

Slingshot dynamics for self-replicating probes and the effect on exploration timescales

Arwen Nicholson and Duncan Forgan

Scottish Universities Physics Alliance (SUPA), Institute for Astronomy, University of Edinburgh, Blackford Hill, Edinburgh EH9 3HJ, UK
e-mail: arwen.e.nicholson@gmail.com

Abstract: Interstellar probes can carry out slingshot manoeuvres around the stars they visit, gaining a boost in velocity by extracting energy from the star's motion around the Galactic Centre. These manoeuvres carry little to no extra energy cost, and in previous work it has been shown that a single Voyager-like probe exploring the Galaxy does so 100 times faster when carrying out these slingshots than when navigating purely by powered flight (Forgan *et al.* 2012). We expand on these results by repeating the experiment with self-replicating probes. The probes explore a box of stars representative of the local Solar neighbourhood, to investigate how self-replication affects exploration timescales when compared with a single non-replicating probe. We explore three different scenarios of probe behaviour: (i) standard powered flight to the nearest unvisited star (no slingshot techniques used), (ii) flight to the nearest unvisited star using slingshot techniques and (iii) flight to the next unvisited star that will give the maximum velocity boost under a slingshot trajectory. In all three scenarios, we find that as expected, using self-replicating probes greatly reduces the exploration time, by up to three orders of magnitude for scenarios (i) and (iii) and two orders of magnitude for (ii). The second case (i.e. nearest-star slingshots) remains the most time effective way to explore a population of stars. As the decision-making algorithms for the fleet are simple, unanticipated 'race conditions' among probes are set up, causing the exploration time of the final stars to become much longer than necessary. From the scaling of the probes' performance with star number, we conclude that a fleet of self-replicating probes can indeed explore the Galaxy in a sufficiently short time to warrant the existence of the Fermi Paradox.

Received 2 April 2013, accepted 24 May 2013, first published online 3 July 2013

Key words: Fermi Paradox, Galaxy, probes, Solar System.

Introduction

The Fermi Paradox is one of the oldest and most important problems in the search for extraterrestrial intelligence (SETI). The Paradox is defined as the absence of extraterrestrial intelligence (ETI) in our Galaxy despite calculations suggesting that galactic colonization should be feasible within the age of the Galaxy (see reviews by Brin 1983; Webb 2002; Cirkovic 2009).

The characteristic time for colonization of the Galaxy (either in person or through conventional or self-replicating probes) is known as the Fermi–Hart timescale, calculated by Hart (1975) to be:

$$t_{\text{FH}} = 10^6 - 10^8 \text{ years.} \quad (1)$$

As the accepted age of the Earth, as an object of roughly present-day mass, is $t_{\oplus} \sim 10^9$ years, the large discrepancy between these timescales leads us to the Fermi Paradox.

The first serious proposal to use probes to explore the Galaxy is credited to Bracewell (1960) as an alternative to interstellar radio communication between civilizations spread through the Galaxy, surmounting the potentially high-energy

cost of transmitting an easily receivable signal over very large distances (Benford *et al.* 2010a, b), and the problems associated with synchronicity of the civilization lifetimes (Forgan & Nichol 2010; Horvat *et al.* 2011). These probes would lurk in potentially habitable planetary systems, to wait for signs of intelligent life.

It is argued by Freitas (1983b) that probes may have visited our Solar System and that we have been too optimistic regarding our ability to detect the probes (Freitas 1983a). Freitas argues that objections to the existence of ETI based on the Fermi Paradox are invalid, as they are based on unsupported assumptions that ETI or their artefacts are not currently present in our Solar System. He argues that humanity's ignorance to potential evidence of ETI is not appreciated (but see Loeb & Turner 2012), and notes that any objects sent by aliens not intended to be found, will not be found. He proposes a scenario where a probe camouflages itself so as to set up a threshold test of the technology or intelligence of the recipient species, where the test must be met before the species is allowed to communicate with the device. Evidence in the form of 'spent' or destroyed probes is less likely, as any civilization attempting interstellar exploration are presumably

skilled engineers, and would send probes with the ability to self-repair, due to the large travel distances and times required for such a task, giving the probes a very long lifespan (Freitas 1983b).

