

Tiling Motion Patches

Manmyung Kim, Youngseok Hwang, Kyunglyul Hyun and Jehee Lee

Seoul National University




Figure 1: Multiple characters are interacting with each other. A diversity of human behavior, such as jumping, crawling, pushing, bowing and slapping, are simulated in the automatically-generated, interpenetration-free animation.

Abstract

Simulating multiple character interaction is challenging because character actions must be carefully coordinated to align their spatial locations and synchronized with each other. We present an algorithm to create a dense crowd of virtual characters interacting with each other. The interaction may involve physical contacts, such as hand shaking, hugging, and carrying a heavy object collaboratively. We address the problem by collecting deformable motion patches, each of which describes an episode of multiple interacting characters, and tiling them spatially and temporally. The tiling of motion patches generates a seamless simulation of virtual characters interacting with each other in a non-trivial manner. Our tiling algorithm uses a combination of stochastic sampling and deterministic search to address the discrete and continuous aspects of the tiling problem. Our tiling algorithm made it possible to automatically generate highly-complex animation of multiple interacting characters. We achieved the level of complexity far beyond the current state-of-the-art animation techniques could generate, in terms of the diversity of human behaviors and the spatial/temporal density of interpersonal interactions.

Categories and Subject Descriptors (according to ACM CCS):
I.3.7 [Computer Graphics]: Three-Dimensional Graphics
and Realism—Animation

1. Introduction

Simulating multiple characters interacting with each other is an emerging research topic in computer graphics. Modeling “interaction” between characters necessitates appropriate actions to occur in a carefully coordinated manner in

both space and time. Precisely coordinating the spatial location, timing, and action choices of many interacting characters is challenging and often requires prohibitive computation and/or memory resource.

We are particularly interested in creating a dense crowd of virtual characters that interact with each other. The interactions may occur between characters in close vicinity and often involve physical contacts, such as hand shaking, pushing, and carrying a heavy object collaboratively. Each character

may perform an arbitrary sequence of actions, which must be spatially aligned and temporally synchronized with other characters. Given a collection of action choices, each character can have exponentially many action sequences and only a small number of action sequences may satisfy stringent spatiotemporal constraints posed by interpersonal interactions. Finding such action sequences simultaneously for many interacting characters is extremely demanding.

We address the problem by collecting episodes of multiple characters. Each episode describes an interesting event featuring a small number of characters for a short period of time. Tiling episodes seamlessly across space and time would generate a crowd of densely interacting characters for an extended period of time. Each episode weaves a collection motion fragments to form a *motion patch* [LCL06]. The motion patch has its *entries* and *exits*, which correspond to the beginning and end frames, respectively, of its motion fragments. Two motion patches can be stitched together by carefully aligning their entries and exits in both space and time. Our major challenge is tiling patches tightly such that the tiling does not have any dangling entries or exits, which cause sudden appearance and disappearance of characters in the crowd.

The tiling is a combinatorial problem of layering a collection of motion patches. It also involves continuous optimization of patch locations for packing patches densely while avoiding collisions. Feasible solutions are rare and exhaustive search is not feasible. Our tiling algorithm uses a combination of stochastic sampling and deterministic search to address the discrete and continuous aspects of the problem. Our motion patches are deformable to add flexibility in stitching patches and thus alleviate the difficulty of tiling. The deformable patch allows its motion fragments to warp smoothly while maintaining the coherence of interpersonal interaction captured in the patch.

Our tiling algorithm made it possible to automatically generate highly-complex animation of multiple interacting characters. We achieved the level of complexity far beyond the current state-of-the-art animation techniques could generate, in terms of the diversity of human behaviors and the spatial/temporal density of interpersonal interactions. The spatiotemporal nature of our tiling algorithm allows us to deal with both static and dynamic virtual environments. We can also control the generation of animation by specifying when and which actions to occur at any specific location of the environment.

2. Related Work

There exists a vast literature on crowd simulation in computer graphics. Crowd simulation refers to the process of simulating the movement of a number of pedestrians, which are equipped with action skills mostly for navigating in virtual environments and capability of replicat-

ing human collective behavior. Interaction behaviors, such as collision avoidance [NGCL09, GCK*09], visual perception [OPOD10], and forming groups/formations [JCP*10] to name a few recent ones, have been explored. Multiple robot coordination pursues a similar goal with much emphasis on collision avoidance [LaV06, GMLM09]. Multiple robots are provided with start and goal configurations and search for collision-free paths simultaneously. It should be noted that our problem is different from crowd simulation and multi-agent path planning. We assume that interpersonal interactions would occur frequently and thus we search for a distribution of interactions rather than moving paths between start and goal configurations. Our characters are provided with a diversity of actions (beyond the scope of locomotion) to choose from. Every action should be precisely coordinated to meet alignment and synchronization posed by frequent interactions.

Research on multiple character interaction has recently emerged. Interactive manipulation techniques for editing multiple character motions has been studied with a varying level of control specification [KLLT08, KHKL09, HKT10]. A lot of researchers have focused on interaction between two characters. Liu et al. [LHP06] studied an optimization method to synthesize two character motions with physics-based objectives and constraints. Kwon et al. [KCPS08] simulated competitive interaction between two Taekwondo players using dynamic Bayesian networks. Komura and his colleagues addressed a similar problem using min-max search and state-space exploration [SKY10]. Wampler et al. simulated two-player adversarial games based on game-theoretic behavior learning [WAH*10]. The state-exploration and learning-based approaches often suffer from the curse of dimensionality. Adding extra participants causes exponential increase in either computation time or memory storage usage. Thus, the techniques that worked well with two characters do not scale easily to deal with many characters.

