

VARIABLE-SPEED OBJECT MOTION IN 3D VISUALIZATIONS OF DISCRETE-EVENT CONSTRUCTION SIMULATION MODELS

SUBMITTED: April 2006

REVISED: March 2007

PUBLISHED: April 2007 at <http://itcon.org/2007/20/>

EDITOR: C. Eastman

*Vineet R. Kamat, Assistant Professor of Civil and Environmental Engineering
University of Michigan, USA
email: vkamat@umich.edu <http://pathfinder.engin.umich.edu>*

*Julio C. Martinez, Associate Professor of Civil and Environmental Engineering
Purdue University, USA
email: julio@purdue.edu <http://www.ezstroke.com>*

SUMMARY: *This paper presents research that addresses the problem of describing the accurate, variable-speed motion of simulation objects on realistically-shaped trajectories (i.e. paths) in animations of discrete-event construction simulation models. The research puts in place techniques that simulation modellers can use to instruct virtual simulation objects to follow any arbitrarily-shaped velocity profiles while adhering to fixed motion completion times while they traverse along any defined path trajectories. A computation scheme that allows simulation models to define the general shapes of relevant velocity profiles and then heuristically scales those profiles to accommodate communicated activity instance durations is presented. While allowing animated simulation objects to be moved with any arbitrarily shaped velocity profiles, this technique ensures that an object's spatial and temporal control rests entirely with the underlying simulation models.*

KEYWORDS: *animation, construction operations, discrete-event simulation, verification, validation.*

1. INTRODUCTION

Visualization of modelled operations can be of significant help in the verification and validation of discrete-event simulation (DES) models (Law and Kelton 2000). This is especially true in construction where typical decision makers are experts in their domain but are not generally proficient in simulation itself. Visualization can also provide decision makers with valuable insight into subtleties of planned construction operations that are otherwise non-quantifiable and non-presentable.

The authors' recent research efforts have focused on designing automatic, simulation-driven methods to visualize modelled construction processes and any evolving products in smooth, continuous, dynamic 3D virtual worlds. Methods have been designed to describe animated 3D worlds that show how simulated processes are carried out, using simple parametric text statements and references to 3D CAD models (Kamat and Martinez 2003a, 2004, 2005). This simple text animation description language, called VITASCOPE, is meant to be written out by end-user programmable tools such as DES systems and allows a computer to create a dynamic 3D virtual world that shows people, machines, and/or materials interacting as they perform the modelled processes.

1.1 Importance of the Research

Process simulation-driven 3D virtual worlds are spatially and temporally faithful to the underlying DES models that author the visualizations. Notwithstanding, the 3D visual representations (i.e. visualizations) of several modelled processes digress in time and space accuracy from the corresponding real-world operations due to the inherent characteristics of DES.

In DES, the state of a running model changes only at discrete, but possibly random sets of simulated time points (Schriber and Brunner 2001). These time points are typically the start or end of the model's activities, and it is only then that a running DES model can communicate with other processes, or perform other actions such as output to an animation trace. A DES model is only concerned about the time instants at which instances of modelled activities begin or end, and chooses to ignore everything (e.g. rate of activity performance) that happens in between. The information that a running DES model can communicate to external 3D animation

methods is thus limited to the start times and durations of all activity instances that occur in any simulation run. The animation methods must then use these limited pieces of discrete information to generate a smooth, continuous, dynamic 3D virtual world representation that depicts the modelled activities being performed.

In any modelled activity that involves motion of simulation entities (e.g. a hauling truck, an airplane taking off), the only kinetic property that can be computed from the pieces of communicated activity instance information (i.e. start time and duration) is the average speed of the simulation entity (e.g. truck, airplane) in that particular activity instance (e.g. haul dirt, take off). This is precisely the computation existing 3D animation methods perform in describing the smooth, continuous motion of simulation entities in virtual worlds.

In particular, the motion of a simulation entity is depicted by transforming (i.e. moving) the pertinent instantiated CAD model of the entity at the computed average velocity. The simulation entity-representing CAD model is smoothly and linearly interpolated on a 3D trajectory that represents that entity's motion path (e.g. haul road, run-way) in the communicated activity instance. Thus, a virtual airplane taking off appears to travel down a runway at a constant (average) speed for the activity instance's sampled time duration (i.e. runway occupancy time). Similarly, a loaded truck hauling dirt on a virtual earthmoving jobsite travels at a constant speed for the entire duration of the haul regardless of the grades on the 3D motion trajectory (virtual haul road) it travels on.

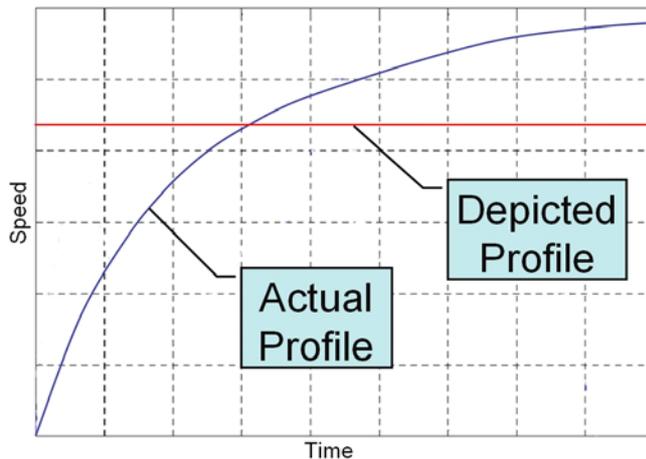
The temporal and spatial accuracy of simulation model-generated dynamic 3D virtual worlds is commensurate with the detail of the communicated information from which the visualizations are created. The existing animation scheme is faithful to DES in that animated activity instances inside a 3D virtual world begin and end at the exact time instants dictated by the visualization-authoring simulation models, with smooth, continuous, "constant-speed" intermediate motion of the involved virtual entities (i.e. simulation objects). Such visualization of simulation objects performing the modelled (and communicated) tasks at constant speed is often sufficient to verify and validate several DES models (Kamat and Martinez 2003b).

