

Method of Computing Simulated Surfaces for Animation Generation and Other Purposes

CROSS REFERENCE TO RELATED APPLICATIONS

This application claims the priority benefit of U.S. Provisional Patent Application No.

- 5 62/983,534 filed February 28, 2020, and U.S. Patent Application No. 17/184,336 filed February 24, 2021, hereby incorporated by reference in their entirety as though fully set forth herein.

FIELD

- 10 [0001] The present disclosure generally relates to simulation for animation and other purposes and more particularly to efficient computation of surfaces of interacting simulated objects.

BACKGROUND

- 15 [0002] Visual representations of scenes intended to reflect real-world scenarios are common in animation and other fields. For example, a computer-generated imagery scene could be created by having an artist manually draw a sequence of frames to form a video sequence. For simple cartoons, for example, this is a feasible approach. However, as viewers have come to expect more complex visuals, there is a need for computer-driven imagery generation. Some of that computer-driven imagery generation might rely on simulation.
- 20 [0003] Computer simulation that is used for imagery generation has been used to animate natural phenomena as well as natural movements of characters, such as by using a physics engine to output movements of an articulated character that are consistent with real-world physics and joint constraints. In some ways, this is often a simple problem – how to determine natural-looking movements of at most a few dozen attached body parts. For other simulations, such as those with flexible objects, fluids, and the like, the number of degrees of
- 25 freedom of individual units is much greater and typically computer simulation requires a trade-off between realism, resolution, and amount of computing resources available. Because of this trade-off, efficient computer simulation techniques can be important as they might allow for an increase in realism and/or resolution without requiring significant increases in computing resources. Simulation computations involving fluid surface features and other
- 30 fluid interactions can often involve such trade-offs.

[0004] Imagery (e.g., animation frames, still images, etc.) generated to represent simulated objects, and the simulated objects themselves, can be created in computer-readable form, either procedurally or manually. For example, an image of a sphere might be generated procedurally from a user input of a center location, a radius parameter, and a color. In
5 another example, a more complex object might be generated procedurally from a set of parameters and a set of procedural rules for object creation. Objects might also be created manually, such as by having an artist draw into a computer system the shape of an object. In some cases, an artist might manually adjust a procedurally-generated object. In many instances, it might be preferred to maintain a procedural representation of objects further into
10 an imagery generation process so that editors have more opportunities to make adjustments at a procedural level rather than manual touch-ups.

[0005] A weighted principal component analysis process has been proposed in Koren, et al. (Koren, Y. and Carmel, L., "Visualization of Labeled Data using Linear Transformations", in Proceedings of IEEE Information Visualization, vol. 00:16, 2003) herein "[Koren]").

15 [0006] Weighted covariance matrices have been proposed in Yu, J., and Turk, G., "Reconstructing Surfaces of Particle-Based Fluids Using Anisotropic Kernels", ACM Transactions on Graphics (TOG), No. 5 (February 2013)(hereinafter "[Yu]").

[0007] It is an object of at least preferred embodiments to address at least some of the aforementioned disadvantages. An additional or alternative object is to at least provide the
20 public with a useful choice.

SUMMARY

[0008] A simulation system as described herein can process a procedural representation of objects simulated by particles to provide representations of object surfaces as described herein.

25 [0009] One general aspect includes a computer-implemented method for modeling surfaces of simulated objects under the control of one or more computer systems configured with executable instructions. The method includes receiving positions and sizes for each particle of a plurality of particles including a first object; identifying a subset of the plurality of particles whose positions are proximate to a first implied surface of the first object;
30 deforming particles of the subset of the plurality of particles to form a set of deformed particles by generating an anisotropy matrix or covariance matrix for each particle of at least some particles of the plurality of particles, where the anisotropy matrix or covariance matrix corresponds to a diagonal stretch matrix operated on by a non-linear function and where the

diagonal stretch matrix is a diagonal stretch matrix of neighboring particles for at least some particles of the plurality of particles; and computing a new shape or position of the first implies surface of the first object based on the set of deformed particles. Other embodiments of this aspect include corresponding computer systems, apparatus, and computer programs
5 recorded on one or more computer storage devices, each configured to perform the actions of the methods.

- [0010]** Implementations may include one or more of the following features. In some embodiments, the non-linear function includes a square root. In some embodiments, at least some of the particles of the plurality of particles are spherical. In some embodiments, at least
10 some of the deformed particles are ellipsoidal. In some embodiments, deforming the particles of the subset of the plurality of particles involves elongating axes of the particles of the subset that are parallel to the first implied surface of the first object. In some embodiments, the first object is solid. In some embodiments, the first object is semi-solid or fluid. In some embodiments, a shape or position of the first implied surface of the first object
15 changes based on an interaction of the first object with a second surface of a second object. In some embodiments, the computer-implemented method further includes generating a visual representation of the second surface of the second object. In some embodiments, the computer-implemented method further includes receiving masses for each particle of the plurality of particles including the first object. In some embodiments, the computer-
20 implemented method further includes computing a computed pressure, density, or force acting on at least one particle of the plurality of particles including the first object. In some embodiments, the computer-implemented method further including changing the positions of the at least one particle based on the computed pressure, density, or force. In some embodiments, the positions and sizes of the particles including the first object are received
25 from a simulation. In some embodiments, the positions and sizes of the particles including the first object are received from a scanning device. In some embodiments, identifying the subset of the plurality of particles whose positions are proximate to the first implied surface of the first object involves computing a weighted mean and the covariance matrix of each particle of the plurality of particles based on a user-defined search radius.
- [0011]** One general aspect includes a computer system for modeling surfaces of simulated objects, the system including: at least one processor; and a computer-readable medium storing instructions, which when executed by the at least one processor, causes the system to carry out the method. Some embodiment include a non-transitory computer-readable storage medium storing instructions, which when executed by at least one processor of a computer

system, causes the computer system to carry out the method. In some embodiments, computing the new shape or position of the first implied surface of the first object based on the deformed particles involves computing an isosurface of a scalar field. In some embodiments, the computer-implemented method further includes generating a visual representation of the first implied surface of the first object. Implementations of the described techniques may include hardware, a method or process, or computer software on a computer-accessible medium. In some embodiments, a scalar field that is computed is a signed distance field, wherein a value at a location in a scene can be returned that represents a distance from an implied surface.

- 10 **[0012]** The term ‘comprising’ as used in this specification means ‘consisting at least in part of’. When interpreting each statement in this specification that includes the term ‘comprising’, features other than that or those prefaced by the term may also be present. Related terms such as ‘comprise’ and ‘comprises’ are to be interpreted in the same manner.
- 15 **[0013]** This Summary is provided to introduce a selection of concepts in a simplified form that are further described below in the Detailed Description. This Summary is not intended to identify key features or essential features of the claimed subject matter, nor is it intended to limit the scope of the claimed subject matter. A more extensive presentation of features, details, utilities, and advantages of the surface computation method, as defined in the claims, is provided in the following written description of various embodiments of the disclosure and
- 20 illustrated in the accompanying drawings.