In computer animation, metaphorical patches have been used to encapsulate either pre-computed or motion capture data. Tiling patches spatially and temporally generates animation of larger scales. Cheney [Che04] discussed the use of square tiles embedding flow patterns. *Motion patches* by Lee et al. [LCL06] refer to a collection of building blocks of a virtual environment, which are annotated with motion data available on the blocks. Shum et al. [SKSY08] generalized the idea further to encapsulate short-term interactions between two characters into *interaction patches* and combined them to synthesize larger-scale interaction among many characters. Yersin et al. [YMPT09] tiled *crowd patches* to generate an arbitrarily large crowd of pedestrians in an on-the-fly manner. Our work is on the extension of these patch-based synthesis approaches. Given a collection of irregular patches, the notorious problems of tiling are gaps (temporal discontinuity) and overlaps (interpenetration between characters). The temporal gaps between dangling

entries and exits would cause characters to stay frozen over the gap or suddenly disappear and reappear at the boundaries. Our goal is tiling arbitrary irregular patches densely without gaps and overlaps.

The problem of packing irregular shaped pieces in a fixed area has been widely studied in geometric optimization and computational geometry [DD95] [Mil99]. The shape packing idea has been exploited in various graphics applications. Kim and Pellacini [KP02] generated image mosaics by filling an arbitrarily-shaped container with image tiles of arbitrary shapes. Fu et al. [FLHCO10] generated a set of quad-mesh to produce a tiled surface of computer-generated buildings. Merrell et al. [MSK10] presented a method for automatic generation of residential building layouts. Our approach were inspired by the research on geometric shape packing and we address further challenges that are unique in character animation. The challenges include the control of complex human behavior, handling multi-character interaction, identifying deformable motion patches as basic units of tiling, and the four dimensional tiling of deformable patches.

3. Deformable Motion Patches

Multiple characters interacting with each other create interesting episodes, which can serve as basic building blocks of complex multi-character scenes. Each episode can be modeled as a motion patch, which describes the actions of characters for a short period of time. Specifically, each motion patch includes a collection of motion fragments and their associated characters. Interaction between characters poses constraints between motion fragments. For example, hand-shaking indicates that the relative position and direction of two characters are constrained with each other while they are shaking their hands together. The motion patch may also include environment features and props that provide context to the characters' actions. The beginning and end frames of a motion fragment are called *entries* and *exits*, respectively, of the patch. We can stitch two motion patches if an exit of one patch matches an entry of the other. The character's position, direction, fullbody pose and timing should match within user-specified thresholds such that the character can make a smooth transition from one motion fragment to the other.

A *tiling* is a collection of motion patches that are stitched with each other. The tiling is *perfect* within a spatiotemporal region if it does not have any dangling entries or exits that do not have matching exits or entries in the tiling. More specifically,

Patch Deformation Our motion patches are deformable to allow flexibility in tiling. The deformation is computed in two steps: Path deformation and fullbody motion refinement. At the first step, we consider a web of motion paths (see Figure 2). Each *motion path* is the trajectory of the skeleton root in a motion fragment and undergoes defor-




Figure 2: The deformation of a motion patch. The manipulation handles and interpersonal interaction are formulated as linear constraints. In this example, the interpersonal interaction is the relative body location and direction of two characters when their hands hold each other.




Figure 3: Stitching deformable patches at their matching entries and exits. Each patch is depicted as a convex polygon that encloses its motion paths projected on the ground.

mation as if it is an elastic string. A collection of motion paths are constrained with each other using linear interpersonal constraints. The position/direction/time of entries and exits operates as if it is a manipulation handle, which can also be formulated as linear constraints. The manipulation handle is spatiotemporal in the sense that it specifies both body position/direction and the timing of the character's action. As a manipulation handle is dragged, a web of motion paths undergo deformation in an as-rigid-as-possible manner while maintaining a collection of linear constraints including both internal interactions and external manipulation. Specifically, we use multi-character motion editing by Kim et al. [KHKL09], which solves for the spatiotemporal deformation of multiple motion paths simultaneously subject to linear constraints based on Laplacian equations. The character's fullbody poses along motion paths are refined later to maintain hand-object contacts and rectify foot sliding artifacts.