However, the depicted constant velocity profiles of moving virtual simulation objects are not an accurate representation of reality. For instance, an airplane taking off on a runway obviously does not travel at a constant velocity. Instead, as Fig. 1(a) presents, the airplane continuously accelerates as it races down a runway and takes-off. Similarly, as Fig. 1(b) presents, the typical velocity profile of a truck that hauls dirt is a function of several disparate factors such as engine power, load being hauled, and the rolling resistance and grade of haul road segments. Due to the limited operational information available from underlying DES models, existing methods of animating simulated processes are, however, unable to adopt such realistic velocity profiles in describing the motion of virtual simulation entities. Instead, commensurate with the available pieces of information, simulation objects are assumed to be moving at constant, average speeds with straight-line velocity profiles such as those superimposed on Figs. 1(a) and 1(b).

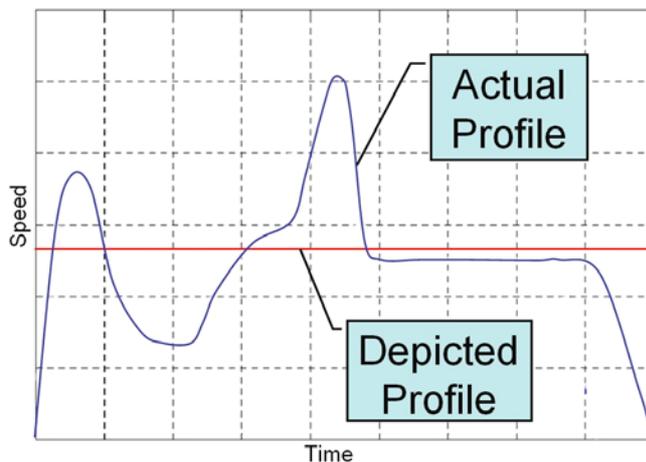
This assumption (and portrayal) of constant-speed motion can often hinder the validation of modelled processes in cases where the relative segmental speeds of moving simulation entities and/or their acceleration/deceleration influence their evolution and inter-object interactions in a modelled and visualized system (e.g. airport, earthmoving jobsite). In addition, such constant-speed visualization of modelled processes can frequently fail to elicit credibility for simulation models, especially from domain experts and decision makers who are not familiar with the mechanics of DES and/or are skeptic about simulation analyses beforehand.

2. RELATED WORK

Several special-purpose animation technologies exist to visualize simulated operations in other domains. Many proprietary simulation tools used in manufacturing display 3D animations of simulated operations. Examples of such tools include Delmia's Quest and AutoSimulations' AutoMod. These tools have modelling elements and built-in 3D templates of common manufacturing environments and equipment that enable users to simulate and visualize common manufacturing operations (Rohrer 2000, Donald 1998). An example of a non-proprietary tool designed to animate simulated manufacturing operations is CASUS (Luckas and Broll 1997). In all these systems, animation of simulated operations is achieved by communicating elemental resource motions (i.e. geometric transformations) to visual elements, and support constant speed resource motion as described above.



(a) Airplane on Take-Off Roll



(b) Dumptruck on Haul Road

FIG. 1: Typical Velocity Profiles of Real Objects.

Several researchers in the past have created and experimented with operations-level visualizations that depict specific scenarios of short duration. Examples of such work include (e.g. Li et al. 2003, Tseng et al. 2000, Fukuchi et al. 1999, Tsay et al. 1996). Many of these works were achieved with software tools and libraries that are available for creating such scenarios with a significantly large ratio of ‘time invested to develop the specific case’ to ‘actual duration of the visualization experience’ (e.g. WorldToolKit, 3dsmax, Bentley Dynamic Animator, etc). Though not related directly to the automated 3D visualization of simulated construction operations discussed in this paper, these works manifest the same critical issues of having to communicate with resources using a primitive geometric transformations based vocabulary. In the case of the Bentley Dynamic Animator, for instance, each elemental motion that a resource (e.g. crane) must perform in animating a construction task must be manually defined and stored for later replay (Bentley 2003). The resulting motion of such objects then evolves with constant speed between two defined time points.

3. DESCRIPTION OF VARIABLE SPEED OBJECT MOTION

The only temporal information about an activity-instance that a running DES model can communicate to external processes such as 3D animation methods is that instance’s start time and its duration. Due to inherent modelling features, a DES model can provide no information on the rate at which the task(s) in any activity instance are performed. This is unlike Continuous Simulation, where the state of a model (and hence the rate of activity performance) is continuously monitored at every time instant using differential equations of motion (Law and Kelton 2000). The performance rate of an activity, however, provides precisely the temporal information needed to describe the velocity profile(s) of simulation object(s) that move (i.e. travel) while performing the task(s) in a particular activity instance.

3.1 Hypothetical Hybrid Animation Approach

In order to describe variable-speed motion of virtual DES objects then, the first possible technique that was explored was to describe a parallel continuous simulation system that would tightly integrate with the methods of animating DES models in 3D. In particular, the authors considered the possibility of externally formulating an in-context simulation object's pertinent kinetic properties (e.g. nominal acceleration/deceleration, maximum permitted velocity) and using that information along with each communicated activity instance (i.e. start time and duration). During visualization, the integrated continuous simulation mechanism would use the formulated kinetic properties to compute the involved simulation object's velocity profile (i.e. its temporal evolution) beginning at the indicated activity instance start time.

