

Contact-Implicit Direct Collocation with a Discontinuous Velocity State

Stacey Shield, *Student Member, IEEE*, Aaron M. Johnson, *Senior Member, IEEE* Amir Patel, *Member, IEEE*

Abstract—Collisions between rigid bodies are impulsive events, meaning they create near-instantaneous, finite discontinuities in the velocity state of the dynamic system. If velocity is assumed to be continuous, that requirement may come into conflict with the desire for bodies to not interpenetrate, or with Coulomb friction (creating a condition known as the *Painlevé paradox*). This paper presents a contact-implicit framework for trajectory optimization using direct collocation that combines an unplanned impulsive contact model with an implicit, high-order numerical integration scheme. We demonstrate that this framework can resolve problems that other implicit formulations cannot, namely elastic or partially elastic collisions, and impacts without collision – the established resolution of the Painlevé paradox. We then evaluate its applicability to legged robotics problems and test its ability to discover the intricate sequence of varied collision types necessary to execute a skateboarding trick.

I. INTRODUCTION

Contact allows robots to interact with objects and terrain, making it an essential component of almost any useful task they might be expected to perform. Robots are often modelled as tree-like systems of rigid bodies, which allows their dynamics to be described using a finite set of generalized coordinates and associated ordinary differential equations (ODEs) [1]. However, contact introduces discontinuities in the directions of allowable motion, corresponding to unilateral inequality constraints that preclude the use of numerical methods for simulating ODEs. For the assumption of rigidity to hold, collisions must be impulsive: relative motion must cease at the instant of contact, under the action of infinitely-large forces.

Modelling discontinuity is challenging in the context of direct trajectory optimization, where the use of constrained nonlinear programming prohibits the use of conditional statements. Consequently, collisions are often modelled using continuous-velocity approximations. This allows the velocity to evolve from its value at the moment of impact to zero over one timestep, as shown in Fig. 1. If implicit numerical integration is used, a body is not guaranteed to have its true final velocity at the moment of impact. This makes it impossible to model partially elastic collisions reliably, as the restitution law is applied at the velocity level. These problems can be avoided by using partly-implicit integration instead [2], but that sacrifices the accuracy and stability for which implicit integration is typically favoured.

An option that preserves velocity discontinuity is the *hybrid* dynamic model or multi-phase approach [3]. This

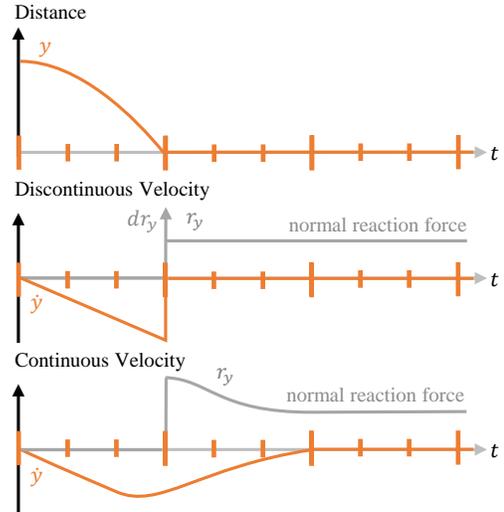


Fig. 1. Comparison between implicit direct collocation with discontinuous velocity state, and continuous velocity states. In the continuous formulation, the velocity must transition to zero over a full timestep under the action of finite forces. This causes the instantaneous impact velocity to be smaller than its true value, making it impossible to model partially elastic collisions accurately.

method explicitly designates which contacts will be active on which timesteps, and then treats the associated unilateral constraints as bilateral when they are active, e.g. [4, 5]. The problem with these methods is the need to specify the sequence of contact states upfront, which is not practical for many problems. An alternative is to use *contact-implicit optimization*, where the contact states are included as additional decision variables either through activation functions [6], smooth approximations of the boundary behaviour [7], or complementarity conditions [8, 9]. These methods allow a wider variety of solutions but, so far, they have either used less accurate first-order integration or assumed a continuous velocity state.

This paper presents an implicit collocation formulation of arbitrary order that overcomes these challenges by implementing complementarity-based contact activation with a discontinuous velocity state, combining the advantages of hybrid and contact-implicit models. This is the first implicit scheme that allows for inelastic, elastic, or partially elastic collisions to occur without a predetermined contact sequence, extending our prior work in [10]. Section II establishes context by expanding on the motivation for allowing velocity discontinuities, the mathematical background, and the basic direct collocation formulation. We then explain the contact model: frictionless impacts are described in Section III, and

S. Shield and A. Patel are with the African Robotics Unit at the University of Cape Town. A. M. Johnson is with the Department of Mechanical Engineering at Carnegie Mellon University. This work is funded by the National Research Foundation of South Africa (grant no. 117744).

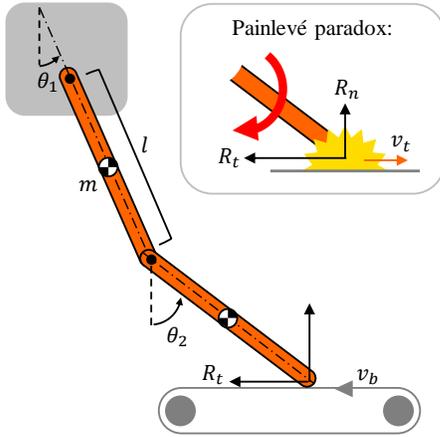


Fig. 2. Two-link pendulum resting on a conveyor belt. The Painlevé paradox will occur if the coefficient of friction between the belt and pendulum is sufficiently large.

Coulomb friction is added in Section IV. The remainder of the paper evaluates its performance with respect to the trajectory optimization of complicated locomotion and object manipulation problems.

