ROUTING AND ACTION

MEMORANDUM

ROUTING

TO:(1) Mathematical Sciences Division (Wilkins, Leonard)

Report is available for review

(2) Proposal Files   Report No.:   Proposal Number: 68084-MA.4

DESCRIPTION OF MATERIAL

CONTRACT OR GRANT NUMBER: W911NF-16-1-0118

INSTITUTION: George Washington University

PRINCIPAL INVESTIGATOR: Evan Drumwright

TYPE REPORT: Final Report

DATE RECEIVED: 10/11/17 12:32AM

PERIOD COVERED: 4/1/16 12:00AM through 12/31/17 12:00AM

TITLE:

ACTION TAKEN BY DIVISION

(x) Report has been reviewed for technical sufficiency and IS [x] IS NOT [ ] satisfactory.

(x) Material has been given an OPSEC review and it has been determined to be non sensitive and, except for manuscripts and progress reports, suitable for public release.

(x) Perfomance of the research effort was accomplished in a satisfactory manner and all other technical requirements have been fulfilled.

(x) Based upon my knowledge of the research project, I agree with the patent information disclosed.

Approved by SSL\JAY.WILKINS on 11/29/17 8:31AM

ARO FORM 36-E

REPORT DOCUMENTATION PAGE

Form Approved OMB NO. 0704-0188

1. REPORT DATE (DD-MM-YYYY)	2. REPORT TYPE	3. DATES COVERED (From - To)
11-10-2017	Final Report	1-Apr-2016 - 31-Dec-2017

4. TITLE AND SUBTITLE	
W911NF-12-R-0012-03: Adaptive Integration of Nonsmooth Dynamical Systems	5a. CONTRACT NUMBER W911NF-16-1-0118
	5b. GRANT NUMBER
	5c. PROGRAM ELEMENT NUMBER 611102
6. AUTHORS	5d. PROJECT NUMBER
	5e. TASK NUMBER
	5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAMES AND ADDRESSES	8. PERFORMING ORGANIZATION REPORT NUMBER
George Washington University 2121 I Street NW Rice Hall Suite 601 Washington, DC 20052 -0086	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS (ES) U.S. Army Research Office P.O. Box 12211 Research Triangle Park, NC 27709-2211	10. SPONSOR/MONITOR'S ACRONYM(S) ARO 11. SPONSOR/MONITOR'S REPORT NUMBER(S) 68084-MA.4

12. DISTRIBUTION AVAILIBILITY STATEMENT

Approved for public release; distribution is unlimited.

13. SUPPLEMENTARY NOTES

The views, opinions and/or findings contained in this report are those of the author(s) and should not contrued as an official Department of the Army position, policy or decision, unless so designated by other documentation.

14. ABSTRACT

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	15. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			Evan Drumwright
UU	UU	UU	UU		19b. TELEPHONE NUMBER 202-994-1057

Standard Form 298 (Rev 8/98)
Prescribed by ANSI Std. Z39.18

RPPR Final Report
as of 29-Nov-2017

Agency Code:

Proposal Number: 68084MA **Agreement Number: W911NF-16-1-0118**

INVESTIGATOR(S):

Name: Evan Drumwright
Email: drum@gwu.edu
Phone Number: 2029941057
Principal: Y

Organization: **George Washington University**
Address: 2121 I Street NW, Washington, DC 200520086
Country: USA

DUNS Number: 043990498 EIN: 530196584
Report Date: 31-Mar-2018 Date Received: 11-Oct-2017
Final Report for Period Beginning 01-Apr-2016 and Ending 31-Dec-2017
Title: W911NF-12-R-0012-03: Adaptive Integration of Nonsmooth Dynamical Systems
Begin Performance Period: 01-Apr-2016 **End Performance Period:** 31-Dec-2017
Report Term: 0-Other
Submitted By: Evan Drumwright Email: drum@gwu.edu
Phone: (202) 994-1057
Distribution Statement: 1-Approved for public release; distribution is unlimited.

STEM Degrees: 1 **STEM Participants:** 4

Major Goals: The project seeks to produce accurate solutions of differential complementarity problems, which are used to model nonsmooth systems, through use of adaptive integration and simulation techniques. The proposal suggested using "guard
functions", functions that yield a safe, generally conservative integration step size such that the trajectory of the system from t0 to t0+dt lies strictly within the admissible region for all unilateral constraints. Put another way- we aim to avoid integrating over nonsmooth events. Through such avoidance, we can find intervals where the system trajectory is smooth. Then tools for smooth systems apply, e.g., higher order integration techniques and integration with error control.

Major goals, as proposed originally, consist of:

1. Incorporating guard functions when unilateral constraints are strictly satisfied (i.e., g(x,t) = 0); in this case, the guard function will return zero, so a simulation would be unable to proceed using a purely guard function strategy (i.e., one could use the regula falsi numerical method to identify an integration step dt > 0 such that g(x(t0+dt), t0+dt) >= 0, but the guarantee that the system lies strictly within the admissible region over [t0, t0+dt] is lost).

2. Determining properties of good guard functions. The smaller the discrepancy between the "true" safe integration step size and that produced by the guard function, the faster a DCP can be solved.

3. Integration with error tolerances. This will allow DCPs to be solved- hopefully efficiently- to a desired level of accuracy.

4. Benchmarks for verifiability. Since closed form solutions are not available for many nonsmooth systems, we seek to produce a set of benchmarks for testing the speed and accuracy of techniques for solving differential complementarity problems in our target application (multi-body dynamics simulation with hard contact).

Accomplishments: Our objectives, beyond this project even, have been to simulate contacting mechanical (rigid body) systems with quantifiable accuracy and with greater speed and accuracy than existing approaches.
This project fits into that goal by aiming to solve differential complementarity problems accurately. A significant subset of these problems corresponds to mechanical systems undergoing contact, for which the state of the art in methods for solving initial value problems for these systems are time stepping algorithms. The problem with time

stepping algorithms is that they are first-order accurate *at best*: they are "zeroth order accurate" until the integration time step becomes sufficiently small. The techniques representing the prior state of the art, piecewise differential algebraic equation (DAE) solutions, suffer from lack of solutions for modeling contacting rigid bodies, an extremely common application, in certain states. Those methods also crawl to a stop when tens of bodies begin interacting, as each contact causes the simulator to expend considerable effort to isolate the event in time.

The PI has studied compliant contact models, for which higher-order ODE integrators with error control can be applied and lack of solutions are not a problem, as a potential alternative; reliable sources indicate that this is the approach that MSC Adams takes. The ordinary differential equations resulting from application of these models are often computationally stiff, and thereby call for computationally expensive implicit integration techniques. Nevertheless, we attempted using a low-order implicit integrator (implicit Euler)- which has error-control capability- to integrate such stiff ODEs. Detailed results will be described and disseminated in the near future, but our conclusions will be summarized here: the problem space contains sufficiently many non-stiff problems such that implicit integrators appear to be an inefficient solution method, requiring orders of magnitude more computation for non-stiff aspects of the multibody dynamics problems. The implicit Euler integrator the PI developed during this process is of extremely high quality, is released under the Apache v2 open source license, supports automatic differentiation and first- and second-order finite differencing, and can be used easily in non-Drake event-driven simulators (this is not the case with, e.g., the existing RADAU5 integrator).

The inefficiency inherent in the compliant model- which introduces computational stiffness to an otherwise non-stiff problem- motivates the differential complementarity problem approach. We began this research project by devising an algorithm to isolate event times for time stepping methods for mechanical systems. While users of our software appreciated the ability to control the error of our simulations (compared to contemporary time stepping methods), the drop in simulation throughput was not. The message we learned, that compliance can be computationally beneficial, was surprising because compliant models of contact have a deservedly poor reputation with respect to computation in several scientific communities (games and animation, nonsmooth dynamics, robotics, among others). We found that Lacoursiere made a similar observation in his 2007 dissertation when he argued from a physical perspective that the computational effort expended to solve rigid contact constraints was wasted.

Major activities

- Devised an algorithm to isolate event times for "time stepping methods" (i.e., methods that discretize the differential equations in a set of differential algebraic equations) for mechanical systems undergoing contact with polyhedral geometries: polyhedra are capable of approximating any solid, so the resulting algorithm is very general. The resulting drop in throughput led us to our second discovery, listed immediately below.

- Developed a model for pseudo-rigid contact between bodies with polytopic geometries. Our idea was to minimize the number of throughput-throttling event isolations by introducing a thin, inertialess compliant layer with a rigid interface around the rigid body cores: events would only need to be isolated as bodies came into contact or left contact, rather than for any time in which the contact manifold changed. Work on this model is ongoing, but we have already proven that the kissing configurations for two or more contacting pseudo-rigid bodies can be determined by solving a convex optimization problem and that each intersecting configuration maps to exactly one deformed, kissing configuration.

- Developed verification benchmarks for contacting rigid bodies: Developing such benchmarks has been the thorniest problem we have encountered to-date. If contact forces vary by 5% from the "true" solution, the effect on the resulting motions may be insignificant or even undetectable to the human eye. Many game simulation libraries are capable of modeling phenomena, like falling buildings, with very compelling results. Yet robotic manipulation has proven far more challenging. Why? We postulate that the entropic scenarios (e.g., falling buildings) are far easier to simulate plausibly than disentropic scenarios (e.g., placing a peg into a hole). Nevertheless, we have constructed at least one challenging to model, yet simple to describe, benchmark problem: a 2D rod contacting a half space. According to Paul Mitiguy (author of SimWise 3D), SimWise is unable to model this example without energy gain.
We have developed time stepping (rigid contact) and compliant models of this problem (see http://drake.mit.edu/doxygen_cxx/classdrake_1_1examples_1_1rod2d_1_1_rod2_d.html); a piecewise DAE approach version of this problem exists in a development branch of our codebase and will be merged into the primary codebase within the next few months.

Significant results

Our significant results include open source (Apache 2 license) software with commercial quality, scientifically peer-reviewed documentation. These contributions include: 1) an event-driven simulator, allowing solutions to initial value problems for piecewise DAEs (see http://drake.mit.edu/doxygen_cxx/classdrake_1_1systems_1_1_simulator.html); 2) a series of ODE integrators (including a symplectic first-order integrator, an implicit first-order integrator, and a third-order error-controlled Runge Kutta integrator)- see http://drake.mit.edu/doxygen_cxx/classdrake_1_1systems_1_1_integrator_base.html ; 3) a solver for dynamical systems with arbitrary unilateral and bilateral constraints (the key component of the time stepping systems)- see http://drake.mit.edu/doxygen_cxx/classdrake_1_1multibody_1_1constraint_1_1_constraint_solver.html, http://drake.mit.edu/doxygen_cxx/structdrake_1_1multibody_1_1constraint_1_1_constraint_accel_problem_data.html, and http://drake.mit.edu/doxygen_cxx/structdrake_1_1multibody_1_1constraint_1_1_constraint_vel_problem_data.html); and 4) provision of dynamical systems with closed form or trusted numerical solutions (a mass-spring-system, the Robertson stiff ODE test, 2D rod contacting a halfspace with friction). The PI developed all of these codes in C++ by himself (albeit with help from code reviews).

Stated goals not met

After noting that all of the event isolation was throttling our simulation throughput- regardless of guard function efficiency- we stopped looking for more efficient guard functions. In place of that direction, we have been spending the last year drawing from our experience with compliant contact models- which have the ability to map two or more bodies intersecting in their undeformed configurations to deformed versions of those same bodies in a kissing configuration.

