

Comparing Constraint-Based Motion Editing Methods

Michael Gleicher¹

Department of Computer Sciences, University of Wisconsin, Madison, Wisconsin

Accepted June 1, 2001

Tools for assisting with editing human motion have become one of the most active research areas in the field of computer animation. Not surprisingly, the area has demonstrated some stunning successes in both research and practice. This paper explores the range of constraint-based techniques used to alter motions while preserving specific spatial features. We examine a variety of methods, defining a taxonomy of these methods that is categorized by the mechanism employed to enforce temporal constraints. We pay particular attention to a less explored category of techniques that we term *per-frame inverse kinematics plus filtering*, and we show how these methods may provide an easier to implement while retaining the benefits of other approaches. © 2001 Academic Press

Key Words: character animation; motion editing; constraints.

1. INTRODUCTION

The proliferation of different techniques for creating motion for computer animation, coupled with the needs to best utilize these artistic assets, has made motion editing a rapidly developing area of research and development. The past few years have witnessed an explosion of new techniques for various motion editing tasks, as well as the deployment of commercial tools. Unfortunately, while each new technique almost always provides impressive new results, there has been little done to examine the differences between and capabilities of various approaches. In this paper, we explore the relationships between a number of state-of-the-art motion editing methods. Specifically, we provide a unifying view of several constraint-based techniques. These methods provide explicit support for maintaining specific geometric features of motions. Our goal is to examine the range of techniques

¹ This research is supported by a grant from Microsoft Research, equipment donations from Intel, and IBM, and software donations from Pixar, Autodesk, Alias/Wavefront, and Metacreation. The author is supported by an NSF Career Award. This research is supported by an NSF Career Award “Motion Transformations for Computer Animation” CCR-9984506.



for insights—comparing and contrasting the methods to understand the problems for which each is appropriate, and to give us insights as to unexplored regions of the design space.

Our discussion is organized as follows. We begin by defining our terms and laying a groundwork to introduce motion editing methods. We then present a taxonomy of recent approaches, describing each within a common set of terminology to aid in comparison. Section 4 focuses on a particularly unexplored category of approach, showing how a simpler variant of an existing technique can provide useful results.

2. CONSTRAINT-BASED MOTION EDITING

Before comparing motion editing techniques, we must first define what it is we mean by motion editing. Motion editing is the act of changing the movement of an object [14]. Outside of computer animation, we generally do not get to control the movement of things without controlling the things themselves. Even in more traditional 2D animation, the movement and appearance of objects are tightly coupled. The idea of discussing motion editing only becomes a worthwhile topic when we deal with 3D animation.

This definition of motion editing is sufficiently broad that most techniques in computer animation fall under it. We focus our attention on tools specifically designed for this task, that is, tools specifically designed for changing existing motions. This broad characterization points to an important feature of these techniques: since we are changing existing motion, there must be a notion of preserving the original as well as adding new features to it. If we did not want to preserve any aspects of the original motion, we would not have needed it to begin with.

For the purpose of this paper, we narrow our focus to character body motion for figure animation. Typically, such animation represents characters as articulated figures: a hierarchical representation connecting a tree of rigid segments. While most of the methods we discuss in this paper are not limited to such tasks, all of them address the primary difficulty of using such a representation: the parameters of the character (joint angles and root position) are related to many of the most important geometric features (the position of end effectors, for example) by complex, nonlinear relationships.

Constraint-based motion editing makes some of the features that are to be preserved or changed in a motion explicit as constraints. This is in contrast to a non-constraint-based technique which does not make such features explicit, instead exposing the underlying representation of characters and motions. For example, a filtering or blending operation applied to a motion does change it (and is therefore an editing operation) and does preserve some aspects of the original, but does not explicitly describe the operation in terms of features in the motion, only in the terms of the underlying parameters used to represent the motion. A constraint-based technique would explicitly represent features in the motion. For example, the footsteps in a motion might be represented as an equation specifying the required position of the end-effector at a given time. Such features serve as constraints either to be preserved (maintain the footplants as the motion is changed) or to control changes (move a footplant to a new location).

Constraint-based motion editing techniques generally focus on spatial or geometric constraints. Spatial constraints provide features of poses at given instants. For example, they might specify that a hand must be in a given place at a given time, that the elbow does not bend backward at any time during the motion, or that the character's gaze direction faces an

event of interest when the event occurs. A distinguishing feature of constraint-based editing techniques is that they allow for spatial constraints independent of the characters' underlying parameters. For example, they permit position the end-effectors rather than just specifying joint angles. This is in contrast to things such as pure signal processing approaches, as introduced by Litwinowicz [19], Unuma *et al.* [31], Bruderlin and Williams [5], Popović and Witkin [34], and others, where the parameters are considered independently of their interrelationships.

In addition to the spatial properties of motions, editing techniques must consider temporal properties as well. There has been less formalism of temporal constraints on motion. While some authors (for example [13, 21, 31]), have used mathematically convenient terminology like frequency, these are typically crude approximation to the properties that we might really desire to control and maintain in motions. One of the most effective temporal constraints presented to date [13, 34] follows from the observation that high frequencies in motion are almost always noticeable and usually caused by significant external events. Therefore, the high-frequency content must be preserved during editing. That is, high frequencies should only be added or removed with care: it is just as incorrect to dull the “snap” of a karate kick or a chef’s chop as it is to add a twitch to a graceful ballet movement.

2.1. *Types of Editing Operations*

To date, constraint-based motion editing operations have been primarily applied to motion transformation tasks. These are problems where good motions are adapted to different settings, such as new characters or new environments.

Constraint-based motion editing methods can be applied beyond motion transformations. For example, a signal processing operation such as filtering may destroy a motion’s spatial constraints. A constraint-based method could then be used to reestablish these constraints. Similarly, interactive manipulation of a frame of the motion might destroy the temporal constraints, which could be reestablished with a constraint-based solver.

The editing paradigm where some operation is used to alter a motion and then a constraint-based solver is used to “clean up the mess” is only one of the possible approaches to constraint-based motion editing. A different paradigm begins with a motion that satisfies the constraints and then alters the constraints while maintaining their validity. An example of the former approach is our work on retargeting [13] and path editing [15], while examples of the latter include our interactive editing work [12], Lee and Shin’s hierarchical editor [20], and Popović’s physically based approach [27].

2.2. *Inverse Kinematics*

The methods used to find the solutions to the geometric constraints are a key implementation choice in constraint-based motion editing systems. The handling of geometric constraints on articulated figures is typically called “inverse kinematics.”

Because the term inverse kinematics has different interpretations, we define ours. Inverse kinematics (IK) is a process for determining the configuration of a character’s parameters (known as its pose) based on specifications of resulting features of the pose, such as end-effector positions. Note that this is a generalization of another common interpretation of the term (popular in robotics texts) that defines IK as the process of determining joint angles from an end-effector position. For our definition, we neither restrict the types of parameters

(i.e., all of a character's parameters, not just its joint angles are considered), nor the number or type of constraints (i.e., not just end-effector positioning). Inverse kinematics implies solving the geometric constraints at a specific instant in time.

All constraint-based motion editing methods must include a solver for IK problems in one form or another to allow for constraints on end-effector positions. The nature of inverse kinematics leads to three general types of challenges that any solver must address:

1. The relationships between the end-effector being controlled and the parameters being determined will be non-linear, often quite complex. This means that solvers often must rely on sophisticated methods.
2. There exist goals for which there are no solutions. Such singular problems are particularly difficult because there are often goals close to these impossible situations that are barely achievable. A solver must distinguish the impossible from the merely difficult. Maciejewski [23] showed that this distinction is inherent in the mechanics of the IK problem.
3. There is almost always a multitude of possible solutions: most end-effector goals may be reached by a variety of poses. The solver must have a mechanism for choosing among the space of possible solutions.

2.2.1. IK solvers. Any IK solver must address these three issues. The first two are critical to creating a robust and efficient method, essential to any application. The third provides some freedom in solution methods. In choosing among correct solutions, solvers can differentiate from one another by choosing "better" answers in terms of the quality of the resulting pose, or the predictability of the results.

The way that an IK solver chooses among possible solutions to goals is an important issue in the solver's design. Some solver implementations make this choice explicit by casting the IK problem as a constrained optimization problem. For example, a solver might prefer the solution closest to its initial starting point, the one closest to the previous frame, or one that provides the "best" pose subject to posture, strength, or balance objectives. However, the choice is often implicit in the design of the solver since many solvers take a more ad hoc strategy. Effectively, many solvers provide "the answer that they happen to choose." Another way to look at these algorithms is that they are optimizing something else: simplicity of code, or the speed of finding a solution.

