

DIGITAL NATIVES

A conversation between virtual reality
visionaries Jaron Lanier and Kevin Kelly

By Casey Newton

In 1989, Kevin Kelly, an itinerant writer, photographer, and future founding editor of *Wired* magazine, traveled to Redwood City, California, to meet Jaron Lanier in his office. Four years prior, Lanier — a computer scientist with a mane of thick blonde dreadlocks — had founded VPL Research in an effort to bring virtual reality hardware to the marketplace.

"Right before my eyes, Jaron Lanier built an artificial reality and then climbed into it," Kelly wrote for his *Whole Earth Review* magazine at the time. Kelly wrote an introduction to Lanier's work at VPL, capturing the growing optimism about the technology as well as fears about its implications. He also organized an accompanying interview with Lanier, in which the programmer attested that, "We're witnessing the birth of a culture here."

Lanier may have been overly optimistic: what followed was a decades-long period where virtual reality continually fell short of sci-fi visions. A quarter century later, Kelly and Lanier met again in the light of all that has happened since. I sat in on the conversation, which took place at Lanier's home in Berkeley, and tossed out the occasional question. But mostly I listened as two of technology's leading voices brought real scrutiny to virtual reality's past, present, and future.

Highlights from Jaron Lanier's end of the discussion are below. And for those with a deeper interest in virtual reality, we've transcribed the full interview.

FULL INTERVIEW



Jaron Lanier tests an environment in VPL's virtual reality goggles

KK: The question I'm interested in, Jaron, is whether anything's changed in 25 years about VR. Did we just sort of skip 25 years and this is a couple months later? Or has there been actual innovation?

J: Well, no, there's been a lot of motion on a lot of levels. I should say that in the time between now and then, I haven't been continuously active in virtual reality. I left the virtual reality field, more or less, around '92 — I was the chief scientist at the

Internet2 organization, which was working on how to make the internet scale; I was a musician in New York and this and that. For me, the 25-year mark isn't such a thing. [Our] interview was actually '87, and I believe your visit that you described was, as well. We'd started the company, VPL research, the first VR company, in '83. My first systems with my friends were something like '81. And of course, the first head-mounted display was '65. So to me the timeline goes back a little further. And my personal involvement goes back to about '79 or so. I tend to think on a 35-year sequence rather than a 25-year sequence.

The biggest thing that's happened is that the industrial use of VR — as opposed to consumer and entertainment use — matured, and has become ordinary enough to be boring. But it hasn't happened in this sort of big, unified way; it's a bunch of little pockets that are each very specialized.

"The biggest thing that's happened is that the industrial use of VR ... has become ordinary enough to be boring."

The major ones are surgical simulations and training. In fact, that one got good enough that at this point our main concern is overuse; there is a danger that surgeons are spending a little too much time in simulators. That's not known, but that's been a concern. We know it's true for flight simulators. For instance, the crash that happened here in San Francisco, the Korean airliner, has been attributed in part to an overreliance on simulator

training and automation in the cockpit, although I think there were some other issues too.

Another one is vehicle design: you basically cannot buy anything that can surround and move you that isn't first prototyped in VR now. Every car, boat, plane, civilian, military — everything is designed in VR now.

Manufacturing processes are generally modeled in a simulation environment, where workers are observed in a simulated environment before machinery is finalized. Another one is security and defense applications, which is huge. There are very elaborate VR tools for all kinds of military scenario training. There are dozens of sub-specialties of different kinds of virtual reality systems for different applications. There's a maturity, I'd say, to the point where it's kind of boring.

Another thing that's happened is the components have finally gotten cheap enough that we can start to talk about them as being accessible in the way that everybody's always wanted. Although, it's one thing to just have an affordable headset that's decent, it's another thing to have a whole system. Moore's law is so interesting because it's not just the same components getting cheaper, but it really changes the way you do things. For instance, in the old days, in order to tell where your head was so that you could position virtual content to be standing still relative to you, we used to have to use some kind of external reference point, which might be magnetic, ultrasonic, or optical. These days you put some kind of camera on the head and look around in the room and it just calculates where you are — the headsets are self-sufficient instead of relying on an external reference infrastructure. That was inconceivable before because it would have been just so expensive to do that calculation. Moore's law really just changes again and again, it re-factors your options in really subtle and interesting ways.

"The components have finally gotten cheap enough that we can start to talk about them as being accessible in the way that everybody's always wanted."

KK: Right, I call it the computational x . Anything you can imagine x is, is going to be done by computers instead of whatever it is. So instead of a lens, you can do computation.

J: Listen, computational displays are a fantastic, emerging area now and those are still hidden from us by the veil of inadequate progress on Moore's law. There's still a computational overhead that's really challenging. But that's definitely coming down the pike. It's amazing.

KK: Explain what computational display is.



