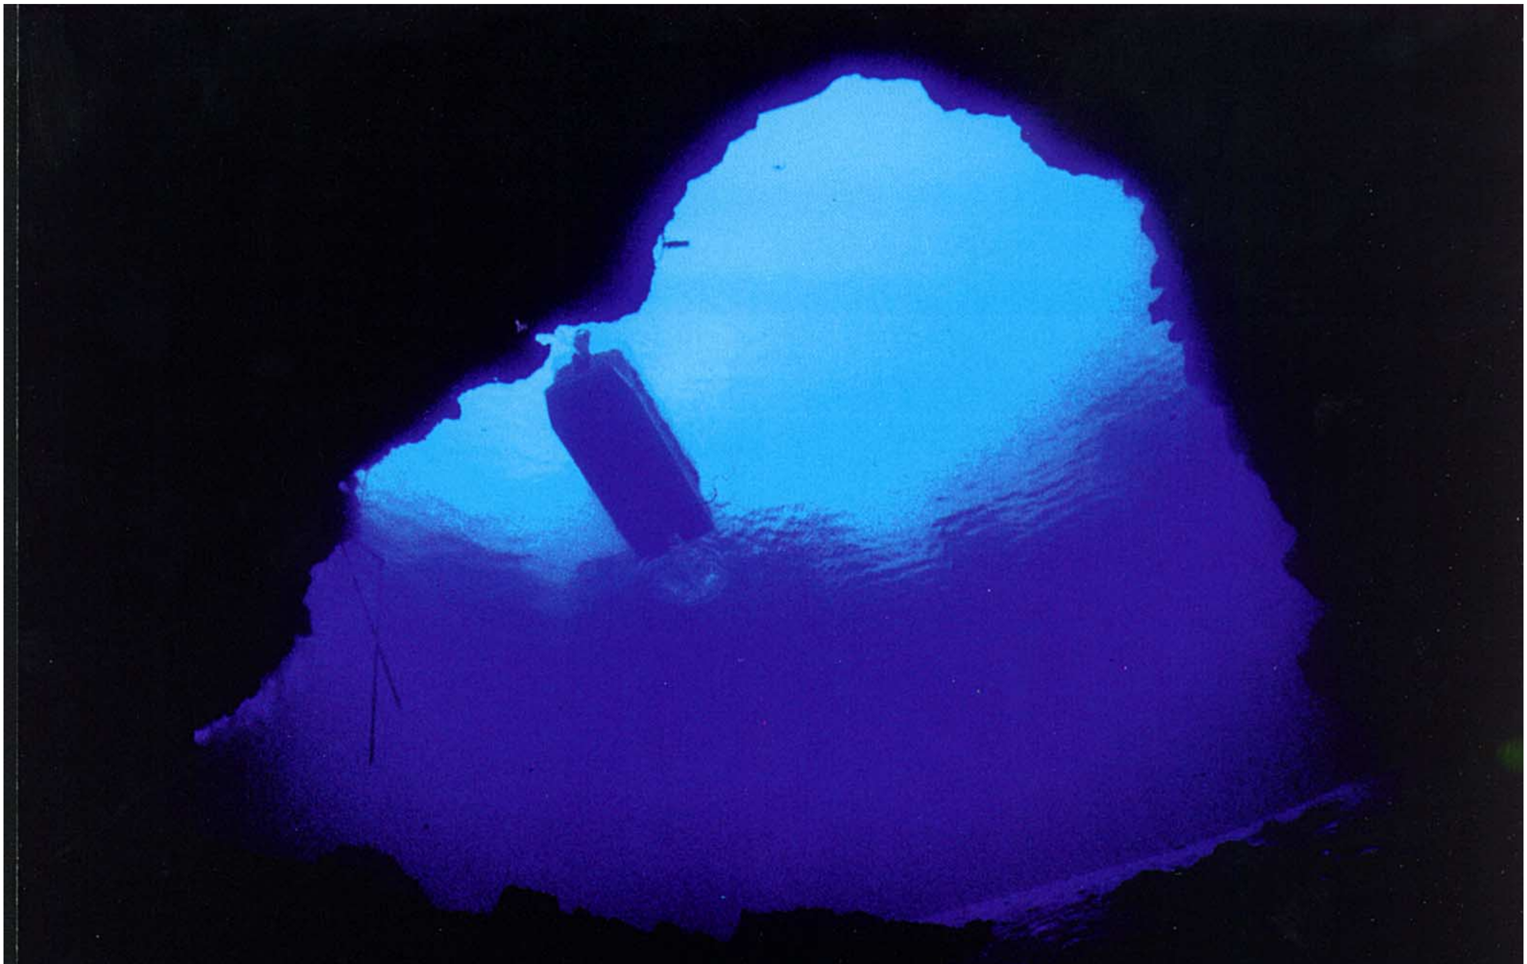

3D SONAR MAPPER

Wakulla 2 Team Tests the Latest Caving Technology

TEXT BY DR. BILL STONE AND DR. BARBARA AM ENDE

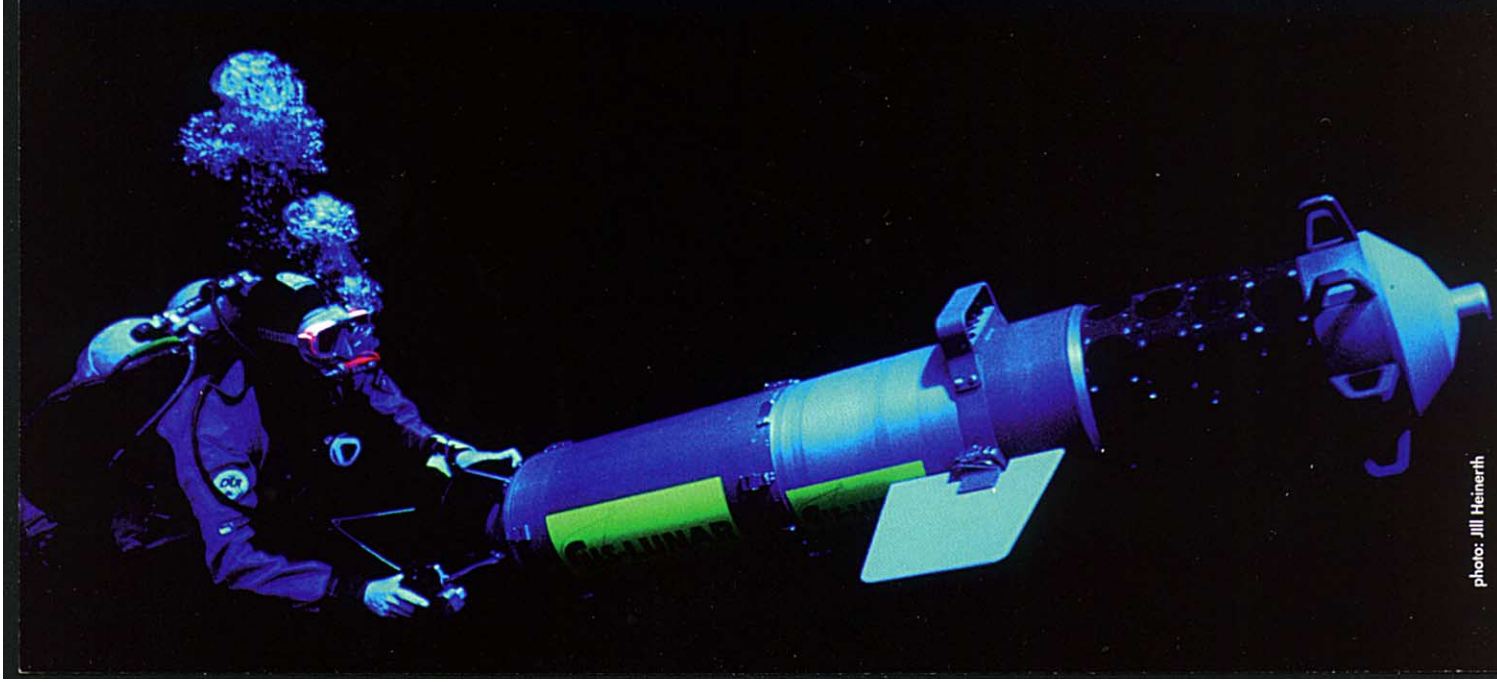

photo: Jill Heinerth

The concept has been around for a long time in the dry caving community: wouldn't it be nice to walk through an unexplored passage with a Star Trek-like scanner and later tour what you had mapped? And do it all in exquisite detail on your computer screen in the dry comfort of your home.

Anyone who has spent any significant time surveying under a pounding waterfall, swimming through underground lakes or crawling through mud-filled tubes eventually has this vision. The reality of cave surveying today is that you end up with, at best, a line plot and field sketch, with a great deal of artistic imagination being added to bring the map to life. This lack of precision persists in air-filled caves despite having the luxury of looking around, taking extra measurements and actually sketching plan, profile and cross-section views of the passage while still underground.

The situation is worse for cave divers. The limitations of life support equipment (even rebreathers) and the desire to reduce decompression don't permit time for a relaxed look around and a thoughtful, detailed sketch of the surroundings. Everything happens fast – usually at about 1.3 meters per second – the speed of the DPV. There is yet another problem with subaquatic speleocartography: you might not even be able to see the walls, floor nor ceilings, due to reduced visibility. During high water conditions in north Florida, tea-colored tannic water flows into the sinkholes and floods into springs reducing visibility in the underlying channelized aquifers to near zero. Sometimes the caves stay "tannic" for months.

It was with that in mind that Paul DeLoach, one of the leading cave divers in the U.S. throughout the 1980s, mused during the fall of 1990 that of all the gadgets he wished he could have, "Tannic Vision" was the highest on the list. He, John Zumrick and Bill Stone were standing at the edge of Whiskey Still sink at the time. Peering down into that 50-meter-deep (165 feet) shaft of brownish-black liquid, Paul's concept was easy to grasp: an electronic mask that somehow scans the tunnel ahead of you, even in zero visibility, and projects a computer-generated mesh ahead of the mask that represents the boundaries of the underwater tunnel as it leads off into the unknown. And while you are at it, why not store the data as well so you won't have to survey – since survey time just means more decompression. Well. What cave diver wouldn't want one of these?

