Viability and Predictive Control for Safe Locomotion

Pierre-Brice Wieber

► To cite this version:

Pierre-Brice Wieber. Viability and Predictive Control for Safe Locomotion. IROS 2008 - IEEE-RSJ International Conference on Intelligent Robots & Systems, Sep 2008, Nice, France. pp.1103-1108, 10.1109/IROS.2008.4651022 . inria-00390555

HAL Id: inria-00390555

<https://hal.inria.fr/inria-00390555>

Submitted on 2 Jun 2009

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Viability and Predictive Control for Safe Locomotion

Pierre-Brice Wieber

Abstract—The problem of safe locomotion of legged and wheeled robots, when trying to avoid falling, tipping over or hitting obstacles, appears to be a problem of viability and not of Lyapunov stability. Theoretically speaking, viability and Model Predictive Control are unquestionably related, but both can quickly lead to untractable numerical problems. We present here a promising approach for the problem of avoiding to fall in the case of legged locomotion that elegantly solves this difficulty. We propose then a brief discussion about what makes this approach successful with respect to the approaches proposed for the other problems where viability is at stake. This paper should be considered therefore mostly as a prospective reflection on the general problem of safe robotic locomotion.

I. INTRODUCTION

Wheeled and legged systems entirely depend on the mechanical interactions between the wheels or the feet and the ground for generating and controlling their displacements [18]. The problem is that the corresponding contact forces are physically limited, what induces strong limitations on their ability to control their displacements [17]. We'll focus here on their ability to avoid falling, tipping over or simply hitting obstacles, properties that appears to be only loosely related to classical stability concepts such as Lyapunov stability, properties that will appear in fact to be better expressed in terms of viability [1].

The question then is how to design a feedback control law that induces this viability whenever possible. We propose to explore here the possibilities of Model Predictive Control schemes. Such schemes have already been applied to wheeled and legged systems [10], [2], but we propose to analyze more closely here their possible connections with the viability theory. We'll focus more precisely on a promising approach that appears to solve the problem of avoiding to fall in the case of legged locomotion. The particular strength of the resulting control scheme is that it doesn't focus explicitly on the viability property, which is obtained in a sense as a side effect.

We propose then a brief and largely open discussion comparing this approach with the more classical ones that have already been proposed for avoiding tipping over or avoiding hitting obstacles in the case of wheeled locomotion [10], [9], [7], [16], [4].

II. GENERAL BACKGROUND

A. A Differential Inclusion

Let's begin by describing with a vector q the configuration of a dynamical system which has locomotion abilities, a

legged or a wheeled robot for example, separating explicitly the description q_2 of its global position and orientation in the space from the rest of the description of its configuration q_1 [18]. In the case of a humanoid robot, that amounts for example to describing the position and orientation of its pelvis on one side, the rest of its posture on the other side. In the case of a car, the position and orientation of the main body on one side, the position and orientation of the wheels and other moving parts due to the mechanical action of the suspension and the steering wheel on the other side.

Unless this dynamical system is equipped with actuators that can directly act on its position and orientation, thrusters for example, the corresponding Lagrangian dynamics can be easily shown to have the following structure,

$$\begin{bmatrix} M_1(q) \\ M_2(q) \end{bmatrix} (\ddot{q} + g) + \begin{bmatrix} n_1(q, \dot{q}) \\ n_2(q, \dot{q}) \end{bmatrix} = \begin{bmatrix} u \\ 0 \end{bmatrix} + \begin{bmatrix} C_1(q)^T \\ C_2(q)^T \end{bmatrix} \lambda, \quad (1)$$

where simply appears the fact that the actuation forces u of the system can't have a direct influence on the part of the dynamics that corresponds precisely to the global position and orientation q_2 [18]. In the case of a humanoid robot, this corresponds to the fact that joint actuators only move joints, in the case of a wheeled robot, the fact that motors only move wheels, and that none of these actuators have any influence on the position of the system *until the system comes in contact with its environment*.

