

THE ART TO KEEP IN TOUCH

The “good use” of Lagrange Multipliers

Antoine Jonquet, Olivier Nocent

Department of Computer Science, Institute of Technology, Reims University, rue des crayères, 51687 Reims cedex 2, France

Yannick Remion

Department of Computer Science, Institute of Technology, Reims University, rue des crayères, 51687 Reims cedex 2, France

Keywords: Physically-based animation, constraints, contact simulation.

Abstract: Physically-based modeling for computer animation has become a wide spread framework because it allows to produce more realistic motions in less time without requiring the expertise of skilled animators. But, a computer animation is not only a numerical simulation based on classical mechanics since it follows a precise story-line. One common way to define aims in an animation is to add geometric constraints. There are several methods to manage these constraints within a physically-based framework. In this paper, we present an algorithm for constraints handling based on Lagrange multipliers. After few remarks on the equations of motion that we use, we present a first algorithm proposed by Platt. We show with a simple example that this method is not reliable. Our contribution consists in improving this algorithm to provide an efficient and robust method to handle simultaneous and antagonist constraints.

1 INTRODUCTION

For about two decades, the computer graphics community has investigated the field of physics in order to produce more and more realistic computer animations. In fact, physically-based modeling in animation allows to generate stunning visual effects that would be extremely complex to reproduce *manually*. On one hand, the addition of physical properties to 3D objects automates the generation of motion just by specifying initial external forces. On the other hand, physically-based animations are even more realistic than traditional key-framed animations that require the expertise of many skilled animators. As a consequence, the introduction of physically-based methods in modeling and animation significantly reduced the cost and production time of computer generated movies. But, one main drawback of this kind of framework is that it relies on heavy mathematics usually hard to tackle for a computer scientist. A second main disadvantage concerns the input of a physically-based animation: in fact, forces and torques are not really *user-friendly* since it is really difficult to anticipate a complex motion just by specifying an initial set of external forces.

A computer animation is definitely not a numer-

ical simulation because it follows a story-line. According to Demetri Terzopoulos (Terzopoulos et al., 1989), an animation is simulation plus control. One way to ensure that the objects fulfill the goals defined by the animator is to use geometric constraints. A constraint is an equality or an inequality that gathers different parameters of the animation like the total time elapsed, the positions or the orientations of the moving objects. In a less general way, mechanical simulations also benefit from the use of constraints in order to prevent interpenetration between physical objects for example.

There are several methods to handle constraints, summarized in a survey paper by Baraff (Baraff, 1993). But, since our research work is mostly devoted to mechanical simulation, we decided to focus on the use of *Lagrange multipliers* to manage the geometric constraints. In fact, numerical simulations require robust and reliable techniques to ensure that the constraints are never violated. Moreover, with this method we are also able to measure the amount of strain that is necessary to fulfill a given constraint. In this paper, we present a novel algorithm to manage efficiently several simultaneous active geometric constraints. We begin by detailing the physical equa-

tions that we use before presenting Platt's algorithm (Platt, 1992) that is the only algorithm of this type based on Lagrange multipliers. With a simple example, we demonstrate that this algorithm is not suitable for handling simultaneous active constraints. We then introduce our own contribution in order to show how to improve Platt's algorithm to make it reliable, robust and efficient.

2 LAGRANGE EQUATIONS OF MOTION

Lagrangian dynamics consist in an extension of Newtonian dynamics allowing to generate a wide range of animations in a more efficient way. In fact, Lagrange equations of motion rely on a set of unknowns, denoted as a state vector \mathbf{x} of *generalized coordinates*, that identifies the real degrees of freedom (DOF) of the mechanical systems involved. Within this formalism, the DOF are not only restricted to rotations or translations. For example, a parameter $u \in [0, 1]$ which gives the relative position of a point along a 3D parametric curve can be considered as a single generalized coordinate.

2.1 Unconstrained Motion

The evolution of a free mechanical system only subject to a set of external forces is ruled by the Lagrange equations of motion (1).

$$\mathbf{M} \ddot{\mathbf{x}} = \mathbf{f} \quad (1)$$

\mathbf{M} is the *mass matrix*. $\ddot{\mathbf{x}}$ is the second time derivative of the state vector. Finally, the vector \mathbf{f} corresponds to the sum of external forces. For more details concerning this formalism, we suggest to read (Goldstein, 1980) and (Arnold, 1989).