In May of 1995, at the Cave Diving Section workshop in Branford, Florida, Jim King, John Zumrick, Wes Skiles, Paul and Jill Heinerth, Larry Green, Barbara am Ende, and Bill Stone, were discussing a return to Wakulla Spring in the fashion of the 1987 expedition – a three-month, full-time stay at the spring. In addition to habitats and rebreathers, the Tannic Vision concept re-surfaced. Could it really be done? It was agreed that it would probably be unrealistic to expect success with a mask-projection system at that time, due to the physical size of the required electronics. But it might be possible to build a digital 3D mapper that could be strapped to the nosecone of a DPV. That afternoon, in the parking lot of Nell's Diner, Jim King and Bill Stone were at the tail gate of Bill's truck, talking about project budgets. Bill was preparing to head north, after having spent the past two months exploring in Mexico. Jim was pensively rubbing his chin when he turned and said, "can you really build it?"

"The mapper,?" Bill said. "Well, yes, I believe so. But we aren't likely to find a sponsor to cover that kind of development anytime soon."

Jim allowed a wide Tennessee grin to spread across his face. "You've got one now," he said, "let's get going on it."

That single decision began a three-year odyssey into the complex world of autonomous mapping not only for Stone, but others who would soon join the team.

At the outset it was realized that whatever form the 3D mapper took, tremendous computational capabilities would be needed inside the device to keep up with the flood of raw data being generated as the device moved through the tunnel. That would require a large number of onboard computers that were all synchronized and communicating with each other. At the other end, such a ponderous load of data would have to be processed in some fashion to make it easy to manipulate and view.

The 3D cave maps of Wakulla Springs shown below are the first sonar-produced images of their kind. They show the entrance funnel, A Tunnel and ►

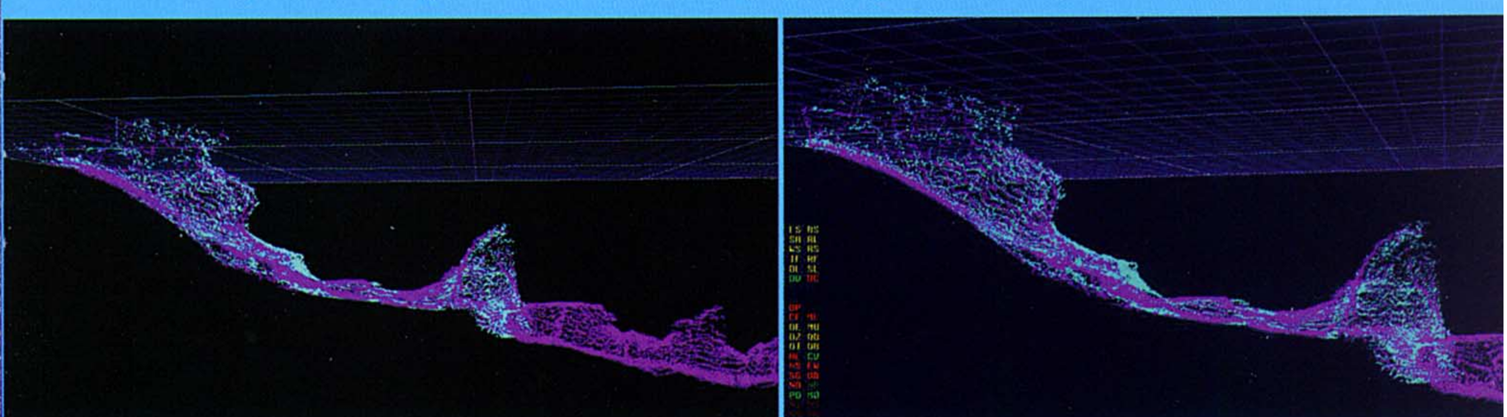

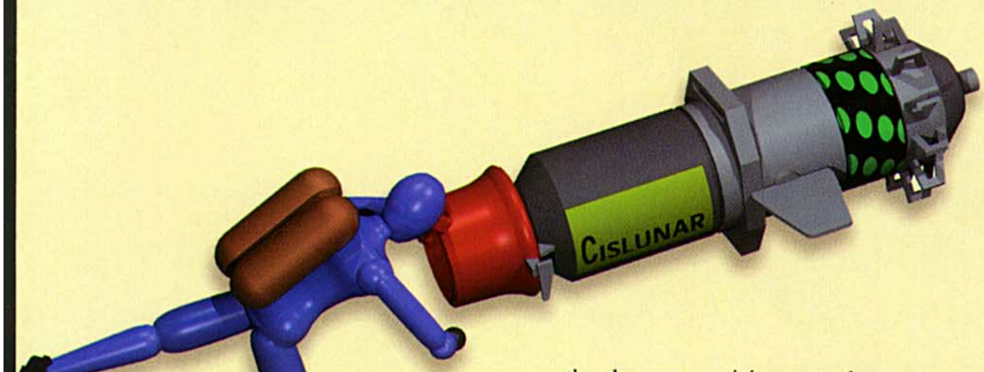

Artist's rendering of the Fat Man DPV with the mapper nosecone.

During the following year a number of people became involved with the project. Nigel Jones and Mike Stevens, from the Cis-Lunar MK5 development team, provided conceptual designs for the onboard computational hardware and software. Fred Wefer, from Mitre Corp., and Barbara am Ende, with the University of Maryland, began to investigate requirements for scientific visualization code to display the data collected by the mapper. Bill Stone worked up the first systems sketches for what the vehicle might look like and defined the required sensors. The objective was to achieve three-dimensional realism in the acquired data. Numerically, that meant tens of millions of data points to define the tunnel walls, ceiling and floors, just for the passages known at that time.

