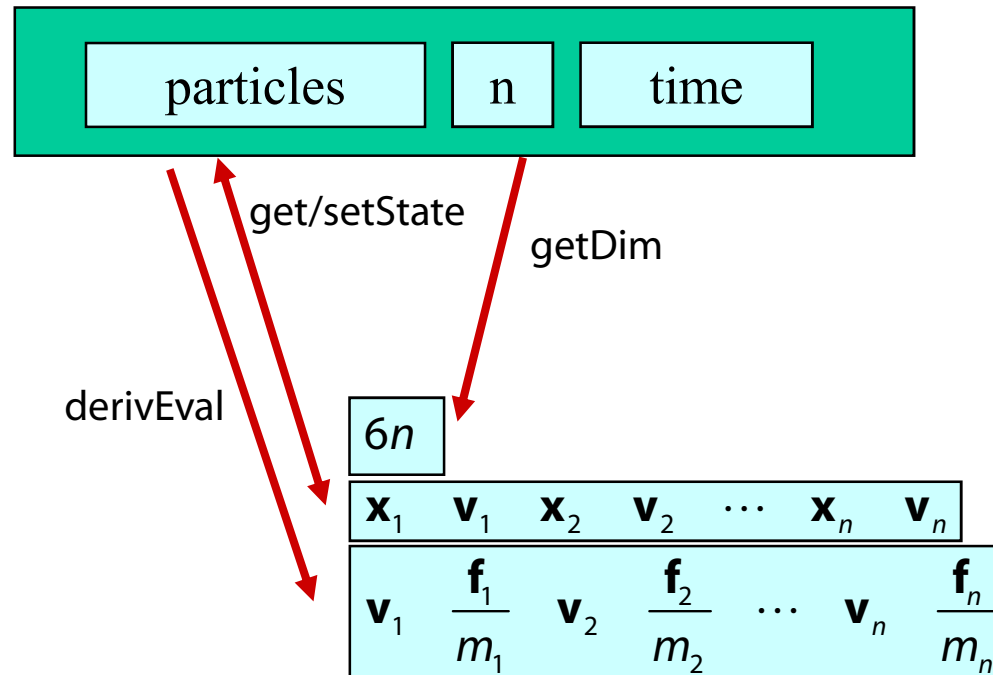
In general, we have a particle system consisting of n particles to be managed over time:



$$\begin{bmatrix} \mathbf{x}_1 \\ \mathbf{v}_1 \\ \mathbf{f}_1 \\ m_1 \end{bmatrix} \begin{bmatrix} \mathbf{x}_2 \\ \mathbf{v}_2 \\ \mathbf{f}_2 \\ m_2 \end{bmatrix} \dots \begin{bmatrix} \mathbf{x}_n \\ \mathbf{v}_n \\ \mathbf{f}_n \\ m_n \end{bmatrix}$$

Particle system solver interface

For n particles, the solver interface now looks like:



Particle system diff. eq. solver

We can solve the evolution of a particle system again using the Euler method:

$$\begin{bmatrix} \mathbf{x}_1^{i+1} \\ \mathbf{v}_1^{i+1} \\ \vdots \\ \mathbf{x}_n^{i+1} \\ \mathbf{v}_n^{i+1} \end{bmatrix} = \begin{bmatrix} \mathbf{x}_1^i \\ \mathbf{v}_1^i \\ \vdots \\ \mathbf{x}_n^i \\ \mathbf{v}_n^i \end{bmatrix} + \Delta t \begin{bmatrix} \mathbf{v}_1^i \\ \mathbf{f}_1^i / m_1 \\ \vdots \\ \mathbf{v}_n^i \\ \mathbf{f}_n^i / m_n \end{bmatrix}$$

\uparrow \uparrow
getState() + dt.derivEval()

Forces

Each particle can experience a force which sends it on its merry way.