2.2 Constrained Motion

By convention, an equality constraint will always be defined as in equation (2) where \mathcal{E} is the set of indices of all the equality constraints.

$$g_k(\mathbf{x}) = 0 \quad \forall k \in \mathcal{E} \quad (2)$$

Constraints restrict the set of reachable configurations to a subspace of \mathbb{R}^n where n is the total number of degrees of freedom. As mentioned before, there exists three main methods to integrate constraints in equation (1)

The projection method consists in modifying the state vector \mathbf{x} and its first time derivative $\dot{\mathbf{x}}$ in order to fulfill the constraint. This modification can be performed with an iterative method like the Newton-Raphson method (Press et al., 1992). Even if this method is very simple and seems to ensure an instantaneous constraint fulfillment, it is not robust enough: indeed it can not guarantee that the process converges in the case of simultaneous active constraints.

The penalty method adds new external forces, acting like virtual springs, in order to minimize the square of the constraint equation, considered as a positive energy function. The main advantage of this method is its compatibility with any dynamic engine since it only relies on forces. But this method leads to inexact constraint fulfillment, allowing interpenetration between the physical objects. In order to diminish this interpenetration, the stiffness of the virtual springs must be significantly increased, making the numerical system unstable.

The Lagrange method consists in calculating the exact amount of strain, denoted as the *Lagrange multiplier*, needed to fulfill the constraint. This method guarantees that constraints are always fulfilled. Since the use of Lagrange multipliers introduces a set of new unknowns, the equation (1) must be completed by a set of new equations, increasing the size of the initial linear system to solve. But we consider that this method is the only suitable to manage efficiently the geometric constraints.

For all the reasons mentioned above, we choose the Lagrange method to manage our geometric constraints. According to the principle of virtual work, each constraint g_k adds a new force perpendicular to the tangent space of the surface $g_k(\mathbf{x}) = 0$. The Lagrange multiplier λ_k corresponds to the intensity of the force related to the constraint g_k . With these new forces, the equation (1) is modified as follows:

$$\mathbf{M} \ddot{\mathbf{x}} = \mathbf{f} + \sum_{k \in \mathcal{E}} \lambda_k \frac{\partial g_k}{\partial \mathbf{x}} \quad (3)$$

We add new equations to our system by calculating the second time derivative of equation (2), leading to equation (4).

$$\sum_{i=1}^n \frac{\partial g_k}{\partial x_i} \ddot{x}_i = - \sum_{i,j=1}^n \frac{\partial^2 g_k}{\partial x_i \partial x_j} \dot{x}_i \dot{x}_j \quad \forall k \in \mathcal{E} \quad (4)$$

In order to correct the numerical deviation due to round-off errors, Baumgarte proposed in (Baumgarte,

1972) a constraint stabilization scheme illustrated by the equation (5). The parameter τ^{-1} can be seen as the speed of constraint fulfillment.

$$\sum_{i=1}^n \frac{\partial g_k}{\partial x_i} \ddot{x}_i = - \sum_{i,j=1}^n \frac{\partial^2 g_k}{\partial x_i \partial x_j} \dot{x}_i \dot{x}_j - \frac{2}{\tau} \sum_{i=1}^n \frac{\partial g_k}{\partial x_i} \dot{x}_i - \frac{1}{\tau^2} g_k \quad (5)$$

When we mix the equations (1) and (5), we obtain a linear system where the second time derivative of the state vector \mathbf{x} and the vector of Lagrange multipliers Λ are the unknowns.

$$\begin{bmatrix} \mathbf{M} & -\mathbf{J}^T \\ -\mathbf{J} & \mathbf{0} \end{bmatrix} \begin{pmatrix} \ddot{\mathbf{x}} \\ \Lambda \end{pmatrix} = \begin{pmatrix} \mathbf{f} \\ -\mathbf{d} \end{pmatrix} \quad (6)$$

\mathbf{J} is the jacobian matrix of all the geometric constraints and \mathbf{d} corresponds to the right term of equation (5).