This hybrid animation approach is, however, impossible to achieve in a DES framework. In particular, such a strategy would only work if a DES model communicated an activity instance's start time and enforced no restrictions on when it ended. That information (i.e. activity instance end time) could then be determined in real-time as the involved simulation object's temporal evolution was computed continuously during visualization. A DES model must, however, explicitly enforce a communicated activity instance's end time (i.e. its duration). This is obvious because the start of instance(s) of other simulation model activities (often involving the object that is in context in the current activity instance) is explicitly tied to the completion of the current communicated activity instance. Once instantiated, any instance of such a successive activity will attempt to exclusively manipulate the in-context virtual simulation object to visually describe the performance of that latter communicated task.

The hypothetical continuous simulation system computing the simulation object's temporal evolution in the instance of the previous activity would, however, be unable to guarantee the completion of its motion at the exact precise instant at which the successive activity starts. Stated differently, it is mathematically impossible to externally formulate kinetic object properties and compute a unique, valid, continuous velocity profile using a set of differential equations if the motion start and end times (i.e. the lower and upper bound of the integration interval) and the distance traversed (i.e. the area under the resulting curve) are both explicitly enforced.

This is however the case in DES. As such, any computation scheme (for visualizing simulated processes) that wrests the temporal and spatial control of simulation objects away from the underlying DES models cannot portray the modelled operations correctly in dynamic 3D virtual worlds. The description of any arbitrary velocity profiles to be applied to mobile simulation objects during animation must thus be sought from the DES models that author visualizations. However, a DES model obviously does not encapsulate any such information (e.g. an object's kinetic properties) simply because the rate of performing any modelled activity is generally irrelevant to the model from the simulation analysis perspective.

3.2 Time-Based Scaling of Velocity Profile Shapes

In order to visually describe variable speed motion of animated simulation objects then, the authors devised a unique computation scheme that prudently shares the responsibility of describing a moving simulation object's arbitrary velocity profile between the underlying DES model and the 3D animation methods. In particular, the general shape of the velocity profile to be applied to a moving object is sought from the visualization-authoring DES model. The shape of this general profile can be explicitly defined (and input into a DES model) by a modeller, or it can be the result of computations performed within a running model.

Then, at each communicated activity instance, the animation methods heuristically scale (up or down) the previously defined velocity profile in such a way that the mobile simulation object to which it is applied traverses an indicated motion path in a time interval that is exactly equal to the communicated activity instance duration. This flexible technique allows for two vital things. 1) Any arbitrarily shaped velocity profile resulting from any DES model-defined or externally performed computation can be explicitly applied to a moving simulation object, and 2) the temporal and spatial control of all simulation objects remains entirely with the underlying simulation model since the animation methods merely scale a defined velocity profile to fit the duration of a communicated activity instance.

For any defined motion path trajectory, a simulation model defines the shape of a desired object velocity profile by specifying an arbitrary number of velocity-distance pairs. The specified velocity values can span any positive numerical range and the corresponding indicated distances are the percentile (0 to 100) arc lengths along the path. The shape of the profile is deduced by plotting the path's percentile arc distance on the abscissa and the

corresponding velocity values on the ordinate. This is graphically presented in Fig. 2. No limitations are placed on this definition except that the distance value in the last specified velocity-distance pair must equal 100 percent (i.e. the current path's total arc length).

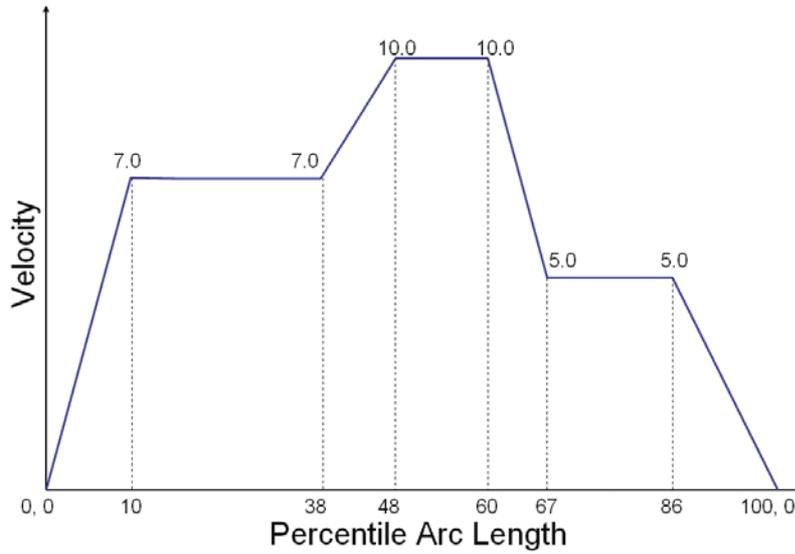


FIG. 2: Definition of Velocity Profile Shape.

