Security

Scientific American (October 2009), **301**, 80-84
doi:10.1038/scientificamerican1009-80

Privacy and the Quantum Internet

Seth Lloyd

Courtesy of some of the weirdest laws of physics, we may someday be able to search and surf the Web without anyone collecting our data.

KEY CONCEPTS

- Current Web searches, even when anonymized, can still reveal personal information about the user.
- Over a quantum version of the Internet now being developed, search engines could return queries back to the users with the answers frut—and with the assurance that no one has saved or copied the data.
- Quantum searches will require search engine databases to use a new kind of memory storage, which is already being demonstrated in the laboratory.

—*The Editors*

ADDITIONAL IMAGES AND ILLUSTRATIONS

what is this Quantam 'internet'?

- why is it secure

[MEMORY FOR QUANTUM SEARCH ENGINES]: Pssst! Keep This to Yourself

Privacy is hard to come by these days, particularly on the Internet, where every time you Google something your desires are recorded for posterity—or at any rate, for advertisers.

Internet search companies say they protect their clients' privacy by encrypting personal information and by using numbers instead of names to give their users anonymity. The problem is that anonymization is not always effective. AOL user number 4417749 found this out the hard way in 2006 when AOL decided to publish online a list of 20 million Web searches, including hers and those of 657,000 other users. Reporters were able to track down the 62-year-old widow in Lilburn, Ga., by analyzing the content of her searches. Luckily, Thelma Arnold was relatively unembarrassed by the revelation of her identity and intimate interests. How many of us could say the same?

The laws of physics, however, could come to the rescue. Communication over special "quantum channels" already enables banks and other institutions to send data with virtually unbreakable encryption. Thus, the technology already exists to hide your searches from eavesdroppers who might intercept your queries. But in the future a new "quantum" version of the Internet may enable you to send queries and receive answers with the assurance that no one—not even Google—knows what questions you have asked. Moreover, the same technologies that will guarantee private searching could also guarantee privacy during the entire online experience.

Of course, search engines save and analyze users' data so that they can display targeted ads. That is how they cover their expenses and make a profit. If they decide to keep the users' data private, the search engines will need a new business model. And users may have to decide if they are willing to pay for searching or if they would rather do it free and give their searches away.

Nonclassical Listeners

It was with some sense of misplacement that in the spring of 2004 I found myself attending a billionaires' banquet during a conference in Monterey, Calif. My role, as near as I could make out, was to be the guy who jumps out of the cake—in other words, to entertain guests interested in quantum technology. The legitimate billionaires at the event included Sergey Brin and Larry Page, the founders of Google. To my surprise, Brin and Page knew a lot about quantum information. After some wild speculation on how quantum physics might change the way people interact with the Internet, I suggested that I would work with my colleagues to investigate "quantum Internet search," whatever that might prove to be.

The ability of quantum physics to supply complete privacy stems from a simple fact: systems in the quantum realm (which includes anything from elementary particles to molecules) can exist in multiple states. At any particular time, an atom can be in several different places; a particle of light, or photon, can be polarized both vertically and horizontally; an electron's magnetic moment can point up and down, and so on. As a consequence, whereas classical (as opposed to quantum) data bits register either the value 0 or the value 1, quantum bits can register 0 and 1 at the same time. Also, whenever a quantum bit takes on the values 0 and 1 simultaneously, you cannot make an exact copy of that quantum bit, and any attempt to do so will change the state of the bit. This rule, known as the no-cloning theorem, also applies to strings of quantum bits, which, for example, can represent words or sentences. As a consequence, someone eavesdropping on a quantum channel—typically an optical fiber carrying photons in multiple polarization states—will not be able to "listen" to the communication without disturbing it, thus revealing the intrusion.

Several different quantum encryption techniques exist to exchange data in complete privacy thanks to no cloning. Yet such techniques presume that the addressees be allowed to read the data you sent them: merely sending Google an encrypted search query would not help. Last year, however, my colleagues Vittorio Giovannetti of the Scuola Normale Superiore di Pisa in Italy and Lorenzo Maccone of the University of Pavia in Italy and I discovered that the no-cloning theorem also makes private queries possible. In the protocol we devised, a user must be able to send the search engine a "quantum question"—a string of quantum bits that simultaneously contains the true question and another one. (It does not matter what the second question is: your computer could even supply a random one automatically.)

The search engine searches its database for the answers to your multiple questions and combines questions and answers into a new quantum package, which it sends back to you. If the search engine makes a copy of the questions for its records, you will be able to tell that your privacy was violated because the quantum state of your original questions will be perturbed in a way that your computer can detect. Crucially, the search engine can provide answers without physically detecting (let alone cloning) the string of bits that encodes the questions and thus without knowing what the questions were.

Although such magic is impossible with current computers, databases and networking hardware, we realized that it is not technologically out of reach. The first requirement for quantum private queries is a rudimentary quantum Internet. The technology to exchange quantum messages along a dedicated line already exists and is in use for secure communication. A full-fledged quantum Internet, however, will have to be not just a line between two points but a network whose nodes route data packets so that any user can reach any other user or any Web server. It turns out that routing data without making temporary copies of them—and thus without suffering the consequences of the no-cloning theorem—is a nontrivial task and requires a sophisticated technology now at the experimental stage, called a quantum router. A prototype of such a network may become available within five to 10 years.

