



The long standing problem of PBD is that the stiffness depends on iteration count. Regardless of how you set the stiffness coefficients in PBD, given enough iterations it will converge to an infinitely stiff solution, effectively overriding the stiffness values.

For example in a cloth model, we have bending and stretch constraints. If we need to increase iterations to reduce stretching, we inadvertently increase resistance to bending. This makes it difficult to set up relative stiffness in simulation assets, and difficult to create reusable assets.

In addition, stiffness does not have a physical basis in PBD. It does not correspond to a traditional engineering stiffnesss. A connection to real-world physics would be nice in order to match real-world material parameters.

Finally, PBD does not have a concept of constraint force magnitude. Although constraints may be solved effectively, we have no idea how much work they have done. We need this information to drive force dependent effects such as haptic feedback, particle system spawning, motors and motor limits.



This example shows how PBD's behavior changes with iteration count. Ideally these two pieces of cloth should behave exactly the same, but we have much stiffer cloth at high iteration counts despite having used a low stiffness coefficient.



I want to start by deriving PBD from a more principled basis. We start here with our equations of motion, Newton's second law, F=Ma. Here F is derived from an energy potential U, which we define as a quadratic form in terms of the constraint functions C.

We prefer to work in terms of compliance which is defined as inverse stiffness. Here alpha is a diagonal, or block diagonal compliance matrix.



We can further re-write the equations of motion by introducing a new variable lambda. This is our Lagrange multiplier. Plugging it into the equations of motion we obtain the familiar constrained equations of motion, with the addition of this compliant alpha*lambda term in the second equation. This can be thought of as a regularization term that makes our constraints correspond to the energy potential we defined initially.



Now we have our equations of motion we need to solve them. This is a differential algebraic equation of order three (DAE3), a notoriously difficult problem to solve well. We take a simple approach of performing an implicit time-discretization on positions. Here we have introduced two new variables, tilde(x) which is sometimes called the inertial, predicted, or unconstrained position. This is the state the system would obtain if we integrated it forward in time without considering the constraints. tilde(alpha) is simply a time-step scaled version of compliance that we use to make the equations more compact.



Now we have our discretized equations we need to solve them. These are non-linear equations, and the standard way to solve non-linear equations is Newton's method. We will apply Newton's method, but first we re-label the equations as g() and h() respectively.



Newton's method involves first linearizing the equations to obtain a linear system that we solve for deltaX and deltaLambda. These are then applied directly to the system, and the process repeated until the deltas fall below some tolerance value.

This method works well, although it can be difficult to implement. Computing K involves the second derivatives (Hessians) of the constraint functions, and it may require a line-search or trust region method to converge. We now seek a simpler approach based on PBD, but to do that we need to make some approximations.



The first approximation is to replace K with M. This drops some terms on the order of dt², it may affect converge rate, but it doesn't change the solution of our Newton iteration. Sometimes this is called a quasi-Newton method.

The next approximation is to treat g(x), lambda) as zero for all iterations. We make two observations to justify this approximation. First, if we initialise the Newton iteration with x0 = tilde(x), and lambda0=0, this means g(x0), lambda0) is automatically zero, so the first Newton iteration is unchanged. In addition, g() typically remains small because it is related to the change in the constraint gradient direction. If constraints are linear, then g() is always zero. In practice g() is close to linear for small time-steps.



Some further intuition can be obtained by looking at this from a variational point of view. We can view implicit Euler as trying to find the closest point on the constraint manifold to the predicted or internal position, tilde(x). In contrast, PBD tries to find a close, but not necessarily the closest point on the manifold. Each iteration solves a slightly different minimization, starting from the current iterate, this leads to this 'curved path' back to the manifold. As soon as PBD finds a point on the manifold it stops, however there is no guarantee that this is the closest point. In practice the difference is small, but this is the consequence of our approximation that g() = 0.



We can now make a connection back to PBD by taking the Schur complement (just plug the expression for deltaX into the second equation). This gives a familiar constrained system, with the addition of the compliance terms. This is equivalent to a compliant version of Goldenthal et al's Fast Projection algorithm.

So far we have written things out in terms of the system as a whole, but one of the benefits of PBD is that it lets us work in terms of single constraints. So let's see what that looks like..



Here we have the original "scaling factor" from PBD, and the new deltaLambda term from XPBD. We can see that when alpha == 0 (infinite stiffness, or zero compliance) that the two equations are the same. From this we can see that s in PBD is actually an incremental Lagrange multiplier change for the infinitely stiff case. This makes sense, as it's the behavior we observe in PBD as we increase iteration count.

When alpha > 0 we have some additional terms in the numerator and denominator which act to regularise the constraint.



So here is the final XPBD algorithm. It is identical to the original PBD method with the addition of these two changes to the deltaLamba update, and then rather than throw away that value, we accumulate it in a per-constraint auxiliary variable.



We compare our method to a non-linear Newton method applied directly to the implicit equations of motion and find that it produces very close results. This graph shows the restoring force at the attachment for a chain swinging under gravity. XPBD approachs the reference solution given enough iterations. Here our Newton solver has used an LDLT decomposition as its linear solver backend.

Our MethodImage: Distribution of the state of the





SUMMARY

- Extend PBD to simulate well defined elastic and dissipative potentials
- Maintains simplicity and robustness of original method
- Requires storing a a single scalar auxiliary variable representing a Lagrange multiplier
- Can be used to compute total constraint force magnitude for e.g.: motor limits, haptic devices, etc
- Reduces to PBD when alpha=0, too few iterations means artificial compliance
- Contact, assume zero compliance

🕺 NVIDIA.



<section-header><list-item><list-item><list-item><list-item>