Indeed, if we look more closely to the inner structure of this dynamics, we can observe that this part not directly influenced by the actuators u corresponds in fact to the Newton and Euler equations of the whole system [18]. The Euler equation can be subtle and convey nonholonomic couplings between the rotation of the system and the rest q_1 of its configuration. Moving joints can influence your orientation, this is how cats manage to always fall back on their feet. The Newton equation however is more straightforward, it is integrable: with the sole action of the forces u , the position of the Center of Mass (CoM) of the system is constant. There's no way to move around by only moving joints [18].

What appears then is that for moving around, the system needs to interact with its environment through forces λ . In the case of mechanical interactions, that would amount to pushing or pulling on the environment. Now, unless the system is tightly fastened to the environment, through a strong hold or any clamping device, it can only push on the environment, and doing so, make use of friction. This typically describes the mechanical interaction between feet or wheels and the ground. In that case, the interaction forces λ are limited by physical laws, what can be generically

Fig. 1. The viability kernel gathers all the states from where the mobile system can avoid falling, colliding obstacles or tipping over through proper action. Outside of this set, these outcomes are unavoidable because the physical constraints (4) or (5) make any alternative impossible.

expressed with the following inclusion,

$$\lambda \in \Lambda(t). \quad (2)$$

Considering possible limitations of the actuation forces in the form of a second inclusion

$$u \in \mathcal{U} \quad (3)$$

that will be of secondary importance here, the dynamics (1) can be rewritten as a Differential Inclusion

$$\begin{bmatrix} M_1(q) \\ M_2(q) \end{bmatrix} (\ddot{q} + g) + \begin{bmatrix} n_1(q, \dot{q}) \\ n_2(q, \dot{q}) \end{bmatrix} \in \begin{bmatrix} \mathcal{U} \\ 0 \end{bmatrix} + \begin{bmatrix} C_1(q)^T \\ C_2(q)^T \end{bmatrix} \Lambda(t). \quad (4)$$

Focusing on the second part of this dynamics, that gathers the Newton and Euler equations of the whole system, we obtain a necessary condition for a movement $q, \dot{q}, \ddot{q}$ to be realized, the Differential Inclusion

$$M_2(q)(\ddot{q} + g) + n_2(q, \dot{q}) \in C_2(q)^T \Lambda(t). \quad (5)$$

B. Viability Theory

The implication of the previous Differential Inclusions, this dependance on interaction forces which are physically limited, is that the mobile system we're considering has limitations in its ability to control its displacements. For a humanoid robot, limitations in its ability to keep its balance [17], for a wheeled robot, limitations in its ability to avoid tipping over during aggressive maneuvers on a rough terrain [13], [16], or more simply limitations in its ability to avoid obstacles [7], [4]. Let's define therefore the set $\mathcal{A}(t)$ of configurations q that the system should definitely avoid, configurations where the robot has fallen, tipped over or hit an obstacle:

$$\forall t, q(t) \notin \mathcal{A}(t). \quad (6)$$

We can observe that there are also states $(q, \dot{q})$ where the collision, the fall, the tipping over, although not having occurred yet, is going to happen no matter the future actions

of the system because the physical constraints (4) or (5) make any alternative impossible. Such states where the problem has become unavoidable, where it is too late to react, will be labeled as *non viable* [1], [9], [17] and gathered in a set $\overline{\mathcal{V}}(t)$ defined in the following way,

$$(q(t_0), \dot{q}(t_0)) \in \overline{\mathcal{V}}(t_0) \Leftrightarrow \forall u(\cdot) \in \mathcal{U}, \forall \lambda(\cdot) \in \Lambda(\cdot), \exists t \geq t_0 \text{ s.t. } q(t) \in \mathcal{A}(t). \quad (7)$$

In the case of obstacle avoidance, these states have regularly been called inevitable collision states [4], [7].