BRIEF DESCRIPTION OF THE DRAWINGS

- [0014]** Various embodiments in accordance with the present disclosure will be described with reference to the drawings, in which:
- 25 **[0015]** FIG. 1A illustrates an example of objects interacting that might be the subject of simulations.
- [0016]** FIG. 1B illustrates an example of objects interacting that might be the subject of simulations.
- [0017]** FIG. 1C illustrates an example of objects interacting that might be the subject of simulations.
- 30 **[0018]** FIG. 2 is a two-dimensional view of an interaction of an object and a surface of a fluid.
- [0019]** FIG. 3 illustrates a particle representation of a fluid object falling into another fluid.
- [0020]** FIG. 4A illustrates movement of particles.

[0021] FIG. 4B illustrates an implicit surface that may be computed based on the locations and surfaces of a plurality of particles.

[0022] FIG. 4C shows a modified surface that has been generated by the surface computation method.

5 [0023] FIG. 5 illustrates an example scene that may for example be the output of a simulation, a sensor, or a human artist, in accordance with at least one embodiment of the present disclosure.

[0024] FIG. 6 illustrates the same scene at a later stage of development, wherein the point cloud has been replaced with a cloud of overlapping ellipsoids, in accordance with at least
10 one embodiment of the present disclosure.

[0025] FIG. 7 illustrates the scene with a splatted or rasterized surface in place of the cloud of overlapping ellipsoids, in accordance with at least one embodiment of the present disclosure.

[0026] FIG. 8 illustrates the scene with a conditional point cloud, wherein an image
15 parameter is dependent on another variable, in accordance with at least one embodiment of the present disclosure.

[0027] FIG. 9 illustrates the scene with a conditional ellipsoid cloud, in accordance with at least one embodiment of the present disclosure.

[0028] FIG. 10 illustrates the scene with a splatted or rasterized surface computed from the
20 conditional ellipsoid cloud of FIG. 9, in accordance with at least one embodiment of the present disclosure.

[0029] FIG. 11 illustrates an example user interface, in accordance with at least one embodiment of the present disclosure.

[0030] FIG. 12 is a diagram of a data flow through a simulation system.

25 [0031] FIG. 13 illustrates an example visual content generation system as might be used to generate imagery in the form of still images and/or video sequences of images.

[0032] FIG. 14 is a block diagram illustrating an example computer system upon which computer systems of the systems illustrated in FIG. 12 or FIG. 13 may be implemented.

DETAILED DESCRIPTION

30 [0033] In the following description, various embodiments will be described. For purposes of explanation, specific configurations and details are set forth in order to provide a thorough understanding of the embodiments. However, it will also be apparent to one skilled in the art that the embodiments may be practiced without the specific details. Furthermore, well-

known features may be omitted or simplified in order not to obscure the embodiment being described.

[0034] In many of the examples described herein, inputs to a computer simulation system include parameters about the virtual material/object/fluid/etc. being simulated, and the output of a computer simulation includes the positions/mass/movement/etc. of the virtual material/object/fluid/etc. Such an output might be an input to an animation system, which can provide for rendering computer-generated imagery of the virtual material/object/fluid/etc. present in a scene in a virtual space. The computer-generated imagery might be still images, stereoscopic images, video sequences, and/or stereoscopic video sequences.

[0035] In some cases, the computer simulation of virtual elements seeks to match what would happen with corresponding real-world elements, but in other cases, artistic or other inputs are used in the computer simulation to create effects that do not correspond to anything in the real-world, or at least anything in available physical environments. For example, in a given simulation, an operator of a simulation engine might provide an input that corresponds to gravity “turning off” for a short period of time, which can be simulated but has no real-world correspondence. In another example, a block of water might be created that prior to an initial instant is constrained to be a particular shape, such as water being held behind a virtual dam or a cube of water with nothing holding it in that shape, and then the simulation begins with the water moving according to gravitational or surface tension constraints.

[0036] The present disclosure aids substantially in the computation and rendering of object surfaces, by improving the appearance of surfaces and reducing the resources required to render and store them. Implemented on a computing system in conjunction with a simulation system or 3D scanning system, the surface computation method disclosed herein provides a practical means of computing surfaces for objects that are modeled as collections of particles.

This improved surface computation transforms a less realistic, more resource-intensive modeling and rendering process into a more realistic, less resource-intensive process, without the normally routine need for an artist or animator to apply manual corrections. This unconventional approach improves the functioning of the simulation and rendering system.

[0037] The surface computation method may be implemented as a software program, at least some outputs of which may be viewable on a display, and operated by a control process executing on a processor that accepts user inputs from a keyboard, mouse, or touchscreen interface, and that may be in communication with a rendering process and a simulation or 3D scanning process. In that regard, the software program performs certain specific operations in response to different inputs or selections made at different times. Certain structures,

functions, and operations of the processor, display, sensors, and user input systems are known in the art, while others are recited herein to enable novel features or aspects of the present disclosure with particularity.

[0038] These descriptions are provided for exemplary purposes only, and should not be considered to limit the scope of the surface computation method. Certain features may be added, removed, or modified without departing from the spirit of the claimed subject matter.

[0039] FIG. 1A illustrates an example of objects interacting that might be the subject of simulations. In this example, a first vehicle 102 is moving to the right and a second vehicle 104 is moving to the left. A simulation might be performed to determine how the vehicles interact when they eventually collide, and a visualization might include a rendering of the surfaces of those vehicles as the surfaces are deformed due to the collision.

[0040] FIG. 1B illustrates an example of objects interacting that might be the subject of simulations. In this example, a boat 110 interacts with, and floats on the surface of, a body of water 112. A simulation might be performed to determine the propagation and interaction of waves along the surface of the water 112, and the rolling, pitching, and yawing movements of the boat 110 under the influence of these waves. A visualization might include the rendering of the complex surface of the water 110 as the waves reflect from, and refract around, the boat 110.

[0041] FIG. 1C illustrates an example of objects interacting that might be the subject of simulations. In particular, FIG. 1C illustrates an object 120 intersecting a cube of fluid 122 at a surface 124. In each case, representative visualizations of interactions of objects that change their surfaces might use a simulation to determine the interactions, and then a visualization would use the resulting surface changes for visualizing the interaction. The object 120 might for example be a single solid object (e.g., a car body or boat hull), or might be made up of a fluid (e.g., water, sand, etc.) or semisolid (e.g., deformable metal, etc.) that may be represented as a collection of particles.

[0042] FIG. 2 is a two-dimensional view of an interaction of the object 120 and the surface 124 of the fluid 122. If the fluid 122 is maintained in place by gravity, the view of FIG. 2 might be rotated to illustrate the object 120 falling into the fluid 122. Either way, a simulation may consider the interactions of the elements of the simulation, and a visualization may accurately render the resulting complex surfaces.

[0043] FIG. 3 illustrates a particle representation of the object 120 falling into the fluid 122. A simulator might generate particle representations of objects, such as a particle representation 302 of the object 120 and a particle representation 304 of the fluid 122. In this

example, object 120 is treated as being a fluid, as for example, water falling into water, albeit shown having an arbitrary shape in the figure for purposes of illustration. The simulator can compute particle-particle interactions to determine how the particles might move.

[0044] FIG. 4A illustrates movement of the particles 302 and 304 of FIG. 3. In this example, the simulator simulates movement of particles of the particle representation 302 and the particle representation 304 as a result of a collision or other interaction. From the moved particles, it is often desired to determine resulting surfaces, such as a surface 402 of the object 120 once deformed and a surface 404 of the fluid 122 once deformed. FIG. 4A illustrates where these smooth surfaces 402 and 404 might be at a particular moment during the collision, wherein the positions or shapes of the surfaces have changed as a result of the collision or other interaction.