2.3 Inequality Constraints Management

By convention, an inequality constraint will always be defined as in equation (7) where \mathcal{F} is the set of indices of all the inequality constraints.

$$g_k(\mathbf{x}) \geq 0 \quad \forall k \in \mathcal{F} \quad (7)$$

For a given state vector \mathbf{x} , we remind the following definitions:

- the constraint is said to be *violated* by \mathbf{x} when $g_k(\mathbf{x}) < 0$. This means that the state vector \mathbf{x} corresponds to a non allowed configuration.
- the constraint is said to be *satisfied* by \mathbf{x} when $g_k(\mathbf{x}) \geq 0$.
- the constraint is said to be *active* by \mathbf{x} when $g_k(\mathbf{x}) = 0$. In this case, the state vector \mathbf{x} belongs to the boundary of the subspace defined by the inequality constraint g_k .

The management of inequality constraints is more difficult than the management of equality constraints. An inequality constraint must be handled only if it is violated or active. In fact, the algorithm is a little more complicated as we explain in the next sections.

That is why we define two subsets within \mathcal{F} : \mathcal{F}^+ is the set of indices of all handled inequality constraints and \mathcal{F}^- is the set of indices of ignored inequality constraints. Finally, we have $\mathcal{F} = \mathcal{F}^- \cup \mathcal{F}^+$. The jacobian matrix of constraints \mathbf{J} of the equation (6) is built from all the constraints g_k where $k \in \mathcal{E} \cup \mathcal{F}^+$.

3 PREVIOUS WORK

Within the computer graphics community, the main published method devoted to inequality constraints management using Lagrange multipliers, known as “Generalized Dynamic Constraints”, was proposed by Platt in (Platt, 1992). In his paper, he describes how to use Lagrange multipliers to assemble and simulate collisions between numerical models. This method is an extension of the work of Barzel and Barr (Barzel and Barr, 1988) that specifies how constraints must be satisfied. Moreover, Platt proposes a method to update \mathcal{F}^+ (the set of handled inequality constraints) during the animation. This algorithm can be compared to classical *active set methods* (Björck, 1996; Nocedal and Wright, 1999).

We do not focus on collision detection that is a problem by itself. We are aware that this difficult problem can be solved in many ways, we encourage the reader to refer to the survey paper by Teschner et al. (Teschner et al., 2005). During the collision detection stage, we assume that the dynamic engine may rewind time until the first constraint activation is detected. In any case, this stage ensures that constraints are never violated. But, it is possible that several constraints are activated simultaneously. The main topic of this paper is to provide a reliable algorithm to handle these multiple active constraints in an efficient way.

At the beginning of the animation, Platt populates the set \mathcal{F}^+ with all the active constraints.

Algorithm 1– Platt’s algorithm.

- 1 Solve equation (6) to get $\ddot{\mathbf{x}}$ and Λ
 - 2 Update \mathbf{x} and $\dot{\mathbf{x}}$ (numerical integration)
 - 3 **for each** $k \in \mathcal{F}$ **do**
 - 4 **if** $k \in \mathcal{F}^+$ **then**
 - 5 **if** $\lambda_k \leq 0$ **then**
 - 6 k is moved to \mathcal{F}^-
 - 7 **else**
 - 8 **if** $g_k(\mathbf{x}) < 0$ **then**
 - 9 g_k is moved to \mathcal{F}^+
-

For each time step, according to algorithm 1, we solve equation (6) and update the state vector in order to retrieve new positions and velocities at the end of the current time step. We then check the status of each inequality constraint. If a constraint g_k is active, it is still handled until its Lagrange multiplier is negative or null, that is to say that the Lagrange multiplier corresponds to a force that prevents from deactivation. According to the new values of the state vector \mathbf{x} , if

the previously inactive constraint g_k is now violated ($g_k(\mathbf{x}) < 0$), the constraint must be added to \mathcal{F}^+ in order to prevent the system to entering in such a configuration.




Figure 1: A simple example with two simultaneous active constraints.