Stitching Patches. We can stitch two deformable patches if there exist good matches at their corresponding entries and exits, through which characters can make smooth transitioning from one patch to the other. Let $\{(\mathbf{p}_i^A, t_i^A) \in \mathbb{R}^2 \times \mathbb{R} \mid i =$

$0, \dots, N\}$ be a motion path in patch A , where \mathbf{p}_i^A is the two-dimensional position of the character projected onto the ground surface and t_i^A is its timing. Note that the frames of motion data are usually sampled uniformly in time, though it is not necessary. The timing of motion or the interval of each individual frame may change as the motion data undergo deformation. The first frame $(\mathbf{p}_0^A, \theta_0^A)$ is an entry of patch A and the last frame $(\mathbf{p}_N^A, \theta_N^A)$ is its exit. Assume that the dissimilarity between the configuration (the position, direction, timing, and fullbody pose) at an exit of patch A and the configuration at an entry of patch B is below user-provided thresholds. Concatenating the Laplacian equations of two patches into a single larger system of linear equations and adding three boundary constraints would deform two patches to stitch them seamlessly.

$$\mathbf{p}_N^A = \mathbf{p}_0^B \quad (1)$$

$$\mathbf{p}_N^A - \mathbf{p}_{N-1}^A = \mathbf{p}_1^B - \mathbf{p}_0^B, \quad (2)$$

$$t_N^A = t_0^B \quad (3)$$

which enforce smooth transiting in position, direction and timing, respectively. If there are more than one matching entry/exit pairs, we specify three transitioning constraints for each matching pair.

4. Patch Construction

Achieving a perfect tiling is challenging even though the flexibility of patches alleviates its difficulty to a certain degree. It is important that a collection of motion patches should be designed carefully for better connectivity. The motion patch is supposed to encapsulate an interesting episode that may entail complex spatiotemporal relationships among multiple characters. On the other hand, it is expected that the boundary (entries and exits) of the patch should be as simple as possible. The simplicity of the patch boundary facilitates the process of patch stitching.

Given a collection of raw multi-character motion data, we would like to construct a collection of motion patches to be used for tiling. It involves following steps (see Figure 4):

- Identify interesting episodes from raw motion data.
- Specify interpersonal and person-object constraints to preserve the spatiotemporal integrity of multi-character interaction under patch deformation.
- Determine the boundary of motion patches.

Though there is no limitation on the size of motion patches, we usually design each patch to have less than five characters. If individual patches are too large, finding a perfect tiling could be difficult. Small patches including a single character are useful because they fit easily into narrow space. In practice, it is favorable to have a mixture of small, single-character patches and large, multi-characters ones.




Figure 4: The procedure of patch construction

4.1. Episode Detection

The episode refers to an event among multiple characters that is important or unusual. We use three criteria to detect such an event from raw motion data.

Contact. The physical contact between characters, such as handshaking, pushing, highfive, and punching, indicates the occurrence of interesting events.

Proximity. Even though there is no physical contact, it is noteworthy if two characters come close to each other and either of them reacts to the approach of the other. The presence of reaction is detected by measuring how much the character moves in a window of consecutive motion frames. The estimated variation of motion at frame i within window $[i - w, i + w]$ is:

$$V_i = \sum_{j=-w}^w \sum_{k=j+1}^w \|\mathbf{P}_{i+j} \ominus \mathbf{P}_{i+k}\|^2 \quad (4)$$

where $\mathbf{P}_{i+j} \ominus \mathbf{P}_{i+k}$ is the difference between two fullbody poses ignoring their horizontal translation and rotation about the vertical axis [LCL06]. The value of V_i above a certain threshold implies there might be a noticeable action at frame i . The threshold has been chosen empirically by trial and error.

Synchronicity. Even though two characters do not come close to each other, it is also noteworthy if two characters perform interesting actions simultaneously and they are aware of each other's action. Assume that motion data include two characters A and B interacting with each other. Character A has its motion data with index i , and character B has its own data with index j . The i -th frame of A and the j -th frame of B are concurrent in the global, reference time. The synchronicity between two characters is:

$$S_{ij} = V_i^A \cdot V_j^B \cdot e^{-g(A,B)} \cdot e^{-g(B,A)}, \quad (5)$$

where $g(A, B)$ is the angle between the gaze direction of character A and the vector from the face of A to the upperbody of B . This term estimates the possibility of character A being aware of the action of B at the moment.

The detection of a noticeable event between two characters results in imposing interpersonal constraints between relevant motion data. The constraints enforce the relative position, direction, and their timing of actions remain unchanged

even though their motion data undergo deformation. If we detect multiple events occurring nearby spatially and temporally, we merge them into a single, larger event, which may include interaction among more than two characters.

4.2. Boundary Determination

Once interesting events are detected, the next step is to determine their boundaries (entries/exits) to form a collection of tilable motion patches. There are several requirements. First, we would like to have the boundaries eventless. Any interaction/event occurring at boundaries would make the stitching process complicated. We want to have patches as compact as possible. Smaller patches tend to allow more flexibility in tiling and denser scenes to be created. Lastly, we want to have as little diversity of fullbody poses at boundaries as possible. Having a wide variety of fullbody poses at entries and exits would lead to difficulty of finding good matches in patch stitching. Note that we determine the boundaries of all patch collection simultaneously to satisfy the requirements.

Raw motion data include a multiplicity of fullbody poses at frames. Our algorithm for boundary determination begins by clustering fullbody poses into groups. Specifically, we use agglomerative hierarchical clustering such that the dissimilarity between any poses in a group is below a user-specified threshold. The threshold is chosen to allow smooth, seamless transitioning between motion fragments if the poses selected from a single group are at the matching entry and exit.