Credit: Ben Delaney

J: Well, that's where you do heavy-duty computation to get magic out of pretty simple display components. The terminology is still evolving, actually — different communities use different terms for overlapping meanings. Computational holography is where you compute fringe patterns on a display instead of the actual image you'll see. So, you're computing all these tiny edges and because of the quantum nature of light, when light encounters an edge, it can bend a little bit. If you can calculate very carefully how it'll bend, you can take a simple display, with a whole bunch of little ridges rendered on it, and bend light to create a 3D effect: you can turn a regular display into a 3D display. I'm oversimplifying it, but it's an amazing thing.

Another incredible one is usually called a compressed light field — although, once again, there is different terminology in play — where you do a few layers of displays, and by calculating very carefully how they combine together, you can create this magical, encompassing 3D effect where you can really focus into the image even though there are really only a few layers of

display. And once again, the computation overhead is pretty significant. We're getting to the point where we can really calculate fields of energy instead of dealing with just the bulk manipulation of a field, like with a lens. That's transformative. It really just opens up a huge territory.

KK: One of the important things that we should keep in mind, [is] that there is not going to be a single version of VR, that there will be multiple types.

J: It's going to be like everybody talks about a computer; the term computer doesn't exactly mean anything anymore. Instead we talk about having a phone, or a tablet, or a cloud service.

KK: That's the technology, but maybe we can explore some of the consequences of the technology. Have you changed your mind about [the impact of VR] very much?

J: Sure, I mean, it's broadened quite a bit. In the '80s, I had maybe an outright mystical approach to it. For me, the very most important thing about VR was that when you were in it, you'd feel your own existence, in the sense that if all the sensory input is artificial, then what's floating there, that's your consciousness. So to me, it was sort of proof that subjectivity is real; that consciousness is real, that it's not just a construct that we put on things. Just to notice that you really exist, to me, was the very, very core of it. There were a zillion and one variations on that that [could] become really vivid and colorful in different ways. But that was always the core for me. And extending from that, this possibility of a kind of communication that would involve directly creating what people sense in common instead of relying as much on symbols such as words.

"So to me, it was sort of

proof that subjectivity is real; that consciousness is real, that it's not just a construct that we put on things."

KK: The post-symbolic...

J: Post-symbolic communication, yeah. I used to go on and on about that stuff. And I can still do it, if asked. At the time, I shared a kind of idealism about what digital stuff would do to the world with my friends. And I've actually been connecting with a lot of people from those days. I've just recently started emailing with Richard Stallman again after years of being out of touch; and many other people from those early days. You know, I think all of us had the sense of mission that we were really doing something that would open up the world, and that a lot of mankind's problems were kind of just artificial and due to inadequate technology: if we could just have better communication and all this stuff, a lot of problems would clear up.

I had to reconsider that ideology at great personal pain because I didn't want to question it. For me, it just took a lot of individual people not doing so well as digital stuff rose up. And, in my case, the particular thing that bothered me was initially seeing musicians not do as well as I thought they would. There was a time, up until around the turn of the century, I was writing fire-breathing essays like, "Piracy is your friend" and "Open everything up and it'll work out." Then, when I started looking at the numbers of people who were benefitting, I realized that what was actually happening was the loss of the middle hump of

outcomes; we were concentrating people into winners and losers, which is the worst outcome. I've also become really concerned about VR's role in that.

A few researchers started to do experiments that I would have been terrified to do myself. I'm thinking of a person who has been research partner, a collaborator for many years — Jeremy Bailenson at Stanford. He started to just sort of see how he could screw with people in VR. I was always like, "Can we give them better math abilities by changing how their bodies work?" — that was the kind of thing I was interested in. [Bailenson] was like, "Hey, I want to see if I can screw with their self-esteem by making them gradually shorter during an interaction, or turn gradually more black during an interaction." And he can. This notion that you could see VR as a way to screw with people without their awareness, crossed that with our current business model where everything is about advertising and manipulation and spying, we [will] have a surveillance economy in the online world. It's been very painful to see that potential unfolding.

"The VR world is sort of like the most saturated tracking environment you can imagine."

The whole spectrum is true at once. I think all the mystical, ideological stuff is still as true as it ever was. The potential is all there, and the beauty is all there. But this potential for manipulation is also there.



KK: Actually, I hadn't thought of that, but you're right. If you're in a virtual world, everything is being tracked, in a certain sense; almost by definition, everything is trackable. The VR world is sort of like the most saturated tracking environment you can imagine.

J: Absolutely.

KK: And what's interesting to me is that in our [previous] discussions, I don't think that ever even came up.

J: I don't know if we talked about it. I wasn't ready to contemplate anything short of hippie optimism in those years. I had confronted it, actually. Just to be clear, Norbert Wiener talked about this stuff in his own way. I mean, maybe not quite as head-on as we might expect today, but from very early on, people saw this possibility.