Training Opportunities: Simulating mechanical systems with contact constraints arising from interactions between rigid bodies with simple geometries (e.g., sphere vs. half-space, box vs. half-space, sphere vs. sphere, etc.) works remarkably well. In contrast, one of the greatest challenge applications in simulating Differential Complementarity Problems is that of simulating mechanical systems with contact constraints arising from interaction between rigid bodies with *arbitrary* geometries. The challenge arises because stepping over exact event (contact) times implies that an important physical invariant- zero volume interactions between rigid bodies- is violated.

Funding from this grant supported a M.S. student, Bjoern Cheng Yi, who investigated this problem. Since polyhedra are capable of approximating virtually any geometric shape, Bjoern began by studying whether there was a reasonable means to correct constraint violation for a subset of these shapes, convex polyhedral bodies. We set forth conditions on the convergence of the constraint correction (also known as "constraint stabilization") process: namely, the dot product between the signed distance of interpenetrating bodies and the Cartesian direction orthogonal to the constraints imposed on these bodies should always be positive. Bjoern showed that the algorithm/implementation, the Expanding Polytope Algorithm, used by multiple simulation software libraries for constraint stabilization violates this condition. Bjoern provided a correct algorithm based on the Minimum Translational Distance (MTLD) metric.

Bjoern's funding continued to result in fruitful research because computing MTLD requires a evaluating a quadratic number of features of two intersecting convex polyhedra, and that Cartesian product often has immense cardinality. A hexic number of features (O(n^3m^3)) is required to perform the same computation for non-convex polyhedra; thus, that process is intractable. Most importantly, my discussions with Bjoern resulted in the pseudo-rigid polyhedral model described elsewhere in this report.

Funding from this grant also supported two undergraduate students, Brad Canaday and John Shepherd, during Summer research projects. Brad was a Freshman and John was a sophomore at the time. I continue to communicate with John Shephard since leaving GWU- he is seeking to continue his education through entering a competitive Ph. D. program in Computer Science.

RPPR Final Report
as of 29-Nov-2017

Results Dissemination: Our primary form of dissemination since late 2016 has been through the Drake codebase (https://drake.mit.edu). Drake includes a state of the art multibody dynamics simulation in addition to other advanced software tools for dynamical systems analysis, estimation, prediction, and control. We hope ARO agrees with our assessment that this software is superior, in this particular field, to peer reviewed papers (which we have authored as well). Drake is the only open source engineering class multibody dynamics simulator, and there is no better proof that our methods work or do not than free, usable software.

All software contributions have been extensively documented and peer reviewed. For an example that demonstrates the level of quality achieved, (see http://drake.mit.edu/doxygen_cxx/classdrake_1_1examples_1_1rod2d_1_1_rod2_d.html#details). This example is that of a two dimensional rod contacting a halfspace, developed as a verification benchmark, with three solution methods: (1) time stepping (already in the Drake codebase); (2) a compliant model (already in the Drake codebase); and (3) a piecewise DAE approach (staged to be included in the Drake codebase in the next quarter).

This project has also resulted in two student authored research publications. [1] treats other sources of computational stiffness (beside the typical one of contact) in the particular problem of dynamic robotic simulation. [2] describes how we were able to use the robust simulations that result from the error-controlled time stepping method to interactively design running robots.

[1] John Shepherd, Samuel Zapolsky, and Evan M. Drumwright, “Fast multi-body simulations of robots controlled with error feedback”, IEEE International Conference on Simulation and Programming for Autonomous Robots (SIMPAR), 2016.

[2] Bradley Canaday, Samuel Zapolsky, Evan M. Drumwright, “Interactive, Iterative Robot Design”, IEEE International Conference on Robotics and Automation (ICRA). Singapore, May 2017.

Finally, we have two invited, peer reviewed book chapter submissions in progress for the Springer Robotics Encyclopedia. They are titled "Contact Dynamics" and "Dynamic Simulation", and both will acknowledge support from ARO.

Honors and Awards: Nothing to Report

Protocol Activity Status:

Technology Transfer: Nothing to Report

CONFERENCE PAPERS:

Publication Type: Conference Paper or Presentation **Publication Status:** 4-Under Review
Conference Name: Workshop on the Algorithmic Foundations of Robotics
Date Received: 26-Aug-2016 Conference Date: 20-Dec-2016 Date Published: 20-Dec-2016
Conference Location: Berkeley, CA
Paper Title: True Rigidity: Interpenetration-free Multi-Body Simulation with Polytopic Contact
Authors: Evan Drumwright
Acknowledged Federal Support: **Y**

Publication Type: Conference Paper or Presentation **Publication Status:** 1-Published
Conference Name: IEEE Intl. Conf. on Simulation and Programming for Autonomous Robots
Date Received: 13-Sep-2017 Conference Date: 03-Dec-2016 Date Published: 03-Dec-2016
Conference Location: San Francisco, CA
Paper Title: Iterative, Interactive Robot Design
Authors: Bradley Canaday, Samuel Zapolsky, Evan Drumwright
Acknowledged Federal Support: **Y**

RPPR Final Report
as of 29-Nov-2017

Publication Type: Conference Paper or Presentation **Publication Status:** 0-Other
Conference Name: IEEE Intl. Conf. on Simulation and Programming for Autonomous Robots (SIMPAR)
Date Received: 13-Sep-2017 Conference Date: 03-Dec-2016 Date Published: 03-Dec-2016
Conference Location: San Francisco
Paper Title: Fast Multi-Body Simulations of Robots Controlled with Error Feedback
Authors: John Shepherd, Samuel Zapolsky, and Evan Drumwright
Acknowledged Federal Support: **Y**

Elasto-polyhedra: modeling contact between bodies with polyhedral geometries

Part I: Motivation and Model

Evan Drumwright and Sean Curtis

This is a natural extension of an idea I've been quietly working on for many years. Since 2005, I thought it was critical for roboticists to have a robust dynamical simulator (graceful degradation in accuracy).

I was funded by ARO to work on the problem of computing accurate solutions to multibody problems with contacts. This research is the main product.

Segue: how did I get into this problem?

Part IA: The motivation

Started working in simulation after attempting to use software like this to test software running on my robots.

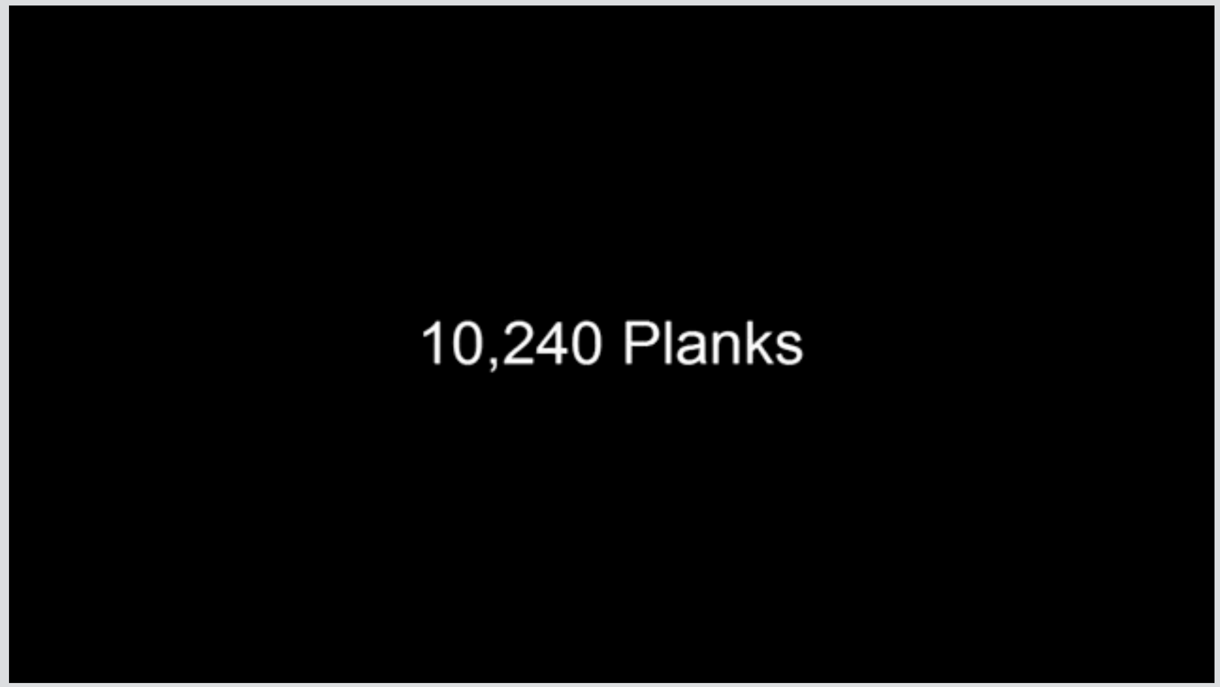

Started working in simulation after attempting to use software like this to test software running on my robots.

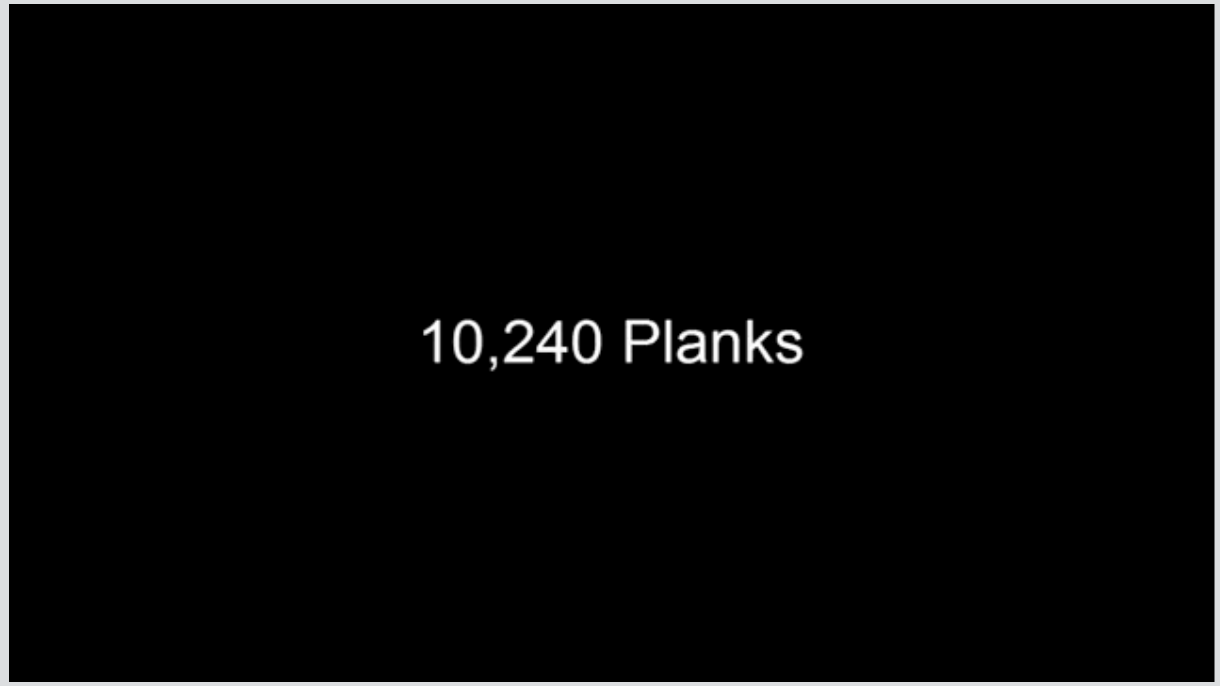

The libraries that produce these beautiful results have failed at simulating robotic manipulation.

