Technology trends in simulation

Dr Carlo Kopp

THE USE OF SIMULATIONS HAS GROWN CONSIDERABLY SINCE THE FIRST MODERN FLIGHT SIMULATORS were introduced decades ago. Today, simulators permeate almost every aspect of military training. Evolving technology and specifically the booming computer games market has seen dramatic improvements in the capability of graphics hardware, software and display technology, decreasing cost per installation. The long-term trend is clearly toward cheaper and better simulation technology.

To appreciate the scale of this improvement, it is worth exploring what it takes to produce a high fidelity simulation and present it to personnel in a training environment.

Simulator fidelity is determined as much by software as it is by hardware. In a typical simulator, the design can be broadly divided into several components, each of which plays a specific role in producing the imagery, sounds and other effects to the user.

One component is the model of the battlespace or other operational environment being simulated. It calculates the relative motion and immediate state of the entities being simulated, such as friendly and hostile platforms, and their relative positions and orientations against the surrounding environment or terrain. This component is the heart of the simulation, and determines in particular how simulated hostiles and friendlies or bystanders behave in the simulation. How computationally intensive this component might be depends largely on the sophistication of the modelling used and the number of entities simulated in the software. For instance, a flight simulator for a helicopter or high performance jet fighter can become quite complex – if it is to faithfully reproduce the finer points of vehicle dynamics and handling through a large flight envelope. A simulator for a platform with simpler kinematics, such as an armoured vehicle or warship, is a much less complex affair. For the simulation to produce the illusion of reality, this model must be recomputed typically several times per second – a truly smooth simulation designed to deceive the user well may generate updates tens of times per second.

The continuously updated output from such a model is essentially a collection of state information, and positional information, including location and orientation, for every entity in the simulation. This output must then be presented in a form useful to the user. That will involve software that calculates the graphical rendering, graphics hardware to render display imagery, and display hardware to present the imagery to the user.

Software that renders simulation imagery will mostly be proprietary to a simulator but more than often layered upon standard rendering libraries such as OpenGL.

This software will generate rendering commands for a graphics engine, the device that actually generates the imagery for a display device. The video signal output from the graphics engine is then fed to one or more display devices. The latter may vary between projectors at one end of the spectrum, to plasma or LCD panels at the other.

High fidelity scene imagery produced by a CAE graphics engine.

Top end simulation products use expensive display technologies, including collimated displays, dome displays, and rear projection displays, to create the illusion of reality. Depicted are CAE simulators used for aircrew training.

If a simulator is merely presenting a radar or sonar scope image, the task of presentation is relatively simple. If the output is a high fidelity audio-visual presentation with a large field of view, the task of presentation becomes much more complex, and expensive.

Historically, most of the cost incurred in professional quality simulations has been the result of producing sufficiently lifelike presentation to convince a human mind well enough to get useful training effects.

A good example of the expanding use of simulators for roles other than platform crew training is the Bohemia Interactive VBS2 simulator, used by the USMC for training.

At the top end, the flight simulator market involves a six-axis motion system moving a lifelike (internally) mockup of a crew station or cockpit, elaborate high resolution and very fast graphics rendering and expensive display technology, even providing collimated presentation (focused at infinity) and covering often as large a visual field of view as can be observed by the crew. Air combat simulators often use hemispherical domes to permit the crew to observe the environment within the full field of view of a bubble canopy. In this respect, flight simulators are the most technologically challenging to build, compared to land and naval warfare platforms.

At the bottom end of the market, single or multiple fixed computer displays are used, and the technology differs little from consumer commodity gaming products.

Most of the hardware cost in modern simulators is incurred in the display component of the simulator, and to a lesser extent in the graphics rendering and computational hardware. It is precisely in this area that the commodification of many technologies is now producing impact, and will further change the simulation market in coming years.

Computational Costs – Moore's Law and Clustering

Gordon Moore's 'Moore's Law' defined during the 1960s appears to a constant metric of computing power growth per dollar expended over time. While Moore's Law was initially defined to relate the density of chips versus cost over time, it has proven to be durable metric when applied to computing power, as computing power has in recent times been largely dependent upon how many transistors designers can cram into a single microprocessor chip.

