Moore's Law and its Implications for Information Warfare

Carlo Kopp, BE(Hons), MSc, PhD, PEng,
Computer Science & Software Engineering,
Monash University,
Clayton, 3800, Australia
Email: carlo@mail.csse.monash.edu.au
W3: http://www.csse.monash.edu.au/ carlo/

January 6, 2002

Invited Paper

The 3rd International AOC EW Conference

Disclaimer:

This document was compiled wholly from available public domain sources. No classified documents were employed in its production.

Distribution:

Permission is granted for distribution free of charge for educational purposes. Distribution for other purposes is not approved without the prior written consent of the author.

Revision Status:

Invited Paper, The 3rd International Association of Old Crows (AOC) Electronic Warfare Conference, Conference Proceedings, Zurich, May 20-25 2000.

```
$Id: moore-iw.tex,v 1.7 2002/01/06 15:55:39 carlo Exp carlo $
```

Contents

1 Abstract

Moore's Law, defined in the sixties, predicts a monotonic increase in available computing power with time. With the commodification of high performance microprocessors, very large amounts of computing power are now readily available in the open market, at very modest cost. This paper will explore some of the implications of this phenomenon for IW in the coming decades, and propose some guidelines for IW strategists and planners.

2 Introduction

When Gordon Moore formulated "Moore's Law" during the mid 1960s, he could have neither anticipated the longevity, nor the longer term implications of this simple empirical relationship. Like many of the pervasive observations of twentieth century science, Moore's Law has become, in a rather mangled form, part of the popular culture. This is in a sense unfortunate, since mangled lay interpretations tend to frequently obscure the scientific substance of the matter. Scholars and practitioners of Information Warfare (IW) and its manifold sub-disciplines should carefully consider the longer term impact of Moore's Law, as it has the potential to significantly change the global economic, military and computing technology environment.

This paper will explore Moore's Law as an artifact of twentieth century technological history, discuss its most likely medium term implications in terms of computing performance, and then explore the areas in which it is most likely to change the science and practice of IW. Finally, a number of ground rules for the development of IW strategies, doctrine and related technologies will be proposed.

3 Moore's Law

Gordon Moore's empirical relationship is cited in a number of forms, but its essential thesis is that the number of transistors which can be manufactured on a single die will double every 18 months. The starting point for this exponential growth curve is usually set at 1959 or 1962, the period during which the first silicon planar transistors were designed and tested.

We now have four decades of empirical data to validate Moore's argument. Figure 1 depicts a sampling of microprocessor transistor counts, against the basic form of Moore's Law. Clearly the empirical data supports the argument well, even allowing for considerable noise in the data set.

Similar plots for high density Dynamic Random Access Memory (DRAM) devices yield a very similar correlation between Moore's Law and actual device storage capacity [2][1].

Two important questions can be raised. The first is that of "how do we relate the achievable computing performance of systems to Moore's Law". The second is the critical question of "how long will Moore's Law hold out". Both deserve careful examination.

The computing performance of a computer system cannot be easily measured, nor can it be related in a simple manner to the number of transistors in the processor itself. Indeed, the

---

[1] Current projections indicate 16 Gigabit DRAM devices by 2007.

only widely used measures of performance are various benchmark programs, which serve to provide essentially ordinal comparisons of relative performance for complete systems running a specific benchmark. This is for good reasons, since the speed with which any machine can solve a given problem depends upon the internal architecture of the system, the internal microarchitecture of the processor chip itself, the internal bandwidth of the busses, the speed and size of the main memory, the performance of the disks, the behaviour of the operating system software and the characteristics of the compiler used to generate executable code. The clock speed of the processor chip itself is vital, but in many instances may be less relevant than the aggregated performance effects of other parts of the system.
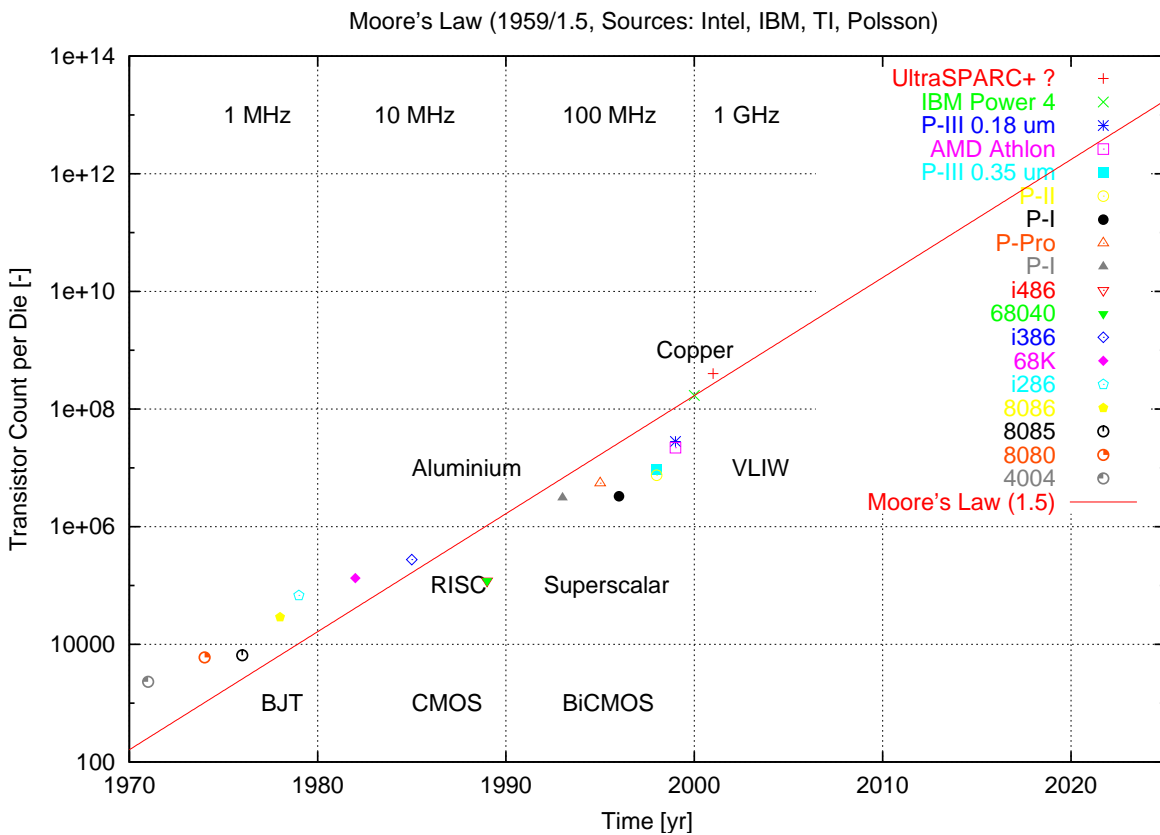