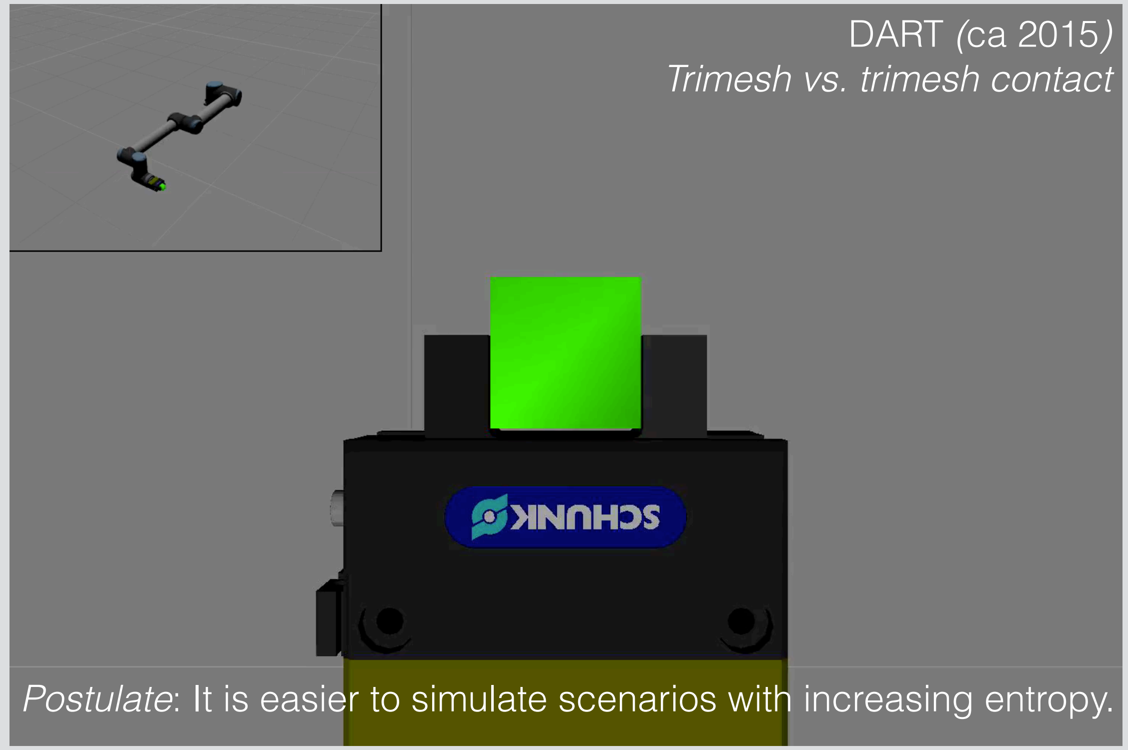

Manipulation is *disentropic* (it’s like Maxwell’s demon).

The libraries that produce these beautiful results have failed at simulating robotic manipulation.

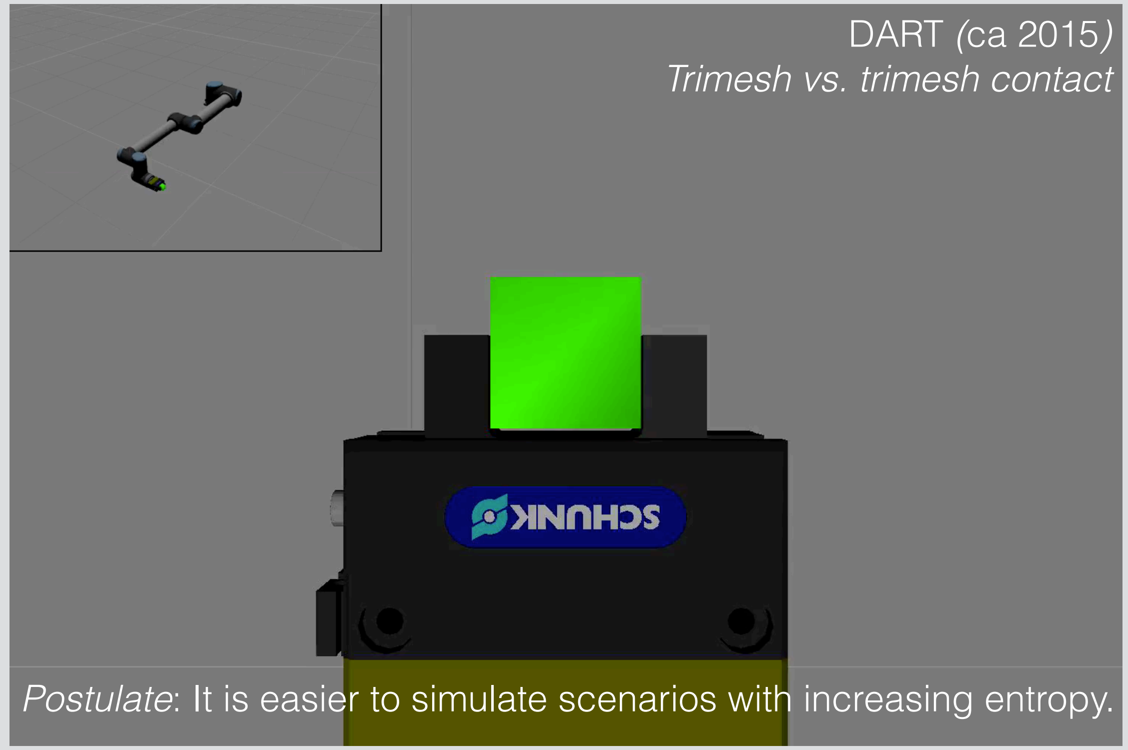

Postulate: It is easier to simulate scenarios with increasing entropy.

Manipulation is *disentropic* (it’s like Maxwell’s demon).

Some of my first simulations were of prehensile manipulation.

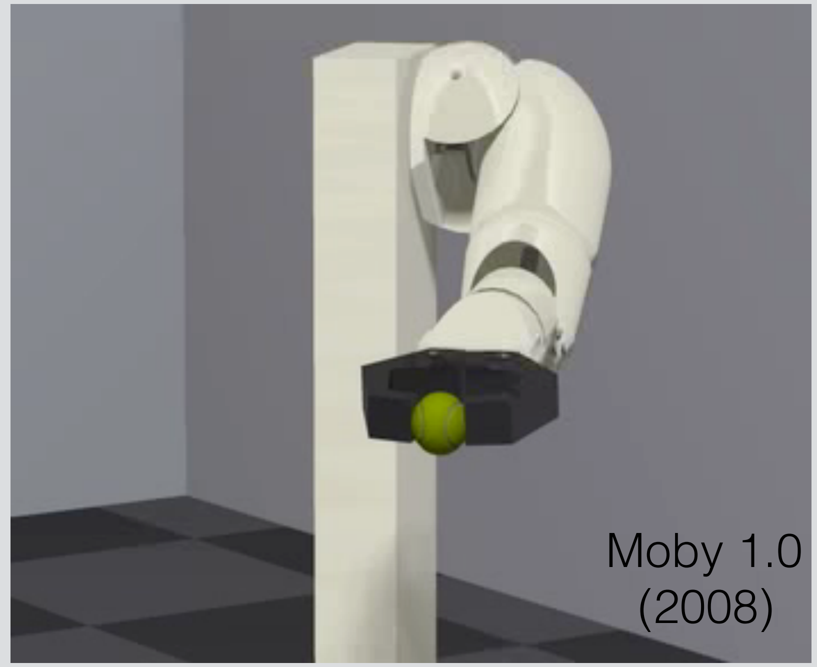

Segue: why were my simulations successful where these others failed?

Some of my first simulations were of prehensile manipulation.

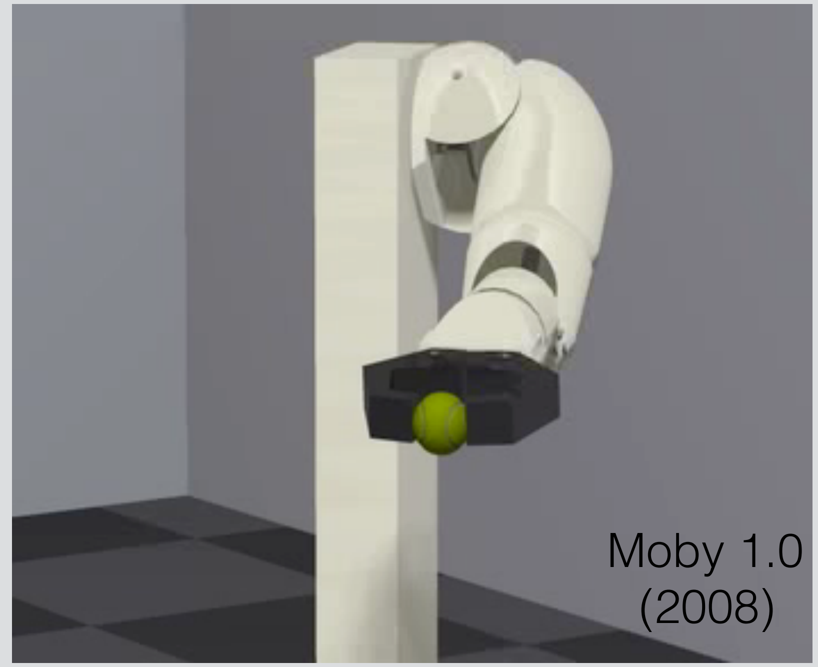

Segue: why were my simulations successful where these others failed?

Possibilities for why those libraries fail to simulate prehensile manipulation:

Possibilities for why those libraries fail to simulate prehensile manipulation:

1. They use low-order integrators.

My thoughts on low order accuracy in robotic simulation

- Good control policies use less-than-perfect models of their plants (robots)
- and the simulation can be expected to change little from step to step when going from first order to a higher order.

These factors are mitigated in the presence of models susceptible to *grazing bifurcations.*

Point #1: if your controller isn't able to control a robot just because higher order accuracy effects in the model aren't reproduced, you probably have a brittle controller.
Point #2: if your controller isn't good enough to handle differences as small as those coming from first-order vs. second-order would expected to be, you probably have a brittle controller. Also, "step to step" must be small, as robot controllers typically are expected to run at a high rate (< 0.01s).

DART has reduced coordinates and their grasping example broke.

Possibilities for why those libraries fail to simulate prehensile manipulation:

1. They use low-order integrators.

DART has reduced coordinates and their grasping example broke.

TOYOTA
RESEARCH INSTITUTE

Possibilities for why those libraries fail to simulate prehensile manipulation:

1. ~~They use low-order integrators.~~

2. Joints are allowed to separate.

DART has reduced coordinates and their grasping example broke.