II. BACKGROUND

A. Impacts Without Collisions

Continuous-velocity approximations of contact are acceptable in many situations, but there are some problems where the presence of Coulomb friction makes it impossible to find a solution.

For example, consider the problem of a two-link pendulum resting on a conveyor belt, as illustrated in Fig. 2. When the belt moves backwards relative to the pendulum, friction (r_t) is related to the normal force (r_n) by $r_t = \mu r_n$. If its coefficient of friction (μ) is sufficiently large, the torque on the lower link produced by friction will be larger than the torque produced by the normal force, creating an angular acceleration that directs the end of the link down into the conveyor belt. Of course, it is not possible for the link to move in this direction, so the problem appears to have no solution. This apparent conflict between the interpenetration constraint and Coulomb friction is called the Painlevé paradox, and has been an important topic of discussion and driver of theoretical development in rigid-body dynamics since its description near the turn of the 20th Century [11]. Its history, consequences and the ongoing questions it raises are well-documented in a review by Champneys and Várkonyi [12]. Besides the inconsistent case described, the paradox might also result in an indeterminate case, where multiple solutions are possible [12].

Stewart [13] points to the assumption of finite reaction forces (and, by implication, the assumption of a time-continuous velocity state) as a key flaw leading to inconsistency: “In particular, it rules out the possibility that the horizontal component of the velocity (v_t) could be brought to zero *instantaneously* by impulsive contact force”. If v_t is immediately brought to zero, it is no longer required that

$r_t = \pm \mu r_n$, so a solution becomes possible. This demands that contact models allow not only impulsive forces during collisions, but *impacts without collision* (IWCs) – instantaneous jumps in the tangential velocity occurring without a change in the contact state. Experiments by Zhao et al. [14] using an apparatus similar to the pendulum and conveyor belt problem in Fig. 2 indicate that IWCs are not just a convenient *patch* for a *bug* in the mathematics, but rather a representation of a real physical phenomenon, as tangential shocks were observed when the apparatus was arranged in paradoxical configurations.

It may be tempting to dismiss the paradoxes as niche cases happening only at unrealistically high coefficients of friction, but with unfortunate contact geometry or mass distribution of the bodies involved, they can come about under more typical conditions [15]. For instance, analysis of two widely-used passive dynamic walking models suggests that they are far from unlikely in legged locomotion [16]. A further contribution of the model we describe in this paper is that it is the first higher-order collocation scheme to allow IWC resolution of frictional paradoxes.

B. Mathematical Background

The discontinuous nature of rigid-body systems with contact constraints can be conceptualized by changing the equations of motion from ODEs of the form

$$\dot{x}_i(t) = F_i(t, \mathbf{x}(t), \mathbf{u}(t)) \quad (1)$$

where t is time, x_i is a state variable, \mathbf{x} is the state vector and \mathbf{u} is the input vector, to *measure differential inclusions* (MDIs). The MDI is a generalization that allows the right-hand side of the differential equation to be a combination of continuous and impulsive parts:

$$\dot{x}_i(t) = F_i(t, \mathbf{x}(t), \mathbf{u}(t)) + \sum_{j \in \mathbb{N}} \eta_j \delta(t - t_j) \quad (2)$$

Here, $\delta(t - t_j)$ is a unit impulse occurring at the instant t_j and η_j is the magnitude. Although $\delta(t)$ is often referred to as the Dirac δ -function, it is not really a function of time at all, but a *measure* – a function that acts on a set, and may be thought of as something closer to a distribution. The important assumptions we are making about this solution are that there are countably many discontinuities, and that $\dot{x}_i(t)$ has *bounded variation* over the trajectory (the difference between the left and right values of $\dot{x}_i(t)$ at each discontinuity can be assigned a finite value, η_j).

MDIs are the cornerstone of the mathematical framework developed by Moreau to handle a class of unilaterally-constrained mechanical problems he termed *sweeping processes* [17–19]. The immediate ancestors of our model are time-stepping methods based on Moreau’s theory [2, 20, 21]. For further reference on these ideas, see [2] for a concise and accessible introduction, [15] for a broader review of numerical simulation methods, and [22] or [23] for a more comprehensive text.

C. The Direct Collocation Problem

The goal of trajectory optimization is to find the dynamic trajectory of a system over a time interval that minimizes some cost function. Direct collocation transcribes the trajectory optimization problem to a constrained nonlinear programming problem (CNLP) [3]. The specialized mathematical knowledge required to solve these CNLPs may once have been an inextricable challenge, but algebraic modeling languages provide an accessible way for non-specialists to transcribe these problems and pass them to general-purpose solvers. The examples in this paper were written using the Python optimization toolbox, Pyomo [24, 25], and solved using the open-source Interior Point Optimization algorithm (IPOPT) [26] equipped with the Harwell linear solver ma97 [27].

For a general rigid-body mechanical system subject to implicit contact constraints, the decision variables will include the state variables (that is, the positions and velocities of the generalized coordinates), their derivatives, the actuator forces, and the reaction forces at the contact points. The typical constraints required to describe the problem include the *equations of motion* (EOM) that model the system's dynamics, *numerical integration constraints* relating the state and derivative variables from one time interval to the next, the *contact model*, and *task constraints* specifying the activity being simulated.