Figure 1: *Moore's Law for microprocessor transistor counts assuming a starting point of 1959 and doubling time of 18 months (Author).*

What can be said is that machines designed with similar internal architectures, using similar operating systems, compilers and running the same compute bound application, will mostly yield benchmark results in the ratios of their respective clock speeds. For instance, a compute bound numerically intensive network simulation written by the author was run on three different generations of Pentium processor, and a mid 1990s SuperSPARC, all running different variants of Unix, but using the same GCC compiler. The time taken to compute the simulation scaled, within an error of a few percent, with the inverse ratio of clock frequencies. Indeed, careful

study of published benchmarks tends to support this argument[2] [10].

The empirical observation that computing performance in like architecture machines scales approximately with the clock frequency of the chip is useful, insofar as it allows us to relate achievable performance to Moore's Law, with some qualifying caveats. Mead observes that clock speeds scale with the ratio of geometry sizes, as compared to transistor counts which scale with the square of the ratio of geometry sizes [14]. The interpretation of Moore's Law used in Figure 2 assumes this square root dependency, and incorporates a scaling factor to adjust the frequency to measured data. The plot in Figure 2 shows good agreement with Mead's model[3].

The conclusion that we can draw is that we will see a direct performance gain over time, proportional to the square root of the Moore's Law exponential, in machines with a given class of architecture. Since most equipment in operational use today is through prior history locked into specific architectures, such as Intel x86, SPARC, PowerPC, Alpha, MIPS and other, the near term consequence is that we will see performance increase exponentially with time for the life of the architecture.

However, major changes in internal architecture can produce further gains, at an unchanged clock speed. Therefore, the actual performance growth over time has been greater than that conferred by clock speed gains alone. Higher transistor counts allow for more elaborate internal architectures, thereby coupling performance gains to the exponential growth in transistor counts, in a manner which is not easily scaled like clock speeds.

This effect was observed in microprocessors with the introduction of pipelining in the 1980s, superscalar processing in the 1990s and will be soon observed again with the introduction of VLIW architectures over the next two years. Since we are also about to observe the introduction of copper metallisation during this period, replacing aluminium which has been used since the sixties, we can expect to see a slight excursion above the curve predicted by Moore's Law.

This behaviour relates closely to the second major question, which is that of the anticipated valid life time of Moore's Law. Many predictions have been made over the last two decades in relation to its imminent end. However, to date every single obstacle in semiconductor fab processes and packaging has been successfully overcome.

It is not difficult to observe that the limits on the scaling of transistor sizes and thus the

---

[2] A quick study of the Intel website, www.intel.com, comparing Spec series benchmarks for the Pentium Pro, Pentium II and Pentium III is an excellent example, since all three processors employ variants of the same microarchitecture, but are implemented in progressively faster processes, and span historically a five year period.

