CUSTOM READY TIMES: A MILESTONE FOR HARPOON

By **Dimitris V. Dranidis**

On March 2003, Jesse Spears publicly released version v3.5.9 of the Harpoon 3 air/naval simulation. Concurrently, the Harpoon Headquarters released the v6.4 version of its famous DB2000 database. Combined, these two releases break one of the longest-standing moulds on all computer versions of Harpoon: the dreaded "30 mins for all" turn-around time. Realistic ready-times for aircraft loadouts are finally here.

Custom Ready What….?

The new Custom Ready Times (CRTs) are the result of the years-long frustration of Harpoon users with the arbitrary setting of 30 minutes of turn-around time for all aircraft. This has been a fundamentally flawed approach for two primary reasons:

1) The arbitrary 30-minute setting is unrealistically small for most tasks where significant pre-flight preparation is required. This is particularly true of strike/bombing missions where particular attention and effort must be dedicated to the coordination of the air platforms with other cooperative elements (escorts etc.)

2) The result of very small turn-around times has been the unrealistic hyperactivity of air assets that has been plaguing computer Harpoon ever since 1989. Aircraft have been too omni-present, too omnipotent for other scenario units (ships, subs and land units) to matter. It should be remembered here that Harpoon is a wargame of modern combined tactical/operational air & naval warfare, not Flight Commander (as an example) with a naval element strapped-in as an afterthought.

It is noteworthy here to elaborate a bit on the negative effect of this air-hyperactivity on all versions of computer Harpoon until now.

In the original Harpoon Classic, as well as its modernized reincarnations (Harpoon 97, Harpoon 2002) and the online version (HOL), the small turn-around times combined with unlimited aircraft weapons to make airpower a near-invincible tool. Almost all scenarios were determined on which side had the most formidable air assets - never mind what other non-air forces were available to either side. A common Blue-side tactic was to simply shuttle your carrier-borne strikers back and forth in a "take off - get to launch distance - launch Harpoons - get back to carrier" cyclic motion until the Red CVBG/SAG ran out of defensive missiles and countermeasures. No real strategy, no real effort, no complex time-on-target planning or attempt for a surprise attack or a multi-axis saturation, just a pre-determined exercise in attrition. One could almost lay back and calculate/predict when his air assets would break through the defences.

The same held true for the Red side, in scenarios where the Soviets had lots of missile-armed bombers and some good escorts - they could simply do the "take off - launch stand-off ASMs - land" dance all day long. This essentially meant that the side with the stronger air assets would win almost every time. (While many airpower enthusiasts may find nothing abnormal to this result, actual combat operations have proven to be a bit more complex than that).

In Harpoon 2 (and by extension, past versions of H3) this situation was partially remedied by enabling aircraft logistics. This optional feature limits the ordnance available to air assets, and thus solves the unfair situation of unlimited airborne weapons versus finite surface-borne weapons. However, while this improvement does (to a large extend) eliminate the mindless attrition mindset, it still does not solve the problem that, even with limited ammo, aircraft still remained hyperactive. This had two implications:

1. The frantic "take off - launch missiles - land" dances of Harpoon Classic continued, the difference now being that at some point they would eventually stop when running out of weapons.

2. Airpower, while great for delivering ordnance, is not limited to this role: equally useful are its surveillance and reconnaissance capabilities, and these are limited only by sortie rates, not by ammo figures. Thus the aircraft logistics option did not prevent air assets in H2 from having an unrealistically high "time in the air" and being far more omnipresent (and thus providing much more complete area coverage) than they would do in real-life operations: For example, with two or three Tu-95RT Bear-D aircraft, and with the standard 30-min turn-around time, the Soviet player in all versions of Harpoon can be guaranteed virtually round-the-clock coverage of a general sea area. In reality, even with 30 such precious aircraft, plus a very elaborate network of shore-based HF/DF stations, sea-going SIGINT spy-trawlers, HUMINT assets, tattle-tale destroyers and even satellites, the Soviet Navy still had considerable coverage gaps even close its operating areas, gaps which could be exploited.

In Harpoon until now, it has been almost impossible for your surface assets (land & sea) to avoid being monitored by aircraft almost continuously – in real life, it is perfectly feasible to avoid both aircraft and satellite surveillance, simply exploiting the gaps in their availability.

So, what was needed was a mechanism to tackle the fundamental problem of unrealistic aircraft availability. A way to show the end-business difference between a hangar-queen and a mud-fighter. A way to show why many aircraft designers go to extraordinary lengths and make important performance sacrifices in their designs in order to make them more serviceable/maintainable. And a way to show how many "invisible" factors like ground-crew proficiency, base ground equipment, airframe/systems complexity, built-in maintainability features etc. etc. can critically affect the end result: how often the aircraft will be "up there", doing their job.

