


Interview with Ron Buck from TYZX

Ron Buck is president and CEO of TYZX. He has more than 20 years of senior management experience in the computer industry, including several successful start-up ventures. Previously, Buck was president and CEO of Verano, a company providing Internet-based supply chain management solutions. He organized Verano, secured two rounds of financing, and successfully managed the company through to its first revenues. Prior to Verano, Buck was vice president of marketing and business development at nCUBE, the first provider of interactive multimedia streaming services. During his four years there, he managed several high-profile trials of multimedia streaming systems and grew nCUBE's revenue from \$4 million to more than \$34 million. Prior to that, he spent six years at Pyramid Technology Corporation (later acquired by Siemens) in senior development and marketing roles. <http://www.tyzx.com>



Please give us some background information about TYZX. TYZX 3D technology is based on more than 15 years of research conducted at Harvard, Interval Research, MIT and Stanford. We have developed a platform of hardware, software and services for building affordable products that see and interact with the world in three dimensions. And now, we are delivering high volume, cost-effective 3D vision solutions to the automotive, consumer electronics, robotics and security markets. TYZX was founded in 2002, and is based in Menlo Park, California. We're privately funded, and last year we received a large investment from Takata (Japan), one of the largest manufacturers of complete automotive safety systems in the world.

The name of your company is a bit unusual; is there some meaning behind the name? The name "TYZX" is an acronym of the dimensions used to compute and measure 3D vision, specifically: **T**=Time, **Y**=Height, **Z**=Depth, and **X**=Width. Also, we like the name because it rhymes with "physics".

Please provide an overview of your technology. TYZX is about making "systems that see". That means enabling computer-based products to visually understand and interact with the things they are looking at. This might be a car seeing an obstacle and braking, or an entertainment system following a hand gesture. Whatever the visual application, it is easier to develop and have it perform more reliably if it sees in 3D, i.e., has a sense of depth perception like we do. Computer vision researchers have understood this for a long time, but computing depth from a pair of stereo images was too slow and expensive to be practical. Now, by using the same commodity CMOS imagers we all have in our cell phones and custom ASICs and software developed by TYZX, real-time depth perception is possible - and affordable - in consumer products.

Tell us about "Woodfill's Law". John Woodfill, Ph.D, is our CTO. He invented the Census stereo algorithm (together with Ramin Zabih, Ph.D.) and is our guru in all things stereo. His observation is that stereo computation gets harder out of proportion with respect to bigger images. Specifically, if N is the number of rows or columns in an image, the amount of computation required is order N^3 . It means that stereo computation on larger images will continue to outstrip the performance capacity of conventional CPUs because it increases faster and therefore custom hardware is appropriate.

Your product portfolio falls under the "DeepSea" brand; is there some significance to that branding? Nothing other than the pun on "Seeing Depth" and the fact that the CEO's hobby is scuba diving.

How does your SEER software work? The SEER umbrella covers a variety of functionality. At the most basic level it includes the libraries and tools you use to set up and control the DeepSea hardware components and how you want the results returned to your application. It includes facilities like high-speed data capture allowing users to capture source and range data at frame rates. There are related applications like SEER Calibrate that are used for camera calibration (if users are experimenting with lenses in their cameras). Then there is the category of 3D middleware – APIs like ProjectionSpace and PersonTrack that provide 3D services on top of basic 3D depth maps

so that developers don't have to re-invent a computer vision application from scratch. SEER is available under Linux and WinXP, and is also designed to make the transition from our development system to our embedded system relatively easy.

Tell us about your stereo cameras. We have spent quite a lot of time on the design of our stereo cameras to make sure they are accurate and robust. Currently we have the broadest line of stereo cameras available with 3 cm, 6 cm, and 22 cm baselines available as standard offerings. We offer both color and monochrome cameras and 75-degree HFOV or 45 degree HFOV lenses. We also do custom cameras for unique applications such as our 33 cm camera with an 11-degree HFOV C-mount lens. We've taken into consideration a number of system design issues. For example, our cameras connect directly to the DeepSea processor via a custom high-speed cable so that the raw image data doesn't need to pass over the PC's system bus, which would reduce its available bandwidth. Each camera stores its own unique calibration parameters in a small local non-volatile memory so that users can swap many cameras around between different systems and have it all just work. Our primary imager from Micron uses a full frame shutter permitting much better performance in moving or vibrating cameras. Our camera design has also been extensively temperature and vibration-tested to be both reliable and to maintain its calibration in harsh use. Of course, each camera is tested and calibrated at the factory before shipping.

Who manufactures your stereo cameras for you? Our imagers are from Micron and Kodak. We use a number of different lens suppliers, but otherwise we design and build our own cameras.

Please tell us about your development of stereo processors. We believe that custom ASICs are the best path to high-performance, low-cost, low-power stereo solutions. The DeepSea stereo correlation architecture is particularly efficient leading to very small die size, which is a significant cost advantage in volume. We also have some broad patents granted in the area. Custom ASICs are instrumental in allowing us to achieve our volume price/performance goals; 3D vision can't go mainstream without a cost-effective source of data.

Describe how the SEER software, stereo camera, and stereo processors combine to create your system-level offering. At a high level, we have two offerings: the DeepSea V2 Development System and the DeepSea G2 Embedded Vision System. The DeepSea V2 uses a stereo camera, a DeepSea stereo processor board and SEER to support an application developer's work at a Windows or Linux workstation. There they have access to keyboards and displays, and have full ability to experiment with new applications. Once they have an application they want to try in the field, they can move it to the G2 Embedded Vision system. The G2 is really a self-contained Linux stereo system. It combines the same camera technology, the same stereo processor and SEER software with a PowerPC CPU and DSP into one small box – about the size of a hardcover book. You hook it up to your standard Ethernet network and you're running. You can even run it with Power over Ethernet (PoE) so you have just one cable into the unit. The PowerPC and DSP are there to run your application.