[3] The "dip" in the empirical data against the process curve during the early/mid 1990s is arguably a result of reduced commercial pressure for further clock speed gains, resulting from the introduction of superscalar architectures. As all industry players adopted superscalar techniques, pressure to gain performance through process improvements resumed.
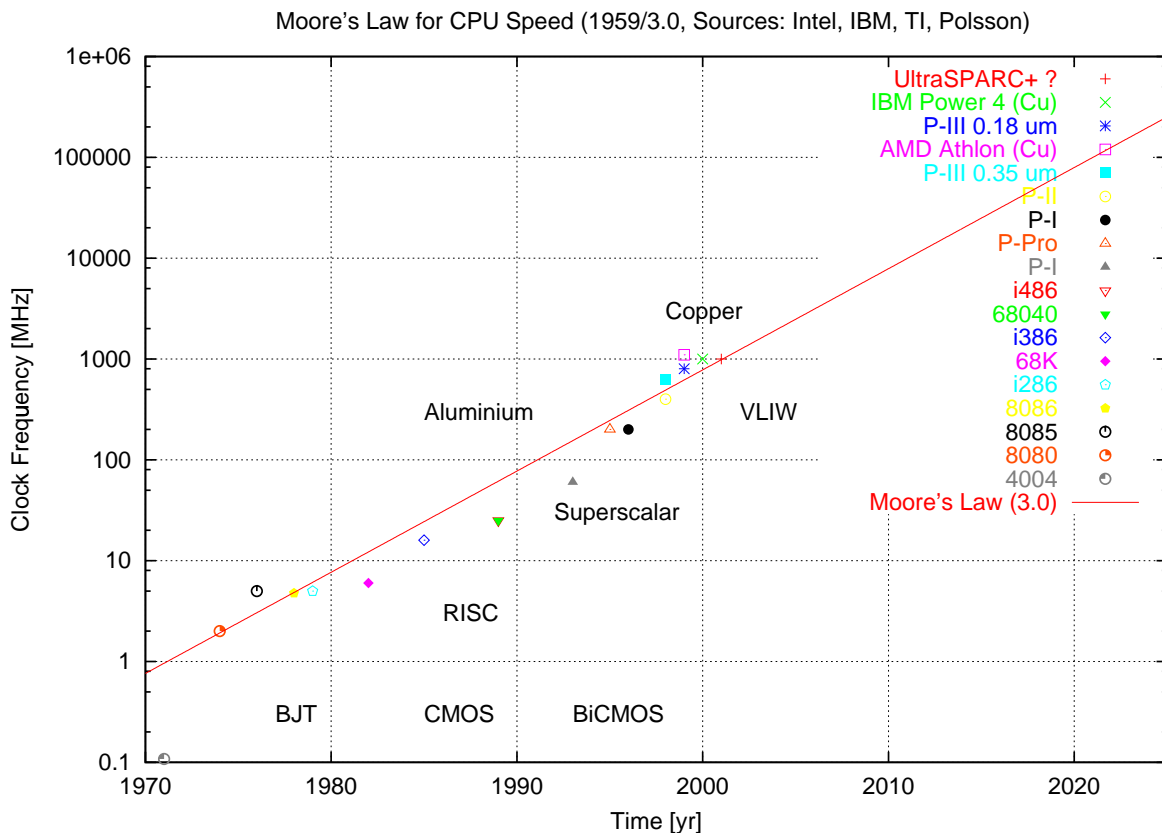
Moore's Law for CPU Speed (1959/3.0, Sources: Intel, IBM, TI, Polsson)

Figure 2: *Moore's Law for microprocessor clock frequencies, assuming a starting point of 1959 and doubling time of 36 months (Author).*

achievable transistor counts per die are bounded by quantum physical effects. At some point, carriers will tunnel between structures, rendering transistors unusable and insulators leaky, beyond some point the charge used to define logical states will shrink down to the proverbial single electron.

Birnbaum cites Motorola projections which indicate that the number of electrons used to store a single bit in DRAM will decline down to a single electron at some time around 2010, thereby limiting further density improvements in conventional CMOS DRAM technology [2]. The bounds for microprocessors are less clear, especially with emerging technologies such as Quantum Dot Transistors (QDT), with sizes of the order of 10 nm, as compared to current MOS technology devices which are at least twenty times larger[4].

It follows that extant and nascent semiconductor component technologies should be capable of supporting further density growth until at least 2010. We cannot accurately predict fur-

---

[4] Good introductions on QDT technology are posted at www.ibm.com and the Department of Electrical Engineering, University of Minnesota.

ther process improvements beyond that time, which accords well with "Mead's Rule" and its projection of 11 years [1].

However, there is considerable further potential for performance growth in machine architectures. To date most architectural evolution seen in microprocessors has been little more than the reimplementation of architectural ideas used in 1960s and 1970s mainframes, minicomputers and supercomputers, made possible by larger transistor counts. We have seen little architectural innovation in recent decades, and only modest advances in parallel processing techniques.

The achievable performance growth resulting from the adoption of VLIW architectures remains to be seen. While these offer much potential for performance growth through instruction level parallelism, they do not address the problems of parallel computation on multiple processors, and it is unclear at this time how well they will scale up with larger numbers of execution units in processors.

It would be naive to assume that we have already wholly exhausted the potential for architectural improvements in conventional Von-Neumann model stored program machines. Indeed if history teaches us anything, it is that well entrenched technologies can be wiped out very rapidly by new arrivals: the demise of the core memory under the onslaught of the MOS semiconductor memory is a classical case study, as is the GMR read head on the humble disk drive, which rendered older head technologies completely uncompetitive over a two year period [1].

It may well be that quantum physical barriers will not be the limiting factor in microprocessor densities and clock speeds, rather the problems of implementing digital logic to run at essentially microwave carrier frequencies will become the primary obstacle to higher clock speeds. The current trend to integrate increasingly larger portions of the computer system on a single die will continue, alleviating the problem in the medium term, however a 2020 microprocessor running at a 60 GHz clock speed is an unlikely proposition using current design techniques. Vector processing supercomputers built from discrete logic components reached insurmountable barriers of this ilk at hundreds of MegaHertz, and have become a legacy technology as a result [9].

Do other alternatives to the monolithic Silicon chip exist? Emerging technologies such as quantum computing and nano-technology both have the potential to further extend performance beyond the obstacles currently looming on the 2010-2020 horizon. Neural computing techniques, Mead argues, have the potential to deliver exponential performance growth with size, rather than speed, in the manner predicted for as yet unrealised highly parallel architectures [13], [9].

It follows that reaching the speed and density limits of semiconductor technology may mean the end of exponential growth in single chip density and clock speeds, but it is no guarantee

that the exponential growth in compute performance will slow down.

What we can predict with a high level of confidence is that Moore's Law will hold for the coming decade, and exponential growth is very likely to continue beyond that point.

4   Information Warfare in the Coming Two Decades

The next two decades will see significant growth in capabilities, both in general computing, networking and military technologies reliant upon digital processing. Predicting exactly the manner and form of such growth is fraught with some risk, as history indicates conclusively, however many changes are clearly predictable and these provide a useful basis for conclusions in the domain of IW[5].

