Seti@Home

Either we are alone in the universe, or we aren't – either possibility is difficult to comprehend.

In hope of the latter, many our best scientific minds spend their time analyzing an endless stream of radio telescope data in the Search For Extra-Terrestrial Intelligence – otherwise known as SETI. Although the search is intellectually lucrative, its success depends mainly upon powerful tools rather than powerful minds.

What makes radio telescopes so useful for SETI researchers? According to most scientists, radio is the best and perhaps only chance we have at interstellar communication. They base this consensus on two factors: absorption and noise.

Absorption

Because the Earth's atmosphere absorbs most of the electromagnetic spectrum (frequencies that we can communicate on), we can detect only three from the ground. These are: radio, visible light, and the near infrared. To make the choice even simpler, interstellar gas and dust severely absorbs most of the visible light and near-infrared energy.

Noise

Compared to the other frequencies, the radio spectrum is generally the least noisy. Of the trillions of stars in the universe, most tend to be very hot (emitting quite a bit of near-infrared energy) and emit huge amounts of light. Our sun is a just one of the many "noisy" stars in our universe.

The SERENDIP Project

Nestled deep in the mountains of Puerto Rico, the massive Arecibo radio telescope covers an area of about 20 acres and gulps an astounding 35 GB (gigabytes) of data per day.

Although its primary use is for traditional radio astronomy, the SETI researchers at UC Berkley have developed an ingenious system to use the telescope as though it were dedicated to their cause. By intercepting (but not interfering with) the radio stream as it enters the telescope, they are able to "piggy back" their experiment on the experiments of others. This "spectrum analyzer" (their data gathering instrument) relieves researchers from the burden of competing for tightly budgeted telescope time.

Because of its single, dedicated task, the spectrum analyzer performs a shallow analysis of the data stream in *real time*. Using its 40 spectrum analyzer boards (that split the task up amongst them), the machine performs a very respectable 200 billion-instructions-per-second. We traditionally measure supercomputers' speeds in *trillions of operations per* second (teraflops,) so this is the same as an 0.2 teraflops per second supercomputer. In comparison, the world's fastest supercomputer runs at a top speed between 1 and 2 teraflops per second. At the end of each day, though, the 35 GB DLT (Digital Linear Tape) that the analyzer records is mailed to UC Berkley to be analyzed by the most impressive supercomputer ever assembled: SETI@Home.

SETI@Home

Although the Serendip spectrum analyzer is incredibly powerful, its computing power is still limited. In order to keep up with the massive stream of data, the analyzer must look only for the strongest signals. There is also a whole subclass of signals that it simply does not have the power to look for.

One option that researchers have to look deeper into the data is a trade-off. Rather than build a larger supercomputer, they can simply let a smaller computer take longer to analyze the data. Unfortunately, with 35 GB of information being created every day, the information simply cannot be processed fast enough.

As a solution, the researchers at UC Berkley called on one of the more recent (and powerful) advances in computer science: distributed computing.

The distributed solution

Distributed computing operates on a very simple principle. First, break a complex problem into many sub-problems. Second, *distribute* these sub-problems to separate computers and let them solve the problems independently. Finally, merge the answers of the sub-problems into the solution of the complex problem.

The SETI@Home solution

Using this distributed approach, the SETI@Home team has helped to build the largest distributed supercomputer ever. In their words:

> "The UC Berkeley SETI team has discovered that there are already thousands of computers that might be available for use. Most of these computers sit around most of the time with toasters flying across their screens accomplishing absolutely nothing and wasting electricity to boot."

By enlisting the help of almost half a million *active* generous home computer users, UC Berkley is applying the equivalent processing power of a **11 teraflop** supercomputer to the mountains of data streaming out of Arecibo. With this power, they search for signals 10 times weaker than those searched for by SERENDIP. Keeping in mind that the world's fastest computer performs approximately 2 teraflops (with a price tag of 85 million dollars), the SETI@Home project breaks ground both in astronomy and computer science.

How does SETI@Home work?

In order for the project to be as non-intrusive as possible, the SETI@Home designers imbedded the calculation software into a flashy screen saver. Because screen savers are only active when the user isn't, users don't have to worry about the software bogging their system down while they are using it.

How does the screen-saver work?

To understand what each computer actually contributes, we must first consider exactly what the telescope at Arecibo records.

As mentioned earlier, one benefit of using radio waves for transmission is the galaxy's relative silence in the radio spectrum. Even within the radio spectrum, though, there is a window (between 1418 MHz and 1421 MHz) where the galaxy is the most quiet. Because these frequencies are between the absorption bands of Hydrogen and Water, scientists jokingly call it "The Water Hole." It seems to be a fitting place for two civilizations to meet! Since this is the most logical window for interstellar communication, SETI@Home analyzes this 2.5 MHz wide band.

However, analyzing a band this wide requires enormous computing power (our initial problem), so UC Berkley breaks this band into 256 chunks of about 10 kHz each. They then send approximately 107 seconds of this 10 kHz data (called a "work unit") to the end user and let the SETI@Home software do the rest.