[0045] FIG. 4B illustrates an implicit surface 408 that may be computed based on the locations and surfaces of a plurality of particles 406. The surfaces of the particles may for example be spherical surfaces based on a particle radius, or other surfaces determined by a particle's shape, size, and orientation. As illustrated, the surface 408 appears bumpy due to the particle representation, and a visualization might require a smoother surface, such as those illustrated in FIG. 4A. A smoother surface may for example make the rendering of the surface more realistic, more visually appealing, and/or computationally simpler.

[0046] In a simulation, a simulator might generate particle representations of objects, where each particle is treated as a solid object having a position, a mass, and a radius, among other parameters. The simulator can then run a simulation of particle interactions. While representing an object with many small particles might make resulting surfaces smoother, such an approach may come at a computational cost. With a smaller number of larger particles, the resulting computed surfaces may be bumpier, but can be smoothed by changing representations of particles as described herein. In some embodiments, such changes are handled in a way that does not introduce visual or computational artifacts.

[0047] In one simulation process, shapes of a particle in a particle representation are adjusted based on whether the particle is near a surface or is well within the object (e.g., well surrounded by neighbor particles). In an adjustment process, particle shapes and possible positions may be changed, and an implicit surface can be determined from the modified particle representations. Preferably, the resulting shapes correspond closely to what they would be if the simulation were not a particle simulation, or if it were a particle simulation involving many more particles. For example, it might be desirable that the resulting surfaces from the simulation be touching if they would be touching at full resolution, and that the

particle modifications to smooth the surfaces do not undesirably modify the object shapes or volumes. Once the simulated surfaces or implicit surfaces are determined, they can be stored or provided to an animation system. More details of an example modification process to smooth surfaces is described below.

- 5 **[0048]** The simulator, or a portion thereof or add-on thereto, reconstructs surfaces of objects (which can be solid, semi-solid, fluids, etc.) from particle-based simulations. This can be used in animation. Prior to simulating an interaction between two objects (or an object and a fluid, or two fluids, etc.), the simulator determines a particle representation of each object and then simulates those particles interacting and determines where the surfaces of those objects
10 would be, based on particle movements. In some cases, interaction is between a solid object and a fluid, so the particles of the solid object might not move relative to each other, in which case its surface does not need to be reconstructed and the particles of the fluid move relative to each other based on forces imposed by gravity, the fluid itself and the solid object, etc. With each particle having parameters such as radii, masses and locations, density and
15 pressure can be calculated and forces on particles determined. Thus, a task to be performed by the surface computation method is to determine a modified particle representation from which a surface 408, implicit or otherwise, can be computed.

- [0049]** In one such modified representation, particles 406 are simulated as spheres, but then are transformed based on a neighborhood and a kernel. This might be done using a weighted
20 principal component analysis process, such as that proposed in [Koren] to compute a weighted covariance matrix C for each particle, with a zero empirical mean and a set of eigenvectors representing particle orientations, and/or a weighted covariance matrix C as proposed in [Yu].

- [0050]** In this approach, a weighted mean and a covariance matrix of each particle is
25 computed. This might be as shown in Equations 1 and 2, where w_{ij} is an isotropic weighting function and x_i is a position of particle i .

$$C_i = \sum_j w_{ij} (\mathbf{x}_j - \mathbf{x}_i^w)(\mathbf{x}_j - \mathbf{x}_i^w)^T / \sum_j w_{ij} \quad (\text{Eqn. 1})$$

$$\mathbf{x}_i^w = \sum_j w_{ij} \mathbf{x}_j / \sum_j w_{ij} \quad (\text{Eqn. 2})$$

- 30 **[0051]** The function w_{ij} might be as in Equation 3, with r_i being a user-defined search radius.

$$w_{ij} = \begin{cases} 1 - (\|\mathbf{x}_i - \mathbf{x}_j\|/r_i)^3 & \text{if } \|\mathbf{x}_i - \mathbf{x}_j\| < r_i \\ 0 & \text{otherwise} \end{cases} \quad (\text{Eqn. 3})$$

[0052] A singular value decomposition of C_i corresponds to transformation of spherical particle i (e.g., a particle 406) into an ellipsoidal particle 40 to account for being near a surface of the particle representation. From this, an orthogonal rotation matrix, $\mathbf{R}_i$, and diagonal stretch matrix, Σ_i , could be computed through a singular value decomposition such that $C_i = \mathbf{R}_i^T * \Sigma_i * \mathbf{R}_i$.

[0053] An improved anisotropy matrix or covariance matrix, $\mathbf{G}_i$, may then be computed for each particle i , according to Equation 4, wherein $\Sigma_i^{-1/2}$ is one over the square root of a diagonal stretch matrix of neighboring particles (e.g., particles that are neighbors of particle i). The anisotropy matrix or covariance matrix may for example include units of distance squared, such that the square root of the matrix yields units of distance. In some embodiments, other non-linear functions may be used instead of or in addition to a square root, to operate on the diagonal stretch matrix in order to compute the anisotropy matrix or covariance matrix.

$$\mathbf{G}_i = \mathbf{R}_i (\Sigma_i^{-1/2}) \quad (\text{Eqn. 4})$$

[0054] Using the above, a spherical particle i is transformed linearly by $\mathbf{G}_i$ to form an ellipsoid to smooth the object's surface 408, thus yielding a smoothed surface 410. From the computed values of $\mathbf{G}_i$, anisotropic kernels, the scalar W_i , can be computed according to Equation 5, where P is a spherical kernel, such as the positive scalar value $P(x) = 1 - x/2$, for $0 \leq x < 2$ and 0 otherwise.

$$W_i(\mathbf{r}, \mathbf{G}_i) = \sigma P(\|\mathbf{G}_i \mathbf{r}\|) \quad (\text{Eqn. 5})$$

[0055] A representation of the modified (e.g., smoothed) surface 410 can then be computed as an isosurface of the scalar field according to Equation 6.

$$\phi(x) = \sum_j W(x - \bar{x}_j, \mathbf{G}_j) \quad (\text{Eqn. 6})$$

[0056] In an alternative embodiment, a representation of the modified surface 410 might be instead computed according to Equation 7.

$$\phi(x) = \max_j (W(x - \bar{x}_j, \mathbf{G}_j)) \quad (\text{Eqn. 7})$$

[0057] In some embodiments, a module that receives inputs representing details of a scene might return a representation of a surface, such as an isosurface. In some embodiments, the module might receive an input argument representing a position in a scene and return a signed distance field value, wherein the signed distance field value represents a distance from

the input argument position to a computed surface. In such embodiments, the surface might not be computed directly, but can be inferred from input argument positions that return a signed distance field value of zero or near zero. In some examples, a signed distance field value is near zero for points near an implied surface, is negative at points deemed to be inside an object for which the implied surface is a surface, and is positive at points deemed to be outside that object's surface.