In the area of general computing, which is the application of technology to commercial, engineering, scientific and consumer computing, it is clear that we will see the direct effect of exponential growth in computing speed and storage capacity.

A desktop commodity computer with multiple GigaHertz clock speed processors, Gigabytes or tens of Gigabytes of RAM, and hundreds of Gigabytes of disk storage is a certainty. Therefore the performance traditionally associated with supercomputing systems and large commercial multiprocessor mainframes will be widely and cheaply available.

This has implications in many areas, some of which will be directly relevant to the IW community.

The first area of interest is that of databases and "data mining". Once a database is loaded into such a system, it will be possible to perform very rapid searches and analyses. Therefore it will be feasible to significantly automate and accelerate many aspects of intelligence gathering and analysis. This is especially true of tasks which have traditionally incurred prohibitive expenses in computer time and expensive manpower, such as imagery analysis. The availability of large historical databases of intelligence imagery, and the increasing availability of commercial very high resolution satellite imagery, will provide the capability to set up substantially automated analyses, in which the analyst sets up search parameters and monitors the process.

A side effect of this area of growth is that the substantial monopoly in this area held by governments will be eroded. If the imagery is available, at a price, and the computing assets available as commodity products, virtually any player will be able to acquire such a capability, with a modest effort and expenditure.

---

[5] This section derives substantial portions of its argument from [11], [12].

The same will be true of the analysis of communications intercepts, electronic intelligence (ELINT) signal intercepts and databases, and archived media output. It is entirely feasible that such a system could be connected into a stream of video traffic, such as CNN, and "watch" for information of interest, against some pre-programmed search parameters.

The implications of such technology are indeed manifold. Civilian applications such as environmental and weather monitoring, agriculture, mining and law enforcement are likely to the big beneficiaries, and the commercial drivers of development. However, in the hands of undemocratic regimes, such technology provides a means of economically implementing an Orwellian measure of social control. For the cost of commodity computer equipment, and some software, such a government could monitor all national telephone traffic, all national Internet traffic, and using video cameras and microphones, invasively monitor the workplace and home.

Preventing such abuses of the technology may not be easy, given the unrestricted availability of computing power.

Large amounts of cheap desktop compute power are likely to exacerbate extant problems in cyberwar, since substantial automation may be introduced into the mechanics of the process, thereby making opportunities available to those who lack the high levels of computer literacy characteristic of the incumbent cracker community. We may yet see the Gibsonian paradigm implemented.

Information security will become an increasingly important issue. Leakages of information into the public domain, which today may be lost in the Gigabytes of garbage information continuously posted, published and broadcast, may no longer be safely assumed to be hidden in noise. If every bit of published material in a specific domain is carefully and comprehensively sifted for specific keywords or images, it will become significantly cheaper to profile the activity of individuals, organisations and governments. Any telltale signatures such as purchases, funds transfers and movements will be increasingly difficult to hide. Secrecy and individual privacy stand to be eroded, unless very disciplined operational practices are followed. The model of "strategicheskaya maskirovka" practiced so diligently by the Soviets during the Cold War may become an essential operating procedure for governments, commercial players, organisations and individuals.

The availability of vast amounts of commodity computing power will be a double edged sword in the field of cryptography. It is clear that weaker ciphers in wider use will become wholly ineffective. As shown by the successful brute force attacks performed against DES under the RSA Challenges, the length and robustness of ciphers in military, government and commercial use will become an ongoing issue[6]. It will be essential that cryptography standards be designed from the outset for extensibility, ie the notion of a fixed key length in a standard will become

---

[6] The best discussion of the "DES Cracker" is on the EFF website at www.eff.org, under "Frequently Asked Questions (FAQ) About the Electronic Frontier Foundation's 'DES Cracker' Machine".

obsolete. The one-time-pad may yet see a revival, in the longer term.

We can also expect to see a significant reduction in the cost of performing fast encryption, which will become increasingly commonly used at every level of society. The potential which data mining techniques offer, whether in conjunction with legally or illegally acquired information, means that significant pressures will exist for virtually all traffic to become encrypted.

The infotainment industry, encompassing the media and entertainment industries, is showing an increasing trend to consolidation and an increasing appetite for high quality digital special effects. If the "Titanic", "Starship Troopers" or the latest "Star Wars" were impressive, it is worth considering the longer term implications of the capability to produce such materials quickly and cheaply on a commodity desktop machine, in conjunction with further refinements of the software involved.

For a propagandist feeding "black propaganda" into an indiscriminate and sensation hungry infotainment industry, such technology is a gift from the gods. Why wait for a cruise missile to go astray and hit an orphanage, if you can synthesize the video evidence on your desktop? Why wait for the opposing political leader to make an inflammatory statement, if you can synthesize it on your desktop?

If current technological trends are continued, within the next two decades it will be feasible to synthesize on the desktop forgeries of video footage which are going to be very difficult for experts to distinguish from a genuine original, and impossible for the lay observer. Distinguishing the truth may become very difficult indeed.

As shown by recent media coverage of issues such as the TWA 747 crash, or the "Tailwind Affair", there are strong commercial incentives for infotainment industry players to accept a poor standard of truth to create a sensation. Given the ample availability of individuals with peculiar personal agendas, extremist special interest groups, fundamentalist religious groups and extremist ethnic lobbies, there is no shortage of potential users for such systems.