Let $\mathbf{G}_1, \dots, \mathbf{G}_k$ be k largest pose groups sorted in descending order (typically, $k = 10$ in our experiments). We intend to choose a small number of groups that suggest candidate frames to be entries and exists. Then, the boundary may be determined to surround each event tightly at frames that belong to the groups. The number of groups used is directly related to the diversity of fullbody poses, so we would like to choose as fewer groups as possible. To do so, we consider all possible subsets, $\{\mathbf{G}_i | 1 \leq i \leq k\}$ of cardinality one, $\{(\mathbf{G}_i, \mathbf{G}_j) | 1 \leq i < j \leq k\}$ of cardinality two, and so on. We examine all possible subsets and choose the one that minimize the weighted sum of the cardinality of the subset and the total size (number of motion frames) of motion patches. The coefficient weighs the diversity against the compactness of motion patches.

5. Tiling Algorithm

The goal of tiling is packing patches in a user-specified spatiotemporal region without any dangling entries or exits while achieving the desired density of multi-character interactions and avoiding interpenetration between them. Specifically, our tiling algorithm is based on simulated annealing, which is known to be an effective sampling method for combinatorial optimization problems [PTVF07]. Though simulated annealing has the ability to find a global optimal solution, it takes a prohibitive computation time to remove dan-

gling entries and exits completely. Our algorithm is a variant of simulated annealing, consisting of two phases. The first phase of the algorithm is intended to sample patches rapidly in the spatiotemporal region, while the second phase is a deterministic search and jittering procedure dedicated to removing dangling entries/exits.

Algorithm 1: Stochastic Sampling for tiling motion patches

\mathbb{P} : A tiling of motion patches;

$R(0, 1)$: generate a random number between 0 and 1;

```

1  $T \leftarrow T_{\max}$ ;
2  $i \leftarrow 0$ ;
3 while  $i < i_{\max}$  do
4    $\mathbf{P}_{\text{new}} \leftarrow$  (randomly choose a patch);
5    $\mathbb{P}' \leftarrow \text{Stitch}(\mathbb{P} \cup \mathbf{P}_{\text{new}})$ ;
6    $\Delta E \leftarrow E_{\text{tiling}}(\mathbb{P}') - E_{\text{tiling}}(\mathbb{P})$ ;
7   if  $e^{-\Delta E/kT} > R(0, 1)$  then
8     if  $\text{IsValid}(\mathbb{P}')$  then
9        $\mathbb{P} \leftarrow \mathbb{P}'$ ;
10       $i \leftarrow 0$ ;
11   $T \leftarrow T - \Delta T$ ;
12   $i \leftarrow i + 1$ ;
```

5.1. Stochastic Sampling

Given a collection of motion patches, the first phase of our tiling algorithm is based on stochastic sampling of patches (see Algorithm 1). The position/direction/timing of patch instances are randomly chosen in a virtual environment and added to the tiling one by one (line 4–5). The sampling procedure minimizes an energy function

$$E_{\text{tiling}} = \frac{N_d}{N_c + N_d} + \alpha \left(\frac{1}{N_c + N_d} \right), \quad (6)$$

where N_c and N_d are the number of connected and dangling entries/exits, respectively. The first term penalizes the occurrence of dangling entries/exits, while the second term encourages rapid sampling of patches. As the number of patches in the tiling increases, it tends to have more dangling entries/exits. α is a constant that balances between the density of patches and the number of dangling entries/exits. Each sample of a motion patch is accepted with certain probability based on simulated annealing (line 6–7). When the simulated temperature is high, the samples are easily accepted even though they do not improve the tiling energy much. As the temperature cools down, the system favors steady downhill traces to reach a stable solution. The Boltzmann parameter k and the annealing schedule (T_{\max} and ΔT) controls the rate of convergence. The algorithm terminates when the sampling procedure fails too many times (line 3).

Stitching, Deformation and Collision. As explained in the

previous section, stitching two patches leads to a large system of linear equations that concatenates two systems of Laplacian equations. Adding a new patch to the tiling of patches involves all patches in the tiling and simply concatenating all linear systems into a single huge one is often infeasible (line 5). We assume that the patches directly adjacent to the new one is deformable and all the others are pinned down. This assumption allows a practical trade-off between the computation time and the quality (in terms of flexibility and deformation) of tiling. The new patch is valid if it connects well to the tiling without incurring interpenetration or excessive deformation (line 8). We use a simple bounding box test on individual frames to check if there is interpenetration between two patches. We measure the degree of deformation on a per-frame basis. In our experiments, we allowed motion path deformation upto positional stretching of 3cm, bending of 7 degrees, and time shifting of 0.018 second per frame.

Algorithm 2: Deterministic search for removing danglings

\mathbb{P} : A tiling of motion patches;
 \mathbb{D} : A set of dangling entries and exits ;
 $N(\mathbb{P})$: the number of dangling entries and exits in \mathbb{P} ;

```

1 while  $\mathbb{D} \neq \emptyset$  do
2   foreach  $\mathbf{d} \in \mathbb{D}$  do
3     foreach ( $\mathbf{P}_{\text{new}}$  that has an entry/exit matching  $\mathbf{d}$ ) do
4        $\mathbb{P}' \leftarrow \text{Stitch}(\mathbb{P} \cup \mathbf{P}_{\text{new}})$ ;
5       if  $N(\mathbb{P}') \leq N(\mathbb{P})$  then
6          $i \leftarrow 0$ ;
7         while  $i < i_{\text{max}}$  do
8           if  $\text{IsValid}(\mathbb{P}')$  then
9              $\mathbb{P} \leftarrow \mathbb{P}'$ ;
10            break;
11            $\mathbb{P}' \leftarrow \text{Stitch}(\mathbb{P} \cup \text{Jitter}(\mathbf{P}_{\text{new}}))$ ;
12            $i \leftarrow i + 1$ ;
13       if  $\mathbf{d}$  is not dangling then
14         break;
15       if  $\mathbf{d}$  is dangling then
16          $\mathbb{P} \leftarrow \mathbb{P} \setminus \mathbf{P}(\mathbf{d})$ ;
  
