Eclipse Phase Space Warfare

Anders Sandberg

April 26, 2011

Abstract

noindent This essay analyses the physics of spacecraft and space combat in Eclipse Phase. Based on the technological assumptions explicitly and implicitly made in the game together with known physics, various constraints on space warfare can be concluded. In general, the space battlefield is extremely high-energy, high-loss and dominated by the forces that can estimate the locations of enemy assets accurately despite massive interference. FTL quantum entanglement communication provides a big but not decisive advantage to forces able to afford it.

Contents

List of Tables

1 Introduction

Eclipse Phase is a mostly hard SF role-playing game with a setting stretching across the solar system (and some exoplanets). Space travel is not uncommon and in the past there have been military conflicts in space. While personal combat is described space combat is not; the game states that:

> "Spacecraft have few stats in Eclipse Phase, as they are primarily handled as setting rather than vehicles. Note also that no stats are given for spacecraft weaponry. It is highly recommended that space combat be handled as a plot device rather than a combat scene, given the extreme lethality and danger involved."[1]

While this is sensible advice for a roleplaying game, it is a direct challenge for players who like to consider how space combat might work, the strategies involved and how these might affect the characters inside the game. For example, is space combat heavily stealth oriented or requires ships with massive armor barging through enemy point defenses? What are fighter craft (and fighter pilots) useful for? How much of an advantage does a space habitat have against an attacker?

In the following I will analyse what follows from the assumptions made in the game, as well as extrapolations from known physics and technology.

Technology in Eclipse Phase has achieved dense energy sources allowing fast spacecraft, fast optical processing running human-level cognition, and antimatter weaponry. But space strategy is limited by lightspeed delays (somewhat modified thanks to quantum entanglement FTL communications), the limitations of materials based on molecular bonds, high visibility of accelerating major crafts and finite reaction mass resources. Many technologies in the game are close to the limits set by physics, which simplifies analysis somewhat.

1.1 Acknowledgements

This essay was inspired by past discussions at the Eclipse Phase Forum[2], where many bright ideas were suggested.

The Atomic Rocket page[3] and Rocketpunk Manifesto[4] have been an important source of inspiration, references and opinions influencing this essay.

The basic ship performance numbers were kindly supplied by JSnead. I would also like to thank him for having done proper design calculations behind the scenes, simplifying this work immensely.

[1] Eclipse Phase core book, p. 346.

[2] `http://www.eclipsephase.com/forum`

[3] `http://www.projectrho.com/rocket/index.html`

[4] `http://www.rocketpunk-manifesto.com/`

Table 1: Ship engine types

Engine	Acceleration G	Acceleration m/s^2	Isp s
Hydrogen-Oxygen Rocket	4+	39.3	450
Metallic Hydrogen	3	29.5	1,600
Plasma Rocket	0.01	0.1	20,000
Fusion Rocket	0.05	0.5	100,000
Anti-Matter	0.2	1.96	200,000
Rocket Buggy	0.5	4.9	

2 Ship performance

The key ability of spaceships is that they can accelerate long and strongly enough to reach high velocities or change their velocity (Δv) radically. This requires expelling reaction mass at a high velocity. A key value is the specific impulse (I_{sp}), how much momentum each kilogram of reaction mass can impart on the ship. The faster the reaction mass is emitted, the higher the I_{sp}. However, this does not necessarily mean a higher acceleration. Available engines typically have a high acceleration for *low* I_{sp} and vice versa: the high acceleration engines are less able to achieve high Δv since they waste much fuel, while high I_{sp} engines cannot produce high accelerations due to energy limitations.

The achievable velocity change is

$$\Delta v = v_e \log\left(\frac{m_0}{m_1}\right)$$

where v_e is the exhaust velocity ($v_e = gI_{sp}$, g is the Earth surface gravity), m_0 the initial total mass of the spacecraft and m_1 the remaining payload mass after all reaction mass has been used. Higher velocities require either higher exhaust velocities or exponentially more fuel.

Ships with $\Delta v > 80$ km/s typically do not have to worry about launch windows, while slower ships need to plan their trajectories so that the origin and destination are in the right alignment.

Ships are limited by how much remaining reaction mass they retain when making course corrections (especially defensive ones) en route. I will assume ships (especially warships) keep a fraction of their reaction mass budget in reserve, giving them a fraction of the total Δv for defensive or offensive course changes[5].

Table 1 describes the performance of the basic ship engine types. Table 2 lists the basic Eclipse Phase spaceship properties.

3 Energy requirements

Some estimates of the powers available can be gained from considering spaceship performance. A ship thruster requires $P = (1/2)m'v_e^2$ W of power, where m' is the mass flow in kilograms

[5] The Eclipse Phase core book (p. 233) suggests many ships burn a quarter to a third of the reaction mass during the initial burn. In practice this is rather costly, as reaction mass not used will make the trip longer and require extra reaction mass for the develeration burn. Saving this much fuel is rational only if drastic course corrections may be needed, or the fuel very cheap compared to the cost of arriving later. Many commercial ships likely retain very low fuel margins, and may rely on 'tugships' that help them slow down at the destination.

Table 2: Ship properties

Ship[1]	Size (m)	Cross-section (m^2)	Fully Loaded Mass (tons)	Empty Mass (tons)	Engine type	Total Δv km/s	Max acceleration (Gs)	Acceler. (m/s^2)	Max power
Destroyer	150x50x50	7,500	40,000	25,000	AM	800	0.2	1.96	800 GW
Fast Courier	75x13x13	975	1,000	300	AM	1,600	0.2	1.96	20 GW
Bulk Carrier	150x25x25[2]	3,750	90,000	4,500	Fusion	40	0.002	0.001964	9 GW
Standard Transport	150x25x25[3]	3,750 / 12,000	10,000	4,500	Fusion	400	0.02	0.01964	10 GW
Fighter	4.5x3x3	14	7	3	MH	11	3	29.46	16.8 MW
SCUM Barge	300x70x70	24,000	180,000	80,000	Plasma / fusion	80 / 400	0.003 / 0.015	0.02946 / 0.1473	5.4 GW / 135 GW
LLOTV HO HI	25x16x16	400	450	26	HO	11	2	19.64	203 MW
LLOTVO HO LO	25x16x16	400	450	26	HO	7	2	19.64	203 MW
LLOTV MH HI	19x12.5x12.5	237.5	450	26	MH	17	2	19.64	720 MW
LLOTV MH LO	19x12.5x12.5	237.5	450	26	MH	8	2	19.64	720 MW
SLOTV HO HI	17x11x11	187	150	11	HO	11	2	19.64	67.5 MW
SLOTV HO LO	17x11x11	187	150	11	HO	7	2	19.64	67.5 MW
SLOTV MH HI	13x8.5x8.5	110.5	150	11	MH	17	2	19.64	240 MW
SLOTV MH LO	13x8.5x8.5	110.5	150	11	MH	8	2	19.64	240 MW
General Exploration Vehicle (GEV)	6x2.2x2	13.2	5.5	3	MH	3.6	0.1	0.982	0.44 MW
Missile[4]	1x0.1x0.1	0.1	0.02	0.002	MH	31	200[5]	1,964	3 GW

[1] Abbreviations: AM = Antimatter, HI = HIgh velocity configuration, HO = Hydrogen-Oxygen chemical rocket, LLOTV = Large Lander and Orbit Transfer Vehicle, LO = LOw velocity configuration, MH = Metallic Hydrogen rocket, SLOTV = Small Lander and Orbit Transfer Vehicle.
[2] Plus externally mounted cargo pods.
[3] 80 m wide and high with rotating booms fully extended.
[4] Own design. Payload can be a small (≈ 1 kt) nuclear warhead, a 3 megaton antimatter warhead, kinetic impactor projectiles or attack nanotechnology.
[5] Short burst launch or evasion acceleration.

Table 3: Energy storage

Energy source	Specific power	Power density	Specific energy	Energy density
Fission	2.5 kW/kg	12.5 MW/m^3		
Fusion	200 kW/kg	1 GW/m^3		
Antimatter	372 kW/kg	1.86 GW/m^3	4.5 PJ/kg	22.5 EJ/m^3
Chemical fuels			10 MJ/kg	20 GJ/m^3
Nuclear isomers			10 GJ/kg	100 TJ/m^3

per second and v_e is the reaction mass speed. The thrust force $F = m'v_e$ produces an acceleration $a = m'v_e/M$ on the spacecraft (of mass M). For a given acceleration and reaction mass speed the power is $P = (1/2)Mav_e$. Expressed in terms of I_{sp}, $F = I_{sp}m'g$ (where $g \approx 9.82$ m/s^2 is the gravitational acceleration at Earth's suface) $a = gI_{sp}m'/M$, $m' = aM/gI_{sp}$, and $P = (1/2)gaMI_{sp}$.

This allows us to estimate the power requirements of spacecraft if we know their mass. Consider a fusion-powered spacecraft. It has $a = 0.05G \approx 0.5$ m/s^2, $I_{sp} = 100,000$ s. A fully loaded bulk carrier with total mass $90 \cdot 10^6$ kg will hence require an output of 9 GW during full thrust, an energy output of 100 W per kg of ship mass. In practice bulk carriers likely sacrifice exhaust velocity and power for economy, so the output will be far below this level. Antimatter-powered ships produce comparable energy outputs (at least in terms of propulsion, since the mechanism is relatively similar). The Destroyer, weighing 40,000 tons, implies a reactor power of 800 GW.

As a comparision with existing technology, a Nimitz-class aircraft carrier produces 190 MW and a major nuclear power plant can reach 8 GW. It is probably safe to assume that large ships have reactors that can produce power up to the terawatt range Most of this energy is likely only available for propulsion rather than powering weapons due to the problems of converting it to electricity. There are also going to be efficiency losses leading to large amounts of waste heat. Assuming 90% efficiency still requires hundreds of gigawatts of cooling. Hence warships are unlikely to use their reactors at full power during battle, in order to avoid having to unfold very noticeable and vulnerable large radiator surfaces. The amount of available weapon power is still going to be very large.