TOYOTA RESEARCH INSTITUTE

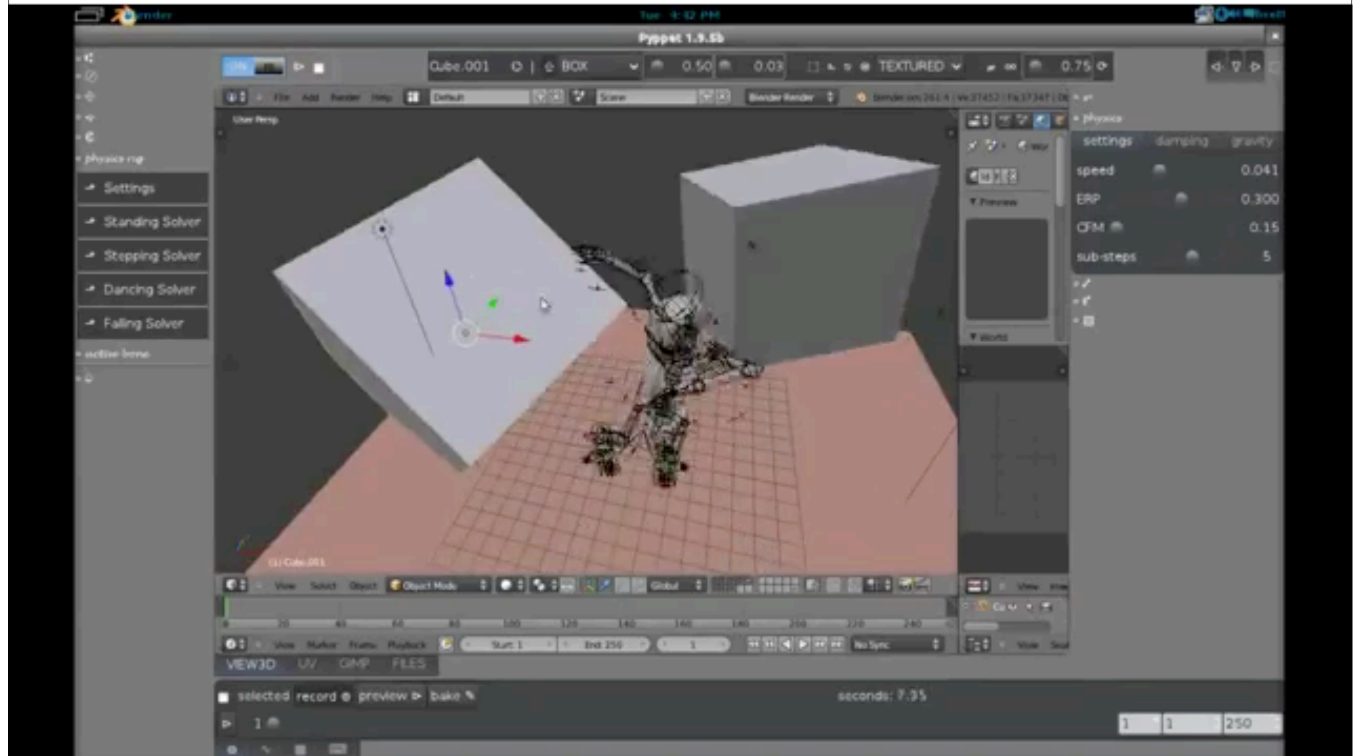

Blender (using Bullet)

DART has reduced coordinates and their grasping example broke.

Possibilities for why those libraries fail to simulate prehensile manipulation:

1. ~~They use low-order integrators.~~

2. Joints are allowed to separate.

Possibilities for why those libraries fail to simulate prehensile manipulation:

1. ~~They use low-order integrators.~~

2. ~~Joints are allowed to separate.~~

3. The libraries are determining contact data incorrectly.

C*ontact determination:* computing contact points, normals, signed distance, curvature, etc.

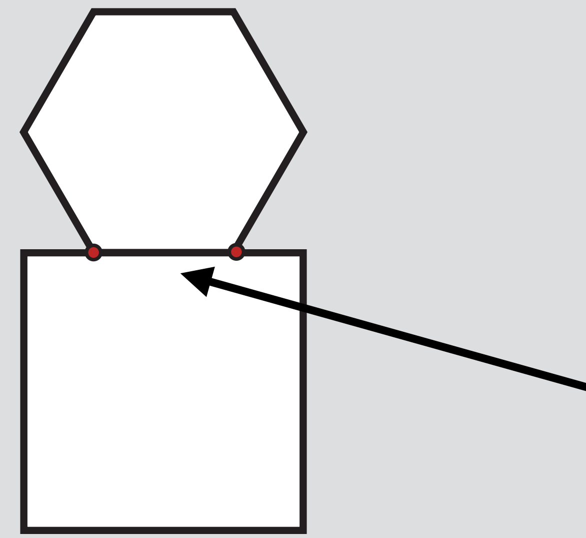

How hard can calculating *that* be?

* This requires us to step right to time of contact, so we can examine bodies in a “kissing” configuration.
* Most newer simulators do not do this: very hard to implement and have low throughput with many events.
* Slowly being implemented in Drake.

Time stepping discretizes dynamics, so all events are treated as occurring at one side of the interval.

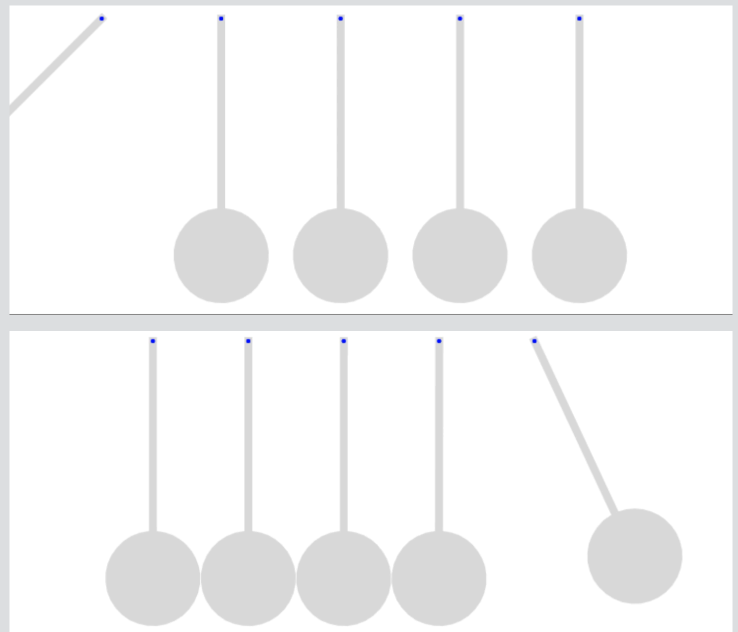

The true model (a modified Newton's cradle): impacts are closely spaced in time.

The time stepping version: one simultaneous impact between all five balls.

Try to ignore the spacing between the balls- it was a necessary artifact.

Time stepping discretizes dynamics, so all events are treated as occurring at one side of the interval.

The true model (a modified Newton's cradle): impacts are closely spaced in time.

The time stepping version: one simultaneous impact between all five balls.

Try to ignore the spacing between the balls- it was a necessary artifact.

At the time of contact, the configurations of the balls look like this:

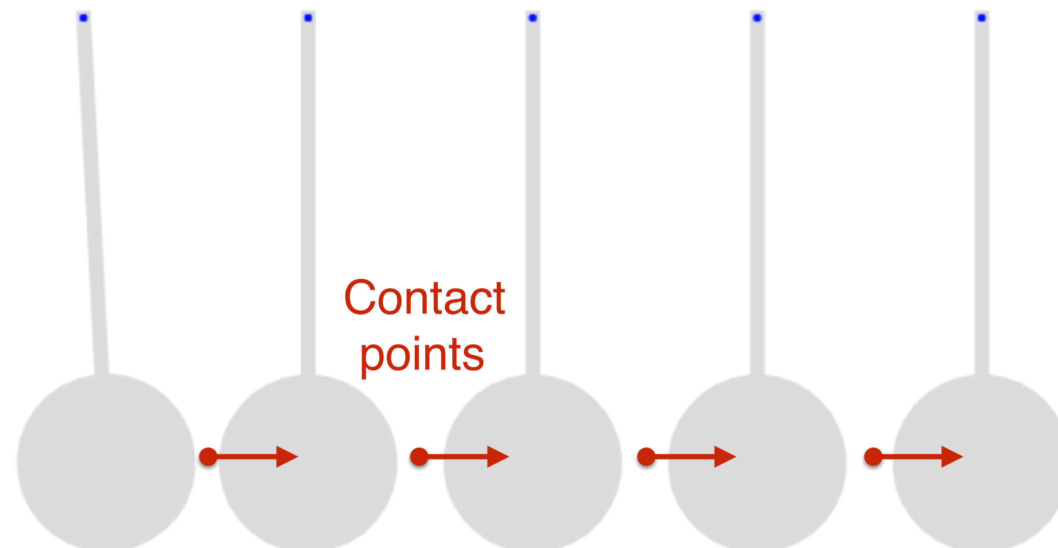

When bodies are sufficiently close or intersecting, constraints are introduced.

Popular "time stepping" methods are:

- 1st order (without events)
- Can model rigid contact and "softened" rigid contact
- Simple to implement.
- Formulated to work without event isolation.

Negatives:
0th order accurate with events.
No error control (can "blow up").

Still a nice tool to have in our toolbox.

Segue: that brings us back to the contact determination problem

Ability to model contact between "bodies" of various shapes is key to successful simulation of manipulation.

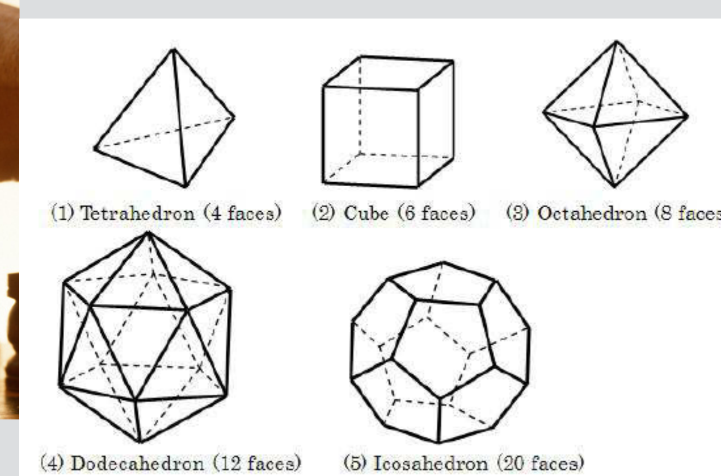

Polyhedra can approximate nearly any solid.

TOYOTA RESEARCH INSTITUTE

Time stepping requires us to answer: Which way should these bodies be separated?

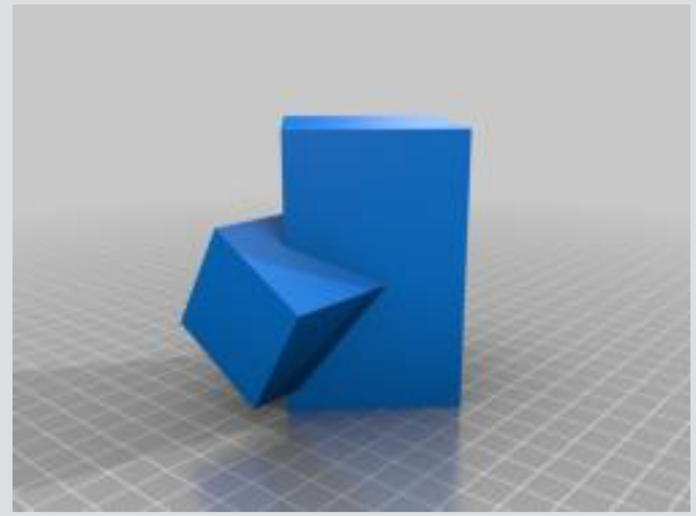

Answer: when the bodies are rigid, looking at a snapshot like this isn't sufficient. When bodies are compliant, we can try to find configurations that yield this picture.