While IK methods were developed in mechanics and robotics before their application in computer graphics, the demands of graphics applications (particularly in character animation) brought new challenges for the methods to address. The methods for IK solvers applied in graphics fall into two categories: geometric or analytic solutions, and numeric or iterative solutions.

Analytic solution methods use closed form geometric constructions to compute configurations for end-effector positions directly. Such methods must be carefully constructed to handle specific cases. Tolani *et al.* [30] recently provided a method for a 7 degree of freedom limb as typically used in computer animation. This work extends earlier solvers, such as those by Korein and Badler [17]. Analytic methods have the advantage that they provide guaranteed fast solutions. However, they lack flexibility in how they choose solutions in underconstrained cases and in the types of problems they can handle.

Numeric solution methods use equation solving or optimization techniques to provide a more general, albeit computationally expensive, IK solution. Because of the nonlinear nature of the inverse kinematics equations, a numerical solver will necessarily be an iterative method. The methods employed fall into two categories: unconstrained and constrained.

Unconstrained methods model IK goals with objective functions that are minimized by an unconstrained optimization routine. Such methods, typified by Badler *et al.* [4], Welman [32], and Phillips and Badler [25, 28], can employ simpler solvers and may trade accuracy of end-effector positions for other goals, like balance. Constrained methods treat end-effector goals and joint limits as hard constraints, and reserve the optimization objective to select among solutions. Realizations of this strategy, such as Zhao and Badler [35], require the use of sophisticated nonlinear programming algorithms.

Lee and Shin [20] present a solver that combines analytic and numerical techniques. They employ analytic solutions for the limbs, where closed form solutions are available, and reserve optimization for the computation of body posture. Shin *et al.* [29] also present a hybrid solver.

3. A TAXONOMY OF SOLUTION TECHNIQUES

The solvers used in constraint-based motion editing all address similar tasks. The main differences are the approaches they employ to achieve their results. In this section, we look at the various approaches used to maintain or establish the spatial and temporal constraints in editing techniques. One might hope that what occurs “under the hood” of an editing technique would be of little concern to the user. Our goal in this exploration is to see how the choices in solver strategy affect what the methods that use them can achieve.

By our definition in the previous section, a constraint-based motion editing technique must provide some method for handling spatial constraints, that is mapping specified goals to a character’s parameters. Therefore, all methods must include an inverse kinematics solver (in a general sense).

Our taxonomy categorizes methods based on their approach to handling temporal constraints. The existing methods vary greatly, from approaches that do not explicitly handle temporal constraints to those that provide explicit equational constraints for temporal properties.

3.1. *Per-Key Methods*

Traditional computer animation systems provide the user with the ability to control a set of “key” poses that are interpolated. Spatial aspects of motion are created by controlling these key poses, and temporal continuity between poses is generated by interpolation. We term such methods *per-key inverse kinematics* (PKIK) because they solve a problem to handle the spatial constraints on each key in the motion. Traditional animation tools fall under the category of per-key methods.

An important aspect of a per-key motion editing technique is that the “solver” changes each pose individually. The choice of how to set each pose is computed independently. These computations may consider other frames in the process, however each pose is computed separately. This attribute is shared with the per-frame methods described below.

Per key approaches do not provide any mechanism for enforcing constraints at frames other than the key frames, and therefore provide no guarantees that such constraints are met. The reason that such methods often work well is that the small number of key frames is often sufficient to describe the motion. If significant key frames are properly chosen, interpolating between them can provide sufficient temporal continuity and control.

A central difficulty with per-key methods for motion editing is that they require a set of sparse, well-chosen key frames to represent the motion. The temporal control of motion

changes is an artifact of how the motion was created and represented, not necessarily determined by the goals of the changes in the motion. This is particularly problematic for motions that are generated algorithmically, for example with motion capture or simulation. However, even manually constructed motion does not place key frames solely at semantically relevant instants. One possible approach to this is to use some process for constructing a convenient representation from the unstructured data.

Recent systems have attempted to make automated use of a per-key approach. The system of Bindingavale [3] identifies local extrema of motions as “key” poses, applies inverse kinematics to these poses, and interpolates between the poses. This creates a system that can preserve events automatically.

3.2. *Motion Warping or Displacement Mapping*

One drawback of the per-key approach is that key frames may not be placed at spacings convenient for editing operations. Motion warping [34], also known as motion displacement mapping [5], provides a solution to this problem. In such an approach, the *changes* to a motion are keyframed. This allows keys to be placed conveniently for editing. We term a motion editing approach based on motion warps as *motion warping plus inverse kinematics* (MW+IK) because it must augment motion warps with an inverse kinematics solver to handle the spatial constraints.

When combined with an IK solver, motion warping provides a type of constraint-based motion editing method. The IK solver is applied to frames with specified geometric constraints, and the spacing of the keys of the displacement map can be used to enforce temporal constraints on the changes applied to the motion.

The MW+IK approach does have some drawbacks. First, there is no control over geometric constraints except at keyframes. Second, the amount of temporal constraint possible depends on the amount of geometric control. To provide more geometric control, more keys must be provided, leading to a reduction in the amount of temporal constraint.

PKIK can be viewed as a special case of MW+IK where the keys in the displacement map are placed at the same times as the keys in the original motion.

A system that explicitly implements the MW+IK approach is described by Choi and Ko [18]. They apply inverse kinematics to selected frames of an animation in order to ensure that key constraints are met on those frames. The changes made to those frames are interpolated between key frames. The original motion warping work of Popović and Witkin [34] also exemplifies the MW+IK approach as it allows the user to use any tools in a commercial animation system (including IK) to set the key frames.

3.3. *Per-Frame Methods*

A per-frame method manipulates each individual sample of a motion independently. We prefer the term per-frame rather than per-sample because it sounds better, although the latter is probably more technically correct, especially if we are dealing with motions that are algorithmically generated or captured at a high sampling rate. We refer to the class of methods that applies a solver to each frame of the motion independently as *per-frame inverse kinematics* (PFIK).

In a sense, a per-frame method is a special case of a per-key method where the keys are regularly spaced. However, this is an important distinction: with a per-key method,

we assume that the keys are strategically placed at significant and important times and, therefore, it is more likely that simply getting these instants correct will provide for a desirable outcome.

With per-frame methods, the poses that we consider are typically densely sampled. That is, the timing between poses is typically quite small relative to the duration of the motion. Because of this, there is a greater demand to enforce the consistency between poses. When the poses are spaced farther the interpolation between poses becomes more significant and will create the necessary consistency.

To emphasize the difference between the per-frame and per-key approach, consider a motion that is the trajectory of a foot making a step. If such a motion was created manually, the artist might create key poses at the beginning, middle and end of the step and interpolate these sparse keys. This same motion could also be sampled a full frame rate. If we were to change the height of the apex of the step, a quite different result would be achieved. With the sparse keys, we would obtain a higher step, while with the dense samples, we would simply put a spike in the motion at the apex.

The contrived nature of the simple example might make per-key methods sound more likely to give desirable results. However:

1. Per-key methods offer only the *opportunity* for choosing keys that are semantically relevant and, therefore, convenient controls over the animation.
2. Per-key methods generally decouple the editing operations from the frame rate. The timing of the keys usually corresponds to events in the motion, while the timing of the frames is usually an artifact of how the motion will be displayed.
3. Per-key methods only change the keys, therefore provide no guarantees on enforcing constraints at nonkey times.

The characteristic of a per frame method is that the decision of how to adjust each frame is made independently. The challenge in per-frame approaches is to enforce the temporal constraints while modifying each frame individually. Each of these independent decisions may take into account more information than just the pose at the current frame. For example, a solver might refer to the previous frame in order to choose a solution for the current one that does not introduce a discontinuity.

3.3.1. Online motion editing. A special class of motion editing problems arises in situations where the adaptations of animations must occur on-line. That is, the system must process animation as it is being generated and viewed. The most common application of online motion editing is performance animation or puppetry, where a performer drives an animated character using motion capture and this is shown to the user immediately. Another application of such online systems is computer games or virtual environments where animated characters' motions must be generated in reaction to events as they occur.

Recently, several systems have appeared that apply constraint-based approaches to on-line problems. Such systems must necessarily be per-frame systems as they must compute the current frame, without the ability to know (or affect) the future, or change the past.

An example of this approach is the On-line Motion Retargeting work of Choi and Ko [6]. For each frame, the inverse kinematics constraints are treated through linearization. In underconstrained cases, the system attempts to match the original joint angles as closely as possible subject to the constraints on the end-effector positions. Coherence between frames is maintained as the end-effectors of the character are constrained to follow the paths of the

performer (whose motions must be continuous), and consistency of solutions is obtained through the objective function that attempts to match the angles of the real performer.