As shown below, there are good tactical reasons for wanting to accelerate quicker than allowed by fusion or antimatter drives. This can be achieved by using high g-thrusters such as metallic hydrogen: while the main engines aim at a very high exhaust velocity to keep reaction mass requirements down while achieving a high Δv, these engines are intended to use low velocity reaction mass to make a few brief but strong changes in ship velocity. The total Δv is negligble since they cannot be sustained for long, but they allow rapid evasive maneouvers. Also, main engines can in some cases be supplied with more reaction mass than normal to produce short bursts of acceleration.

3.1 Reactor sizes

What is the size of capital ship reactors? For fission reactors the specific mass is around 40 kg/kW, although advanced vapor core reactors might go down to 0.4 kg/kW. According to R.W. Bussard[6] the fusion reactor specific mass could be 0.05 kg/kW. For the transport's 10 GW reactor the weight would be 500 tons.

The destroyer is antimatter powered, and assuming the whole reactor is about the size of the containment system we get a specific mass of 0.0025 kg/kW for antimatter power.

Assuming a density of the reactor to be about 5000 kg/m^3 a 10 GW fusion reactor would be 100 cubic meters, or a $4.6 \times 4.6 \times 4.6$ cube. In practice the reactor will be far more extended, since these estimates mainly deal with the core. Containment, control, cooling systems etc. will probably be at

[6] http://www.askmar.com/Fusion_files/FusionElectricPropulsion.pdf

least ten times as large (but also of lower density). Similarly, older reactors will also be much larger.

It should be noted that small vehicles like fighters and missiles likely use other energy sources. An antimatter-powered 1 ton reactor could provide up to 370 MW of power, but for maximal acceleration various chemical fuels (such as burning metallic hydrogen) would be more effective. Given the available technology high energy-storage densities are possible, probably on the order of ≈ 10 MJ/kg (corresponding to computed explosives such as dinitroacetylene, octanitrocubane and octaazacubane and still smaller than lithium-fluorine combustion). For even higher densities nuclear isomer storage might be possible.

This table sums up the above considerations. The specific power of chemical fuels and nuclear isomers depends on the rate of reaction and can in principle become very high when released explosively.

3.2 Exhaust temperature

A sizeable fraction of the energy output is going to be present as heat in the expelled reaction mass. Rocket nozzles of chemical rockets can achieve 60-70% efficiency as heat engines converting heat into velocity. The remaining energy will largely be carried away by reaction mass. If the efficiency is η and the power is P, then the exhaust will have temperature

$$T_{exhaust} \approx (1-\eta)P/Cm'$$

where C denotes its specific heat capacity in J/kg K.

For the Destroyer, expelling 40 kg/s hydrogen at 60% efficiency $T_{exhaust} = 560,000$ K. This is a hard UV source and not far from a particle beam weapon. At distance d, assuming the exhaust radiates spherically, the incident energy is $(1-\eta)P/4\pi d^2$ W/m^2. In this case it is 25 MW/m^2 at one kilometer distance, enough to vaporise the surface of steel[7]. This is why cooling is so essential for accelerating spaceships.

Jon's law: "Any propulsion system powerful enough to be interesting, is powerful enough to be a weapon."

3.3 Cooling

A key problem for all spacecraft with high power is cooling since space is a perfect thermal insulator. While some engines (metallic hydrogen, fusion, antimatter) carry away a sizeable fraction of the power as heat in the exhaust, most ship power plants will produce vast amounts of waste heat that must be removed[8]. As a rough approximation, assuming 50% efficiency, the same amount of power the reactor produces for the engines and other forms of usable work is also produced as waste heat. Radiators radiate waste heat into space but have an upper acceptable temperature T_{max}. This requires a total radiator area larger than $A_{rad} = P/\epsilon\sigma T_{max}^4$ where ϵ is the emissivity (likely chosen close to 1) and $\sigma = 5.67 \cdot 10^{-8}$ W/m^{2}K^4 is Stefan-Boltzmann's constant.

Using liquid lithium as a coolant gives $T_{max} = 1600$ K. For $P = 800GW$ A_{rad} =2.15 million m^2, requiring multi-kilometer fins (or droplet radiators, where sheets of droplets of molten metal are allowed to drift from emitters to collectors). The fighter requires 45 m^2 of radiators when using full power, not too different from fighter plane wings. Using thermal conduction in 3000 K tungsten the fighter can reduce the radiator area to 3.6 m^2.

[7] `http://panoptesv.com/SciFi/DamageAverage.html`

[8] There will also be separate radiators for cooling low-temperature sections of the ship such as the life support system, but they are negligible compared to the main radiators.

In practice radiators of military ships are foldable, fully used only when doing main accelerations and completely folded back when in battle mode for maximum protection and minimum emissions. Since most battles are very short this is usually just a minor problem, but if the situation persists overheating begins to become an issue (see section on cooling and internal storage of heat). Damage that prevents full unfolding after a battle severely limits course changes.

Calculations by Nemtos[9] for more realistic tapering cooling fins find that the total flux that can be radiated is $(4/5)L\epsilon\sigma T^4$ where L is the width of the fin. This gives a total mass of $\rho L\delta_0/3$, where $\delta_0 = (8/5)L^2\epsilon\sigma T^3/\lambda$ is the thickness of the central condensing channel, ρ is the density and λ is the thermal conductivity of the fin material. Using liquid potassium as coolant ($T = 1000$ K), pyrolytic graphite ($\lambda = 400$ W/m, $\rho = 2200$ kg/m^3) as a heat conductor and white ceramic ($\epsilon = 0.95$) as a surface material allows radiating away 43 kW/m^2. Smaller fins were much more efficient in terms of energy release per weight but require more extended pipe systems since they need to be extended further out.

Liquid droplet radiators were estimated to require about 50% of the cooling fin mass. Since the droplets loose energy faster when they are hot it is more effective to build compact radiators with a flight time around a second. Given some of the limitations of droplets screening each other it was concluded that it could radiate away 20 kW/m^2 when using liquid tin at $T = 1000$ K. Droplet radiators are hence preferable over cooling fins when mass budgets are an issue, such as in small ships.

Another cooling method is to heat coolant and dump it into space. Using hydrogen $14.30 \cdot 10^3$ J per kilogram and Kelvin can be removed[10]. Heating metallic hydrogen to 10,000 K would remove 143 MJ per kilogram. As an example, the 800 GW Destroyer would need to vaporize 2,797 kg/s for cooling. This is impractical for normal cooling, but acceptable during combat where radiators are not available (and the detectable coolant emissions are overshadowed by the main engines). The fighter requires 0.06 kg/s (assuming 50% efficiency). It only got about 4 tons of MH fuel, so it gets just 18 hours of cooling even if it doesn't use any hydrogen for propulsion.

Cooling through expelling coolants is particularly useful for cooling lasers and railgun weapons during battle, especially if they are disposable. It is worth remembering that if you need X Joules to harm your enemy's ship you will have to dissipate $X(1-\eta)$ Joule of heat at home, where η is the conversion efficiency of the weapon.

[9] `http://nemtos.ouvaton.org/techfiles/Cooling_Systems.pdf`

[10] I am ignoring the complications of different thermal capacity at different temperature here.

4 Detection

Detection of enemy spacecraft and assets is of central importance. Unlike on a planet there is no intervening material that blocks radiation emissions, but the sheer volume of space complicates things. There are often tradeoffs between how wide part of the sky can be scanned, how exact positions can be determined and the time it takes.

4.1 Sensors

Passive sensor systems collect energy from particular directions, sum the total energy and try to tell whether there is a statistically significant difference from the background.

The signal to noise ratio (S/N) for a sensor collecting photons such as a CCD sensor is

$$S/N = \frac{Ft}{\sqrt{(Ft + Bn_pt + B_tn_p + Dn_pt + R^2n_p}}$$

where F is the average photon flux from the source (photons per second per square meter), t is the time interval of the measurement, B is the flux per pixel per second from the sky background, B_t is the flux per pixel from the telescope itself, D is the dark current flux (due to the CCD array itself), R is the readout noise per pixel and n_p is the number of pixels[11]. Typically useful observations begin to be possible at $S/N > 5$, although guessing that something might be there is possible at $S/N = 2$ or 3 (but estimates of energy and other properties will have 50% errors). In Eclipse Phase sensors can be assumed to be nearly perfect - B,B_t,D, and R are small, and every photon is caught and turned into a measurable electron.

For sensors with very low noise, a dark sky background and a relatively bright source the S/N ratio is $\sqrt{Ft}$. The time needed to reach a useful S/N is on the order of $1/F$ seconds. If the sensor looks for photons of wavelengths in an interval $\Delta\lambda$ (its bandwidth), has a collecting radius r and the source has a flux of F_s photons/m$^2/\mu$ m/s then the time is on the order of $1/\pi F_s r^2 \Delta\lambda$. As r increases the sensor can tell whether there is something there faster.

As an example, for infrared light $\lambda = 1$ μm (most sensitive to a 3000 K body), a bandwidth of $\Delta\lambda = 1$ μm (a broadband sensor) and radius $r = 1$ m the sensor needs about 300,000 photons in order to detect the target in one second.

If the sky background is significant compared to detector noise, then the time needed scales as Bn_p/F^2. This tends to scale as $1/r^2$ for large targets and $1/r^4$ for targets so small they are limited by diffraction: larger photon collectors are significantly faster. Typical backround sky fluxes in the solar system are between $10^{-9} - 10^{-6}$ Watt/m^2 per steradian in the microwave to UV range of interest to spacecraft detection[12]. For lukewarm (300 K) objects the heat radiation in the zodiacal light is the main confusing factor, while hotter (3,000 K) objects are confused by the background of reflected sublight, faint stars and galactic cirrus. In order to be visible against the backround the flux density from the target needs to be above 12,000 - $3 \cdot 10^8$ photons per square meter at the detector. [develop!]

The time it takes to achieve a given

[11] `http://www.physics.mq.edu.au/current/undergraduate/units/ASTR278/10_ASTR278_JL_5_Sensitivity.pdf` See also *The Design and Construction of Large Optical Telescopes*, ed. Pierre Y. Bely, Springer 2003

[12] Ch. Leinert, S. Bowyer et. al. The 1997 reference of diffuse night sky brightness, Astron. Astrophys. Suppl. Ser. 127, 1-99 (1998)

S/N is

$$t_{detect} = (S/N)^2 \frac{1}{F} \left(1 + \frac{B}{F}\right)$$

The Planck radiation law gives that a spherical object of radius r and temperature T at distance R will produce a spectral flux of

$$F_\lambda = \frac{2\pi c}{\lambda^4 (e^{hc/k\lambda T} - 1)} \left(\frac{r}{R}\right)^2 \Delta\lambda$$

wavelength λ photons per square meter (here I assume $\Delta\lambda$ is small enough; for broadband detectors F needs to be integrated over all sampled wavelengths). If there are several parts of the object of different temperature their spectral fluxes can be added together. Note that for best performance the detector needs to look at a small angle of the sky, since the background flux will grow with the angle.