How it's now done with time stepping: the gap function (computing signed Euclidean distances between bodies in arbitrary configurations).

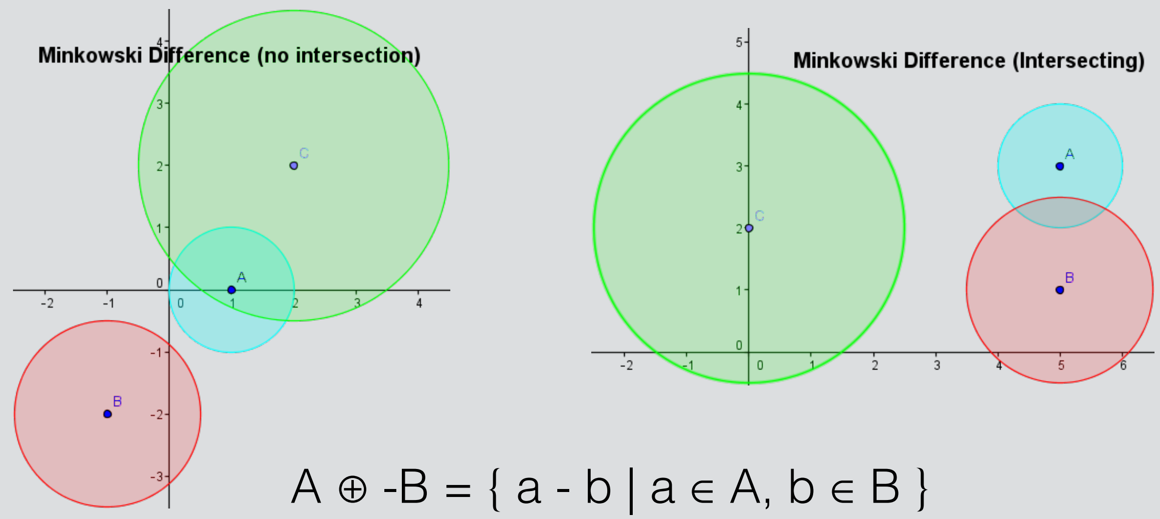

$$A \oplus -B = \{ a - b \mid a \in A, b \in B \}$$

This Minimum Translational Distance (MTLD) can be computed by the Minkowski Sum

This is a mathematical formulation, not a mechanical one.

Computing signed distance between *convex* polyhedra requires $O(m^{3/4+\varepsilon} n^{3/4+\varepsilon} + m^{1+\varepsilon} + n^{1+\varepsilon})$ time (for some $\varepsilon > 0$).

Developers and researchers have used heuristics instead.

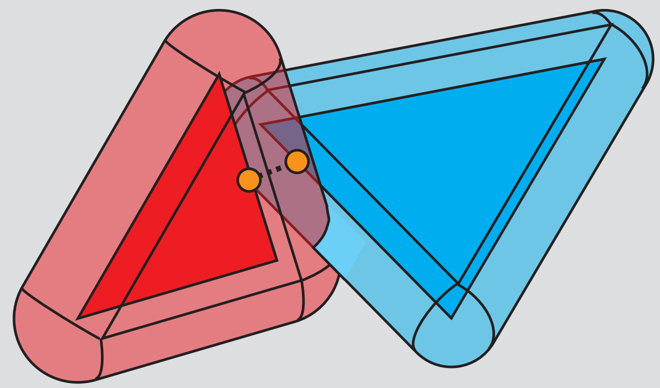

One heuristic:
use "thick" triangles in the triangle mesh.*

* Hauser, K, *Robust contact generation for robot simulation with unstructured meshes*, Springer Tracts in Advanced Robotics, vol 114 (2016)

O(1) time if bodies are disjoint.

Why use heuristics before understanding their impact on the true solution?

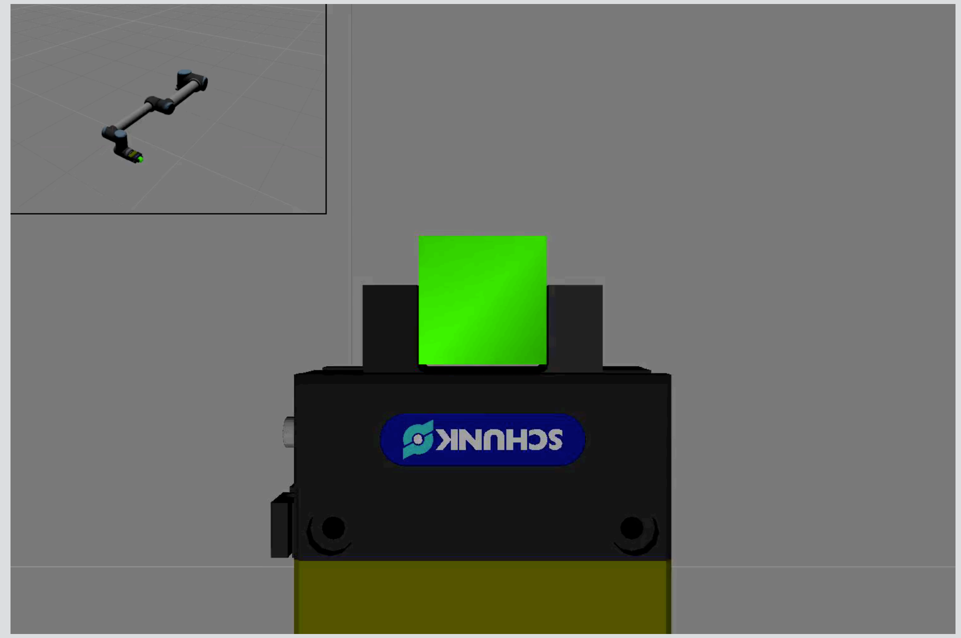

This is why I believe DART failed to robustly simulate grasping.

Segue: a couple of years ago, I wanted to try a different approach. My first approach to this problem tried to prevent interpenetration (because rigid bodies don't interpenetrate). I was funded to do this by ARO early last year.

Why use heuristics before understanding their impact on the true solution?

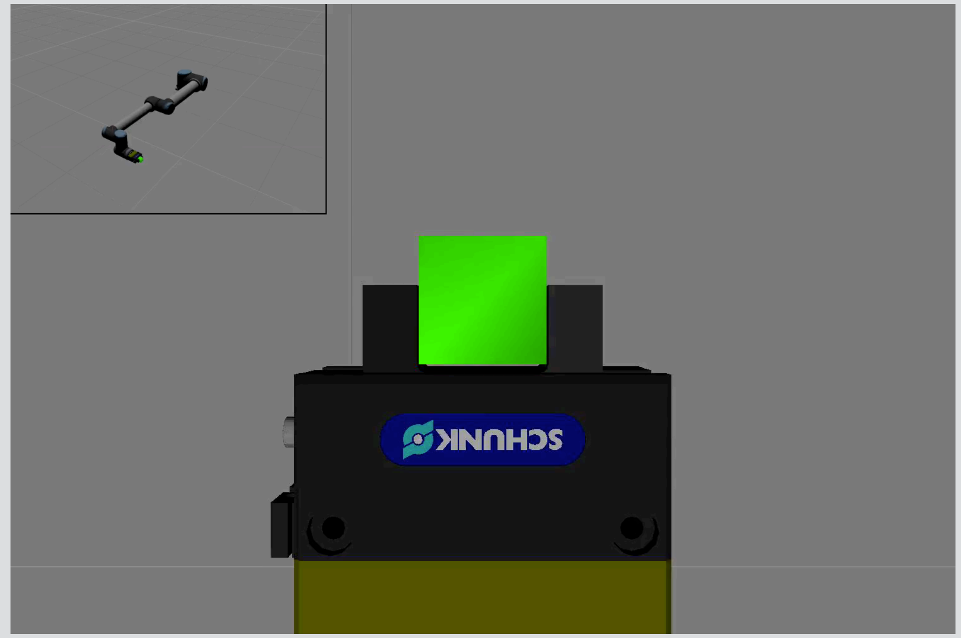

This is why I believe DART failed to robustly simulate grasping.

Segue: a couple of years ago, I wanted to try a different approach. My first approach to this problem tried to prevent interpenetration (because rigid bodies don’t interpenetrate). I was funded to do this by ARO early last year.

Time stepping with event isolation

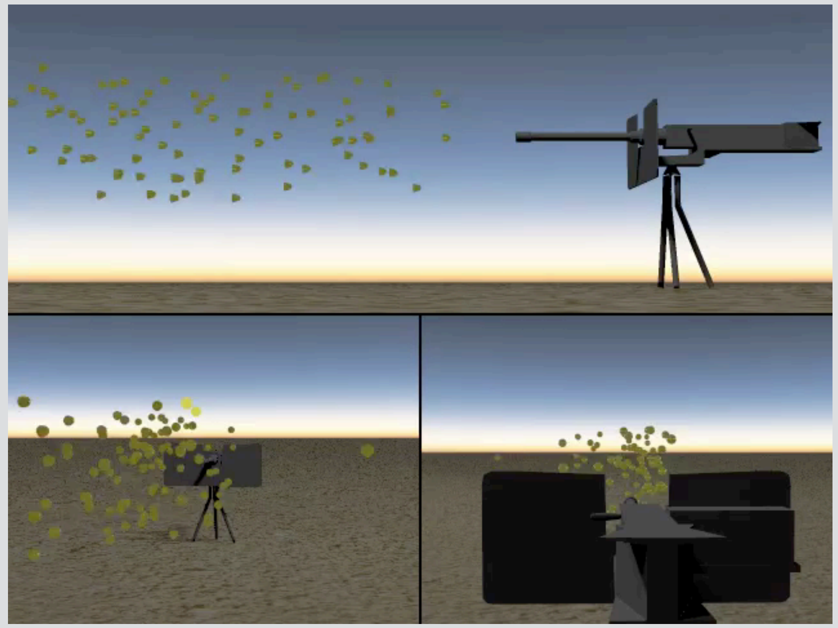

Using the time stepping equations, step directly to changes in the contact manifold.

Nice feature of this approach: it prevents “tunneling”

E. Drumwright. True rigidity: Interpenetration-free multi-body simulation with polytopic contact. arXiv, 2016.

Segue: all of the event processing can really limit simulation throughput.