Even if this algorithm seems to give a reliable solution for inequality constraints handling, some problems remain. We set up a simple scene as in figure 1 to illustrate the insufficiencies of Platt's method. A particle of mass m is constrained to slide on a 2D plane. It starts from an acute-angle corner modeled by two linear inequality constraints $g_1(\mathbf{x}) = x - y$ and $g_2(\mathbf{x}) = y$. Finally, this particle is subject to a single external force $\mathbf{f} = (2, -1)$. In this particular case, the state vector \mathbf{x} is composed of the 2D coordinates (x, y) of the particle. According to equation (1), the generalized mass matrix for this system is defined by:

$$\mathbf{M} = \begin{bmatrix} m & 0 \\ 0 & m \end{bmatrix} \quad (8)$$

As the geometric constraints $g_1(\mathbf{x})$ and $g_2(\mathbf{x})$ are linear, their first and second time derivative do not produce any deviation term defined in equation (5):

$$\mathbf{d} = \begin{pmatrix} 0 \\ 0 \end{pmatrix} \quad (9)$$

According to the initial value of the state vector $\mathbf{x} = (0, 0)$, the two constraints $g_1(\mathbf{x})$ and $g_2(\mathbf{x})$ are active, so their indices are inserted in \mathcal{F}^+ and \mathbf{J} , the jacobian matrix of constraints, is defined as follows:

$$\mathbf{J} = \begin{bmatrix} 1 & -1 \\ 0 & 1 \end{bmatrix} \quad (10)$$

From equations (6) (8) (9) and (10), we obtain a linear system whose unknowns are the second time derivative of the state vector \mathbf{x} and the two Lagrange multipliers λ_1 and λ_2 associated with g_1 and g_2 :

$$\begin{bmatrix} m & 0 & -1 & 0 \\ 0 & m & 1 & -1 \\ -1 & 1 & 0 & 0 \\ 0 & -1 & 0 & 0 \end{bmatrix} \begin{pmatrix} \ddot{x} \\ \ddot{y} \\ \lambda_1 \\ \lambda_2 \end{pmatrix} = \begin{pmatrix} 2 \\ -1 \\ 0 \\ 0 \end{pmatrix} \quad (11)$$

The solutions are $\ddot{\mathbf{x}} = (0, 0)$ and $\Lambda = (-2, -1)$. Because \ddot{x} and \ddot{y} are null, the particle does not move during this time step. But, since λ_1 and λ_2 are both negative, their corresponding constraints are moved to \mathcal{F}^- . This means that, for the next time step, the system will be free of any constraints. As the force remains constant, the next value of $\ddot{\mathbf{x}}$ will be equal to $(2m^{-1}, -m^{-1})$. These values will lead to an illegal position of the particle, under the line $y = 0$.

The amount of violation of the constraint $g_2(\mathbf{x}) = y$ mainly depends on the ratio between the mass m of the particle and the intensity of the external force \mathbf{f} . The section 5 offers different results and comparisons.

4 OUR CONTRIBUTION

4.1 A First Approach

The problem of Platt's method relies on the fact that it keeps inequality constraints in \mathcal{F}^+ abusively. In fact, the condition $g_k(\mathbf{x}) < 0$ used to populate \mathcal{F}^+ with inequality constraint is not well suited and an other way must be proposed. A solution would be to replace the condition $g_k(\mathbf{x}) < 0$ by a violation tendency condition expressed as $J_k \ddot{\mathbf{x}} < d_k$. An active constraint that does not fulfill the violation tendency condition will be satisfied during the next time step and does not have to be handled.

At the beginning of the animation, we solve equation (1) to get $\ddot{\mathbf{x}}$ and we then populate the set \mathcal{F}^+ with the active constraints that fulfill the violation tendency condition. It is clear that we handle less constraints than Platt because our criteria is more restrictive.

Algorithm 2– Platt's improved algorithm.

- 1 Solve equation (6) to get $\ddot{\mathbf{x}}$ and Λ
 - 2 Update \mathbf{x} and $\dot{\mathbf{x}}$ (numerical integration)
 - 3 **for each** $k \in \mathcal{F}$ **do**
 - 4 **if** $k \in \mathcal{F}^+$ **then**
 - 5 **if** $\lambda_k \leq 0$ **then**
 - 6 k is moved to \mathcal{F}^-
 - 7 **else**
 - 8 **if** $J_k \ddot{\mathbf{x}} < d_k$ **then**
 - 9 k is moved to \mathcal{F}^+
-