In order to calculate whether detection is possible we need some estimates of the thermal emissions of spaceships.

4.2 Thermal emissions

A spacecraft or other object radiating at power P uniformly in all directions will produce a total flux of $F_{tot} = P/4\pi d^2$ W/m^2 at distance d. A ship of temperature T and surface area A will radiate $\epsilon\sigma A T^4$ W of thermal radiation, or $F_{tot} = \epsilon\sigma A T^4/4\pi d^2$.

Typically a ship on full power will have extended radiators at temperature T_{max} (and enough area to handle the power). For reasonable T_{max} between 1,000 and 3,000 K the peak flux is between 1-3 μm IR radiation.

A ship that is merely coasting will have an energy output much below these levels, but still significant. Each biomorph onboard produces around 100 W, not to mention life support. A rough guess at the energy dissipation is about 1 kW per crew member. This would put the Destroyer minimal energy dissipation at $9 \cdot 10^4$ W and the fighter at 1 kW. These limits can probably be pushed for short spans, especially by using heating internal cooling reserves (see below). The internal environment will also be maintained at a temperature around 300 K through heating or cooling, and this will likely contribute to a harder to shield surface temperature.

4.3 Exhaust emissions

An accelerating ship will be leaving a long trail of energetic hydrogen, chemical exhaust or plasma behind it, and this will have detectable black-body radiation. Even if the ship itself is perfectly caumoflaged the thermal emission (and its doppler shift, allowing a calculation of relative velocity to the observer) will be detectable.

Exhaust temperatures go down as

$$T_e(t) = \frac{1}{\sqrt[3]{3A\epsilon\sigma(t - C)/K}}$$

where t is the time, A is the area of a one second parcel of exhaust, K is the thermal capacity of it (J/K kg) and $C = K/A\sigma T_0^3$ is a constant set so that at time 0 the temperature is the initial exhaust temperature T_0. For high temperature exhaust it is a good approximation to treat it as releasing nearly all its energy instantly at temperature T_0. This tends to dominate other radiation sources, especially for short-wavelength emissions.

4.4 Sunlight

Internal heat becomes dominant roughly around the orbit of jupiter. Inside that orbit reflected sunlight is a significant source of radiation of total power $P = SA/R^2$ W at solar distance of R AU, sunward ship area A and solar constant $S = 1.366 \cdot 10^3$ W. If the ship or object has albedo α it will reflect $P_{reflect} = \alpha SA/R^2$ W, which is can be very visible. The reflected sunlight might however be made highly directional, for example by placing a plane mirror in front of the object.

The absorbed sunlight, $P_{absorb} = (1-\alpha)SA/R^2$ kW, will turn into heat that is emitted largely homogeneously. The resulting temperature due to the absorbtion (assuming equilibrium and a spherical object of radius r) will be

$$T = \left[\frac{S(1-\alpha)}{4\sigma\epsilon R^2}\right]^{1/4}$$

where ϵ is the emissivity (typically close to 1 for dark bodies and down to 0.02 for polished silver). In Earth orbit for a high reflectivity object with $\alpha = 0.9999$ with low emissivity $\epsilon = 0.02$ the equilibrium temperature is 74 K. For a shiny metal object with $\epsilon = 0.75$ and $\alpha = 0.9$ $T = 168$ K.

Inner system civilian spaceships and equipment are often made white or reflective to keep solar heating down, while outer system ships can be any color they like. Military ships will have reflective mirrors and dark, absorbent coloring in order to be stealthy during certain phases of battle.

4.5 Sensor sizes and sky scanning

Ideally sensors should cover the entire sky and watch continously, but they are limited by the conflicting demands of having large apertures that can collect many photons (a large "light bucket"), a narrow angular field of view to avoid too much background noise and the physical practicalities of where to attach sensors to spacecraft. For habitats and defense systems it is possible to put large numbers of big sensors in place covering most of the sky, but a spacecraft and in particular a mobile asset such as a missile will not have much space.

As an example, the WMAP satellite has a 52.8 arcminute beam size from its 2.24 m^2 sensors: this corresponds to 1 part in 67,827 of the entire sky. If it were to make a quick one second scan of each part it would need 18 hours to do a full sky scan.

For a large ship in Eclipse Phase affixing a few square meter size sensors does not appear to be a major problem; assuming 1% of the ship surface is used for sensors would allow the Destroyer to have 75 m^2 sensors and the fighter 0.14 m^2. If we assume 8 sensors on the Destroyer (one scanning each octant of sky) they would have a 9.4 m^2 area each. These would be the high sensitivity deep scan sensors: in direct battle more numerous, disposable sensors would be deployed for point defense control against incoming missiles.

Full sky scans at long distance do not have to be instant. Assuming a full sky scan takes an hour, standard transports can move 1,400,000 km, destroyers 2,900,000 km, fighters 40,000 km, fast couriers 5,800,000 km and missiles 110,000 km. This is more than enough to detect them before they can get close. Once detected sensors can track them more intently, estimating their true speed, course and other properties.

A sensor covering one octant of the sky in one hour will watch 0.000436 steradians of sky per second, a field of view of about 1.2 degrees side.

One issue is detecting that something is an interesting target and not just random debris, a space habitat or a remote star. A first step is to compare the position with a detailed object catalogue, which allows the sensor to ignore all known objects. The next step is to compare the spectrum to possible ship profiles. Active spaceships have different emissions from debris - infrared from internal energy production, short wavelength emissions from the drive, hot radiators, doppler effect from high velocity etc. The scanner will distribute its scanning time between investigating objects of interest at length and jumping past empty or blocked directions.

4.6 Ship spectra and detectability distance

Putting the above considerations together, we get the following approximate fluxes from the ships of Eclipse Phase.

I here assume the target ship is 1 AU from the sun, 1 AU from the detector, has albedo 0.5, that the radiators are fully extended and the ship is accelerating maximally. The emissions consist of thermal emissions from the radiators, emissions from the ship (largely due to absorbed sunlight), emissions from the exhaust and reflected sunlight. The sensor is assumed to be a 0.000436 steradian sensor with 1 m^2 area, bandwidth 1μ and desired $S/N = 5$. Background noise is assumed to be 10^-6 W/m^2 sr.

Manned ships can be detected over interplanetary distances using an IR sky scan that takes one hour. Typically ships have a characteristic multi-temperature spectrum: one peak for the hot engine exhaust, one for the radiators, one for the reflected sunlight.

Reflected sunlight and the radiators are the biggest contributions: a ship that has folded most of the radiators, reflects away sunlight with a mirror, powered down to a minimal level and cooled the surface to a few Kelvin is significantly harder to detect. The Destroyer is just barely detectable at 0.09 AU distance (13,000,000 km) in this case, and the fighter at 60,000 km. This means that stealth can be profitable, at least in the light of cloud combat: a silent approach allows a surprise launch of attack assets in a cloud large enough to be hard to evade for the enemy. The fighter is not quite stealthy enough to reach a target without being discovered, but it has a chance to close to a very short range and launch a barrage of missiles.

4.7 Scanning limitations

These estimates assume overload-free surroundings. As soon as the bullets start to fly any sensors this sensitive will be blinded (quite possibly destroyed) if they look anywhere close to detonations.

Note that not all directions are available for scanning: thermal sensors pointing at the sun or nearby planets will be blinded. In deep space this is a minor problem, but in the vicinity of planets distributed sensors are necessary to keep watch over local space. The number of objects is also far larger, turning scanning into more of a pattern recognition problem than a detection problem.

A ship that is accelerating is mostly blind backwards due to the exhaust cloud. The blind angle is determined by how fast hot particles from the exhaust spread laterally relative to how fast they move backwards. It is a few times[13] $\alpha = 2\arctan(\sqrt{kT/mv_e^2})$ where T is the exhaust temperature, m the mass of the exhaust particles and v_e the average exhaust speed. For for a hydrogen-oxygen engine at 3000 K with $v_e = 450,000$ m/s $\alpha \approx 0.3°$ while for an antimatter rocket with a million K hydrogen plasma moving back at 10^8 m/s $\alpha \approx 0.1°$. Hence the blind angle is a few degrees across.

Sensors must also be placed so they are not blinded by unfolded radiator fins. These cover a far larger part of the sky but can be avoided by for

[13] Since the Maxwell-Boltzmann distribution extends beyond its average value.

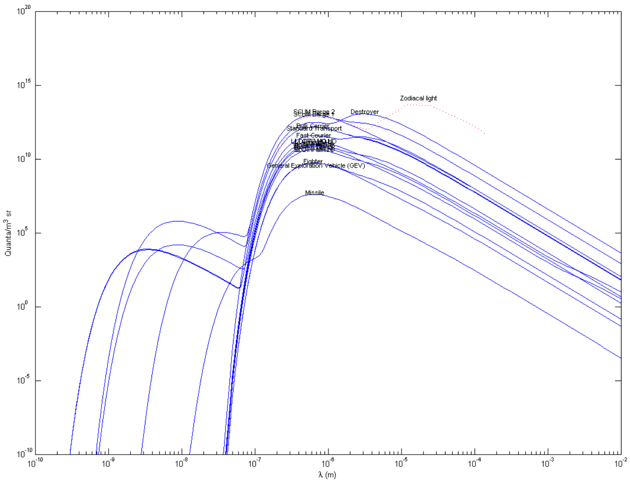

Figure 1: Quanta received by a 1 m^2 sensor at 1 AU from different ship types at full acceleration. The red line is the zodiacal light background, giving a rough estimate of the noise.