What Moore's Law says is that per dollar, available computing power will double every 18 months. Plotting representative chip performance since 1970 indicates that despite peaks and dips when new technologies are introduced, or are late in introduction, the trend has held well over almost four decades. Moore's Law has yet to hit the long anticipated brick wall of device physics, upon which the 18-month doubling rule collapses.

Since commodity microprocessor chips are at the heart of all simulators, and with the performance now available in the market in a consumer PC outstripping a supercomputer of the 1960s or early 1970s, it is fair to observe that no matter how complex a simulator might be computationally, the cost of computing hardware will not be an obstacle.

While the enormous growth in performance per chip since 1990 has had the greatest impact, other technologies have also enhanced this effect.

One technique that emerged during the 1990s is clustering, whereby a large number of computers are hooked together using a high-speed local network, and large computational problems are split across a multiplicity of cheap computers. Often described as 'supercomputing on a shoestring', clustering will over time penetrate much of the top end in the simulation market, since most simulation problems can be readily decomposed into smaller components, which is precisely the style of computation that clusters do best.

Graphics Engines and Moore's Law

Graphics engines or GPUs (Graphics Processing Units) used to render imagery quickly have been a critical driver of simulator cost as well as the fidelity of the imagery produced. Even two decades ago a professional quality flight simulator would have been constructed using a room full of 19-inch rack cabinets, each filled with custom designed and built rendering hardware. Moore's Law now sees single chip graphics engines matching or exceeding the performance of such legacy equipment.

What has produced this effect is the booming consumer gaming market. For better or for worse, children, adolescents, teenagers and adults are addicted to computer games, especially games which simulate adventurous real life activities, be they modern combat, historical combat, urban combat and crime fighting, car racing, flying aircraft, engaging with attractive women, or playing out roles in fantasy worlds. This market has been cut throat competitive on costs versus rendering performance, and has seen the rise of manufacturers such as NVIDIA and ATI, who periodically battle for the top slots in the market.

The cinema industry's insatiable hunger for advanced digital special effects must also be factored into this equation, as many of the sophisticated graphics rendering techniques now used in games were developed to provide special effects, and migrated over time into gaming products. The historical pattern for such hardware was typically that of GPU cards built for digital special effects providing the vanguard in new rendering techniques, but usually lacking the computational performance for anything other than offline non-real-time rendering, a non-issue in the special effects business. These techniques then migrated down into gaming products where speed is paramount.

Internally, modern GPUs have evolved in a like manner to what we have seen in general purpose CPU chips. In the latter, increased chip density has allowed the packing of the internal architecture of historically much larger computers into single chips. Most of the 'advanced' architectural features seen in a modern Pentium or Athlon chip were first used in the 1960s mainframes or supercomputers.

A modern GPU follows this pattern exactly, with what used to be multiple racks of hardware condensed into single chips or chipsets.

A state of the art ATI or NVIDIA GPU engine on a PCI card typically contains internally dozens of individual graphics-specialised processing modules for rendering graphics primitives, with specialised hardware for shading and other visual effects. The GPU will be coupled to a very high-speed memory, currently of up to a Gigabyte in size, to support the rendering calculations. The manufacturer's claims of lifelike quality are believable, judging from sample imagery now available.

Long term, this technology will dominate the market since it is cheaper to build a professional quality simulator by using dozens of the very same GPU chips the suburban gamer uses on their desktop, rather than developing custom hardware.

Display Technologies

By far the most expensive component of professional quality simulators has been the display technology, and at the top end of this market that is unlikely to change. At the bottom end of the market, home theatre technology will displace custom products.

At the top end of the market, simulators are typically equipped with multiple projectors, and often complex optical arrangements to achieve desired field of view and focus.

Collimated displays produce the best effect but are the most complex and expensive to construct, as additional optics (mirrors and lenses) are required to manipulate the focus. Much less expensive are rear-projected displays where images are projected onto segmented screen panels to create the illusion of the outside world. The drawback of projected displays – whether using frontal or rear projection – is that the image is not at infinity and thus less able to replicate the real world. Finally full or partial dome displays remain widely used, and are more than often frontal projection arrangements.