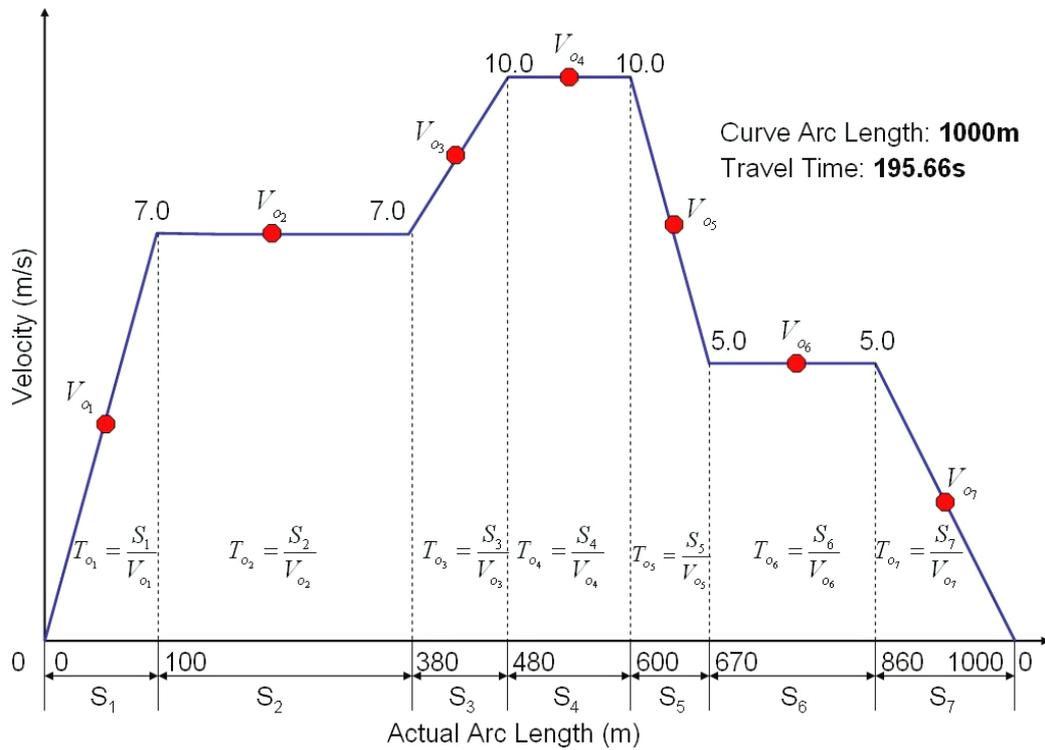
As a first pre-processing step, the defined velocity versus percentile distance profile is converted to a velocity versus actual distance curve. As Fig. 3(a) depicts, this is accomplished by simply replacing the percentile arc lengths on the abscissa by the corresponding actual arc distances for the current path. The specified velocity values are left unchanged. The total time T_o , required to traverse this converted profile in its unmodified form can then be given by the following equation:

$$\begin{aligned}
 T_o &= \frac{S_1}{V_{o_1}} + \frac{S_2}{V_{o_2}} + \frac{S_3}{V_{o_3}} + \dots + \frac{S_n}{V_{o_n}} \\
 &= T_{o_1} + T_{o_2} + T_{o_3} + \dots + T_{o_n}
 \end{aligned} \tag{1}$$

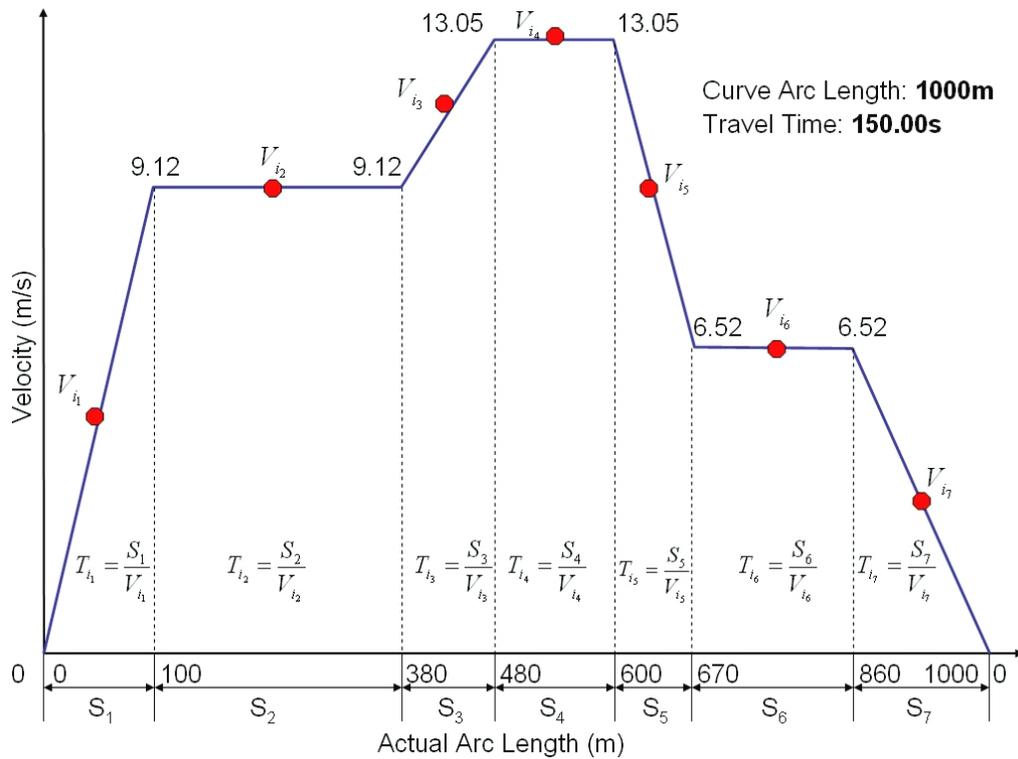
As Figs. 3(a) indicates, $V_{o_1}, V_{o_2}, V_{o_3}, \dots, V_{o_n}$ are the average velocities at which the respective path segments