Recently, computational power has increased sufficiently to allow the creation of detailed simulations of probes exploring a sector of the Milky Way. Bjørk (2007) investigated the timescale for a given number of space probes to explore 40 000 stars in a box, and then modelled 260 000 of these 40 000 stellar systems all located in the Galactic Habitable Zone (GHZ) (Lineweaver *et al.* 2004). In these simulations the speed of the probes was assumed to be $0.1c$. This speed was selected as it is low enough to be able to ignore effects of general relativity, but high enough that the time travel is of the order of years for travel between stars. Their results found that with eight probes with eight subprobes each, 4% of the Galaxy could be explored in 2.92×10^8 years. With 200 probes again with eight subprobes each the exploration time was reduced to 1.52×10^7 years. To put these numbers into perspective: it took 2% of the age of the Universe for eight probes with eight subprobes to explore $\approx 3.85\%$ of the Galaxy so exploring the entire Galaxy with this method would take a huge amount of time.

Cotta & Morales (2009) extended this analysis by attempting to optimize the trajectory of the probes, an instance of the NP-hard ‘travelling salesman problem’ (Golden & Assad 1988; Toth & Vigo 2001) in an attempt to determine the possible number of extraterrestrial technological civilizations (ETCs) that could be exploring the Galaxy while still giving a high probability that we would not have been visited yet. They found that it was unlikely for more than 10^2 – 10^3 ETCs would be exploring the Galaxy in a given Myr. This result assumes a constant probe velocity of $0.1c$, and a probe lifetime of 50 Myr, with contact evidence lasting 1 Myr. The number of ETCs drops considerably if the contact evidence is assumed to last for 100 Myrs. In this case, the upper-bound for the number of exploring ETCs goes down to about 10. The method for exploration was again using a number of probes with subprobes. They found that the exploration time was again very high; around 12 Myrs for four probes and 3 Myrs for eight probes to explore a quarter of the Galaxy. They conclude by saying that only with a large number of ETCs exploring the Galaxy would it become improbable for us to have no evidence of such an exploration. Cartin’s (2013) modelling shows that even in attempts to explore the Solar neighbourhood, with a fixed number of probes launched as a fleet from the Solar System, can be frustrated by failure of probe components. Further, these issues cannot be completely avoided by simply increasing the number of probes in the initial fleet.

This implies that for the Fermi Paradox to truly hold, a very large fleet of probes would be necessary – the most efficient means of generating such a fleet invokes the concept of a self-replicating probe (SRP). So far, no simulations of this type have been carried out using SRPs. Also known as Von Neumann probes, SRPs are spaceships capable of producing an exact copy of themselves under full automation. If probes require a great deal of self-repair, then it is not unlikely they may be able to achieve the next step of self-replication.

A civilization may send out one to a few SRPs, which will be programmed to choose the next star they travel to according to some decision-making algorithms. Once they reach the new star system, they scan for signs of life, and create a copy of themselves. The parent and child probe each pick a new star (not the same star) to travel to, and the process repeats itself.

Freitas (1980) outlines the specifics of an SRP, estimating the mass of each component of the spacecraft, and the material the spacecraft would need to accumulate to be able to produce a copy of itself and provide fuel. They suggest that a spacecraft may capture comets and use atmospheric mining to acquire these materials. They state that in their minds there is ‘little doubt that such a machine can, in theory, be designed’.

Wiley (2011) presents many of the arguments given against the use of SRPs, concluding that the arguments are insufficient for SRPs to be excluded from attempts to answer the Fermi Paradox. Sagan & Newman (1983) proposed that ETI would be unlikely to use SRPs due to the risk of mutations in the replication process resulting in dangerous unforeseen consequences. For example, a mutated probe may no longer possess the same goals as a ‘healthy’ probe, malfunctioning in a process analogous to cancer in biological systems, filling the Galaxy with malignant probes and potentially destroying that which it was designed to find.