We discretize the trajectory into N timesteps (referred to as *finite elements*) each consisting of P collocation points. The numerical integration method we use is a P^{th} -order implicit Runge-Kutta scheme where the collocation points are located within the element according to the roots of a Legendre-Gauss polynomial [28]. Unless otherwise specified, we use $P = 4$ in all examples described subsequently. The value of each x_i at the p^{th} point of the n^{th} element is linked to the other points within that element by

$$x_i[n, p] = x_i[n, 0] + h[n] \sum_{j=1}^P \Omega_{p,j} \dot{q}[n, j] \quad (3)$$

The coefficients $\Omega_{p,j}$ are derived from integrals of Lagrange interpolating polynomials [28]. The duration of the element is given by a variable timestep $h[n]$ with maximum duration h_m and minimum duration $0.1h_m$ to allow more flexibility in when the contact state can change. The timesteps may be constrained to add up to a desired total duration T . The Legendre-Gauss collocation points do not include the boundaries, so two additional *mesh points*, denoted by $[n, 0]$ and $[n, P + 1]$, must be defined for each element. The final value is calculated using the Gaussian quadrature rule

$$x_i[n, P + 1] = x_i[n, 0] + \frac{1}{2} h[n] \sum_{j=1}^P w_j \dot{q}[n, j] \quad (4)$$

with weights w_j derived from the Legendre polynomial.

If the variable's trajectory is assumed to be continuous, $x_i[n, 0] = x_i[n-1, P+1]$. We apply this continuity constraint

to the position vector \mathbf{q} , but modify it for the velocity $\dot{\mathbf{q}}$ to allow it to jump by some finite value $\eta_i[n]$:

$$x_i[n, 0] = x_i[n-1, P] + \eta_i[n], \quad x_i \in \dot{\mathbf{q}} \quad (5)$$

The resulting integration scheme is similar to a previous adaptation of arbitrary-order orthogonal collocation to hybrid dynamic problems [5].

The instantaneous velocity change is brought about by an impulsive contact reaction, $d\mathbf{r}$. In our transcription, the velocity jump and contact impulse are represented by the acceleration, $\ddot{\mathbf{q}}$, and reaction, \mathbf{r} , variables at $[n, 0]$. They are related by the impulsive equations of motion:

$$\mathbf{M}(\mathbf{q}[n, 0])\ddot{\mathbf{q}}[n, 0] = \mathbf{J}_c^T(\mathbf{q}[n, 0])\mathbf{r}[n, 0] \quad (6)$$

where \mathbf{M} is the inertial matrix, and \mathbf{J}_c is the contact Jacobian.

III. FRICTIONLESS CONTACT

We use complementarity to control the activation of the contact reactions $\mathbf{r}[n, p]$. We will use the notation $A \perp B$ to denote a complementarity constraint between two variables A and B . It means that they satisfy the following conditions:

$$A \geq 0, \quad B \geq 0, \quad AB = 0 \quad (7)$$

In practise, the constraint becomes much easier to solve if the equality constraint is replaced with

$$AB \leq \varepsilon \quad (8)$$

where ε is a penalty variable that can be brought within an acceptable tolerance by reducing its upper bound incrementally over multiple solve attempts, or by minimizing it as a term in the cost function [29]. In our subsequent experiments, we use the cost function approach, and consider solution to be feasible if $\varepsilon \leq 1e-4$ at all points.

We assume that a well-defined tangent plane exists between the bodies involved in each contact, so \mathbf{r} can be decomposed into a normal component r_n and tangential components \mathbf{r}_t . The tangential components will be zero in the frictionless case. We will assign the coordinate y to the distance between bodies in the normal direction, and use coordinates x and z to describe the tangent plane. The fundamental complementarity relationship that determines the contact state is between r_n and y .

The role of the impulsive reaction component is to produce a jump in the normal velocity \dot{y} that satisfies $\dot{y}^+ = -e\dot{y}^-$, where e is the coefficient of restitution. We cannot complement $r_n[n, 0] \perp \dot{y}[n, 0] - e\dot{y}[n-1, P+1]$ directly, however, as the right-hand side of this expression will not always be positive when the contact is inactive. We therefore implement the contact complementarity at the initial mesh point using positive auxiliary variables a^+ and a^- as follows:

$$a^+[n] - a^-[n] = \dot{y}[n, 0] + e\dot{y}[n-1, P+1] \quad (9a)$$

$$r_n[n, 0] \perp y[n, 0] + a^+[n] + a^-[n] \quad (9b)$$

$$a^+[n] \geq 0 \quad a^-[n] \geq 0 \quad (9c)$$

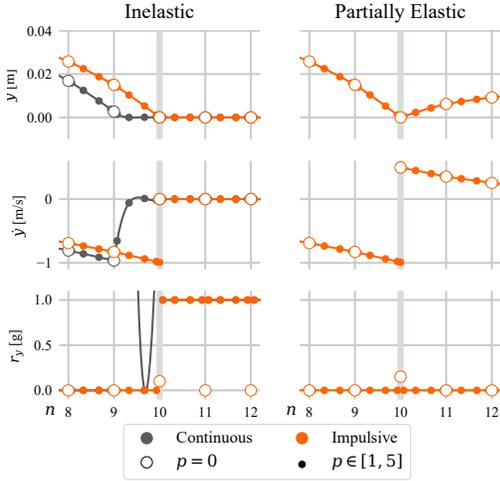


Fig. 3. Trajectories of inelastic and partially elastic ($e = 0.5$) collisions between a point mass and the ground. The normal impulse dr_n is depicted at $p = 0$. A continuous velocity model is included for comparison on the inelastic case, showing the change in velocity spread out over an entire finite element.