It appears therefore that on top of avoiding the set $\mathcal{A}(t)$, the mobile system should also absolutely avoid the set $\overline{\mathcal{V}}(t)$. Let's define then the *viability kernel* $\mathcal{V}(t)$ [1], [9], [17] gathering all the *viable* states, from where the system can still avoid falling, colliding obstacles or tipping over through adequate actions (Fig. 1):

$$(q(t_0), \dot{q}(t_0)) \in \mathcal{V}(t_0) \Leftrightarrow \exists u(\cdot) \in \mathcal{U}, \lambda(\cdot) \in \Lambda(\cdot) \text{ s.t. } \forall t \geq t_0, q(t) \notin \mathcal{A}(t). \quad (8)$$

This set appears to be the complementary set to $\mathcal{A}(t) \cup \overline{\mathcal{V}}(t)$. It can be easily shown to be controlled invariant, i.e. it could be made invariant under adequate control [3]. It is in fact the largest controlled invariant set not intersecting the set $\mathcal{A}(t)$ [1].

Going out of the viability kernel immediately implies that a fall, a collision or a tip over is going to happen unavoidably. Staying inside it secures the possibility to avoid such issues, what can be considered therefore as the primary goal of any mobile system. One problem however with this viability kernel is that it is hard to connect to practical numerical solutions. Indeed, establishing whether a given state is viable or not amounts to determining whether *there exists* a way to avoid the undesirable states. It is always possible to detect specific classes of viable states, for example equilibrium points and cyclic movements (Figure 1), but in the general case, with a dynamics as complex as the one of a humanoid robot or a car, it is globally impossible to establish numerically such a property.

C. Model Predictive Control

Model Predictive Control (MPC) [12] is a feedback control scheme that globally amounts to repeatedly solving online a series of Optimal Control problems, always taking into account the latest observation of the real state of the system. It usually takes the form of minimizing at every time t_k a cost function $\mathcal{L}(\dots)$, considering a prediction of the dynamics (4) over a horizon of length T (finite or infinite):

$$\min \int_{t_k}^{t_k+T} \mathcal{L}(q(t), \dot{q}(t), \ddot{q}(t), u(t), \lambda(t)) dt. \quad (9)$$

The control $u(t)$ that results from this optimization is applied then to the system until the next observation time t_{k+1} (which is supposed to arrive before the end of the prediction horizon $t_k + T$).

Among the traditionnal advantages of this scheme is its flexibility in formulating the control objectives and its

capacity to make use of the real limitations of complex dynamical models such the Differential Inclusion (4), what can lead to significant gains in performance. Its traditional limitation is its potentially high computational demand when having to solve online and repeatedly the underlying Optimal Control problems.

Even though probably not the only solution, Model Predictive Control obviously appears as an interesting approach here because it is supposed to explicitly work with predictions of the future outcome of decisions made in the present. One could express directly then the constraint that this future outcome shouldn't lead to any undesirable state. Easy to say, not easy to do: before even considering the minimization of the cost function (9), that would imply to be able to check whether a given state is viable or not, and we have already seen that this is not computable in the general case: we need to be more subtle.

III. AVOIDING TO FALL IN THE CASE OF LEGGED LOCOMOTION

A. The Point Mass Model

The Newton equation that appears in the necessary condition (5) exclusively describes the motion of the Center of Mass of the whole system [18]: in this equation, the dynamics of the system is strictly equivalent to that of a point mass. On the contrary, in the case of legged locomotion, every detail of the postural motion can have an effect on the Euler equation through nonholonomic couplings already discussed in section II-A [18]. For this reason, not taking into account the whole posture of a legged system when analyzing its locomotion will always lead to an approximation.

The simplest approximation is to consider the whole legged robot only as a point mass, not taking into account any rotational effects [8]. For being able to consider at least simple rotational effects such as the bending of the torso which is sometimes helpful for keeping balance, it has been proposed to consider the whole legged robot as a rigid body with constant or varying inertia parameters instead of only a point mass [15], [11].