A counter argument to this can be found in Tipler (1980). Biological evolution is a process with no design or foresight, whereas self-replicating machines will be the result of careful intelligent design. Evolution has developed many measures to protect against malignant mutations just as human technology has built in safeguards against corruption in electronic systems, so it is reasonable to assume that SRPs would be protected in a similar and possibly better fashion. While we feel that Wiley’s (2011) argument against dangerous mutation is overly stated in places (e.g., he compares the number of probe replications to the number of cell divisions in the human body, which we feel is an incorrect comparison, as the probes are likely to have a large number of cell-analogue components, as well as complex firmware/software), we agree that the concerns of mutations are insufficient to exclude all civilizations from building SRPs.

A related argument against SRPs includes the ‘predator and prey’ scenario (Chyba & Hand 2005) where mutated probes abandon their original mission of exploring the Galaxy and start to prey on normal probes, decreasing the number of explorer probes and greatly increasing the exploration time. Wiley’s counter argument makes the obvious point that although biology gives us many examples of how predators greatly influence the numbers of prey, the SRPs imagined do not behave as normal prey. The probes are continuously travelling at maximum speed and constantly dispersing radially. This means that as the predator probes would have the same maximum velocity as non mutated probes, they would have a hard time catching their prey. It seems implausible to ignore SRPs based on this argument, although Wiley (2011) notes that once a section of the Galaxy is explored, the lurking explorer probes leave themselves open for consumption by the predators, which may prevent some target species from being contacted.

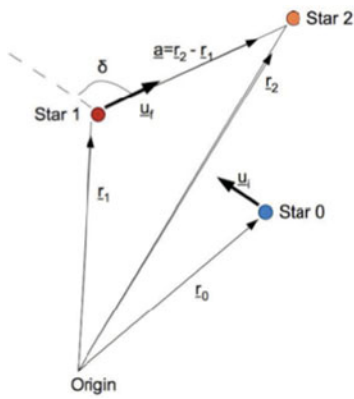


Fig. 1. The slingshot manoeuvre. The probe changes its direction by an angle δ while the magnitude of its velocity, in the frame of Star 1, u , remains constant. Diagram taken from Forgan *et al.* (2012).

Recent advances in 3D printing suggest self-replicating machines of sufficient complexity may be within our grasp in the coming centuries. With that in mind, when considering extra-terrestrial civilizations potentially a thousand times older than ours (especially given that terrestrial planets are likely to be 1 Gyr older than the Earth, Lineweaver 2001), it is not unlikely that they would have such technology.

In this paper, we present results from the first models to simulate the colonization of a sector of the Milky Way using SRPs with realistic probe dynamics. We build on previous Monte Carlo realization (MCR) simulations by Forgan *et al.* (2012), which considered a single non-replicating probe using orbital slingshot manoeuvres to provide velocity boosts instead of relying purely on powered flight, as has previously been the case. Using slingshots inside gravitational potential wells allows a probe to produce relatively large Δv and alter its trajectory without expending fuel, boosting its speed relative to the rest frame of its starting position (Surdin 1986; Gurzadyan 1996).

In the Method section, we describe the construction of the numerical model; in the Results section, we show the results of the simulations; in Discussion section, we note some issues and limitations of the model; and in Conclusions section, we summarize the work.

Method

The simulation domain is composed of stars uniformly distributed with a uniform density of 1 star per cubic parsec. To simulate the rotation curve of the Milky Way the stars were set in a shearing box configuration; however, for convenience the stars remain fixed in position even though they have velocity vectors. We created ten MCRs for three scenarios, investigating four different values for the total number of stars (see below), giving a total of 120 realizations. We average each set of ten realizations to characterize the effect of random fluctuations. The maximum velocity of the probes was chosen to be $3 \times 10^{-5}c$, where c is the speed of light in a vacuum),

which is approximately the maximum velocity obtained by unmanned terrestrial spacecraft such as Voyager 1¹.