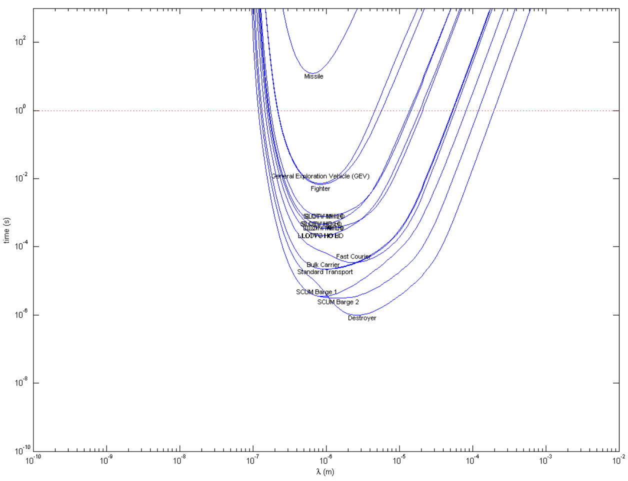

Figure 2: Time until a 1 m^2 sensor at 1 AU can confidently detect different ship types at full acceleration. The red line is 1 second scanning time.

example placing sensors at their ends (gaining parallax information in addition).

4.8 Particle emissions

Antimatter annihilation does not just produce the desired gamma photons, they also produce pions and muons that decay while radiating neutrinos. Fusion reactors also produce neutrinos for some fusion reactions (pure helium 3 reactions avoid it, but reactions with hydrogen may release neutrinos). This means that even if a ship hides its plasma tail it will radiate a neutrino signature. Given Eclipse Phase technology such as emergency farcasters we know that neutrinos can be detected over interplanetary distances. However, it may be hard to get a good position from neutrinos, so the enemy will just know there is an active antimatter reactor *somewhere*. Note that muon detectors would be effective at detecting active antimatter annihilation over distances of a few kilometers, helping missiles to zoom in on active antimatter reactors.

4.9 Railgun projectiles

Railgun projectiles can in principle be detected by their heat emissions. When accelerating a projectile to velocity v, $(1/2)mv^2$ J of work is done. A fraction f ($\ll 1$ but > 0) of this will turn into heat. The temperature becomes $T = (1/2)fv^2/C$, where C is specific heat capacity ($\approx$ 500 J/kg K for metal). If $f = 1\%$ and $v = 10$ km/s, the result is bright 1,000 K projectiles. For $v = 100$ km/s f must be much less, since otherwise the projectile would be a vaporized mess. If $f = 10^{-4}$ the faster projectile will also be 1,000 K.

The only thing making them hard to see is their small area. Assuming the visible area is $\approx$ 10x10 cm, then they can be detected 4,487,940 km [check, update - old equation instead gives "'13.4*0.1*500 = 670 km away. That gives you 67, 6.7 or 0.67 seconds to point defence them. At least for projectiles slower than 100 km/s this is pretty OK for the defender."']

current railguns have plasma flashes 23,000-35,000 K, blackbody radiation 1.6-8.1 MW/cm2

4.10 Active sensors

Active sensors are a dead give-away of where you are (unless they manage to mimic "'natural"' EM activity), but the sensors can be put on an expendable buoy (and triangulation of targets from dispersed sensors is significantly more accurate). Stealthing against radar/twave/lidar/Xdar on all wavelengths is not going to be practical.

Unfortunately active sensors have a shorter range than passive sensors since the radiation emitted decreases as the square of the distance and then the reflected radiation also decreases with the square of the distance, giving a return signal that scales like $1/d^4$. In order to double the range the power has to be increased 16-fold. The radar equation describes the limiting distance where an active sensor can detect a target:

$$d_{radar} = \left[\frac{P_S G^2 \lambda^2 \sigma}{64\pi^3 P_m}\right]^{1/4}$$

where P_S is the power emitted, G is the antenna gain, λ is the wavelength used, σ is the radar cross section of the

[14] A way around this is to use bistatic radar, where the signal emitter and receivers are in different locations: sensors close to the target will receive a stronger signal. This requires that the sensor cloud is at least as large as the basic radar range to work.

[15] `http://www.fas.org/spp/military/program/track/pavepaws.htm`

target and P_m is the minimum received power that can be detected[14].

A current missile and satellite tracking system like PAVE PAWS[15] uses a peak power of 582 kW, a frequency of 420 Mhz ($\lambda = 1.4$ m) and has a range of ≈ 5500 km. Scaling it up to a 1 GW radar would increase the range about 6-fold, to 35,000 km. Note that shorter wavelengths have shorter ranges: a teraherz version would have a range of just 113 km, while one using 30 kHz would have a range of 650,000 km.

While having a long range may appear useful, it also includes more potential targets and more clutter. In a space battle active sensors are more useful for pinpointing nearby incoming projectiles and direct defensive fire on them.

There is also a trade-off between range and resolution. The angular resolution is $1.22\lambda/L$, making the kHz radar useless for detecting direction. The range resolution is $c/2B$ where B is the signal bandwidth, $\propto 1/\lambda$. Shorter, more high frequency pulses have higher bandwidth; a radar with 1 m range resolution needs a frequency in the 150MHz band. Point defense radar needs very accurate position and Doppler measurements and will hence have a short range.

4.11 Stealth

Reducing the profile of a ship or asset requires reducing emissions that can be seen with passive sensors, and preventing signals from active sensors from bouncing back with revealing information.

Stealthing against active sensors works if you can absorb the signal well enough or reflect it in a safe direction, reducing the radar cross section. Thanks to metamaterials and advanced materials science this is often possible - for particular frequencies. It is generally not possible to stealth against all frequencies, so if the enemy uses the wrong sensors the invisible object will be obvious. Some military ships can reconfigure the surface metamaterials to adapt to expected opponent strategies but the process is not instantaneous, taking minutes to hours.

4.11.1 Anisotropic radiation

Averaged over time the total power radiated by an object must equal the total power generated. It is possible to cool a ship surface (at an energy cost) and radiate the heat into particular directions, and to store heat into tanks for a while. However, these stealth methods have serious limitations.

A ship of power P that emits its power as a blackbody will have a surface temperature $T = [P/\sigma A]^{1/4}$ where A is the total surface area. Using only a fraction f of this area increases the temperature of the hot surface by $f^{-1/4}$ and the flux will be P/f W/m^2.

[extend]

4.11.2 Cooling

Cooling the surface using a cold reservoir at temperature T_C has maximum theoretical efficiency $\eta = T/(T - T_C)$. The amount of work needed to reduce the heat of the surface is $\Delta W = \eta \Delta Q$ where ΔW is the work if the heat pump and $\Delta Q = K\Delta T$ is the change of heat in the surface (K is the thermal capacity). Putting this together the energy cost of cooling from T_{hot} to T_{cool} is

$$\begin{aligned} W &= K \int_{T_{cool}}^{T_{hot}} T/(T - T_C) dT \\ &= K \left[T_C \ln \left(\frac{T_{hot} - T_C}{T_{cold} - T_C)} \right) + T_{hot} - T_{cold} \right] \end{aligned}$$

A typical spacecraft temperature in Earth orbit varies between 173 K and

393 K depending on light and shadow. Using an average of $T_{hot} = 283$ K and cooling to $T_{cool} = 3$ K using a reservoir at $T_C = 1$ K would take $W = K[2.15+280]$ J of energy. Assuming a radius 20 meter spherical ship with specific thermal capacity around 1 kJ/kgK, density 1.2 g/cm3 and that only the top cm need to be cooled, K is about 60,000 and the total energy cost for cooling down to stealth temperature is 17 MJ.

During stealth mode the power has to be fed directly to the cooling tank. It will last for $K_C M/P$ seconds before increasing in temperature by 1 Kelvin (where K_C is the specific thermal capacity of the cooling substance, M the tank mass). Since T_C is so low it will be enough for a few Kelvin's of increase to get above the desired stealth temperature T_{cold} (and the efficiency η of cooling drastically decreases). Using hydrogen, which has the best specific heat capacity $K_C \approx 12$ kJ/(gK) and $P = 30$ kW (power of a one-man helicopter), the ship uses up 2.5 kg of coolant per second. A one hour stealth episode would require 1,440 kg (21 m^3) of coolant. Assuming a more energetic ship of $P = 140$ MW (Boeing 747) the rate is 12,000 kg/s, and the above radius 20 ship would at most (assuming it to be entirely filled with coolant) last 187 seconds.

[cooling lasers are too inefficient - compare heat capacities. Electromagentic thermal radiation has effective volumetric heat capacity of $32\pi^5 k^4 T^3/15(hc)^3$.]

4.12 Parallax

Determining the position of something in space will be dependent on resolving its location. The resolving power of a telescope is $\delta\theta = 1.22\lambda/D$ where λ is the wavelength and D the diameter of the telescope. Parallax distance errors are $\delta d \approx \delta\theta/\theta^2$. For the ideal case of a target at orthogonal distance d to a spaceship of length L with two telescopes at the sides, $\theta \approx 2L/d$ and we get $\delta d \approx 0.305\lambda d^2/DL^2$.

For a baseline of 100 m, looking at $\lambda = 10^{-5}$ (300 K blackbody radiation), $D = 1$ m and $d = 10$ km the uncertainty in distance is about 3 mm. A target at 1000 km distance will however have uncertainty of 300 meters and at 10,000 km the uncertainty is more than 30 km - far too much for any useful targeting of even a ship-sized object.

Turning the formula around, assuming u_{max} to be the maximum acceptable distance uncertainty, the maximum range where targets can be hit is

$$d_{parallax} = L\sqrt{Du_{max}/0.305\lambda}$$

Note that increasing L increases d_{max} proportionally: having separate sensors imaging the same target from widely separated vantage points greatly extends the range from which it could be hit. By using multiple sensors and data fusion this range can be improved further to some degree. Larger telescopes and shorter wavelengths are much less effective.

4.13 Hiding from many eyes

If the enemy is known to be watching from a particular direction it might be possible to reduce detection probability by carefully aligning a mirror to hide the emissions of the ship (and avoid reflecting other bright sources), avoid accelerations and use stealthing against active sensors. Similarly it is sometimes possible to approach (or depart) from an observer along the direction towards the sun, planets, stars or even hiding in the zodiacal light (assuming very low emissions; it has a power around 0.0005 W/m^2).