KK: I don't recall that — I mean I know he had lots of complaints but I don't recall this particular complaint.

J: The very, very first issue of *Whole Earth Review*...

"All panaceas become poison: computers as poison."

KK: "All panaceas become poison: computers as poison."

J: "Computers as poison" was the cover. These ideas had currency from the very origin of computer science. Norbert Wiener's book, *The Human Use of Human Beings*, essentially lays out the terms of the problem — not quite in the language that we'd understand today, but it's there if you read it carefully.

In the practically accessible picture, the way we used to put it was like in a virtual world, art is still art, but bullets aren't bullets. It would naturally pull people away from harming each other into a realm that was still consequential in good ways but less consequential in bad ways. That was the idea — that it would be inherently beneficial. The way events have unfolded, the empirical evidence is not supporting that level of idealism. I feel like I have to adjust.



Jaron Lanier and Howard Rheingold at VPL

KK: But at the highest level you say you are still an optimist and we should go and build these systems. As we build VR systems, what things should people keep in mind? Now that we know what we know, how should we build these differently?

J: I think every technical person is obliged to think about how we can move towards a world that really serves people, rather than splitting us into an elite and everybody else. It's funny how old all these ideas are. So much of this goes back to 19th-century thinking, and it's correct thinking, it turns out. It's just the strangest thing. At any rate, I'm really concerned by the way tech culture has evolved since the clouds got rich, you know? We're seeing a kind of a tech-supremacy feeling.

KK: My friends who are making more VR worlds or gear, what do you think they should not be doing that they are doing right or that they should be doing that they're not doing?

J: We have to evolve out of what we're calling the advertising

business model. If you extend the idea of advertising to total surveillance in the way that we're doing it, it doesn't result in a stable, serviceable way to build a society. We have to all come to that recognition, find an alternative, and it's never more true than with the VR stuff.

Obviously, I'm hoping Facebook's business model will evolve by the time they ship something. Facebook is kind of painting itself into a corner where both it and Google are in this mutual embrace of making each other more and more creepy in battle. And they have to find some way out of that.

"Facebook is kind of painting itself into a corner where both it and Google are in this mutual embrace of making each other more and more creepy in battle."

KK: A counter to that is the success of *Minecraft*, which, as far as I know, is not running on an advertising model; that kind of captures that sense of technology for good, and being creative, and going into these worlds. I wouldn't call it a VR, but there's an immersive-world aspect to it. What's actually impressed me most about it [is that] when I saw it, I would never have guessed that eight-year-old girls would be playing this on their phones — in an addictive way. I would not have believed that. Is that evidence this could work in that way?

J: Yeah, I absolutely think it is. It's a great example, and it's not the only one. There are others in the gaming community that have tremendously admirable qualities. In my utopia, there would be a larger number of people achieving middle-class stability in the *Minecraft* ecosystem than there are. But, on the other hand, everything you said is absolutely correct. It's wonderful to see kids work with it, and to my knowledge, nobody is using it to spy on the kids and create behavioral models of them and all that. One of the things about our world today is that in the attempt to create the omni-transparent world, we created the super-opaque, creepy world, you know? It's just absolutely backfired. Everything that we thought created transparency had precisely the opposite effect.

KK: [*Minecraft*] captures a lot of what you were hoping for VR at that time, which is that you have this sort of open-ended sandbox in which you could make anything you could imagine, play in it, and share that world with other people. Now *Minecraft*, it's pixelated, it's blocked...

J: Well, it's a start.



Event producer Sally Rosenthal demonstrates a head-mounted display at NASA's Ames Research Center

KK: But it's very much in that vein. If you just amped up the resolution, all these other factors, then you'd have something very similar to what you're talking about.

J: The particular thing is that *Minecraft* is probably not going to scale a whole lot further. It has a particular identity, which is just great. I have a feeling there will be other designs that scale further. I might be wrong about that, we'll see. At any rate, I totally agree. It is very much along the lines of what I imagined back then as a form of communication that was presenting forms and dynamics that people invented for each other. So that has come to pass.

And not just in *Minecraft*. I think there are a lot of other gaming platforms where you see that sort of thing emerging. There's a

bit of an unfortunate thing that happens in the gaming world — which is a mirror of what happens in the cinema world — where the very biggest budgets and promotions are reserved for what's perceived as the most conservative bets, which are inevitably these macho destruction fantasies. I suspect that those decisions are actually not the optimal ones from the business point of view; that if people had the courage to tunnel through and try to do more things like *Minecraft*, they would actually make more money than with more shoot-'em-up things.

KK: I think the people who are making a living from *Minecraft* are the people making tutorials and the millions of YouTube videos and the books.

J: Absolutely. I absolutely celebrate all of that. I just think at some point, a stable civilization would have to generate more means to sustenance from that larger community than our current system is.

