by Dr Carlo Kopp, PEng

IS THE JOINT STRIKE FIGHTER RIGHT FOR AUSTRALIA?

PART 2 – JSF V RISK FACTORS

The F-35 Joint Strike Fighter (JSF) is one of the most technologically ambitious aircraft development programs ever seen, in many respects more ambitious than the TFX program which realised the F-111.

This ambition offers the promise of a battlefield interdiction and close air support optimised fighter with survivability and lethality well beyond that of the F-16C, A-10A, F/A-18A-D, AV-8B and UK Harriers it is designed to replace. The flipside of this payoff is that a considerable number of risk factors come into play, potentially affecting costs, timelines and the ultimate capabilities of the production JSF.

For Australia these risk factors combine with the deeper and more fundamental issues arising from the intended use of a survivable battlefield interdiction and close air support fighter in the more challenging roles of 'air dominance fighter' and 'deep strike fighter', missions which impose their own unique needs on combat aircraft. As Sukhoi numbers grow across Asia, Australia will face over coming decades the most competitive region worldwide, with the statistically newest fleet of third generation fighters in service worldwide.

There can be no doubt the strategy of early commitment to a new fighter has its merits as an ambit claim to lock down future defence funds, which otherwise could be gobbled up by competing programs from the Army and Navy. Buying into SDD – System Development and Demonstration – provides some sectors of Australia's industry, especially in component manufacture, access to a potentially huge market. Australia also gets to sit in on development team meetings, gaining an opportunity to learn much about the technology base used in the F/A-22A and JSF.

The early commitment strategy however has its drawbacks as well. The first is that the RAAF must politically defend a massive burst of single service expenditure in the 2012 to 2020 timeframe – with early outlays beginning post 2006. In the face of intense inter service budgetary competition, other parts of the RAAF could suffer badly as a result, sacrificed to protect the JSF. To what extent the early F-111 retirement is a result of this is yet to be known.

A second problem is the degree of access Australia actually gets by SDD buy-in, especially in key areas like stealth, engine hot end technology, AESA (Active Electronically Scanned Array) radar and software. Unless personnel with suitable engineering/science backgrounds and experience are engaged to exploit the gathered data in depth, it may contribute little useful value.

The industry benefit may also prove illusory, in that the highest value added systems integration and software sector of the industry gets a much smaller bite than the hardware manufacturing sector, who in turn must compete against overseas peers to retain their workshare. The worst case outcome – a risk in its own right – is that the manufacturers end up with very little, the Commonwealth with little technology transfer, and the RAAF gets stripped to the bone over the next decade fending off Army and Navy demands for budget.

The RAAF has lost out in the internal budgetary game in recent times – last year's Defence Capability Review saw the RAAF lose the F-111 for no gain in AEW&C, tankers or other 'tier one' assets. The Army gained Main Battle Tanks, the Navy's air warfare destroyers and support ships were confirmed, but the RAAF lost the F-111.

At the most fundamental level the RAAF faces two key challenges in replacing the F-111 and F/A-18. The first is in choosing technology which is relevant 40 years hence, effectively ruling out evolved third generation fighters like the Rafale, Eurofighter, F-15E and F/A-18E. The second is in maintaining the relative advantage Australia enjoyed over the broader region for the last 20 to 30 years, by virtue of the F-111 and F/A-18A in its earlier life. In an increasingly competitive region aiming for a low target capability in replacing the existing fleet will guarantee an inferior strategic position in one to two decades' time, if not earlier.

STEALTH CAPABILITY ISSUES

The JSF is the first ‘stealth fighter’ intended for export, and we can expect that production F-35s will be delivered in ‘high stealth’ (US) and ‘low stealth’ (export) configurations, differing in the performance and application of radar absorbent and lossy materials. In an environment where every ally is clamouring for the ‘high stealth’ model, it might be politically very tricky for Australia to get access to the full stealth potential of the aircraft when other US allies are barred from doing so.