In this area of technology, projectors have seen the most evolution in recent years. Historically, the preferred projector technology were high brightness Cathode Ray Tube displays, usually one per colour channel with a lens arrangement integrated in the package, and all three tubes mechanically collocated to illuminate a single area

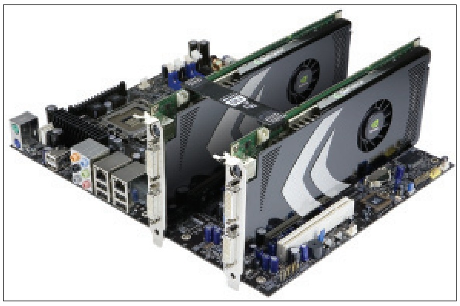

NVIDIA's 8800GT GPU is a good example of the state of the art in low cost rendering technology, intended to produce lifelike imagery.

with all three colour channels. The drawback of such projector CRTs was in limited life, due to the exhaustion of phosphors and cathodes – a more severe manifestation of CRT exhaustion seen in consumer displays.

Since then LCD light valve technology, micro mirror technology has emerged, and most recently, laser projector technology. The computer industry and more recently the home theatre industry have been voracious consumers of these technologies, driving down costs and improving optical performance. The CRT is unlikely to survive longer term in this marker niche.

At the bottom end of the simulator market, historically occupied by desktop computers and larger CRT monitors designed for computer applications, the home theatre industry has changed the game forever, and the CRT is rapidly heading for extinction.

The two technologies responsible for this are the LCD display panel, and the Plasma Display Panel (PDP), both now dominating consumer television and home theatre sales, both in standard definition (SDTV) resolutions, and high definition (HDTV) resolutions.

In comparison, the PDPs generally deliver better brightness, contrast ratio, colour fidelity, and useful life than LCD panels, but LCD panels are more popular in the consumer market. Both display technologies have problems with 'burn in' of static images, PDPs due to localised phosphor exhaustion, LCDs due to charge retention. Standard sizes in the consumer market are now at 30", 32", 37", 40", 42", 50", 58" and 65" diagonals, typically with resolutions of 1920 x 1080 (HDTV native), 1366 x 768 (HDTV resampled), or 852 x 480 (SDTV), with the latter usually including hardware to resample HDTV resolution signals, often with remarkably good quality. The smaller displays retail in Australia for under $1,000 the larger at $5,000 or more.

The computer industry has also been riding on this home theatre driven boom, but with more modest sizes due to demands in response time (6 milliseconds versus 20 milliseconds for TV applications). The top end of this market are 30-inch LCD monitors at 2560 x 1600 resolution, and 27-inch LCD monitors at 1920 x 1200 resolution.

While the resolutions and aspect ratios of home theatre and computer monitor displays differ, all typically share DVI or HDMI interfaces, permitting integrators to use all of these display types in simulator products. Where viewing distance is not critical HDTV products will often be much more practical than specialised OEM computer displays.

What the home theatre boom will result in the longer term is increased display capability in low cost simulation products. A simulator which provided a limited field of view with several CRT monitors clustered together can now provide vastly better field of view and image quality, using fewer or equal numbers of Plasma or LCD display panels, which take little effort to directly integrate due to common interfaces.