```

5.2. Deterministic Search

Even though the energy function used in the first phase tends to minimize the number of danglings, it usually converges very slowly (see Figure 5). The second phase of the algorithm is based on combinatorial exploration and randomized jittering in a continuous domain (see Algorithm 2). It examines every possible patch sequences from each dangling entry/exit until there remains no more dangling (line 2). Such an exhaustive search is infeasible in large free space. However, the patch tiling obtained from the first phase of the al-




Figure 5: The plots of the tiling energy and the number of dangling entries/exits with respect to the computation time. The red line shows the time when the second phase of the algorithm begins.

gorithm covers the target domain uniformly because of the nature of random sampling. The free space is fragmented into narrow spots, in which patch connection is strictly constrained. In the narrow, fragmented spaces, even exhaustive search is not computationally expensive any more.

For any dangling, the algorithm iterates over all possible patch choices that have an entry/exit matching the dangling (line 3). We first check if the number of danglings is increased by the new patch (line 5). Otherwise, we check further if the new patch incurs interpenetration or excessive deformation (line 8). If the new patch is not valid, we explore a continuous search domain by randomly jittering its spatial location, direction and timing (line 11). In our experiments, the jittering range is initially set as a circle of radius 0.6m in XZ-plane, an angle interval of -15 to 15 degree, and a time interval of -0.3 to 0.3 seconds. As the iteration continues, the jittering range expands gradually to explore a wider range. The maximum range is a circle of radius 1.0m, an angle interval of -25 to 25 degree, and a time interval of -0.5 to 0.5 seconds. If the dangling entry/exit has no patch to connect with it, the patch with the dangling entry/exit is removed from the tiling (line 14–15). The removal of the patch makes yet another dangling entries/exits. As the algorithm iterates, it will examine other possibilities that has not been explored before to eventually remove danglings completely (line 3). We keep track of sampling history so that patches removed in line 15 is not chosen again in line 3.

5.3. User Control

Though our tiling algorithm is capable of automatically generating a crowd of interacting characters, it can also allow the user to have direct and indirect control over the generation of animation. The user control may influence the mo-

tion of each character indirectly by specifying a variety of environment features or the ratio/preferences of actions. The user control can be more direct and immediate to specify which action to occur when and where in the environment. We annotate each patch with labels, such as slow/fast and friendly/hostile, and allows the user to decide how frequently each class of actions to occur in the animation. More specifically, we define a preference function

$$F = e^{-f_1(x,y)/\sigma_1} \cdot e^{-f_2(\theta)/\sigma_2} \cdot e^{-f_3(t)/\sigma_3} \cdot e^{-f_4(l)/\sigma_4}, \quad (7)$$

where (x, y) is the position of a patch instance, θ is its orientation, t is the timing in the reference time, and l is its label. f_i controls the tendency of patch distribution and σ_i weighs the significance of the terms. For example, $f_1(x, y) = (x_d - x)^2 + (y_d - y)^2$ if we want to control the characters to move to their target location (x_d, y_d) . If we do not have any preference over labels, then $f_4(l) = 0$ for any l . f_i can be defined in various ways to provide the user with flexibility and versatility in control specification. The preference function applies differently to the aforementioned algorithms. In the random sampling process (line 4, Algorithm 1), the normalized preference function is used as a probability distribution of sampling \mathbf{P}_{new} . In the deterministic search (line 3, Algorithm 2), the preference function prioritizes the candidates to determine which patches to examine first.

Interactive Control. The second phase of the algorithm can be easily modified to deal with interactive control of a small number (less than a dozen) of characters. Two minor modifications are required. One is the way patch instances are sampled in time. While the original algorithm has a fixed time horizon, the interactive version of the algorithm advances the time horizon incrementally and patch instances are sampled continuously at runtime to fill up the horizon. The other is the way the tiling is deformed when a new patch instance is added. The original algorithm allows both the new patch and existing patches in the tiling to undergo deformation, while the interactive version allows only the new patch to deform. The rationale is that patch instances in the tiling is already in the past and should not change.

6. Experimental Results

The timing data in this section was measured on a 2.93GH Intel Core i7 computer with 8Gbyte main memory and an ATI Radeon HD 4550 graphics accelerator. In our experiments, parameter α , k , T_{max} and ΔT are 0.2, 0.03, 200 and 1, respectively, for all examples. We do not need to tune parameters for each individual example.

Motion Patch Construction. We captured two subjects performing a variety of interactive actions, such as dancing, tagging, and bumping each other. Some actions involved a lot of physical contacts, while our subjects sometimes exhibited interaction without touching each other. We also designed some patches manually using our motion editing system we

developed based on the idea of Kim et al. [KHKL09]. We constructed multi-character patches by putting together several pieces of single-person motion capture data and specifying interpersonal constraints.

