

Dean Chang, Hunter Hoffman, Saiki. Kevin Montgomery, Phil Harter

Game Environment

Dean Chang:

I'm Dean Chang and I'm the CTO and VP of Gaming of Immersion Corporation. We develop and license haptic technologies, and if you ever played a video game, where you felt the Door Shock and the Play Station controller or the forces in a steering wheel, that's our technology. If you've used one of our CathSim products or endovascular or endoscopy products, simulators, that's also our technology. I'm glad to be here.

Stan Saiki

I'm with the VA-DOD Pacific Telehealth and Technology Hui. This is a Federal Technology Development partnership in Honolulu. We've been working with Dale Alverson in New Mexico on VR simulations and Medical Education for about four years now.

Kevin Montgomery

I'm Kevin Montgomery. I'm the Technical Director of the National Biocomputation Center at Stanford University. I have the distinct honor and pleasure to work with some of the people that are here. I'm more of a technologist, so one of my focuses is going to be about how these technologies can improve what it is that people want to do and want to learn.

Hunter Hoffman

I'm Hunter Hoffman from the Human Interface Technology Lab. I do research on medical applications of Virtual reality for distracting burn patients during wound care and treating phobias, and post dramatic stress orders. One stress disorder that we are treating is for survivors of terrorist's attacks. A lot of the research I do is manipulating interfaces to try to improve the effectiveness of the clinical treatment and measuring the outcomes at the Univ. of Washington, Seattle.

Phil Harter

I'm Phil Harter, Emergency Physician. I work at Stanford, and have been working with SUMMIT. I am not a gamer. I have an X-box at home for 2 reasons. One is I have a teenage son, and the other is I need something to play DVD's. So I'm the guy for which you are building this to some extent. My position in general is to come in with the subject matter, but to make sure that whatever comes out of this, at least I can play and understand and it makes sense to me. I work with the SUMMIT group in the VR Emergency Department, and do the debriefing of the trainees.

LeRoy Heinrichs:

How are we going to use interfaces that we have and what's coming?

Dean Chang:

I was intrigued by Butch's talk as he talked about how Super Monkey Ball had a 63% correlation with skills for laparoscopy and I was thinking how did they measure that, and what about the other 37% that's uncorrelated; how can we get closer to 100%? When your playing Super Monkey Ball your playing with your thumbs on these little thumb sticks and it occurred to me, "is this an opportunity to make a peripheral that actually has that one foot long laparoscopic tool with the little scissor handles stuck onto a Nintendo game pad, and now the correlation factor goes up to 89% and instead of a 42% reduction there, so we have an 80% reduction area?" So you've got me thinking about basically, "Is there an opportunity for new types of peripherals that take advantage of these X-box and PS2 platforms, that suddenly would make it a lot easier for one to practice the feel part of many of these procedures that we've been talking about?" So far we've been, more often than not, talking about the visuals and sometimes the sounds, and ignoring how the procedure actually feels. As every body knows, their sense of touch is such a critical part of so many aspects, particularly in surgical procedures, so it's important to capture that part too. A lot of people brought up the distinction of decision making and familiarity versus skills, and I agree with that 100%, except that when I view skills, I'm thinking in terms of the manual dexterity frame of mind of skills; that's the part that a standard controller doesn't provide. In general I think there's a unique opportunity to have new peripherals that really let you feel what its like to do some of these procedures that aren't prohibitively expensive as it was 10 years ago when just the SGI itself cost 10grand.

Stan Saiki:

As I listen today, I'm struck by the different approaches of gamers versus educators. As an educator, I'm not really that interested in the fidelity of the game. I'm not trying to shoot the drill sergeant, because I'm trying to teach something else in a different agenda. From the educators' perspective, the development of the application itself has many factors involved, one of which is what the subject matter experts come together (about). Our procedure is to use subject matter to define a curricular agenda, and then secondarily, what can we afford? Ben's comment about, "what's in your budget?", from an educational prospect, we have to worry about production budget, but also about the deployment budget. When we are talking about 60, 80 100 students, it's really difficult to use head mounted displays that cost \$25,000 each. Our particular experience has been that it's much more appealing to be able to do something on a flat screen even though the fidelity may be lower.