The DeepSea G2 Vision System is a self-contained, intelligent, real-time 3D sensor designed for production deployment of 3D applications in security, robotics and interactive systems

An important consideration in the G2 system design was how to process the large amounts of 3D data we are generating with a low-power embedded processor. It's fine to send high-resolution 3D point clouds at your high-end Intel Workstation at 60 frames per second, but it's way beyond what an embedded processor can handle. Actually it's often beyond what x86 can handle. To begin to address this, we came up with ProjectionSpace. It is

also implemented in hardware in the G2 and performs a number of common, expensive processing operations on the range data before sending it to the embedded processor. Using ProjectionSpace and other hardware features with the embedded processor, we were able to more than double the frame rate of our tracking system as compared to a PIII x86 running with our stereo processor alone.

Can your SEER software be run in connection with other cameras and processing systems? Yes, the DeepSea V2 can run from images passed to it via the system bus, though there are often custom integration tasks that need to be performed.

Do you consider TYZX to be a company focused on developing 3D vision systems or a company that develops the components necessary to make such systems? We deliver 3D vision systems and know-how to our customers, who then use them in their 3D applications. We provide software technology such as PersonTrack to users who might use it to develop a security system, an entertainment system or something we haven't even thought of yet. If our customer's market is large and specialized, we are happy to enter into a components and licensing arrangement with them so they can develop the right product for their needs.


Where do you see the biggest market opportunities for your technology? Generally, we seek opportunities that take advantage of our real-time performance and low latency. If your application is OK with running captured video through servers overnight, then it's probably not interesting to us. Automotive applications have been good for TYZX - they require the real-time control and the advantages a vision-based system can offer. Person tracking for security, marketing and entertainment purposes is also a large opportunity. One that we believe has tremendous future potential is robotics.

It seems that real-time tracking would require a substantial amount of data processing power, at least when you are tracking many targets. In this regard, are there any limitations to your system? In general, no; we are not very sensitive to the number of people being tracked. However, it's true that real-time tracking requires a large amount of processing. Our approach to this problem was to effectively build a large, distributed computer with a network of G2s. Each G2 automatically learns its precise 3D position relative to the other G2s it shares space with and tracks everything in its local field of view internally. It then forwards just the person tracks to PersonTrack Moderator (the central control system). All Moderator has to do is combine the tracks from each G2 into a single view. So all the heavy processing is done in the G2s out in the network, and much of that is boosted by custom hardware.

Your focus seems to be on stereo tracking – how do you envision your technology might apply to stereo display systems? Most standalone stereo displays rely on assigning image regions to depth planes (one of 16 or 32 for example). Doing this assignment based on stereo range data would be fairly straightforward. Of course, it only gets really interesting to us if you want to do this in real time...




This is an implementation of TYZX technology in an agricultural environment, where a vehicle can move independently based on the onboard cameras and sensing tools. The image on the left is what the vehicle "sees" as compared to an actual image on the right



Here's a security application that enables constant locational awareness of individuals within a premises

In the past, you have hinted at implementations of your technology in such things as vacuum cleaners and lawn mowers. Tell us more. Have you demonstrated such devices? The application is autonomous navigation. Right now, our work in the space has more to do with NASA and DoD than it does with vacuum cleaners. But the work is directly applicable. The focus now is for TYZX to continue to drive down the cost.

In previous editions of the 3rd Dimension, we showcased your implementation for the Electroland Target Breezeway at the top of Rockefeller Center in New York City. This appears to be really fun – was it also fun to put together or were there some difficult challenges? It was an absolute blast. I visited the site during an update we performed (the only one in 18 months of continuous operation). No doubt the site was challenging; there were glass doors opening to the rooftop observation deck on both sides of the room and thousands of watts of high output lamps changing constantly – both tracking people and in random “full on” flashes. The lights would go so bright you had to squint – but we still tracked. Watching people interact with the display late at night when visitors were, ahem, less inhibited was really fun. Electroland did a great job designing the experience.



Installation of the TYZX technology at the top of the Rockefeller Center, both an actual photo and a graphical representation of the sensing arrangement

Describe how your team might typically help a customer install the TYZX technology. It really depends on the application. We make an honest effort to ascertain whether the customer’s application is reasonably suited to our technology and ask a few questions to help them understand the skill level required for them to do an implementation. We have some customers that take the system and do their entire application with little or no input from us. We have other customers that engage us to help them with their development. In these cases, they usually purchase a development system from us and capture some typical data they would expect to work with. We then spend a few hours consulting with them to review their data and walk through the high level aspects of their application system design. At that point they can usually determine what aspects of the development they are comfortable with themselves, and what aspects they would like to engage TYZX for.

What can we next expect to see from you? Better, faster and cheaper, of course! We continue to make improvements in calibration (we have one researcher on staff dedicated solely to this important task) for better accuracy and greater FOV. We will also be adding to our repertoire of 3D middleware, adding APIs — some with new hardware acceleration — to help speed our customers’ applications to market.

Please describe what you think TYZX will look like three years from now. Bigger, of course, and at the nexus of an exploding industry that is making the gift of sight a ubiquitous feature in all aspects of daily living from the car you drive to the robot that assembles your products and vacuums your floor. One of the great things about this job is you never know what cool new application is going to come in the door next.