The role of the finite reaction component is to prevent interpenetration when bodies are in contact for longer than an infinitesimal instant. It is not sufficient to complement $r_n[n, p] \perp y[n, p]$, as this would also allow non-impulsive collisions that do not obey the specified coefficient of restitution. Preventing these is effectively the same as preventing collisions from happening in the middle of finite elements, which is also necessary to ensure an accurate spline approximation of the dynamics. We do this by complementing the normal force with the distance at all points in the finite element:

$$r_n[n, p] \perp \sum_{i=0}^{P+1} y[n, i] \quad (10)$$

If the velocity is assumed to be continuous, only the above constraint controls contact activation. It is not possible to satisfy this constraint perfectly during the continuous approximation of touchdown, as $y[n, 0]$ must be greater than zero to prevent interpenetration. In our previous work with continuous high-order formulations [10], we effectively applied the constraint at the velocity level by complementing r_n with the sum of y over the *next* finite element instead, but that approach is less accurate, as it does not minimize the distance between bodies when r_n begins to act.

A. Example: Falling Point Mass

To demonstrate impact, we simulated a one-dimensional point mass experiencing (A) an inelastic collision ($e = 0$) and (B) a partially elastic collision ($e = 0.5$) with the ground. $N = 20$ and $h_m = 0.02$, and the initial height $y[1, 0] = 0.1\text{m}$. The only objective was minimizing the complementarity penalties described in (8). The inelastic collision was also simulated using the equivalent continuous-velocity formulation, to illustrate the differences between the two approaches. The partially elastic case is inadmissible for

the continuous formulation, which is a key advantage of the proposed approach.

Fig. 3 shows the resulting position, velocity and ground reaction force trajectories. The models behave as expected: the velocity in the impulsive problem jumps upon contact with the ground in accordance with the coefficient of restitution. A normal impulse occurs at the moment of impact, following which the normal force prevents the mass from falling through the ground in the inelastic case. The continuous approximation of touchdown takes place over a single finite element, with the mass not quite grounded when the normal force begins to decelerate it.

IV. CONTACT WITH COULOMB FRICTION

Using Coulomb's law, the set of possible friction forces in a two-dimensional contact plane forms a disc with a radius of μ around the contact point. We refer to the combination of this disc with the set of possible normal forces as the *friction cone*. The reaction force acting at a stationary contact can fall anywhere in the interior of the cone, while the force at a sliding contact must lie on the boundary of the cone in the opposite direction to the relative tangential velocity, \mathbf{v}_t .

It is possible to simulate reaction forces falling anywhere on the friction cone [30, 31], but these techniques involve a nonlinear transformation of \mathbf{v}_t and \mathbf{r}_t into a polar representation of the contact plane. A more computationally-efficient option is to convert this to a linear complementarity problem by working with a polyhedral approximation of the friction cone, where the disc of possible tangential forces is replaced by a polygon [8, 10, 20]. This polygon is the convex hull of k evenly-spaced direction vectors of length μ . For planar problems, $k = 2$. The minimum number of vectors for a spatial problem is usually $k = 4$, giving a set of direction vectors that coincides with the positive and negative directions of the x and z axes describing the contact plane. We can write the friction force as

$$\mathbf{r}_t = \mu r_n \mathbf{d}^k \alpha^k \quad (11)$$

using a set of k unit vectors, \mathbf{d}^k . The vector α^k consists of k activation variables, each having a value between zero and one.

Unless the direction of \mathbf{v}_t falls precisely between two of the direction vectors, only one element of α^k should be nonzero at any point where sliding occurs. If $\mathbf{v}_t = \mathbf{0}$, α^k must take on the values required to oppose the net tangential force acting on the bodies. To manage these activation variables, we introduce an auxiliary variable $\gamma \geq 0$ and relate it to each activation variable $\alpha_i^k, \forall i = 1, 2 \dots k$ with the complementarity constraint:

$$\alpha_i^k[n, p] \perp \gamma[n, p] - \mathbf{d}_i^k \mathbf{v}_t^T[n, p] \quad (12)$$

The requirement that $\gamma - \mathbf{d}_i^k \mathbf{v}_t^T \geq 0$ means that γ will equal the magnitude of the largest projection of \mathbf{v}_t onto one of the unit vectors in \mathbf{d}^k – that is, the magnitude of the projection of \mathbf{v}_t onto the unit vector that best matches its direction. The right-hand side of (12) can only be zero for the constraint

corresponding to this nearest direction vector, so only the associated activation variable can have a nonzero value.

To ensure that its value will be one if $|\mathbf{v}_t| \geq 0$ (and hence, if $\gamma \geq 0$), we complement

$$\sum_{p=0}^{P+1} \gamma[n, p] \perp 1 - \sum_{i=1}^k \alpha_i^k \quad (13)$$

We sum γ over all points to prevent changes between sticking and sliding states from happening mid-element.

A. Example: Tangential Impact

Because the proposed formulation enables impulsive reaction forces to act at the boundaries of any finite elements where contact is active, it permits the resolution of the Painlevé paradox through a tangential impact without collision (IWC). To demonstrate this, we simulated a paradoxical situation in a planar system based on the two-link manipulator and conveyor belt apparatus that Zhao et al. used to investigate IWCs experimentally.

The model is shown in Fig. 2. We assigned the links unit mass and a length of 0.5 metres, and assumed uniform mass distribution. Collisions were assumed to be perfectly inelastic. The height of the top link was selected such that the end of the double pendulum rests on the conveyor belt at initial angles of $\theta_1[1,0] = 0$ and $\theta_2[1,0] = 0.25\pi$ rad. The initial velocity was fixed to zero. The belt was initially stationary, but its velocity was abruptly stepped up to $v_b = 0.5$ m/s (that is, -0.5 m/s relative to the pendulum) at point [6,0]. We selected a very high coefficient of friction, $\mu = 2$, so Painlevé’s paradox would be induced at this instant. The timing parameters were $N = 10$, $h_m = 0.02$ s and $T \geq 0.1$ s.