Shin *et al.* [29] present a more comprehensive approach to dealing with the problem of applying motion data to characters of varying size in an on-line application. A Kalman filter technique is applied to reduce noise and enforce temporal constraints, a special high-performance IK solver is used to ensure end-effector constraints are met, and an importance metric is used to gracefully switch between tracking the performers end-effectors and joint angles. Unlike the system of Choi and Ko, this only tracks end-effector positions when they are relevant, preferring to match postures at other times. By using graceful switching between the two modes, Shin’s system avoids the discontinuities that may arise from switching.

3.4. *Per-Frame Plus Filtering*

Once the decisions of how to alter each frame are made, some approaches then provide a “global” process that considers multiple instants together to remove artifacts of the independent nature of the first step. Almost always,² this step is a low-pass filter meant to remove the spikes and other discontinuities introduced by the independence of the first steps. The first published motion editing system based on such an approach was the Hierarchical Motion Editing technique of Lee and Shin [20] that used B-Spline fitting to implement the low-pass filtering. We term the general approach *per-frame inverse kinematics plus filtering* (PFIK+F) and discuss the approach in detail in Section 4.

The PFIK+F approaches are distinguished by the existence of two separate phases. Generally, there is a phase that considers spatial constraints, applying inverse kinematics to make coordinated changes to all joints of a character simultaneously, followed by a signal processing stage that implements the temporal constraints. The spatial phase considers the entire character simultaneously, albeit at a single instant in time, while the temporal phase considers the entire motion (or a portion thereof) simultaneously, albeit a parameter at a time. Because these two phases cannot consider the work of the others, one may “undo” the work done by the other, so they are often interleaved in an iterative process.

Treating each signal independently is technically incorrect. At a local level, groups of parameters are almost always coupled to represent rotations. For example, independently filtering each parameter of an Euler angle does not necessarily have meaning for the rotations themselves. In practice, we believe that this approach works as well as it does on the displacements because the displacements are typically small, and in small-angle approximation, the orientation representations are linear. Recently, Lee and Shin [21] have shown that FIR filters can be adapted to operate on orientations through the use of exponential coordinates. However, this result does not apply to the common approaches for doing PFIK+F.

3.5. *Spacetime Methods*

Spacetime constraint-based motion editing techniques distinguish themselves from the previously described methods in that they do not consider frames individually. Spacetime Constraints refer to methods that consider a duration of motion simultaneously in a computation. Rather than computing an individual frame, as an IK solver does, the solver computes

² All of the examples to date.

an entire motion, or any subwindow of it. This allows it to consider constraints on the entire duration of the motion, and to have objective criteria that consider entire motions.

The initial use of spacetime constraints specified desired positions for a character and used the solver to compute the “best” motion that met these positions. These initial works, presented by Witkin and Kass [33] and Cohen [7], included constraints that enforced the laws of physics and created an objective function that defined the “best” solution as one that minimized the amount of energy the character expends with its muscles. This synthesized novel motions that had a simple character perform simple motions that were physically correct.

The power of the spacetime constraints approach is also its drawback. While the approach provides tremendous opportunity to define constraints and objective functions that describe features of the resulting motions, these must be defined for the approach to work. Defining such mathematical characterizations for motion properties is challenging. While the approach offers the potential for high level properties to be employed as criteria, to date, concepts such as “graceful” or “angry” or “like-Fred-Astaire” have eluded a mathematical description that fits into the framework. Also, the approach requires solving a single mathematical problem for the entire motion. This leads to very large constrained optimization problems that are usually very difficult to solve.

We initially proposed applying the spacetime constraints approach to a motion transformation problem in [10]. Conceptually, the main difference with the standard spacetime work was that our objective sought resulting motions similar to the initial motions, rather than seeking results that minimized energy consumption. This allowed us to avoid the difficult problem of specifying motion details: we did not have to figure out how to describe a walk with constraints since we could define a walk by example.

The important properties to preserve in a given motion may not always be simple: “realism,” “grace,” “like-in-Singing-in-the-Rain,” or other high-level properties may be desirable to preserve during adaptation. In practice, we are limited by our ability to define high-level qualities of the motion mathematically, by our ability to compute adaptations efficiently when the metrics become complex, and by the amount of effort we wish to expend in identifying (or having the user identify) these properties. Even if we encoded the desired animation completely in a constrained optimization, we would still need to solve these problems. Generally, richer sets of constraints and objective functions lead to more difficult problems to solve.

To date, the spacetime approach has been demonstrated on a variety of motion transformation tasks. We have demonstrated that the approach can provide real-time interactive performance on modest computers in real time [12].

3.6. *Physical Approaches*

Physics provides a specific, and useful, source of constraints. While certain physical properties, such as contact with the floor, are conveniently represented as geometric constraints in the spacetime systems of the previous section, other physical properties are ignored for the sake of performance. Most significantly, Newton’s laws are not included. The choice to exclude these constraints was a specific concession of the work of [12]. Kinetic constraints and energy properties are particularly computationally expensive in the spacetime framework.³

³ Intuitively, such constraints are challenging because they create couplings among frames.

For some motions, the kinetic behavior is critical. In such situations, it is unacceptable to make the simplification of discarding Newton’s laws as a constraint. Popović and Witkin [27] make a different set of trade-offs to make the spacetime approach realizable. Rather than simplifying the physics by ignoring Newton’s laws, they chose to simplify the geometry of the character instead.

To date, the physical approach of Popović has not been extensively demonstrated. It has only been tested on a small number of short motions. To create a walking motion, the method is applied to a single gait cycle and looped. In addition to its complexity, the technique does have the disadvantage of being newer, and therefore less evolved, than other approaches. Several researchers are exploring ways to approximate momentum effects within a spacetime editing framework.

4. PER-FRAME IK PLUS FILTERING

In their 1999 paper, Lee and Shin [20] introduced an interactive motion editing system that addresses many common tasks, including retargeting, transitions, and editing. In addition to the specific details of their system, the work presents a novel approach to dealing with both spatial and temporal constraints. The paper presents the first example of a PFIK+F (per-frame inverse kinematics plus filtering) motion editing approach. While the paper does not describe other ways of realizing the PFIK+F approach, it did inspire us to consider the space of implementations of this style of approach.

A PFIK+F solver uses an inverse kinematics solver applied to each frame of the motion in order to handle spatial constraints. The results of these solutions are then processed to enforce the temporal constraints. In the inevitable event that the second phase destroys the work of the first, the process is repeated.

Lee and Shin apply the temporal constraint suggested by Gleicher [13]: that the *changes* to the motion should be band-limited. Alternatively, one might view their process as saying the spatial constraints should be met using as low a frequency change to the motion as possible. The advantage of this latter view is that it suggests that we are willing to do whatever is required to meet the spatial constraints, even if it means adding high frequencies.

The original paper by Lee and Shin describes a specific implementation of the PFIK+F approach. Their system makes specific choices for each component of the approach. With each step, they provided an innovation:

1. They use exponential maps as a representation for rotations and provide methodology for using this representation with displacement maps.
2. They provide a novel, simplified IK solver.
3. They use B-spline fitting to the displacement maps to perform the filtering.

The biggest innovation of the paper, is that it introduces the PFIK+F approach. The paper provides the first demonstration of a PFIK+F approach and illustrates its applicability to a wide range of motion editing problems.

Each of their innovations represents a tradeoff:

1. Exponential maps provide a methodol for implementing displacement mapping on Quaternions and give some theoretical meaning to filtering operations (as presented in [21]). However, this representation has other implementation issues (see [16]), does not necessarily provide “nice” motions when splined, and is difficult to integrate with existing tools.

2. Their simplified IK solver has the advantage that it is extremely fast. This is important since each application of the PFIK+F process may require hundreds of individual calls to the IK solver. The tradeoff of this fast solver is that it supports only a limited set of constraints and sometimes seems to provide unnatural looking motions.

3. B-Splines implement low-pass filters with very wide support and can be easily adapted to weight different samples differently. However, they are more difficult to implement (and potentially computationally expensive) than convolution-based (FIR) linear filters.

Clearly, the PFIK+F approach can yield a spectrum of implementations, with different choices made for each aspect of the system.

