LOOKING FORWARD

anced so that if there ever was a contention it could be managed.

Cloud computing customers want low latency message passing. Software managed page level memory coherency can provide global memory coherency when needed but not fine grained L2 coherency. Each tile can be run at a different frequency and voltage as well. There is power management over the I/O and memory controllers.

In addition to the high-speed on-chip network, the chip has power management techniques that allow all 48 cores to operate at as little as 25 watts when idle, or at 125 watts when running at maximum performance.

Intel plans to gain a better understanding of how to schedule and coordinate the many cores of this experimental chip for its future mainstream chips. To do that, the company has plans to engage industry and academia next year by sharing 100 or more of these experimental chips for hands-on research in developing new software applications and programming models. Professor Wen-Mei Hwu at the University of Illinois will be getting one, as will the parallel processing lab at UC Berkeley, Microsoft, and of course Intel's favorite lab in Saarlandes Germany.

In another meeting, Intel's CTO Justin Rattner asked, "Could you replace a rack of servers used for cloud computing with a many core, SCC?" No, he told us. This is never going to be a product. It is an experimental platform. It uses a simple in-order Pentium processor like an Atom core design. There are 1.3 billion transistors in the chip, and it was designed by 40 people, on three continents. There was only one significant bug due to the re-use of a design using an existing DDR3 memory controller and an inverter got added by the CAE program. It took a layer change to correct it, and according to Intel's researchers, the SCC seems stable now and they have booted several operating systems on it.

And, there is a third-generation design running in simulation in the Santa Clara lab now.

What do we think?

The obvious question is, is this a Larrabee prototype? Held says no for several reasons, it's not an SIMD construction like Larrabee; it doesn't have the MMX or any vector processing, and uses a message passing mesh not a ring. Also, the floating point processor in the SCC is just a simple x87 FPP.

The SSC represents the latest in a five-year effort in terascale computer research, and is the second-generation experimental processor. Intel saw the projection of Moore's law and believed the first experience with computers was textual, then with the web and multi-cores it became visual, and that in the next era making use of many cores will be social, citing gestures, and the clothing example mentioned. "Machines," said the affable Rattner, "will be capable of sensing the world around them, much like humans do." Is Rattner suggesting the singularity?

Intel will be disclosing more details at the upcoming February ISSCC meeting.

Canesta's 3D sensor

Seeing the way we do, but with one eye

Founded in 1999 in Santa Clara, Canesta has developed a sensor and associated processing algorithms they call a new electronic perception technology that promises low cost devices (mobile phones, PCs, gesture controls, etc) to "see" and understand their environment as objects. Canesta's single-chip solution forms three-dimensional images of its nearby surroundings in real time.

Perceiving specific objects with just a 2D sensor array of a million or two pixels is a profoundly difficult problem. Thousands of algorithms have been tried in the last 50 or so years that have attempted such identification, and usually depend on either multiple cameras and controlled light sources, or some advance knowledge of the elements or motion in the scene—and all the techniques require substantial amounts of computing power.

This is Canesta's secret sauce. The company has made a breakthrough in this regard with its development of a low-cost electronic perception system. The technology includes a new type of chip-based image sensor, similar in size, complexity and cost to the commodity-priced video camera chips. The difference is that Canesta's design is able to resolve the three-dimensional features of a scene. That is, in addition to the brightness of a specific color of light reflected from objects, the distance from the illuminated object to each pixel is also determined. Since each pixel in the sensor is illuminated by a different feature in the scene being viewed, the result is a true 3D representation. In addition, the sensors can operate at over 30 fps, making real-time applications possible such as automotive sensors that can accurately judge the size, shape, and position of a passenger, security systems that identify individuals from a 3D view of their face, or game player in front of a TV or computer. Canesta says unlike distance data derived mathematically from 2D cameras though the use of geometric algorithms, distance and size are now first order parameters, measured directly by their sensor. This, says the company, leads to a much more straightforward interpretation of the scene, free from assumptions that are implicit in distance data derived from 2D cameras.

Time of flight

The light illuminating each pixel in an image sensor comes from a different feature in the scene being viewed. Canesta recognized that if you could determine the amount of time that light takes to reach each pixel from the corresponding feature in the "scene," you then could then calculate with certainty the exact distance to that feature. In other words, you could develop a three-dimensional "relief" map of the surfaces in the scene. In 3D, objects previously indistinguishable from the background, metaphorically "pop" out. This approach proves extremely helpful in reducing the mathematical and physical complexity that has plagued previous attempts in 3D computer vision.

In Canesta's U.S. patent (#6,580,486) entitled "Methods for CMOS-compatible three-dimensional image sensing using quantum efficiency modulation," The company describes several of its inventions for "timing" the travel time of light to a unique, new class of low-cost sensor chips.

Fundamentally, the chips work in a manner similar to radar, where the distance to remote objects is calculated by measuring the time it takes radio waves to make the round trip from a transmitting antenna to a reflective object. In the case of these chips, however, an IR source is used. Canesta says the chips are not fooled by ambient light.