$S_1, S_2, S_3, \dots, S_n$ are traversed. In the case of this converted, unmodified curve, the total original travel time works out to be 195.66 seconds. Now, for any communicated activity instance of duration T_i (say 150 seconds), the described velocity profile is segmentally scaled up or down such that a simulation object traveling that path with the resulting modified (i.e. scaled) velocities reaches the end of the path in exactly T_i time units. This is graphically depicted in Fig. 3(b). This scaling procedure can be described as:

$$\begin{aligned}
 \frac{(T_{i_1} + T_{i_2} + T_{i_3} + \dots + T_{i_n})}{T_i} &= \frac{(T_{o_1} + T_{o_2} + T_{o_3} + \dots + T_{o_n})}{T_o} \\
 \therefore (T_{i_1} + T_{i_2} + T_{i_3} + \dots + T_{i_n}) &= \frac{T_i}{T_o} (T_{o_1} + T_{o_2} + T_{o_3} + \dots + T_{o_n}) \\
 &= T_{o_1} \left(\frac{T_i}{T_o} \right) + T_{o_2} \left(\frac{T_i}{T_o} \right) + T_{o_3} \left(\frac{T_i}{T_o} \right) + \dots + T_{o_n} \left(\frac{T_i}{T_o} \right)
 \end{aligned} \tag{2}$$



(a) Original Velocity Profile



(b) Adjusted Velocity Profile

FIG. 3: Derivation of Time-Scaled Velocity Profiles.

The underlying assumption that is made in segmentally scaling the original converted velocity profile (Fig. 3(a)) to accommodate the currently specified activity instance duration (Fig. 3(b)) is that:

$$T_{i_1} = T_{o_1} \left(\frac{T_i}{T_o} \right), T_{i_2} = T_{o_2} \left(\frac{T_i}{T_o} \right), \dots, T_{i_n} = T_{o_n} \left(\frac{T_i}{T_o} \right) \quad (3)$$

With this assumption, the equation that distributes the communicated activity instance duration over the different velocity segments can be written as:

$$\begin{aligned} T_i &= T_{i_1} + T_{i_2} + T_{i_3} + \dots + T_{i_n} \\ &= \frac{S_1}{V_{i_1}} + \frac{S_2}{V_{i_2}} + \frac{S_3}{V_{i_3}} + \dots + \frac{S_n}{V_{i_n}} \end{aligned} \quad (4)$$

The fact that:

$$T_{i_1} = \frac{S_1}{V_{i_1}}, T_{i_2} = \frac{S_2}{V_{i_2}}, \dots, T_{i_n} = \frac{S_n}{V_{i_n}} \quad (5)$$

provides the basis for computing the average segmental velocities $V_{i_1}, V_{i_2}, V_{i_3}, \dots, V_{i_n}$, from which the scaled, modified, activity instance-specific velocity profile can be constructed (Fig. 3(b)). A simulation object that follows this modified velocity profile is guaranteed to traverse the path in the exact communicated activity instance duration T_i .

The following section presents the implementation of the described computation scheme. The implementation is a powerful tool that allows engineers to accurately describe the 3D motion of virtual DES objects on realistic-looking, smoothly-curved motion trajectories.

4. SOFTWARE IMPLEMENTATION

The algorithms designed to define flexible, smooth, curved motion path trajectories and then move virtual simulation objects on those trajectories with desired velocity profiles are implemented as a software tool named PathFinder. This tool has been designed as an extension to the VITASCOPE visualization system. VITASCOPE is a user-extensible 3D animation language designed specifically for visualizing simulated processes (particularly construction operations) in smooth, continuous, dynamic 3D virtual worlds.

In particular, the PathFinder add-on defines parametric text animation language statements that allow the definition and manipulation of smooth, curved motion trajectories of arbitrarily complex shapes. PathFinder also implements statements that allow simulation models to 1) specify a default per-path arbitrary velocity profile shape, and 2) override (if necessary) the default velocity profile shape for any communicated instance of simulation object motion.

In addition, PathFinder also implements terrain-following techniques that help orient moving objects correctly in 3D, and the velocity profile scaling methods that heuristically resize the current velocity profile shape to accommodate the duration of a communicated activity instance. PathFinder thus implements the techniques that allow users to accurately describe the 3D motion of virtual discrete-event simulation objects on realistic-looking, smoothly-curved motion trajectories.

Fig. 4 presents an animation trace with statements that define a smooth motion trajectory and then move an object over it with an indicated velocity profile. Motion trajectories are defined in 2D planar resolutions by specifying a series of 2D (i.e. x and z coordinate only) control points and then manipulating (if needed) the shape of the resultant spline curve by adjusting the tension, continuity, and/or bias of one or more knots (i.e. control points). A default velocity profile (if any) can also be specified as part of a path's definition. In the absence of an explicitly indicated velocity curve for a path, PathFinder assumes a default, constant-speed profile for objects that traverse that path. The defined 2D trajectory is then superimposed on the 3D terrain model that represents the underlying simulated system's landscape.