Kevin Montgomery

One of the panels was talking about the appropriate levels of realism and in the context of interfaces as well, there's an appropriate level of interaction and interfaces required. Some of these technologies are really great but may not necessarily be what are really required for learning the skill at hand. Along those lines, some of the technologies that we are going to be talking about are display technologies from both immersive and non-immersive, collaborative and non-collaborative, is a good way to look at it. From a simple screen perhaps a screen augmented with stereo glasses to head mounted or wall mounted displays, caves, other systems that might be non-immersive, but collaborative, as well as distributed systems which maybe both collaborative and immersive.

About distributed spaces: Michael Zyda, one of the earlier panelists, has done a lot of work in distributed collaborative spaces, as have a number of other researchers, and that starts to be relevant as well. Also the means by which the will of the user is related to the computer, through either tracking or haptics devices; trackers, meaning there are various technologies optical, mechanical, acoustic, inertial, all of them are different technologies, and all of these I've mentioned, these tracking technologies, even the haptic technologies. They're all just tools. It comes down to what you want to use them for. You can drill a screw into this table using a hammer, but it's not going to give you a good result. Similarly if you use the wrong tools for visualization or interaction, it may be either wasting money or may not be appropriate. Along the lines of haptic devices, there are stylus type devices, laparoscopic or endoscopic tools, catheter-type tools, and other more specialized devices too. But again, you have to pick the right tool for the job. A personal experience with the X-Box to which we tried to port our simulation software some time ago failed because Microsoft was disinterested, even obstructionistic. So, economics is the important part. If you lose money on every game console and you only make up for it in the licensing fees by selling at least 100,000 of these games, "how are you going to talk a company into providing you support when you're going to max out your market at 10,000 units?"

Finally, are haptics really necessary? That panel will talk about "when and where" haptics are really appropriate, and whether it's better to have more lower-cost, well integrated simulations in some areas depending on what you're trying to do. One can have lower costs per seat, as Dr. Saiki was just mentioning, and when it's really necessary to have the more expensive, more difficult program for interfaces, the higher cost is figured into the equation.

Hunter Hoffman

I'm a big fan of using immersive VR helmets even though this technology is very, very expensive, but there are a number of new display technologies on the fairly near horizon that have potential for dramatically reducing the cost of the helmets. As we've seen with a lot of the hardware dropping in price, like scanning fibre or virtual retinal display technologies being displayed at the HIT Lab where I work, the costs are coming down. Tactile cues are important – when we were treating spider phobics, we had one group which was able to physically touch the virtual spider by using mixed reality. For the comparison group, we had a toy spider, and they reached out with their cyber hand to touch the virtual spider, while their real hand simultaneously touching this furry spider, producing the illusion of physically touching the spider. Incredibly, on objective measures of treatment success, in other words in their ability to approach a live tarantula after their therapy, the ones who physically touched the spider were able to approach twice as close. They were able to approach within six inches, but the other group only could approach to two and a half feet, so that tactile cues can be very helpful with learning. So the nice thing about putting someone into a helmet, it blocks their view, producing the illusion of being in a computerized world that pulls a huge focus of their attention into the virtual world, such a huge draw that they can lure their attention away from pain signals during wound care for severe burns. So anybody who studies memory, attention is helpful for creating memories.

Phil Harter

We've talked a lot about different environments today. Everything from dropping black eyed peas into a cup, to holodeck v. 001 alpha1, and everything in between. It really is going to be driven by what your learning objectives are for your particular scenario, what ever it is you're trying to get, it's going to vary. I don't need to learn how to work the scissor handle laparoscopic stuff for my job, so what I'm going to pick is a whole different environment of virtual reality in order to teach my students and my residents how to practice emergency medicine or trauma care or dirty bomb blast or whatever. A lot of the skills that we are going to be able to assign to one of these types of environments are going to be transferable to other types of scenarios. A dirty bomb is something I will probably never see in my lifetime, God willing, but the skills that I would need to be able to manage that (situation) might help me take care of other types of disasters, those which I'm much more likely to see; e.g., multiple car crash or all kinds of different natural disasters, e.g. earthquakes. So whatever environments we develop, they really need to be designed for whatever the learning objectives are going to be for that particular group of students.

LeRoy Heinrichs:

Dean, you've designed a number of interfaces and controllers. How do you approach the design application area?

