DISTRIBUTED VIRTUAL PHYSICAL INTERACTIONS (HOW TO SCRATCH THE SURFACE)

Gregory B. Tackett
U S Army Aviation and Missile Command
Research, Development, and Engineering Center (AMRDEC)
AMSAM-RD-SS-AA
Redstone Arsenal, AL 35898 USA

Keywords: Grand Challenge, haptics

Abstract

It is inevitable that any discussion of Modeling and Simulation (M&S) Grand Challenges will eventually include the serious and zealous promotion of a "holodeck" capability, the gist of which is an immersed virtual environment made -to-order upon command, and indistinguishable from reality. This paper is not that promotion.

It is valuable, however, to take the ubiquitous concept of the holodeck and use it as a departure point for describing an immersive environment that does not require the suspense of any laws of physics to design and utilize, and yet is so distant from our current way of interacting that it constitutes a Grand Challenge. This is the scope of this discussion, and the concept behind distributing virtual physical interactions.

INTRODUCTION

For the sake of discussion and example, I would like to propose a simple set of conditions: Let's assume that I have an immersive portal into a virtual environment, and so do you, and I'd like us to do a few simple things, "together" in a virtual sense. Let's limit our input devices to gadgets that sense our complete body movements and articulations in three-dimensional space (a challenge in itself, but probably not a Grand one). We also need to presum e devices that provide force feedback (truly a Grand Challenge in it's own right, clearly a technology barrier, but I leave that to someone else to explore). To make our environment interesting, we will assume that both of us have taken the time to create models of some of our favorite things, our virtual houses for example, or our virtual cars. We also need to assume that we have not spent the past six months or the last six minutes working out interfaces with each other, developing Federation Object Mod els, or purchasing common software from any particular visualization system vendor.

The issues regarding input devices and force feedback, in M&S life cycle terms, fall into the category of "executing the model of the system", which is interesting and tempting to pursue, but this paper will focus on a more fundamental issue, through discussion of a previous life cycle phase; "representing the system as a model". In this case, however, the system under discussion is "Newtonian physics" and the model is a distributed virtual environment. This problem space obviously applies across a general span of M&S applications rather than a single vertical niche.

So here are a few simple tasks that I'd like to perform with you in our distributed virtual space:

I'd like to ring your doorbell.
I'd like to drive your car.
I'd like to shake your hand.
I'd like to arm-wrestle.
I'd like to Sumo-wrestle.

IS THIS A GRAND CHALLENGE?

For the purposes of this forum, a Grand Challenge problem has been defined to exhib it at least the following characteristics [1]:

1. The problem must be demonstrably hard to solve. Ideally, the problem requires (demonstrably) several orders-of-magnitude improvement in our capability in one or more areas.
2. The problem must not be known to be unsolvable. If it probably cannot be solved, then it can't be a Grand Challenge. Ideally, quantifiable measures that indicate progress toward a solution are also definable.
3. The solution to a Grand Challenge problem must have a significant economical and/or social impact.

The Problem Must Be Hard to Solve – Limitations in Current Art

Current art is based upon a legacy of computer games and military trainers, or virtual prototypes, that rely on either general purpose input devices (such as a joystick), or higher-fidelity controls and displays (such as an aircraft cockpit).

The signals and digital input from these devices are then interpreted in the simulation, or published in the case of a distributed simulation, to interact with other players and objects in the simulation environment.

The difference between our Grand Challenge case and current art is that in current art, the elements to be physically manipulated are not immersed in the environment, but consist of real hardware, which by its nature is not physically accessible across a distributed link.

While there is an emerging technology for instrumenting the human body for direct input, including products such as data gloves and head-trackers, and while the discussion in this paper relies on the continued evolution of those products to provide the assumed body input, the existence of these products alone is not sufficient to create the needed virtual physical interactions in a distributed fashion. Current use of these devices typically involves some combination of abstractions, such as touching certain finger-thumb combinations to grab or manipulate, or other combinations to move around in the environment. And even these artificial interactions are not published in any way that could be used to simulate a distributed physical interaction between combinations of bodies or devices.

Current art in interplay between objects and players also has its basis in legacy military applications, and can be summarized by the simple concept that things hit things. Missiles hit targets. Bullets hit dirt. (Remember, I'm a missile guy.) Vehicles hit buildings or streams. Dismounted soldiers hit walls.