The stealth capability in the JSF is designed for low cost and maintainability, rather than best possible stealth performance at any cost. Stealth is achieved by a combination of shaping, detail design and absorbent/lossy materials. While detail design and materials can evolve over the life of a design, and be upgraded incrementally to match an evolving threat, airframe shaping is fixed and whatever limits it imposes are unchangable.

The JSF’s stealth design is optimised by shaping for the ‘narrowband’ X-band and Ku/K/Ka-bands, which fits the most likely threats US operated JSFs will encounter – highly mobile battlefield air defence weapons and fighter air intercept radars. The serrated nozzle and inlet design reflect this optimisation – with increasing radar wavelength both will progressively lose effectiveness. The inlet tunnels use S-bending and absorbent materials, while the tailpipe is claimed to use a blocking structure, both most effective against the X-band. The planform and edge alignment is much less disciplined than that in the F/A-22A or YF-23A, again less critical for an X-band threat confined mostly to the fore/aft sectors.

US Air Force thinking is that the JSF is used to demolish battlefield ground targets once the F/A-22As have broken the back of the air defence system and opposing fighter force – in effect the long range S-band, L-band, UHF and VHF radars have been killed off by F/A-22As, as have the opposing L-band or S-band AEW&C systems.

In this environment the greatest risk is presented by opposing fighters hunting with minimal or no ground radar or AEW&C support, and mobile AAA and SAM systems like the Roland, Crotale, Rapier, 2K12/9M9 (SA-6), 9K33 (SA-8), 9M37M (SA-11), Tor M1 (SA-15) and ZSU-23-4P. Such SAM/AAA systems typically use the C, X and Ku bands for their search and engagement radars, and X or Ku bands for missile guidance. For such ‘shoot and scoot’ high mobility surface threats and fighter threats the JSF’s stealth optimisation will work very nicely.

The technological design features of a fighter can be divided by the rate at which they evolve over time. The smartest long term choices are always those which put the highest priority on design features which cannot be altered once the aircraft is in service, accepting that rapidly changing technologies will be replaced over the life of the aircraft. The most attractive aspects of the JSF are all in areas which rapidly evolve, whereas its least attractive aspects are in areas which cannot evolve. From a technological strategy perspective the JSF is a very poor choice long term compared to the F/A-22A (Author).

Lowest Priority – Can Change Quickly By Upgrades	Rapidly Evolving Design Features	Slowly Evolving Design Features	Fixed Design Features	Highest Priority – Cannot Change Once Acquired
	Computer Hardware	Engine Technology	Airframe Size/Weight	
	Computer Software	Stealth Materials	Aerodynamics	
	Cockpit Displays	Fuel Systems	Stealth Shaping	
	Weapons	Hydraulic Systems	Int Fuel Capacity	
	ElectroOptical Sensors	Structural Materials	Radar Aperture	
	Datalinks/Nav/Comm	ECS/Cooling Systems	EO Apertures	
	Radar Processing	Electrical Pwr Systems	Engine Massflow	