The second requirement for private Web searching is that users and data servers possess rudimentary quantum computers, meaning computers that are able to store and handle quantum bits. Unfortunately, quantum bits are notoriously fickle and tend to spontaneously lose their multiple quantum states within a fraction of a second. Experimental quantum computers that store quantum bits in the magnetic states of single ions suspended in a vacuum, for example, can store only eight bits or so at a time so far. A full-fledged quantum computer would require hundreds if not thousands of quantum bits and is probably many decades away, even as a laboratory demonstration. Fortunately, though, for the purpose of quantum private searches, only 30 quantum bits or so will be sufficient: if properly coded, a 30-bit query can pull an answer out of a database with more than a billion entries. Such 30-bit "quantum microprocessors" might also become available in five to 10 years.

Not So Random

So far everything looks good: quantum private searches seem to require only very simple quantum computers and quantum communication systems. Now the hard part comes. To answer a user's multipronged quantum question, a search engine's database must be able to supply the answer to each component of the question simultaneously. Doing so will require a new type of data storage called quantum random-access memory, or quantum RAM.

RAM is just a device for storing data, arranged in a treelike structure. Each piece of data is a sequence of eight bits, or a byte, and has an address that is itself a sequence of bits. Bytes are like the leaves on the tree; the address controls the route from the trunk to the particular leaf. The first bit of the address specifies which of two branches to take at the lowest level of the tree, the second bit controls the second-level branching, and so on. The branches double at each level, and in a traditional RAM with 30-bit addresses, retrieving data requires throwing 2 (more than one billion) switches.

One could design a quantum version of traditional RAM. The only difference is that the switches that route information through the binary tree must now be capable of routing information through two different branches simultaneously, because each bit of a quantum question can specify two different routes. Such quantum switches can be built using existing technology, such as semitransparent mirrors that "split" photons making them follow two different paths at once. The problem is that quantum circuits are exquisitely sensitive to noise and errors: if just one of the switches is messed up, the privacy of the corresponding bit is lost. Because a typical address bit controls a huge number of switches, the chances of losing privacy are very high.

Giovannetti, Maccone and I came up with a different design for addressing RAM (both quantum and classical), in which far fewer switches are thrown for each memory call. The secret is to route address bits along the same tree branches that data are to follow, rather than through separate addressing lines. Because the address bits are passed sequentially through the array, we call this a "bucket brigade" RAM [*see box on opposite page*].

The bucket brigade architecture requires throwing just one switch at each level of the array, whereas conventional RAM throws every switch at every level. The savings are striking: a bucket brigade RAM with a billion memory slots throws 30 switches for each memory call, compared with a billion switches thrown for each memory call in conventional RAM. And the benefits of the bucket brigade architecture, in terms of both error rate and energy savings, grow exponentially with the number of bits.

The Solace of Quantum

At first, we thought the bucket brigade idea had the potential to revolutionize the industry of classical RAM, and visions of dollar signs began to dance through our heads. But we soon found out that others had thought of a similar design before and that anyway the design was too slow for classical RAM (although it could be an energy-saving solution for nonvolatile memories such as those used in digital cameras).

But the bucket brigade design would be crucial for quantum searches, because its architecture can tolerate an error rate of one in 30, rather than one in a billion. The memory medium for a quantum RAM could consist of a conventional physical support. For example, the data could be stored in billions of tiny mirrors like those that make up the surface of a conventional CD. The truly quantum part of the quantum RAM is the switching array, which could be constructed from switches that each can shunt quantum bits along both branches at the same time. Such quantum switches already exist and attain error rates sufficiently low to build a quantum RAM with a billion slots or more.

Of course, assembling quantum switches to make a large quantum RAM will likely prove to be hard, not to mention the problem of connecting quantum RAM to quantum communication channels to implement quantum private queries. But none of the difficulties seem insurmountable. Incidentally, we realized recently that the data-routing techniques of our quantum RAM design could be applied to switching networks of the quantum Internet as a whole. People could surf the Web in complete anonymity, without revealing not only what they are searching for but also what Web sites they are visiting.

A few months after my colleagues and I had nailed down the details of how to build quantum RAM and make quantum private queries, I ran into Brin and Page at a conference in Napa, Calif. In a hot tub beneath fig trees,

as the stars wheeled overhead, I described how quantum queries worked and what their benefits might be. Their first response was that Google's business model was to keep the information about all queries and to use it to prioritize advertising and future search results. Not retaining the information about queries had not occurred to them. When I put to them the evident advantages of a new, quantum business model—based on charging customers for search results—they thought some more. "Okay," they said, "Let's see if you can build it."

Recently Francesco De Martini and his group at the University of Rome "La Sapienza" have done just that. Using lasers, polarizers and photon detectors, they built a simple quantum RAM and demonstrated our search protocol on a small database. Quantum private queries are thus a real possibility. If someday we will have larger quantum RAMs or a viable quantum Internet, what happens then is anybody's guess.

THE AUTHOR

Seth Lloyd is professor of mechanical engineering (or a "quantum mechanic," as he likes to describe himself) at the Massachusetts Institute of Technology and director of M.I.T.s W. M. Keck Center for Extreme Quantum Information Theory. He developed one of the first theoretical models for quantum computation and is working with several groups to construct quantum computers and quantum communications systems. Lloyds book *Programming the Universe* was published by Knopf in 2004.

MORE TO EXPLORE

Best-Kept Secrets. Gary Stix in *Scientific American*, Vol. 292, No. 1, pages 79–83; January 2005.

Quantum Random Access Memory. Vittorio Giovannetti, Seth Lloyd and Lorenzo Maccone in *Physical Review Letters*, Vol. 100, No. 16, pages 160501–160504; April 25, 2008.

Quantum Private Queries. Vittorio Giovannetti, Seth Lloyd and Lorenzo Maccone in *Physical Review Letters*, Vol. 100, No. 23, pages 230502–230505; June 13, 2008.

The Quantum Internet. H. J. Kimble in *Nature*, Vol. 453, pages 1023–1030; June 19, 2008.

Scientific American ISSN 0036-8733