Where do these forces come from? Some examples:

- ◆ Constant (gravity)
- ◆ Position/time dependent (force fields)
- ◆ Velocity-dependent (drag)
- ◆ N-ary (springs)

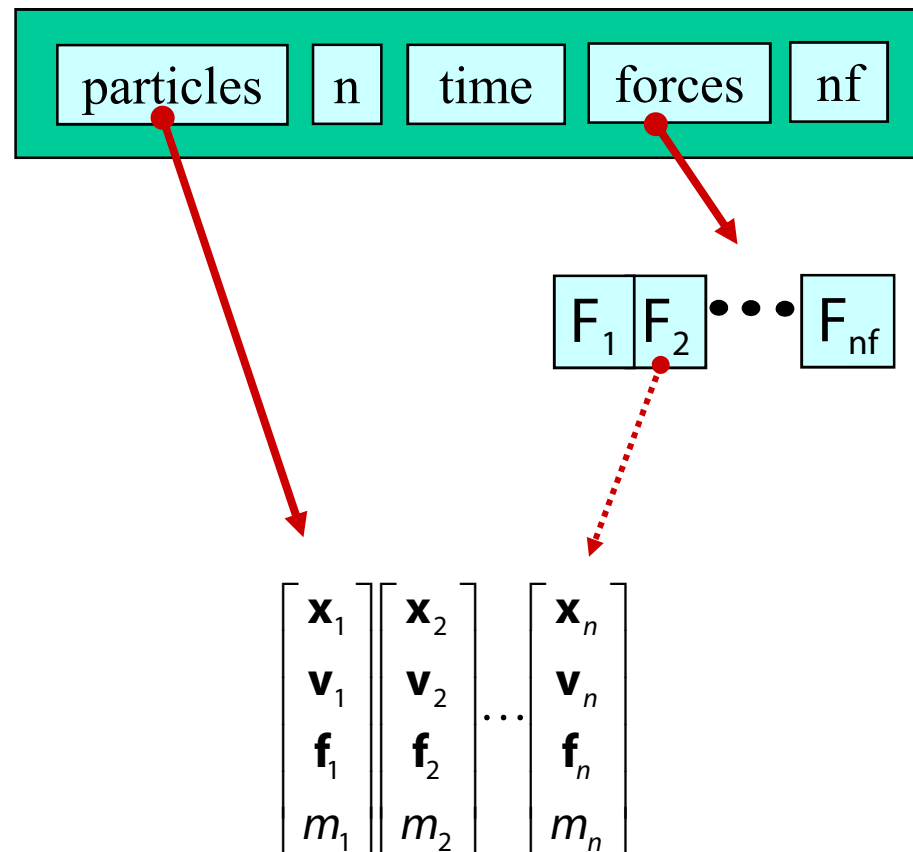
How do we compute the net force on a particle?

Sum them up
"principle of superposition"

Particle systems with forces

Force objects are black boxes that point to the particles they influence and add in their contributions.

We can now visualize the particle system with force objects:



Gravity and viscous drag

The force due to **gravity** is simply:

$$\mathbf{f}_{grav} = m\mathbf{G}$$

$$\text{p} \rightarrow \mathbf{f} \quad += \quad \text{p} \rightarrow m \quad * \quad \mathbf{F} \rightarrow \mathbf{G}$$

Often, we want to slow things down with **viscous drag**:

$$\mathbf{f}_{drag} = -k_{drag} \mathbf{v}$$

$$\text{p} \rightarrow \mathbf{f} \quad -= \quad \mathbf{F} \rightarrow k \quad * \quad \text{p} \rightarrow \mathbf{v}$$

$$\begin{aligned} \mathbf{f}_{terminal} &= m\mathbf{G} - k_{drag} \mathbf{v}_{term} = 0 \\ \mathbf{v}_{term} &= \frac{m\mathbf{G}}{k_{drag}} \end{aligned}$$

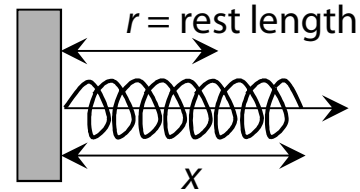


Damped spring

A spring is a simple examples of an “N-ary” force.