Behaviors are also scaled to this limited level of interaction. Complex geometries are simplified into "bounding volumes" for use in collision avoidance algorithms and computationally burdensome line-of-sight calculations. An entity may bounce off another object, may get stuck, may climb over it, either sticking to a vertical wall or "clamping" on the top surface, all based on interactions having no basis with the physics behind the two objects.

Current art in three-dimensional modeling follows the same trend, and is driven by visualization, not physical contact. Often the same algorithms that determine what polygons intersect for visual purposes are not accessible to determine the intersection of two complex geometries to arbitrate a physical interaction.

In summary, current art and industry don't lend themselves to a solution space. The technology to design an individual simulation or two to interact with one another's explicit complex geometries in a way that physical attributes can determine the result may not be hard enough to be considered a Grand Challenge. However, moving the state of the industry in this new direction, with standards and repositories and all the resources that would allow this kind of simulation to proliferate is in the realm of very difficult.

The Problem Must Not Be Unsolvable – Defining Physical Interactions

Before we can even begin to apply simple Newtonian physics to an interaction, we have to have enough geometric fidelity and resolution to realize contact. A simple algorithm to detect the intersection of two bounding volumes, as mentioned above, is certainly inadequate to determine the contact, touch, or intersection of two complex geometries. Even contact between distributed object geometries as simple as a finger and a doorbell are beyond currently demonstrated capability. While this particular example could be demonstrated with little development effort, the construct of individual special cases does not satisfy the requirement of interacting without any previous custom integration work.

Once we have established contact, we can begin to worry about applied force. Is mere contact adequate to simulate a doorbell press with a simple intersection of polygons? Probably. But if I try my next simple task, driving your virtual car, you might want to know how hard I'm pressing the accelerator or the brake (unless you have virtual insurance). Just opening the door, turning the ignition, then gripping and turning the steering wheel, each of these interactions demands an exchange of data about force at specific geometric angles, and the associated torque and friction, yet are the very basic principles in our everyday manipulation of real-world controls.

Once we've mastered interaction with hard surfaces, we can begin to improve our resolution and fidelity to allow elastic/plastic deformation in geometries, and perhaps get the feel of a good handshake. Once we prove our handshake, arm-wrestling should be a straightforward extension. Add a little sweat, and mass/momentum information, and our ultimate goal of virtual Sumo-wrestling is within our grasp (puns intended).

The progression described above could serve as a set of "quantifiable measures that indicate progress toward a solution" as stated in Grand Challenge characteristic #2.

As a first step, prototype virtual control suites, virtual machines, virtual keyboards, could be developed using only virtual push-button interfaces requiring nothing more than algorithms that determine intersection of complex geometries. For example, a doorbell model could sit passively in virtual space and simply monitor the location, orientation and shape of published entities to detect contact with the button area, and trigger a bell whenever there is an intersection. Like its real-life counterpart, the button model would not care if the intersection derived from a finger model, a nose, an elbow, or the barrel of a tank. This level of performance requires nothing from the publishing entity that is not already considered in a typical “entity state” construct, although the complete three-dimensional geometric models of the objects must be shared in advance, at some time prior to a potential interaction during simulation execution.

To get to the next level of performance, the publishing entity would be required to expand its entity state to include force vector and frictional coefficient information. With this added data, grip, friction, and torque can come into play. It is at this state, for human-in-the-loop simulations, that force feedback technology begins to be of value. When applying force between articulated or moveable parts, a high-speed feedback loop must integrate between the applied forces and motions of parts. Once this basic information is available, however, it enables any possible combination of virtual manipulations of rigid, articulated objects.

Another intermediate capability step could be provided by simple inclusion of gravitational and inertial effects at this point, allowing entities to push, pull, and carry rigid objects in a virtual environment, and allowing explicit walking, running, and other human motions without resorting to visual animation techniques.

By coupling these physical parameters with pliant geometry models, ideally consisting of continuous, curved mathematical surfaces rather than polygons, a simple, realistic distributed handshake would prove to be a significant quantifiable measure of success. The final demonstrations of various degrees of human-to-human interactions such as wrestling should be natural extensions of the previous capabilities, the fidelities of these interactions only being limited by the quality of the force feedback technology.