[0058] FIG. 4C shows a modified (e.g., smoothed) implicit surface 410 that has been generated by the surface computation method of the present application. As described above, particles 406 that are near a surface of the object (for example, particles that are not surrounded by neighbors on all sides) are transformed from spheres into ellipsoids 407. For example, the ellipsoids 407 may have axes 409 that are parallel to the surface of the object. These axes 409 may be elongated by the surface computation method, such that (a) the particles 406 near the surface 410 occupy a greater portion of the surface, (b) the gaps between the ellipsoids 407 are smaller than the gaps between the unmodified particles 406, and/or (c) the outer surfaces of the ellipsoids present a flatter, less peaked profile with respect to the object surface. As a result, an implicit surface 410 that is derived from the ellipsoids 407 may be smoother than an implicit surface 408 that is derived from the unmodified particles 406. A smoother surface may in some cases be advantageous, for example by being more realistic, more visually appealing, less computationally intensive or memory/storage intensive, etc.

[0059] The process of determining a splatted or rasterized surface (e.g., an implicit surface 408 as shown for example in FIG. 4B, or a smoothed surface 410 as shown for example in FIG. 4C) is described with reference to simulation, but it might also be the case that the particle representation for which a surface is to be determined is obtain for other than simulation purposes. For example, an object (fluid or otherwise) might be represented by a point cloud obtained from a scanning device that generates point clouds from physical objects, such as a 3D laser scan, and the methods described herein could be used to reconstruct a surface of such an object.

[0060] FIG. 5 illustrates an example scene 500 that may for example be the output of a simulation, a sensor (e.g., a 3D camera), or a human artist, in accordance with at least one embodiment of the present disclosure. At this stage of development, the scene 500 comprises a point cloud or plurality of particles describing three different fluid example shapes: a wave 520, a vertical spout 530, and a cube. In general, the vertical spout 530 is a convoluted 1D feature, whereas the wave 520 is a convoluted 2D feature, and the cube 540 is a 3D feature.

Scenes may include any combination of 1D, 2D, and 3D objects or features, viewable and manipulable within a Graphical User Interface (GUI). By showing these features as a point cloud or plurality of particles, the system enables an artist to see the shapes, positions, orientations, and other attributes of the particles, as well as interactions of the various objects or features, or the voxel size of the simulation, without the computing time required to (1) compute ellipsoidal shapes for each point in the point cloud, (2) compute a surface to fit the ellipsoids, or (3) do a full render of the image. Thus, if adjustments to the size, shape, position, or interaction of the objects are made at this stage, the process can be quick and convenient. It is understood that the point cloud may be displayed as points or particles, or as spheres, polyhedrons or other shapes of selectable size.

[0061] FIG. 6 illustrates the same scene 500 at a later stage of development, wherein the point cloud 510 has been replaced with a cloud or plurality of overlapping ellipsoids 610, in accordance with at least one embodiment of the present disclosure. The dimensions and orientation of each ellipsoid depend on the number and direction of its neighbors within a given search radius of the ellipsoid (as described above for example in Figure 4C), such that, for example, the edges 642 of the cube 540 and the edges 622 of the wave 520 show larger ellipsoids than other areas of the cube or wave, while the corners 644 and 624 show still larger ellipsoids. In this example, the vertical spout 530 consists of very large ellipsoids, because the spout 530 comprises only a single 1D structure, with each ellipsoid having only 1 or 2 neighbors. As described above, this anisotropy helps generate smoother surfaces. There is a burden of computing time associated with computing the positions, orientations, and sizes of these ellipsoids. However, edits made to the scene 500 at this stage (for example, moving, reorienting, or resizing objects, or adjusting the output voxel size, related to the sizes of the ellipsoids) will nevertheless permit the artist to adjust the scene 500 without the computational burden of computing a surface to fit the ellipsoids, or do a full scene render.

[0062] FIG. 7 illustrates scene 500 with splatted or rasterized surfaces 710 in place of the cloud of overlapping ellipsoids 610, in accordance with at least one embodiment of the present disclosure. In the example of FIG. 7, parameters have been selected (e.g., through a user interface as described below in FIG. 11) such that the splatted or rasterized surface is smooth in all locations. However, this may not always be desirable. For example, a wave 520 in the real world may for example include smooth portions as well as rough, foamy, aerated portions. It is therefore desirable for the graphical user interface (GUI) to enable an artist to vary parameters of the scene 500 based on position, or on other factors. Once the

artist is happy with the appearance of the scene, a full or partial rendering of the scene can be performed.

[0063] FIG. 8 illustrates scene 500 with a conditional point cloud 810, wherein an image parameter is dependent on another variable, in accordance with at least one embodiment of the present disclosure. In the example shown in FIG. 8, the color of each point in the conditional point cloud 810 is dependent on its position along the X- axis, such that points toward the right side of the scene 500 appear darker, whereas points toward the left side of the scene 500 appear lighter. In other examples, other parameters may be conditional on position, or maybe conditional on other variables. For example, a color variable may be conditional on an X, Y, or Z position variable, and a voxel size may be dependent on the color variable. Such conditional effects may be used for example to permit some portions of a wave to be white, while others are blue or transparent.

[0064] FIG. 9 illustrates scene 500 with a conditional ellipsoid cloud 910, in accordance with at least one embodiment of the present disclosure. In the example shown in FIG. 9, the sizes of the ellipsoids are dependent on their position along the X-axis, such that ellipsoids toward the right side of the scene 500 are smaller, whereas ellipsoids toward the left side of the scene 500 are larger. Since the size of the ellipsoids affects the smoothness and continuity of any surface generated from them, an artist editing image parameters at this stage can have a good understanding of how the image will look, without the computational overhead of computing the surface and rendering the image.

[0065] FIG. 10 illustrates scene 500 with a splatted or rasterized surface 1010 computed from the conditional ellipsoid cloud 910 of FIG. 9, in accordance with at least one embodiment of the present disclosure. In the example of FIG. 10, portion 1022 of wave 520 and portion 1042 of cube 540 appear foamy, frothy, or aerated, with gaps appearing in a relatively rough surface. Similarly, portion 1024 of wave 520 and portion 1044 of cube 540 appear as individual points, droplets, or spray, each completely surrounded by air and not connected to its neighbor points. An artist making changes at this stage can see the calculated surface, but without the computational overhead of actually rendering the scene. Thus, a person of ordinary skill in the art will appreciate that the present disclosure advantageously provides artists with the ability to view and edit a scene at various stages of development, without the need to generate subsequent visual steps until the results of the present step are deemed satisfactory. This improves the efficiency of artists and reduces the need for computing time.

[0066] FIG. 11 illustrates an example user interface (UI) 1110, in accordance with at least one embodiment of the present disclosure. The surfacing node takes in particles, and

performs the neighbor analysis based on Input parameters, which say how far around to a look to analyze each particle's neighborhood. If the artist or operator knows the voxel size the simulation has been run with, a good search radius is on the order of that voxel size, e.g., between 0.5 and 3.0 times the voxel size. Artists also often vary simulation point radius (i.e. instead of simulating with uniform particle size they change it per particle for whatever reason then have). Sometimes it is desired to make the search radius proportional to particle size, so that's what the "Override with pscale" flag 1125 is for. This conveniently permits the artist or operator to proceed without re-tuning the search radius every time.