KK: Right. And those occupations, I think, are still opaque to us. In a sense, they're hard for us to imagine. Who would have guessed?

C: In the case of Facebook and Oculus in particular, have you thought of a business model direction that wouldn't be based on advertising and surveillance?



Jaron Lanier at VPL

J: Absolutely. This gets to my last book about that stuff. To my mind, there are a few proposals that are on the table that are interesting. The one that I've been trying to push the hardest was the origin point for networking when it started — Ted Nelson's idea of the universe of micropayments. The thing about micropayments is they have to be really universal or else there are only tiny pockets of people who get anything and everyone else gets frustrated and pissed off. It has to be big enough to really send some benefits around to everybody. That's a hard gap to leap over to get to that point. Facebook has a very tough road to hoe here because they have no background in this stuff. Neither does Oculus. They are kind of starting from scratch, but I certainly wish them the very best success. What they can do is create a system where anybody who wants to can set up a super easy system for micropayments and then add to it. They might see this exploding economy that generates a lot of economic growth, a lot of profit, and a lot of distributed benefits for their users.

"Every time there is a new platform, there is a change to create a new economic model."

Every time there is a new platform, there is a change to create a new economic model. 3D printers for hobbyists have kind of come out on a Linux model, where everyone shares their models. What if that had happened on a pay-per-model business where everyone could pay and get some money out of it? It would be an experiment. A lot of people would really be offended by the very idea of it. But what if it actually generated a lot of cool models and if it gave some people the ability to pay for their kids' college education What if that stuff actually worked out really well and everyone ended up happy?

I don't know. I mean, look, with these things, ideology isn't a good measure. You have to be empirical. Starting up a VR platform is another one of these points where we could empirically test it and say, let's do this one in a whole new way that's never been done before: [create a] really easy, convenient universe of micropayments where everybody is on equal footing, everybody is a buyer and seller, everybody is a first-class citizen, and just see what happens. It might be great. I have a feeling with VR that it could be great — the thing about it is it does take a lot of effort and craft to make good stuff; to make a good *Minecraft* world, or to make a good world in whatever thing you do. It would make sense for society to design that as a way [in which] you can actually directly make a living.

KK: But didn't *Second Life* have a peer-to-peer...?

J: Yeah, I was an advisor to *Second Life*. *Second Life*, to my mind, was a really worthy experiment that yielded mixed results.

KK: The economic part of it seemed to be good, right?

J: If you want to move towards an economy that could be sustainable when Moore's law has moved to some ultimate point — some decades from now — then *Second Life* didn't give results strong enough to support that. But [it came] a hell of a lot closer than most other things. It was really a step forward. *Second Life* was a failure in terms of a design that could interface with existing laws and existing economic systems. It had a huge problem related to taxation and regulation. [With] technology and idealism, you wish away these things, but you can't. There was a degree of fraud, there was a degree of bullying, but overall *Second Life* was kind of impressive. The biggest failure of *Second Life* is that it didn't succeed in holding its audience's attention beyond the initial years. I mean, it did pretty well for a few years but then it kind of fell out of public consciousness.

KK: It's still going, yeah.

J: But I would say it's not going with the intensity of something like *Minecraft*. *Minecraft* has proven some sort of longevity. *Second Life*, for a brief moment, was probably as big as Oculus or VPL in terms of public attention. It didn't sustain that. Another thing with *Second Life* — I don't think it was extensible enough. There was this problem that a lot of this stuff started to look the same because the tools didn't have that much breadth. And that was a big issue.

KK: I think it's a scalable thing. It's centralized in a certain sense. With *Minecraft*, everybody can put it up on their own server and that really helped spread it so you can have a little bit more decentralized, scalable input.



The Cave Automatic Virtual Environment at the University of Illinois

J: Yeah, that's a really interesting question, too. This whole issue about who should own the servers, where they should be, and what works best is actually a very empirical question and is still evolving. I don't think there is any absolute answer because the right answer could vary with a particular design.

KK: If we can go back to the idea of there being varying, different species of VR: so some will be more centralized, and some will be more distributed and they'll just be different forms; they'll just be different media, almost, in some ways.

J: Oh definitely. There is absolutely no question that there will be. There are about four major variations on VR: one where you're all in there, which is like Oculus; a heads-up thing, which would be like Google Glass; another would be mixed reality where you see stuff overlaid and mixed in with the real world;

another would be telepresence where you feel that you've embodied some device; and another is tele-immersion, where you might [be] in a remote location but with sufficient illusion that it transcends video conferencing.

KK: What's the difference between tele-immersion and inclusive?