But this rigid body approximation still doesn't convey any of the nonholonomic couplings which are necessary for relating the posture of the robot and its orientation: the rotations of a rigid body are absolutely impossible to relate to those of an articulated body, there is an irreducible discrepancy between them. The most obvious example of this discrepancy is that a rigid body can't always fall back on its feet as cats easily do thanks to postural motions. A more subtle example is that a non-zero mean momentum of rotation can be observed in walking motions [18]: reproducing a walking motion with a rigid body would imply therefore a continuously rotating rigid body, not very easy to relate to a constantly upright torso!

The gain in using a rigid body approximation instead of a point mass approximation is therefore not clear: it is more complex and it may not be more meaningful. The only way to capture properly the relationship between the posture and the orientation of a legged robot would be a multi-body model

that would convey at least minimal nonholonomic couplings, for example a multi-point masses model, at least one for the torso and one for each limb.

Focusing back on the simple point mass approximation, the deviation between this approximation and the complete model during "standard" humanoid walking motions has been observed to be "acceptable" with respect to the necessary condition (5). For example, the difference between the Center of Pressure in one case and in the other appears to be generally less than 2 cm , what can be easily taken into account by introducing proper safety margins [19]. This surprisingly low discrepancy may be understood under the light of the observation that the momentum of rotation of humanoid walking motions appears to be quite low [14]: rotational effects can therefore be temporarily put aside as long as we consider "standard" walking motions.

This point mass approximation appears therefore to be a good choice for an initial analysis of walking motions. In that case, with x_G the position of the Center of Mass, the left-hand side of the necessary condition (5) boils down to [17]

$$m \begin{bmatrix} \ddot{x}_G + g \\ x_G \times (\ddot{x}_G + g) \end{bmatrix}. \quad (10)$$

B. Quick Viability analysis of the Cart-Table Model

Let's consider now the specific case of legged locomotion on a flat ground, with all the points of contact with the environment located on the same horizontal surface with the same friction coefficient μ_0 , and with the Center of Mass moving strictly horizontally at a constant altitude h_G above this surface. In this case, we can extract from the generic necessary condition (5) for the point mass model (10) two independant necessary conditions. The first one is related to friction,

$$\|\ddot{x}_G\| \leq \mu_0 g, \quad (11)$$

and appears to be of secondary importance so we won't consider it any longer here. The second one is related to unilateral contact, the fact that the feet can only push on the ground,

$$x_G - \frac{h_G}{g} \ddot{x}_G \in Z(t), \quad (12)$$

where $Z(t)$ is the convex hull of the contact points on the ground at time t [17].

One can recognize that the left-hand side of this inclusion is nothing else but the Center of Pressure (CoP) of the contact forces, sometimes called the Zero Moment Point [17]. This second necessary condition states therefore very classically that this CoP lies within the convex hull of the contact points. This linear Differential Inclusion is sometimes called a Linearized Inverted Pendulum model since this is exactly how the CoM behaves with respect to the, or the Cart-Table model where the CoM is seen as a Cart moving on a horizontal Table, the feet of which correspond to the contact points on the ground [8] (Figure 2).

Fig. 2. The Linearized Inverted Pendulum model on the left, the Cart-Table model on the right, two ways of seeing the same linear approximate model of a legged robot.