However, as the parallax section shows, there is a great advantage in having multiple sensors scattered over a long distance. If the sensors are a distance L apart and the angular diameter of the background object is θ, then the hiding ship must be more distant than

$$d_{hide} = (L/2)\tan(\pi/2 - \theta/2)$$

to hide from both. In the case of the Sun in Earth orbit, for $L = 100m$ d_{hide} =10.7 km, while a 10 km baseline forces a hiding distance of more than 1,000 km. At Jovian distance the distances are about 5.2 times larger and at Saturn 9.2 (as a rule of thumb, just multiply by the distance to the sun in AU). Hiding in planetary light is harder: using Venus or Jupiter to mask an approach near Earth has a $d_{hide} \approx 300$ km for 100 m baseline and a 30,000 km distance for a 10 km baseline. Using red giant stars will not work within 400,000 km even for near-sighted $L = 100$ m spaceships[16].

Multisensor detection depends not just on failure to keep away from backgrounds that produce contrast, but also whether the sensor network can detect a discrepancy and flag it as interesting. The larger the number of sensors the higher chance one of them will detect something interesting, but at the same time the amount of data to be processed and the number of errors will increase. If there are N sensors and each has a probability p per second of generating a false positive ("'crying wolf"') the probability per second of avoiding false alarms will be just $(1-p)^N$. In practice very sophisticated data fusion algorithms can be used to get robust estimates, handling faulty or even suborned sensors (c.f. the Brooks-Iyengar algorithm, which works up to $N/3$ faulty sensors). However, the best algorithms also require significant network bandwidth as each sensor needs to communicate with every other.

[fake blackbodies] [metamaterials] [occultation probability] [directional radiation] [changing direction]

4.14 Spoofing and jamming

While correctly imitating the exhaust plume from an accelerating spacecraft is hard (the luminosity, spectrum and doppler shift need to match the original, and this requires essentially the same engine and performance as the original) in the high-noise environment of a space battle it is likely possible to produce distracting or apparently similar phenomena. This might mislead sensors, targeting systems or point defenses.

[spoofing doppler] [spoofing background] [confusing sensors]

It is easy to clutter radar and IR by releasing chaff that reflects signals strongly or in the right wavelengths. Exactly how well chaff works depends on the signal processing abilities of the enemy.

In particular, sensors are easily blinded by bright detonations or deliberate scorch attacks with beam weapons. This either permanently damages them or leaves them blind until they recover. Having replacement sensors that can be opened when the current one are down will be necessary, but still introduces a short delay of observation. Assuming a fully functional C3I system one side can time closing sensor ports with the arrival of energy from their own detonations, giv-

[16]Of course, the ship or object has to have an angular diameter much less than the background for this trick to have a chance.

[17]As a simplistic example, sensors could be watching during even seconds and attacks timed to occur during odd seconds. In practice the pattern would have to be pseudorandom and take lightspeed delays into account.

ing itself an advantage. This is easier with QE signalling, but even without it some synchronization is possible[17].

[overwhelming point defenses with large clouds of objects

If point defenses can "'drill"' X meters of object per second, then attacking with more than that (per second) will allow a hit on the ship.]

5 Weapons

The forms of weapons available for ranged space battle are lasers, railguns, particle weapons, missiles and fighters. Nanoweapons can be included in railgun and missile payloads.

5.1 Independent weapon buses

Weapons can either be placed on the ship, or on free-flying weapon buses. Placing offensive assets outside the ship has the advantage of allowing shots at closer range and without risk to the main ship, but also limits the amount of energy they can provide since they are too small to contain reactors.

Chemically stored energy is on the order of 20 MJ/kg (≈20 GJ/m^3) while nuclear isomers could reach 10 GJ/kg (≈100 TJ/m^3). A volume V weapon with energy density ρ and efficiency η (likely < 0.5) will be able to fire $N = \eta\rho V/E_{attack}$ attacks of energy E_{attack}. The remaining $(1-\eta)\rho V$ energy becomes waste heat; the resulting temperature ($T = (1-\eta)\rho/C$ where C is the heat capacity in J/m^3) of the weapon will scale proportional to the energy density of storage. Using large amounts of energy makes the weapon heat up significantly, making it hard to hide (in addition to energy release from orientation thrusters). Most weapons are hence intended to fire just a single burst attack and then coast away on the recoil.

Missile buses also have this disadvantage, but to a lesser degree: missiles can be launched quietly from a stealthed bus using springs or cold gas and their engines activated at a distance from the bus. The main use of a missile bus rather than individually drifting missiles is that the bus can be equipped with maximal stealth, hopefully drifting close to a tactically important volume before releasing the missiles[18] (at the price of risking losing all missiles to an unexpected attack: some bus designs have an emergency eject function that releases the missiles prematurely if the bus comes under attack).

[railguns in space can be made long]

[snap open phased arrays - since disposable device, little need keep low profile after firing]

5.2 Lasers

[Expand, update]

Laser weapons work by either heating a target to an unsustainable temperature (which requires a long lock on the same spot providing more than 10^8 W/m^2 over a second or more), a rapid energy impulse causing a local plasma detonation (requires on the order of $10^{13} - 10^{14}$ W/m^2) or drilling through the outer shell (requires an energy density of LE_{eap} W/m^2 where L is the desired drilling distance per second and E_{vap} is the energy needed for vaporisation per cubic meter).

5.2.1 Sources

Laser beams can be generated using single laser cavities or phased arrays. Laser cavities contain a gain medium where atoms, molecules or free electrons are placed in an excited energy state and then stimulated to decay to a lower energy state, releasing electromagnetic waves that trigger other de-

[18] There is some potential for game theory in whether to launch all missiles or leave one as a surprise later launch. A revealed missile bus is easy to hit and hence has little value to store a surprise in, but in a situation where long-range defenses are busy (since there are approaching missiles) it would be a low priority target. This leads to a mixed strategy equilibrium where the missile side randomly leaves surprise missiles and the defender side randomly decides whether to shoot at the "worthless" target.

cays and shoots out as a beam. This is relatively simple but has the problem that to function the cavity needs to be resonant: the waves must be able to bounce between the front and back to produce a resonance, and this means the cavity itself will need to resist the laser power. Worse, putting the medium into an excited state involve big energy losses that also heats the medium. Hence large amounts of cooling are needed.

Phased arrays make use of many smaller lasers or antennas, producing a beam by combining many small components accurately. For lasers phased arrays need to be manufactured using nanotech metamaterials.

5.2.2 Beam optics

(Gaussian) laser beams have a divergence angle of $\lambda/\pi w_0$ where λ is the wavelength and w_0 its smallest width. If the laser is produced by a lens of width L it will produce a spot size at distance d $w = \lambda d/2\pi L$ and with intensity $\pi P L^2/\lambda^2 d^2$ W/m^2 if P is the total beam power. Short-wavelength lasers remain sharp over longer distances than long-wavelength lasers, and in space it is possible to go all the way down the the "‘vacuum frequencies"’ of UV around 10^{-8} m that are strongly absorbed by air and other matter. Large lenses allow tightly focused beams. However, with the nanotechnology available in Eclipse Phase phased array lasers are possible: many small elements producing parts of the beam, possibly focusing it closer to the source if needed (this also allows higher power densities at the target than at the source, always a nice thing for a weapon). With a size L array focused on the distance d target the focal width is $w_0 = 2\lambda d/\pi L$ and the length of the focal region is $\approx L^2/4\lambda^2$.

[example, showing that the region is usually long]

When firing a laser the ideal beam width at the target is $w_{opt} = 2(\delta v)d/c$, the uncertainty in target lateral velocity δv times the lightspeed delay between the target and laser (if there are QE-linked sensors shortening the delay this is reduced further down to a minimum of $(\delta v)d/c$). If the velocity measurement is perfect there is still a $\delta v < 2ad/c$ due to unknown accelerations since last observations (with QE the factor 2 approaches 1). So the optimal beam width will be

$$w_{opt} = 2(\delta v)d/c + 4ad^2/c^2$$

However, close to the laser the actual beam width will be limited by the focal width of the beam, giving a beam width at the target of $\min(w_{opt}, 2\lambda d/\pi L$. The distance beyond which velocity uncertaity dominates is $d = (c^2\lambda/2\pi a L) - (c(\delta v)/2a)$. Typically this depends on the beam wavelength, with IR lasers being uncertainty-limited and UV lasers diffraction-limited over combat distances.

Now, this suggests another reason you don't want to fight close to planets: stationary defence stations can easily set up pretty big phase array lasers, and then they can blast you very well. Ships could in principle unfold big arrays too, but I expect it is hard to both power them and keep them accurately pointed while dodging incoming lasers, projectiles and missiles.

This kind of laser arrays still have the problem that if you are uncertain of exactly where the enemy is (and we are talking about meters here) you will miss him. So my previous calculations still apply - the Titan moon lasers can vaporize nearly anything, but if you are more than a few thousand kilometres from a sensor that pinpoints you and flying evasively, they will not be able to hit you.

[Energy requirements, size requirements, range]

[calc neergy needed for 1 m penetration - kills smaller ships, scale up 10 to hurt big ships]

[cycling time] [flash cooling]

5.2.3 Heating beams

A beam pulse of length t_p that has a circular radial spread at the target r arrives after a delay t_d from observation. During this time the target has accelerated with acceleration a, having an unknown relative velocity at_d, which will have become $a(t_d + t_p)$ at the end of the pulse. Assuming the beam hits, it will move across the surface since the firing system is not properly taking the unknown components of target motion into account. The total area heated will be $\pi r^2 + 2ra(t_d + t_p)$.

The velocity smear will reduce the heating, and preclude damage if the area is more than K times the intended. This gives the requirement

$$t_d + t_p < (K-1)\pi r/2a$$

.

An 1 m^2 object that of mass M and specific thermal capacity K that is heated by power P Watt/s for t_p seconds will reach temperature Pt_p/MK if it cannot radiate away the heat. The time needed to reach a damaging temperature T_dam is roughly $t_p = MKT_dam/P$ shorter pulse lengths require proportionally higher power, up to the limits set by the firing array.

In reality the target will reach an equilibrium temperature where the influx equals the thermal radiation. $P = \epsilon\sigma T^4$ $T_eq = [P/\epsilon\sigma]^{1/4}$ If T_eq is too low there will not be any real damage. The required power is $P > \epsilon\sigma T_dam^4$ (per square meter).