Enter the CRTs.

So, how does it work?

In the discussions between the H3 developer J.Spears and members of the Harpoon community, it became evident that simply assigning a unique ready-time for each aircraft would not suffice. This is because the very same aircraft can have drastically different preparation times depending on its mission. It's not simply a matter of physically preparing the aircraft itself for the mission: depending on the task, the crews must rest, be briefed (often a protracted event!), coordinate their activities with other assets for the same package etc. Therefore, having just one figure for each aircraft would end up being too simple.

A more elegant solution was to base the ready-time on the loadout selected for the aircraft, and thus indirectly the mission profile that is reflected by the loadout stores. In all versions of computer Harpoon, each aircraft has several loadouts available. Each of them is a different predefined package of weapons, sensors, extra fuel etc., and represents a different mission profile: for example, one loadout might be heavy on air-to-air armament and represent a typical CAP load, another might emphasize long range and precision strike capability (PGMs and fuel tanks), yet another may go all-out on short-range CAS (rockets & cluster bombs on all racks) etc. etc. It is therefore fairly straightforward to determine the mission that an aircraft is destined for, by checking on the loadout configuration. It thus makes good sense to base the ready-time on the loadout itself.

Programming-wise, a number of modifications were required in order to implement the new feature:

- A new field ("ReadyTime") was added on the DB2000 database, on the tLoadout table[1].

- Jon Reimer's superb MS-Access-based DATed database editor was upgraded to version 1.2 to support the new DB field. You can get the new version here: http://jlreimer.home.netcom.com/h2dbb/h2dbb.htm

- The Harpoon 3 executable itself was modified accordingly, to take into account the new DB field when assigning the ready-times to aircraft being prepared for a sortie.

What actually happens during the simulation is pretty straightforward. When the player selects a loadout for any aircraft, the simulation engine looks for the ReadyTime field on the DB record for that loadout. If it finds a valid value, it assings that time for the aircraft. If it does not find a valid value or the ReadyTime field is not found (this would indicate an older DB), it defaults to 30 mins. This is done in order to preserve compatibility with older scenarios built around the "30 mins for all" assumption.

Real-life factors

So, what factors affect the preparation time (and thus sortie rate) for aircraft missions?

- Aircraft maintenance: the aircraft must undergo the normal maintenance procedures following the last flight. This also includes the replacement of spare parts who have reached their limit of flight hours, as well as repairs on avionics & powerplant modules that have failed (a situation especially common with older/Soviet components, who exhibit low MTBF rates).

- Re-arming & refuelling: Stores and fuel have to be loaded to the aircraft. This is additionally complex where multiple or special carriage schemes are involved (rotary launchers, multi-level bomb racks etc.)

- Crew rest: The aircraft crew must receive a minimum amount of rest between missions in order to perform effectively. This can be skipped if the aircraft are hot-shared between multiple crew teams; however, the ratio of crews-to-aircraft is rarely large enough to make this a common practice.

- Briefing: This can very often be a long and elaborate process, particularly for offensive strike operations where large packages are assembled.

- Miscellaneous procedures: Getting aircraft in and out of protective shelters or revetments or below-deck hangars (typically under tow), starting-up (either with an internal APU or an external power cable), uploading mission data on the aircraft avionics suite, assembling on the flight line, last-minute checks & inspections etc. Each of these actions normally does not take long on its own, but the adding-up can make a lot of difference.

Most of these procedures take even longer if the number of aircraft in a given airbase increases, as the base personnel must divide its material & manpower resources amongst more airframes (this is the second reason against concentrating a large number of airframes in any single base). They also vary widely between nationalities and even different branches of the same country's armed forces.

[1] The DB structure described is the one present on Jon Reimer's superb MS-Access-based DATed database editor. It is not necessarily the way it is internally interpreted by the H2/3 engine – but since DATed handles all the ugly details of the conversion to the cryptic H2/3-readable format, this is not a concern.

So you simply add the numbers, right?

Well, not quite. Although a simple addition of the times for most of the typical procedures would suggest a quite short amount of downtime for most combat aircraft, recent real-life operations suggest a different picture. For example, even the most heavily overtasked offensive airframes during Desert Storm in 1991, USAF's F-4G "Wild Weasel" defence-suppression aircraft, had regular breaks between sorties of a minimum of 6-8 hrs. Most tactical aircraft during that conflict achieved 2-2.5 sorties per day during their peak availability. Similarly modest sortie rates were exhibited both during the Vietnam conflict and the various operations in the 90s, particularly during Allied Force and Desert Fox. It is therefore misleading to simply trust an add-up figure without reflecting on actual operational uptime data.