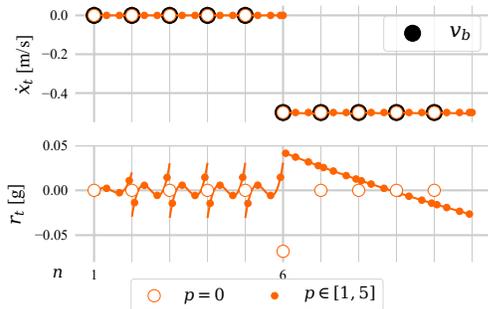


Fig. 4. Solution to the pendulum and conveyor belt problem shown in Figure 2, demonstrating resolution of the Painlevé paradox via tangential impact at the start of the sixth finite element.

The results of the optimization are displayed in Fig. 4. While the contact state never changes, an impulse occurs that instantly increases the tangential velocity of the pendulum’s end to match the velocity of the belt. The end of the link therefore remains stationary relative to the belt throughout, so the magnitude of the friction force is allowed to be $< \mu r_n$, and a downward acceleration is not created at the contact point. This is precisely the resolution of the Painlevé paradox via IWC described before [14].

V. APPLICATION I: INELASTIC CONTACTS

The impulsive formulation has a clear benefit in trajectory optimization problems involving partially elastic contact and IWC, as shown above, but we are also interested whether it has any advantage over the continuous formulation when perfectly-inelastic contact is assumed, as is typical in legged locomotion tasks. It is possible that the ability to resolve frictional paradoxes could allow it to explore the solution space more effectively, and discover solutions that would be infeasible for other formulations, but this could be outweighed by an increased computational load similar to adding another collocation point to each element.

To evaluate its effect on solver performance and solution quality, we compare continuous and impulsive versions of the fourth-order Legendre-Gauss scheme (LG4I and LG4C, respectively) over two trajectory optimization tasks. For benchmarking purposes, we also compare the solving times to the established first-order (FO) contact-implicit method described in [8] – the most direct predecessor to our approach. This method was assigned $4N$ timesteps to give a similar problem size. More recent contact-implicit optimization schemes often reduce the computational load by modeling the robot as a quasistatic system [32], but we exclude these methods from this comparison because this is not a feasible approach for our intended application, as the quasistatic assumption becomes highly inaccurate for high-speed manoeuvres [33].

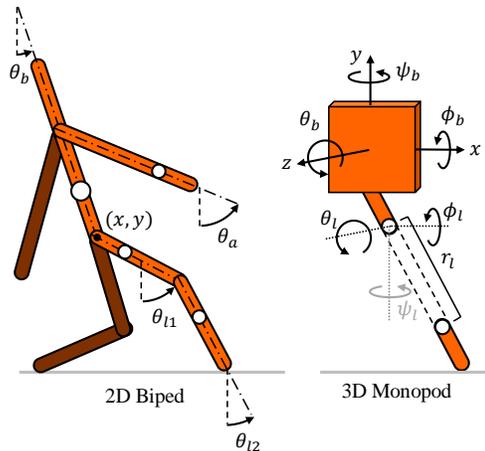


Fig. 5. Models of planar biped and spatial monopod used in performance comparison tests.

The models for each test are illustrated in Fig. 5. Both models effectively have nine degrees of freedom (DOF); although the monopod is modelled using 10 generalized coordinates to produce more tractable equations of motion [33], the leg is constrained so it cannot yaw relative to the body. Each model includes frictional contacts at the feet, and frictionless contacts acting as hard stops at the upper and lower limits of all joints. We ran each experiment with two coefficients of friction: $\mu = 0.6$, and $\mu = 1.6$ – a high value that is more likely to produce Painlevé paradoxes.

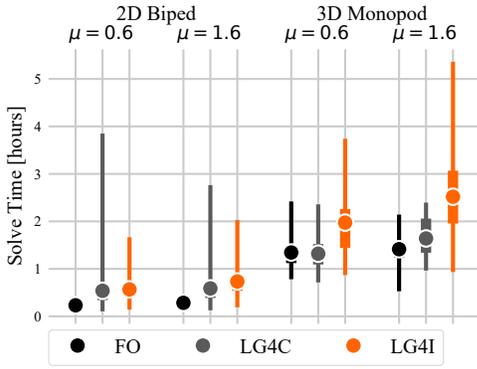


Fig. 6. Combined solve times for the feasible and optimal stages on the 2D biped and 3D monopod using: First Order (FO) [8], Legendre-Gauss 4^{th} order continuous contact (LG4C), and impulsive contacts (LG4I). The circle indicates the mean, the narrow line is the range, and the wider line is the interquartile range.

A. Planar Bipedal Stopping Problem

To simulate a stopping maneuver, the initial condition for the biped was sampled from the midstance phase of a sprinting trajectory, and the final condition at $[N, P + 1]$ required grounded feet, no forward translational or rotational velocity (that is, $\dot{x} \leq 0$ and $\dot{\theta}_b \geq 0$), and all other velocities to have magnitudes within five percent of their initial values. We minimized the stopping distance by creating a variable upper bound $x[n, p] \leq x_m$ at all $[n, p]$, and then minimizing x_m . We allocated $N = 25$ finite elements of maximum duration 0.1 s to the task.

B. Spatial Monopodal Turning Problem

The monopod was required to start at rest in an upright position, and travel 2.5 m in the x direction without exceeding $z \leq 0$, followed by 2.5 m in the z direction. The final state was not specified beyond the requirements that $x = z = 2.5$ m, and that the yaw of the body (ψ_b) be displaced 90 degrees from its starting point. We minimized the sum of the squared actuator forces and torques over all joints and collocation points. This task was assigned $N = 25$ finite elements of maximum duration 0.1 s.