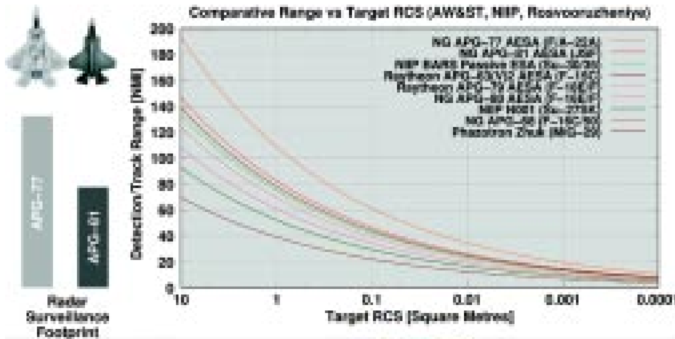

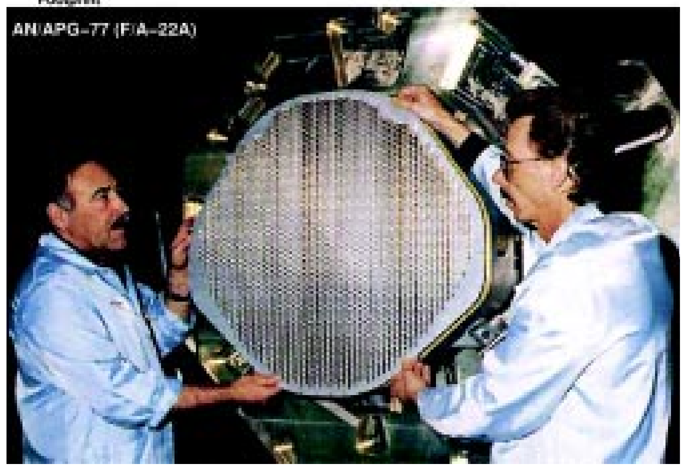

This chart compares publicly available performance figures for a range of current radars, including intended performance for the JSF's APG-81 AESA. While the higher power rating of the JSF radar makes it highly competitive against the older technology passive array in the current Su-30, the introduction later this decade of active array technology in the Sukhoi will tip this balance decisively. The F/A-22A's APG-77 has an unassailable lead which it will retain longer term (Author).

For the RAAF, which intends to use the JSF to replace the F-111 in its ‘deep strike’ (strategic land strike) role and the F/A-18 in the air combat role, the X-band oriented optimisation of the JSF is a poor fit. In both roles this optimisation will frustrate opponents using X-band engagement and fire control radars, but leaves a major vulnerability in the lower bands, occupied by static or semi-mobile early warning, ground control intercept and acquisition radars, as well as AEW&C radars.

The availability of Russian beyond visual range missiles with very modern infrared seekers and heatseeking adaptations of area defence SAMs like the SA-6 presents a situation where the JSF could be engaged at a respectable distance, despite its good X-band stealth capability. Sukhoi Su-27/30 fighters could be vectored into a firing position without having to light up their X-band radars, or SAM sites cued in a similar fashion.

This is the pitfall of economy ‘narrowband’ stealth – it can defeat upper band radars used for the engagement control, but is much less effective in defeating the long range systems used to acquire targets. If an Su-30 can be positioned close enough, it can engage the JSF regardless of stealth, and with a kinematic and missile performance advantage the odds are unlikely to favour the JSF.

While having any real stealth always beats having no stealth, Australia should not develop unrealistically high expectations of the JSF’s stealth capability, especially in relation to the principal regional capabilities like the Su-27/30, A-50 AEW&C, S-300 and supporting long range radar systems. The only fighter optimised for that threat environment at this time is the F/A-22.

The big wildcard in longer term US Air Force force structuring will be the FB-22A, currently a theoretical concept for a stretched delta wing F/A-22A derivative heavy strike fighter. Sized around the F-111, with a 1500nm (2780km) class radius, the FB-22A would achieve a high level of commonality with the basic F/A-22A. At the recent AFA symposium Gen J P Jumper, US Air Force CAS, presented a scenario in which FB-22A development would start in FY 2004, initial deliveries happening in FY 2011, and full rate production in FY 2016, with an initial build target of 150 FB-22As to supplement the currently planned 381 F/A-22A strike fighters – all 381 now counted as strike assets (Author/USAF).

AVIONICS CAPABILITY ISSUES

The JSF builds extensively upon the experience gained with the F/A-22's JIAWG (Joint Integrated Avionics Working Group) core avionics system, an implementation of the Pave Pillar model. It is built around three liquid cooled fault tolerant Raytheon Common Integrated Processors (CIP), each originally using a mix of DoD VHSIC custom processors and i960 chips on SEM-E format modules. The system effectively absorbs all of the processing tasks historically distributed across boxes in the radar, EW equipments, comm/nav equipment, main mission computers and cockpit display processors where used.