A rough calculation of the time needed to reach this temperature is $T_eq = Pt_p/MK$ $t_p = MK[1/\epsilon\sigma P^3]^{1/4}$ t_p needs to be shorter than this in order to avoid large energy losses.

If the vaporization energy per kg is E_{vap}, in time t_p it can vaporize to a depth $z = Pt_p/pir^2rhoE_{vap}$ (assuming a $t_dt_p << z$ - otherwise the movement of the target will prevent drilling) If a certain depth z is required for doing useful damage, this leads to the condition $r^2 > t_p[P/pirhoE_{vap}z]$

Diagram of $r - t_p$ plane for fixed intensity r must be larger than diffraction limit r must be smaller than limit of damaging power t_p limited by heating of array - too fast gets too hot t_p limited by equilibrium temperature

$r^2 << P/piat_drhoE_{vap}$ - drilling ability $r^2 > t_p[P/pirhoE_{vap}z]$ - able to get to damage depth

constraint due to targeting

5.2.4 Explosive beams

A very high power does damage not by heating the target but by vaporising the surface layer, creating a pressure wave fed by the beam.

5.3 Railguns

Rick Robinson's First Law of Space Combat: "An object impacting at 3 km/sec delivers kinetic energy equal to its mass in TNT."

Railgun projectile speeds: currently a few kilometers per second, comparable to normal inter-spacecraft velocity differences (even running into a "'stationary"' pebble will do significant damage to a ship). It is plausible that in Eclipse Phase ship-launched railgun projectiles will be significantly faster, between 10 and 1000 km/s. Railgun projectiles need to be tens or hundreds of km/s in order to hit fleeing ships, but can often move more leisurely.

The kinetic energy from a 10 km/s 1 kg impact is 50 MJ. At this point the kinetic energy starts to become bigger than any (chemically) explosive force that can put in the projectile. 100

km/s is 5 GJ (about one ton of TNT) and 1000 km/s projectiles release 500 GJ (100 tons of TNT).

A railgun projectile of mass m moving at speed v takes $(1/2)mv^2/\eta$ J to launch at efficiency η. Conversely, with an energy E the projectile gets velocity $v = \sqrt{2\eta E/m}$.

multi-barrel railguns can be made shorter 6 km/s, 20000 G accel, 100 kg projectile 90 m, 4 barrels, temperature increase 450 K sleeve thickness 10 mm, sleve radius 0.09 m, outside radius 0.07m

They will penetrate to a distance about equal to the projectile length times the ratio of projectile to armour density (Newton's penetration law). This is actually a problem: the attacker wants to deposit all the energy inside the ship, so they must tune the projectile length to the target armour. Too heavy projectiles will go straight through the ship. Sometimes having no armour at all is the optimal strategy (just hope they do not hit any antimatter containers). Too light projectiles and all energy gets deposited outside the armour. [Merge with armor section discussing this?]

[Energy requirements, size requirements, range]

5.4 Particle weapons

Particle weapons produce beams of heavy relativistic particles. Given the existence of personal particle beams and fusion engines (which are essentially propelling proton-electron plasma in a beam) larger particle beam weapons are plausible. The advantage of particle beams is that they deposit their energy deeper into the target, producing a stronger detonation and depositing Bremsstrahlung and secondary particles into sensitive nanosystems.

Unlike lasers they are hard to focus and will tend to disperse over long distances. Beam divergence angles are on the order of $\theta = 4.5 \cdot 10^{-8}\sqrt{T/Z}$ where T is the beam temperature and Z is the atomic number of the beam particles[19]. If a beam power P and initial radius r is directed at a target at distance d the intensity will be $\approx P/4\pi(r + k\sqrt{T/Z})^2 \approx PZ/4\pi k^2 T d^2$ W/m^2. The halving distance (where the energy per square meter has declined to half) for a proton beam ($Z = 1$) is 280 km, for a mercury ion beam 2,500 km.

[check this, compare destroyer engine]

[Energy requirements, size requirements, range]

5.5 Missiles

Missiles can presumably accelerate at least a few 100 G.

missiles that burst into clouds of shrapnel as they approach or if hurt [calc velocity needed to cover target]

number of fragments that can be zapped on approach

[calc saturation with how many meters laser point defenses can burn through - indicates how many missiles can be shot down per second and what the overwhelming number is]

can act as vector denial system

[Energy requirements, size requirements, range]

5.5.1 Kinetic impactor rods

Kinetic missiles are little more than mines, making use of ship delta v rather than their own acceleration. Missiles typically have small delta v (a few km/s), but this is enough to do serious damage.

either straight - armor piercing sideways - higher imact area, good for

[19] `http://www.projectrho.com/rocket/rocket3x1.html#particle`

weakly armored targets usually built so that a hit disperses rods

5.5.2 Nukes and antimatter

Both nuclear warheads and antimatter have roughly the same effect: a massive release of gamma rays (and some particles, especially from enhanced radiation warheads), with a small expanding cloud of plasma. The kinetic impact, heat blast and EMP that occurs in planetary environments is absent making the kill radius small, on the order of kilometers. A high precision hit is hence required.

An x kt detonation provides an energy of $4.19 \cdot 10^{12} x/4\pi d^2$ W/m^2 at distance d in about a microsecond. In order to produce impulsive shock damage (the vaporisation of material moves faster than the speed of sound in the material) on the order of 10^{13} to 10^{14} W/m^2 is required, producing a minuscule kinetic kill distance of

$$d_{kkill} = 5.77\sqrt{x}$$

meters. The most likely effect is to melt part of the facing surfaces, possibly vaporising a thin layer. This may or may not be disabling.

However, the amount of particles is likely enough to kill biomorphs and sensitive nanotechnology at a longer distance. Using the loose estimates in Nuclear Rocket[20] a conventional nuclear weapon produce an X-ray fluence of $2.6 \cdot 10^{27} x/d^2$ and neutron fluence of $1.8 \cdot 10^{23} x/d^2$, and a unshielded biomorph will receive a dose of $1.78 \cdot 10^9 x/d^2$ Grays acute x-ray dosage and $7.2 \cdot 10^8 x/d^2$ Grays of neutron dosage[21]. In order to exceed 20 Grays (immediate disorientation) the ship needs to be closer than

$$d_{rkill} = 9433\sqrt{x}$$

m. This does not take radiation protection from the spaceship into account: if the ship armor absorbs a fraction f_X of x-rays and f_n of neutrons the distance changes to $d_{rkill} = \max(9433\sqrt{f_x x}, 6000\sqrt{f_n x})$ m. For a ship with 5 cm of steel and 5 cm carbon in the hull $f_x \approx 0.62$ and $f_n \approx 0.85$[22], giving $d_{rkill} = 7428\sqrt{x}$ (still dominated by the x-ray damage in this case).

Using nuclear shaped charges ("Casaba-Howitzer"') jets of plasma travelling at 10,000 km/s can be generated, transmitting up to 5% of the total energy of the detonation as kinetic energy in a cone with half-angle 0.1 radians[23]. Assuming it can be directed accurately at the target, this would have a

$$d_{kkill} = 667\sqrt{x}$$

meters and (unshielded) radiation kill out to

$$d_{rkill} = 21,000\sqrt{x}$$

m. Since the beam diameter is $0.2d$ m, it has a fairly broad cross-section for hitting a spaceship.

Nuclear warheads have the advantage of guaranteed stability, but antimatter packs significantly more punch per weight. Theoretical limits on fusion warhead mass are ≈ 1 kg per kiloton (current warheads closer to 500 kg per kiloton). A warhead requires a 45

[20] http://www.projectrho.com/rocket/rocket3x1.html#nuke

[21] Antimatter weapons would instead deliver nearly all energy as gamma rays.

[22] $f \approx 2^{-\sum_i t_i/v_i}$, where the sum is over all materials, each with thickness t_i and half-value thickness v_i for the relevant kind of radiation. This is at best an approximation since it ignores the effects of secondary radiation, Bremsstrahlung from charged particles and other complications.

[23] Winterberg, Thermonuclear Physics, p.41, 122.

[24] http://www.projectrho.com/rocket/rocket3x1.html#nuke

kg/kt missile with a volume of 0.036x m^3/kt (based on extrapolations from the US trident missile[24]). Antimatter warheads have 43 megaton yield per kg of mass, requiring 10 times the amount of shielding. Using previous numbers implies a missile mass of 0.01 kg/kt and $8 \cdot 10^{-6}$ m^3/kt. So small chemical missiles are likely not possible, but small antimatter "'candles"' might be possible.

Given the stability risks of antimatter thermonuclear devices are likely preferred on ships not using antimatter engines and not expecting significant combat. Antimatter warheads however have the advantage that if point defenses destroy them at close range they will go off anyway, doing damage the less volatile nuclear warheads will not do.

`http://www.5596.org/cgi-bin/nuke.php` has a nuclear effect calculator.

5.5.3 Rocks

It is entirely feasible to boost small asteroids into orbits that will impact targets. The upside is that the mass can be significant, the damage enormous, and the asteroid is hard to stop (especially if equipped with some point defenses). The downside that the impact can usually be predicted weeks or months ahead. Rocks represent a relatively minor threat to space habitats and aerostats, which can be moved out of the way. They represent a significant threat to planetary or asteroid habitats that are not defended by defense arrays. However, since they would tie up the defense array (and produce fragments) as they get within range they also provide an ideal time for a conventional attack.

5.5.4 Relativistic missiles

More of a strategic weapon than a realistic space battlefield weapon, relativistic missiles attempt to hit targets at long distance using projectiles with a sizeable relativistic mass. Their main benefit is the impossibility of protecting oneself from them, since they would not be detectable from the target until shortly before impact and there is no way of stopping a very large mass moving fast (it would pass through armor as per Newton's impact law, and any effect that splinters it will not be able to make the fragments deviate very far before it hits).

Fortunately, given the energy requirements of launching relativistic missiles launch is very visibile, and QE allows monitoring of possible launch sites giving information long before the missile reaches its target. There are also some doubts on whether relativistic missiles are feasible or not in Eclipse Phase.

5.5.5 Leashes

A leash is a warhead that is attached to the hull of the enemy ship, ready to explode on command or if tampered with. In principle just placing a missile with a functioning warhead on or near the ship works, but in practice the extra reassurance from tamperproofed antimatter containment is preferred. Actually leashing a ship is quite a coup, and allows the leasher to dictate terms to the leashed.