Dean Chang:

I think I'll borrow something from what Kevin said, and that is to start off with what your budget is. When we work with licensees, company's such as Logitech – for instance in 1995-1996, they wanted to make a force feedback controller for games, because whether your shooting a gun or flying a plane or driving a car, feeling what's happening in the road through your controller, produces much more realistic experiences. We found informally that people playing driving games like Grand Turismo III on a Play Station actually get faster lap times with force feed back than without, and the reason is that when you drive a real car, you know that the faster you go around a sharper turn, the more centripetal resistance you feel trying to stay in that turn. So with the force feedback steering wheel, you feel the same forces; if the tires are about to loose contact with the road, about to 'break loose', and you're about to fishtail, you'll feel that in your wheel before you visually see it. If you wait till you visually see that you started to loose it, it's too late, your gone – your back-end's out. So there are a lot of similar situations that aren't just about driving in which the combination of all the inputs, sight, sound and touch, they all add benefit. That's why Logitech, Microsoft, and others want to introduce gaming peripherals; their target price was \$200. That's about the most they believed they could charge for a gaming peripheral, so we had to work around that cost boundary, and we also had some pretty well defined, then current gaming products to use as starting points, joysticks and so on. In medicine it's much more wide open. You're not limited to game pads, steering wheels and joysticks. We've tended to look at which are the procedures where our technology, by adding the sense of touch, can really get that much closer to the real life experience, and have a large enough market that can benefit enough people to do warrant the initial R&D and so on. Price tags are obviously much different. Products start from 10 grand and go all the way up to \$100,000 and so on, so going back

to what Kevin said, it really starts with the budget. Then we work with the content experts, in this case the peripherals, and the experienced users to make it 'right'. With CathSim for instance, it's a catheter simulator for practicing how to insert needles, whether its for intravenous catheterization, or phlebotomy, or so on; for that procedure, we worked with doctors who actually implemented real needles, measuring the force it took to break through the skin and to break through the vein wall, and so on. Using that as a starting point and working with people who'd done tens of thousands of procedures, to really fine tune it, we got not just one generic case representing the entire human population, we now have all kinds of different modules, e.g., for elderly people because their veins are more brittle, and so on. So working really closely to bring our expertise, which is in simulation haptics, and marrying that with the subject matter expert is the key, whether it's in gaming or in medicine.

LeRoy Heinrichs:

I'm looking forward to an amphitheater in which each seat has a couple of chopsticks and a group of students can all be working on software simultaneously, and in an independent, and then in a coordinated way. And then we must consider another budget, and that's the time for learning. To how many students can one teach the same lesson at once? The professor can't come back and teach it 80-160 times? So that's a budget we have to think about, too.

(Audience – unidentified)

There is somewhat of a difference between the cognitive versus the skills based aspects. In my optic view being a surgeon, we don't worry about the cognitive part; we just do the skills base (as some say!). The interface between me holding an instrument that I'm trying to teach a resident to use during a lap-choly, versus asking them to begin experiencing an immersive environment, as if they truly are in the ER, is an entirely different set up. With our set up, we have the video trainers which are like Butch showed, with the (real) beans that you move around, but you're using real instruments. Our residents prefer the video trainer. They don't really like the virtual reality set up, and that's predominantly an interface issue – not the task itself. And that's an area that I'd like for you all to address is "whether putting it on X-box is an intermediate step? It doesn't have to be the best in the world. An ideal approach is, "I'd put in the CD of my patient that projects them into the virtual reality world, and I'm going to be doing the exact operation as the day before". We are a long way from that. But, "is there some way we can get to having an X-Box that gives us at least the next step to where residents are able to train in that kind of environment?" Where do we go from here?

I'm speaking about the 'pea-trainer', and going back to the face validity concept, residents are using real instruments, a real laparoscope in a box trainer, (so they are not seeing it), and they're looking at a regular display as they always would in the OR. It has been shown in several studies that they prefer that interface because of the face validity.

Hunter Hoffman

In that situation, I would not recommend using a head-mounted display because the desktop display is much more similar to that used in surgery.

Kevin Montgomery

As a technologist, as much as I love the technical challenge of replicating reality and the haptics, doing all this cool research, no matter what, you've got to start with the learning objectives. What are you trying to teach people? And frankly, use the lower cost solution – that's great. If you can use an orange to teach someone how to do an injection – great! They're cheap! Infinite frame rate, and re-usable! And, they don't scream in pain! Then, after starting with the learning objectives, pick the required technologies to help teach that? Don't go the other way around! Don't use surgical simulation as an excuse to work on the fun stuff that you want to work on. Be Honest! I love replicating reality, so I'm going to work on research that way. If you really want to build surgical simulators, great! Build surgical simulators, but build them to teach people to do what it is that they need to be doing.