[0067] The Controls block allows the users to manipulate the neighborhood information gathered on each particle. The control mode can either be Global or Per-particle, and that's where the convenience of user-defined masking comes in. Users may employ custom parameters they have previously assigned to particles (like "@alpha" in the example shown in FIG. 11) to modulate the description of particle neighborhoods manifesting themselves as particular shapes of ellipsoids. Each neighborhood/ellipsoid is described by a set of values per particle, such as rotation matrix, and singular values. The artist or operator either manipulates those values directly (e.g., by using a masking control to manually select or define a neighborhood, and one or more artistic controls to define one or more attributes of the particles within the neighborhood, wherein the masking and artistic controls may be known in the art or may be customized to or unique to the system disclosed herein), or uses controls expressed as attributes: @maxElongation, @positionSmoothing etc, which are used by the algorithm internally to modify the ellipsoids, and may be affected using code-like arithmetic expressions as shown.

[0068] The output can be set to a fast Ellipsoids mode, so the artist or operator can directly view the particle data as a cloud of colored ellipsoids (as shown for example in FIG. 6), avoiding expensive splatting to surface (as shown for example in FIG. 7). Once the user is happy with the result they can flip it to Signed Distance (and wait longer, since it is more expensive) to obtain and display the splatted or rasterized surfaces of the final renderable volume.

[0069] In order to achieve the effects and image development stages described above and shown in FIGS. 5-10, the UI 1110 includes a first slider 1110 for selecting a voxel size employed by the particle simulation, and a second slider 1120 for selecting a "gather radius" or "search radius" that helps to determine how neighboring particles affect one another. The UI 1110 also includes an output type selector, that permits the ability to select the voxel type (spheres, ellipsoids, etc.), as well as a size multiplier slider to determine how output voxel

size if affected by neighbor distance. The UI 1110 also includes a control type selector that permits the artist or operator to select how particles are affected by a given variable (e.g., as individual particles, as neighborhoods, globally across the entire point cloud, etc.), and an expression window 1150 that includes code-like expressions to make certain variables (e.g.,
5 elongation, flatness, smoothing, volume, and global scaling of the ellipsoids) dependent on another variable (e.g., position, color, etc.). The code-like expressions may for example include variables, constants, or arithmetic operators. With UI 1110, an artist or operator is able to generate the effects and image development stages shown in Figs 5-10, much more efficiently than if it were necessary to generate ellipsoids, generate the surface, and render the
10 scene each time a change is made.

[0070] In some embodiments, other options may be available, such as the ability to render a scene only in black and white, or the ability to render the scene with a maximum number of vertices per traced light path. As with the other UI options, this permits the artist to make changes at an earlier development stage, while having a good idea what they will look like in
15 the final render, but with much lower computational overhead. Thus, the generation of fluid surfaces can be performed more efficiently and with more realistic results. This combination of improved realism and reduced computing time represents a substantial improvement over the related art.

[0071] FIG. 12 is a diagram of a data flow through a simulation system 1200 when the
20 system 1200 is performing a process such as one of those described herein. An input to a simulation system 1204 is object representations 1202, interface details 1206, and parameter values 1212, perhaps obtained by artist or user 1216 input via a computer 1214. An output 1208 of a simulation might be represented in memory and provided for example to an animation creation system 1210.

[0072] FIG. 13 illustrates the example visual content generation system 1300 as might be
25 used to generate imagery in the form of still images and/or video sequences of images. Visual content generation system 1300 might generate imagery of live action scenes, computer generated scenes, or a combination thereof. In a practical system, users are provided with tools that allow them to specify, at high levels and low levels where necessary,
30 what is to go into that imagery. For example, a user might be an animation artist like artist or user 1216 illustrated in FIG. 12) and might use visual content generation system 1300 to capture interaction between two human actors performing live on a sound stage and replace one of the human actors with a computer-generated anthropomorphic non-human being that behaves in ways that mimic the replaced human actor's movements and mannerisms, and

then add in a third computer-generated character and background scene elements that are computer-generated, all in order to tell a desired story or generate desired imagery.

- [0073]** Still images that are output by visual content generation system 1300 might be represented in computer memory as pixel arrays, such as a two-dimensional array of pixel color values, each associated with a pixel having a position in a two-dimensional image array. Pixel color values might be represented by three or more (or fewer) color values per pixel, such as a red value, a green value, and a blue value (e.g., in RGB format). Dimensions of such a two-dimensional array of pixel color values might correspond to a preferred and/or standard display scheme, such as 1920-pixel columns by 1280-pixel rows or 4096-pixel columns by 2160-pixel rows, or some other resolution. Images might or might not be stored in a compressed format, but either way, a desired image may be represented as a two-dimensional array of pixel color values. In another variation, images are represented by a pair of stereo images for three-dimensional presentations and in other variations, an image output, or a portion thereof, might represent three-dimensional imagery instead of just two-dimensional views. In yet other embodiments, pixel values are data structures and a pixel value can be associated with a pixel and can be a scalar value, a vector, or another data structure associated with a corresponding pixel. That pixel value might include color values, or not, and might include depth values, alpha values, weight values, object identifiers or other pixel value components.
- [0074]** A stored video sequence might include a plurality of images such as the still images described above, but where each image of the plurality of images has a place in a timing sequence and the stored video sequence is arranged so that when each image is displayed in order, at a time indicated by the timing sequence, the display presents what appears to be moving and/or changing imagery. In one representation, each image of the plurality of images is a video frame having a specified frame number that corresponds to an amount of time that would elapse from when a video sequence begins playing until that specified frame is displayed. A frame rate might be used to describe how many frames of the stored video sequence are displayed per unit time. Example video sequences might include 24 frames per second (24 FPS), 50 FPS, 140 FPS, or other frame rates. In some embodiments, frames are interlaced or otherwise presented for display, but for clarity of description, in some examples, it is assumed that a video frame has one specified display time, but other variations might be contemplated.
- [0075]** One method of creating a video sequence is to simply use a video camera to record a live action scene, i.e., events that physically occur and can be recorded by a video camera.

The events being recorded can be events to be interpreted as viewed (such as seeing two human actors talk to each other) and/or can include events to be interpreted differently due to clever camera operations (such as moving actors about a stage to make one appear larger than the other despite the actors actually being of similar build, or using miniature objects with other miniature objects so as to be interpreted as a scene containing life-sized objects).

[0076] Creating video sequences for story-telling or other purposes often calls for scenes that cannot be created with live actors, such as a talking tree, an anthropomorphic object, space battles, and the like. Such video sequences might be generated computationally rather than capturing light from live scenes. In some instances, an entirety of a video sequence might be generated computationally, as in the case of a computer-animated feature film. In some video sequences, it is desirable to have some computer-generated imagery and some live action, perhaps with some careful merging of the two.

[0077] While computer-generated imagery might be creatable by manually specifying each color value for each pixel in each frame, this is likely too tedious to be practical. As a result, a creator uses various tools to specify the imagery at a higher level. As an example, an artist might specify the positions in a scene space, such as a three-dimensional coordinate system, of objects and/or lighting, as well as a camera viewpoint, and a camera view plane. From that, a rendering engine could take all of those as inputs, and compute each of the pixel color values in each of the frames. In another example, an artist specifies position and movement of an articulated object having some specified texture rather than specifying the color of each pixel representing that articulated object in each frame.