5.5.6 Nanoweapons

Nanomachines encased in diamondoid shells can resist accelerations up to $10^8 - 10^{10}$ G, making delivery by high velocity impactors possible if they can be slowed just before impact. This can for example be done by placing them far back in a penetrating warheador having an explosive charge accelerate them in the opposite direction just before impact. The main problem with nanoeweapons is that the amount that can be transfered is relatively small compared to the amount of defensive nanosystems onboard (see section 7) and that hives generally cannot survive impacts. This makes direct attacks using disassemblers weak. The main nanoweapon payloads tend to be sabotage nanites or proteans constructing a position tracking transmitter.

5.6 Fighters

essentially a roving missile bus

5.7 Networked warfare

Generally, I expect space battles to involve extended networks of decoys, drones, munitions, sensors and whatnot. The communications issues are serious: I expect it to be worth the money to use FTL quantum communication to keep everything linked, untraceable, fast and unjammable. Qubits are a very strategic resource (incidentally, I doubt they can be stored without some very good nano - bad news for the Jovians, who probably have to break a few rules to get it). Primitive forces that have lightspeed-limited networks are at a serious disadvantage, and must also ensure that the enemy cannot detect *where* the cloud sends its messages (OTP encrypted neutrino broadcasts in all directions instead?)

Burnside's Zeroth Law of space combat: "Science fiction fans relate more to human beings than to silicon chips." While the polities of the solar system have very good reasons not to like AIs, AGIs and infomorphs controlling weapon systems the military benefits often outweigh these considerations. On the strategic level there is a tension between avoiding developing systems similar to the ones implied in the Fall and keeping a military edge.

6 Armor

If we assume that the cost of drive systems is large compared to the cost of weapon systems, then it is clear that warships will tend to be as protected as possible. This can be achieved through heavy armoring or making them able to evade (through stealth, maneouverabilityor point defenses) attacks.

Armor needs to be able to handle hypervelocity impacts and beam weapons.

6.1 Kinetic impact

Any impacts faster than a few km/s will produce inertial stresses much larger than the material strength (even for extremely hard materials like diamond): for all practical purposes the armor behaves like a liquid. This is not just a military problem but also an issue for any ship moving at orbital velocity, since impacts with micrometeorites and other debris is potentially dangerous (at speeds of 800 km/s 1 gram impacts correspond to explosions of 76 kg of TNT).

One way to handle this is Whipple shields: a relatively thin outer bumper shield placed a distance from the wall of the spacecraft. The incoming particle will be shocked and may disintegrate, spreading its energy across a larger area of the inner wall. Shields can be stacked, further dispersing the energy, and the space between them filled with shock or radiation absorbing material. They are light but increase the spacecraft size; some ships use unfoldable Whipple shields. The outer shields tend to be damaged by micrometeors and conflict, but are expendable and can be cheaply repaired afterwards.

Newton's law of impact depth states that a projectile of length L with density ρ_P impacting a target with density ρ_A will penetrate to a depth $L(\rho_P/\rho_A)$. This is applicable for hypervelocity impacts where material cohesion can be ignored. While high density projectiles are possible (osmium achieves 22610 kg/m^3, 2.9 times steel) just extending the projectile into a spear guarantees deep penetration. This will also tend to penetrate Whipple shields, since the desintegrating parts will just make way for more incoming spear.

Heavy armor needs to be thicker than spear projectiles unless it is significantly denser. This is hard to achieve across a whole ship since the mass becomes prohibitive: even a heavily armored ship will just armor vital systems (reactor, cooling, antimatter confinement, possibly crew battlestations) and rely on redundant or repairable systems elsewhere. Another approach is to use thick low-density shielding in the form of water or reaction mass tanks that would anyway be present. A hit will destroy the compartment but dissipate the energy.

Another approach is to have very light armor and allow projectiles to pass through, relying on redundant ship systems to survive the damage and nanoswarms to repair it. However, high velocity impacts turn into rapidly expanding clouds of shrapnel. A velocity v impactor of mass m will release $0.5mv^2$ J of kinetic energy. If all of it turns into kinetic energy of fragments (and assuming them to have a Maxwell-Boltzmann distribution[25]) the average lateral velocity of the fragments will be $v/\sqrt{2}$, corresponding to a $approx 35°$ cone. Hence this strategy works best for very long and thin ships, where little of the volume can be

[25] There are m/m_f fragments, where m_f is the fragment mass. Each gets $E = 0.5m_fv^2$ kinetic energy if it is divided evenly. The average speed in the lateral direction is $\sqrt{E/m_f}$, producing the above formula.

affected.

6.2 Beam impact

[Flash damage, impulse kill, drilling The energy needed to drill through an object is within a factor 2 of the heat of vaporisation of the object best vaporization energy for mass carbon, 29.6 kJ/g (boron even better) des 5 g/cm2 (increase density?), burning a 1cm2 hole requires 148 kJ and 20 milliseconds combat conditions need larger spot to remain focused - 10 cm2 spots accepts uncertainty velocity 5 m/s]

6.3 Conclusions

One approach is to add maximal armor, making the target hard to damage. [how handle laser effects?]

Self-repair using nanomachines can restore Whipple shields fairly rapidly after a battle, but is too slow to matter during a battle.

7 General strategy

The conflict between different military assets can be described using the Lanchester ("square"[26]) model:

Let B represent the number of assets of side "blue" and k_b is the destruction efficiency, how many enemy assets one blue asset can neutralize per unit of time. Let R and k_r represent the "red" side. The both sides will suffer attrition like

$$B' = -k_r R$$
$$R' = -k_b B$$

Blue will win by attrition (i.e. reduce R to 0) if $k_b B^2 > k_r R^2$. There is a sizeable advantage in having more offensive assets, bigger than merely lethal assets. Concentrating forces so they can focus their fire gives a big advantage (but the concentration needs to be diffuse enough so that no volume effect weapons can strike at them, something the model does not include). The side with fewer assets still have a chance of winning if they can make strikes that disrupt the command and control structure of the other side, for example taking out the main communications links (again outside the basic attrition model).

In practice the Lanchester equations will be just a crude approximation, since there exist weapons that destroy groups of enemies, some weapons have synergistic effects (detonations that temporarily blind enemy sensors yet provide active sensor information about enemy asset locations), there are sensors and other assets that do not attack yet are valid targets, and there are different kinds of weapons that work differently well against different targets. The key point likely remains: the side that can manage to rapidly reduce the enemy assets early on has a large advantage.

In traditional warfare the side with smaller numbers can improve its chances if it splits up into hard-to-find units and make local raids where it has numerical superiority[27]. This is hard to do in deep space combat, since visibility is high, but an approximation is to use missile buses and fighters to rapidly deploy local superior force when needed. In orbital warfare it is easier to do "guerilla" tactics, making surprise ambushes and retreats into unmonitored volumes.

[26]The name comes from the square in the winning criterion, the equations themselves are linear. Ironically there are also the "linear Lanchester model" that describes conflicts where the loss rate is proportional to the product of the force sizes (and hence forms a pair of nonlinear equations). This model describes situations where combat occurs between pairs of units, and here the advantage in number is reduced to a linear relation.

[27]S.J. Deitchman, A Lanchester model of guerilla warfare, *Operations Research* 10:6, 818-827

8 Cloud combat

Much of deep space battles occurs between distributed assets drifting through space, forming two (or more) clouds passing through each other at multikilometer per second speed. If friendly assets have not cleared enough enemy assets ahead of the ship, it will likely be hit.

Assets can be made very low-emission by cooling before launch, mechanical low-energy launching (springs, slow railguns, even bola launching), internal heat storage, low emission electronics, being folded together and equipped with stealth covering. This can also cause a hard-to detect change in ship velocity that could grow into positional uncertainty before the battle.

A cloud launched at time T before combat at velocity v_s will have radius $v_s T$. The passage through the cloud of a ship of velocity v takes $2v_s T/v$ seconds. v_s is presumably restricted to less than 1 km/s if silent launch is used. If $v = 10km/s$ and $T = 1$ day the passage takes 4.8 hours, while for $v = 100km/s$ passage takes 29 minutes and for $v = 400$ km/s just 7 minutes.

However, after the first side has launched their cloud the other side can choose to do a maneouver with $\Delta v > v_s$, avoiding the cloud and being able to launch a cloud of protective assets. Obviously the first side can themselves change the velocity to make this cloud worthless, and so on. If continued this turns it into a pure ship-to-ship battle. This is the preferred strategy for a side with limited spaceborne assets and plenty of Δv to burn.

If a ship moving with velocity v towards another ship launches an asset cloud spreading with velocity v_s when at distance d, the cloud will have radius $v_s d/v$ when it reaches the other ship. To evade it, the ship needs to change its velocity by at least v_s (favoring rapid, possibly detectable launches). However, the density of assets in a wide cloud will be lower. If assets have a range r and N assets are launched in the cloud, on average a ship will be within range of $3Nr^2v^2/4v_s^2d^2$ assets if passing straight through the cloud. This favors a late and slower launch, producing a denser cloud. If M assets within range are needed to achieve a win, they should be launched within distance $\sqrt{3N/4M}(rv/r_s)$, but if the enemy is expected to have a high enough Δv budget earlier and sub-optimal launches are needed - much hinges on accurately estimating how much fuel the enemy can spend on evasion and the performance of their engines.

8.1 Weapons vs. sensors

A cloud of density ρ has average distance to the nearest asset $< d >= k\rho^{-1/3}$, where $k \approx 0.5622$.

If there is resources C to produce assets per unit volume, and sensors cost 1 and weapons x, $C = \rho_s + x\rho_w$ where ρ_s is the sensor density and ρ_w is the weapon density. Assumine a fraction f of the resources are used to make weapons, $\rho_w = fC$ and $\rho_s = ((1-f)/x)C$. The average delay between a sensor detecting a target and a weapon has a chance to hit it is:

$$t = k\left[\frac{\rho_s^{-1/3} + \rho_w^{-1/3}}{c} + \frac{\rho_w^{-1/3}}{v}\right]$$

where v is the weapon velocity. The middle term represents a delay as information is transmitted at lightspeed to the nearest weapon. If a QE link is present it is zero.

$$t = kC^{-1/3}\left[\frac{x^{1/3}(1-f)^{-1/3} + f^{-1/3}}{c} + \frac{f^{-1/3}}{v}\right]$$

To minimize t the expression in the bracket has to be minimal, leading to the optimal weapons allocation

$$f = \frac{1}{1 + (1 + c/v)^{-3/4} x^{1/4}}$$

If x is very low (cheap weapons) most assets should be weapons. For weapons as costly as sensors $\approx 50\%$ should be weapons if $v \ll c$ while $\approx 62\%$ if they are lasers: a railgun or missile-based system will tend to have numerous weapons, hoping to be close to a detected target. For more expensive weapons a smaller fraction is optimal.