Let's focus now on the closed convex hull of all the present and future contact points at each time t ,

$$\tilde{Z}(t) = \overline{\text{conv}} \bigcup_{\tau \geq t} Z(\tau).$$

We can legitimately consider that every locomotion starts from an equilibrium state and ends at an equilibrium state. Quickly investigating therefore the equilibrium states of the Differential Inclusion (12), we can trivially conclude that they correspond to states where the CoM lies above the support polygon, and therefore above the set $\tilde{Z}(t)$,

$$x_G \in Z(t) \subset \tilde{Z}(t).$$

Let's see now what happens when the CoM leaves this set. More precisely, let's consider a state of the system at a time t_0 where the CoM is lying on the edge of this convex hull, with a velocity pointing outwards. Since this convex hull is convex, we can define a line by a vector a and a constant b such that the whole set $\tilde{Z}(t_0)$ lies on one side of this line (Figure 3),

$$\forall z \in \tilde{Z}(t_0), \quad a^T z + b \leq 0,$$

the CoM (on the edge of the convex hull) lies exactly on this line,

$$a^T x_G(t_0) + b = 0,$$

and the outward pointing velocity corresponds to

$$a^T \dot{x}_G(t_0) > 0.$$

Combining the first inequality with the Differential Inclusion (12), we obtain an Ordinary Differential Inequality

$$\forall t \geq t_0, \quad \frac{h_G}{g} a^T \ddot{x}_G \geq a^T x_G + b$$

that can be solved analytically to give us that

$$a^T x_G(t) \geq a^T \dot{x}_G(t_0) \sqrt{\frac{h_G}{g}} \sinh \left(\sqrt{\frac{g}{h_G}} (t - t_0) \right) - b,$$

Fig. 3. An example of convex hulls $Z(t)$ in dotted line and $\tilde{Z}(t)$ in plain line for five steps of biped walking. When the CoM lies on the edge of the set $\tilde{Z}(t)$, it lies on a line (in red) defined by an orthogonal vector a (in blue) such that the whole set $\tilde{Z}(t)$ lies on one side of this line.

leading to the unavoidable conclusion that the CoM is diverging infinitely away from the set $\tilde{Z}(t)$ in the direction of the vector a ,

$$\lim_{t \rightarrow +\infty} a^T x_G(t) = +\infty.$$

Every locomotion can be considered to start with a CoM at rest within the set $\tilde{Z}(t)$, and as soon as this CoM leaves this set, it diverges infinitely away from it. We can observe that such a divergence to infinity of this Cart-Table approximate model corresponds in fact to a fall for the original legged robot (Figure 2): avoiding the legged robot to fall means therefore avoiding the Cart-Table model to diverge to infinity. From the point of view of the Viability analysis discussed in section II-B, the undesirable states for this Cart-Table model appear therefore to be the states “at infinity”.

This viability analysis could be led further, and approximations of the corresponding viability kernel may be obtained, but it is time now for Model Predictive Control to come into play.

C. Minimizing the derivatives of the position of the Center of Mass

One point of interest in the previous analysis is that when diverging to infinity, the CoM diverges at a pace greater than exponential, what implies that all its derivatives, velocity, acceleration, jerk, etc... diverge at the same pace, greater than exponential. Integrals of norms of these derivatives such as

$$\int_{t_k}^{+\infty} \|x_G^{(n)}(t)\|^2 dt \quad (13)$$

appear therefore to have infinite values for all diverging motions, and finite values for all motions ending up in finite time at an equilibrium state, the kind of locomotion we're generally interested in. Minimizing such integrals implicitly selects therefore a non-diverging motion whenever possible, what solves exactly the problem discussed in section II-B of selecting viable states out of non-viable states. This is quite a striking discovery that minimizing any derivative of the position of the CoM always generates a safe legged locomotion whenever possible!

On top of that, the exponential rates in the divergence of the CoM also imply that the values of the same integrals but

over finite time intervals,

$$\int_{t_k}^{t_k+T} \|x_G^{(n)}(t)\|^2 dt, \quad (14)$$

raise exponentially as well, allowing to discard most of the non-viable states with the same minimization process even on finite time intervals. Of course, the longer the time interval T , the more non-viable states can be discarded.