For this reason, the HarpoonHQ crew participating in the development of the DB2000 have performed an extensive research on relevant available information. This included everything from official sources to pilot handbooks to chit-chats with pilots to decade-old Usenet posts to multi-inch thick encyclopedias to NATO/WP doctrine & planning documents…….a long list of sources by anyone's standards. Here is a small but representative sample of the findings:

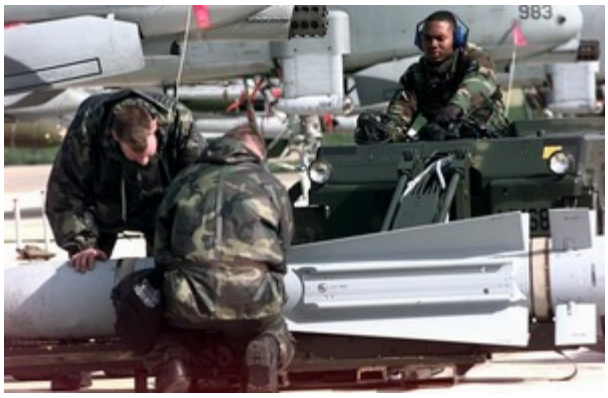

- Modern strike aircraft typically fly two to (at most) three strike missions per day. This holds true for both carrier-based as well as for land-based tactical aircraft, including the teen series (F-14, F-15, F-16, F/A-18), A-6E, Mirage 2000, Rafale, Su-24M and Su-34. Older jets like the F-4E Phantom or (Soviet-employed) MiG-23 were normally able to produce only 2.5 sorties per day, as their avionics and other subsystems were generally less reliable and more maintenance-intensive.

- Several airframes of exceptional complexity and/or unreliability had/have their uptime suffer significantly as a result. For example the F-111D with its super-sensitive avionics was not typically able to generate more than 1.5 sorties/day. Early (Block 10) models of the B-2A could only manage one sortie per 48h due to the need to re-process their delicate radar-absorbent coatings after each flight.

- Many air forces are known to not being capable of maintaining high sortie-rates. This is a result of various factors: poor ground-crew proficiency, lack of miscellaneous base equipment in useful numbers, doctrinal harnesses (particularly in Soviet-trained third-world forces), low pilot-to-airframe ratios, difficulties with complex hardware & sub-systems etc.

- The F-117A only flies at night, which results in one sortie per 24h.

- Strategic bombers like the B-52G/H and B-1B (as well as the later Block 20/30 versions of the B-2A) normally only fly one mission per day, while in-theatre bombers like the Tu-16 or Tu-22M Backfire can do 1 to 1.5.

- CAS aircraft benefit greatly from (typically) simple avionics, exceptional built-in maintainability and ruggedness and short, simple briefings (since most of the tasking is typically performed in the air, in cooperation with forward air controllers). Thus they can often produce 5-6 sorties per day.

These findings, together with a myriad other specific sub-cases, have then to be enforced and applied on the database. Now, considering the loadout dataset currently consists of literally thousands of entries, this is easier said than done. The "right" thing to do would be to examine each and every loadout case separately, correlate it with a real-world quote and assign the relevant value. This would produce a hyper-realistic dataset…..which would be ready for release around 2012. Clearly, a different approach was needed in order to at least have something solid to begin with, in a reasonable timeframe. Thus, a system of phases was introduced:

- In **Phase I**, a general set of rules is enforced, based on the research findings: Typically, modern strike aircraft will be able to fly three sorties per

day (ready time is six hours), older aircraft two sorties (10 hours ready time), long range bombers can fly one or one-and-a-half sorties (12-18h ready time), while CAS aircraft have a 2h-ready time. Air-to-air missile loadouts and most other configurations (incl. guns, torpedo, simpler A/G loadouts) will have 30- to 120-minute ready time. This phase was completed well before the public release of DB2000 v6.4.

- **Phase II** is the second-time scan of the dataset: The HHQ crew has run through the list repeatedly refining the data in each turn, to suit specific aircraft, loadouts, different countries, different timescales, different tactics and doctrines, special cases (F-111D, F-117, B-2 etc.) and so on. And with each subsequnet pass, supported by the hitherto performed research, these figures get closer and closer to the actual real-life figures. This phase began before the public release of v6.4 and is close to being completed.

- **Phase III** commenced as the v6.4. dataset was nearing public release, and covers the longer-term refinements. The HHQ crew realize that many of the numbers assigned during Phase I/II may be well off the mark (although still much closer to reality than the "30 mins" figure), and thus are committed to continuously refining them. The research for more realistic time figures is far from over; in fact, it has only just begun. All the usual means of research are being mobilized to provide data on each distinct loadout configuration as well as hints on variations of it.