The Solution Must Have Significant Impact – the Value in Representation and Integration

Note that even in the final complex form, these interactions are explicitly state-driven. There is no “interaction” data exchange required, because all the action takes place in response to direct contact and applied force as published in the expanded entity state vector. This approach allows for detailed interactions between discovered entities with no previous coordination, and no knowledge of the other simulated entity beyond its geometry and published state. This method more accurately reflects real-world interactions, where objects respond to the forces applied against them, regardless of the source of the force. This is great news for those of us who spend too much of our time working on integration and interoperability. It allows me to place my 1974 Camero model in the virtual environment and let anybody with the right key drive it around, without investing in an integration activity and hosting a FOMARAMA. In a more practical sense, I can push the controls of my missile fire control system into the virtual environment so that soldiers can operate it from their home stations rather than traveling to Redstone Arsenal.

This immersion of controls (and displays, though not discussed here) can provide a significant cost reduction, far above the simple savings of not traveling to a hardware facility. Much current virtual prototyping effort is only cost-effective when compared to high-cost prototypes of products like aircraft and weapon systems. In fact, in our zeal to find applications for virtual prototyping, we are often spending significant funds to do things in virtual space that are more cheaply and effectively done in the back yard. Reconfigurable hardware controls help mitigate the cost for lower-end products, but at the technical expense of reduced fidelity. But as virtual controls are used and re-used, allowing cockpit production to become a non-recurring cost, then simpler virtual tasks become more cost effective.

Given this approach, the construct of a virtual prototype could be a mere compilation of the three-dimensional Computer Aided Design (CAD) model of the system and the interfaces to actual control software where system digital control exists. The only unique software development, which could still be highly standardized and reusable, would be digital representation of analog controls. This highly automated link to the tools used in the system design process and the actual final system software products, provides not only cost savings, but rapid virtual prototyping cycles that become an integral part of an iterative design process.

Though the integration/cost payoff alone warrants pursuit of this approach, the largest benefits are in the representations and activities it allows in an immersed environment. With this approach, it becomes possible to develop virtual laboratories, virtual terrain, virtual cities that come to life with functioning machines, people, synthetic creatures, and fantasy characters. What has previously been limited to visualization in our “look but don’t touch” virtual environments, can be expanded into true physical interaction, allowing the creation of virtual worlds worth visiting for work and entertainment. These worlds could be opened for population by any combination of developers, pioneers, and virtual emigrants, so that in our simulation spaces, we might create the kind of parallel universes that have only existed in science fiction.

These other worlds need not be constrained by the proper physics and Earth-normal gravity that we incorporate in our geo-specific universe. Alternate standards or physics-servers could be created to define outlandish interactions in order to extrapolate to develop worlds corresponding to other gravity conditions, or other physical domains. Sub-atomic, extraterrestrial, and hyper-thermal virtual worlds could be developed to support analysis of rare or hard-to-observe physical behaviors in extreme environments. Or simple scaled force data could be applied to enhance human behavior in a virtual environment (which is the only way I could win that arm-wrestling match).

CONCLUSIONS

In science fiction, a holodeck design requires real-world technological miracles such as force fields, teleportation, and anti-gravity to provide its necessary functionality. Short of waiting for those supernatural tools, which clearly violate Grand Challenge characteristic #2, we might still strive towards natural physical interactions in immersed virtual environments through the development of an underlying software infrastructure to publish, arbitrate and simulate contact and force.

Pursuit of this Grand Challenge could provide intermediate value as the technology improves. This is not a trip to the moon, with final payoff at the end of a long development, it is a trip down a road paved with useful products, a trip that could pay for itself along the way through reduction of integration costs and increases in simulation value.

The Grand Challenge: to expand our concept, approach, and execution of models to fully characterize their physical state, enabling interactions to occur in the virtual world as they occur in the real one. The Grand Payoff: realistic virtual worlds, and an infinite variety of fictional worlds where those same miracles mentioned above could be as commonplace as a handshake.

Acknowledgments

The author extends thanks and appreciation to Mr. Dell Lunceford and Dr. Ernest Page for inviting this article.

Reference List

[1] Guidelines to Authors, First International Conference on Grand Challenges for M&S, 2002

Author Biography

GREGORY B. TACKETT is the Distributed Simulation Manager for the U.S. Army Aviation and Missile Command Research, Development, and Engineering Center (AMRDEC). He received a B.S in Physics from Mississippi College in 1982. He is a member of the Simulation Interoperability Standards Organization (SISO), the Military Operations Research Society (MORS), and the Military Applications Society (MAS). He sits on the Defense Modeling and Simulation Office (DMSO) Architecture Management Group (AMG), the Army Materiel Command (AMC) Research, Development, and Engineering Center (RDEC) Federation, and the Distributed Simulation Working Group (DSWG), and is an occasional contributor to various simulation-related conferences and publications.