At Berkley's end, each work unit is tracked in a massive database. In this database, work units that are sent out are marked as "in progress," and work units that are received back from the user are marked as "completed." After being sent out, if the work unit is not received within a few weeks Berkley simply assigns it to another user. This is where the distinction arises between "active users" and "users": "active users" have completed a work unit within the last two weeks.

What SETI looks for

In SETI research, scientists have to make several assumptions about what form the message will arrive in.

Because the Arecibo telescope does not move to track the stars, the first assumption is that the signal will not appear to be steady. Since the stars are appearing to move past the telescope (and the planet sending the signal moves past the telescope's focus), the assumption is that a signal will appear dim at first, rise in intensity, and then fade again. This rise, then fall, in intensity is called a "Gaussian."

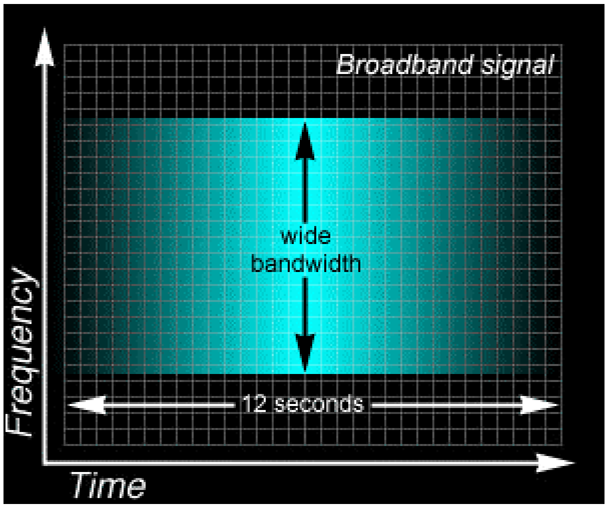

The second assumption is that the signal will be sent as a “narrow band signal” – in narrow frequency range.

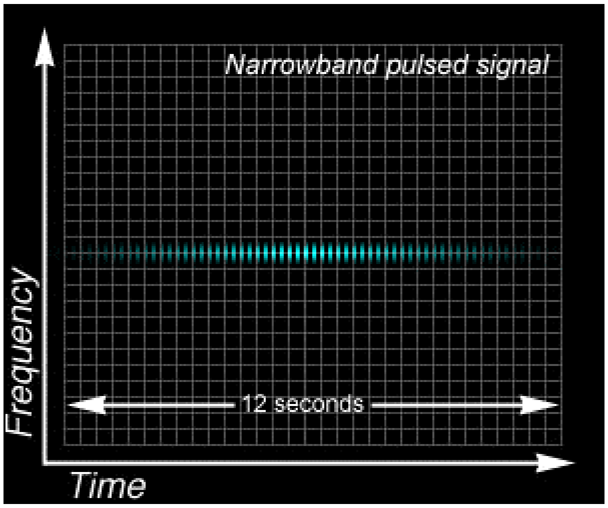

Because the “broad band signals” are more difficult to separate from the background noise, it would be logical to communicate in as narrow a frequency as possible. SETI researchers are also assuming that an incoming communication would contain a message as well – so watch for signals that appear “pulsed.”

Finally, scientists think that the signal will probably be “chirped” as well. Since our galaxies move with respect to each other, the pitch of the signal will probably rise or fall throughout the Gaussian.

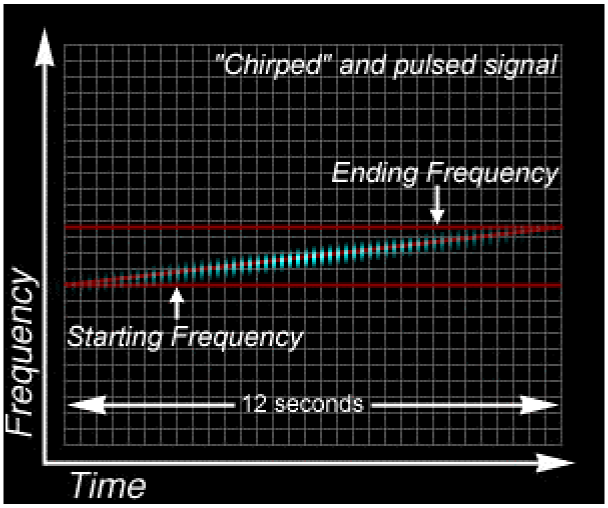

In the event of *The Discovery*

Once a candidate for an extraterrestrial signal is received, it is subjected to a gauntlet of tests. First, it is checked against a database of known “Radio Frequency Interference” (RFI) sources. This check eliminates 99.9999% of all the candidate signals. Second, scientists take a second observation from the same part of the sky and compare the results. If these signals seem to agree, the SETI@Home team asks another SETI team to take a look. This other team uses different computers, telescopes, and receivers to verify the observation.

Once confirmed, the SETI@Home team will announce the discovery to the International Astronomical Union, and then to the public. A multinational effort will then be co-ordinated, and the historic process of “first contact” will begin.

Make yourself part!

Becoming part of the SETI@Home project is as easy. Visit http://setiathome.ssl.berkley.edu/download.html to download the program that will run on your computer. There are versions for Windows (95/98/NT), Macintosh, and many of the different flavours of Unix. Once you have installed the client, you can be proud to be part of the largest computation ever performed!