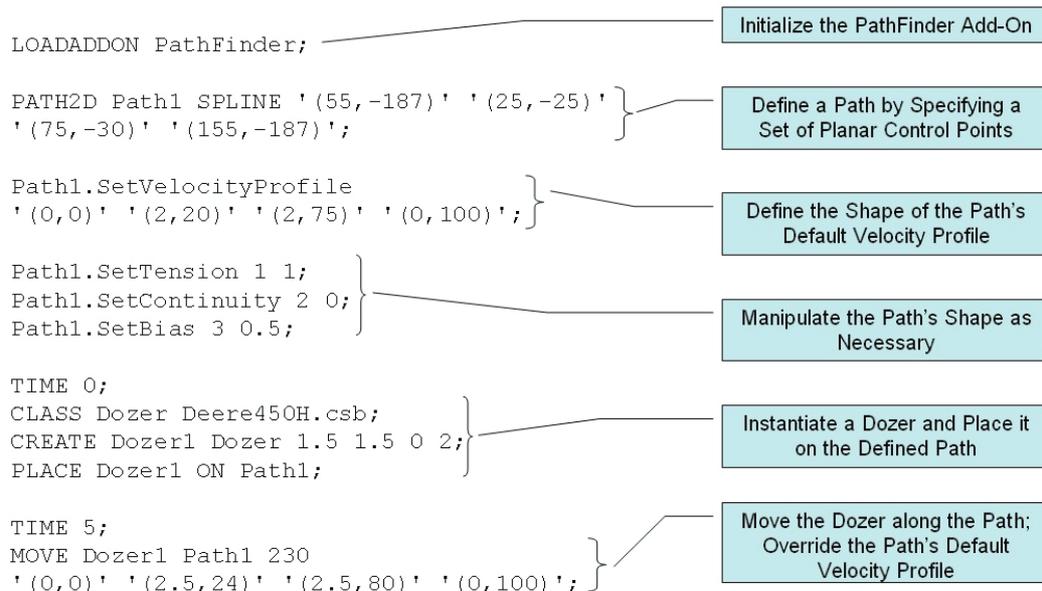


FIG. 4: Moving Simulation Objects on Defined Motion Paths.

Simulation objects can obviously override a path's default profile during an activity instance (i.e. motion) communication. This can, for instance, allow simulation models to specify a unique velocity profile that is a function of the properties (e.g. engine power, loaded mass) of the in-context simulation object (e.g. dumptruck) in a communicated instance of an activity (e.g. haul dirt). When a simulation object is instructed to move on a particular path in the specified simulation time units, the current (path default or object overridden) velocity profile is appropriately scaled such that the time to traverse the path with that profile is equal to the communicated activity instance duration. The simulation object then traverses the superimposed 3D path trajectory while following the terrain surface as closely as possible.

Fig. 5 presents a strip of animation snapshots taken during the visualization of the motion statement from the animation trace in Fig. 4. The 3D projection of the described 2D trajectory passes through highly uneven terrain. However, as the snapshots depict, PathFinder's terrain-following algorithms ensure that the object (i.e. dozer) is correctly oriented on the terrain as it travels the path with the scaled velocity profile. The snapshots presented are not successive computer frames observed during visualization. The discretely captured frames are displayed in a filmstrip format merely to depict a sense of motion. The smooth motion of the dozer during visualization and its non-constant velocity cannot be fully captured in static snapshots. Only the animation can convey that information.

4.1 Describing Object Motion with Terrain-Following

Defined and modified 2D motion path trajectories are converted to their 3D representation by projecting them onto the 3D terrain model that describes the simulated system's landscape. At any instant, a moving object's current, adjusted velocity profile dictates the downstream distance on a path at which it is currently located. Given the current downstream distance, the object's planar position (i.e. x and z coordinate) is determined by parametrically retrieving (on the defined 2D path) the point which corresponds to that arc length (i.e. downstream distance). Then, the object's current yaw is calculated by simply computing the tangent (in the 2D plane) to the curved path trajectory at that determined position (i.e. x and z coordinate).

To determine the 3D position of the object inside the virtual world, the computed planar position point is projected on the 3D terrain model of the simulated system's landscape i.e. on the superimposed 3D path. This is accomplished by retrieving the elevation (i.e. height) of the terrain model at that horizontal plane location i.e. the y coordinate of the terrain point that corresponds to the x and z values of the determined 2D position is retrieved. This describes the travelling object's current 3D position at that time instant. Since this computation is performed dynamically at visualization run-time, the 2D path's 3D projection always drapes the terrain surface even if its shape deforms during animation i.e. the procedure always retrieves the current terrain elevations (heights) below 2D position points.