Gutenberg's printing press played a pivotal role in the propaganda wars between the reformation clerics and the Papacy. Today we are seeing analogous use of the Internet by lobbyists and individuals with agendas, wishing to propagate their particular cause.

Commodity computers of the coming two decades will without doubt become a potent capability for such players. It is worth contemplating the implications of a "Super-Eliza", programmed to connect itself to a Usenet newsgroup and argue a particular agenda. Judging from my own experience in Usenet debates, many individuals I encountered could be readily replaced by such a "propaganda engine". The inability to pass the "Turing Test" need not be a hindrance in such games !

5   Embedded Military Systems

We can expect to see developments in military technology which parallel those in the civilian domain. Current trends in military acquisition are away from dedicated military computer standards and architectures, to allow the exploitation of Commercial Off The Shelf (COTS) processing power [16], [3].

For many military applications, commercial hardware can be readily adapted by ruggedisation, allowing military systems to keep pace with the commercial world. However, many military environments are so harsh and demanding of space and cooling, that ruggedised COTS technology is not a feasible proposition.

The recent "Bold Stroke" and "Oscar" trials of "militarised" VMEbus hardware in Class II airborne applications have yielded very encouraging results, using late generation IBM/Motorola PowerPC RISC processors [8]. The release of the Intel Pentium I architecture for use in Milspec chipsets is another very encouraging step.

Despite much effort expended in recent years, we are still seeing a significant time lag between the introduction of a microprocessor in the commercial market, and its integration into embedded military systems rated for harsh environments. This time lag must be substantially reduced, since the exponential nature of Moore's Law means that even with a constant time lag of N years, the gap between commercial and military computing capabilities will continue to grow.

The principal obstacles to the use of microprocessor dies designed for commercial use in embedded military systems are the robustness of the electrical, mechanical and thermal designs, and the lack of prior usage histories required to establish reliability behaviour. The traditional approach to this problem is to redesign the die where necessary, to meet the more demanding Milspec environmental specifications[7]. Given the cost of such a redesign, and the need to qualify the component, it is questionable whether this strategy can be sustained in the longer term.

The increasing disparity between the life cycles of commercial microprocessors and embedded military systems exacerbates the problem. A particular variant of a microprocessor, using a given die design, fabrication process and package may remain in production for a period of less than a year, before it is replaced with a different variation on the theme. The only portion of the design which retains a measure of stability is the basic architecture, which is bound to the established base of software. On the other hand, military platforms tend to remain in service for decades: the 1950s technology B-52 may remain in front line service past 2030 ! For many such platforms, basic software systems may remain in operation, with ongoing upgrades, for decades.

---

[7] See Chapter 4 in [3].

Figure 3: *The USAF's new F-22A ATF employs up to three Common Integrated Processor (CIP) multiprocessors, built around the JIAWG backplane bus. This allows most of the functions performed traditionally in hardware to be implemented as software running on the CIPs. The CIPs employ liquid cooled variants of the Intel i960 RISC microprocessor and VHSIC DSP chipsets, with optical fibre links used for transferring data between the CIPs and sensors (Lockheed-Martin/Boeing).*
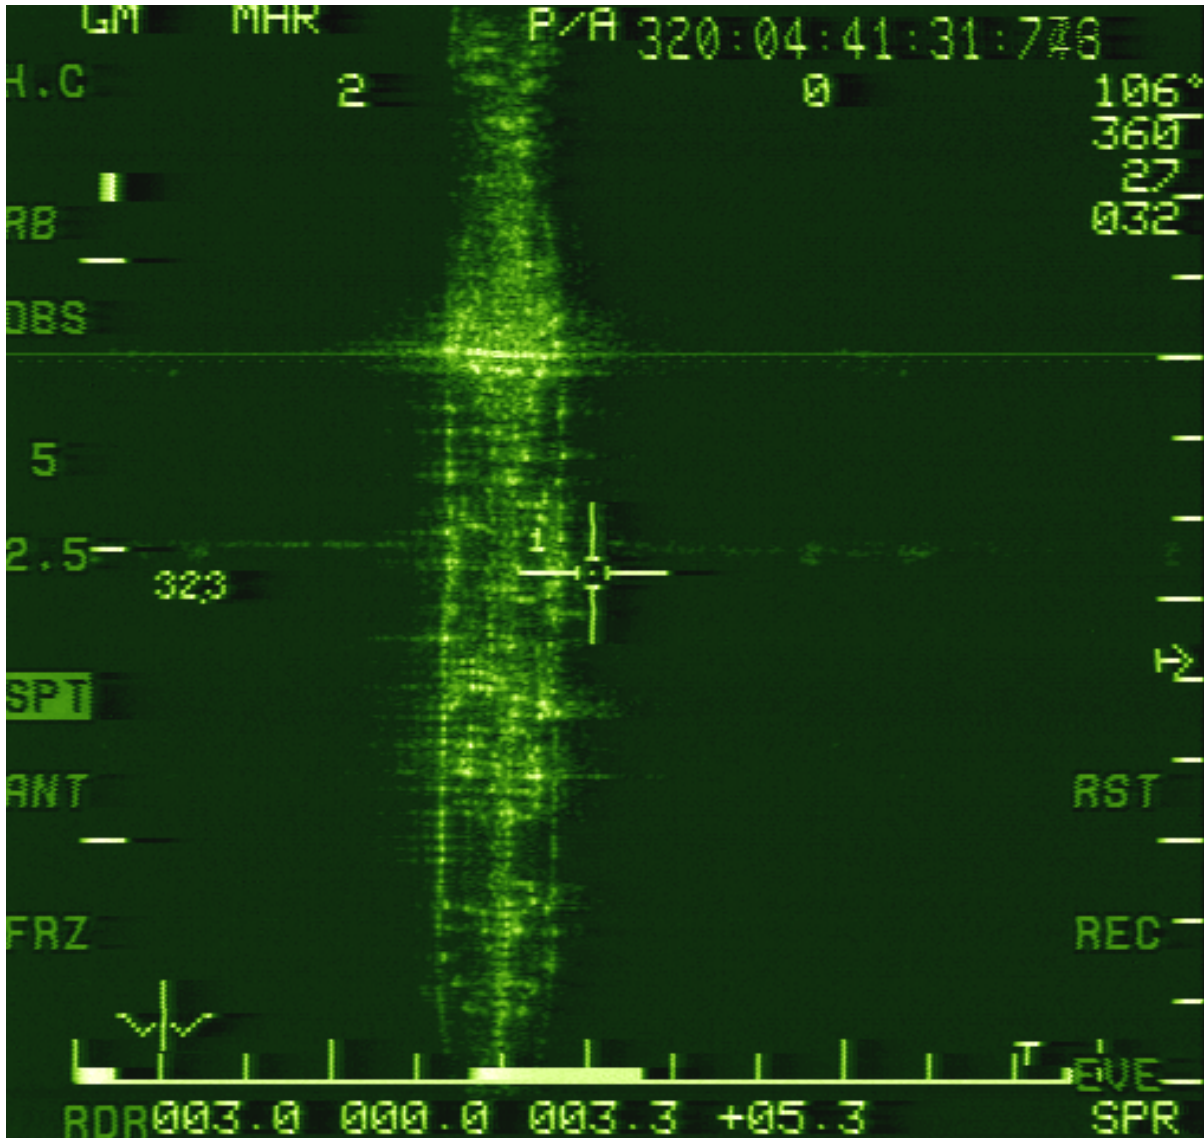
Figure 4: *High resolution Synthetic Aperture Radar imagery produced by the APG-76 multimode attack radar. This radar, designed for an F-4E retrofit, relied on older generation processors, yielding a total weight of 625 lb. A modern reimplementation using current generation processors would be not only more capable, but substantially smaller (Northrop-Grumman ESSD).*