The aim of this model was to produce a system which could be rapidly upgraded in processing power by the addition or replacement of standardised processing modules, yet providing the ability to flexibly allocate processing power as needed by specific system functions, all implemented in software. The F/A-22A system set a record for software complexity in a fighter, with around 2.5 million lines of software source code cited. The system departed from the historical use of low speed Mil-Std-1553B busses, using the high speed Fibre Channel-Avionics Environment (FC-AE) serial bus for high speed internal interconnects.

The F/A-22A is the first aircraft to exploit this highly flexible and powerful avionics model, one which is inherently designed to ride on the back of Moore's Law (of processor speed doubling every three years). It has also been the first design to fall foul of processing chip evolution outrunning the system's development cycle, and the sheer complexity of the software creating major delays to production in its own right.

The recently redesigned 'CIP 2000' configuration uses up to 66 commercial based Motorola/IBM PowerPC RISC (ie Apple Mac compatible) and Intel i960MX processor chips and is aimed at cost reduction and supportability, with a follow on upgrade planned to further increase computing power. Since the 'G4' variant, PowerPC chips typically include an embedded 'Altivec' short vector processor which is exceptionally well suited to signal processing tasks, as found in radar, comms and EW processing.

The JSF avionics suite is built around an evolution of the F/A-22A model, but is much more complex in implementation due to the additional, and extensive, electro-optical suite and digital 'soft' cockpit. Its liquid cooled Integrated Core Processors (ICP) are intended to be a cheaper equivalent to the F/A-22A CIP, relying to a greater extent on commercial packaging technology. Like the F/A-22A, the JSF is expected to use high speed FC-AE serial buses (replacing the originally planned IEEE SCI/RT – a commercial flop) in the JAST Pave Pace model, supplemented by Firewire bussing (also used in Apple computers) in the Vehicle Management System (VMS).

For JSF System Development and Demonstration, the Mercury RACE++ Powerstream processor will be used for signal processing and I/O processing functions (this is a 9U VME format packaged multiprocessor, built around PowerPC RISC processors – essentially a bigger and faster cousin to the 6U VME packaged PowerPC processors now being used in F-15E, F/A-18E/F and F-111C Block C-4).

The core avionics system, centred in the Integrated Core Processors and their software, will present some significant development risks. While VME packaged PowerPC hardware is now widely used, it has not been used on the massive scale of the JSF to date. The large number of interconnects, density of hardware, and the demanding thermal cycling and vibration environment has the potential to produce reliability problems, especially of the intermittent variety, in the ICP subsystem. This may not become statistically obvious until a good number of systems are operationally deployed – cyclic wearout problems in printed circuit boards and connectors often resemble the

The JSF's Electro-Optical Sensor System (EOSS), comprising the ventral Electro-Optical Targeting System (EOTS) and spherical coverage Distributed Aperture System (DAS), coupled via digital processing to the Helmet Mounted Display System (HMDS) and single panel cockpit display, represents the most comprehensive – and complex – electro-optical package ever installed in a combat aircraft. While the EOTS is a repackaged Sniper XR pod derivative, conceptually closest to the F-117A's DLIR/FLIR package, the EO DAS is entirely new. Its aim is to provide spherical day/night IR coverage to facilitate target acquisition and evade threats, especially heatseeking missiles. The EOSS is primarily aimed at close air support and lower altitude battlefield interdiction roles, a result of US Air Force and Marine Corps inputs to this traditionally dangerous regime of operations (LM/CMC/VSI).

From a simple risk perspective, the much more mature F/A-22A presents far fewer headaches than the JSF does – both in terms of meeting long term capability needs, and in terms of program stability post 2010. Currently in low rate initial production, most of the initial build of around 300 F/A-22As will be completed in the 2012 to 2015 timeframe. At the time of writing the F/A-22A had just been cleared for Dedicated Initial Operational Test & Evaluation (DIOT&E), with deliveries underway to the first operational squadron, at Tyndall AFB, Florida. (LM)

behaviour of airframe fatigue damage and will not manifest until some number of cycles is accrued.