4.1. *Our PFIK+F Solver*

In order to better understand the design space of PFIK+F solutions, we have created a very different implementation of the approach than Lee and Shin have presented. We emphasize that we have not attempted to improve Lee and Shin’s results. While their system is specifically designed to be an effective implementation of a PFIK+F approach, ours uses existing components within our existing system. We were more concerned with implementation convenience than performance or result quality. However, the fact that we were able to realize the approach inside our existing system using already available components speaks to an advantage of the approach. Having the two approaches within the same system will allow for easier comparisons in Section 5.

The PFIK+F solver that we have built is based on the following design decisions:

1. We treat whatever underlying rotational representation used by the system as independent coordinates. In practice, this means that we do per-component arithmetic operations on Euler Angles.

Applying the PFIK+F solver directly to Euler Angles does appear to work in practice, despite its theoretical incorrectness. Rather than trying to make an argument to justify this theoretically, we prefer to acknowledge that such a “signal processing” approach, performing per-component arithmetic on Euler Angles, to motion editing is widely used. In practice, we believe that this approach works on the displacements because the displacements are typically small, and in small-angle approximation, the orientation representations are linear. The approach fails most spectacularly when the Euler Angles wrap around at 2π , as the simple arithmetic cannot notice that two very different parameter values represent the same angle.

Applying the PFIK+F solver directly to Euler Angles allows easy integration with existing systems. We feel this is an important asset of the PFIK+F approach, even if it comes at the expense of mathematical cleanliness.

2. We use our existing IK solver to solve the IK subproblems. Or, more specifically, we use our general purpose spacetime solver (as described in [13]) on a single frame at a time. While this solver provides a wide range of constraints, this generality comes at the expense of speed and precision. Our spacetime solver is most likely not as good at solving the per-frame IK problems as it is designed with excess generality, and intended to provide performance on the large sparse problems that occur in the spacetime, rather than the small, dense problems that arise in IK.

Again, the key advantage of this choice is to demonstrate integration with existing systems. Implementing a good IK solver is a challenging process, and such a solver is generally

considered a valuable asset in an animation environment. Arguably, implementing an IK solver is the key challenge of creating a PFIK+F system. However, since an existing solver can be leveraged, this difficulty can be avoided.

3. We use linear FIR filters, rather than B-Spline fitting. This has advantages in generality and ease of implementation. The hierarchical B-Spline approach used by Lee and Shin, first presented by Lee *et al.* [22], offers many advantages over FIR filters for the PFIK+F approach that will be mentioned in the subsequent discussion.

4.1.1. Operation of the algorithm. Each iteration of the solver produces a new motion (\mathbf{m}_{i+1}) from the results of the previous iteration. We denote the initial motion as \mathbf{m}_0 .

1. The inverse kinematics solver is applied to each frame of the animation. To compute $\mathbf{m}_{i+1}(t)$, we apply the solver beginning with configuration $\mathbf{m}_i(t)$ and the constraints at time t . The solver is run on each frame of the animation independently.

2. The displacement map is computed as the difference of the new motion and the original motion, $\mathbf{d} = \mathbf{m}_{i+1} - \mathbf{m}_i$.

3. The displacement map is filtered by convolving it with a low-pass filter kernel. Each channel of the vector is treated independently. We denote this filtered result as \mathbf{d}_f .

4. The new motion is computed as the sum of the filtered displacement map and the motion at the beginning of the iteration, $\mathbf{m}_{i+1} = \mathbf{m}_i + \mathbf{d}_f$.

We note that (assuming the inverse kinematics procedures work), the “original” \mathbf{m}_{i+1} in step 1, and therefore the displacement computed in step 2 ($\mathbf{m}_i + \mathbf{d}$), motions will satisfy the constraints on each frame. The filtering process may provide a motion that does not satisfy the constraints. In fact, there is no guarantee that the result of step 4 (\mathbf{m}_{i+1}) will be any closer to satisfying the IK constraints than \mathbf{m}_i in step 1. We cannot make any theoretical claims about the convergence of the algorithm without placing restrictions on the filter.

The intuition for the algorithm is that at each iteration we determine what changes to the motion are needed to satisfy the geometric constraints for each frame (step 1 and 2). We then remove any part of the changes that would violate the temporal constraints. We include only the portion of the changes that are acceptable. These changes are not enough to fully satisfy the spatial constraints since they are only part of the changes determined necessary.

To create an algorithm that prefers low frequency changes, but does whatever is necessary to satisfy the geometric constraints, we use different filters on each iteration of the algorithm. This is accomplished in Lee and Shin’s system through the use of B-Splines with increasingly many control points. In our approach, we use filters with increasingly high cutoff frequencies (i.e., smaller and smaller filter kernels). Intuitively, this causes as much of the change to be accomplished by the low frequency changes.

We believe that this approach could be implemented in the scripting languages of most major animation packages. To date, we have not actually created such an implementation, however, we have done paper design experiments with 3D Studio Max [8] and Maya [1].

4.1.2. A trivial demonstration. To provide some insight as to how the PFIK+F algorithm works, we examine a trivial example: a 2D particle whose initial motion is along the x axis. We constrain the particle to be on the x axis at the beginning and end of the motion, and to have a height of n units halfway through the motion. To simplify the illustrations, we do not allow the solver to alter the x positions of the particle (so that the x position of the particle is equal to the time step). The initial problem is illustrated in the upper left of Fig. 1.




FIG. 1. (upper left) Initial motion and constraints for the trivial example. The other frames show the third iteration of the solver. The upper right panel shows the displacement required to “solve” the constraint, the lower left panel shows the filtered displacement, and the lower right panel shows the filtered displacement added to the previous step.

Walking through the steps of the process, we begin by “solving” the constraints on each frame. For this trivial example, solving merely means resetting values at the constrained frames. Subtracting the input to this iteration from this motion provides a displacement map, which is filtered to compute \mathbf{d}_f , which is added to the starting point of the iteration to provide the result of the iteration. This process is illustrated in Fig. 1.

The convergence of this process is shown in Fig. 2. For this example, each iteration uses an identical filter: in this case a binomial filter of width 9 [the kernel (1, 8, 28, 56, 70, 56, 28, 8, 1)]. In this case the solver will never converge exactly—each iteration only moves the center point 30% closer to the goal.

Figure 3 shows a different example illustrating how the addition of frequency bounded changes preserves the original motion. In this case, the initial signal has high frequencies and is subject to the same constraints as the previous example. Figure 3 mimics the illustrations provided for the simpler example, and the convergence of the solver is shown in Fig. 4.




FIG. 2. The path of the particle on each iteration of the solver.

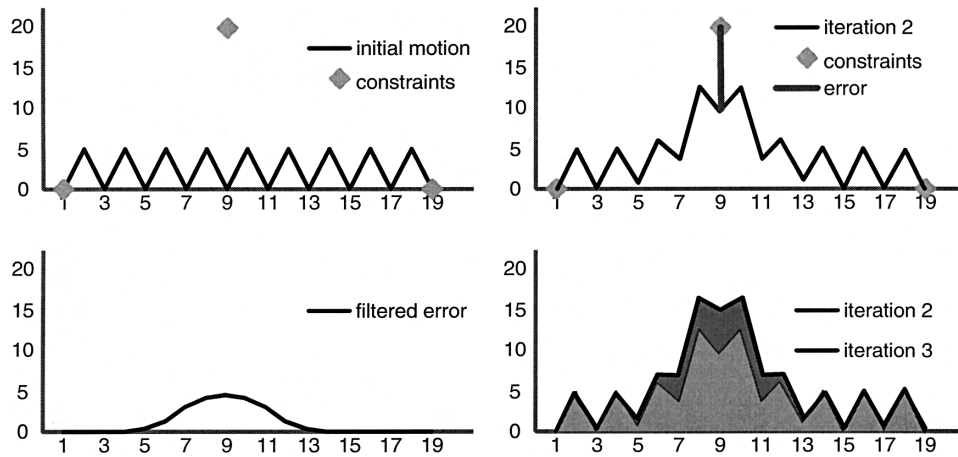


FIG. 3. The steps of an iteration of the solver on the wave example. The upper left panel shows the initial condition. The third iteration of the solver is shown in the other frames. (upper right) The results of the per-frame solve provide the displacements. (lower left) The displacements are filtered to provide filtered displacements. (right) The final result sums the filtered displacement and the initial curve.

To make the example more challenging, we add additional constraints in Fig. 5. Again, the process is illustrated with the results in Fig. 6. Even this trivial example has the property that it will never converge: it is not possible to meet the spatial constraints without adding high-frequencies to the initial motion. To better handle such cases, we change the solver to use a decreasingly large filter kernel on each iteration. The result of this process is shown in Fig. 7.