Each of those points would also have to somehow be "registered" relative to a surface benchmark – that is, we would need to precisely establish the longitude, latitude and depth of each data point on the wall. How do you do that? On the surface one can use phase differential GPS (and other methods) to track a moving vehicle with great precision. But GPS (or any other known precision metrology system - such as laser ranging) fails utterly if so much as a piece of paper is placed in the line-of-sight between the active elements of the measurement system. Take your hand held GPS receiver into a dense wood or into your house and you'll immediately see the problem.

Early on we recognized that there was no way we could individually record all of those wall locations, even if it were physically possible to do so – bottom time limitations, even using rebreathers, would preclude the development of a precise map. However, the process could be automated if it were possible to infer the wall coordinates relative to the mapper vehicle. To do this, two crucial things had to happen: (1) we had to know the location and orientation of the vehicle at all times and (2) we had to be able to sweep the cross section of the cave wall continuously as we moved through the tunnel. There were lots of other issues – how to control the system underwater (the diver interface); how to store the data; what environmental sensors to include; how to integrate the disparate electronic systems; how to "talk" to the onboard computers, etc. But it was the first two issues that really controlled the design since it was those two sensor systems that defined the accuracy of the overall device, and hence the final map.

Getting the passage cross section on a continuous basis meant using some form of sonar to ping the walls and return distances from the wall to the mapper. Thus, our original thinking focused on side scan sonar, since it was known that these could be used for imaging sea floor profiles on a continuous basis. All of the units we investigated had a limited field of view which measured no more than about one-third of the tunnel cross section on any given scan cycle. And so the concept of using a rotating side scan head came about. It was this concept that was presented to the State of Florida as part of the formal proposal for the Wakulla 2 project in early 1996.

Meanwhile, there was still the issue of registering the sonar distances to a known location.This meant knowing the exact coordinates and orientation of the vehicle at all times as it weaved through the cave passages. In aerospace terms, we needed to track the vehicle trajectory and attitude. A device for doing this, without the need for any other references (such as radio location), is known as an inertial measurement unit (IMU). Early versions were developed for the autopilot systems of ICBMs. We began investigating IMUs over the following year and finally settled, in the fall of 1996, on a ring-laser, gyro-based device since this offered the best accuracy within the allowable size constraints.

During this time calculations were performed which showed that the use of sidescan sonar led to significant problems with "smearing" of the data, due to long update times between each scan cycle of the sonar head. A typical sidescan unit will sweep a 120-degree swath of ocean floor in about 15 seconds. That's just a third of a passage cross section. To go all the way around, you have to re-position the sonar head (by rotating it), settle down, and then do the next scan. By the time you've gone once around the clock it's been a minute or more. The problem is that during that time your DPV has moved around 240 feet down the tunnel. The result is a helical data stripe that conveys little information concerning the cave. Now you have a real

illustration: Curt Bowen

the Grand Canyon Dome. The grid at the top represents the surface and each square denotes a 100-square-meter (330-square-feet) area. The smaller

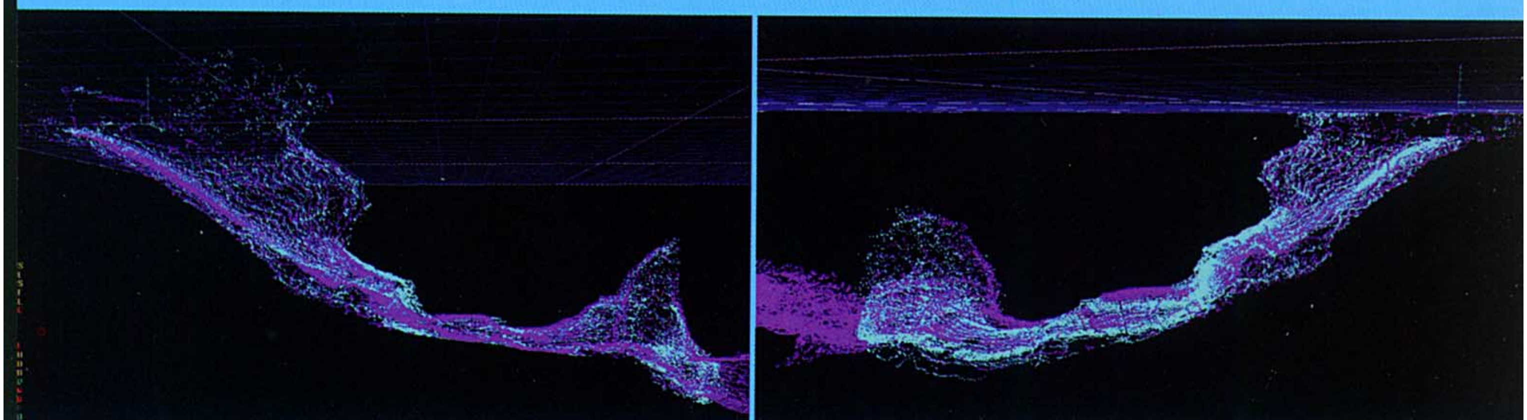

dilemma: to create 3D realism, you need passage cross sections about every foot. And each cross section must be referenced to a United States Geological Survey benchmark on the surface for accuracy.