Static Environments. Our first example was constructed in a static environment of $12\text{m} \times 12\text{m}$ square. We generated an animation of 600 frames (20 seconds). In the first phase of the algorithm, we assigned preferences values to the patches such that larger (more characters, larger spatial extent, extended period of time) patches are more likely to be sampled. The second phase of the algorithm added mostly small, single-character patches that can easily squeeze into narrow spots (see Figure 6). The computation time increases pseudo-exponentially with the density of patches. Constructing dangling-free tiling with 173, 374, 490, 657, and 907 patches took 0.3, 2, 12, 17, and 103 minutes, respectively (see top row in Figure 7).

Dynamic Environment. Our spatiotemporal tiling algorithm is inherently capable of dealing with dynamically-changing environments. Whenever a new patch instance is sampled, the interpenetration with (either static or dynamically-moving) obstacles and other patch instances are detected and avoided. The environment may have sources and sinks where characters are newly created or disappear from the scene. The alignment of sources and sinks generates the flow of character's movements. The environment in the second example has two dynamically-moving obstacles (see left image, middle row in Figure 7). One of the obstacles is tall and the other is shorter. The characters can jump over the shorter obstacle, but cannot jump over the other. It took about 90 minutes to construct a tiling with 528 patches featuring 66 characters that playbacks for 20 seconds.

Chicken Hopping. We captured two subjects chicken-hopping, which is a play between human players. The player must stand straight, lift up one of his/her leg, hold the ankle with hands, and hop on the other leg. The players bump each other to make the opponent loose balance. We recorded 50 seconds of the play and our patch construction algorithm identified 6 two-player patches from the data. Many of the patches exhibit two players bumping and some patches shows one player dashes and the other evades. We manually picked 15 single-player patches that allow each individual character to navigate and steer. We made two demos using the data. The first demo shows two teams of players (red and blue) aligning at opposite sides, rushing towards each other, and bumping. Each team has 15 players. The player bumps the opponent, but avoids the fellow. The animation is 22 seconds long. The generation of the animation took 7 minutes. The second demo shows interactive control of small groups of characters. Each team has four players and the user selects target locations interactively to control them. Our interactive algorithm run on the average at the rate of 45 frames per second.

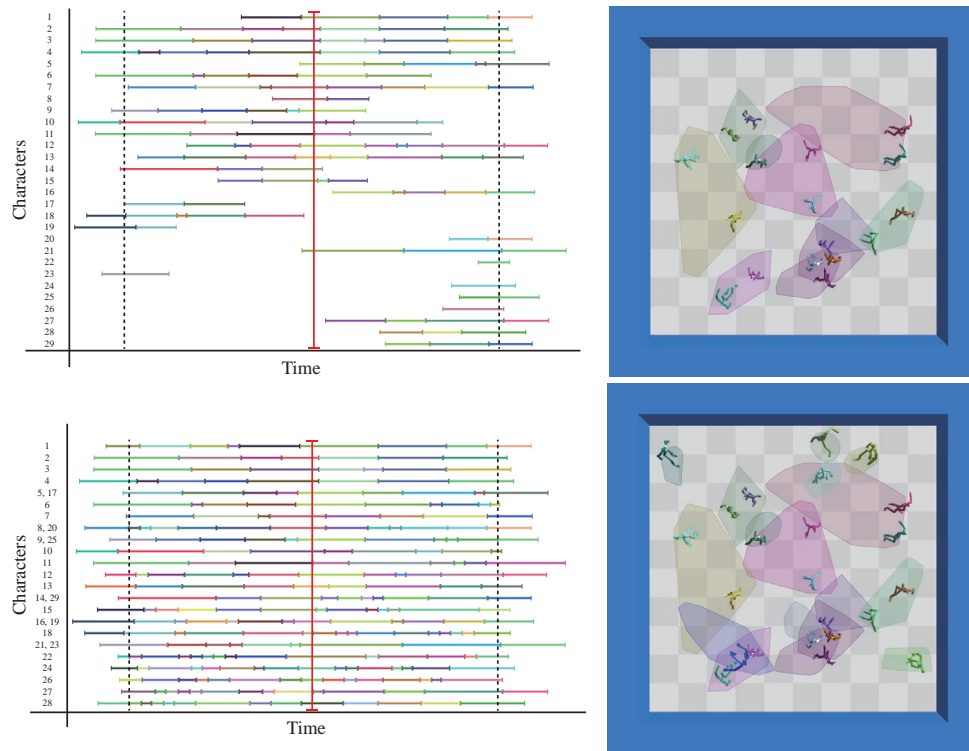


Figure 6: A small example for illustrating the tiling procedure. The spatiotemporal volume consists of a square spatial region of $10\text{ m} \times 10\text{ m}$ and a time interval of 600 seconds. The time interval is depicted as dashed lines. We cropped the animation at dashed lines. The screenshots on the right side were taken at the frame the red lines indicate. (Up) The first phase of our algorithm generated 67 patches featuring 29 characters. (Down) The second phase of the algorithm added new patches to fill in gaps. The final result included 197 patches featuring 23 characters. Note that the number of characters decreased, even though it has more patches. It is because the new patches filled in the gaps between the trajectories of characters 5 and 7, 8 and 20, 9 and 25, 16 and 29, and 21 and 23 to merge them.