The F/A-22A's Milspec hardened SEM-E packaged system was reported to have had a number of hardware reliability problems, initially misdiagnosed as software faults – the JSF's more complex and softer commercial derived ICP has the potential to do the same on a larger scale.

A less obvious issue for the JSF will be achieving genuine 'open systems' standards compatibility throughout the ICP package and bussing. There will be a temptation to get better performance by using proprietary enhancements to commercial standards, opening a Pandora's box of longer term support issues with single source Silicon and interfaces embedded in the system.

Software has proven to be the single biggest headache in the F/A-22 development program, and the JSF with twice as much, is apt to make for twice or more the headache, regardless of lessons learned in the F/A-22. Large realtime systems on multiprocessing computers present some interesting theoretical and practical problems, especially in scheduling computing tasks and guaranteeing shared data consistency and synchronisation – many are considered analytically intractable (the author has both practiced in industry and lectured at university level real time software system design, software/systems reliability engineering, and computer internal architectural design).

Sheer complexity is a problem in its own right, typically software bug counts in systems of this complexity increase at a rate faster than the increase in the size of the code, as more software components have opportunities to interact adversely. While cockpit control, radar signal processing, EW processing, and comm/nav functions are likely to be less troublesome, the big question will be the bugginess or otherwise of the DAS (Distributed Aperture Systems) functions, data fusion functions, and offboard data networking software. Additional difficulties will arise in testing technique to validate the system. Odds are the software will be one of the biggest sources of development cost and time overruns in the latter phase of SDD and LRIP.

A related risk factor will be whether Australia is permitted access to the full software functionality, and whether source code and development systems will be provided for local enhancements and bug fixes.

The primary sensors, the APG-81 AESA radar and EOTS (electro-optical targetting system) present much lesser risks as they ride on the back of the F/A-22A APG-77, F-16E/F APG-80 and F-16/F-15E Sniper XR programs – the bigger issue for both is long term growth potential. Aperture size in the EOTS will set bounds on growth in long range detection performance. For the AESA, the bigger issue for growth will be the aircraft's cooling capacity – the physics of high linearity RF amplifier design in AESAs result in around 55% or more of the power pumped into the AESA coming out as waste heat via the liquid cooling system. Waste heat management has been an ongoing and frequently reported issue in the JSF program. Significant detection range improvements, or X-band jamming power improvements, may well be limited by the aircraft's systems rather than available AESA technology.

The X-band jamming capability planned for the APG-81 may run into similar issues as expected with the X-band optimised stealth capability – most key regional threat systems may sit well outside the frequency band coverage of the antenna design.

AIRFRAME AND PROPULSION ISSUES

As with the avionics suite and stealth capability, the airframe and propulsion package of the JSF faces some technological risks in implementation, yet concurrently the role specific optimisations of the design may not mesh well with the much broader range of roles to be performed by the RAAF using JSFs.

In terms of the airframe, the biggest development issue will be in containing the empty or basic weight of the aircraft (refer March AA Newsdesk). Excess dead weight will exact penalties in performance, be it agility, range or weapon payload at range. Techniques for reducing excess weight can include reductions in structural weight, at the expense of G-limits or airframe fatigue life, reductions in internal fuel payload at the expense of range/endurance, or reductions in the size of the avionics suite. All essentially amount to reductions in aircraft capability.

The alternate path is the use of stronger, more exotic and expensive structural materials to retain capability at the expense of cost. Both the Su-27/30 and F/A-22 use large amounts of titanium alloy for this reason.

US reports published late last year indicated that a worst case 2270kg (5000lb) excess weight could have arisen – during that period aggressive weight reduction measures are claimed to have slashed 1545kg (3400lb) of excess weight. One weight saving measure cited was achieved by changing the assembly technique, at the expense of increased assembly time and cost in production.