Recall the equation for the force due to a 1D spring:

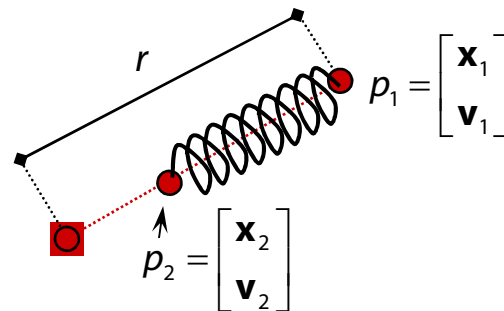
$$f = -\underline{k_{spring}}(x - r)$$



With damping:

$$f = -[\underline{k_{spring}}(x - r) + \underline{k_{damp}}v]$$

In 2D or 3D, we get:



$$\Delta \mathbf{x} \equiv \mathbf{x}_1 - \mathbf{x}_2$$

$$\Delta \hat{\mathbf{x}} \equiv \frac{\Delta \mathbf{x}}{\|\Delta \mathbf{x}\|}$$

$$\Delta \mathbf{v} \equiv \mathbf{v}_1 - \mathbf{v}_2$$

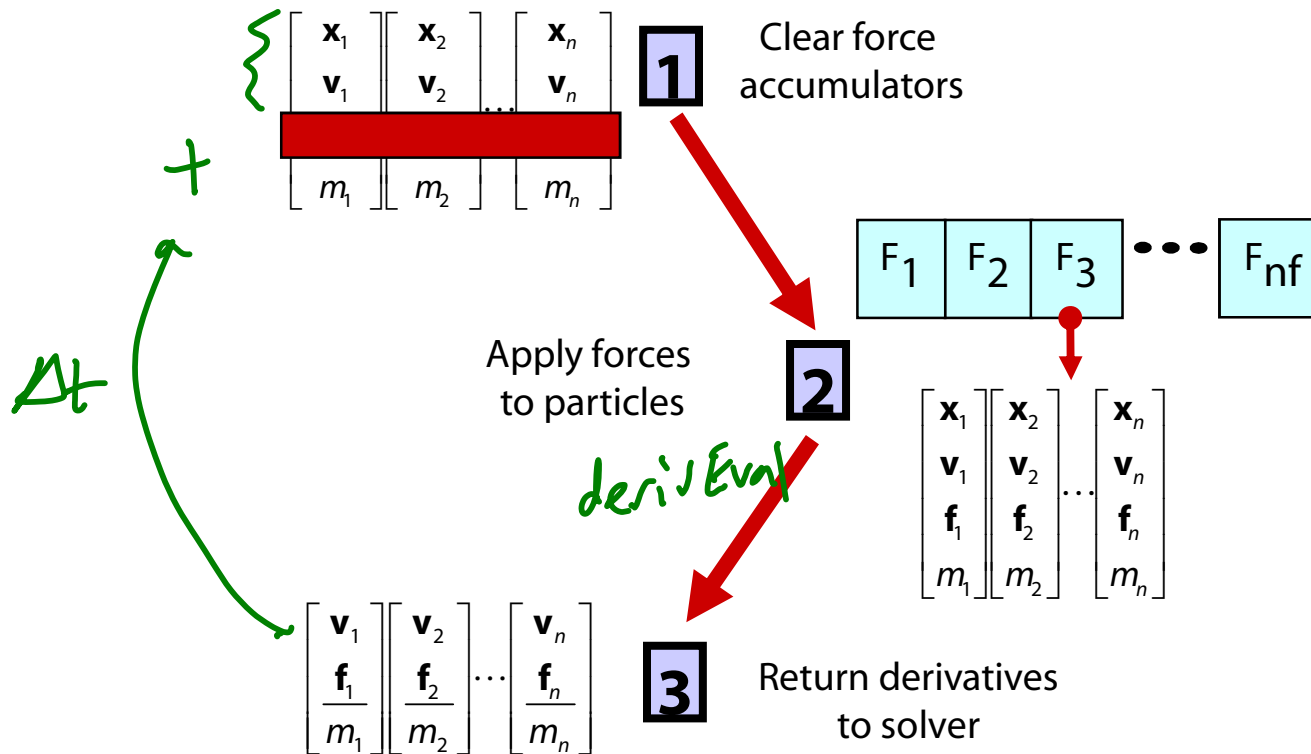
$$f_1 = -\left[k_{spring} (\|\Delta \mathbf{x}\| - r) + k_{damp} (\Delta \mathbf{v} \cdot \Delta \hat{\mathbf{x}}) \right] \Delta \hat{\mathbf{x}}$$

$$f_2 = -f_1$$

Note: stiff spring systems can be very unstable under Euler integration. Simple solutions include heavy damping (may not look good), tiny time steps (slow), or better integration (Runge-Kutta is straightforward).