Time stepping with event isolation

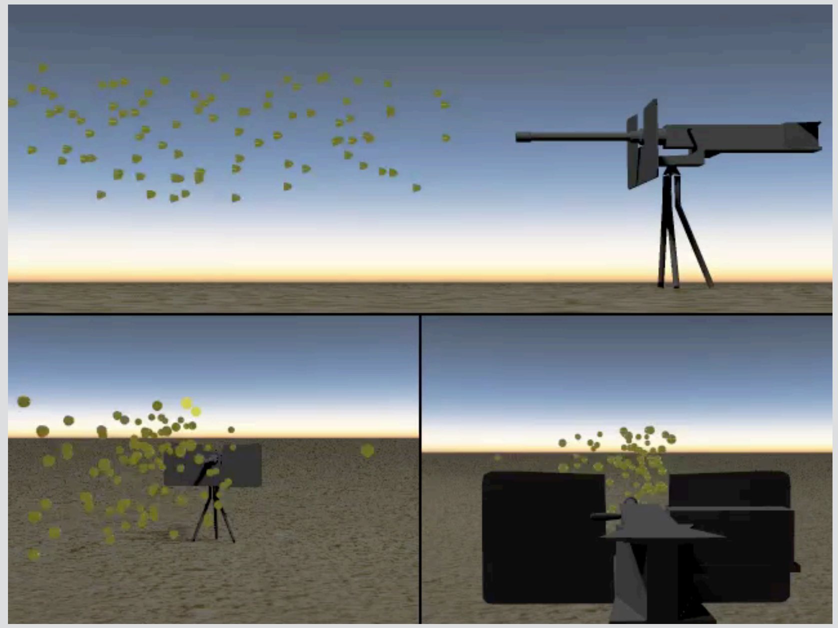

Using the time stepping equations, step directly to changes in the contact manifold.

Nice feature of this approach: it prevents "tunneling"

E. Drumwright. True rigidity: Interpenetration-free multi-body simulation with polytopic contact. arXiv, 2016.

Segue: all of the event processing can really limit simulation throughput.

If events are isolated, you can do time stepping with error control.

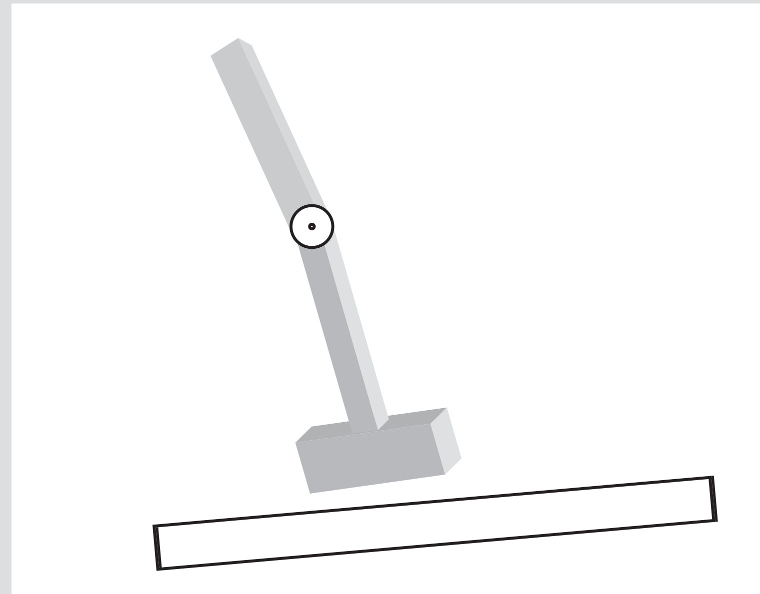

Downside: for robotic applications, contact manifold can change frequently.

Elaboration on the downside: the event processing can really limit simulation throughput- even when there aren't many bodies. Still a huge pain to keep track of everything.

Segue:
How can we speed this up? Use compliance (compliance has been trendy in robotics lately, even though it's hard to model).

Part IB: The model

Method for model pseudo-rigid bodies with convex polyhedral geometries that does not require isolating event times.

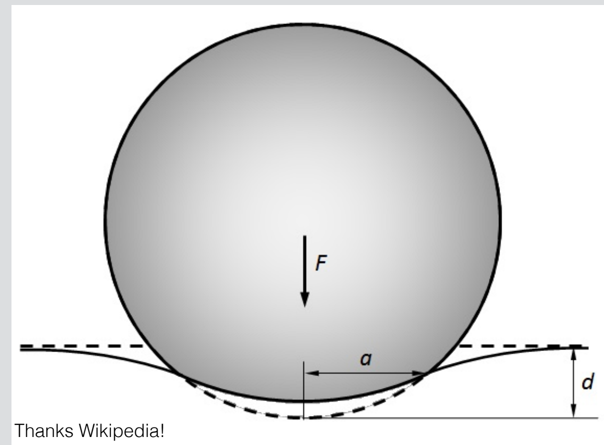

Analogy: Hertzian contact maps intersection between undeformed ball and halfspace to kissing contact between deformed versions.

Like Hertz, our contact forces will be based on a mechanical model (not a mathematical one like MTLD).

Rules of the game:

- Each configuration of intersecting undeformed bodies must map to a unique kissing configuration using deformation.
- We must be able to construct an (imperfect) physical specimen of the ideal mechanism:
  - Mass and lengths to the limit zero are fair game.
  - Mass and lengths to the limit infinity are not.
  - Mechanism vs. environment and mechanism vs. mechanism geometric intersections cannot be ignored.

Pseudo rigid model (*elasto-polyhedron*): rigid core + pistons + spring-dashpot + rigid interface with negligible inertia

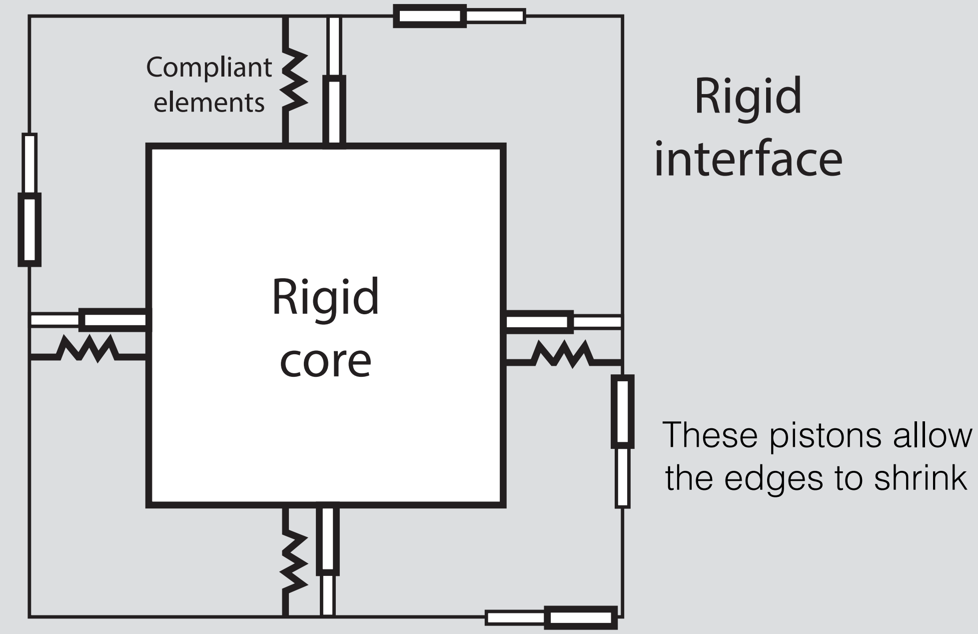

We're starting in 2D to make things easier. 3D will be introduced gradually.

Pistons have hard limits.

A pseudo-rigid body compressed by a rigid body at a face.

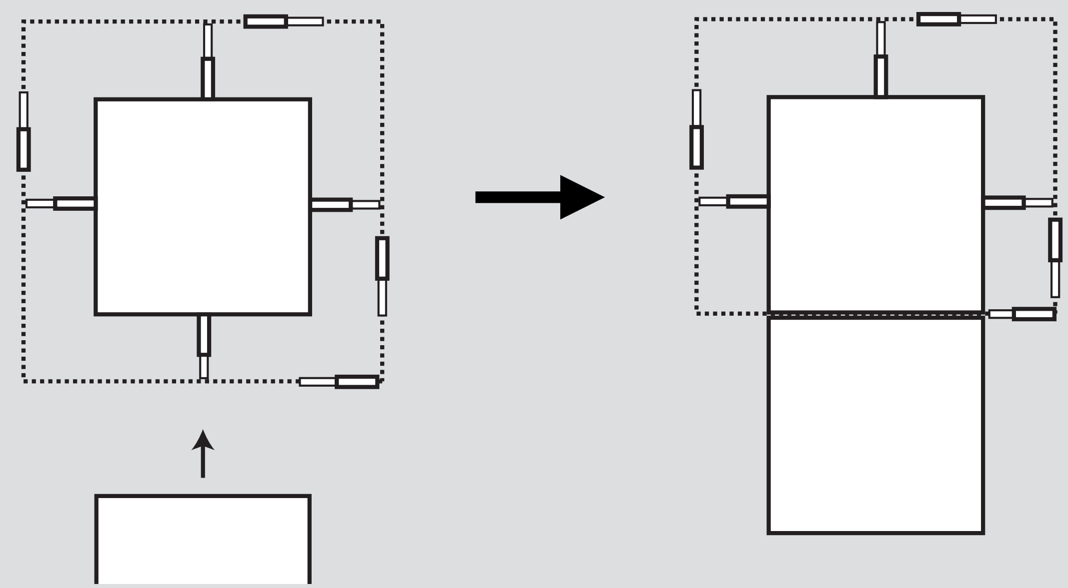

First I’m showing you contacts with purely rigid bodies so that the interface compression is predictable.

A pseudo-rigid body compressed by a rigid body at a vertex.

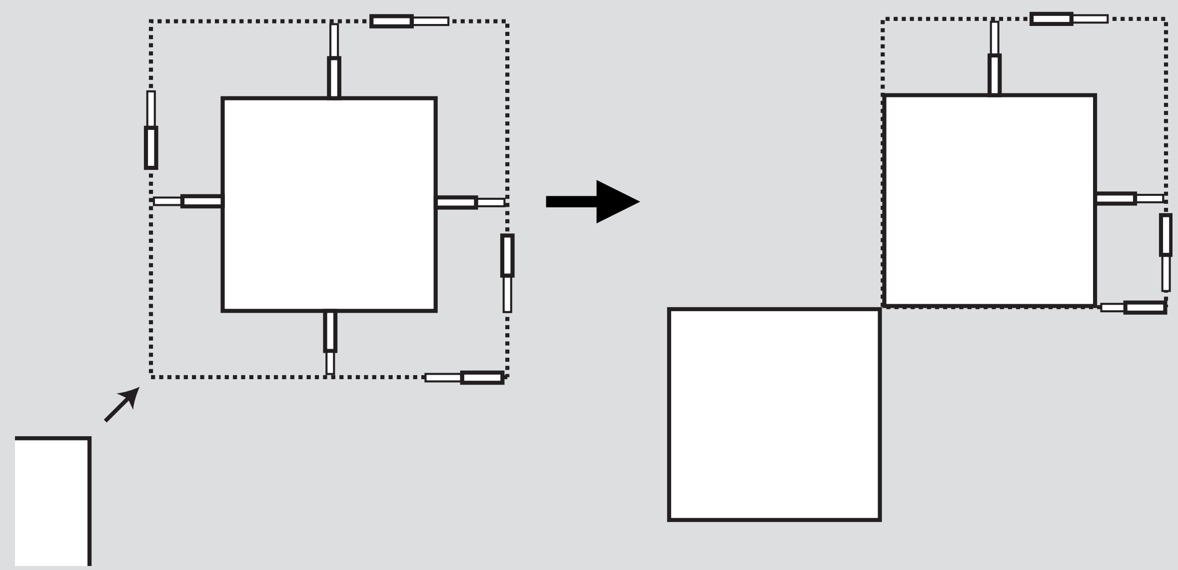

A demo

How to keep the solids closed as faces compress and expand?