[0078] In a specific example, a rendering engine performs ray tracing wherein a pixel color value is determined by computing which objects lie along a ray traced in the scene space from the camera viewpoint through a point or portion of the camera view plane that corresponds to that pixel. For example, a camera view plane might be represented as a rectangle having a position in the scene space that is divided into a grid corresponding to the pixels of the ultimate image to be generated, and if a ray defined by the camera viewpoint in the scene space and a given pixel in that grid first intersects a solid, opaque, blue object, that given pixel is assigned the color blue. Of course, for modern computer-generated imagery, determining pixel colors – and thereby generating imagery – can be more complicated, as there are lighting issues, reflections, interpolations, and other considerations.

[0079] As illustrated in FIG. 13, a live action capture system 1302 captures a live scene that plays out on a stage 1304. Live action capture system 1302 is described herein in greater detail, but might include computer processing capabilities, image processing capabilities, one

or more processors, program code storage for storing program instructions executable by the one or more processors, as well as user input devices and user output devices, not all of which are shown.

5 **[0080]** In a specific live action capture system, cameras 1306(1) and 1306(2) capture the scene, while in some systems, there might be other sensor(s) 1308 that capture information from the live scene (e.g., infrared cameras, infrared sensors, motion capture (“mo-cap”) detectors, etc.). On stage 1304, there might be human actors, animal actors, inanimate objects, background objects, and possibly an object such as a green screen 1310 that is designed to be captured in a live scene recording in such a way that it is easily overlaid with
10 computer-generated imagery. Stage 1304 might also contain objects that serve as fiducials, such as fiducials 1312(1)-(3), that might be used post-capture to determine where an object was during capture. A live action scene might be illuminated by one or more lights, such as an overhead light 1314.

15 **[0081]** During or following the capture of a live action scene, live action capture system 1302 might output live action footage to a live action footage storage 1320. A live action processing system 1322 might process live action footage to generate data about that live action footage and store that data into a live action metadata storage 1324. Live action processing system 1322 might include computer processing capabilities, image processing capabilities, one or more processors, program code storage for storing program instructions
20 executable by the one or more processors, as well as user input devices and user output devices, not all of which are shown. Live action processing system 1322 might process live action footage to determine boundaries of objects in a frame or multiple frames, determine locations of objects in a live action scene, where a camera was relative to some action, distances between moving objects and fiducials, etc. Where elements have sensors attached
25 to them or are detected, the metadata might include location, color, and intensity of overhead light 1314, as that might be useful in post-processing to match computer-generated lighting on objects that are computer-generated and overlaid on the live action footage. Live action processing system 1322 might operate autonomously, perhaps based on predetermined program instructions, to generate and output the live action metadata upon receiving and
30 inputting the live action footage. The live action footage can be camera-captured data as well as data from other sensors.

[0082] An animation creation system 1330 is another part of visual content generation system 1300. Animation creation system 1330 might include computer processing capabilities, image processing capabilities, one or more processors, program code storage for storing

program instructions executable by the one or more processors, as well as user input devices and user output devices, not all of which are shown. Animation creation system 1330 might be used by animation artists, managers, and others to specify details, perhaps

programmatically and/or interactively, of imagery to be generated. From user input and data
5 from a database or other data source, indicated as a data store 1332, animation creation system 1330 might generate and output data representing objects (e.g., a horse, a human, a ball, a teapot, a cloud, a light source, a texture, etc.) to an object storage 1334, generate and output data representing a scene into a scene description storage 1336, and/or generate and output data representing animation sequences to an animation sequence storage 1338.

- 10 **[0083]** Scene data might indicate locations of objects and other visual elements, values of their parameters, lighting, camera location, camera view plane, and other details that a rendering engine 1350 might use to render CGI imagery. For example, scene data might include the locations of several articulated characters, background objects, lighting, etc. specified in a two-dimensional space, three-dimensional space, or other dimensional space
15 (such as a 2.5-dimensional space, three-quarter dimensions, pseudo-3D spaces, etc.) along with locations of a camera viewpoint and view place from which to render imagery. For example, scene data might indicate that there is to be a red, fuzzy, talking dog in the right half of a video and a stationary tree in the left half of the video, all illuminated by a bright point light source that is above and behind the camera viewpoint. In some cases, the camera
20 viewpoint is not explicit, but can be determined from a viewing frustum. In the case of imagery that is to be rendered to a rectangular view, the frustum would be a truncated pyramid. Other shapes for a rendered view are possible and the camera view plane could be different for different shapes.

- [0084]** Animation creation system 1330 might be interactive, allowing a user to read in
25 animation sequences, scene descriptions, object details, etc. and edit those, possibly returning them to storage to update or replace existing data. As an example, an operator might read in objects from object storage into a baking processor 1342 that would transform those objects into simpler forms and return those to object storage 1334 as new or different objects. For example, an operator might read in an object that has dozens of specified parameters
30 (movable joints, color options, textures, etc.), select some values for those parameters and then save a baked object that is a simplified object with now fixed values for those parameters.

[0085] Rather than requiring user specification of each detail of a scene, data from data store 1332 might be used to drive object presentation. For example, if an artist is creating an

animation of a spaceship passing over the surface of the Earth, instead of manually drawing or specifying a coastline, the artist might specify that animation creation system 1330 is to read data from data store 1332 in a file containing coordinates of Earth coastlines and generate background elements of a scene using that coastline data.

- 5 **[0086]** Animation sequence data might be in the form of time series of data for control points of an object that has attributes that are controllable. For example, an object might be a humanoid character with limbs and joints that are movable in manners similar to typical human movements. An artist can specify an animation sequence at a high level, such as “the left hand moves from location (X1, Y1, Z1) to (X2, Y2, Z2) over time T1 to T2”, at a lower
10 level (e.g., “move the elbow joint 2.5 degrees per frame”) or even at a very high level (e.g., “character A should move, consistent with the laws of physics that are given for this scene, from point P1 to point P2 along a specified path”).

- [0087]** Animation sequences in an animated scene might be specified by what happens in a live action scene. An animation driver generator 1344 might read in live action metadata,
15 such as data representing movements and positions of body parts of a live actor during a live action scene. Animation driver generator 1344 might generate corresponding animation parameters to be stored in animation sequence storage 1338 for use in animating a CGI object. This can be useful where a live action scene of a human actor is captured while wearing mo-cap fiducials (e.g., high-contrast markers outside actor clothing, high-visibility
20 paint on actor skin, face, etc.) and the movement of those fiducials is determined by live action processing system 1322. Animation driver generator 1344 might convert that movement data into specifications of how joints of an articulated CGI character are to move over time.

- [0088]** A rendering engine 1350 can read in animation sequences, scene descriptions, and
25 object details, as well as rendering engine control inputs, such as a resolution selection and a set of rendering parameters. Resolution selection might be useful for an operator to control a trade-off between speed of rendering and clarity of detail, as speed might be more important than clarity for a movie maker to test some interaction or direction, while clarity might be more important than speed for a movie maker to generate data that will be used for final
30 prints of feature films to be distributed. Rendering engine 1350 might include computer processing capabilities, image processing capabilities, one or more processors, program code storage for storing program instructions executable by the one or more processors, as well as user input devices and user output devices, not all of which are shown.