The latest reports indicate that the design remains around 450kg (1000lb) above intended weight targets. In an interview published last September, Rear Admiral S L Enewold, deputy program director of JSF, indicated that weight reductions would be achieved by reducing the performance envelope, ie “take some corners of the envelope and shave them off”. This is consistent with the Cost As Independent Variable (CAIV) design approach, in which capability is traded down to maintain a target unit cost.

For US users of the JSF, who will task it mostly with battlefield interdiction and close air support, reductions in the aircraft’s performance envelope, especially speed and agility, will be of marginal relevance – a stealthy equivalent to the F100 powered afterburning A-7 Corsair II interdictor prototype will be more than adequate.

If the USAF F-35A CTOL JSF ends up a 7.5G rated, Mach 1.3 dash speed fighter with a sea level wet thrust/weight ratio of 0.9:1, the aircraft will still be a major improvement over the types it replaces in this role. Recent statements by US Air Force Secretary Roche indicate many US Air Force JSFs may be delivered in the least agile STOVL (USMC/RN) configuration. An aircraft in this performance bracket would not be competitive in air combat roles in the Asia-Pacific environment of post 2010.

To date there have been no adverse reports on the P&W F135 and GE F136 engines, both using enhanced derivative cores from the respective F/A-22A engines, the F119-PW-100 and YF120. Both of these ‘supercooled’ engines have the hottest running cores to date, even hotter than the F119-PW-100 which has yet to accrue significant operational hours.

The big issue for the JSF engines will be durability – not designed for dry supercruise, the JSF will need to use afterburner in combat more frequently than the F/A-22A, presenting a more aggressive thermal cycling environment - durability of the F/A-22A’s engine hot end could be a poor indicator of JSF hot end durability. Historically more aggressive operating cycles proved to be a major issue for durability in the hot end of the F-15A and F-16A F100 engine, with a number of hot end fires and written off aircraft.

If durability issues arise, they may not become apparent until low rate initial production aircraft are in early service, and the typical measure to deal with this is derating the engine. This costs top end performance, again a non critical issue for US users, yet a problem for Australia. An issue in its own right will be the durability of any stealth coatings used in the nozzle and tailpipe areas.

External and especially internal munitions clearances could also present risks, and problems may not be solved until late in the program. The drag increasing pylon toe-out in the F/A-18E/F presents a good example. Internal release of smaller weapons like the GBU-39/B or GBU-38 500lb JDAM can be challenging, as ejection velocities in excess of 20ft/sec could be required. While the use of pneumatic ejectors will address this for the basic payload of eight GBU-39/Bs, growth configurations may present genuine problems.

JSF GROWTH POTENTIAL ISSUES

For Australia another key long term issue will be the growth potential of the JSF design. Additional engine thrust for a given core technology is usually achieved by increasing engine massflow – informed sources indicate the current inlet design has only a very modest growth margin in available massflow. Whether a 50,000lb (222kN) class F135/F136 derivative can be used with this inlet has not been disclosed to date.

Another growth issue will be available internal volume for avionics, and especially waste heat management capacity. Any increases in ICP capacity and AESA power rating will be reflected in significantly greater waste heat to be dumped from the systems, already reported to be an issue at this stage. Again, for US users targeting interdiction and support roles avionics growth limits may be largely irrelevant – more radar range and a larger information gathering footprint are not critical factors. For Australia, competing with Sukhois in air combat roles, and using the JSF to provide ISR and long range strike capabilities, growth will be a decisive issue.

