Lowering the cost of high performance computing

A new microchip from NVDIA promises to let you have high performance computing in your desktop PC for under $10,000.

Imagine if high performance computing was so cheap that you could have your own super powerful computer sitting on your desk, rather than queue up your high performance computing jobs in a data centre.

This is getting closer and closer, thanks to a new microchip being released by Californian company NVIDIA.

For £4,000 / USD 8,000, you can install four NVIDIA graphics processing units (GPUs) in your desktop computer, which gives you about 100 times more processing power than would have been available in a 1990s mainframe computer, which would have cost $80,000 to $100,000, according to NVIDIA's Andy Keane, general manager of Tesla computing products.

You can even install them in your laptop (although it is 10 inches by 4 inches in size, there is probably a limit to how many you can fit in).

Some users are thinking about switching off their data centres, because people have enough computing power on their desks to do what they want now, Mr Keane says. "The GPU is a room sized computer in one chip."

Many people in the oil and gas industry believe that if they had better access to faster computers, they could probably find more oil and gas, according to a recent Microsoft survey; but they haven't been given access to it, presumably because a decision was made that the outcome would not justify the expense.

But if the expense reduces, then the decision to provide high performance computing gets much easier.

Steve Purves, technical director of ffA, a company which provides high performance seismic analysis software, observes that there seems to be a psychological bar at around $10,000 / £5,000 at many oil and gas companies for a personal high performance computer. If the price drops beyond that level, the decision to give someone their own powerful computer seems to get much easier.

"Not everyone is going to get a $100,000 computer," says Mr Purves.

By having faster computers, you can see your results straight away, rather than create processing jobs, send them to a centre, and have a look at the results a few hours later. You can create views of the subsurface which you can walk around and look at from different angles, rather than wait several hours for the computer to draw a new view.

And besides, everybody knows what is going to happen to the amount of data people have to deal with in coming years – it will grow exponentially. That means we will need faster and faster computers to deal with it.

"Oil companies are trying to scale up exploration – and this means processing more data and generating more prospects," says Mr Purves. "That's where software tools come in."

The amount of information continues to increase, as resolution of surveys increase. "Interpreters will need to find a way to deal with it," he says.

Mr Keane observes that big computing centres have got bigger and faster over the past few years, and mobile phone technology is also racing ahead, but technology for the humble workstation has not moved so far ahead.

Latest chips

The latest chip to be launched is the Tesla 10 series.

The Tesla T10P processor has 1.4 billion transistors, 1 teraflop of processing power (double the speed of its previous chip, the Tesla 8), and 240 processing cores.

It has 4 GB of memory (compared to 1.5 GB on the Tesla 8).

The standard computer package is 1 unit in size and contains 4 GPUs, so you get 4 teraflops altogether 16 GB of memory, and 960 cores. It uses just 700W. A single GPU will use just 160 W.

The previous chip, the Tesla 8, was launched in late 2006.

The chip has been widely used in many different applications, including medical imaging, astrophysics, weather forecasting, clothing design and finance. The US National Center of Atmospheric Research believes it can halve the time taken to make a weather forecast using the chips.

It can perform many different tasks in parallel, something that can't be done on a standard central processing unit (CPU) you might find on your normal desktop.

You have to write code especially for the chip, but it uses standard coding languages, such as C. "We stick to the language standards everyone knows," says Mr Keane.

The NVDIA Graphics Processing Unit (GPU): turns your laptop into a supercomputer

It is possible to have a GPU and CPU on the same computer, and then use them separately – the GPU for special tasks, and CPU for general running of the computer.

Big computing centres

If four GPUs in your desktop is not enough and you want to build a computing centre, you will get much more processing for your dollar, and kilowatt, by using GPUs than standard CPUs (as found in most PCs and servers).

US financial consultancy Hanweck estimates that 12 GPUs would have equivalent processing power to 600 CPUs, but take up just 6 units of rack space (compared to 54). The hardware would cost $42,000, instead of $262,000, and have an annual cost of $140k instead of $1.2m (calculated on a storage cost of $1,800 per unit per month, for rack and power charges).

Using its 12 GPUs, Hanweck Associates regularly scans the entire US options market in under 10 milliseconds, something which would normally take 60 traditional 1 unit servers, the company believes.

If you want to build a 100 teraflop data centre, you could do it with 1429 CPU servers, each providing 0.07 teraflops, with four teraflop CPUs in each one (at estimated cost of $3.1m and power consumption of 571 Kw); or you could do it with 25 servers each with 4 GPUs, providing 4 teraflops per server, with a total cost of $310,000, and total power consumption of 27kW.

There is also a physical limit to how fast you can get using CPUs, because the time taken for all the communications be-

tween the chips weighs things down. With GPUs you can get much faster.

Computer games

NVIDIA started life as a manufacturer of chips for computer games – in 1995 it made the graphics card for Sega Saturn gamepads.