[0089] Visual content generation system 1300 can also include a merging system 1360 that merges live footage with animated content. The live footage might be obtained and input by reading from live action footage storage 1320 to obtain live action footage, by reading from live action metadata storage 1324 to obtain details such as presumed segmentation in captured images segmenting objects in a live action scene from their background (perhaps aided by the fact that green screen 1310 was part of the live action scene), and by obtaining CGI imagery from rendering engine 1350.

[0090] A merging system 1360 might also read data from rulesets for merging/combining storage 1362. A very simple example of a rule in a ruleset might be “obtain a full image including a two-dimensional pixel array from live footage, obtain a full image including a two-dimensional pixel array from rendering engine 1350, and output an image where each pixel is a corresponding pixel from rendering engine 1350 when the corresponding pixel in the live footage is a specific color of green, otherwise output a pixel value from the corresponding pixel in the live footage.”

[0091] Merging system 1360 might include computer processing capabilities, image processing capabilities, one or more processors, program code storage for storing program instructions executable by the one or more processors, as well as user input devices and user output devices, not all of which are shown. Merging system 1360 might operate autonomously, following programming instructions, or might have a user interface or programmatic interface over which an operator can control a merging process. In some embodiments, an operator can specify parameter values to use in a merging process and/or might specify specific tweaks to be made to an output of merging system 1360, such as modifying boundaries of segmented objects, inserting blurs to smooth out imperfections, or adding other effects. Based on its inputs, merging system 1360 can output an image to be stored in a static image storage 1370 and/or a sequence of images in the form of video to be stored in an animated/combined video storage 1372.

[0092] Thus, as described, visual content generation system 1300 can be used to generate video that combines live action with computer-generated animation using various components and tools, some of which are described in more detail herein. While visual content generation system 1300 might be useful for such combinations, with suitable settings, it can be used for outputting entirely live action footage or entirely CGI sequences. The code may also be provided and/or carried by a transitory computer readable medium, e.g., a transmission medium such as in the form of a signal transmitted over a network.

[0093] According to one embodiment, the techniques described herein are implemented by one or more generalized computing systems programmed to perform the techniques pursuant to program instructions in firmware, memory, other storage, or a combination. Special-purpose computing devices may be used, such as desktop computer systems, portable
5 computer systems, handheld devices, networking devices or any other device that incorporates hard-wired and/or program logic to implement the techniques.

[0094] FIG. 14 is a block diagram that illustrates an example computer system 1400 upon which the computer systems of the systems described herein and/or visual content generation system 1300 (see FIG. 13) may be implemented. Computer system 1400 includes a bus 1402
10 or other communication mechanism for communicating information, and a processor 1404 coupled with bus 1402 for processing information. Processor 1404 may be, for example, a general-purpose microprocessor.

[0095] Computer system 1400 also includes a main memory 1406, such as a random-access memory (RAM) or other dynamic storage device, coupled to bus 1402 for storing information
15 and instructions to be executed by processor 1404. Main memory 1406 may also be used for storing temporary variables or other intermediate information during execution of instructions to be executed by processor 1404. Such instructions, when stored in non-transitory storage media accessible to processor 1404, render computer system 1400 into a special-purpose machine that is customized to perform the operations specified in the instructions.

20 [0096] Computer system 1400 further includes a read only memory (ROM) 1408 or other static storage device coupled to bus 1402 for storing static information and instructions for processor 1404. A storage device 1410, such as a magnetic disk or optical disk, is provided and coupled to bus 1402 for storing information and instructions.

[0097] Computer system 1400 may be coupled via bus 1402 to a display 1412, such as a
25 computer monitor, for displaying information to a computer user. An input device 1414, including alphanumeric and other keys, is coupled to bus 1402 for communicating information and command selections to processor 1404. Another type of user input device is a cursor control 1416, such as a mouse, a trackball, or cursor direction keys for communicating direction information and command selections to processor 1404 and for
30 controlling cursor movement on display 1412. This input device typically has two degrees of freedom in two axes, a first axis (e.g., x) and a second axis (e.g., y), that allows the device to specify positions in a plane.

[0098] Computer system 1400 may implement the techniques described herein using customized hard-wired logic, one or more ASICs or FPGAs, firmware and/or program logic

which in combination with the computer system causes or programs computer system 1400 to be a special-purpose machine. According to one embodiment, the techniques herein are performed by computer system 1400 in response to processor 1404 executing one or more sequences of one or more instructions contained in main memory 1406. Such instructions
5 may be read into main memory 1406 from another storage medium, such as storage device 1410. Execution of the sequences of instructions contained in main memory 1406 causes processor 1404 to perform the process steps described herein. In alternative embodiments, hard-wired circuitry may be used in place of or in combination with software instructions.

[0099] The term “storage media” as used herein refers to any non-transitory media that store
10 data and/or instructions that cause a machine to operation in a specific fashion. Such storage media may include non-volatile media and/or volatile media. Non-volatile media includes, for example, optical or magnetic disks, such as storage device 1410. Volatile media includes dynamic memory, such as main memory 1406. Common forms of storage media include, for example, a floppy disk, a flexible disk, hard disk, solid state drive, magnetic tape, or any
15 other magnetic data storage medium, a CD-ROM, any other optical data storage medium, any physical medium with patterns of holes, a RAM, a PROM, an EPROM, a FLASH-EPROM, NVRAM, any other memory chip or cartridge.

[0100] Storage media is distinct from but may be used in conjunction with transmission media. Transmission media participates in transferring information between storage media.
20 For example, transmission media includes coaxial cables, copper wire, and fiber optics, including the wires that include bus 1402. Transmission media can also take the form of acoustic or light waves, such as those generated during radio-wave and infra-red data communications.

[0101] Various forms of media may be involved in carrying one or more sequences of one or
25 more instructions to processor 1404 for execution. For example, the instructions may initially be carried on a magnetic disk or solid-state drive of a remote computer. The remote computer can load the instructions into its dynamic memory and send the instructions over a network connection. A modem or network interface local to computer system 1400 can receive the data. Bus 1402 carries the data to main memory 1406, from which processor
30 1404 retrieves and executes the instructions. The instructions received by main memory 1406 may optionally be stored on storage device 1410 either before or after execution by processor 1404.

[0102] Computer system 1400 also includes a communication interface 1418 coupled to bus 1402. Communication interface 1418 provides a two-way data communication coupling to a

network link 1420 that is connected to a local network 1422. For example, communication interface 1418 may be a network card, a modem, a cable modem, or a satellite modem to provide a data communication connection to a corresponding type of telephone line or communications line. Wireless links may also be implemented. In any such implementation, communication interface 1418 sends and receives electrical, electromagnetic, or optical signals that carry digital data streams representing various types of information.

[0103] Network link 1420 typically provides data communication through one or more networks to other data devices. For example, network link 1420 may provide a connection through local network 1422 to a host computer 1424 or to data equipment operated by an Internet Service Provider (ISP) 1426. ISP 1426 in turn provides data communication services through the world-wide packet data communication network now commonly referred to as the “Internet” 1428. Local network 1422 and Internet 1428 both use electrical, electromagnetic, or optical signals that carry digital data streams. The signals through the various networks and the signals on network link 1420 and through communication interface 1418, which carry the digital data to and from computer system 1400, are example forms of transmission media.