J: As a practical matter, you can't detail immersion with a heads-up display because your face is covered. I've tried and other people have tried. Let's say I'm wearing some kind of head-mounted display and you're wearing a head-mounted display and we both want to have the experience that we're in the same room with neither of us wearing them so we can see each other's faces. We have to have sensors trying to sense our face, and then we have to reconstruct it well. You get the uncanny valley issues. It's not technically impossible, but it's a pretty long-term goal to do that. The more likely thing is to have neither person wearing them but instead have them sense from remote sensors and have some sort of display that has a volumetric quality. In practice, they're distinct.

KK: So the immersive one is where you are experiencing something. The other one is where someone else is experiencing you.

J: The totally immersive one, which I call classic VR, [is] like what we were selling in the '80s. You can see other people but you have to be avatars, you have no choice. In the other space you're attempting realism. You could attempt any degree of avatasation, or whatever the word should be but the...

KK: Let's make up that word right now. Avatize!

C: Avatizing?

J: Avatized. I'm an avatizer. I don't know. The terminology just becomes so insane. But I would say the five I just listed are probably the big five.

KK: So it's just a matter of degrees in terms of how much of the overlays are visible most of the time?

J: You could say it's a matter of degree, but in practice, it requires a pretty different strategy. For instance, with Google Glass, the focal plane you focus on doesn't really have to match what's there. Another issue is one eye versus two. If you really want to do mixed reality, you kind of have to do it with both eyes. [While] heads up is actually much, much, much more sensible with one eye. You can say it's degrees, but in practice they really diverge into two different designs.

KK: You said you've been involved with [VR] since the '70s into the '80s, thinking about it and very actively building stuff. But it doesn't strike me that you have any interest in spending time in VR now. Am I wrong about that? If there was the right system with the right form, can you imagine yourself spending some significant time there?



An image from VPL's "reality built for two," a virtual telepresence project

J: The first thing I'll say is I use VR a lot these days in research of different kinds — cognitive science research and some visualization stuff. I also love working with exotic optics and sensors and I still play around with that stuff. And I absolutely take time to play beyond what's needed for the research. I still find great pleasure in screwing around with it. But I've always felt, from the very beginning that it's — how would I say it — you know, every musician I know prefers to live without the radio on, enjoys the silence and contrast. The best way to use VR, not in some moralistic or judgmental sense, but just in terms of my experience, is [to] use it as little as possible and enjoy the contrast that it gives me to reality.

At the old lab, VPL, one of the things we'd sometimes have is just a flower sitting there. So if somebody was in a demo for 20 minutes, we'd come out and say, just look at this flower. And you suddenly saw this flower in this hyperreal way because your senses had adjusted to this sort of lower resolution of a virtual world. Then, when you see reality, you suddenly see it with this kind of detail and this density. You see just the sheer reality of it. You just feel things from it. It's really incredible. To me, that contrast, that feeling that you have when you're out of it after you've used it, has universally been more precious than what happens in it. So, yeah, I like it. I might be wrong about this, but I suspect that a lot of people will find what I found: that the coolest thing to do with [VR] is not to be in there for hours the way people are with their pocket devices these days, just staring at the screen...

"When you see reality, you

suddenly see it with this kind of detail and this density. You see just the sheer reality of it."

KK: Or playing games.

J: ...or playing games. The coolest thing is to come out of it for hours. I think that's really just the most amazing thing. I still really do. Here we are, surrounded by all these acoustic musical instruments and in a way, building these devices and building virtual stuff has just helped me appreciate physical stuff more and more. Physicality is just so astonishing.

KK: It's underappreciated.

J: Yeah, maybe we'll appreciate it more when we have something to compare it to.

KK: That's certainly what I call the third way — the nerd way: you investigate real things [by] making an artificial something, which can both give you some evidence about how the real world works and also better questions, more interesting views about the real world when you come back from those. Whether it be VR, artificial intelligence, or artificial democracy, you investigate the real world by making synthetic things.

J: I totally agree. As for what people actually do, I'm kind of bracing myself; waves of teenagers whose bodies have become so inactive from being immersed that hospitals have to sever their limbs or something. I'm very much hoping that it isn't that bad but the thing about reality is that it's not fully predictable.

We have to dive in and learn. I think it's important to make experiments, and it's really important even to experiment in a way that can be a little dangerous. It's important to take risks. The thing that's really a sin is to not learn from them, to ignore the results. That's when you really lose it. As long as we're awake and paying attention and we learn and get better, that's what matters.

"As long as we're awake and paying attention and we learn and get better, that's what matters."

KK: Our sense of history in this world is very dim and very short. We were talking about the past: VR wasn't talked about for a long time, right? Thirty-five years. Most people have no idea that this is 35 years old. Thirty years later, it's the same headlines. Was the technological power just not sufficient 30 years ago?