We briefly verify that this algorithm gives a correct solution to our example illustrated in figure 1. According to equation (1), $\ddot{\mathbf{x}} = (2m^{-1}, -m^{-1})$. The two constraints g_1 and g_2 are active because $\mathbf{x} = (0, 0)$ but only g_2 fulfills the violation tendency condition as mentioned in equation (12).

$$\begin{aligned} J_1 \ddot{\mathbf{x}} &= 3m^{-1}, & g_1 \text{ is moved to } \mathcal{F}^- \\ J_2 \ddot{\mathbf{x}} &= -m^{-1}, & g_2 \text{ is moved to } \mathcal{F}^+ \end{aligned} \quad (12)$$

In this special case, equation (6) becomes:

$$\begin{bmatrix} m & 0 & 0 \\ 0 & m & -1 \\ 0 & -1 & 0 \end{bmatrix} \begin{pmatrix} \ddot{x} \\ \ddot{y} \\ \lambda_2 \end{pmatrix} = \begin{pmatrix} 2 \\ -1 \\ 0 \end{pmatrix} \quad (13)$$

The solutions of the linear system (13) are $\ddot{\mathbf{x}} = (2m^{-1}, 0)$ and $\lambda_2 = 1$. Finally, the particle will slide along the x -axis without crossing the line $y = 0$ because the constraint g_1 that was not handle did not introduce a false response.

This new algorithm seems to manage multiple inequality constraints in a good way, but we could highlight a problem in this algorithm by using the same example illustrated in figure 1 with a new external force $\mathbf{f} = (-1, -2)$.

At the beginning, since $\mathbf{x} = (0, 0)$, the constraints g_1 and g_2 are active. From equation (1), we obtain that $\ddot{\mathbf{x}} = (-m^{-1}, -2m^{-1})$, and from equation (14) that only the constraint g_2 is handled.

$$\begin{aligned} J_1 \ddot{\mathbf{x}} &= m^{-1}, & g_1 \text{ is moved to } \mathcal{F}^- \\ J_2 \ddot{\mathbf{x}} &= -2m^{-1}, & g_2 \text{ is moved to } \mathcal{F}^+ \end{aligned} \quad (14)$$

According to equation (14), we build the linear system (15).

$$\begin{bmatrix} m & 0 & 0 \\ 0 & m & -1 \\ 0 & -1 & 0 \end{bmatrix} \begin{pmatrix} \ddot{x} \\ \ddot{y} \\ \lambda_2 \end{pmatrix} = \begin{pmatrix} -1 \\ -2 \\ 0 \end{pmatrix} \quad (15)$$

The solutions are $\ddot{\mathbf{x}} = (-m^{-1}, 0)$ and $\lambda_2 = 2$. Then, after the update of \mathbf{x} and $\dot{\mathbf{x}}$, the particle goes through the plane defined by the constraint g_1 and reaches an illegal state. This is due to the fact that the Lagrange multiplier λ_2 pushes the system in an illegal state from the constraint g_1 point of view, which was not previously inserted in the system 6 as its initial tendency acceleration tends to disconnect the particle from the $x - y = 0$ plane.

At the end of the step, we check inequality constraints for the next animation step, as $\lambda_2 \leq 0$, g_2 stays in \mathcal{F}^+ and $J_1 \ddot{\mathbf{x}} < d_1$, then g_1 move to \mathcal{F}^+ . So at the next step, the linear system will be correctly set, but it is too late since the system is already in an illegal state.

4.2 The “Right” Algorithm

The use of the violation tendency condition $J_k \ddot{\mathbf{x}} < d_k$ improves simultaneous active constraints management, since only the appropriate inequality constraints are handled by the equation (6). But we have seen, from the previous example, that it is insufficient to reach the accurate configuration. In fact, the constraints from \mathcal{F}^+ that fulfill the violation tendency condition will produce a vector Λ of Lagrange multipliers that prevent the system from being in an illegal configuration. In the meantime, the acceleration tendency $\ddot{\mathbf{x}}$ of the system could lead to another illegal configuration. The only way to deal with this problem is to use the new acceleration tendencies to test if the other active inequality constraints g_k fulfill the violation tendency condition and have to be handled. We then need an iterative process that computes the acceleration tendencies and test for other active inequality constraints g_k to insert in \mathcal{F}^+ , until the system reaches the appropriate state.