One point is worth repeating here: The loadout times are based on actual wartime sortie rates, NOT on the theoretical minimum time it takes to prepare an aircraft and its crew for a given mission. This subtle difference has been the source of considerable misunderstanding in various discussions between members of the Harpoon community. Hopefully this statement clarifies the situation.

Gameplay effects

It can be safely assumed that robbing Harpoon's aircraft of their hitherto nigh-omnipotence is bound to have drastic repercussions on playing the wargame itself. Here is a sample of some of the effects that can confidently be expected on how the simulation interacts with the user's actions & decisions:

- Aircraft whose designers made important sacrifices on outright performance in order to improve their uptime & availability are now finally being vindicated. Watch simple aircraft as the F-16A, A-4, F-5, MiG-21, Su-25 or A-10 roam over the battlefield repeatedly while their more capable, more sophisticated (and thus more complex and less available) siblings make an occasional and long-awaited appearance, displaying their unique abilities before again withdrawing to their maintenance hangars. Heavy & long-range combat aircraft in particular[2], as well as special-mission and electronic-reconnaisance aircraft, now finally get the "silver bullet" treatment they deserve.

- The vastly reduced sortie-rate of most aircraft, combined with the finite scenario duration time, means that players now have much fewer available aircraft sorties in total. Therefore, they have to make much better and responsible use of them in order to succeed: the margins for errors have just shrunk dramatically. Defensive CAPs now have to be planned with greater care, allowing for cross-coverage between successive screens and maybe even hot-pad reinforcements ready to surge. Likewise, offensive missions have to be planned very carefully, with a renewed emphasis on catching the adversary off-guard and overwhelming him with multi-axis attacks rather than with brute-force repeated visits, as has been the case until now. In some air ops-centered scenarios with very tight time limits (such as R. Emsoy's epic "Clash of the Titans" or Steve Le Blanc's "Operation Babylon"), failure in the first attempt may well not be an option anymore.

- Numerous old Harpoon scenarios (particularly 360's stock CD scenarios) are suddenly becoming attractive to play again. In too many cases a scenario was unbalanced because of one or the other side having a decisive airpower edge which, because of the omnipotence of aircraft, translated into an easy victory. Time to dust-off some of the goldie oldies and see why the scen. designers bothered to put those subs and surface ships in there as well.

- With the gaps between aircraft availability now getting close to their real-life values, one has to be much more careful about covering his air assets while they're on the ground (or on the deck). Pearl Harbor, Midway and

[2] *Until now, placing medium/heavy bombers in scenarios was done very carefully, as their combination of heavy firepower and unrealistic availability made it very possible to win the scenario all by themselves.*

the start of the 6-day war are but three of the numerous examples of an air force being caught quite literally with its pants down. Now it can happen to you too, so keep your guard up. This also has the implication that a decisive "first strike" is going to regain its importance at the expense of mindless attrition tactics.

- Multiplayer match-ups are going to be very interesting to watch. The need for creative air-asset management, combined with the natural innovation of a human brain (as opposed to a scripted AI) promises a wide assortment of tricks and surprises. It is quite fortunate that the auto-screendump (a.k.a movie-maker) feature is now available and mature; Some of the human MP matches are almost definitely going to be begging to be immortalized through a good video-replay ☺

- Dispensing with the air-hyperactivity means getting back to the roots of Harpoon: combined naval and air operations. Ask yourselves this: in a pre-v.3.5.9 scenario, how much of your time & actions do you devote to aircraft ops and how much to your other assets? Exactly.

Not doing the "land-arm-launch" dance every half-hour means having time to plan the overall strategy better, maneuvering the non-air assets, checking one's forces for damages/casualties from the last clash (a tricky thing in large scenarios), preparing the surviving forces better (reloading specific weapons in mounts, forming custom groups etc.), replenishing and re-supplying the surface/sub forces etc. In short, it means having an overall better and more realistic wargaming experience.

An ongoing endeavor

As previously stated, the refinement of the CRTs is a continuous effort. Some of the current figures may well turn out to be wrong, or too generic (or from another POV, not specific enough for aircraft-X with loadout-Y for mission-Z). Here's where the Harpoon community's strong tradition of cooperation can once more be of benefit: If you think you have some piece of information that is not featured in the DB, or contradicts what is already in the dataset, your best bet is to contact the HHQ crew either in the forums (http://www.harpoonhq.com/phpBB2/) or alternatively, directly at the HHQ site. Such communication is, in fact, already happening: CRTs are, for some time now, one of the hottest subjects of discussion in the various computer Harpoon discussion forums. The HHQ crewmembers working on them are permanent dwellers on such hangouts and are in constant communication and exchange of information with fellow pooners, in the quest for perpetually refining the ready-time figures (in addition to all the other DB-related projects).

The first and most important step has been made, with the code-enable of CRTs and the DB modification: the cat is now out of the bag for good, and there's simply no going back.