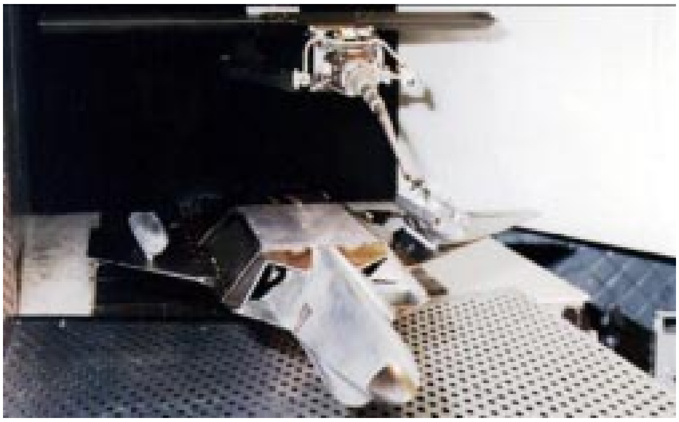

The F/A-22A is not just a dedicated and specialised air superiority fighter – the US Air Force Global Strike Task Force will use its F/A-22A component mostly for trucking smart bombs. Depicted (top) is USAF AEDC wind tunnel testing of a developmental external stores pod for the F/A-22A, intended to reduce the radar signature of additional external bomb payloads. The jack of all trades JSF is part intended to replace the A-10A and AV-8B Harriers in close air support roles (bottom), and is not optimised to fulfil ‘air control’ or ‘air dominance’ roles. (USAF/LM)

The design of the EOTS window fairing and nose radomes will impose hard limits on any aperture size growth in these key sensors, in turn setting bounds on achievable sensitivity growth. This is especially a problem for advanced IRST capabilities, which require also an expensive replacement of the Sapphire windows with a longwave transmissive material.

There are many as yet unresolved technological risks in the JSF, and many of these may not be manifested until later this decade – potentially impairing the performance of the JSF in areas where Australia needs to be highly competitive longer term.

BUILD NUMBERS, TIMELINES AND COSTS

Other major risks will arise in relation to build numbers, delivery timelines and costs. We have already observed a 12 month delay introduced into the program to manage risks, while $US5bn was shifted from the low rate initial production budget into the development budget late last year.

While full scale production is almost a decade away, any schedule slippages will impact on production costs. Flyaway costs of aircraft are highest at the start of full scale production, and progressively reduce as cumulative build numbers accrue, production investment is amortised, and component manufacture matures.

Current Defence planning sees Phase 1/2 JSF deliveries starting around 2012 and ending later that decade. If the JSF production schedule is delayed significantly, Australia buys more expensive JSFs sitting earlier on the production cost curve. In plain dollar terms, buying JSFs in 2020 is cheaper than buying them in 2012.

With similar internal fuel loads in production models (differing from demonstrators), the larger but cleaner F/A-22A (left) provides similar combat radius to the F-35 JSF. Both types will suffer combat radius loss with draggy external payloads, and both types require extensive aerial refuelling support to compete with the existing F-111 in both range/payload and on station persistence. The F/A-22 can however carry more than twice the external fuel payload of the F-35 in drop tanks, giving a total fuel payload 6% greater than the F-111's. (USAF)

Cost related risks fall into three broad categories. The first is that resolution of technological problems drives up the build cost. The second is that schedule delays put any Australian buy into an earlier portion of the cost curve, assuming current schedules for F/A-18A replacement. The third is that US and export clients buy lesser numbers.

The third is potentially the most problematic, as it is driven by overseas budgetary politics and evolving strategic needs. It could manifest itself very late in the program. Since Australia joined SDD we have seen the US Navy and Marines trim back their buys, with the current total sitting around 2500 aircraft. Only the Marines and the UK are technologically locked into the JSF as they use STOVL carriers. The US Navy could bail out and buy more F/A-18E/Fs if the going gets too tough for them at any stage.

The US Air Force is F/A-22A centric in its thinking, for good strategic reasons. The JSF provides a mechanism to drive down the cost of radar, engine and avionic technology used in the F/A-22A, like the high volume F-16A drove down engine costs for the F-15A. No less importantly the JSF presents a big chunk of reserved funding for the ACC fighter fleet, one which might be redirected at a future date into funding more F/A-22As. Given the choice of putting the money into more F/A-22As and FB-22As, or JSFs, there is no contest once the US Air Force has covered its most critical replacement needs in close air support tasked A-10As and older F-16s.