derivEval

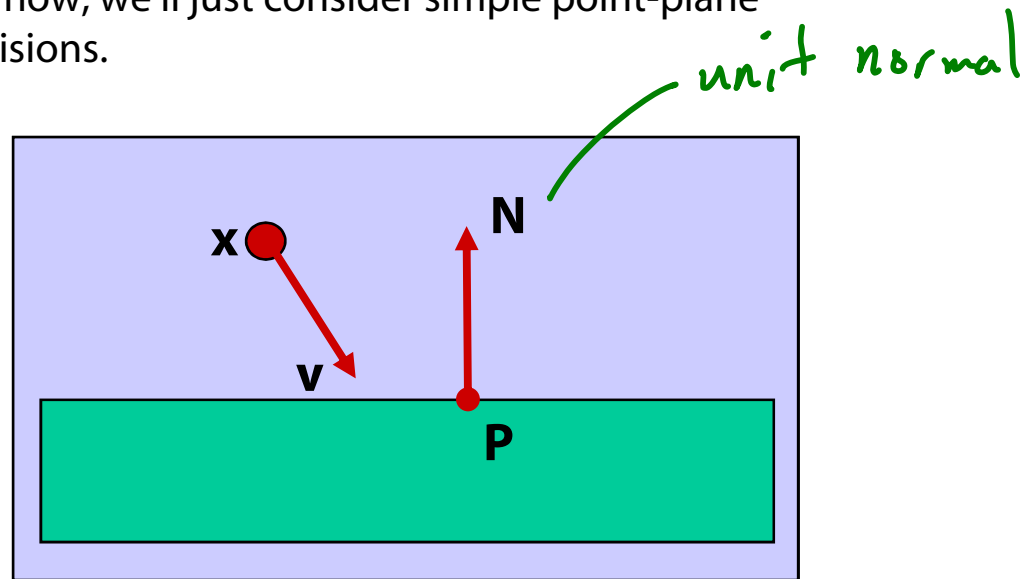
1. Clear forces
 - Loop over particles, zero force accumulators
2. Calculate forces
 - Sum all forces into accumulators
3. Return derivatives
 - Loop over particles, return \mathbf{v} and \mathbf{f}/m



Bouncing off the walls

Handling collisions is a useful add-on for a particle simulator.

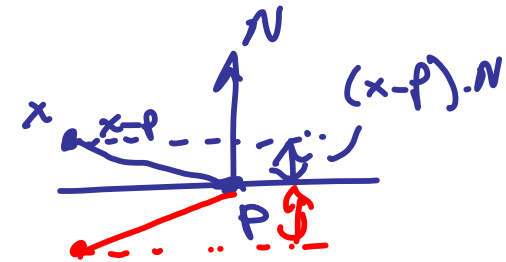
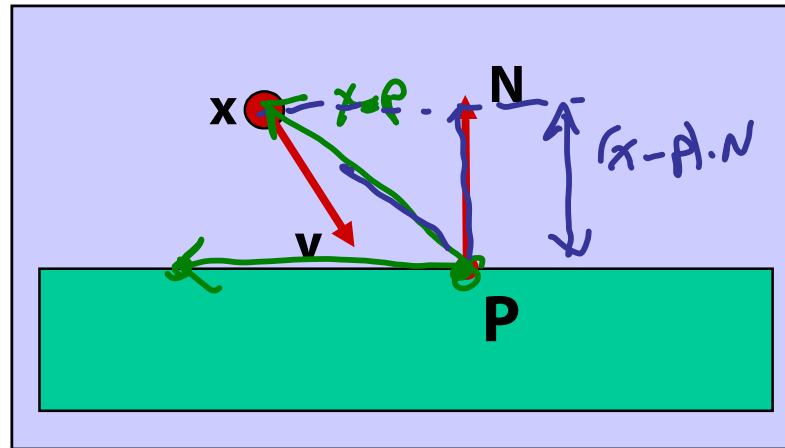
For now, we'll just consider simple point-plane collisions.



A plane is fully specified by any point P on the plane and its normal N .

Collision Detection

How do you decide when you've made **exact** contact with the plane?



