Graphical Models 76 (2014) 1–16

Contents lists available at ScienceDirect

Graphical Models

journal homepage: www.elsevier.com/locate/gmod

Optimal crowd editing ☆

Pierre Allain [a,*], Nicolas Courty [b], Thomas Corpetti [c]

[a] Inria, Rennes, France
[b] Université de Bretagne Sud, Vannes, France
[c] CNRS, Rennes, France

ARTICLE INFO

Article history:
Received 3 April 2013
Received in revised form 6 August 2013
Accepted 24 September 2013
Available online 8 October 2013

Keywords:
Crowd editing
Crowd simulation
Agent steering
Control
Adjoint methods
Optimal control strategy

ABSTRACT

Simulating realistic crowd behaviors is a challenging problem in computer graphics. Yet, several satisfying simulation models exhibiting natural pedestrians or group emerging behaviors exist. Choosing among these model generally depends on the considered crowd density or the topology of the environment. Conversely, achieving a user-desired kinematic or dynamic pattern at a given instant of the simulation reveals to be much more tedious. In this paper, a novel generic control methodology is proposed to solve this crowd editing issue. Our method relies on an adjoint formulation of the underlying optimization procedure. It is independent to a certain extent of the choice of the simulation model, and is designed to handle several forms of constraints. A variety of examples attesting the benefits of our approach are proposed, along with quantitative performance measures.

1. Introduction

Animating a crowd of thousand of individuals is a challenging task. Most of the time, human crowds exhibit very subtle and specific patterns. The variability of crowd dynamics and behaviors is a consequence of the diversity of the persons inside it (age, sex, social and psychological attributes), as well as the spatial configuration of obstacles and lanes. Nevertheless, several models exist, that can either rely on per-individual strategies, usually grouped under the denomination of steering behaviors or microscopic models, or in contrast on a global definition of the crowd flow and properties: the macroscopic models. Assessing the quality of those models is usually difficult, since several criteria can be evaluated: computation performances, presence of emerging behaviors, individual trajectories respecting the least effort principles, etc. In the context of civil engineering, those models can provide a lot of information about potentially dangerous areas or bottlenecks where problems are likely to occur. The context of graphics is slightly different since animators usually have a precise idea of what they want to show, which may differ from what is actually given by the simulation models. A possibility is then to tweak the different parameters of the used model, in a trial-and-error fashion, until the simulation is acceptable. It is easy to understand that this strategy, depicted in Fig. 1a, is likely to fail for most simulation models regarding the complexity of the dynamics inherent to the crowd if the animator tries to obtain a desired kinematic or dynamic pattern at a given time in the simulation process. This calls for automatic procedures that can reach those constraints while preserving the specificities of the crowd models. This notion of controlling a crowd is investigated in this paper. In practice, because of the discrete nature of pedestrians (and thus the highly discontinuous nature of the solution space), the use of simple gradient based methods yields some difficulties related to collision handling between individuals. This issue has already been encountered in the control of multibody dynamics in [33], where the authors suggest to select a solution from a set of samples (computed on-line on a clus-

☆ This paper has been recommended for acceptance by Siome Goldenstein and Peter Lindstrom.

* Corresponding author.
E-mail address: pierre.allain.t@gmail.com (P. Allain).

1524-0703/$ - see front matter
http://dx.doi.org/10.1016/j.gmod.2013.09.001

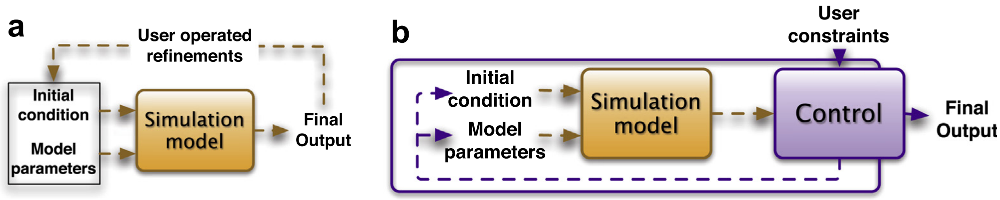

Fig. 1. Differences between a user operated control of a crowd simulation model (a) and the control procedure (b) proposed in this paper.

ter of machines). In this paper, we propose a combination of a gradient based method along with an heuristic to explore several distinct paths in the solution space.