Over the next six months many alternatives were considered for the sonar cross-section capture system. Then we hit upon an analogy to high precision planetary mapping. A satellite is placed in orbit around the planet and uses a high-accuracy point ranging system to measure the distances to the planet's surface. When the satellite makes enough orbits, and given that the planet rotates a slight amount on each orbit, you ultimately end up with a high definition map of the entire planet. In our case, we did not have the luxury of having the cave rotate about the mapper, but we could create multiple data streams from a large number of precision instruments, each of which could be processed independently in real-time. From this concept we developed a design for a massive phased array sonar system that consisted of a helical array of 32 focused high speed sonar transducers. These are spirally wrapped at 11-degree increments around the hull of the mapping device. Driving this beast was a computational nightmare. In the end we wound up with eight onboard computers, all seamlessly linked together and cranking out the equivalent of five Pentium Pro desktop machines.

By mid-May 1997 the sonar subsystem was being fabricated and the main pressure hull had been designed. About this time power estimates for the mapper electronics came in. These indicated that the original battery system (a traditional lead acid gel design) would not be sufficient to power both the mapper and the DPV to any significant range. We had, prior to this point, been discussing the use of nickel metal-hydride batteries. These now became imperative in order to maintain the 5-hour range that we wanted for the mapper. Because of the significant change in battery weight and density it was decided at that time to develop a new, optimized DPV designed around the new NiMH technology. The new DPV, owing to an oversize cross section, became known as the "Fat Man." The physical machine came in at 140 pounds with a 5.6 hour burn time at 240 feet per minute. It could be used either as a stand alone DPV or as the propulsion system for the mapper. Four Fat Men were built for the project.

By the summer of 1997 the sonar subsystem was up and running and tests were conducted at the Space Systems Lab Neutral Buoyancy Facility at the University of Maryland. These conclusively demonstrated the ability of the phased array sonar to capture passage cross sections in real-time. We were able to push the system to 12 Hz update rates while still maintaining wall point accuracy to 1.8 inches – that is, we completely defined the perimeter geometry of the test tank 12 times a second. We also showed that by moving the sensing pod around with an overhead crane we could indeed fill in greater levels of detail. This would have great significance during the actual Wak2 project since every vehicle passing down the tunnel would add new information, particularly if the paths were intentionally moved about to different sections of the tunnel each time. Fabrication of the remaining parts of the mapper continued throughout the fall of 1997. The first in-water tests were conducted in March of 1998 when it was

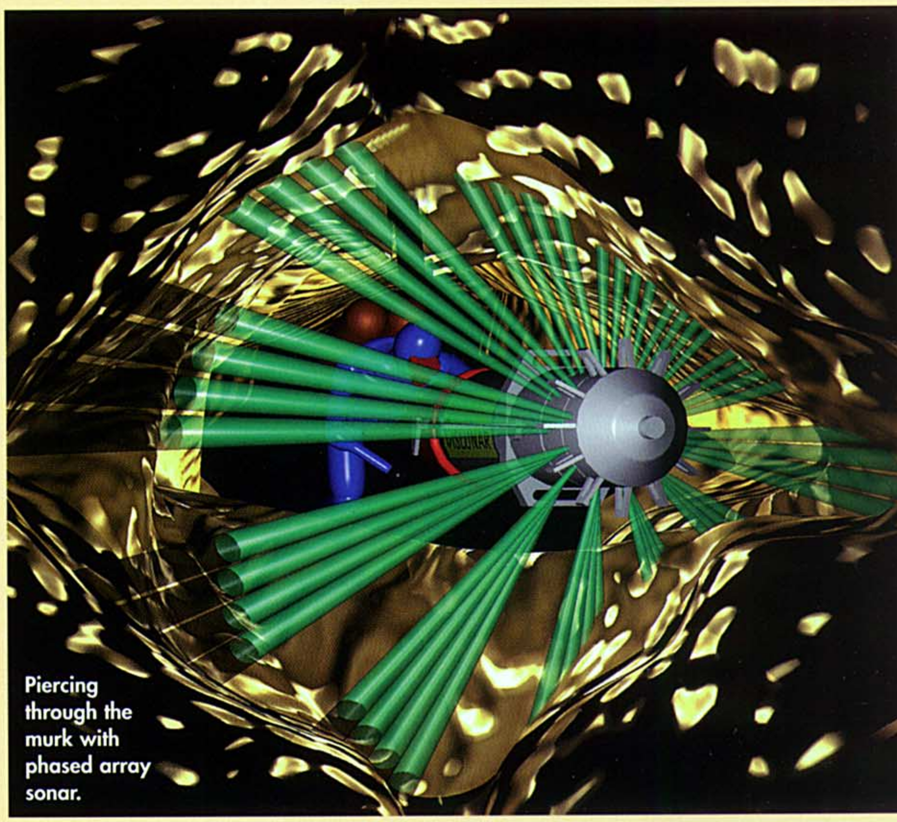

Piercing through the murk with phased array sonar.

illustration: Curt Bowen

grid represents a 10-square-meter (33-square-feet) area. The use of multi-colored points represent different passes made by the mapper, with the ►