Keyframing. The user can directly control the position/direction/timing of specific patch instances in the tiling. To do so, the user selects several patch instances and arranges them in the spatiotemporal domain as he/she wants. We use the arrangement as the initial configuration of the tiling in Algorithm 1, which will add new patch instances around the user-specified instances to make them fully connected. The resultant tiling thus obtained will include the user-specified patches at desired spatiotemporal location. Specifically, we constructed a large patch with ten characters aligning in two rows to bow to each other. We placed three instances of this large patch at 1.5, 11.5, and 21.5 seconds in the initial configuration of the tiling. The result shows the two-row formation is formed and dispersed repeatedly in complex multi-character animation (see the second image, middle row in Figure 7).

Spatially Repeating Tile. We can construct a seamless tile that repeats either spatially or temporal. Given a symmetric spatiotemporal region, the construction of repeating tiles al-

lows patch instances to wrap around the symmetric region. Spatial repetition of tiles can cover an arbitrarily large spatial region seamlessly, while temporal repetition of tiles generates an infinitely long animation. Our $10\text{ m} \times 10\text{ m}$ square tile repeats seamlessly on regular grids (see the third image, middle row in Figure 7). The repetition of square tiles is visible at large scale. It might be possible to construct aperiodic tilings, such as Wang tiles [CSHD03] and Penrose tiles [Gla98], by sampling motion patches on triangular or rectangular regions with proper boundary conditions.

Box movers. In our box mover example, we have four patches that involve props. The character can pick up a box from a stack of boxes, carry it around, pass it to another character, put it down on the ground, and sit down on the box (see right image, middle row in Figure 7). The characters are precisely aligned and synchronized with each other when they pass a box from one character to the other.



Example	# of patches	# of characters	Size (m^2)	Time of Scene (seconds)	Computation time (minutes)	
					Phase1	Phase2
Static	173	16	12×12	20	0.1	0.2
Static	374	34	12×12	20	0.2	1.8
Static	490	48	12×12	20	1	11
Static	657	60	12×12	20	3	14
Static	907	70	12×12	20	58	45
Dynamic	528	66	12×12	20	53	44
Control	628	55	12×12	23	17	14
Repeating	813	55	10×10	27	83	43
Object	526	36	10×10	27	1.5	5

Figure 7: Examples and Statistics.

7. Discussion

We explored a patch-based approach to create a dense crowd of characters interacting with each other. The formulation was reduced to tiling spatiotemporal patches without temporal gaps and spatial overlaps. Our stochastic sampling algorithm demonstrated the power of patch-based approaches in creating very complex animation by tiling a small collection of precomputed motion patches.

There is no guarantee that the stochastic sampling converges while achieving all desired properties, such as the density of the crowd, collision avoidance, obeying user control, and coping with environment features. In practice, the convergence of the algorithm is closely related to the diversity and flexibility of motion patches. Including various single-character patches facilitates the convergence because it allows each individual character to be steered independently.

The collision detection and avoidance is a bottleneck of computation. More than half of the computation time is spent for bounding box tests and jittering-based collision avoidance. Collision avoidance can be more efficient if the penetration depth between two motion patches can be computed. Choi et al. [CKHL11] demonstrated an interactively-controllable

character that is equipped with path planning capability and can navigate through highly-constrained environments in real-time. It was possible because all the obstacles had simple geometry and the interpenetration between the character and obstacles could be easily resolved. Unfortunately, we have no efficient algorithm to compute the penetration depth, which directly indicates the resolution of interpenetration, between four-dimensional spatiotemporal objects. Efficient computation of penetration depth would achieve significant performance gain for our tiling algorithm.

There have been research streams on patch-based motion synthesis [SKSY08, YMPT09]. Specifically, we have made two major improvements over the previous work. The first is the perfect tiling of motion patches. Shum et al. [SKSY08] has a lengthy discussion on why the perfect tiling is difficult and their observation motivated our work. We have found a solution to the problem that remain unsolved in the previous work. The two-phase algorithm and deformable patches played a key role in constructing the perfect tiling. The second improvement is the automatic algorithm for identifying motion patches from a collection of raw multi-character motion data. The overall contribution is a system that generates arbitrarily-large/long animations from a collection of raw

motion data. The baseline algorithms are fully automatic. The animator may intervene to change and control the result by tuning parameters, defining control functions, and adding special features.

Patch-based motion synthesis has demonstrated its promise in controlling animated characters in large scale while coping with the complexity of environments, the diversity of human behavior, and the computation of long-term planning. The data-driven animation community have considered fragments of motion capture data as basic building blocks of human motion synthesis. We advocate motion patches as building blocks of large-scale motion synthesis.