C. Optimization Procedure

We solve the trajectory optimization problems in two stages: first, we minimize the complementarity penalties to obtain a feasible trajectory, then we bound the penalties below a suitable threshold (here, $1e - 4$) and minimize the desired objective using the feasible trajectory as a starting point. For the feasibility stage, we seeded the position and reaction force variables with random values ≤ 0.1 . We gave the same random seed to all models in each test, and repeated this 100 times for each friction coefficient.

D. Results

The LG4I formulation produced solutions slightly slower than the LG4C formulation did in both biped tests, and substantially slower in both monopod tests, as shown in Fig. 6. This suggests that the complementarity constraints

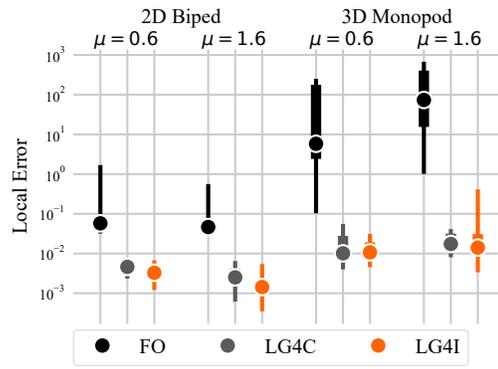


Fig. 7. Mean local error for the optimal trajectories on the 2D biped and 3D monopod using: First Order (FO) [8], Legendre-Gauss 4^{th} order continuous contact (LG4C), and impulsive contacts (LG4I). Error values are relative to the largest magnitude of the state variables. The circle indicates the median, the narrow line is the range, and the wider line is the interquartile range.

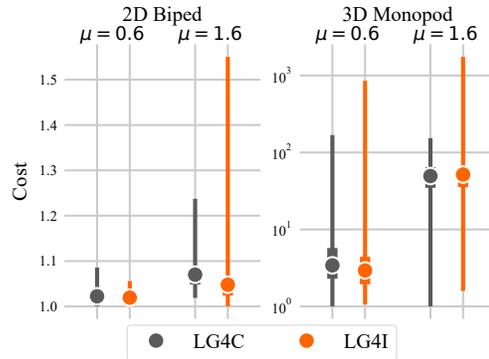


Fig. 8. Cost values for optimal solutions, normalized to the minimum achieved by either model in each test. The circle indicates the mean, the narrow line is the range, and the wider line is the interquartile range.

are the primary contributor to the computation time, as the increase in their number between the LG4C and LG4I formulations is much larger for the 3D problem.

Both LG4 configurations typically took around twice as long as the FO benchmark to solve in the biped tests, but the difference was much smaller for the 3D monopod tests, with the LG4C method even performing slightly better in the lower friction case. They are, however, much more accurate, as shown by the plot of the mean local error in Fig. 7. This error metric measures how well the polynomial approximation of the problem satisfies the EOM between collocation points. It is calculated from the difference between the derivative spline, and the derivative calculated by substituting the state and input splines into the EOM, integrated over each finite element. The error for each state variable was scaled by the largest magnitude of the variable on each element, and the mean value for the trajectory was then aggregated over all finite elements and state variables. Both LG4 approaches were dramatically more accurate than FO, with the impulsive formulation showing slightly better accuracy than the continuous formulation.

Fig. 8 compares the costs obtained, normalized to the lowest value achieved by either model in each case. (We

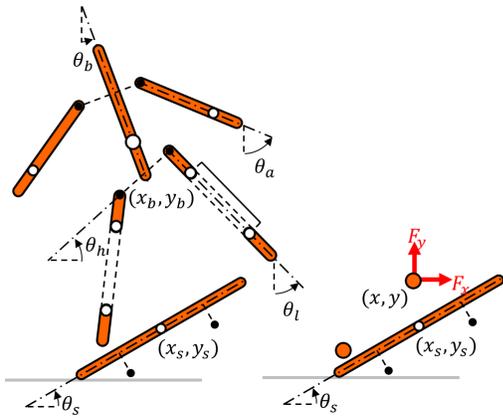


Fig. 9. (Left) Planar model of skateboard and bipedal rider. (Right) Simplified system using two point masses instead of a full-body rider model.

exclude the FO solutions from the cost comparison, as they cannot be meaningfully compared to the fourth-order results due to the differences in accuracy.) The median costs were near-identical in all tests, but the spread of the LG4I results tended to be wider. While its median results were usually slightly better, it also generated the worst solution in most tests. This suggests that it might be able to identify superior strategies that are infeasible for the continuous version, but also that the slightly more cumbersome formulation could be more prone to getting trapped in bad local minima.

Overall, there does not appear to be a clear advantage to using the impulsive configuration for problems with exclusively inelastic collisions. While the accuracy was slightly improved over the continuous formulations, the computation time and consistency were worse.

VI. APPLICATION II: MIXED CONTACTS

In theory, the advantage of contact-implicit trajectory optimization for motion planning is that it allows the model to discover the best possible contact sequence to perform a task, or identify a feasible sequence if one is unknown. In practise, these problems are difficult to solve and prone to falling into local minima, especially on whole-body system models with many degrees of freedom. To evaluate the ability of our formulation to discover solutions to complicated, contact-heavy motion problems, we selected a skateboarding trick called an *Ollie* as a test case.

Although the achievement of *sick air* is not currently regarded as an important priority in robotics research, the Ollie is interesting as an example of a challenging object manipulation problem requiring an intricate sequence of diverse contact interactions to complete successfully. We have used the Ollie to explore trajectory optimization with varied contacts before [34], but required a combination of hybrid and complementarity-based contact schemes to generate the motion. The impact model presented in this paper allows a fully contact-implicit formulation of the problem.