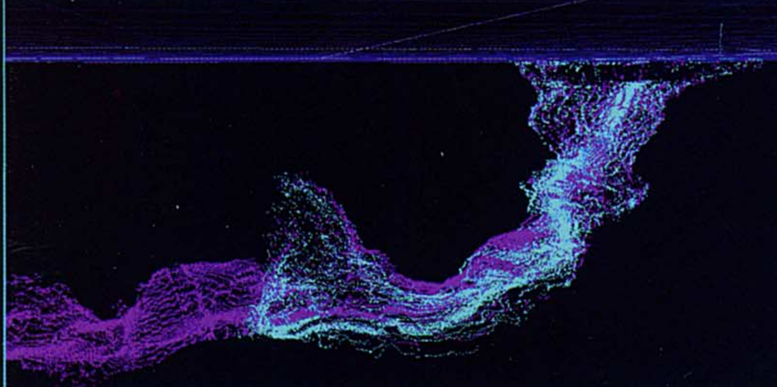

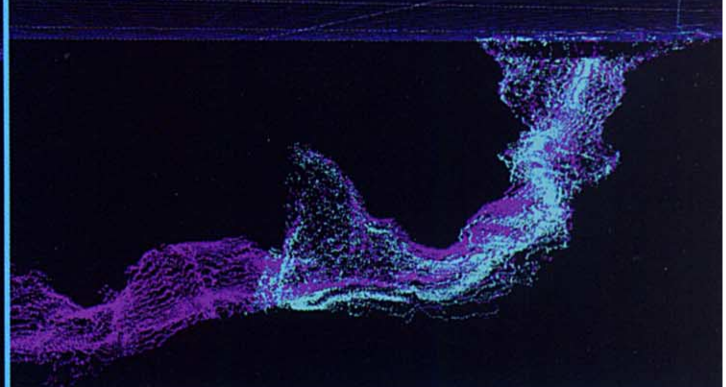

The mapper silhouetted against the sky in Wakulla Springs.

photo: Jill Heinerth

discovered that the vehicle was slightly negative and tail heavy – 320 pounds dry weight; about 10 pounds negative in the water. Replacement designs were installed and tested in late May 1998. With these changes the design goal of a neutral, non-rotating machine was achieved – you could let go of it and it would just hang there in the water.

In parallel with these tests, calibrations were performed on the IMU at the National Geodetic Survey test course in Gaithersburg, Maryland. This was a scene in itself: Jones and Stone with a laptop computer and a quarter million dollars worth of ballistic missile guidance technology in the back seat of a pickup truck cruising around the test course at 240 feet per minute (the speed the mapper would achieve underwater). One important reason for doing the test at this location was that accurate CAD models of the site were available as well as precisely determined benchmarks all along the route. Our runs were typically a half hour long and we would create a closed path back to our starting point. Over that time, because of the finite processing speed of the onboard computers and the resolution of the lasers and accelerometers, the predicted position began to drift to where, depending on how tight and fast we made certain turns, we might end up with several tens of meters of error when we returned to our starting point. The fact that it was only that much (over three kilometers or almost two miles of travel) without any other external reference, was in itself a marvel. Much of the resolution improvement was a result of a microcode change we made inside the IMU. Between each benchmark, the data were extremely good, even though cumulative error was slowly building up. Using this information, and the known locations of the benchmarks, we were able to develop an accurate drift compensation technique that brought the vehicle ground track precisely back into line. It did require the known location of at least one point approximately every 500 meters (1650 feet) along the path in order to maintain accuracy within the survey quality we were seeking for Wakulla, which was that no point within any given kilometer of tunnel would deviate from truth by more than a meter.

During the early development of the mapper, someone mentioned to Stone that he should talk with a fellow in the DC area, named Fred Wefer who is a long-time dry caver and computer graphics professional. Fred had written software on a Silicon Graphics computer that displayed caves in 3D. The data came from some dry caves he mapped in the Dominican Republic. Typical of dry cave surveys, his data was a line traverse with stations generally spaced between five and 30 meters (15-100 feet). At each station four wall distances were guestimated: left, right, up, down. He plotted the traverse line and wall points (four per station) in 3D space, then linked the points together into polygons. These polygons are the basis for most graphic displays. The polygons can be shaded, texture mapped or simply displayed as a wire frame.

Barbara, Bill and Fred spent a weekend together in the spring of 1997 plotting out a strategy for developing 3D software code for the mapper. That changed significantly as time went by. Barbara got the operating system set up to work with the various libraries and generated artificial cave passage wall points for testing the software. Meanwhile, Fred had the much bigger task of determining file formats and modifying his existing code to match the data stream collected by the mapper. To accomplish the job, 7 separate programs were written to: (1) separate the mapper location data (as recorded by the IMU), (2) convert the mapper location data to actual X, Y, Z coordinates in space, (3) thin wall point data to remove excess data that is duplicated under certain circumstances, (4) convert the distances measured by the sonar transducers into X, Y, Z coordinates, (5) adjust the traverse line (i.e., the mapper

data merged to form a more complete picture. Using the grid as a guide, it is easy to see the different orientation of each map relative to the surface.