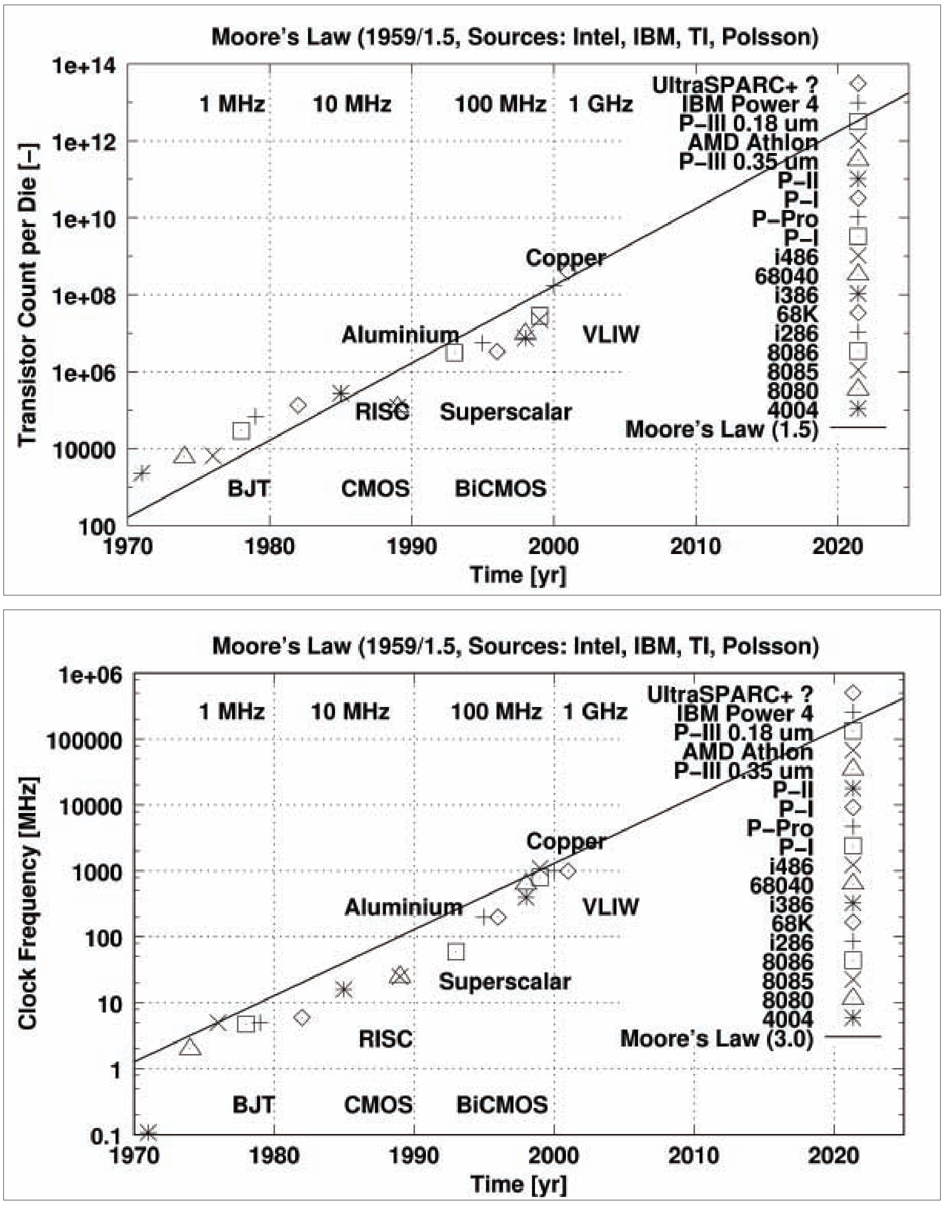

Moore's Law has proven to be a decisive driver, not only in CPU performance, but also in the performance of the critical GPU technology.

Perusing gamer websites and journals it is clear that dedicated gaming addicts in the consumer market are already adopting this approach, for whom 50-inch or larger display panels are a long coveted dream come true.

Conclusions

Observed these days is the confluence of a number of commodity consumer technologies, driven by demand in computer industry, home theatre and home computer gaming markets, and impacting the military (and commercial) simulation technology markets.

Top end high fidelity simulators will gain incrementally due to cheaper and faster GPU hardware for rendering high quality imagery, and because of higher resolution projector technology displacing legacy displays. This will improve the affordability of such simulators only in part since they still carry the large cost burdens of crew space mockups, external optical systems, and where applicable, motion systems.

Bottom end simulators, which have historically been constrained by commodity computer monitor costs and resolutions, are on the threshold of a major boom in the market as the home theatre technology boom drives down costs and drives up sizes and resolutions in displays, and the gaming technology boom drives down the unit costs of rendering GPUs capable of producing professional quality graphics.

What this means in practical terms is that we will see increasing use of simulations in military applications, with penetration into areas where simulations have been historically too expensive to use, or provided inadequate fidelity at the expense of training effects.

Military users need to start thinking now about where to best deploy these technologies to maximise training value and produce best long-term effect. In context, simulators have been used historically for both 'procedural' and 'tactical' training tasks, with the former often performed by low end products, and the latter by high end products.

With increasing graphics and display fidelity in low end products, the opportunities for good tactical training expand considerably.

What is clear is that simulation is entering a golden age.