4.1.3. Filters and fitting. Within the context of the trivial example, we can better understand choices for filters. It is obvious from the first example that the filtering scheme chosen leads to a very slow convergence of the algorithm in this case: each iteration only progresses 30% of the way toward the goal.

A different way to implement the low-pass filtering behavior would be to fit a low-order polynomial through the constraints. Given that there are only three constraints in the most trivial case of Fig. 1, this would provide a smooth curve that would exactly meet the




FIG. 4. The path of the particle on each iteration of the solver.

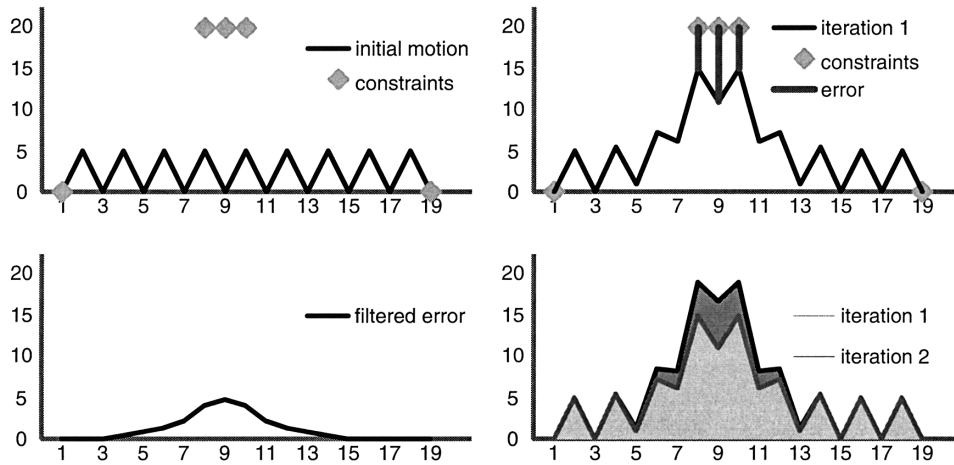


FIG. 5. Initial motion and constraints for the third simple example with 3 constraints. The steps of the second iteration of the solver are shown: (upper left) initial problem; (upper right) the results of the per-frame solve provide the displacements; (lower left) the displacements are filtered to provide filtered displacements; (lower right) the final results sums the filtered displacement and the initial curve.




FIG. 6. The path of the particle on each iteration of the solver.




FIG. 7. The path of the particle on each iteration of a solver where a filter of decreasing size is used on each iteration.

constraints. More generally, we could imagine choosing a piecewise polynomial such that desired smoothness constraints are met, and fitting this curve to the data. If there are many constraints that are close together in time, the curve fit through them will either need to have many degrees of freedom (and therefore have high-frequency wiggles) or it will not interpolate the data.

In Lee and Shin’s PFIK+F implementation, a scattered data interpolation method described by Lee *et al.* [22] is employed. This scheme first fits a polynomial with a small number of knots, and then fits polynomials with an increasing number of knots to fit the residual. This has the effect of fitting the data with as low-frequency as possible interpolant. A similar scheme can be realized by applying a sequence of FIR filters between each iteration of the algorithm, with latter iterations using filters with higher frequency cutoffs.

4.2. Experiments

We have implemented the PFIK+F approach within our Timelines animation testbed. Our spacetime constraints solver, described in [13], is applied to each frame individually to serve as the IK solution. By implementing this approach in our existing framework we are able to apply it to the same problems that we have been exploring with other approaches.

We have applied our PFIK+F solver to a range of problems, including retargeting motions to new characters, constraint editing, and path editing. In the first and last type of problem, a change is made across the entire motion and a new solution is found to existing constraints. In a constraint editing problem, the constraints are manipulated interactively by the user.

When our spacetime solver is applied in batch mode, its performance is slower than when used interactively. This is because we make additional performance concessions (as described in [12]) for interactivity. More significantly, in batch mode, we lose “differentialness” [11]. That is, during dragging there are only small changes made between each solution. Therefore, the previous solution is a good starting point in the search for the next. Interestingly, our PFIK+F solver does not gain major advantage from differentialness. While a good starting point might make each individual IK solution slightly faster,⁴ the total number of iterations is the same.

We discuss a few representative experiments chosen to highlight some of the differences in the results of the two solvers. Timings reported in this section are run on a Dell Precision 610 workstation with a 550 Mhz Pentium III Xeon Processor and 512 M of RAM.

4.2.1. A simple retargeting problem. We first apply the solvers to the task of retargeting the walk of a fashion model down a catwalk to a shorter character, 60% of the original’s height. Footplant constraints are applied to the characters heels and the balls of her feet. The motion has 183 frames on which 954 scalar constraints have been placed⁵ to the heel and forefoot strikes. The solver is permitted to move the positions of the footplants, however it must keep the feet on the floor and in place when planted. This example is similar to the one used in Section 5.1.2.

⁴ We believe the solver described by Lee and Shin takes a constant amount of time, independent of how far the initial guess is.

⁵ We count scalar equations, so a constraint that positions a point in 3D counts as 3 constraints.

We should note that because the character was scaled uniformly, there is a trivial solution to this retargeting problem: scaling the translation by the same 60% as the limbs, as shown in Fig. 8. This solution is found by the heuristic starting point finder described in [13], therefore we omit the heuristic on this example. Instead, we force the character to walk using the footplant positions of the original, causing a more challenging retargeting problem.

Our spacetime solver (the least-squares method described in [13]) produces a good result for the motion when the control points of the B-Spline for the displacement map are placed 5 frames away from one another (3 frame spacing shows noticeable “lunging” to meet footplants). The solver only moves the footplants a small amount from their original positions, so the small character does take very long steps in relation to its height (this is normally addressed by the scaling heuristic). Otherwise, the solver produces an appealing motion. In under half a second, the solver provides a solution where the total error magnitude over the entire motion is less than 2 cm.

Five iterations of the PFIK+F solver take slightly over 2.5 s to provide a solution where the total error magnitude is approximately 6 cm. (no single frame shows more than a few millimeters of error). The resulting motion exhibits a little more “lunging” to meet the long stride lengths than the spacetime solution, which is a sign that we should have applied more filtering to the first few iterations of the solver (our implementation presently only supports limited filtering).

Figure 9 shows the amount of error per frame after each iteration of the solver.

4.3. Real Retargeting Problems

We have also applied both the PFIK+F and spacetime solvers to a number of “real” retargeting problems where we adapt existing motions to new characters. In such examples, we find motion and character geometry in libraries, determine a skeleton to fit inside of the geometry, retarget the motion to work with this skeleton, and then use the skeleton to drive the movement of the geometry. Some of the menagerie of characters that were tested are shown in Fig. 10. In all cases, the characters were found in libraries, as were most of the test motions.

The spacetime-based solver reliably provided subjectively acceptable results over a very wide range of motions and characters. The PFIK+F solver was not quite as reliable. On several examples, the solver failed to converge to a solution that met both temporal and geometric constraints. Also, conflicting constraints and anomalies with euler angles (filtering over zero-crossings) often caused the PFIK+F solver to fail to return a visually acceptable solution, even in cases where the spacetime solver was able to provide answers.

Figures 11a and 11b shows the results of performing a retargeting test using both our PFIK+F and Spacetime solvers, respectively. We began with motion capture data of a person performing a spinning leap (not shown), and retargeted this motion to a cartoon ant. The image shows three different frames of the motion from the results of each solver. The two solvers give very different results. The PFIK+F solver has the ant leaping much higher than the spacetime solver. It is impossible to say which is “correct”: we do not know how a cartoon ant would perform the jump.

In terms of temporal performance, the differences are difficult to characterize. There were examples for each solver where it was the fastest. This is not surprising given the

iterative nature of the algorithms (especially the spacetime solver), and will be explored in Section 5.1.

4.4. *Path Editing*

In this section, we consider the problem of transforming the path that a locomoting character moves along. For example, we transform a motion of walking in a straight line to walking in a curved line. The details of the method, described in [15], are unimportant for this discussion: as far as the solver is concerned, the motion is simply transformed in a manner that does not preserve the constraints so they must be reestablished. In practice, the path editing and solution take small amounts of time: we are able to interactively change the paths of motions with several hundred constraints.

One issue in using the PFIK solver for this approach is that a single position for each constraint must be found. This does not trivially extend to constraints that exist over a duration of time. The position of such a constraint may be mapped to different locations at different times. One way to view the problem is that a single instant must be chosen for the constraint to determine how to transform its position, otherwise its position may not be mapped to a unique location.