The objective of the Ollie is to get all four wheels of the skateboard off the ground. The rider stamps on the tail of the board while jumping up, so it bounces off the ground

and propels the board into the air. Once airborne, the feet manipulate the board to execute further aerial tricks, or just position it for a safe landing.

Our model of a skateboard and humanoid rider is shown in Fig. 9. We also attempted the test using the simplified model on the right of this figure, which isolates the contact problem by replacing the rider with a pair of point masses actuated by external forces. Three different types of contact are present in the system:

- 1) Partially elastic ($e = 0.6$), frictionless contact between the tail of the skateboard and the ground.
- 2) Inelastic, high-friction ($\mu = 1.6$) contact between the feet and skateboard. The position of the contact point with respect to the board is variable, so the feet can connect anywhere along the deck.
- 3) Inelastic, frictionless contact between the wheels of the skateboard and the ground. (The wheels are modelled as simple contact points offset below the deck, as they are only required to support the board in this example.)

We used $P = 2$ and $N = 40$ with $h_m = 0.02$ for this example to allow more opportunities for contact state changes than an equivalently-sized problem with $P = 4$ would. The initial and final conditions have the humanoid standing upright on the board with both wheels grounded, and the system initially at rest. An additional *air* condition requires both wheels to be more than 0.2 m above the ground at point $[n, p] = [20, 0]$.

We tried four seeding approaches to initialize the problem:

- 1) **null seed**: the default initial vector.
- 2) **perturbed null seed**: the position and contact variables are assigned small, random values.
- 3) **perturbed null seed with hint**: same as previous, but the height of the board tail is fixed to zero at the point $[10, 0]$.
- 4) **perturbed solution**: a previous successful result, perturbed by small random values.

Both models reliably generated the trick when given a *hint* specifying the initial tail contact, or a perturbed solution as a seed (strategies 3 and 4). The average solving time for the successful attempts on the full-body model was around 43 minutes. Unfortunately, neither model produced a feasible solution from the null seed, or from 20 randomized null seeds. This highlights a key challenge of contact-implicit trajectory optimization: in many cases, the desired result lies within a small basin of attraction that is exceedingly difficult to discover without some pre-existing knowledge of the contact sequence.

Although seeding with perturbed solutions is not a practical method, as it requires the problem to be solved at least once before, the success of this approach shows that the solution is admissible and that motion discovery is possible, given a seed of sufficient quality. Problem initialization is an important aspect of trajectory optimization, so further research is needed to determine precisely what “sufficient quality” means in this context, and how these seeds can be obtained for truly unknown contact sequences. One

approach involves reducing the whole-body model to its centroidal dynamics [35], and we have also had promising preliminary results initializing legged locomotion problems from procedurally-generated smooth-random gaits we termed *silly walks* [36]. As models become more complicated and detailed, the most effective role for contact-implicit optimization in motion planning is likely to be the refinement of coarse motions generated by a high-level global planner [37], rather than primary motion discovery.

VII. CONCLUSION

This paper introduces a contact formulation for direct collocation that combines the advantages of hybrid-dynamic and contact-implicit approaches. By accommodating finite discontinuities in the velocity state, it can capture behaviour that continuous-velocity formulations cannot, such as partially elastic collisions and tangential “impacts without collision” – the established resolution of the frictional paradoxes identified by Painlevé. However, the more computationally cumbersome formulation leads to longer solve times for 3D or especially contact-heavy problems, and so the approach is currently best suited to problems that require elasticity or impacts without collisions. Further work on problem initialization is required to support the discovery of the varied and complicated contact sequences this formulation theoretically allows.