[check QE term, check other case]

9 Hitting a dodging enemy

Suppose a vessel observes another vessel at distance d and known velocity and fires immediately with a projectile with velocity v ($v = c$ for a laser)[28]. The time it will take for light from the enemy to reach the vessel and for the projectile to reach the vicinity of the enemy is

$$t = d(1/c + 1/v)$$

During this time the enemy will be accelerating with acceleration $a < a_{max}$ in a random direction to escape a hit. After time t it has moved up to

$$r_{max} = (1/2)a_{max}t^2$$

and will be within a sphere of radius r_{max} around the position it would have had if it had not accelerated. The attacking vessel does not know where in the sphere the enemy is, but we will assume it knows r_{max} (for example by observations of past accelerations and ship type).

A projectile will hit anything along its path through the sphere within its effective cross section σ (this includes the area A of the ship and the radius r_p of the projectile, $\sigma \approx A + 2\pi\sqrt{A}r_p$; if N shots are optimally fired σ is multiplied by N). This is equivalent to selecting a point on the cross-sectional disk seen by the launching ship: the area of the disk is πr_{max}^2, and an area σ will be affected. The evading ship will attempt to distribute its probability across the disk uniformly[29]), so the probability of hitting will be

$$p_{hit} = \sigma/\pi r_{max}^2 = \frac{4\sigma}{\pi a_{max}^2 d^4 (1/c + 1/v)^4}$$

When $p_{hit} \ll 1$ the evader has a good chance of avoiding a hit. While a high acceleration ability is useful, increasing distance has a far greater effect.

The critical distance $d_{evasion}$ where p_{hit} approaches unity is

$$d_{evasion} = \left[\frac{4\sigma}{\pi}\right]^{1/4} \frac{1}{\sqrt{a_{max}}(1/c + 1/v)}$$

Inside this distance the probability of hitting is large. The first term varies slowly with σ and is of the order unity. For $\sigma = 1$ m^2 it is ≈ 1.06, for $\sigma = 100$ m^2 it is ≈ 3.36 and for $\sigma = 1000$ m^2 ≈ 5.97.

As an approximation, the enemy is possible to hit if $d < d_{evasion} \approx v/\sqrt{a_{max}}$ for projectiles and if $d < d_{evasion} \approx c/2\sqrt{a_{max}}$ for lasers. For

[28]This is also equivalent to having a sensor and a weapon at equal distance d from the target and a QE link between them.

[29]Since otherwise the attacker could distribute their attack probability across the disk correspondingly and increase the probability of a hit: a uniform distribution is a saddle-point for the game. The direction in which to accelerate should obviously be completely random. The optimal probability distribution of radial distance is $f(r) \propto 1/\sqrt{1-(r/r_{max})^2}$, since this produces a uniform cross-section probability. However, if the evading ship does not know *when* an incoming weapon arrives it cannot pre-calculate r_{max} (which is time dependent) and might hence be distributing itself suboptimally. However, $f(r)$ places most of its probability weight at extreme r, so a good guess is to accelerate close to a_{max} and change direction often.

a capital ship with $a_{max} = 0.1g \approx 1$ m/s^2 lasers will start to have a chance to hit at a distance of 150,000 km. Projectiles moving at 100 km/s will become dangerous at 100 km distance (as a loose rule of thumb, the critical distance corresponds to one second of projectile movement).

[Calculations for hitting fighters, missiles and incoming projectiles.] [The later have $a_{max} = 0$, so they can be hit (assuming low position and velocity uncertainty) if there is enough time to fire and disperse them]

[The FTL case - how big advantage is it?]

In the situation where there is no QE communications between sensor and weapon, and the target, sensor and weapon are at roughly equal distance d the effective distance is increased by 50%. This means $d_{evasion}$ is reduced to 2/3 for beam weapons, while projectile weapons roughly retain their (short) range. Hence low-tech forces will be forced to get in closer, use denser asset clouds, use more projectile weapons, or combine sensor systems with their weapons launchers.

10 Conclusions

Putting together the results above suggests the following overall conclusions about Eclipse Phase space combat:

10.1 Deep space combat

In deep space combat two ships pass each other far from any planets or other objects. Generally encounters will occur with velocity differences at least several kilometers per second, sometimes far higher. Whether the ship with the largest potential Δv can get close to an escaping ship depends on how much fuel they have left, the initial distance and the initial velocity difference. In many cases ships will be limited to relaying their observations to homebase since they are unable to close enough to attack and fulfill their original mission.

[Need for a section on pursuit capture differential games?]

Since detection range is far larger than weapons range, yet doesn't allow fast closing the two ships typically have plenty of time to prepare for the upcoming fight. The ships can adjust their velocities to change the situation, for example by increasing their relative velocity (reduces the length of the confrontation, makes kinetic weapons more deadly if they hit), decreasing it (increases the length of the confrontation, allows more launching of fighters ahead of the ship, makes kinetic weapons less powerful and increases the utility of point defenses), passing ahead of the enemy (allows leaving a cloud of kinetic "'mines"' that it could collide with) or shifting into trajectories that will have other tactical or strategic effects (e.g. require even a successful enemy to make a large and potentially dangerous detour). A late trajectory adjustment can make earlier launched equipment miss the predicted battle volume.

At some time before the battle both sides begin launching assets: sensors, positioning systems, weapons, decoys, interceptors and fighters. A quiet launch early enough is hard to detect and allows the slow dispersion of hard to detect assets over a large volume. A problem is that too early launch will disperse them too widely, reducing their density in the combat volume and making recovery hard or impossible. A too late launch will guarantee a dense cloud that is easier to scan and avoid. Many warships are little more than engines carrying an assortment of disposable military equipment around. Since the engines are typically the largest and most expensive part destroying or neutralizing it tends to be the tactical goal. Even if the engine is lost the assets can still fulfill a successful attack, but they will not be able to change course afterwards.

In deep space combat a fighter is little more than a mobile weapons bus, and its main purpose is to engage in battle before the main ship has arrived and with a higher Δv budget, hopefully unsuring victory and no risk of a successful counterattack on the engine.

Combat usually begins when one side attempts active scanning of the whereabouts of opposing assets or begins beam weapon attacks, and the other side responds by blinding sensors through jamming or detonating some suitable weapon. Henceforth the game is to sense the enemy assets without being blinded while preventing him from seeing where your assets are. If one side is successfully blinded it will be unable to damage the enemy while the enemy might well be able to not just damage the main ship but to force surrender by placing a suitable weapon on or near it. Boarding is rare due to the high Δv demands, but might occur if there is time after surrender.

During the core combat the distributed sensors attempt to pinpoint key targets (the main ship, its distributed weapons, its sensors) and signal to nearby weapons to destroy or neutralize them. The ships and weapons try to avoid being hit by evasive maneouvers and point defense fire against incoming missiles and kinetic weapons. Point defenses must be overwhelmed, which favors ships that can distribute a large number of incoming objects or hit using beams.

After the clouds have passed each other there is seldom a chance for a second pass. The survivors (if any) gather up assets and continue on their way. Expanding debris clouds will remain a navigational hazard for some time but usually disperse within a few days to the degree that they are irrelevant.

The goal of deep space combat is usually to prevent enemy ships from arriving near habitats or other vulnerable locations, proactively defeating them before they can do damage. At the same time luring defenders away from the target is also a valuable strategy.

[Diagram showing the different ranges of sensors and weapons]

[what is range of small weapons?] [what is prob hit? how small can weapons be made?]

10.2 Orbital warfare

Orbital warfare is battles near stationary targets, in particularly planet-moon systems. Unlike in deep space combat everything is within or close to firing range, large volumes are hard to sense, and there are many neutral or irrelevant targets. Some of the apparently neutral or irrelevant targets may also be enemy assets in disguise, making the combat zone full of surprises.

Typical orbital warfare objectives involve achieving orbital dominance (total control over traffic within a volume), trade or transport interdiction (ability to destroy or intercept civilian ships entering or leaving the system) or destroy certain habitats or objectives.

Hitting "stationary" targets like orbiting habitats is very easy but some can have fearsome armor such as beehive habitats surrounded by many meters of astroid regolith or planets with an atmosphere. Tincan, torus, Reagan and Hamilton cylinders are on the hand very vulnerable to attacks: at best they can use point defenses against incoming kinetic weapons and missiles, but they can easily be damaged by beam weapons. The amount of debris caused by this is potentially a major navigational hazard along the orbit of the targeted structure.

Planets, moons and asteroids are potential "fortresses". Not only are they naturally armored, they also have ample space, energy and heat sink capacity for defenses. Places with atmospheres are relatively safe from orbital attacks as long as their forces retain orbital superiority, but conversely have a hard time launching anti-orbit weapons or hitting objects in orbit with beams. Airless locations can use very long baseline sensors and weapons, have essentially arbitrary heat sink capacity and can in principle produce very large amounts of energy to power extremely heavy weapons. Shuttles or atmospheric fighters are generally useless in orbital warfare: they are easily detectable while launching or landing, yet have far too low velocity or evasion capability in this situation to avoid being hit by orbital or ground weapons.

Relative velocities in orbit tend to be lower than in deep space combat, and surviving opposing ships will see each other again after they have rounded the planet. Ships in retogerade orbits or doing high-velocity entry from interplanetary space can however achieve the high velocities favoring short battles. The Oberth effect (and aerobraking in some cases) also allows drastic course changes.

There are hidden volumes due to planetary masses and phenomena such as magnetotails. These can be used for surprise maneouvers, to avoid long-range weapons or approaches. However, the local forces have potentially very large advantages in having numerous hard-to-defeat sensors that can monitor most angles.