Computer games and subsurface simulation have plenty of things in common.

Both create a simulation of the real world; both generate images in 3D; and both have plenty of real time interactivity.

Many geophysicists and geologists have spotted that their children have more powerful computers in their games consoles than they do at work.

A lot of seismic modelling is still done in 2D slices, to reduce the amount of computer power needed. "It's that mainframe mentality," says Mr Purves.

ffA

ffA, which does 3D seismic analysis software, is working together with NVIDIA to develop high performance computing capabilities for its seismic volume imaging and 3D visualization software SVI Pro.

Oil and gas customers include Hess, Total, CGG / Veritas, Headwave, Accelerware, and Seismic City.

ffA believes that the computational performance will help improve its workflows for seismic volume interpretation.

At ffA, the computer chips are being used to put together 3D images, working on data sets as much as 200gb at once.

ffA makes software which can extract faults and complex structures in the data, so it needs all the data processing it can get. It works on the post stack seismic data.

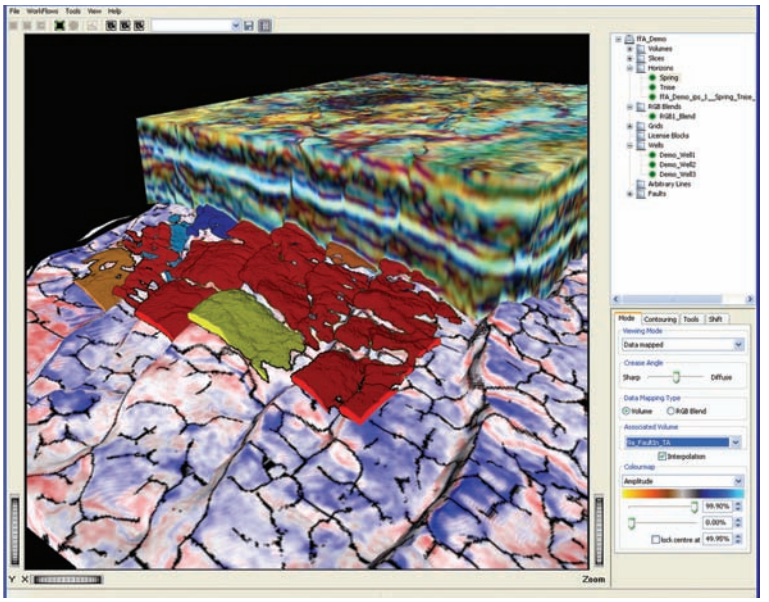

A screenshot from ffA's advanced seismic volume imaging and 3D visualisation software SVI Pro, which will use NVIDIA CUDA to deliver enhanced performance

Using the new chips, it can provide real time processing – in other words, the user can see the results of what they are doing straight away, rather than send the data to a processing centre and wait a few hours. "We can get the results in front of the user in real time," says Mr Purves. "We want to hide the computing."

Or, as NVIDIA's Mr Keane says, "It's like the difference between looking at the photo and looking at a movie."

ffA has been recommending that its customers install NVIDIA chips for 5 years. Mr Purves says that what he likes most about the chip is its scalability – you can start with one, and then add more and more as your needs increase.

He also likes the flexibility- software which runs on NVIDIA chips can also run on other chips. "We've always been cautious about locking ourselves into hardware," he says.

"We will develop for this quite aggressively over the next 18 months," says Mr Purves. "This computing capacity is a way to design software differently. It will change the way you do your work," he says.

Texas Instruments - new projector chip

Texas Instruments has launched a new projector chip, which can create much clearer and brighter projections than others on the market, the company claims.

Texas Instruments has launched a new high resolution projector chip – called WUXGA with a resolution of 1920 x 1200 which is higher than full HD, with its 'Digital Light Processing (DLP) technology, which can provide much clearer images in oil and gas collaborative working centres, and also head much less maintenance than other chips on the market.

It should help companies create much more appealing large displays, which are more pleasant to look at, and which can convey more information, more clearly, and which can combine many different images and windows on the same screen.

The image is created inside the chip, and then the hardware around it projects the image onto a display.

The technology is capable of using up to 6 colours, instead of the standard 3 colours used on most of projectors - it has yellow, cyan and magenta in addition to the normal red green and blue.

"That lets us do more accurate and more colourful image," says John Reder of Texas Instruments. "This is an approach many of the highest photo printers use."

The projector also has the best contrast ratio, he says.

Another advantage is that projector cooling system does not have any filter, and so no filter cleaning is required; standard projectors recommend that the filter is cleaned every 100 hours of use, which can be a hassle.

The projector avoids the need for filters by sending the cooling air along a different path. Light comes in at the top of the chip

and goes out of the top of the chip, so it can be cooled from the bottom.

"With other display technologies, the light passes through the chip, so you have to blow air through the chip to cool it," he says.

WUXGA projectors based on DLP technology available now are projection design F10 and projectiondesign F30.