REFERENCES

- [1] B. Siciliano and O. Khatib, *Springer Handbook of Robotics*. Springer, 2016.
- [2] D. E. Stewart, “Time-stepping methods and the mathematics of rigid body dynamics,” *Impact and Friction of Solids, Structures and Machines*, vol. 1, 2000.
- [3] M. Kelly, “An introduction to trajectory optimization: How to do your own direct collocation,” *SIAM Review*, vol. 59, no. 4, pp. 849–904, 2017.
- [4] A. Hereid, C. M. Hubicki, E. A. Cousineau *et al.*, “Hybrid zero dynamics based multiple shooting optimization with applications to robotic walking,” in *IEEE International Conference on Robotics and Automation*, 2015, pp. 5734–5740.
- [5] D. Pardo, M. Neunert, A. W. Winkler *et al.*, “Hybrid direct collocation and control in the constraint-consistent subspace for dynamic legged robot locomotion,” in *Robotics: Science and Systems*, 2017.
- [6] D. W. Marhefka and D. E. Orin, “A compliant contact model with nonlinear damping for simulation of robotic systems,” *IEEE Transactions on Systems, Man, and Cybernetics-Part A: Systems and Humans*, vol. 29, no. 6, pp. 566–572, 1999.
- [7] I. Chatzinikolaïdis, Y. You, and Z. Li, “Contact-implicit trajectory optimization using an analytically solvable contact model for locomotion on variable ground,” *IEEE Robotics and Automation Letters*, vol. 5, no. 4, pp. 6357–6364, 2020.
- [8] M. Posa, C. Cantu, and R. Tedrake, “A direct method for trajectory optimization of rigid bodies through contact,” *The International Journal of Robotics Research*, vol. 33, no. 1, pp. 69–81, 2014.
- [9] I. Mordatch, E. Todorov, and Z. Popović, “Discovery of complex behaviors through contact-invariant optimization,” *ACM Transactions on Graphics (TOG)*, vol. 31, no. 4, p. 43, 2012.
- [10] A. Patel, S. L. Shield, S. Kazi *et al.*, “Contact-implicit trajectory optimization using orthogonal collocation,” *IEEE Robotics and Automation Letters*, vol. 4, no. 2, pp. 2242–2249, 2019.
- [11] P. Painlevé, “Leçons sur le frottement,” *Comptes Rendu des Séances de l’Académie des Sciences*, vol. 121, pp. 112–115, 1895.
- [12] A. R. Champneys and P. L. Várkonyi, “The painlevé paradox in contact mechanics,” *IMA Journal of Applied Mathematics*, vol. 81, no. 3, pp. 538–588, 2016.
- [13] D. E. Stewart, “Rigid-body dynamics with friction and impact,” *SIAM review*, vol. 42, no. 1, pp. 3–39, 2000.
- [14] Z. Zhao, C. Liu, W. Ma, and B. Chen, “Experimental investigation of the painlevé paradox in a robotic system,” *Journal of Applied Mechanics*, vol. 75, no. 4, 2008.
- [15] B. Brogliato, A. Ten Dam, L. Paoli *et al.*, “Numerical simulation of finite dimensional multibody nonsmooth mechanical systems,” *Appl. Mech. Rev.*, vol. 55, no. 2, pp. 107–150, 2002.
- [16] Y. Or, “Painlevé’s paradox and dynamic jamming in simple models of passive dynamic walking,” *Regular and Chaotic Dynamics*, vol. 19, no. 1, pp. 64–80, 2014.
- [17] J. J. Moreau, “Standard inelastic shocks and the dynamics of unilateral constraints,” in *Unilateral Problems in Structural Analysis*. Springer, 1985, pp. 173–221.
- [18] —, “Unilateral contact and dry friction in finite freedom dynamics,” in *Nonsmooth mechanics and Applications*. Springer, 1988, pp. 1–82.
- [19] M. Kunze and M. D. M. Marques, “An introduction to Moreau’s sweeping process,” in *Impacts in Mechanical Systems*. Springer, 2000, pp. 1–60.
- [20] D. E. Stewart and J. C. Trinkle, “An implicit time-stepping scheme for rigid body dynamics with inelastic collisions and coulomb friction,” *International Journal for Numerical Methods in Engineering*, vol. 39, no. 15, pp. 2673–2691, 1996.
- [21] M. Anitescu and F. A. Potra, “A time-stepping method for stiff multibody dynamics with contact and friction,” *International Journal for Numerical Methods in Engineering*, vol. 55, no. 7, pp. 753–784, 2002.
- [22] V. Acary and B. Brogliato, *Numerical methods for nonsmooth dynamical systems: applications in mechanics and electronics*. Springer Science & Business Media, 2008.
- [23] R. I. Leine and H. Nijmeijer, *Dynamics and bifurcations of non-smooth mechanical systems*. Springer Science & Business Media, 2013, vol. 18.
- [24] M. L. Bynum, G. A. Hackebeil, W. E. Hart *et al.*, *Pyomo—Optimization Modeling in Python*, 3rd ed. Springer Science & Business Media, 2021, vol. 67.
- [25] W. E. Hart, J.-P. Watson, and D. L. Woodruff, “Pyomo: Modeling and solving mathematical programs in Python,” *Mathematical Programming Computation*, vol. 3, no. 3, pp. 219–260, 2011.
- [26] A. Wächter and L. T. Biegler, “On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming,” *Mathematical programming*, vol. 106, no. 1, pp. 25–57, 2006.
- [27] HSL, “A collection of Fortran codes for large-scale scientific computation,” See <http://www.hsl.rl.ac.uk>, 2007.
- [28] L. T. Biegler, *Nonlinear programming: concepts, algorithms, and applications to chemical processes*. SIAM, 2010.
- [29] B. Baumrucker, J. G. Renfro, and L. T. Biegler, “MPEC problem formulations and solution strategies with chemical engineering applications,” *Computers & Chemical Engineering*, vol. 32, no. 12, pp. 2903–2913, 2008.
- [30] J.-S. Pang and D. E. Stewart, “A unified approach to discrete frictional contact problems,” *International Journal of Engineering Science*, vol. 37, no. 13, pp. 1747–1768, 1999.
- [31] D. Pretorius and C. Fisher, “A novel method for computing the 3d friction cone using complementary constraints,” in *IEEE International Conference on Robotics and Automation*, 2021.
- [32] T. Pang and R. Tedrake, “A convex quasistatic time-stepping scheme for rigid multibody systems with contact and friction,” in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 6614–6620.
- [33] A. Knemeyer, S. Shield, and A. Patel, “Minor change, major gains: The effect of orientation formulation on solving time for multibody trajectory optimization,” *IEEE Robotics and Automation Letters*, vol. 5, no. 4, pp. 5331–5338, 2020.
- [34] N. Anderson, S. Shield, and A. Patel, “The ollie: A case study in trajectory optimization with varied contacts,” in *International SAUPEC/RobMech/PRASA Conference*. IEEE, 2020, pp. 1–7.
- [35] H. Dai, A. Valenzuela, and R. Tedrake, “Whole-body motion planning with centroidal dynamics and full kinematics,” in *IEEE International Conference on Humanoid Robots*, 2014, pp. 295–302.
- [36] S. Shield and A. Patel, “On the effectiveness of silly walks as initial guesses for optimal legged locomotion problems,” in *International SAUPEC/RobMech/PRASA Conference*. IEEE, 2020, pp. 1–7.
- [37] J. Norby and A. M. Johnson, “Fast global motion planning for dynamic legged robots,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2020, pp. 3829–3836.