We propose a simple and efficient solution to the inequality constraints handling problem, based on the way physicists work. At the beginning of each time step, all active inequality constraints g_k are detected, and \mathcal{F}^+ is emptied. We then begin an iterative process that runs until there is no new insertion in \mathcal{F}^+ . The acceleration tendencies $\ddot{\mathbf{x}}$ are computed from the equation (6) and the violation tendency condition $J_k \ddot{\mathbf{x}} < d_k$ is tested on each active inequality constraint. For any inequality constraint g_k that fulfills the condition, we insert its index k in \mathcal{F}^+ and start another iterative step.

Algorithm 3– The right algorithm.

```

1 repeat
2   Solve equation (6) to get  $\ddot{\mathbf{x}}$  and  $\Lambda$ 
3   for each active constraint  $g_k$  do
4     if  $J_k \ddot{\mathbf{x}} < d_k$  then
5        $k$  is moved to  $\mathcal{F}^+$ 
6 until  $\mathcal{F}^+$  has not been updated;
7 Update  $\mathbf{x}$  and  $\dot{\mathbf{x}}$  (numerical integration)
    
```

Since we begin with no constraints and that only appropriate inequality constraints are inserted to \mathcal{F}^+ during the iterative process, the last computation of $\ddot{\mathbf{x}}$ and Λ from equation (6) will lead to an accurate configuration of the system. Moreover, this process guarantees the convergence towards an appropriate solution as we begin from an empty \mathcal{F}^+ and only insert constraints in \mathcal{F}^+ .

In a recent communication (Raghupathi and Faure, 2006), Raghupathi presented a method also

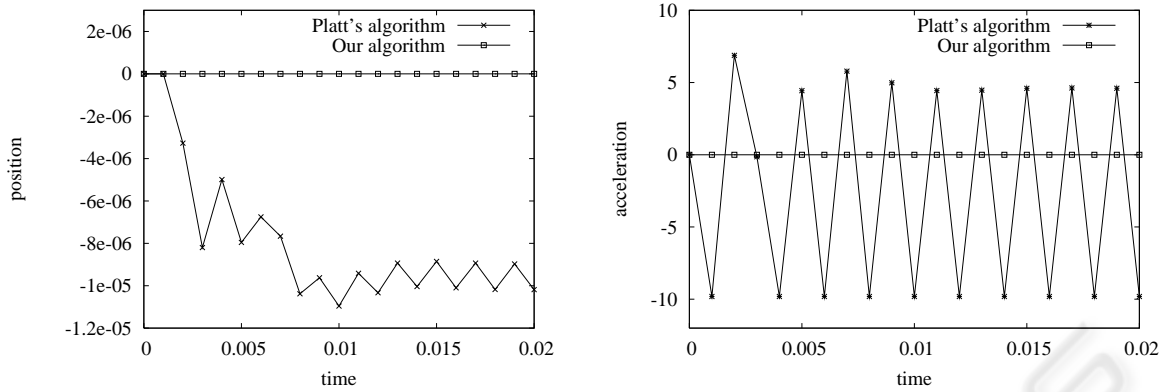


Figure 2: Comparison of Platt's algorithm and our method using the example illustrated in figure 1 with a mass $m = 2$. The numerical values correspond respectively to position and acceleration along the y -axis.




Figure 3: Comparison of Platt's algorithm and our method using the example illustrated in figure 1 with a mass $m = 3$. The numerical values correspond respectively to position and acceleration along the y -axis.

based on Lagrange multipliers. For realtime considerations, they do not allow the dynamic engine to rewind time to get back to the first constraint activation. They have to manage constraints at the end of the time step, trying to find the right accelerations to ensure constraints fulfillment. They also confess that this process is not guaranteed to converge in any situation.

5 RESULTS AND COMPARISONS

We will now compare the results obtained with Platt's algorithm and our method, using the example illustrated in figure 1. Figure 2 and 3 illustrate a comparison of the positions and accelerations along y -axis of a particle of mass $m = 2$ and $m = 3$. We recall that the inequality constraint g_2 forbids negative values for y and that the constant force \mathbf{f} applied to the particle

is equal $(2, -1)$.