Rejected idea: apply spring forces to pull the faces closed.

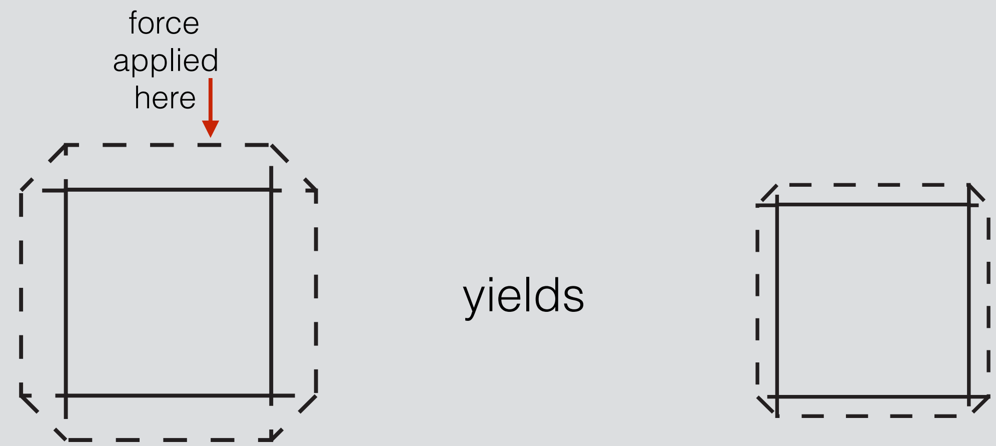

That approach yields global deformation from applied forces.

Segue: instead

The model does no work to “seal” itself as the volume changes.

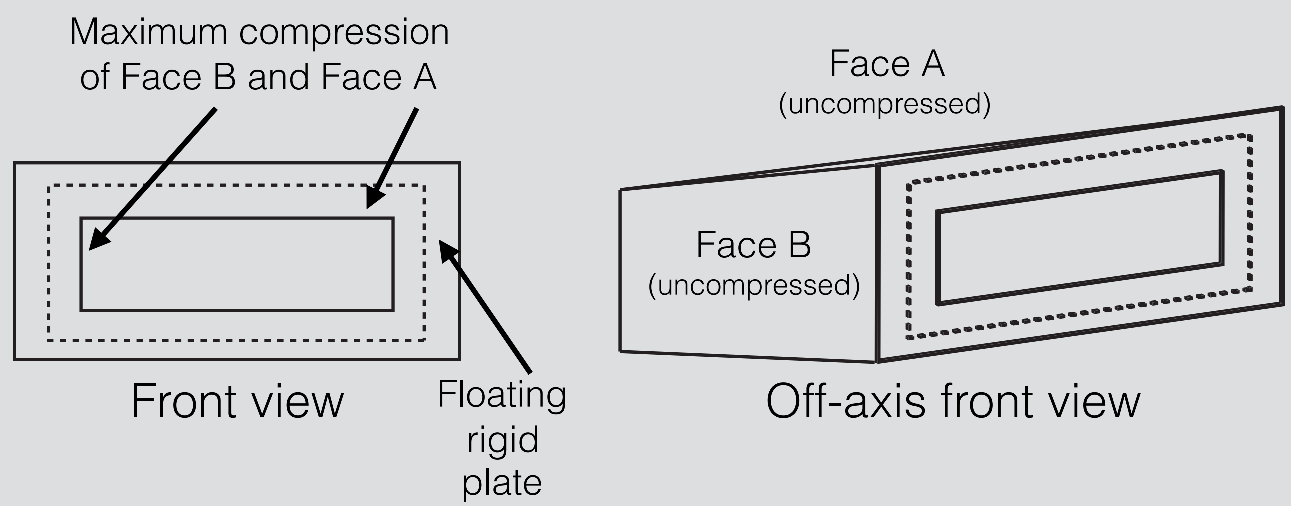

Here is how a physical implementation of the idealized model would look.

Look past my poor 3D drawing

We do not allow coplanar faces.

(i.e., they must be merged)

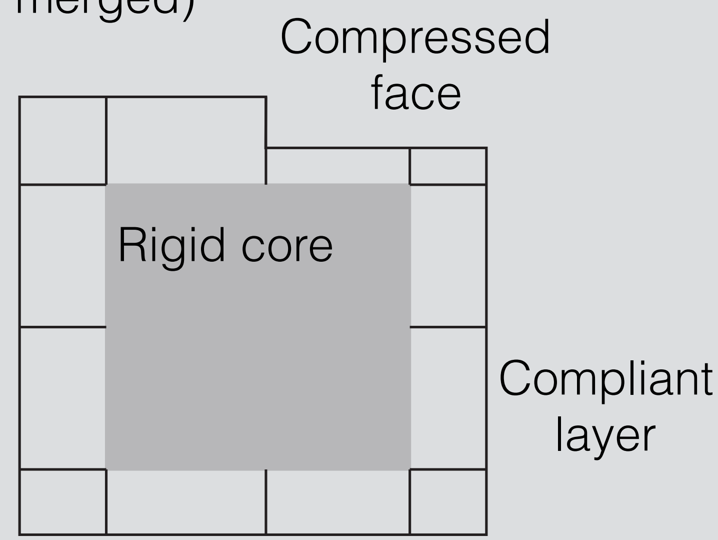

We do not allow coplanar faces.

(i.e., they must be merged)

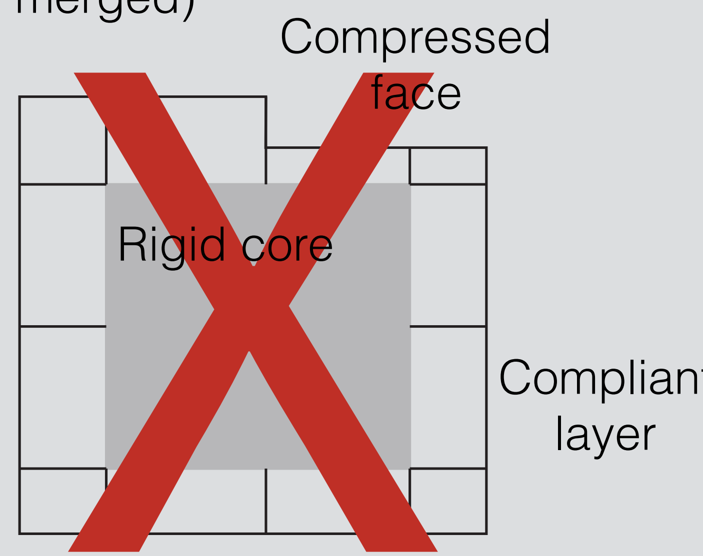

Avoiding coplanar faces ensures signed distances between pseudo-rigid bodies are convex.

Given the kissing deformed configuration, computing contact forces is straightforward.

- Option 1: Compute all forces from compliant elements, *then* compute constraint forces
- Option 2: Use “soft constraints” with piston compressions corresponding to constraint violation.

Option #1 appears to allow us to avoid Painleve type effects too!

One or more pairs of pseudo-rigid bodies contacting.

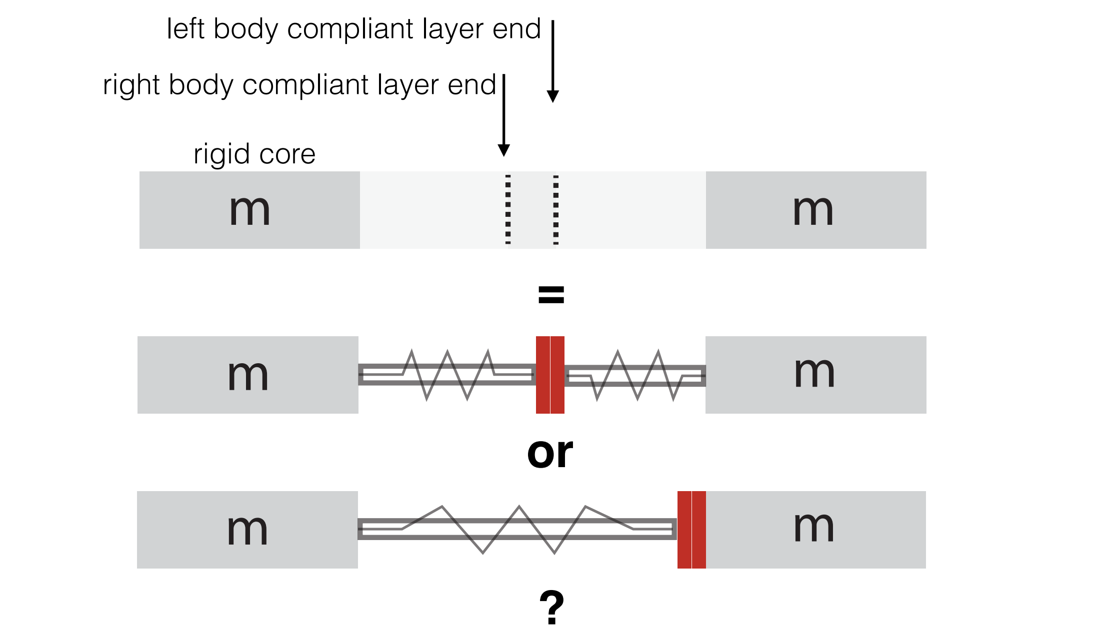

Given the locations of the two masses, where do the interfaces meet?

This principle is for fully deformable models

Principle of minimum potential energy dictates how far the pistons will compress

$$\min_{\mathbf{x}} \; \frac{1}{2}\sum_{i=1}^{n} k(x_{r_i} - x_i)^2$$

$$\begin{aligned} \text{s.t.} \;\; & 0 \leq x_i \leq x_{r_i} \\ & \phi_{i,j} \geq 0, \quad \forall i,j \in \mathcal{S} \text{ s.t. } i \neq j \end{aligned}$$

This principle is for fully deformable models

Principle of minimum potential energy dictates how far the pistons will compress

Piston extension is the optimization variable

$$\min_{\mathbf{x}} \ \frac{1}{2}\sum_{i=1}^{n} k(x_{r_i} - x_i)^2$$

$$\text{s.t.} \ 0 \leq x_i \leq x_{r_i}$$

$$\phi_{i,j} \geq 0, \quad \forall i,j \in \mathcal{S} \ \text{s.t.} \ i \neq j$$

This principle is for fully deformable models

Principle of minimum potential energy dictates how far the pistons will compress

Maximum piston extension

Piston extension is the optimization variable

$$\min_{\mathbf{x}} \frac{1}{2}\sum_{i=1}^{n} k(x_{r_i} - x_i)^2$$

$$\text{s.t. } 0 \leq x_i \leq x_{r_i}$$

$$\phi_{i,j} \geq 0, \quad \forall i, j \in \mathcal{S} \text{ s.t. } i \neq j$$

This principle is for fully deformable models

Principle of minimum potential energy dictates how far the pistons will compress

$$\min_{\mathbf{x}} \frac{1}{2}\sum_{i=1}^{n} k(x_{r_i} - x_i)^2$$

$$\text{s.t. } 0 \leq x_i \leq x_{r_i}$$

$$\phi_{i,j} \geq 0, \quad \forall i,j \in \mathcal{S} \text{ s.t. } i \neq j$$

Piston extension is the optimization variable ($\mathbf{x}$)

Maximum piston extension (x_i)

Signed distance between bodies *i* and *j* ($\phi_{i,j}$)

This principle is for fully deformable models

The objective function is clearly strictly convex.

$$\min_{\mathbf{x}} \ \frac{1}{2}\sum_{i=1}^{n} k(x_{r_i} - x_i)^2$$

$$\text{s.t. } 0 \leq x_i \leq x_{r_i}$$
$$\phi_{i,j} \geq 0, \quad \forall i, j \in \mathcal{S} \text{ s.t. } i \neq j$$

Piston extension constraints are clearly convex constraints.

$$\min_{\mathbf{x}} \; \frac{1}{2}\sum_{i=1}^{n} k(x_{r_i} - x_i)^2$$

$$\text{s.t.} \; \boxed{0 \leq x_i \leq x_{r_i}}$$

$$\phi_{i,j} \geq 0, \quad \forall i, j \in \mathcal{S} \text{ s.t. } i \neq j$$

Sufficient condition for the convexity of the pairwise signed distance constraints

Decreasing extension of one or more pistons must not decrease signed distance between any two pseudo-rigid bodies.