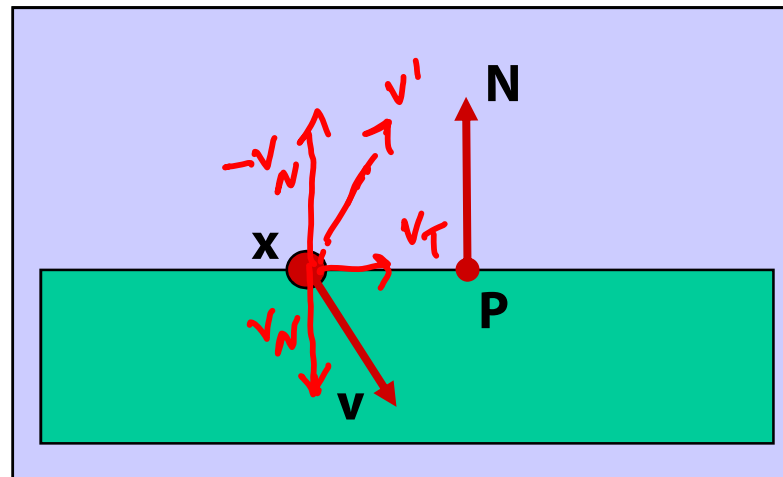
$$\cancel{x-p=0}$$

$$(x-p) \cdot N = 0$$

$(x-p) \cdot N = \text{dist. from plane}$
Signed

Normal and tangential velocity

To compute the collision response, we need to consider the normal and tangential components of a particle's velocity.

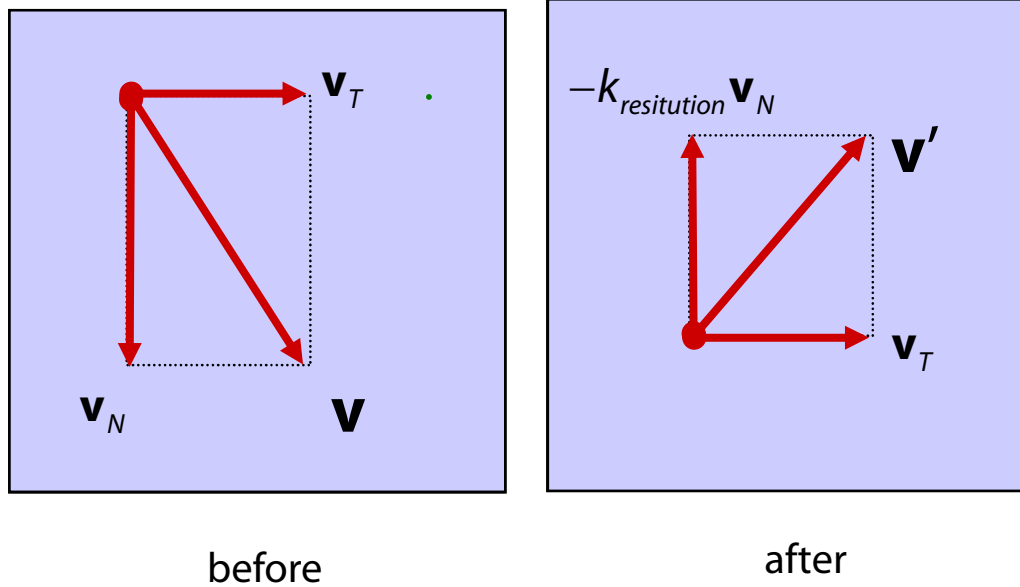


$$v' = v_T - v_N$$

A diagram showing a velocity vector v being decomposed into its normal component v_N and tangential component v_T . The normal component v_N is perpendicular to the surface, and the tangential component v_T is parallel to the surface. The vector v is the sum of v_N and v_T .

$$\mathbf{v}_N = (\mathbf{N} \cdot \mathbf{v})\mathbf{N}$$
$$\mathbf{v}_T = \mathbf{v} - \mathbf{v}_N$$
$$\mathbf{v} = v_N + v_T$$

Collision Response



The response to collision is then to immediately replace the current velocity with a new velocity:

$$\mathbf{v}' = \mathbf{v}_T - k_{\text{restitution}} \mathbf{v}_N$$

$e \in [0, 1]$

The particle will then move according to this velocity in the next timestep.

Collision without contact

In general, we don't sample moments in time when particles are in *exact* contact with the surface.

There are a variety of ways to deal with this problem.