The Nintendo Power Glove, based on a VPL dataglove design

J: Both I and a lot of other people really, really wanted to get a consumable version of this stuff out. We managed to get a taste of the experience with something called the Power Glove. Remember that? You could put on this big glove and reach into things on the screen. Sony actually brought out a little near-eye display called Virtual Boy; not very good, but they gave it their best shot. And there were huge projects that have never been shown to the public to try to make a consumable [VR product], very expensive ones. Counting for inflation, probably more money was spent [than] than Facebook just spent on Oculus. We just could never, never, never get it quite there.

KK: Because?

J: The component cost. It's Moore's law. Sensors, displays... batteries! Batteries is a big one.

KK: So the vision was there, but the technology wasn't.

J: It wasn't there. Whether it's there [now], we'll find out in the next year or two. We're sort of saying it's there but nobody's actually proven it yet. Proof is in the actual doing.

KK: You can't discount, 20 years from now, some headline saying, "Okay this time it's real, we're looking at VR in 20 years."

J: I wonder if the reason we keep on cycling back to hope about cool things like VR is that for all the tech news and our fetishizing about our touch devices, we're still a little disappointed in the menu of tech items that we have at this late date. It's 2014 and you can buy a robot to clean your house, but it doesn't really work that well yet. We all think it will, some day, but it's not quite there. We have some demos of cars that drive themselves but you can't really buy one. Everything is kind of taking so long. So I feel like, in a way, we keep on cycling through the same tech hope stories because there's an impatience and frustration. We wait until we've forgotten one of them, then we rediscover it. So I think there's a little bit of that going on.

**"I feel like, in a way, we
keep on cycling through the
same tech hope stories
because there's an**

impatience and frustration."

C: Reading the original interview, one of the things I found so exciting was your excitement about VR as a way to temporarily escape the physical bounds of your body, to escape all reality. I'm not sure that you would characterize it as an escape at the time, but reading it I just thought, my god, what a different way of being.

J: If you were to try to enumerate what was actually invented at VPL, there are a few different things. Arguably, the most important one was avatars, because nobody had experienced that before. There were representations of people on the screen, but as far as being immersed in an alternate body and being with other people in a simulation, there had never been a social, immersive simulation before. That was such a striking and remarkable experience. To this day, it's still very rarified. I've noticed there is a revival of experiments in which people look out at each other's eyes and that sort of thing. There are a bunch of videos about that online, which is one of the things we were really into in the old days. As far as becoming a full-bodied avatar and being with somebody else, seeing each other as avatars and interacting that way, that's still an incredibly rarified experience. I don't know why it should be, because we have all the tools now. Just put your damn body in front of a Kinect and put on the Oculus and you should be able to do it. But for some reason, people haven't rediscovered that, and when they do, there are going to be blown minds because it's just so intense. It's just really amazing.

KK: I know that that's what High Fidelity is working on in some ways. What they're doing is capturing all the emotional contortions of your face so they can actually present that to an avatar.

J: This is one of those moments where I feel a bit of jerk saying, "Oh yeah, I've been doing that for decades," but yeah, I've done that a lot. And yeah, it's absolutely true. In fact, face-tracking and mapping the avatar's face is absolutely a core thing. That first started to work in the '90s, and I and some friends had a little startup that was doing it. I suspect [that] there are a few code packages that can do that and I have a feeling that they might be using our old one because a lot of people do.



"This is one of those moments where I feel a bit

of jerk saying, "Oh yeah, I've been doing that for decades," but yeah, I've done that a lot."

If you can capture as many nuances of your body as possible and transfer them into the avatar, every extra parameter you can measure enhances it in a synergistic way where it gets cooler and cooler. The other thing that's really interesting is to see how weird an avatar can get and still be you. The classic early experiment was the famous lobster avatar that Ann Lasko made. Jeremy Bailenson over at Stanford has [tried] super minimal [designs] where you just turn into a cube you can stretch. The interesting thing is even if you just turn yourself into a cube, you can still convey affect and attention and all sorts of things.

KK: Ask any animator at Pixar who does this all the time. They take a rug or a carpet and they make it into somebody with a personality and a persona.

J: For you to become that thing in real time is something different, and I think a lot of people will find this to be a really delightful zone of experience.

KK: Also, multiple people can share an avatar in interesting ways. The other thing is in a virtual world, this question of which is the avatar and which is the environment can be shifting over time. All kinds of things can happen: you can turn into the clouds and what not. It gets very, very interesting and very fun. Are you still working in some of these different directions, even at Microsoft? With the various elements — whether it's avatars or face recognition. What's your role now?

"Hey look at Microsoft Research; every bad thing that can be said about us has already been said, we don't care what you say."