Shifting strategic needs could have the greatest impact on US Air Force numbers, as its targeting model is reoriented from predominantly static to mostly mobile ground targets. Even at the JSF's nominal 600nm (1110km) radius, a lot of tanking is required to achieve significant persistence. An F-111 sized FB-22A works much better as a battlefield interdiction asset than a JSF does, and if the FB-22A does materialise it will subsume over time much of the battlefield interdiction role, driving the JSF into the specialised lower altitude close air support role which it is superbly adapted to.

As yet an unknown is the pricing and numbers impact arising from the likelihood of the US Air Force splitting its JSF buy into CTOL and STOVL variants – a proposal revived by SecAF James Roche at the recent Air Force Association symposium in the US and intended to bolster close air support/battlefield air interdiction strength in expeditionary forces. If this occurs, build numbers of the CTOL F-35A JSF will go down, driving up flyaway costs, and build numbers of STOVL F-35Bs go up, driving down flyaway costs. Out of a finite budget a smaller total number of JSFs is bought for the US Air Force, in turn impacting flyaway costs across all three variants. The US Air Force is already hedgeing its bets on JSF timelines by planning engine and avionic upgrades for many A-10As in its fleet.

Long term export numbers for the JSF remain unclear. Many European F-16 operators will simply opt to swap their existing fleets for JSFs, in a truly benign post Soviet local strategic environment.

WHAT NEXT FOR AUSTRALIA?

Australia's interest in using the JSF for air control/air dominance and long range strike roles does not fit well with the basic design optimisations of the JSF, or the outcome of likely cost driven downstream performance/cost tradeoffs in the JSF program. In distant historical terms it is akin to using a P-40 to do the jobs of a Beaufighter and P-38.

In its core role of 'classical' battlefield interdiction and close air support, the production JSF is apt to be a superb performer, more lethal and survivable than the F-16C, F/A-18A-D, A-10A and AV-8B it replaces. Its effectiveness in the air combat role, against the ever evolving capabilities of the Sukhoi fighters and newer Russian missiles, is very much open to debate and clearly problematic. In the long range strike role, around 60 JSFs with generous tanking could match the aggregate punch of the F-111 fleet, but the 'narrowband' stealth optimisations of the design will not provide the kind of unchallenged survivable deep strike capability Australia gained in 1973 with the F-111, pitted against then regional capabilities.

The big question for Australia is whether the JSF is suitable as a single type replacement for the F/A-18A and F-111. Aside from the fractional battlefield interdiction and close air support roles, the JSF falls well short in the prime air control and deep strike roles, compared to the alternative F/A-22A and likely future FB-22A.

Even at this early stage in the New Air Combat Capability/Air 6000 program an overwhelming case can be made for restructuring the program to focus on the F/A-22A rather than JSF, with a decision deferred to 2008. While the F/A-22A is more expensive, it is also more mature and much more capable permitting smaller numbers to achieve better combat effect.

A package of 36 F/A-22As is more lethal and survivable than 72 JSFs, especially in the critical air control and deep strike roles. An 'F/A-22A centric' NACC solution involves a mature production fighter after 2010 and incurs none of the schedule, technology and cost structure risks, or longer term strategic and technological risks associated with the JSF – an 'F/A-22A-centric' NACC is a very safe solution.

The current plan for early retirement of the F-111 is particularly unhelpful in terms of providing long term options for the NACC program. Retention of the F-111s past 2020 would permit spreading the expense of F/A-22A, JSF or mixed buys over a longer timeline, without any capability gaps arising. The current plan simply forces the replacement buys into an earlier and more expensive time window, while incurring a large capability gap and wastage of prior taxpayer's investment.

The stark reality is that whatever aircraft is chosen, Australia will have to live with it into the 2040 timescale. Choices which might look just good enough against the region today will not be competitive two to three decades hence, as a wealthier Asia invests increasingly in modern airpower.