References

- [Che04] CHENNEY S.: Flow tiles. In *Proceedings of ACM SIGGRAPH/Eurographics Symposium on Computer Animation* (2004), pp. 233–242. 2
- [CKHL11] CHOI M. G., KIM M., HYUN K. L., LEE J.: Deformable motion: Squeezing into cluttered environments. *Computer Graphics Forum (Eurographics 2011)* 30, 2 (2011). 9
- [CSD03] COHEN M. F., SHADE J., HILLER S., DEUSSEN O.: Wang tiles for image and texture generation. *ACM Transactions on Graphics (SIGGRAPH 2003)* 22, 3 (2003). 8
- [DD95] DOWSLAND K. A., DOWSLAND W. B.: Solution approaches to irregular nesting problems. *European Journal of Operational Research* 84, 3 (1995), 506–521. 3
- [FLHCO10] FU C.-W., LAI C.-F., HE Y., COHEN-OR D.: K-set tilable surfaces. *ACM Transactions on Graphics (SIGGRAPH 2010)* 29, 4 (2010). 3
- [GCK*09] GUY S. J., CHHUGANI J., KIM C., SATISH N., DUBEY P., LIN M., MANOCHA D.: Clearpath: Highly parallel collision avoidance for multi-agent simulations. In *Proceedings of ACM SIGGRAPH/Eurographics symposium on Computer animation* (2009), pp. 177–187. 2
- [Gla98] GLASSNER A.: Penrose tiling. *Computer Graphics and Applications, IEEE* 18, 4 (1998), 78–86. 8
- [GMLM09] GAYLE R., MOSS W., LIN M. C., MANOCHA D.: Multi-robot coordination using generalized social potential fields. In *Proceedings of IEEE International Conference on Robotics and Automation* (2009), pp. 3695–3702. 2
- [HKT10] HO E. S. L., KOMURA T., TAI C.-L.: Spatial relationship preserving character motion adaptation. *ACM Transactions on Graphics (SIGGRAPH 2010)* 29, 4 (2010). 2
- [JCP*10] JU E., CHOI M. G., PARK M., LEE J., LEE K. H., TAKAHASHI S.: Morphable crowds. *ACM Transactions on Graphics (SIGGRAPH ASIA 2010)* 29, 6 (2010). 2
- [KCPS08] KWON T., CHO Y.-S., PARK S. I., SHIN S. Y.: Two-character motion analysis and synthesis. *IEEE Transactions on Visualization and Computer Graphics* 14, 3 (2008), 707–720. 2
- [KHKL09] KIM M., HYUN K. L., KIM J., LEE J.: Synchronized multi-character motion editing. *ACM Transactions on Graphics (SIGGRAPH 2009)* 28, 3 (2009). 2, 3, 7
- [KLLT08] KWON T., LEE K. H., LEE J., TAKAHASHI S.: Group motion editing. *ACM Transactions on Graphics (SIGGRAPH 2008)* 27, 3 (2008). 2
- [KP02] KIM J., PELLACINI F.: Jigsaw image mosaics. *ACM Transactions on Graphics (SIGGRAPH 2002)* 21, 3 (2002). 3
- [LaV06] LAVALLE S. M.: *Planning Algorithms*. Cambridge University Press, 2006. 2
- [LCL06] LEE K. H., CHOI M. G., LEE J.: Motion patches: building blocks for virtual environments annotated with motion data. *ACM Transactions on Graphics (SIGGRAPH 2006)* 26, 3 (2006). 2, 4
- [LHP06] LIU C. K., HERTZMANN A., POPOVIĆ Z.: Composition of complex optimal multi-character motions. In *Proceedings of ACM SIGGRAPH/Eurographics symposium on Computer animation* (2006), pp. 215–222. 2
- [Mil99] MILENKOVIC V. J.: Rotational polygon containment and minimum enclosure using only robust 2d constructions. *Computational Geometry* 13, 1 (1999), 3–19. 3
- [MSK10] MERRELL P., SCHKUFZA E., KOLTUN V.: Computer-generated residential building layouts. *ACM Transactions on Graphics (SIGGRAPH ASIA 2010)* 29, 6 (2010). 3
- [NGCL09] NARAIN R., GOLAS A., CURTIS S., LIN M.: Aggregate dynamics for dense crowd simulation. *ACM Transactions on Graphics (SIGGRAPH Asia 2009)* 28, 6 (2009). 2
- [OPOD10] ONDŘEJ J., PETTRÉ J., OLIVIER A.-H., DONIKIAN S.: A synthetic-vision based steering approach for crowd simulation. *ACM Transactions on Graphics (SIGGRAPH 2010)* 29, 4 (2010). 2
- [PTVF07] PRESS W. H., TEUKOLSKY S. A., VETTERLING W. T., FLANNERY B. P.: *Numerical Recipes 3rd Edition: The Art of Scientific Computing*. Cambridge University Press, 2007. 5
- [SKSY08] SHUM H. P. H., KOMURA T., SHIRAISHI M., YAMAZAKI S.: Interaction patches for multi-character animation. *ACM Transactions on Graphics (SIGGRAPH ASIA 2008)* 27, 5 (2008). 2, 9
- [SKY10] SHUM H., KOMURA T., YAMAZAKI S.: Simulating multiple character interactions with collaborative and adversarial goals. *IEEE Transactions on Visualization and Computer Graphics*, 99 (2010). 2
- [WAH*10] WAMPLER K., ANDERSEN E., HERBST E., LEE Y., POPOVIĆ Z.: Character animation in two-player adversarial games. *ACM Transactions on Graphics* 29, 3 (2010). 2
- [YMP09] YERSIN B., MAÍM J., PETTRÉ J., THALMANN D.: Crowd patches: populating large-scale virtual environments for real-time applications. In *Proceedings of symposium on Interactive 3D graphics and games* (2009), pp. 207–214. 2, 9