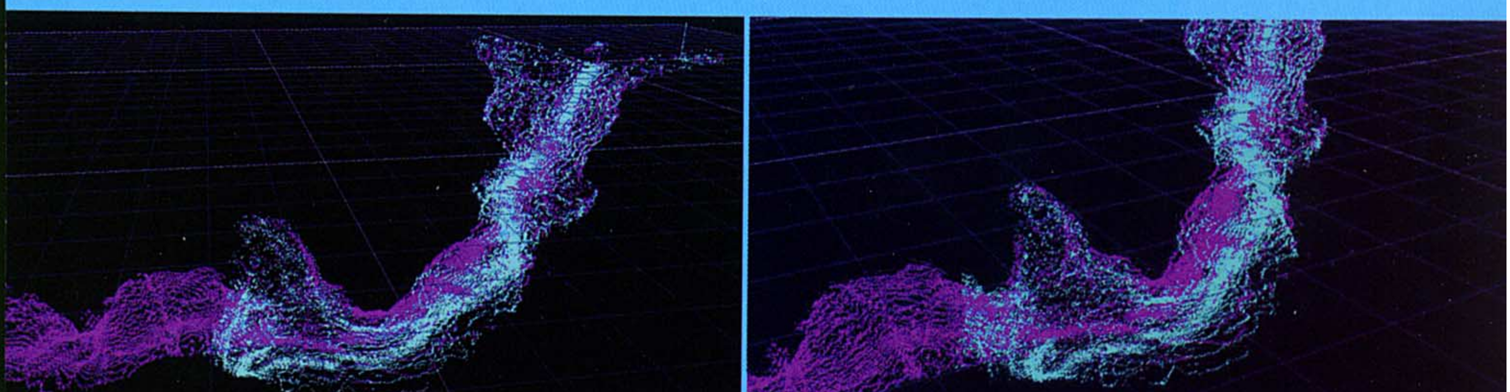

trajectory) and the wall points to compensate for mapper drift by using the fixed waypoints whose locations were determined by magnetic induction (aka "cave radio"), (6) mesh the wall points into triangular polygons, and finally (7) display the data as an interactive 3D cave map.

We had, since December of '95, known that "waypoints" were going to be needed to control the map. And we knew that "cave radio" could probably be used for that purpose. At the May 1997 meeting, Barbara and Fred proposed enlisting Brian Pease in the project. Pease, another veteran dry caver, was an electronics engineer from Connecticut who had a long history in cave radio location projects. He was particularly good at locating places inside caves where people either wanted a spare entrance to the cave (and didn't know what sinkhole to go looking in) or were looking to drill water wells (ranchers and farmers who owned the caves). The principles of cave radio are simple – a pulsed inductive (magnetic) field is generated and collapsed over time by a transmission coil located in the cave; a locator coil is employed on the surface to track into the center axis of the field, which turns out to be the point directly above the coil on the surface. While this sounds easy, there are only a handful of practitioners in the world who truly know how to make it work, largely due to a plethora of unusual noise problems that mask the real data. As an example, the locator unit we were using at Wakulla was so sensitive it could pick up a thunderstorm in Texas. Stone and Pease began working on an underwater version that could be easily deployed from a DPV. The prototype unit was ready and calibrated by April of 1998. A dozen have since been built for the project. One of these every 500 meters (1650 feet) inside Wakulla (they are retrievable on subsequent dives) would be sufficient to lock in the survey and correct any vehicle drift.

On June 13th, 1998 the various personnel involved in the project converged at Madison Blue Spring, Florida where Anna and Mike Bruic generously hosted our team for two days of testing. These dives helped us calibrate the mapper and also, crucially identify several inconsistencies between the provided and anticipated data coming from the mapper. A few very late nights by Fred and Nigel resolved these and the first tenuous data began appearing on the computer screen at the guest lodge. The underwater beacon was also tested and subsequently used to refine the existing map of Madison Blue Spring – once the management realized what it could do they wanted more data.

The evening of June 15th our convoy descended upon Chris and Kathy Brown's homestead in east Tallahassee. There we proceeded to fill the kitchen, living rooms and garage with techno-junk and computers. The garage served as our assembly station for the 320-pound mapper. Once assembled we dropped it horizontally into a pneumatic eight-wheeled custom transport carriage and rolled it into the van for trips to the springs. By pre-arranged permit, we started diving at the Wakulla Spring basin on June 16th. We first set up a control grid of surveyed points around the basin and down the slope leading into the cave, then swept out the basin with the mapper in a wide loop. This was followed by direct flights down into the giant funnel-like entrance to a depth of approximately 50 meters (165 feet). One problem that was identified early on was that one bank of four sonar transducers weren't firing. Bill and Nigel spent a 14-hour day in 104 degree heat tearing apart the mapper and ultimately finding a tiny wire shorted across one of the computer board traces. The next day, all worked fine. On June 18th we performed a formal demonstration of the equipment to the State Park officials. We had a large contingent of support divers including Larry Green, Jim Schlesinger, John Zumrick, Paul Heinerth, Jill Heinerth, Randy McGuire, Mike Bruic, Jim Lockwood, and Mark Meadows. Divers entered the cave with the mapper and the magnetic beacon. Then, on the surface, Brian instructed Sandy Cook, park superintendent, how to use the headphones and digital readout to follow the signal. Within 10 minutes she walked to the exact location on the ground surface above the radio beacon, which was located directly under the Grand Canyon dome. Finally the tour entered the air-conditioned comfort of the lodge to view the cave entrance area in 3D on an SGI computer. What we showed was just a hint of what will come in October.