If the gap between commercial and embedded military processors is to be narrowed, if not eliminated, it will be necessary to adopt a more effective strategy for militarising the components. This will not be an insurmountable problem in the area of packaging design and cooling to withstand Milspec environments, but may prove to be an issue with component geometries on the die proper, especially with very high density components.

The ideal outcome would be the adoption of a new packaging and cooling arrangement which would be readily adaptable for use with volume production commercial dies, these being screened for suitability at the production phase in the traditional manner.

Even should measures to narrow the gap not yield an ideal outcome, it is clear that the exponential growth effects of Moore's Law will benefit many established military applications.

Perhaps the greatest benefit will be seen in signal processing and data processing, especially in sensors such as radar, ESM, RWR, FLIR, IRS&T and communications equipment [6].

The established approach to radar applications has been the use of Milspec rated variants of commercial digital signal processing (DSP) chips, and some custom Milspec DSP chips. Whether this strategy is viable in the longer term remains to be seen, since the performance of commodity general purpose processors is becoming very competitive against specialised DSPs. While a DSP is architecturally optimised for the application and will always perform better than a general purpose processor of similar complexity, running at the same clock speed, the question is whether the economics of production volumes will be able to sustain the DSP in the longer term. Cost pressures in hardware and development tools will increasingly favour the general purpose processor.

The F-22 and JSF avionic architectures depart from the established model of "federating" black boxes over low speed busses, and concentrate all sensor processing on the platform into centralised multiprocessors[8]. Should such multiprocessors be designed from the outset to accommodate board level upgrades of processor and memory modules, it will become increasingly difficult to justify dedicated chip architectures in the longer term.

Vast amounts of computing power will yield important dividends, by allowing the use of algorithms which are too resource hungry for current applications, and also by allowing a higher measure of concurrency in operating modes.

Groundmapping and attack radars have largely shifted from real beam mapping techniques to synthetic aperture techniques, over the last decade and a half. Ample compute cycles allow for real time SAR imaging at very high resolutions, antenna and RF hardware permitting, down to inches. Interferometric SAR offers the potential for 3D imaging of terrain, but hitherto required prohibitive numbers of compute cycles [18], [17]. This barrier is likely to fall during

---

[8] See Figure 3, [16]. The architecture of the CIP is designed for easy extensibility by adding further SEM format processor modules, or replacing such with faster variants.

the next several years, resulting in attack radars with the capability to produce very high resolution 3D SAR imagery.

Another capability which may result from the combination of very high resolution and 3D imaging is the automatic detection and classification of targets by shape. Unless wavelength tuned radar opaque camouflage is used, surface assets will be unable to hide effectively.

Ground Moving Target Indicator (GMTI) radars are beginning to proliferate, mostly in wide area surveillance systems, but also in attack radars (APG-76, Figure 4, also see [18].). Algorithms which can analyse the fine Doppler signature of a ground target for identification and classification have been demonstrated, and with abundant compute power can be significantly improved upon.