[0104] Computer system 1400 can send messages and receive data, including program code, through the network(s), network link 1420, and communication interface 1418. In the Internet example, a server 1430 might transmit a requested code for an application program through the Internet 1428, ISP 1426, local network 1422, and communication interface 1418. The received code may be executed by processor 1404 as it is received, and/or stored in storage device 1410, or other non-volatile storage for later execution.

[0105] Operations of processes described herein can be performed in any suitable order unless otherwise indicated herein or otherwise clearly contradicted by context. Processes described herein (or variations and/or combinations thereof) may be performed under the control of one or more computer systems configured with executable instructions and may be implemented as code (e.g., executable instructions, one or more computer programs or one or more applications) executing collectively on one or more processors, by hardware or combinations thereof. The code may be stored on a computer-readable storage medium, for example, in the form of a computer program comprising a plurality of instructions executable by one or more processors. The computer-readable storage medium may be non-transitory. The code may also be provided carried by a transitory computer readable medium e.g., a transmission medium such as in the form of a signal transmitted over a network.

[0106] Conjunctive language, such as phrases of the form “at least one of A, B, and C,” or “at least one of A, B and C,” unless specifically stated otherwise or otherwise clearly contradicted by context, is otherwise understood with the context as used in general to present that an item, term, etc., may be either A or B or C, or any nonempty subset of the set of A and B and C. For instance, in the illustrative example of a set having three members, the conjunctive phrases “at least one of A, B, and C” and “at least one of A, B and C” refer to any of the following sets: {A}, {B}, {C}, {A, B}, {A, C}, {B, C}, {A, B, C}. Thus, such conjunctive language is not generally intended to imply that certain embodiments require at least one of A, at least one of B and at least one of C each to be present.

5 [0107] Directional references e.g., upper, lower, inner, outer, upward, downward, left, right, front, back, top, bottom, above, below, vertical, and horizontal are used for identification purposes to aid the reader’s understanding of the claimed subject matter and are not intended to create limitations, particularly as to a position or orientation of simulated objects, or use of the surface computation method. The word “comprising” does not exclude other elements or steps, and the indefinite article “a” or “an” does not exclude a plurality. Unless otherwise noted in the claims, stated values shall be interpreted as illustrative only and shall not be taken to be limiting.

15 [0108] The use of examples, or exemplary language (e.g., “such as”) provided herein, is intended merely to better illuminate embodiments of the invention and does not pose a limitation on the scope of the invention unless otherwise claimed. No language in the specification should be construed as indicating any non-claimed element as essential to the practice of the invention.

20 [0109] In the foregoing specification, embodiments of the invention have been described with reference to numerous specific details that may vary from implementation to implementation. The specification and drawings are, accordingly, to be regarded in an illustrative rather than a restrictive sense. The sole and exclusive indicator of the scope of the invention, and what is intended by the applicants to be the scope of the invention, is the literal and equivalent scope of the set of claims that issue from this application, in the specific form in which such claims issue, including any subsequent correction.

25 [0110] Further embodiments can be envisioned to one of ordinary skill in the art after reading this disclosure. In other embodiments, combinations or sub-combinations of the above-disclosed invention can be advantageously made. The example arrangements of components are shown for purposes of illustration and combinations, additions, re-arrangements, and the like are contemplated in alternative embodiments of the present invention. Thus, while the

invention has been described with respect to exemplary embodiments, one skilled in the art will recognize that numerous modifications are possible.

5 [0111] For example, the processes described herein may be implemented using hardware components, software components, and/or any combination thereof. The specification and drawings are, accordingly, to be regarded in an illustrative rather than a restrictive sense. It will, however, be evident that various modifications and changes may be made thereunto without departing from the broader spirit and scope of the invention as set forth in the claims and that the invention is intended to cover all modifications and equivalents within the scope of the following claims.

10 [0112] All references, including publications, patent applications, and patents, cited herein are hereby incorporated by reference to the same extent as if each reference were individually and specifically indicated to be incorporated by reference and were set forth in its entirety herein.

15 [0113] In this specification where reference has been made to patent specifications, other external documents, or other sources of information, this is generally for the purpose of providing a context for discussing the features of the invention. Unless specifically stated otherwise, reference to such external documents or such sources of information is not to be construed as an admission that such documents or such sources of information, in any jurisdiction, are prior art or form part of the common general knowledge in the art.

CLAIMSWHAT IS CLAIMED IS:

1. A computer-implemented method for modeling surfaces of simulated objects, the computer-implemented method comprising:
under the control of one or more computer systems configured with executable instructions:
receiving positions and sizes for each particle of a plurality of particles comprising a first object;
identifying a subset of the plurality of particles whose positions are proximate to a first implied surface of the first object;
deforming particles of the subset of the plurality of particles to form a set of deformed particles by generating a covariance matrix for each particle of at least some particles of the plurality of particles, wherein the covariance matrix corresponds to a diagonal stretch matrix operated on by a non-linear function and wherein the diagonal stretch matrix is a diagonal stretch matrix of neighboring particles for at least some particles of the plurality of particles; and
computing a new shape or position of the first implied surface of the first object based on the set of deformed particles.
2. The computer-implemented method of claim 1, wherein the non-linear function comprises a square root.
3. The computer-implemented method of claim 1, wherein at least some of the particles of the plurality of particles are spherical.
4. The computer-implemented method of claim 1, wherein at least some of the deformed particles are ellipsoidal.
5. The computer-implemented method of claim 1, wherein deforming the particles of the subset of the plurality of particles involves elongating axes of the particles of the subset that are parallel to the first implied surface of the first object.
6. The computer-implemented method of claim 1, wherein the first object is solid.

7. The computer-implemented method of claim 1, wherein the first object is semi-solid or fluid.

8. The computer-implemented method of claim 1, wherein a shape or position of the first implied surface of the first object changes based on an interaction of the first object with a second surface of a second object.

9. The computer-implemented method of claim 1, further comprising:
receiving masses for each particle of the plurality of particles comprising the first object;

computing a computed pressure, density, or force acting on at least one particle of the plurality of particles comprising the first object; and

changing the positions of the at least one particle based on the computed pressure, density, or force.

10. The computer-implemented method of claim 1, wherein the positions and sizes of the particles comprising the first object are received from a simulation.

11. The computer-implemented method of claim 1, wherein the positions and sizes of the particles comprising the first object are received from a scanning device.

12. The computer-implemented method of claim 1, wherein identifying the subset of the plurality of particles whose positions are proximate to the first implied surface of the first object involves computing a weighted mean and the covariance matrix of each particle of the plurality of particles based on a user-defined search radius.

13. The computer-implemented method of claim 1, wherein computing the new shape or position of the first implied surface of the first object based on the deformed particles involves computing an isosurface of a scalar field.

14. The computer-implemented method of claim 1, further comprising generating a visual representation of the first implied surface of the first object.

15. The computer-implemented method of claim 8, further comprising generating a visual representation of the second surface of the second object.

16. A computer system for modeling surfaces of simulated objects, the system comprising:

at least one processor; and

a computer-readable medium storing instructions, which when executed by the at least one processor, causes the system to carry out the method of claim 1.

17. A non-transitory computer-readable storage medium storing instructions, which when executed by at least one processor of a computer system, causes the computer system to carry out the method of claim 1.