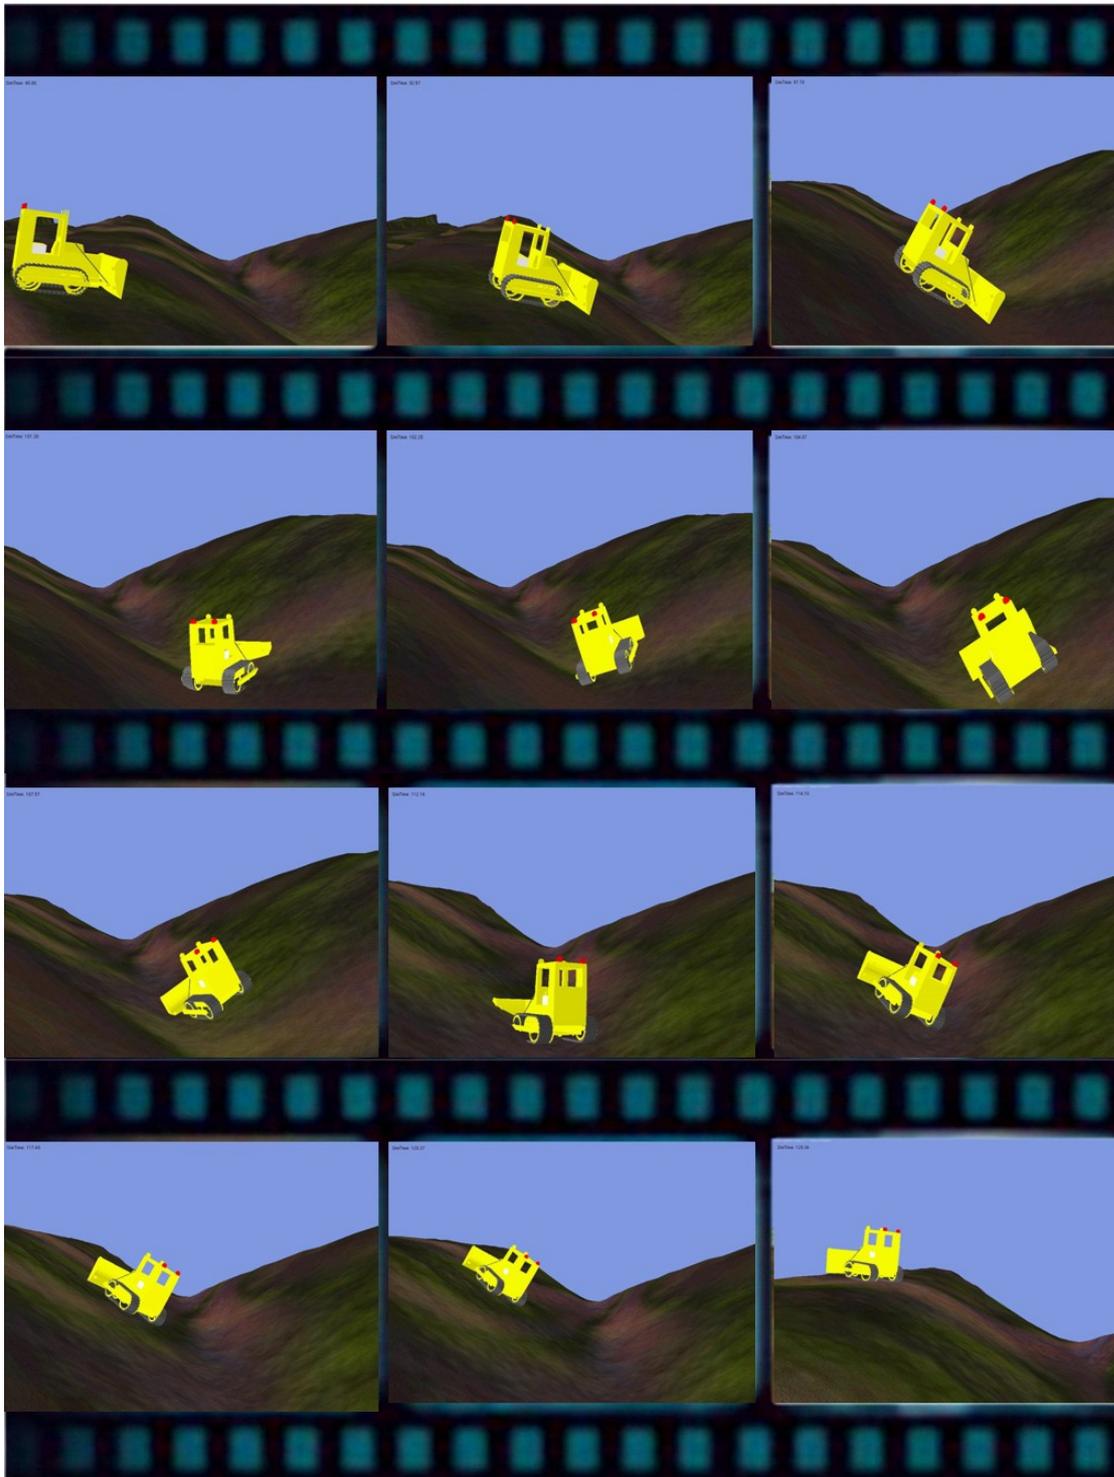


FIG. 5: Animation Snapshots of Terrain-Followed Object Motion.

Given the 3D position and yaw of the moving object on the terrain surface, the goal of terrain-following is to now orient that object correctly (i.e. compute the pitch and roll) such that its virtual contact points (e.g. a truck's tires) all touch the terrain's surface as closely as possible. To keep moving objects correctly oriented on a terrain's surface, it is necessary to find the locations where that object's contact points touch the virtual terrain. This can be done using geometric collision detection techniques. However, this is inefficient because general collision detection inherently involves more complex, CPU-intensive computations than merely computing an object's terrain contact points (Barrus and Waters 1997).

To enforce terrain-following in a moving simulation object, a generalized technique is adopted that systematically computes that object's pitch and roll by projecting its contact points on the terrain model in a manner similar to that used in computing the 3D position. In particular, the 2D positions corresponding to an object's surface contact points (e.g. tires, crawler edges) are projected onto the 3D terrain model to determine the positions where they intersect the surface.

The pitch of the object is calculated first by determining the mean terrain elevation (i.e. height) along the object's front and rear edges. In particular, the mean elevation along the front edge can now be obtained by simply averaging the *y* coordinates (i.e. heights) of the calculated front contact points (e.g. front tires). The mean elevation at the object's rear edge is similarly the average of the *y* coordinates of the rear contact points. The object's pitch at that animated instant is then given by the direction (i.e. vertical orientation) of the 3D vector constructed by joining the computed mean elevation points at the object's front and rear edges.

The side roll of the object is finally calculated using an exactly similar procedure. In particular, the side roll is given by the direction of the vector constructed from the mean elevation positions along the object's left and right edges. These individual computation steps are obviously not visible during an animation. At each animated instant, the object is drawn on the screen in its final, fully-oriented position. PathFinder thus computes a moving object's accurate 3D configuration by prudently synthesizing inputs from the defined 2D motion path trajectory and the terrain model on which the simulated operations are animated.

5. FUTURE WORK

The variable speed motion of simulation objects traversing motion trajectories during the performance of communicated activity instances is purely based on kinetics. In particular, no physical constraints (e.g. mass of an object, its locomotive power, grades of the terrain, gravity) are considered in the computation that describes the simulation object motion. Any depicted 3D process (e.g. a loaded truck continuously accelerating uphill a steep haul road) is a faithful representation of the information communicated by an underlying DES model regardless of whether that process (accelerating when travelling uphill loaded) can be accomplished in real life. Methods that can provide such feedback on physically-impossible simulated processes during visualization can, however, be of significant help within a framework intended to validate modelled (and animated) processes. Future work can explore techniques of designing such methods by incorporating dynamic physical variables in computing the motions of simulation objects.