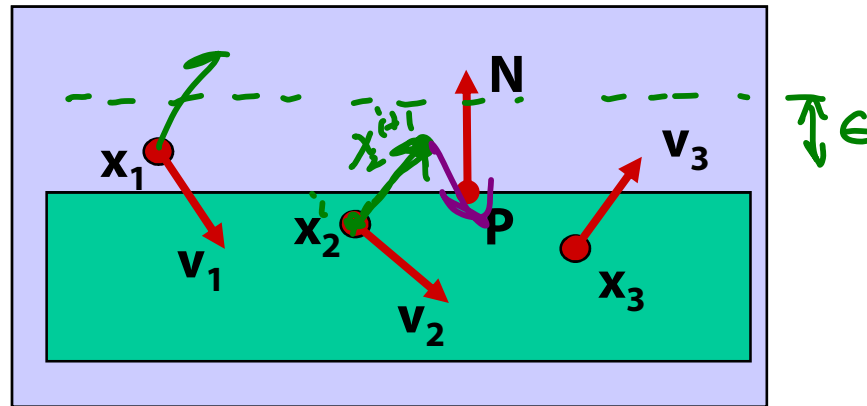
The most expensive is **backtracking**: determine if a collision must have occurred, and then roll back the simulation to the moment of contact.

A simple alternative is to determine if a collision must have occurred in the past, and then pretend that you're currently in exact contact.

Very simple collision response

How do you decide when you've had a collision during a timestep?

$$(x-p) \cdot N \leq 0$$



A problem with this approach is that particles will disappear under the surface. We can reduce this problem by essentially offsetting the surface:

$$(x-p) \cdot N \leq \epsilon$$

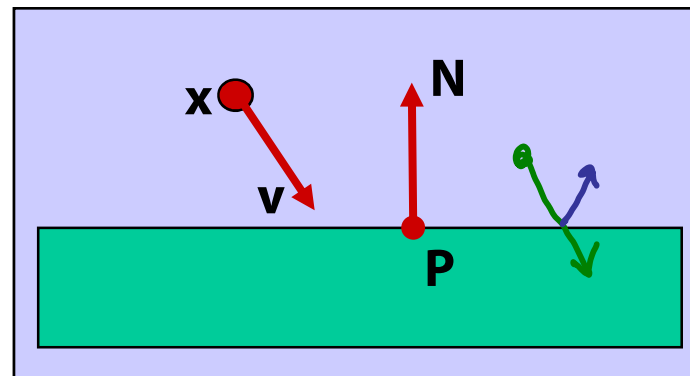
Also, the response may not be enough to bring a particle to the other side of a wall. In that case, detection should include a velocity check:

$$v \cdot N > 0 \Rightarrow \text{don't reflect}$$

More complicated collision response


Another solution is to modify the update scheme to:

- ◆ detect the future time and point of collision
- ◆ reflect the particle within the time-step