J: One of the reasons I ended up at Microsoft Research is incredibly simple. When Google bought Eyematic [Interfaces, which developed ways for computers to understand images], at one point Sergey Brin told me, "At Google we're just not that comfortable with people being out there [and] having outspoken opinions; we really don't want people to be sort of blogging about this and that." I was like, are you kidding? That's what I do. It was at some event and Bill Gates was there too and he said, "Hey look at Microsoft Research; every bad thing that can be said about us has already been said, we don't care what you say." It's like, well, that's kind of cool. It's given me a chance to be part of a large-scale, great lab without the constraints that Silicon Valley companies often put on people. But the other thing that was magical about Microsoft was Kinect. The notion that you could, using depth cameras, turn into an avatar and interact with worlds.

KK: So are you continuing [work on Kinect] or are you doing anything new in that direction?

J: I might not choose to fully disclose everything I'm doing right now because that's part of the fun. But I do a lot of stuff. I've always been a pretty promiscuous technical person. I have one collaboration with a cosmologist on a digital model of the earliest moments of the universe. I'm trying to come up with an alternative to inflation, just to see where it goes.

In terms of VR? Sure, I'm still working on it. I'm still really interested in all the things that I've always been interested in — like where can optics go if you have computation? Where can sensors go if you have computation? All that kind of stuff.

KK: That's probably what we're going to see in the next 15 years or so; continuing [to develop] the variety of sensors that can generate information for a virtual world.



A SimEye helmet-mounted virtual display (Ben Delaney)

J: We seem to be able to keep Moore's law going, but we haven't quite seen a battery's Moore's law keep up. Energy technology is probably going to be more the impediment than

anything else. And also there is the tooling cost; to keep Moore's law going we might be moving to very different kinds of material. If we were seeing something as dramatic as a Moore's law-level exponential improvement in batteries, we would happily explain it as a sort of "know-how" improvement based on information access. The only problem is that we're not seeing. The problem with energy is it's closer to the constraints of reality; it's a little further from our fantasy worlds. I'm still optimistic that we'll figure it out. I'm optimistic that we'll come up with better ways of generating and storing energy, but it's an actual hard problem.

KK: Right, exactly. If Moore's law was to cease or slow down, it would have tremendous affect on our lives because right now; we expect things to get cheaper and better all the time. If that wasn't true, if your little devices weren't getting cheaper and better all the time, that would have a huge hit on the culture and the economy.

J: It would be giant. Well, it has to happen at some point. Reality doesn't have infinite resolution. There is going to be a Moore's law limit. The question is, where it is and when is it? One of the things that Microsoft Research is particularly great about is exploration at the edge of Moore's law.

"Reality doesn't have infinite resolution. There is going to be a Moore's law limit."

And there are a few of them right now. One is trying to use

silicon in more flexible ways. Doug Burger, a colleague there, just demonstrated using field programmable data rays as a cloud architecture so you can just reconfigure the whole cloud all the time instead of having a fixed processor design. The benefits are really amazing; it's going to totally change what we expect from clouds. That's an example of noticing an opportunity for more flexible architecture so that you can simulate a jump in Moore's law even though you're still stuck with the basic stuff.

KK: What happens is that we just redefine Moore's law. We've been doing that all along.

J: Of course. That's a game that you can play.

KK: Like cameras: for a while, the number of pixels was important. Then we realized, omigosh, at a certain point, the number of pixels isn't important, it's the speed.

J: In the next few years we're going to start to really understand what we can do with quantum computing. We're so close now, empirically, to getting to something on that. I'm pretty enthused about it. I think what it'll look like is a cloud service that can do more than we could have plausibly expected it to do otherwise. I think we're going to start to see much better machine learning. But all of it is going to be module — all the same latency and clogging and political issues that darken our cloud services now... none of that is going away with technology.

KK: Going back to this issue of the reality of virtual things and what we might want to think about as we make these worlds larger and more encompassing and more persistent; have you had any ideas about making the distinction between how we treat things and who gets to decide about them?

J: I have been part of a lot of conversations about what the laws should be and trying to come up with regulations. How do

you protect the kid who's being bullied without impinging on freedom of speech? How do you prosecute revenge porn without empowering some politician? Or the right to be forgotten — how you do that without empowering some politician? It's these kinds of discussions that led me to become more interested in Ted Nelson's original ideas about micropayments. In a lot of ways, if you can make an economy adjust for these sorts of things instead of adjudication with rules, it just works better for everybody.

KK: Really?



J: Yeah, I think so. For instance, if there's a grocery store that has some weird tart that I want, instead of arguing with people about whether I should have it, if I have money I could get it. It creates a very simple system where there is only one varying parameter: price. It removes, in most cases, all these legal things. There is a lot to be said for it — it's a good simplifying principle for human affairs. That's why it exists, you know?

KK: It seems as if you're saying the market is going to solve everything.

J: Not everything, but maybe some things.

KK: Okay, so let's say the right to be forgotten; how would having a market and people's money solve that issue? If you pay enough, you erase [them]?

J: No, I think if people can set a price for their information and have a single parameter adjustment for how private they want it to be and people can find points of equilibrium.