Inverse SAR techniques capable of providing coarse images of a target's shape have been available for some time in maritime search radars, with ample compute cycles these can be further improved upon. Importantly, combining such techniques with GMTI signature analysis will provide a substantial capability to detect, identify and classify surface targets. If high resolution 3D SAR techniques are also employed, a surface target will be difficult to hide whether it is on the move, or static[9].
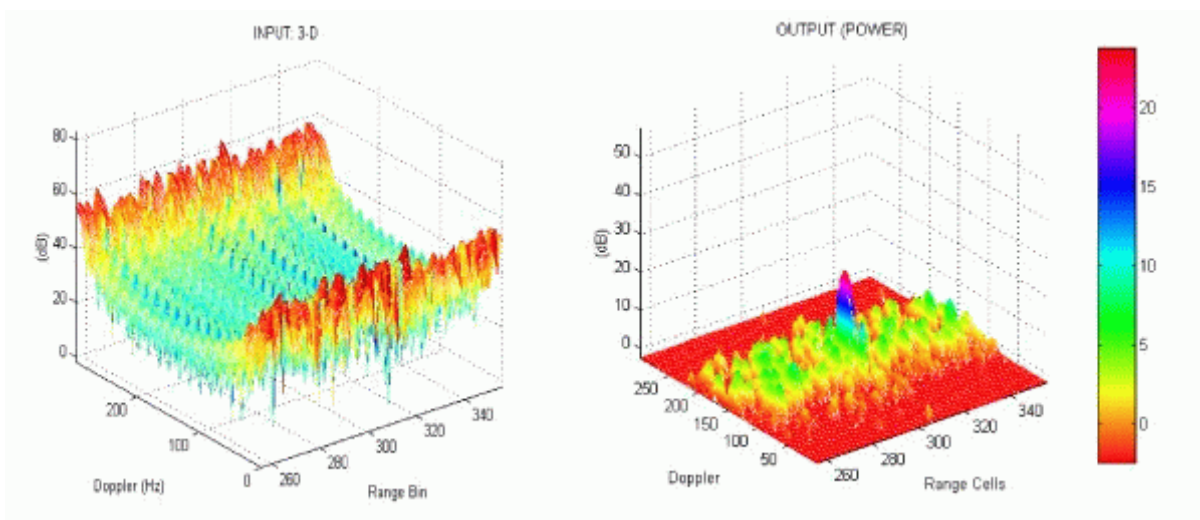

Figure 5: *Space Time Adaptive Processing (STAP) techniques are extremely effective at removing clutter in pulse Doppler radars but have hitherto been impractical due to prohibitive demands for compute cycles. This signal processing technique will become feasible in the coming decade (USN NRL).*

Counter Air operations rely heavily upon Doppler techniques to detect airborne targets against

---

[9] An analogous capability for air-air operations can be provided by Ultra-High Resolution (UHR) techniques, which exploit direct spreading techniques for very precise range cell analysis [5].

clutter backgrounds. Space Time Adaptive Processing (STAP, Figure 5) techniques have the potential to largely remove the effects of clutter, and jamming sources, but have hitherto been impractical due to prohibitive demands for compute cycles. This barrier is likely to also fall over the coming decade, denying the use of clutter to opposing aircraft.

Another important trend seen in recent years is the shift from traditional mechanically steered antennas (MSA) to active electronically steered arrays (AESA) [18], [15], [4], [5]. The AESA offers enormous potential, allowing the antenna to be shared between multiple concurrent modes and allowing its use for tasks such as in-band jamming and high gain interferometric direction finding, as well as pencil beam high data rate communications. All of these capabilities will rely upon the availability of compute cycles, which are the primary enabler within the system design. Importantly, such antennas provide the means of implementing Low Probability of Intercept (LPI) modes[10].

The importance of concurrency should not be undervalued. The ability to correlate multiple target signatures acquired in multiple operating modes of the radar will produce a disproportionate payoff in capability. Concurrency is contingent upon ample computing power.

Surveillance platforms such as the AWACS and JSTARS, and their smaller cousins, will see such technology introduced first, since they can readily accommodate racks of COTS processors.

We can also expect to see commensurate improvements in other sensors, such as FLIR, IRS&T, DIAL Lidar and ESM/RWR, as well as missile seekers, since compute cycles will cease to be the barrier to capability. Indeed, the rate at which algorithms are developed and refined, and sample target signatures acquired, is likely to become the limiting factor in operational deployment.

A great beneficiary of Moore's Law will be sensor fusion, be it of the sensors on a single platform, or networks of sensors on multiple platforms. Sensor fusion techniques are in relative terms, in their infancy, and are limited in speed primarily by the availability of compute cycles. The ability to store large amounts of real time sensor data and reference databases of signatures in massive RAM arrays, coupled with cheap compute cycles, offers much long term potential in this area.

The most general conclusion to be drawn is that the signatures of military platforms will become the deciding factor in their military viability, since Moore's Law will result in sensors and systems of fused sensors capable of sifting targets from background clutter, camouflage and jamming with great effectiveness. Stealthiness will become vital, those platforms which are not stealthy become virtually unusable [6].

Moore's Law will produce other effects. One is the potential for increasing levels of intelli-

---

[10] Such modes incorporate spread spectrum techniques and pseudorandom scan patterns, where appropriate [6].

gence in munitions and autonomous platforms. Weapons and UAVs may acquire substantial abilities to independently seek out targets, and evade detection and engagement. While the UCAV has been proposed as a defacto replacement for manned aircraft, we have yet to see the development of artificial intelligence techniques capable of providing UCAVs with more intelligence than that of an insect. Until this software problem is solved, compute cycles will not confer the sought capabilities, especially flexibility, which we find in crewed platforms. A genuine "robot warrior" will require a major breakthrough in computer science, the timing of which we cannot predict at this stage [7].