As shown on figure 2, Platt's algorithm holds the particle abusively on the corner at the first time step, and releases it at the next time step. As a consequence, the particle evolves in an illegal state during the following steps. With a mass $m = 2$, the error related to the position is less than 10^{-5} with an oscillating acceleration (right column). But if we set the mass m to 3, as shown in the figure 3, errors are much more important, and the particle crosses the line $y = 0$ modeled by the constraint g_2 .

As illustrated, our algorithm keeps the particle along the x -axis within a controlled numerical error value, that is less than 10^{-8} in these examples.

To illustrate multiple contact constraints, we have set a billiard scene with 10 balls placed in triangle in a corner and another ball that comes to them. For each ball, we define two inequality constraints for each side of the corner and one inequality constraint related to the distance between each ball. This example is fi-

nally composed of 11 balls and 77 inequality constraints (figure 4).

6 CONCLUSION

In this paper, we presented a novel algorithm to manage simultaneous inequality constraints. Among all the existing methods to handle constraints within a physically-based animation, we focused on the Lagrange method which provides a reliable way to ensure that constraints are always exactly fulfilled. But, in the special case of several active inequality constraints, we have to take care on how to handle these simultaneous constraints. Platt proposed an algorithm based on Lagrange multipliers but we showed that this method is unable to solve even simple examples. We then explained how to improve this algorithm in order to propose a new reliable and efficient method for inequality constraints handling. Beyond the example illustrated in figure 1, we produced a short movie simulating a billiard game. Some snapshots are gathered in figure 4.

ACKNOWLEDGEMENTS

We would like to thank the Champagne-Ardenne regional council who supports this work, which is a part of the SYS-REEDUC project.

REFERENCES

- Arnold, V. I. (1989). *Mathematical Methods of Classical Mechanics*, volume 60 of *Graduate Texts in Mathematics*. Springer Verlag, New York, 2nd edition. 508 pages.
- Baraff, D. (1993). Non-penetrating rigid body simulation. In *State of the Art Reports, Eurographics '93*.
- Barzel, R. and Barr, A. H. (1988). A modeling system based on dynamic constraints. In *SIGGRAPH '88: Proceedings of the 15th annual conference on Computer graphics and interactive techniques*, pages 179–188, New York, NY, USA. ACM Press.
- Baumgarte, J. (1972). Stabilization of constraints and integrals of motion. *Computer Methods in Applied Mechanics and Engineering*, 1:1–16.
- Björck, Å. (1996). *Numerical Methods for Least Squares Problems*. SIAM, Philadelphia, Penn.
- Goldstein, H. (1980). *Classical Mechanics*. Addison-Wesley, Reading, MA, U.S.A., 2nd edition. 672 pages.
- Nocedal, J. and Wright, S. (1999). *Numerical Optimization*. Springer, New York.
- Platt, J. (1992). A generalization of dynamic constraints. *CVGIP: Graphical Models and Image Processing*, 54(6):516–525.
- Press, W. H., Flannery, B. P., Teukolsky, S. A., and Vetterling, W. T. (1992). *Numerical Recipes: The Art of Scientific Computing*. Cambridge University Press, Cambridge (UK) and New York, 2nd edition.
- Raghupathi, L. and Faure, F. (2006). QP-collide: A new approach to collision treatment. In *Journées du groupe de travail Animation et Simulation (GTAS)*, Annual French Working group on Animation and Simulation, pages 91–101. Institut de Recherche en Informatique de Toulouse.
- Terzopoulos, D., Platt, J., Barr, A., Zeltzer, D., Witkin, A., and Blinn, J. (1989). Physically-based modeling: past, present, and future. In *SIGGRAPH '89: ACM SIGGRAPH 89 Panel Proceedings*, pages 191–209, New York, NY, USA. ACM Press.
- Teschner, M., Kimmerle, S., Heidelberger, B., Zachmann, G., Raghupathi, L., Fuhrmann, A., Cani, M.-P., Faure, F., Magnenat-Thalmann, N., Strasser, W., and Volino, P. (2005). Collision detection for deformable objects. volume 24 of *Computer Graphics Forum*, pages 61–81.




Figure 4: A billiard game session illustrating our algorithm for constraints management (11 balls and 77 inequality constraints).