1.1. Contributions

In the presented article, the control of a crowd is seen as an editing process, in the sense that the trajectories produced by the simulation model are deformed to achieve users constraints (Fig. 1b) while minimizing the discrepancy with the simulation model's dynamics. The types of constraints can be twofold: (i) per-individual constraints, meaning that the user can specify its own properties related to pedestrians (like positions, velocities or even shape-related information) or (ii) macroscopic constraints, such as respecting a given velocity field or higher order dynamical information, like the divergence or rotational components of a velocity field. Those two types of constraints are illustrated in the rest of the paper. Our optimization process uses recipes from optimal control of variational models by formulating the problem with the adjoint theory [16]. It is virtually adaptable to any kind of simulation model provided that it can be analytically described as a variational system. This is usually the case with crowd dynamics model, but not anymore if one considers cognitive modeling of pedestrian steering behaviors. Also, the quality of the produced animation strongly depends on both the realism of the controlled model and on the nature of the constraint imposed by the animator. In that sense our method augments the latent qualities of a given crowd model but is not meant to produce systematically more realistic simulations than advanced techniques. In other words, the resulting scene after control using the model and the users constraints is optimal, but not the model which basically stays the same.

1.2. Outline of the paper

First, a presentation of the related work is performed in Section 2, and an overview of the control process is given in Section 3. The optimal control of dynamical system, along with its requirements are then presented (Section 4). Algorithmic aspects of this control will notably be explained in this part. In Section 5, a variety of control are applied on a generic dynamic crowd model to illustrate the power of our method and quantitative convergence results are also presented, before a discussion and a conclusion end the paper (Section 7).

2. Related work

We first begin by giving the main approaches of crowd simulation, which we dissociate from the idea of crowd control. We also discuss other control works such as fluid ones.

2.1. Crowd simulation

Simulating crowd of individuals has drawn a lot of attention over the past decades for the potential interests of computer graphics, but also for safety engineering or robotics applications. The different models are commonly divided into two categories: microscopic and macroscopic. Microscopic approaches tend to model member of the crowds as agents with specific behaviors. Sophisticated behavior models seek autonomous agents endowed with goals and specific attributes [20,30], but for a somehow limited number of individuals. Oppositely, their motions can be the result of simple laws, such as in the seminal work of Reynolds on flocking [27]. Designing and tuning these laws is now known as the steering problem, for which several solutions exist thanks to different strategies; for examples interacting particles under psychosocial forces [8], reproducing experimental observations [23], principle of least effort [6] or vision based strategies [22]. The most recent methods allow to simulate large scale crowds at interactive framerates with convincing emergent behaviors of the groups. Another recent trend is to capture from the real world heterogeneous behaviors of pedestrians to add variety in the simulation and possibly realism. This is the case in [14,15], where authors capture individual trajectories of pedestrians and reuse them in an online fashion. Conversely, macroscopic models generally consider the crowd as a whole and model its dynamic by means of continuum mechanics equations, allowing analogies with the domain of computational fluid dynamics [10,32,26,21]. This type of modeling works well with dense crowds where the weight of individual decisions is somehow weakened, but fails to describe realistic interpersonal collision avoidance behaviors or heterogeneous crowds with individuals exhibiting distinct goals or motivations.

2.2. Simulation control

Controlling a crowd to achieve a given effect is a rather difficult task, mostly because the only control parameters are those of the simulation model, which are generally not designed for it. Ulicny and colleagues [34] are the first