6   Conclusions

The exponential growth we are observing in computer performance and capabilities is transforming the world around us, in a systemic and unrelenting fashion. The aggregate effect of this transformation will be increasing instability in areas ranging from foreign and domestic politics, across commerce and economics to warfare. In such an environment, the model of fixed medium to long term planning in force structures, capabilities, doctrine and strategy is unsupportable.

It follows that a fundamentally different adaptation is required in order to survive and prevail in such an environment. This adaptation is the ability to evolve faster in technology, and operational doctrine over potential opponents. Indeed it is worth stating this as an axiom:

"The player who can evolve technology and doctrine faster than an opponent, all other things being equal, will prevail" [6].

The implications of this are many, but the essential substance is simple. If we are developing a system or a military platform, it must be designed from the outset to continually incorporate faster computing hardware at the lowest possible incremental cost, over its operationally useful life. Evolution must be part of the initial design specification and functional role definition.

If we are developing a doctrine, it must accommodate the ability to adapt to changes in the underlying technological environment, as well as changes in the broader environment resulting from technological evolution.

Attempting to employ long term fixed doctrine, regulations or legislation to confront changes resulting from evolution is utterly futile. The only thing which will remain fixed is ever increasing change.

Evolution is a part of nature, and to deny the primacy of nature's dictates is to deny reality.

References

[1] Bell C.G. . The Folly/Laws of Predictions 1.0. "The Next 50 Years", ACM97 Conference Talks, Association for Computing Machinery, 1997.

[2] Birnbaum J. Computing Alternatives. "The Next 50 Years", ACM97 Conference Talks, Association for Computing Machinery, 1997.

[3] DY4 . DY4 Harsh Environments COTS Handbook. Comp. of Papers and Tech. Reports, DY4 Systems, Inc., 1998.

[4] Fulghum D.A. New F-22 Radar Unveils Future. *Aviation Week & Space Technology*, 152:6:50–51, February 7, 2000.

[5] Fulghum D.A. F-22 Radar Ranges Pass Expectations. *Aviation Week & Space Technology*, 152:11:26–26, March 13, 2000.

[6] Kopp C. Sensors, Weapons and Systems for a 21st Century Counter-Air Campaign. Invited paper, Proceedings of the "Control of the Air: The Future of Air Dominance and Offensive Strike" conference, Australian Defence Studies Centre, Australian Defence Force Academy, Canberra, 15-16 November, 1999.

[7] Kopp C. Waypoints: New Technologies and War-Fighting Capabilities,. *USAF Air Power Journal*, Fall 1996:111–114, 1996.

[8] Kopp C. Flying High with VME. *Systems: Enterprise Computing Monthly, Auscom Publishing, Pty Ltd, Sydney, Australia*, November:32–40, 1999.

[9] Kopp C. Supercomputing on a Shoestring:Experience with the Monash PPME Pentium Cluster. *Conference Paper, AUUG 99, AUUG Inc, Sydney, Australia*, September 7:(Not included in proceedings due to late submission), 1999.

[10] Kopp C. GigaHertz Intel Processors: Xeon and Athlon. *Systems: Enterprise Computing Monthly, Auscom Publishing, Pty Ltd, Sydney, Australia*, April:TBD, 2000.

[11] Kopp C. Information Warfare: A Fundamental Paradigm of Infowar. *Systems: Enterprise Computing Monthly, Auscom Publishing, Pty Ltd, Sydney, Australia*, February:46–55, 2000.

[12] Kopp C. Information Warfare: Current Issues in Infowar. *Systems: Enterprise Computing Monthly, Auscom Publishing, Pty Ltd, Sydney, Australia*, March:30–38, 2000.

[13] Mead C.A. . Semiconductors. "The Next 50 Years", ACM97 Conference Talks, Association for Computing Machinery, 1997.

[14] Mead C.A., Conway L.A. *Introduction to VLSI Systems*. Addison Wesley, Reading, Massachusetts, 1980.

[15] Nordwall B.D. F-15 Radar Upgrade Expands Target Set. *Aviation Week & Space Technology*, 151:24:40, December 13, 1999.

[16] Shirley F. . RASSP Architecture Guide, Rev.C. ARPA Doc.No.AVY-L-S-00081-101-C, Lockheed Sanders, Inc, Hughes Aircraft, Motorola, ISX Corp., April 14, 1995.

[17] Skolnik M.I. *Radar Handbook, Second Edition*. McGraw-Hill, New York, 1991.

[18] Stimson G.W. *Introduction to Airborne Radar, Second Edition*. Scitech Publishing, Mendham, N.J., 1998.

Biography

Born in Perth, Western Australia, the author graduated with first class honours in Electrical Engineering in 1984, from the University of Western Australia. In 1996 he completed an MSc in Computer Science by research and more recently, submitted a PhD dissertation dealing with long range data links and ad hoc mobile networks. Both theses were completed at Monash University in Melbourne.

He has over 15 years of diverse industry experience, including the design of high speed optical communications equipment, computer hardware and embedded software. More recently, he has consulted to private industry and government organisations, and now lectures on computing topics at Monash University. He has been actively publishing as a defence analyst in Canberra based *Australian Aviation*, since 1980, and more recently in Amberley based *Air Power International*, the UK based *Janes Missiles & Rockets*, and the US based *Journal of Electronic Defense*.

His work on electronic combat doctrine, electromagnetic weapons doctrine, laser remote sensing and signature reduction has been published by the *Royal Australian Air Force* since 1992, he has contributed to *USAF CADRE Air Chronicles*, the *InfowarCon* series of conferences and

the second edition of Schwartau's *Information Warfare*. His most important work to date has been the development of a doctrinal model for the offensive use of electromagnetic weapons in strategic information warfare.