Dean Chang

In all our medical simulation products whether it's for laparoscopy or endoscopy, or endovascular, we start with the real tool. We take the actual laparoscopic tools, and then we interface them with some electrical-mechanical box that lets one feel the same kind of sensation that should be felt; the right amount of force, the right amount of resistance and so on. We have a cyber glove product; this glove is the most accurate hand tracking you will find. It's very refined in how it can track the hand; we have a program demonstrating sign language in which it will 'speak what your signing', so effectively that it can tell the difference between a bend at the first knuckle versus the second knuckle versus the third knuckle of each finger, and so on. But, to use the cyber glove to try to do a laparoscopy simulation, you don't get the feel of holding those tools and how much force one needs to apply. Maybe you can get how much force to apply to close the grip and all that kind of stuff, but just feeling the actual handles, and getting the little creases in your fingers from squeezing them tight, that's an important part of learning how to do the procedure right. So, we've taken the path that medical products that physicians or other medical persons need to hold, they should be the 'real thing', and then we do our magic behind the scenes to make it as realistic as possible; that approach has worked very well for us.

Audience (Pamela Andreatta-Univ. Mich.)

My question is associated with both the development and the design, but in particular related to the duie, and I'm wondering if you could address the advantages and disadvantages of a tactile model where you actually have something physical versus an image driven haptic system; so you have the haptic driving some sort of a display as opposed to an actual tactile model.

Kevin Montgomery

Very briefly, when you have a physical thing, it has one shape and it doesn't change too much – that's kind of the way it is. In the virtual world I can bring up different kinds of

patients and different kinds of anatomy, different pathologies, case of the week, and you can do it over, and over again. So basically if you pay the capital up front, then you can try again and again and again. There are lots of other models that other people in this room can talk about, in which they have instrumented mannequins and non-instrumented mannequins, where you can pop in a piece of anatomy to make it look like a morbidly obese patient or a patient with a difficult airway or what ever; that s difficult to do and for many procedures, it may be destructive – you can't do (them) over very well. The promise of doing them virtually is that you can bring up any patient for which the simulator has been made, but also, any other patient that anyone has ever prepared before, or even a new patient who's come through the door and who has an unusual anatomy. In the future, students and residents will be able to train on such patients, also. So, it gives you the ability of a greater breadth of surgical training by presenting the resident with anatomical variations and pathologies that may take them 20 years before they encounter in the field.

Also, quantitative performance metrics, objective quantitative metrics are a big deal for simulators that you can't get otherwise.

Audience (unidentified):

When you design your systems, the realistic tools as interface, how do you evaluate the systems to know whether what you've designed makes sense?

Dean Chang:

With our CathSim system, we worked with some Ohio doctors at the supercomputer center. We actually start with collecting some real data, but in the end, we are completely dependent on the opinions of people who have done tens of thousands of these procedures to say -“It's just not quite right there, and you need to fine tune it.”

So it's a mix of both science as well as the opinion of somebody who is considered an expert.

Audience:

Have you done quantification to evaluate the skills during use by novices, intermediates, and by experts?

Dean Chang:

Yes, we do some of that. At the very end of our CathSim procedure, you are given an evaluation of how well you performed, and how long the procedure took. It even has some kind of model determining how much pain you inflicted if you couldn't tell already, because, unlike an orange, our simulator actually does cry out when you perform poorly.

If you start a hematoma, not only will you see bleeding, both external and internal bleeding, producing the purple 'skin' appearance from the deoxygenated blood, and so on, but the patient might actually start saying “ouch”! In fact if you are using the pediatric module, on a newborn, as soon as you put the tourniquet on the baby's head, the baby starts crying, and that really unnerves some people; even caused some to feel queasy after using the simulator, which is “great”, because we want that level of reality. We don't want it to be a video game in the end, so “Wow, look at this great high score I've gotten”, so that the first time you're actually confronted with an uncooperative

patient, or even a belligerent patient, you forget everything that you've learned. So there's feedback while you are doing it as to how well you did, but also at the very end there's some objective results of your performance.