The view on the computer looked good, but it wasn't quite as great as what we believed was possible. That day Nigel, the hardware guy, and Fred, the software guy, discovered a minor problem that made all the difference. Nigel was having the mapper record the array of transducers clockwise, while Fred was displaying them counterclockwise! Furthermore, each had chosen a different coordinate system for heading angles. Once these inconsistencies were resolved – with some vitriol concerning whether the mathematician or the engineer was right – Fred stayed up late that night fixing his code. Meanwhile, earlier in the evening, the divers continued mapping the entrance area and the first 250 meters of A-tunnel. On June 19th, our last day, we were able to show the park what we, and they, most wanted to see: a 3D map of the cave that accurately represented what was really inside, including the Grand Canyon dome where roof of the cave soars upward to 50 meters (165 feet) above the floor. Even though the vehicle path had been down near the floor (240 feet) it captured the geometry of the entire dome. That evening (after the glass-bottomed boats finished their tours) the divers

Visitors to Wakulla Springs will one day be able to take a virtual tour of the behemoth cave system using this technology. ■

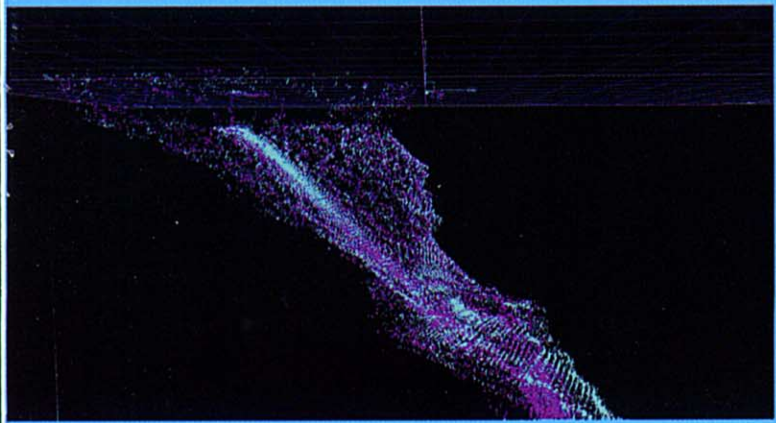

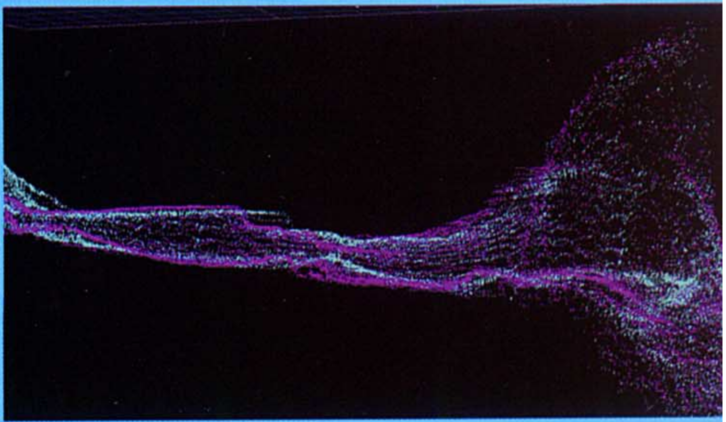

made one last dive into the cave, this time mapping out to Grand Junction, around 400 meters in. The combined result of those two dives, about 12 minutes of total data collection, was nearly 100,000 registered points defining the walls of the cave, in addition to the path taken by the vehicle. The world's first automatically generated 3D cave map had been created.

Back in Maryland, Barbara merged the data sets for the two deepest penetrations into the cave. We were delighted at the precision of overlap of the two data sets. Because the mapper will rarely be at the exact location as on a previous dive, each additional dive will fill in the areas between the wall points recorded on previous dives. Further, because the sonar beams radiate out from the mapper, the closer the mapper is to a wall, the closer the wall points lie to one another. When the wall is very far away, the sonar beams spread out quite a bit. This response characteristic can be used to advantage. By flying the mapper close to one passage wall on the way in and the other wall on the way out, we can merge the data sets and get closely spaced data for both walls. It is a particularly powerful capability when the passage is oddly shaped with either a rise in the floor or a drop in the ceiling.

On July 10th another important milestone for the project was passed: the board of directors of the National Geographic Society, based largely on the digital map of A-tunnel, voted to join the Wakulla 2 expedition. They have subsequently granted it status as one of their principal exploration projects of 1998. As such, the Wak2 website will shortly be accessible through Geographic. The site will go "live" shortly after Mission Control is set up at the springs on October 1.

Between now and October we'll be porting the data to new simulation software that will merge our underground data with surface topography. Various corporate groups are working with us to bring that data to life in 3D immersive theatres. It will soon be possible for the average visitor to Wakulla Springs (as well as park and state officials) to tour the cave without the risk traditionally borne by the handful of explorers who have touched its innermost reaches. Since more than 90% of the park has, until this moment, lain effectively invisible to its stewards, that will be no small accomplishment. DT

Dr. Bill Stone is a long-time cave explorer and will lead the USDCT project at Wakulla Springs in October.