Now, the minimization of these integrals can obviously be put in the form of a Model Predictive Control scheme as discussed in section II-C, leading directly to an exact viable feedback in the case of the infinite integrals (13) (exact with respect to the Cart-Table model) or to an approximate viable feedback in the case of the finite integrals (14). The specificity of this MPC scheme is that the viability property need not be expressed and checked explicitly, it is satisfied implicitly when minimizing the cost functions. And interestingly enough, it appears that with such an MPC scheme, considering short prediction horizons, no more than 1 s, what implies looking only one step ahead, is already enough in preserving viability and generating safe locomotion even in the presence of strong perturbations [19].

The fact that minimizing the jerk of the CoM generates stable walking motions has already been acknowledged in [8], but through an optimal Linear Quadratic Regulator scheme only related to Lyapunov stability and therefore inducing viability only locally. This scheme would be unable therefore to face properly strong perturbations. This minimization of the jerk of the CoM has been generalized then in [19], following the MPC scheme presented here, minimizing the cost function (14) with respect to the dynamics (12). The strong connection between this scheme and the viability property that we presented here wasn't discussed there, but numerical simulations were proposed, demonstrating its ability to face strong perturbations properly.

In these previous works, the foot steps and therefore the convex hulls $Z(t)$ were considered to be fixed in advance and impossible to change. This limits of course the possibilities for the legged system to maintain viability and avoid falling. We can observe in the analysis presented here that the set-valued function $Z(t)$ can be included in the control variables without any other modification of the feedback scheme: modifying foot placement in order to minimize the jerk of the CoM will continue to generate stable walking motions whenever possible, as shown in [5].

IV. DISCUSSION

We have seen in the previous section how an exact viable feedback control law can be designed for an approximate model of legged locomotion through a very simple Model Predictive Control scheme. We have seen that accurate and perfectly tractable approximations can also be obtained very easily by varying the length of the corresponding horizon of prediction.

Viability and Model Predictive Control are not unknown concepts in the field of wheeled locomotion, being for obstacle avoidance, where non viable states are sometimes

called Inevitable Collision states [9], [7], [4], or for avoiding tipping over on rough terrain [10], [16]. Similarly to what has been done in the previous section for legged locomotion, these analyses of wheeled locomotion are generally based on simple approximate models of cars, with the notable exception of [10] where is advocated the fact that incomplete models imply possible unpredicted dynamical effects and therefore possible counter-productive planned actions.

The approximate model common to all these works is generally nothing else but the simple Point Mass model discussed in section III-A, with various approximations of the steering process of similar complexity. When aggressive maneuvers on rough terrain are considered, explicitly taking into account the risk of tipping over or sideslipping, conditions on the contact forces need to be precised in more details, leading to inequalities similar to (11) and (12), but of different use since the position x_G of the CoM is usually considered to be fixed with respect to the contact points and their convex hull $Z(t)$ [16].

The models considered in these analyses are simple enough to allow a brute force approach of the viability conditions, discretizing the state and control spaces and checking the outcome of all possibilities, with proper pruning of the obviously viable or non viable cases to help make computations faster [9], [7], [4], [10], [16]. But it appears that at best, only kinetic effects related to forward speed and curvature of the trajectory are taken into account so far in these viability analyses [16], while other dynamical effects such as those related to acceleration are not, what might be necessary, as advocated in [10]. The problem then is that the brute force approach considered there is very likely to be unable to cope with more complex models when they become necessary because of the exponential increase in the corresponding computational demands.

On the contrary, the viable control law described in the previous section for legged locomotion doesn't focus explicitly on the viability condition: viability can be seen as a side effect of the proposed MPC scheme. And this MPC scheme, which only amounts to minimizing any derivative of the position of the CoM, can be applied without any problem to more complex models, and continue to convey viability as a side effect. Of course, more complex models means more computational demands, but there exist extremely fast numerical algorithms for solving Nonlinear MPC problems [6], and the increase in computations is in no way comparable to what appears with brute force approaches.