Audience:

Have you asked experts what should be quantified by the evaluating system?

Dean Chang:

Yes, that's an ongoing process as well – we've just released a new version of CathSim, and we're making product plans for version v.2.0, so that is an active process.

Kevin Montgomery

I've another example that may be relevant to the question. In cleft lip and palate surgery, it takes residents about four cases to know what they're doing. What it really involves is being able to lay down about 11 different anatomical markers, and once that skill is acquired, the rest is just cutting. But learning to do the anatomical marking is the hard part. It takes about four cases to learn that. We've heard anecdotally from many surgeons that the number is four. So I made a simulator to simulate cleft lip and palate surgery, and first we used force feedback and all that stuff, and got a great paper in CARS about this system. Then we realized that the act of cutting is not what residents need to learn, they need to practice laying down the markers. We took out the haptics and made the simulator to focus on the anatomical markers on regular computers, and had a bunch of lay people, residents and attendings use it. And sure enough, residents asymptoted after four virtual cases. The hope is that they won't harm real patients, given this training. Remarkably, residents performed better than attendings – we interpreted that residents were more familiar with this method of learning. Lay people did poorly, as expected. Later, quantifying "how well did trainees perform actual surgery is important, also", and that's one part that we haven't evaluated. But I agree that one needs to quantify after the training whether it makes better surgeons, or not.

Audience (Janet Kim – Philadelphia)

We have a product that we're considering making haptic, and it's a surgical simulation, however the actual VR part of it focuses more on the surgical steps of the procedure, rather than the precision. My question is, "Can you think of any situation in which haptics might have negative effects on someone trying to learn?"

Kevin Montgomery

Bad haptics is a distraction!

Stan Saiki

When the haptics becomes so expensive that one cannot afford any kind of simulation! Expanding on the dynamic between cognitive and skills training, one must have the proper combination of tools for the proper job, as Kevin says.

Dean Chang

There is NEVER a case when haptics is going to hurt – with a qualifier – when properly implemented! But that also assumes you have the right budget and technology as well. We see this in gaming, too. Step aside from medicine for a moment. There are some people who hate force feedback with the vibration in games, and its not because the technology is not good, its because the developer didn't do a good job of taking advantage of it, for example in games like Metalgear Solid, you are a sniper at one point, and you're really nervous, so your heart begins racing, and as you feel the controller, you feel the rapid heart beat in the controller. It's just going pounding, pounding, and it affects your aim because every time your heart beats, you see your telescope kind of jiggle, so you have to go find this drug that calms you down, and then the pounding subsides. In some part of the game, just to stay alive, there is a button you hit as rapidly as you possibly can for about 3 minutes at a time, and you are 30 seconds off, and you do it for another 3 minutes, and 15 minutes later your arms are just burning, and so next your seeing a doctor and the doctor says; "Oh your arm must really be in pain! Hold your controller up against your forearms!" It turns out that the vibration, in full blast, is like massaging your arm – it feels good. So my point is, when done well, people love the technology and say there's no way I would play without a controller that doesn't have this haptic feedback. And the same thing is true in medicine. When done properly, there's no way you can hinder, but, as people have said, there are certainly cases in which one just don't need the haptics. If it really is about teaching familiarity and decision making, and not about the actual mechanics of cutting, for instance. Then one can argue that it doesn't add much; if done well, it doesn't hurt either, but that's where you have to determine, "what is it that you want to accomplish?"

David Gaba (Stanford Univ.)

Back to the issue about how you know whether something is good based upon expert opinion, I agree that trying to put sensors in the entire world, you still must depend on expert opinion in many cases. One of the things that plagued a number of the efforts in development of the manikin-based simulators is that if one relies on too few experts and too few domains, one can get into real trouble because it turns that a lot of these things are very individual-dependent or school-of-thought dependent, regional-dependent, and domain-dependent. So, company's that are making products that run a spectrum from a frontline emergency responder to emergency room people, to OR people and ICU people, and so forth, have found it very challenging to get expert opinions that all converge, and if they only get single opinions, and I think that this is true of almost anything, there is real risk of getting locked into something that seems great to one expert, but many user populations won't like. User-centered design rather than expert-centered design, needs to happen in parallel.

Rick Severinghouse**LeRoy Heinrichs:**

As you can see, this is a very complex area, one with many areas for creativity. Let's thank our panelists. And, thanks to all of you!