

THE WORLD'S MOST POWERFUL MRI TAKES SHAPE

Medical researchers expect unprecedented resolution



An MRI scanner equipped with a superconducting magnet strong enough to lift a 60-metric-ton battle tank will offer unprecedented images of the human brain when it comes on line a little more than a year from now, say its builders.

The imager's superconducting electromagnet is designed to produce a field of 11.75 teslas, making it the world's most powerful whole-body scanner. Most standard hospital MRIs produce 1.5 or 3 T. A few institutions, including the University of Illinois at Chicago and Maastricht University, in the Netherlands, have recently installed human scanners that can reach 9.4 T. Superconducting magnets used in the Large Hadron Collider, which last year was used in the discovery of the Higgs boson, produce a field of 8.4 T.

The development of the scanner, known as INUMAC (for Imaging of Neuro disease Using high-field MR And Contrastophores), has been in progress since 2006 and is expected to cost €200 million, or about US \$270 million. The project reached a key milestone this summer with delivery of more than 200 kilometers of superconducting cable, which is now being wound into coils that will produce the scanner's magnetic field.

"We're pretty proud of having met all the requirements, plus given them a little extra," says Hem Kanithi, vice president of business development at Luvata, in Waterbury, Conn., which built the superconductor.


Standard hospital scanners have a spatial resolution of about 1 millimeter, covering about 10 000 neurons, and a time resolution of about a second.

THIS ONE GOES TO 11: The main superconducting coil for the INUMAC imager is made from 170 kilometers of niobium-titanium. When bathed in liquid helium and charged, it produces a magnetic field of 11.75 teslas.

The INUMAC will be able to image an area of about 0.1 mm, or 1000 neurons, and see changes occurring as fast as one-tenth of a second, according to Pierre Védrine, director of the project at the French Alternative Energies and Atomic Energy Commission, in Paris. With this type of resolution, MRIs could detect early indications of brain diseases such as Alzheimer's or Parkinson's and perhaps measure the effects of any methods developed to treat those illnesses. It would also allow much more precise functional imaging of the brain at work than is currently available. "You cannot really discriminate today what is happening inside your brain at the level of a few hundred neurons," Védrine says.

High-field MRI could also allow scientists to explore different methods of imaging. Most MRI machines rely on imaging the nuclei of hydrogen atoms, but stronger scanners might gain useful physiological information by looking for weaker signals from sodium or potassium nuclei.

Improved superconducting wire is key to making such a powerful machine. The



wire in the INUMAC magnet is made from niobium-titanium, a common superconductor alloy. But it will experience some uncommon conditions as part of INUMAC. To reach the required field strength, the electromagnet must be able to carry 1500 amperes at 12 T and be cooled by superfluid liquid helium to 1.8 kelvins. That requires specialized manufacturing and precise control of the dimensions of the wire, allowing it to be coiled so the cables are aligned to within a few micrometers of precision. “We are pushing the superconducting material niobium-titanium very close to its limits,” Védrine says.

Another material, niobium-tin, can produce magnetic fields stronger than 20 T, but it was passed over for the job because it’s more expensive than niobium-titanium and very brittle, making it difficult to wind.

Ultimately, Luvata produced 170 km of wire for the main superconducting coil. The company made another 58 km for two secondary coils, which will produce an opposing magnetic field to shield the area outside the machine from stray magnetic fields.

Instead of winding the wire into one long coil, as is standard in systems with lower fields,

engineers are using a “double pancake” design, in which the wire is coiled into two reels that are spliced together, one on top of the other. The whole magnet will consist of 170 of these double pancakes connected in series. Védrine explains that this reduces the chances for error: Making a mistake in the winding phase using a single helical coil could ruin the whole magnet. However, a miswound pancake can simply be swapped out for a new one. The design provides space for the liquid helium bath to reach all of the coil and keep the temperature low, and it also allows engineers to place the best-performing coils at the center of the system, which improves the precision of the magnetic field.

The inner diameter of the magnet will be 90 centimeters, wide enough for a human body. Patients getting scanned will lie entirely inside the machine, but the region where the field is precise enough to get maximum resolution will be only 22 cm long. “The very good field region is only in the middle of the magnet,” Védrine says. While patients could be situated so that other body parts would fit inside that region, “first we are concentrating on the brain,” he says. It would take an even more massive machine to enlarge the high-resolution area.

Védrine expects to deliver the fully assembled magnet by September of next year. Other parts of the imaging system will then be added in and around the magnet, followed by about three months of testing. “Probably we’ll have the first images by the beginning of 2015,” he says. —NEIL SAVAGE

THE CACHE MACHINES

Many-core processors need the flexibility of software to use their onboard memory best



A new process for managing the fast-access

memory inside a CPU has led to as much as a twofold speedup and to energy-use reductions of up to 72 percent. According to its designers, realizing such stunning gains requires a big shift in what part of the computer controls this crucial memory: Right now that control is hard-wired into the CPU’s circuitry, but the substantial speedup came when the designers let the operating system handle things instead.

The CPU uses high-speed internal memory caches as a kind of digital staging area. Caches are a CPU’s workbench, whether they’re holding onto instructions a CPU may need soon or data it may need to crunch. And from smartphones to servers, nearly every CPU today manages the flow of bits in and out of its caches using algorithms built into its own circuits.

But, say two MIT researchers, as computers and portable devices accumulate more and more memory and CPU cores, it makes less and less sense to leave cache management entirely up to the CPU. Instead, they say, it might be better to let the operating system share the burden.

In itself, this idea is not completely new. Some of IBM’s Cell processors, as well as Sony’s PlayStation 3—which runs on Cell technology—allow their applications and OS kernels to fiddle with low-level CPU memory management. What’s new about the MIT technology, called Jigsaw, is its middle-ground approach, which enables software to configure some on-chip memory caches but without requiring so much control that programming becomes a memory-management nightmare.

“If you go back six or seven years, you’ll see that everybody was complaining that they launched the PlayStation 3 and nobody could program it well,” says Daniel Sanchez, the assistant professor at MIT’s Computer Science and Artificial Intelligence Laboratory and one of the inventors of Jigsaw.

Today, CPU hardware typically controls all the on-chip caches. So those caches must be designed to handle every conceivable job, from pure floating-point number crunching (which places a small burden on caches) to intensive searches and queries of a computer’s mem-

The electro-magnet must be able to carry 1500 amperes at 12 teslas while cooled to 1.8 kelvins