Particle frame of reference

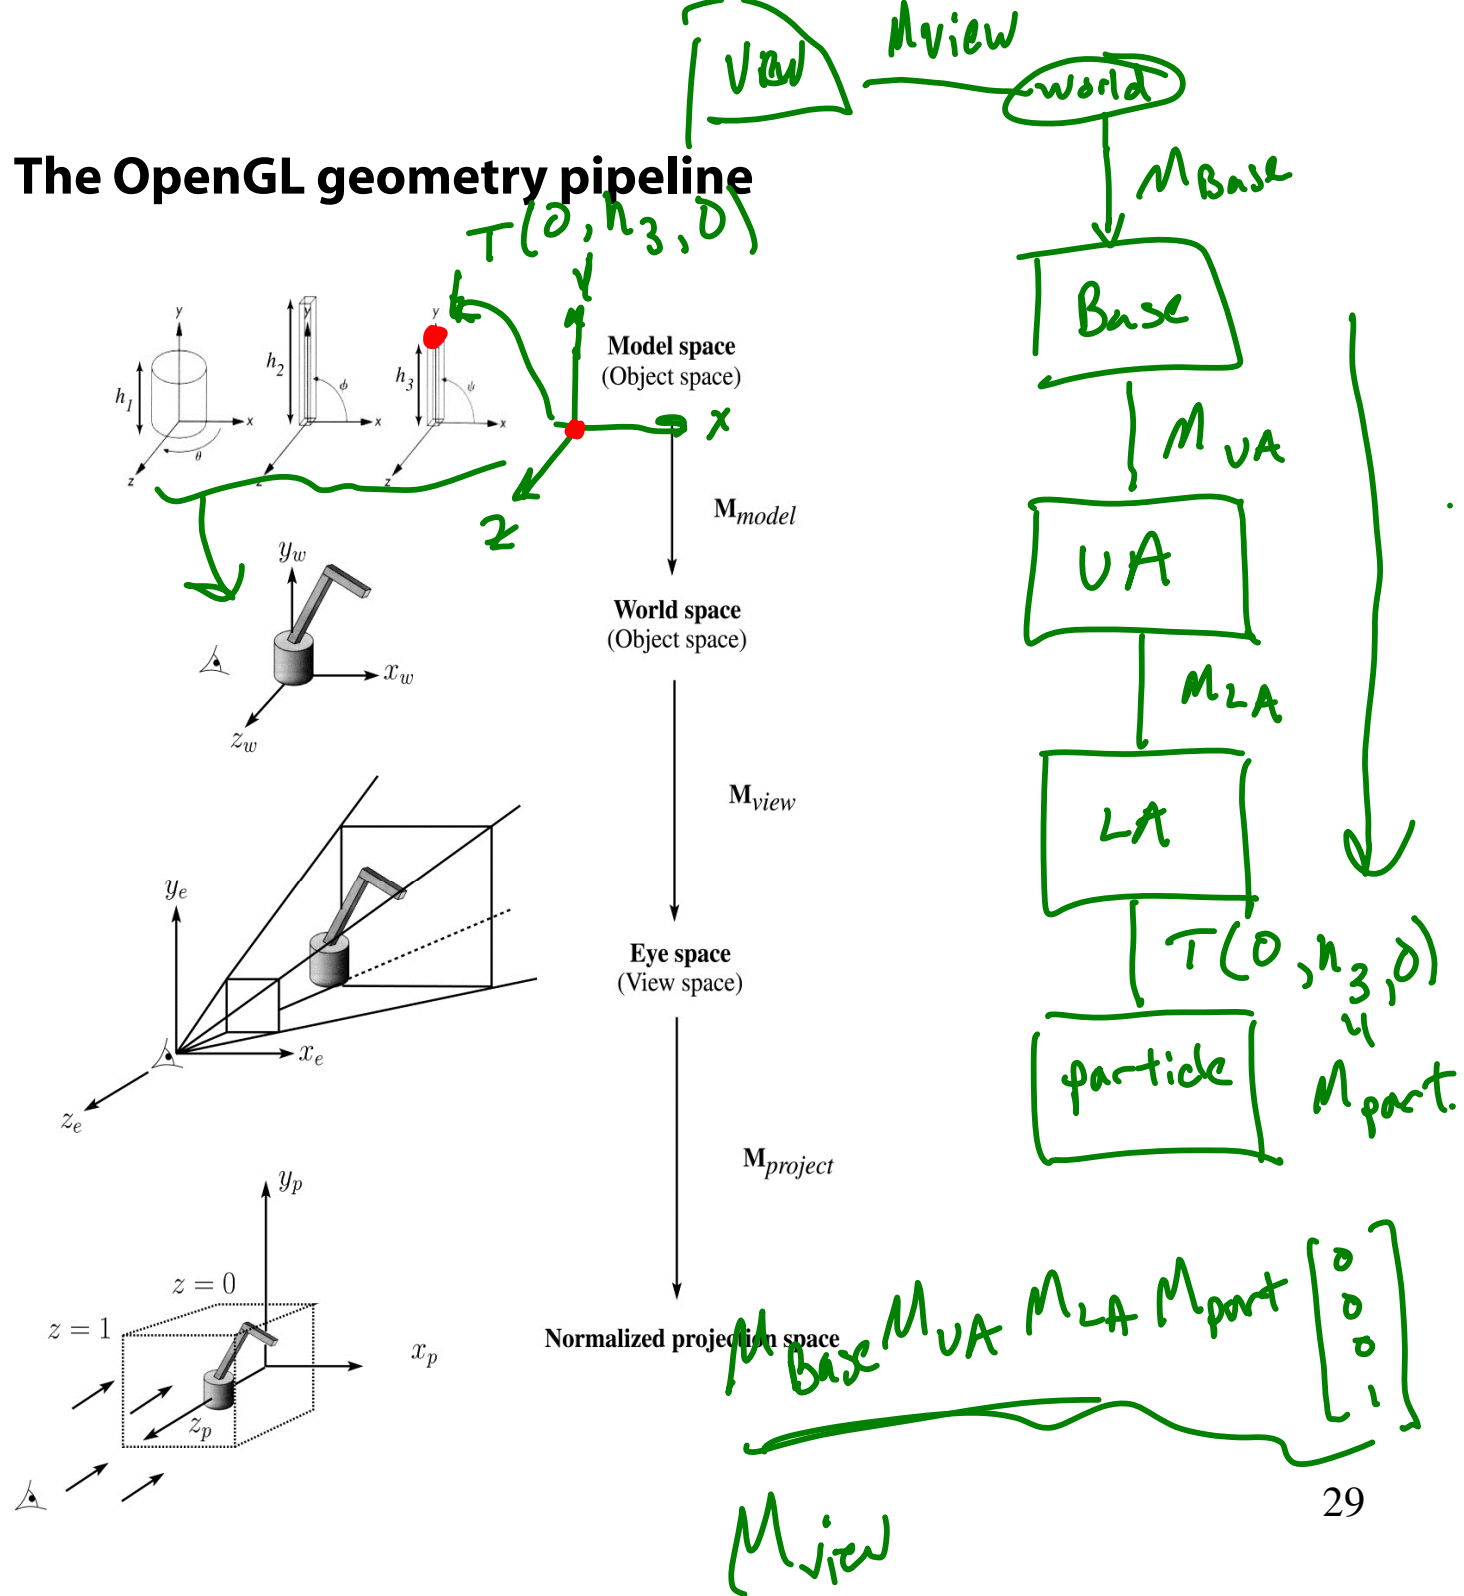
Let's say we had our robot arm example and we wanted to launch particles from its tip.