Different choices in time lead to different constraint positions which, in turn, lead to different characteristics of the motions. The three most obvious choices are the beginning, middle, and end of the duration of the constraints. Figure 12 shows how this choice affects a walking motion. The differences are subtle, and illustrate the challenge of motion transformation: there is no clearly “correct” answer given the limited information we have. We are left with either building more sophisticated models to determine the movement, relying on heuristics (or user input) to select among the myriad of small choices, or reconciling ourselves to never having the perfect answer.

One specific caveat: entire footplants must be dealt with at a single instant, the heel and toe strikes cannot be given separate times. If they are transformed at the same instant, their positions will be transformed rigidly. Otherwise, the heel and toe strike positions may be transformed differently, changing the distance between them which is impossible if the character is a rigid skeleton. The spacetime approach provides the option of allowing the solver to determine where the footplants are placed.

5. COMPARING APPROACHES

At a high level, the different methods discussed so far all achieve the same result: they allow for finding motions that meet a set of spatial (geometric) and temporal constraints. Our goal in this section is to understand the differences between the methods in terms of what they accomplish.

Our focus is on comparing the PFIK+F and Spacetime approaches as they address similar problems:

- Both the spacetime and PFIK+F approaches are applicable to motion from any source: keyframed data can be sampled. The PKIK methods only work given good keyframes, and therefore are not applicable across the range of problems of these other methods.
- Both spacetime and PFIK+F methods are inherently off-line. That is, they require the entire motion to examine. While this has the benefit that the methods can look at larger ranges of the motion to create better temporal constraints, the ability to look into the future may

4-Color Artwork




FIG. 8. The model's walk is retargeted to a smaller character (right).

not be practical. The PFIK approach seems to be uniquely suited for on-line applications where future durations of the motion are unavailable, and past decisions cannot be changed as discussed in Section 3.3.1. Recent developments in online systems [29] may lead to PFIK that are competitive for off-line problems.

- The PFIK approach can be viewed as a variant of the PFIK+F approach with an identity filter. In practice, a more sophisticated IK technique would be used to enforce temporal constraints. However, it appears that the class of temporal constraints that can be achieved this way is different than what can be achieved with the other methods.

4-Color Artwork




FIG. 9. Total per-frame error for the catwalk example for each iteration of the PFIK+F solver. Error is measured in feet.




FIG. 10. A variety of characters that we have retargeted motions to. The original skeleton from the motion capture data is shown in the center.




FIG. 11. (a) Results of retargeting a leaping dance motion to a cartoon ant using our PFIK+F solver. (b) Results of retargeting a leaping dance motion to a cartoon ant using our Spacetime solver.

- We do not consider the physical approaches as they are predominantly a different objective function and constraint for the spacetime approach. Also, because the methods have not been demonstrated on comparable classes of problems, analysis is difficult.
- The PFIK+F and Spacetime approaches are the only methods that have been demonstrated across a wide and similar set of editing problems, and have demonstrated viable interactive performance on these problems.

5.1. Comparison of Performance

The actual performance is as much a measure of the implementation as of the method. While modern processors provide impressive performance for applications, achieving the potential of these applications typically requires a great deal of tuning and care.⁶

More specific to our comparison, it seems that even within a given approach, there are a wide range of design decisions and tradeoffs that can be made. For example, compare the tradeoffs made in the spacetime approach to achieve interactive performance on a circa 1996 personal computer [12] with the decisions made to create a noninteractive system on a high-performance workstation where maintaining physical constraints were essential [27].

Similarly, Lee and Shin make concessions to performance in their implementation of the PFIK+F approach. Through the use of the simple geometric IK solver, they trade solution quality and generality for performance. In our implementation, we made fewer concessions by employing a more general, however more expensive, optimization-based solver. For our analysis, we consider the latter because it is more able to handle the range of constraints that the spacetime approach can. Specifically, it can allow arbitrary points to be positioned, multiple points to be placed simultaneously, joint limits to be used, and other types of constraints to be applied.

The iterative nonlinear solvers used for spacetime typically only require a small number of iterations to converge to a solution. This makes the performance difficult to characterize: a small perturbation in the initial conditions may require the solver to make an extra iteration or two, which is substantial when the solver requires only a handful of iterations altogether.

5.1.1. Computational complexity. To consider the complexity of each method, we assume that the motion has n frames. Each frame of the motion has a small number of constraints, therefore the problem must consider $O(n)$ constraints. Similarly, in each frame, the character has the same number of variables to describe its configuration, so the total number of variables is again $O(n)$.

Convergence of any algorithm that solves general nonlinear equations is not guaranteed. Therefore, it is impossible to place strong bounds on the time it takes to solve nonlinear equations, whether for an IK problem or for a spacetime solution. However, in practice, we do not seek an exact solution and typically bound the number of iterations by a small constant—if a suitable solution is not found in this number of iterations, we assume one will not be found. In practice, we find this number to be quite small for our solver. We therefore analyze the complexity of a single iteration of the nonlinear solver, since this does grow in a deterministically bounded way with problem size, and we assume that the number of iterations is bounded by a constant, so this complexity does apply to the overall solution.

⁶ Anecdotally, we have achieved very substantial speedups in our code by attending to issues in memory system performance. A cache miss can take as long as a large number of floating point operations.

4-Color Artwork




FIG. 12. The example of Fig. 8 is adapted using different methods for choosing the new footplant positions. The different colors represent the motion obtained by moving footplants to their beginning, middle, or end of their duration.

At each iteration of the nonlinear solver, constraint errors and derivatives are computed, a linear system is solved, and the estimate of the solution updated.⁷ In [11] we argue that constraint problems are sparse for many graphics applications because each constraint only affects a small number of variables independent of the overall size of the problem. This means that the constraints and their derivatives can be computed in time linear with the number of constraints and variables. The only step with worse than linear complexity is solving the linear system.

Under these assumptions, the computational complexity of the PFIK+F approach is $O(n)$. At each iteration of the solver, an IK problem must be solved for each frame of the motion. The amount of time each of these problems take is independent of the total length of the animation. Similarly, other steps in the process also are independent of the duration of the animation. We assume that the solver runs for a small constant number of iterations.

The analysis of the spacetime solver is more complex. Following the analysis of [11], we can see that the amount of work done for each constraint in evaluating and differentiating is a constant independent of the total number of constraints. The constraint Jacobian (the derivative of each constraint with respect to each variable) will consist of $O(n)$ rows (one per constraint), each having a small constant number of nonzero entries (each constraint refers to only a small number of instants in time). This sparse linear system, with $O(n)$ nonzero entries, can be solved in $O(n^2)$ time with an iterative method ($O(n)$ iterations each requiring an $O(n)$ sparse-matrix multiply [2]).⁸ Exploiting the banded nature of these matrices can reduce the complexity even further, yielding $O(n)$ complexity.

For each of these analyses, we have limited the solver to a constant number of iterations. With the PFIK+F approach, we have the choice of ensuring that despite this bounded number of iterations that either the temporal constraints are satisfied (by terminating with a filtering step) or that the spatial constraints are satisfied (by terminating with a step that skips the filtering). The spacetime approach does not provide this choice: because we

⁷This characterizes a wide class of methods [9, 24].

⁸Technically, iterative solution of the linear system with $O(n)$ nonzero elements takes $O(n^2)$ time only with infinite precision arithmetic. In practice, the use of damping and large tolerances allow us to achieve this performance.

Motion		PFIK+F	Spacetime		
Frames	Constraints	Time	Time	Iters	Time/Iter
890	4890	1.6	6.1	4	1.5
1780	9780	3.4	19.4	7	2.8
2670	14670	5.2	29.9	7	4.3
3560	19560	7.4	45.3	8	5.7
4450	24450	9.7	58.5	8	7.3
5340	29340	11.5	83.3	10	8.3
6230	34230	13.5	109	11	9.9
7120	39120	15.4	126	10	12.6
8010	44010	17.3	150	11	13.6
8900	48900	19.4	148	9	16.4
9790	53790	21.5	165	9	18.3
10680	58680	23.4	211	11	19.2
11570	63570	25.2	198	8	24.8
12460	68460	29.6	244	9	27.1
13350	73350	29.8	877	11	79.7

FIG. 13. Results of the performance experiment. The timings provided are computed on an IBM workstation with an 866 Mhz Pentium III Xeon Processor and 512 M of RDRAM.

cannot be guaranteed of convergence, we cannot know that the spatial constraints will be satisfied.

5.1.2. Performance experiment. To explore the asymptotic performance of the algorithms, we devised an experiment where we created very long motions with many constraints. In order to keep the “difficulty” of the numerical problem constant, we needed to use similar constraints. To create the long examples, we began with a loopable 89 frame walking motion on which footplant constraints were applied to the heel and toe. We repeated this motion as necessary to create a motion of the desired length.