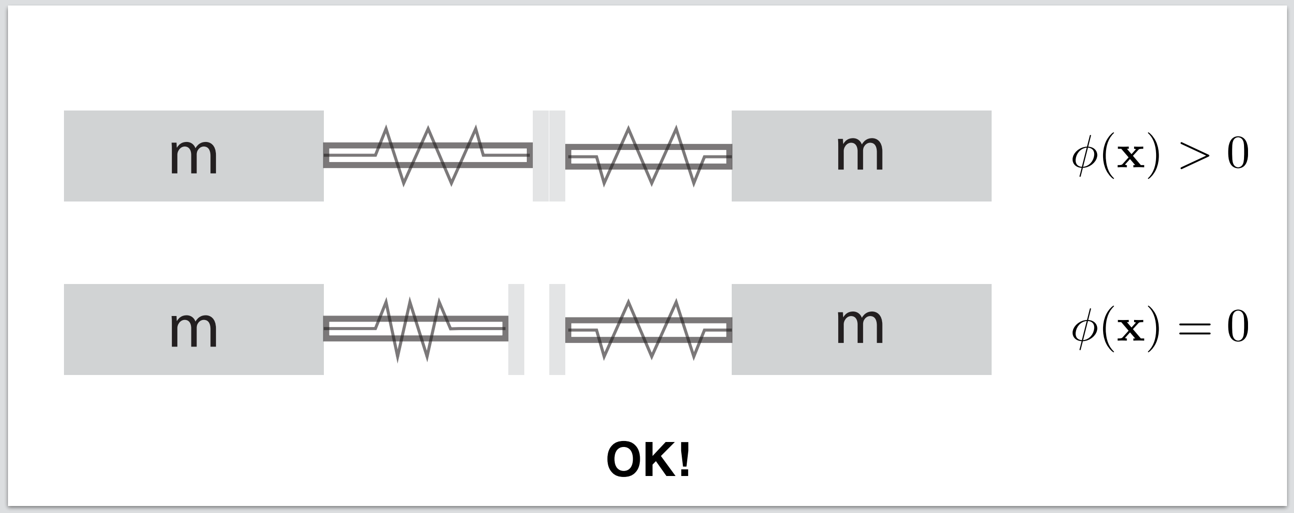

(though clearly this is not a problem in 1D!)

Sufficient condition for the convexity of the pairwise signed distance constraints

Decreasing extension of one or more pistons must not decrease signed distance between any two bodies.

Given this property, we can show that:

$$\phi_{i,j}(t\mathbf{x}^{(1)} + (1-t)\mathbf{x}^{(2)}) \geq 0$$

$$\forall t \in [0,1] \text{ and } i,j \in \mathcal{S} \text{ and } \mathbf{x}^{(1)}, \mathbf{x}^{(2)}$$

$$\text{s.t. } i \neq j \text{ and } \phi_{i,j}(\mathbf{x}^{(1)}) \geq 0, \phi_{i,j}(\mathbf{x}^{(2)}) \geq 0$$

The sufficient condition will be true if the shape at strictly less extension is enclosed by the shape at greater extension.

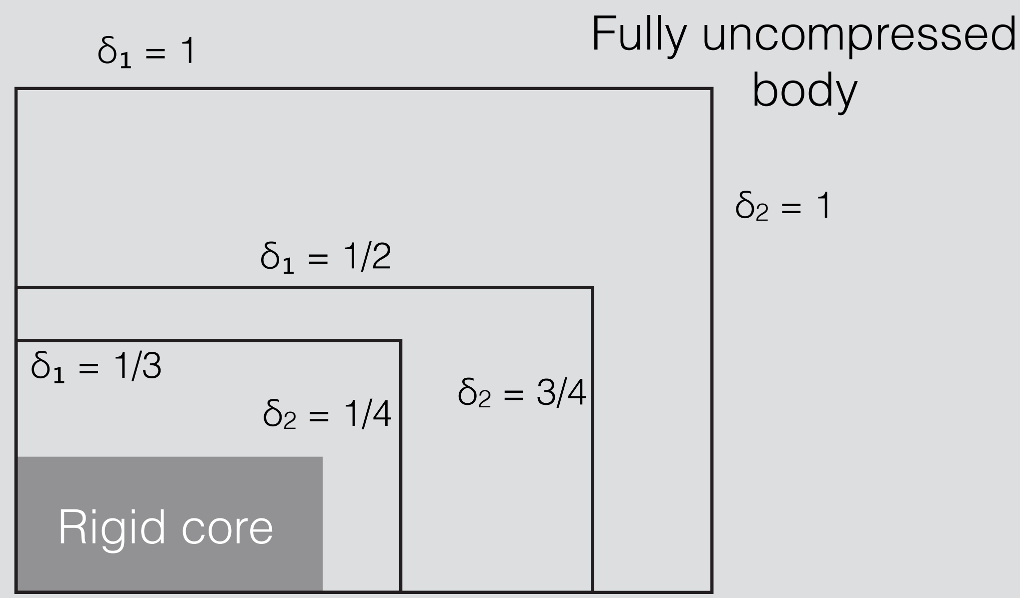

Conclusion / circumspection

1st order accurate, limited event detection, no computational stiffness (from contact)

Conclusion / circumspection

1st order accurate, limited event detection, no computational stiffness (from contact)

High throughput for this →
but using physical principles gives confidence of correct contact modeling between complex geometries

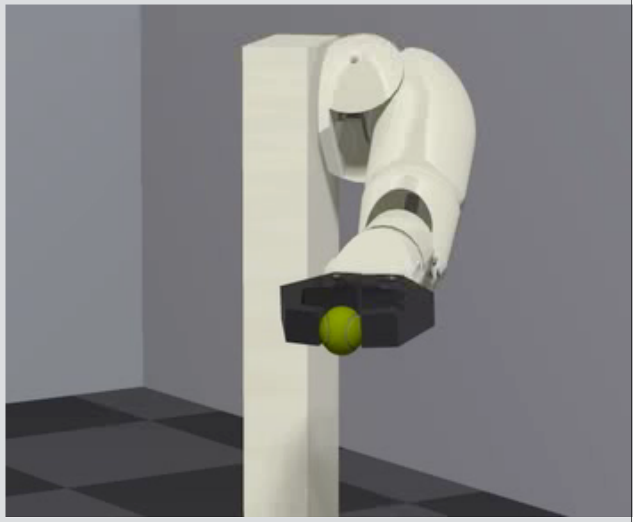

Conclusion / circumspection

Should handle the
average case (rigid
cores not intersecting)
robustly and quickly.

Conclusion / circumspection

Should handle the average case (rigid cores not intersecting) robustly and quickly.

Allow us to visualize deformations.

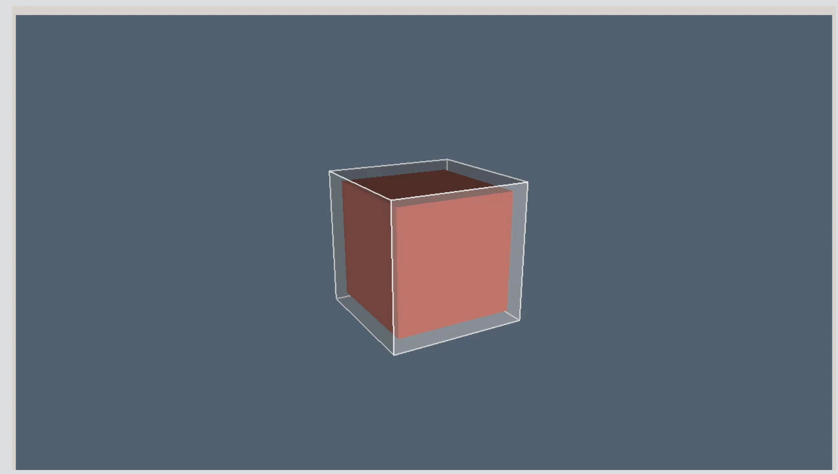

TOYOTA RESEARCH INSTITUTE

Is convexity important?

Model does not have the freedom of fully compliant models but should still be effective for manipulation.

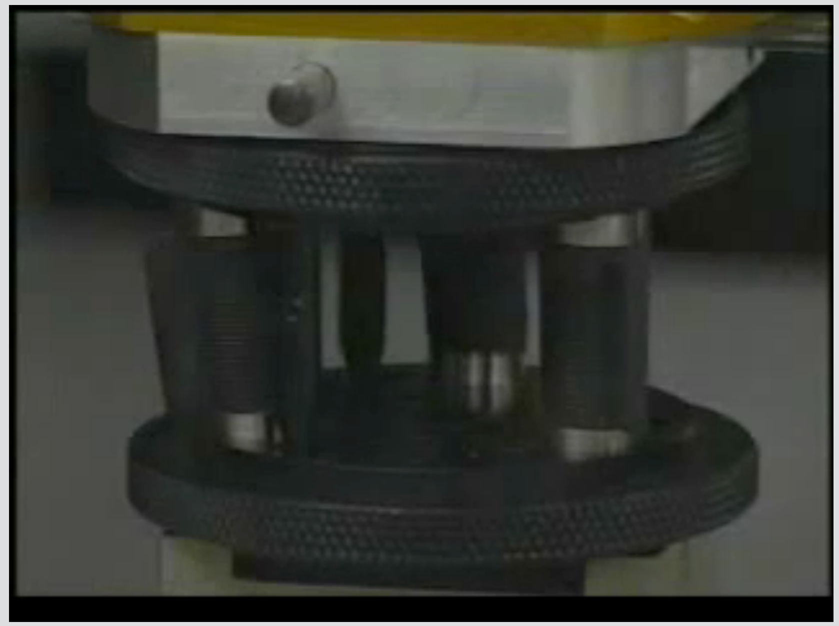

Remote center of compliance tool

I don’t like nonconvex optimization until I understand what the implications of a locally optimal solution are

Is convexity important?

Model does not have the freedom of fully compliant models but should still be effective for manipulation.

Remote center of compliance tool

I don't like nonconvex optimization until I understand what the implications of a locally optimal solution are