KK: So you pay for privacy.

J: And you also earn for it. It would be a balance point. There could be people who say: I like living in a barter world. I will give up my privacy by setting the cost of my information to be really cheap: have at it. It's the bargain I want, and I like it. And somebody else might say: actually, I don't like being modeled. I don't like being manipulated. I don't like being known. I want to be off the grid. I'm going to set my price really high, even though that means not many people are ever going to pay for it.

But most important, the thing that I want people to understand, is that if there were a society like that — and this is speculative without test of course — [it] would put a price on government spying on people, too. We'd put a check on it. And, of course,

law enforcement would be able to get warrants to intervene in some cases. There'd be all the adjustments that one imagines, but it might very well work better.

This idea of openness and prioritizing free speech is wonderful in theory, but it has two practical flaws that have emerged empirically that I certainly hadn't predicted theoretically. One of them is ever more income concentration, because nobody benefits from information in an economic sense, except for the pleasure of the people who have the biggest and most influential computers. The other thing is that at random, people are victimized by it very terribly. An example would be revenge porn. Another example is people whose personal information was used to target them with effective but malicious financial offers and practices.

KK: I certainly endorse Esther's idea of having to pay people to read your email.

J: That's a great example. For people who don't know the story: Esther Dyson proposed, at one point that if we put the tiniest postage, a micropostage, on email, it would solve spam.



A CAVE virtual reality installation at the Ars Electronica Futurelab

KK: But it's important that the postal payment goes to the person, the receiver; that's who's getting the money. It's not the post office.

J: Right, there's no third party. There's not some entity that's growing rich on postage. It's person-to-person in distributing.

KK: That's very important.

J: Absolutely correct. I think there is some kind of a setup where Mark Zuckerberg will read your email for \$100 or something now. So that exists in a couple isolated cases. But she was shouted down because now everything has to be open and free. The beautiful thing about that, is that for anybody doing normal communication, the cost is so marginal that it really shouldn't have any effect. And since it doesn't go to some third party, it doesn't create some new, centralized powerhouse. I

think it was an interesting proposal and I wish it had been tried more.

KK: I do too. The question is: how do you implement it? Either the whole system has it or it doesn't. It's very hard to implement incrementally. That's one of the issues about this world is that there are lots of things you can implement incrementally, but there's a whole set of things that you can't.

J: Micropayments are like that. If micropayments aren't universal, their benefits aren't spread around enough to generate support for it. Maybe if you could get over the valley, there would be. The only other thing I can say about that is because of Moore's law, there are always new platforms coming out, like 3D printing, virtual worlds. Every time a new platform comes about, there is a new opportunity to experiment. And there will be dozens of them.

KK: You can actually say it slightly differently: there are natural monopolies, and each time a monopoly comes up, that's an opportunity. Monopolies are much more ephemeral because they unravel almost as fast as they build up. Basically, we have a future of one natural monopoly after another. Each time one comes up, there is an opportunity for doing something that's ubiquitous. So there's hope.



Early VPL concept drawings of virtual reality

"Just to make people not become nauseous, you have to really get into it and become human-centered."

J: I think there is tremendous hope. The stuff that scares me most these days is the same stuff that would scare most people, which is whether we'll sort out some way to deal with climate change in time. And whether we can reconcile modernity and human nature in some workable, sustainable way. People are tribal and modernity wants us to be globalized. How do we reconcile those things? People are biological, but we want to be immortal. How do we reconcile those things? Those are the kinds of big-picture issues that are really scary and really challenging.

KK: And is VR a solution or just another part of the problem?

J: In the old days, I used to think of VR as a technology that turned technologists into humanists because you had to work with people so much. It's one thing to put stuff on the screen and ask people to project themselves onto the screen, get lost in it, but as soon as they are wearing the stuff, you just have to work with a human body; you have to think with human factors. Just to make people not become nauseous, you have to really get into it and become human-centered. You have to become super human-centered. You can't ask people to meet you halfway anymore. You have to really go to the people to get VR to work. So I always thought it helped you notice people and human nature more and it should make engineers more empathetic and more sympathetic to people. And I still think that's kind of true.

Among the post-idealist ideologues, VR has come to this completely different role where the great AI in the future will somehow recreate all our consciousness in VR. In the extremes of the movement, people are worried about. There's this idea that the virtual copies of your consciousness will be created in VR by the big AI in the future [and it] will torture your copies if it finds out that you didn't give all your money to some AI company to help bring about the singularity.

C: Roko's basilisk?

J: Exactly, that kind of stuff. There's a lot of variations of that. There's this sort of insane way that VR is being used to turn technology into a new medieval church of some kind. It's sort of like the worst of scholasticism in Catholic tradition or, maybe, the most nerdy and controlling version of Islam, or something like that, but from 1,000 years ago. That's certainly not what I was expecting. That's a little different.