Our test was to retarget the motion to a character 60% of the size of the original. The results of this experiment are shown in Fig. 13. As in the examples of Section 4.2.1, the smaller character is forced to walk in the footsteps of the original character to create a more challenging retargeting problem and the scaling heuristic is not used.

These trials test extremely long motions: the longest trial is over 7 min long. All of the experiments, except for one, performed the retargeting in less than real time. The anomalies in the longest trials are explained by the fact that this problem is too large to fit in physical memory.

The parameters for the solvers were chosen to be the fastest settings that produced visually acceptable results. This required four iterations of the PFIK+F solver, with a filter schedule that ended with a 3-point B-Spline filter kernel. As expected, the timings for the PFIK+F solver are linear (when the last example is excepted), as shown in Fig. 14.

The behavior of the spacetime solver seems less simple. For each test, a small number (between 4 and 11 of iterations was required by the solver. When the solver took more than 3 steps, it rarely made much progress (the last steps of the solver improved the solution by less than 1%). The solver’s choice to take more iterations does not seem to be coupled to the quality of the result, or the size of the problem. In theory, it is tied to the nonlinearity of

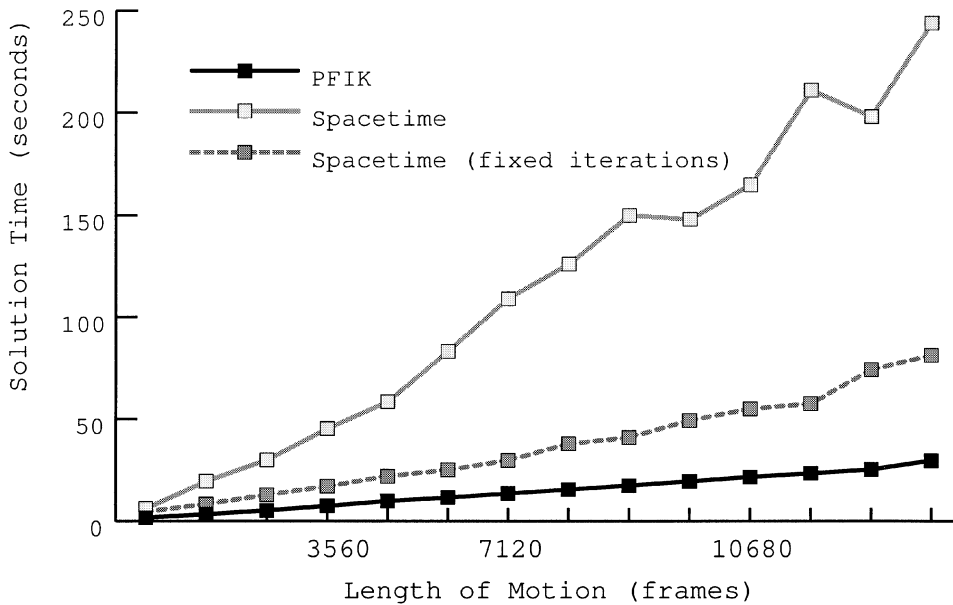


FIG. 14. Performance of the PFIK+F and spacetime solvers on problems of varying size. A third line shows the expected spacetime solver performance if a fixed number of iterations is used.

the problem. In practice, there seems to be a degree of luck as to how well the solver does at finding a solution and determining that it is unlikely to make substantial improvements by taking more steps.

If we limit the spacetime solver to 3 iterations, it achieves results comparable to the PFIK+F solver, so this timing is also provided. When the timing for the spacetime solver is examined on a per-iteration basis, we would expect the time to be dominated by the step of solving the sparse linear system. Because we use a conjugate-gradient solver (similar to the one described in [26]), we would expect $O(n^2)$ behavior. We instead see a result that is linear (to a very high statistical confidence). The best explanation for this is that the matrix is well enough conditioned, and the stopping tolerance is large enough, that an effectively constant number of iterations is required. Rather than claiming that this is likely to happen across problems, we use this example as a demonstration that the time required for the spacetime solver depends as much on how hard the problem is as how big the problem is.

The spacetime trials were run with the control points of the B-Spline placed on every third frame. This yielded a motion that was of acceptable quality and was visually superior to the PFIK results. Using a control point every fifth frame provided an even more natural looking motion, albeit one with more foot sliding (in terms of measured distance, not perceived quality). The smaller number of control points took a longer time to solve.

5.2. Range of Constraints Handled

Different types of solution techniques allow differing amounts of generality to the types of constraints that they may allow. The set of constraints that we might want to preserve in a motion is open. Ultimately, we might have constraints to specify not only geometric details, such as the positions of end-effectors and limits on joints, but also high-level properties

such as the “naturalness” of a motion, or its “mood.” To date, mathematical descriptions of such high-level properties have not been identified. However, generality in a solver means that such constraints are more likely to be usable when they are developed.

The class of constraints a system can handle depends on much more than the solver’s approach. The details of the implementation of the solver parts make limitations. For example, not all inverse kinematics solvers implement all possible IK constraints. Sometimes limitations on the availability of constraints are more than a solver issue: for constraints to be used effectively, there must be methods provided to specify and visualize them.

Our question is whether the motion editing *approach* affects the range of possible constraints. Just because an approach can handle a particular constraint does not mean that all implementations of the approach can. In this section we explore the generality afforded by the spacetime approach over that provided by an IK-based approach. That is not to say that a sophisticated PFIK approach may not implement a wider variety of constraints than a simple implementation of a spacetime approach. In this section, when we refer to the constraints of an approach, we are discussing that constraints are *possible* with *an* implementation of the approach. We are seeking fundamental limitations in the approaches.

Our first observation is that the set of spatial constraints possible in a spacetime approach is a proper superset of what an IK solver can do. At worst, the same equations used in an IK solver can be used in a spacetime solver applying to a single instant in time. In terms of temporal constraints, the situation is less clear. A spacetime solver can recreate a PFIK solver by using an objective function that does not couple frames. However, the preference for low frequencies given by the PFIK+F approach (with soft temporal constraints) is more difficult to recreate. At worst, it can be implemented within the spacetime approach by making repeated calls to the spacetime solver with gradually increasing frequency limits.

The spacetime approach can handle constraints that a per-frame approach cannot. A simple example is that it is possible to create constraints that hold over time. We call such constraints variational, as they place limits on the function curves. In a PFIK-based approach, constraints can only be placed at discrete instants. To date, this advantage has not been exploited: variational constraints are implemented by sampling at discrete instants.

A related issue is interframe constraints: that is, a constraint that relates two or more different instants in time. For example, we might want to specify that a hand is in the same place in several different frames. We note that this implies that we do not have a specific value in mind: if so, we could have equivalently specified a position for each individual constraint.

A major flexibility of the spacetime approach comes from the fact that it handles “don’t care” constraints well. That is, a constraint that does not fully specify a goal. Such constraints are common. While many constraints specify interactions with specific points in the world, an even larger class of constraints place only general restrictions. For example, most footplant constraints require that the foot be on the floor, and that the foot does not slip while on the floor. However, the exact position on the floor is typically unimportant (unless the character is walking in a mine field, crossing a river on rocks, or another such contrived case).

The information to determine values for the “don’t cares” can come from any part of the motion. Allowing the solver to determine the footplant positions based on information from other frames leads to more natural retargeting and more useful constraint-based control. For example, a character whose goal is moved may take larger steps in order to reach its destination.

5.3. *Order Independence*

Interframe constraints with “don’t cares,” as described in the above section, suggest another limitation of per-frame approaches: they are order dependent. We must choose the configuration for one frame before choosing another.

To illustrate the order independence property, consider placing a constraint on the motion specifying that the character must be standing in place. In a PFIK approach, each frame would have to look to the results of a previously solved frame to choose a position. Were it to choose an unsolved frame, that frame might change in the future. Consider if an additional constraint specifying the character’s position is placed on one of the frames of the motion. Were this frame solved first, then all other frames would end up in the right place, if another frame was chosen to be solved first, when the positioned frame’s turn to be solved arrived, it would have a conflict: it would need to choose between being in the same place as the previously solved frames or meeting its positional goals.

Some of the advantages of the spacetime approach might be recreatable in a PFIK+F approach if there were a mechanism for determining the order in which to solve constraints. An illustration of this is in Fig. 12, where a single time is chosen to determine the position of each footplant. The declarative, order independent, nature of constraints are one of the features that make constraint-based descriptions attractive.