In addition, the time-based velocity profile shape scaling techniques designed in this work assume the profiles to be composed of piecewise-linear segments. This guides the procedure used in scaling the original defined profiles to accommodate communicated activity instance durations. However, typical velocity profiles generally exhibit curvature, particularly in the acceleration and deceleration phases of a moving object. Future work can explore techniques of defining such curved velocity profiles and design methods for heuristically scaling them to accommodate different motion completion times.

6. CONCLUSION

The presented research extends the state-of-the-art of scientific 3D visualization of DES modelled processes. The work puts in place techniques that can be used to instruct virtual simulation objects to follow any arbitrarily-shaped velocity profiles while adhering to fixed motion completion times when traversing along any defined motion path trajectories.

Discrete-event simulation models, by their very nature, are unconcerned about the rate at which activities in a model are performed. They only enforce the time instants at which activity instances start and end. Since the temporal and spatial control of virtual simulation objects cannot be wrested away from DES models, any information that defines an object's variable velocity profile must originate within an underlying DES model. A computation scheme that allows DES models to define the general shapes of relevant velocity profiles and then heuristically scales those profiles to accommodate communicated activity instance durations performs well in a framework for animating DES models.

Such a technique not only allows simulation objects to be moved with any arbitrarily shaped velocity profiles, but also ensures that their temporal and spatial control rests entirely with the underlying DES models. The visually accurate animation results obtained prove that this is not only possible, but also very effective in convincingly presenting modelled operations in dynamic 3D virtual worlds.

7. ACKNOWLEDGMENTS

The authors gratefully acknowledge the support of the National Science Foundation (NSF) CAREER (Grant CMS-9733267) and ITR (Grant CMS-0113890) programs to this research. Any opinions, findings, and conclusions or recommendations expressed in this paper are those of the authors and do not necessarily reflect the views of the NSF.

8. REFERENCES

- Barrus, J. W., and R. C. Waters. 1997. QOTA: A Fast, Multi-Purpose Algorithm for Terrain Following in Virtual Environments. In Proceedings of the Second Symposium on Virtual Reality Modelling Language, 59-64, New York, New York: Association for Computing Machinery.
- Bentley Systems. (2003). "Bentley Dynamic Animator", Exton, PA. (www.bentley.com).
- Donald, D. L. (1998). "A Tutorial on Ergonomic and Process Modelling Using Quest and Igrip", Proceedings of the 1998 Winter Simulation Conference, Society for Computer Simulation, San Diego, CA, 297-302.
- Fukuchi, Y., Hirai, Y., Kobayashi, I., Hoshino, Y., and Kazuhiro, Y. (1999). "Application of Collaborative Supported Frame Accurate Animation for Bridge Construction Project", Proceedings of the IEEE International Conference on Information Visualization, Piscataway, NJ, 166-173.
- Kamat, V. R., and J. C. Martinez. 2003a. Automated Generation of Dynamic, Operations Level Virtual Construction Scenarios. *Electronic Journal of Information Technology in Construction (ITcon)*, Vol. 8, 65-84. Available online via <http://www.itcon.org> [accessed April 5, 2006].
- Kamat, V. R., and J. C. Martinez. 2003b. Validating Complex Construction Simulation Models Using 3D Visualization. *Systems Analysis Modelling Simulation*, Vol. 43:4, 455-467.
- Kamat, V. R., and Martinez, J. C. (2004). "Dynamic Three-Dimensional Visualization of Fluid Construction Materials", *Journal of Computing in Civil Engineering*, Vol. 18, No. 3, American Society of Civil Engineers, Reston, VA, 237-247.
- Kamat, V. R., and Martinez, J. C. (2005). "Dynamic 3D Visualization of Articulated Construction Equipment", *Journal of Computing in Civil Engineering*, Vol. 19, No. 4, American Society of Civil Engineers, Reston, VA, 356-368.
- Law, A. M., and W. D. Kelton. 2000. *Simulation Modelling and Analysis*, 3rd Ed. New York, NY: McGraw-Hill.
- Li, H., Ma, Z., Shen, Q., and Kong, S. (2003). "Virtual Experiment of Innovative Construction Operations", *Automation in Construction*, 12(5), Elsevier Science, New York, NY, 527-541.
- Lucas, V., and Broll, T. (1997). "CASUS - An Object-oriented Three-dimensional Animation System for Event-oriented Simulators", Proceedings of Computer Animation '97, IEEE, Piscataway, NJ, 144-150.
- Rohrer, M.W. (2000). "Seeing is Believing: The Importance of Visualization in Manufacturing Simulation", Proceedings of the 2000 Winter Simulation Conference, Society for Computer Simulation, San Diego, CA, 1211-1216.
- Schriber, T. J., and D. T. Brunner. 2001. Inside Discrete-Event Simulation Software: How it Works and Why it Matters. In Proceedings of the 2001 Winter Simulation Conference, ed. B. A. Peters, J. S. Smith, D. J. Medeiros, and M. W. Rohrer, 158-168, Piscataway, New Jersey: Institute of Electrical and Electronics Engineers.
- Tsay, T. C., Hadipriono, F. C., and Larew, R. E. (1996). "Virtual Reality Modelling for Bridge Construction", Proceedings of the 3rd Congress on Computing in Civil Engineering, ASCE, Reston, VA, 63-69.
- Tseng, C. H., Hadipriono, F. C., Duane, J., and Larew, R. E. (2000). "A Three Dimensional Construction Operation Simulation of a Multistory Building in a Virtual Environment", Dept. of Civil and Environmental Engineering and Geodetic Science, The Ohio State University, Columbus, OH. Available: <http://www.siu.edu/~coalctr/paper393.htm> (Nov. 20, 2006).