One can wonder therefore if an approach similar to the one proposed here for legged locomotion, circumventing the numerical difficulties inherent to the concept of viability, obtaining viability as a side effect, wouldn't be possible for the problems of avoiding tipping over, or avoiding obstacles. This is of course an open question today.

REFERENCES

- [1] J.-P. Aubin. *Viability Theory*. Birkhäuser, 1991.

- [2] C. Azevedo, P. Poignet, and B. Espiau. Moving horizon control for biped robots without reference trajectory. In *IEEE International Conference on Robotics and Automation*, pages 2762–2767, Washington, USA, May 2002.
- [3] F. Blanchini. Set invariance in control. *Automatica*, 35(11):1747–1768, 1999.
- [4] N. Chan, M. Zucker, and J. Kuffner. Towards safe motion planning for dynamic systems using regions of inevitable collision. In *Proceedings of the IEEE International Conference on Robotics & Automation, Workshop on Collision-free Motion Planning for Dynamic Systems*, 2007.
- [5] H. Diedam, D. Dimitrov, P.-B. Wieber, K. Mombaur, and Diehl. M. Online walking gait generation with adaptive foot positioning through linear model predictive control. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots & Systems*, 2008.
- [6] M. Diehl, H.G. Bock, and J.P. Schlöder. A real-time iteration scheme for nonlinear optimization in optimal feedback control. *SIAM J. Control Optim.*, 43(5):1714–1736, 2005.
- [7] T. Fraichard and H. Asama. Inevitable collision states, a step towards safer robots? *Advanced Robotics*, 2004.
- [8] S. Kajita, F. Kanehiro, K. Kaneko, K. Fujiwara, K. Harada, K. Yokoi, and H. Hirukawa. Biped walking pattern generation by using preview control of zero-moment point. In *Proceedings of the IEEE International Conference on Robotics & Automation*, 2003.
- [9] M. Kalisiak and M. van de Panne. Approximate safety enforcement using computed viability envelopes. In *Proceedings of the IEEE International Conference on Robotics & Automation*, 2004.
- [10] A. Kelly and A. Stentz. Rough terrain autonomous mobility, part 2: An active vision, predictive control approach. *Autonomous Robots*, 5:163–198, 1998.
- [11] S.-H. Lee and A. Goswami. Reaction mass pendulum (RMP): An explicit model for centroidal angular momentum of humanoid robots. In *Proceedings of the IEEE International Conference on Robotics & Automation*, 2007.
- [12] D.Q. Mayne, J.B. Rawlings, C.V. Rao, and P.O.M. Scokaert. Constrained model predictive control: stability and optimality. *Automatica*, 26(6):789–814, 2000.
- [13] S.C. Peters and K. Iagnemma. An analysis of rollover stability measurement for high-speed mobile robots. In *Proceedings of the IEEE International Conference on Robotics & Automation*, 2006.
- [14] M. Popovic, A. Hofmann, and H. Herr. Zero spin angular momentum control: definition and applicability. In *International Conference on Humanoid Robotics*, 2004.
- [15] J. Pratt, J. Carff, S. Drakunov, and A. Goswami. Capture point: A step toward humanoid push recovery. In *International Conference on Humanoid Robotics*, 2006.
- [16] M. Spenko, Y. Kuroda, S. Dubowsky, and K. Iagnemma. Hazard avoidance for high-speed mobile robots in rough terrain. *Journal of Field Robotics*, 13(5):311–331, 2006.
- [17] P.-B. Wieber. On the stability of walking systems. In *Proceedings of the International Workshop on Humanoid and Human Friendly Robotics*, 2002.
- [18] P.-B. Wieber. Holonomy and nonholonomy in the dynamics of articulated motion. In *Proceedings of the Ruperto Carola Symposium on Fast Motion in Biomechanics and Robotics*, 2005.
- [19] P.-B. Wieber. Trajectory free linear model predictive control for stable walking in the presence of strong perturbations. In *International Conference on Humanoid Robotics*, 2006.