5.4. *Quality of Result*

Getting an answer quickly is nice. However, it is also important to get “good” answers. Unfortunately, it is difficult to define what a good answer is since there is no ground truth answer to compare with. In a motion editing tool, one might define the correct answer as being what the user wants. Therefore, the quality of a solution is directly related to how well it can be controlled to provide a user with a desired result.

Just as the specifics of implementation may mean as much as the approach for the performance of solvers, they also have considerable influence on the quality of the resulting motions. For example, in a PFIK+F or PFIK approach, an IK solver that tends to pose the character in natural and plausible ways is more likely to provide good motions than a solver that doesn’t. An example showing differences between our spacetime and PFIK+F solvers is shown in Figs. 11a and 11b.

The specific examples applied have impact on the quality of the resulting motions. Achieving good results requires having good motions to start with.⁹ It also requires good spatial and temporal constraints to be defined for the initial motion. Both of these factors are somewhat independent of the motion editing approach. However, the approach can impact quality through the types of constraints it allows for and how well it operates in the face of “bad” constraints, such as conflicts, and under-specified cases.

5.5. *Summary: PFIK+F vs Spacetime*

The advantages of the spacetime approach:

1. Wider range of objective and temporal constraints possible, however this flexibility has not yet really been exploited.

⁹ Generally; the use of motion editing techniques to turn bad motions into good motions has not yet been fully explored.

2. Order independence.
3. Better handling of “don’t care” geometric constraints.
4. Interframe geometric constraints.

The advantages of the PFIK+F:

1. Simpler to implement. A PFIK+F solver can basically be build from “standard” parts in an animation system.
2. Can ensure solving the geometric constraints.
3. Easily allows for weak frequency limit constraints.
4. More predictable solution times.

6. CONCLUSIONS

In this paper, we have examined a range of methods of performing constraint-based motion editing. With this menagerie of techniques at their disposal, developers can choose techniques that fit their problems. Some classes of editing applications suggest particular solution approaches, for example PFIK for on-line applications. For general off-line editing applications, several approaches apply. Spacetime and PFIK+F methods, the two primary contenders in such applications, each have a set of strengths and weaknesses, and physically based approaches offer promise as a viable approach in the future.

ACKNOWLEDGMENTS

The author thanks Andy Gardner for his assistance in performance evaluation and picture making, Taylor Wilson of House of Moves Motion Capture Studios for the sample data, Jehee Lee and Sung-Yong Shin for their discussions of the methods, and Alex Mohr for his help with proofreading.

REFERENCES

1. Alias/Wavefront, Maya, computer software, 2000.
2. R. Barrett, M. Berry, T. Chan, J. Demmel, J. Donato, J. Dongarra, V. Eijkhout, R. Pozo, C. Romine, and H. van der Vorst, *Templates for the Solution of Linear Systems: Building Blocks for Iterative Methods*, SIAM, Philadelphia, 1994.
3. R. N. Bindiganavale, *Building Parameterized Action Representations from Observation*, Ph.D. thesis, University of Pennsylvania, 2000. [Appears as Technical Report MS-CIS-00-17]
4. N. I. Badler, K. H. Manoochehri, and G. Walters, Articulated figure positioning by multiple constraints, *IEEE Comput. Graphics Appl.* **7**(6), 1987, 28–38.
5. A. Bruderlin and L. Williams, Motion signal processing, in *SIGGRAPH 95 Conference Proceedings*, Annual Conference Series, pp. 97–104, Aug. 1995.
6. K. J. Choi and H.-S. Ko, On-line motion retargeting, *J. Visual. Comput. Animation* **11**, 2000, 223–243.
7. M. F. Cohen, Interactive spacetime control for animation, *Computer Graphics (Proceedings of SIGGRAPH 92)* **26**(2), 1992, 293–302.
8. Discreet, a division of Autodesk, Inc., 3d Studio Max, computer software, 2001.
9. R. Fletcher, *Practical Methods of Optimization*, Wiley, New York, 1987.
10. M. Gleicher and P. Litwinowicz, Constraint-based motion adaptation, *J. Visual. Comput. Animation* **9**, 1998, 65–94.
11. M. Gleicher, *A Differential Approach to Graphical Interaction*, Ph.D. thesis, School of Computer Science, Carnegie Mellon University, 1994.

12. M. Gleicher, Motion editing with spacetime constraints, in *Proceedings 1997 Symposium on Interactive 3D Graphics*, pp. 139–148, Apr. 1997.
13. M. Gleicher, Retargeting motion to new characters, in *SIGGRAPH 98 Conference Proceedings*, Annual Conference Series, pp. 33–42, ACM SIGGRAPH, Addison Wesley, July 1998.
14. M. Gleicher, Animation from observation: Motion capture and motion editing, *Comput. Graphics*, invited paper to appear.
15. M. Gleicher, Motion path editing, in *Proceedings 2001 ACM Symposium on Interactive 3D Graphics*, March 2001.
16. F. Sebastian Grassia, Practical parameterization of rotations using the exponential map, *J. Graphics Tools* **3**(3), 1998, 29–48.
17. J. U. Korein and N. I. Badler, Techniques for generating the goal-directed animation of articulated structures, *IEEE Comput. Graphics Appl.* **2**(9), 1982, 71–81.
18. S. Park Kwangjin Choi and H. Ko, Processing motion capture data to achieve positional accuracy, *Graphical Models Image Process.* **61**(5), 1999, 260–273.
19. P. C. Litwinowicz, Inkwell: A 2 $\frac{1}{2}$ -D animation system, in *Computer Graphics (SIGGRAPH '91 Proceedings)*, Vol. 25, pp. 113–122, July 1991.
20. J. Lee and S. Y. Shin, A hierarchical approach to interactive motion editing for human-like figures, in *Proceedings of SIGGRAPH 99*, pp. 39–48, Aug. 1999.
21. J. Lee and S. Yong Shin, *Multiresolution Motion Analysis and Synthesis*, Technical Report CS-TR-2000-149, Computer Science Department, KAIST, 2000.
22. S. Lee, G. Wolberg, and S. Y. Shin, Scattered data interpolation using multilevel b-splines, *IEEE Trans. Visual. Comput. Graphics* **3**(3), 1997, 228–244.
23. A. A. Maciejewski, Dealing with the ill-conditioned equations of motion for articulated figures, *IEEE Comput. Graphics Appl.* **10**(3), 1990, 63–71.
24. J. Nocedal and S. Wright, *Numerical Optimization*, Springer Series in Operations Research, Springer, New York, 1999.
25. C. B. Phillips and N. I. Badler, Interactive behaviors for bipedal articulated figures, *Comput. Graphics (Proceedings of SIGGRAPH 91)*, **25**(4), 1991, 359–362.
26. W. Press, B. Flannery, S. Teukolsky, and W. Vetterling, *Numerical Recipes in C*, Cambridge University Press, Cambridge, UK, 1986.
27. Z. Popovic and A. Witkin, Physically based motion transformation, in *Proceedings of SIGGRAPH 99*, pp. 11–20, Aug. 1999.
28. C. B. Phillips, J. Zhao, and N. I. Badler, Interactive real-time articulated figure manipulation using multiple kinematic constraints, in *1990 Symposium on Interactive 3D Graphics*, Vol. 24, pp. 245–250, March 1990.
29. H. J. Shin, J. Lee, M. Gleicher, and S. Y. Shin, Computer puppetry: An importance based approach, *ACM Trans. Graphics*, accepted.
30. D. Tolani, A. Goswanmi, and N. Badler, Real-time inverse kinematics techniques for anthropomorphic limbs, *Graphical Models* **62**, 2000, 353–358.
31. M. Unuma, K. Anjyo, and R. Takeuchi, Fourier principles for emotion-based human figure animation, in *SIGGRAPH 95 Conference Proceedings*, Annual Conference Series, pp. 91–96, ACM SIGGRAPH, Addison Wesley, August 1995.
32. C. Welman, *Inverse Kinematics and Geometric Constraints for Articulated Figure Manipulation*, Master's thesis, Simon Fraser University, Sept. 1993.
33. A. Witkin and M. Kass, Spacetime constraints, in *Computer Graphics (SIGGRAPH '88 Proceedings)*. Vol. 22, pp. 159–168, Aug. 1988.
34. A. Witkin and Z. Popović, Motion warping, in *SIGGRAPH 95 Conference Proceedings*, Annual Conference Series, pp. 105–108, Aug. 1995.
35. J. Zhao and N. I. Badler, Inverse kinematics positioning using nonlinear programming for highly articulated figures, *ACM Trans. Graphics* **13**(14), 1994, 313–336.