The result is an array of "distances"—updated as often as 60 times a second or more—that provide a mathematically-accurate, dynamic "relief" map of the surfaces being imaged. The image and distance information are then handed off to a local processor running Canesta's proprietary perception software that further refines the 3-D representation

LOOKING FORWARD

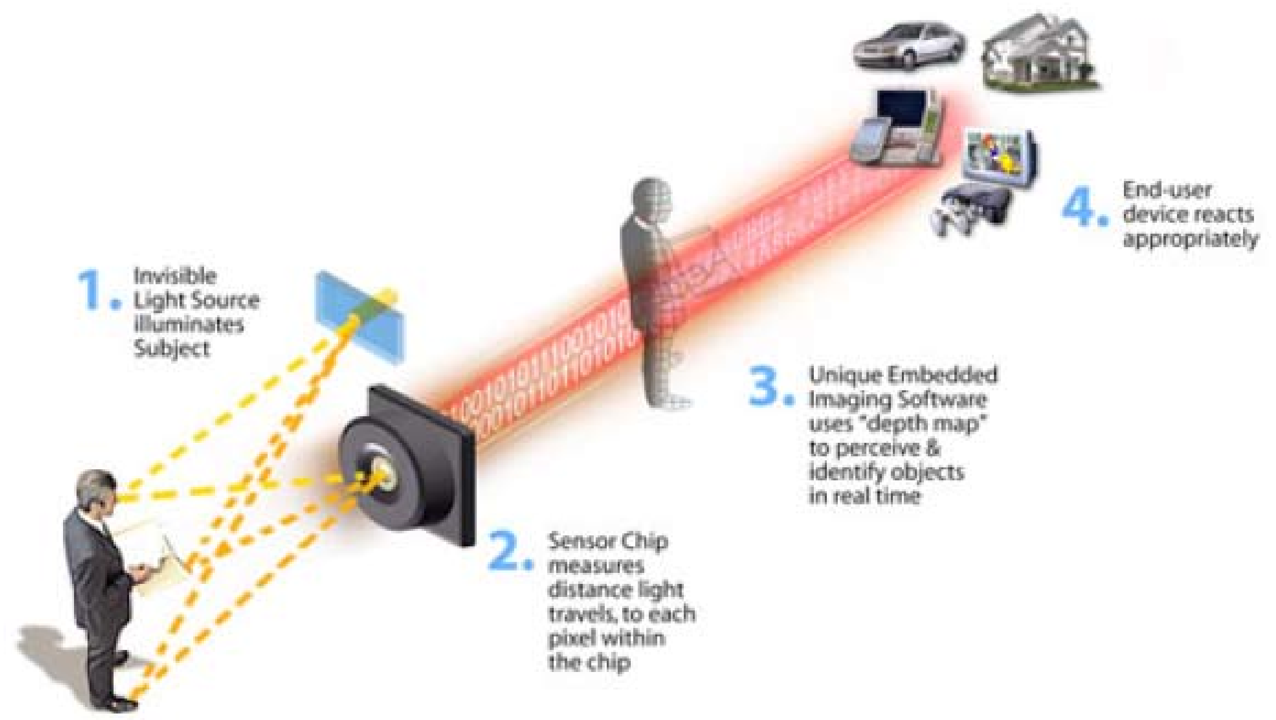

FIGURE 3: Canesta's 3D sensor. *(Source: Canesta)*

before passing it off to the application-layer software, which identifies objects in the scene so the vehicle or device can understand or react as appropriate.

The company has had some heavy hitters investing in them, including Quanta Computer, SMSC, and, Honda, plus VCs like Venrock, Caryle, and Korea Global. Canesta has had four successful funding rounds and has raised $60 Million to date.

What do we think?

Products using this technology should show up next year. In the discussion with Intel about their many-core terascale processor (see this issue) one of the examples of exploiting the massive compute power of the chip is to use a 3D sensor with a laptop for gesture control and/or amusing things like modeling clothing.

And we think this will also open up new opportunities for augmented reality (see story this issue.) Low-cost 3D sensing is one of the things the industry has had a pent up desire for and the Nintendo Wii has heightened that desire. Microsoft is rumored to be using Canesta in their Natal project, and we can envision all types of automotive and safety applications. This looks like a winner to us.

> **FAST FACTS**
>
> - Augmented reality (AR) is a term for a live direct or indirect view of a physical real-world environment whose elements are merged with (or augmented by) virtual computer-generated imagery—creating a mixed reality. The augmentation is conventionally in real-time and in semantic context with environmental elements.

Augmented reality: Total Immersion

Turning the science into consumer delights

Founded in Paris, France in 1999, Total Immersion has been one of the early developers and researchers in augmented reality (AR), Total Immersion has helped auto and aircraft makers visualize their products, car mechanics fix engines, and theme parks surprise their visitors. But they were mostly one-off projects, and sometimes a science project.

Total Immersion asked themselves how they could offer augmented reality to average consumers rather than just professionals. So, after several years of research, the Total Immersion team identified online and mobile as the two best market opportunities.

The team then started developing novel ideas for how a consumer could engage with AR and, in so doing, they discovered they could offer a new experience or relationship for suppliers with their customers—the product becomes part of the show—e.g., the can of Coke becomes a joystick.

One of the examples of a consumer application involving a product is Total Immersion's work with the new *Star Trek* movie (http://tinyurl.com/yzzkwav). And the company has done a deal with Mattel on a line of new action figures based on the movie *Avatar*, and associated trading cards that have 3D figures that come out of the card.

The company also has deals with Coca Cola and McDonalds.

How it works

Total Immersion solutions merge the virtual and real worlds together. Shown below are what is typically required—a video camera, a display, and a standard PC hosting Total Immersion's software. The target can be collateral, trading