How would we go about starting the particles from the right place?

First, we have to look at the coordinate systems in the OpenGL pipeline...

The OpenGL geometry pipeline



Projection and modelview matrices

Any piece of geometry will get transformed by a sequence of matrices before drawing:

$$\mathbf{p}' = \mathbf{M}_{\text{project}} \mathbf{M}_{\text{view}} \mathbf{M}_{\text{model}} \mathbf{p}$$

The first matrix is OpenGL's GL_PROJECTION matrix.

The second two matrices, taken as a product, are maintained on OpenGL's GL_MODELVIEW stack:

$$\mathbf{M}_{\text{modelview}} = \mathbf{M}_{\text{view}} \mathbf{M}_{\text{model}}$$

Robot arm code, revisited

Recall that the code for the robot arm looked something like:

```
glRotatef( theta, 0.0, 1.0, 0.0 );  
base(h1);  
glTranslatef( 0.0, h1, 0.0 );  
glRotatef( phi, 0.0, 0.0, 1.0 );  
upper_arm(h2);  
glTranslatef( 0.0, h2, 0.0 );  
glRotatef( psi, 0.0, 0.0, 1.0 );  
lower_arm(h3);
```

All of the GL calls here modify the modelview matrix.

Note that even before these calls are made, the modelview matrix has been modified by the viewing transformation, \mathbf{M}_{view} .

Computing the particle launch point

To find the world coordinate position of the end of the robot arm, you need to follow a series of steps:

1. Figure out what \mathbf{M}_{view} is before drawing your model.

```
Mat4f Mview = glGetModelViewMatrix();
```

2. Draw your model and add one more transformation to the tip of the robot arm.

```
glTranslatef( 0.0, h3, 0.0 );
```

3. Compute $\mathbf{M}_{\text{model}} = \mathbf{M}_{\text{view}}^{-1} \mathbf{M}_{\text{modelview}}$

```
Mat4f particleXform = getWorldXform(Mview);
```

4. Transform a point at the origin by the resulting matrix.

```
Vec3f particleOrigin = particleXform * Vec3f(0,0,0);
```

Now you're ready to launch a particle from that last computed point!

Summary

What you should take away from this lecture:

- ◆ The meanings of all the **boldfaced** terms
- ◆ Euler method for solving differential equations
- ◆ Combining particles into a particle system
- ◆ Physics of a particle system
- ◆ Various forces acting on a particle
- ◆ Simple collision detection
- ◆ How to hook your particle system into the coordinate frame of your model