Choosing the next star

We explore realizations of the same three scenarios as Forgan *et al.* (2012):

1. Powered flight to the nearest neighbour. Each probe will travel to its closest neighbour at its maximum powered velocity. Δv is fixed by the repeated deceleration and acceleration the probe makes at every star it visits (labelled *powered*).
2. Slingshot-assisted flight to the nearest neighbour. The path is identical to that for scenario 1, but the probe only accelerates to maximum velocity once and does not decelerate. It instead uses slingshot manoeuvres to repeatedly boost its maximum velocity. It is assumed that the Δv required of the probe to make course corrections in order to use a slingshot trajectory is negligible (labelled *slingshot*).
3. Slingshot-assisted probe selecting the next star by seeking the maximum velocity boost. This time the probe seeks the path such that the relative velocity between the current and destination stars is large and negative, meaning the destination star is moving towards the current star. This will result in a larger velocity boost, but in general will require a longer path length to achieve it (labelled *maxspeed*).

We apply these scenarios to self-replicating probes to explore how self-replication affects exploration time for each scenario.

Dynamics of slingshot trajectories

The slingshot trajectories used in these simulations are the same as those described in Forgan *et al.* (2012). A slingshot trajectory uses the momentum of the star it passes to either lose or gain velocity depending on the incident angle of the probe's approach. This means the probe does not need to use the additional energy that would be required to perform a similar trajectory using only powered flight, and may receive a velocity boost. Here, we briefly cover the mathematics of slingshot trajectories. For more detail the reader is referred to Gurzadyan (1996), in particular Chapter XIII, section 4.

The left side of Fig. 1 shows one stage of a probe's trajectory using slingshots. From its starting point at Star 0, the probe accelerates under its own power to velocity u_i (measured in the reference frame where Star 1 is at rest). Once it has arrived at Star 1, the probe achieves a velocity change Δu during the slingshot such that it leaves Star 1 with velocity u_f (again in Star 1's reference frame). δ is the angle between u_i and u_f . The probe follows a hyperbolic trajectory while performing this manoeuvre and the speed of the probe in this frame remains unchanged:

$$|u_i| = |u_f|. \quad (2)$$

Owing to the huge mass difference between the probe and the star, the transfer of momentum from the star to the probe is negligible. In the simulation, we observe the probe in the frame

¹ <http://voyager.jpl.nasa.gov/mission/weekly-reports/index.htm>

where the Galactic Centre is at rest. The initial and final velocities in this frame are v_i and v_f , respectively (right side of Fig. 1). The magnitude of the two vectors in the Galactic frame are no longer equal, and

$$\Delta v = 2|u_i| \sin\left(\frac{\delta}{2}\right). \quad (3)$$

The probe's initial velocity between Star 0 and Star 1 is the maximum velocity and it can achieve under powered flight, but as it travels to more stars using slingshot manoeuvres it can increase its speed. The magnitude of the boost attained by a slingshot manoeuvre is greater if the star's velocity is parallel to that of the probe's trajectory. Thus, it is possible for a probe to choose a course based on the proper motion of stars relative to one another such that it performs slingshots with maximal Δv .

Probe replication

Our model for probe replication is a very simple one. While a probe is travelling between stars, we assume it collects matter from the interstellar medium, and this is used to create a replica probe. We assume that the quantity of material collected during the flight is great enough such that the probe does not have to stop and mine for materials at any time, and thus the process of replication does not affect the journey time between any two stars.

Each new probe is an exact copy of its parent probe, and no mistakes in the building or programming of new probes are ever made. Thus, each new probe behaves in exactly the same way as the original probe; they select the next star to visit in the same way as their parent, and use (or do not use) slingshot trajectories in the same way.

We assume that a probe releases its replica on arrival at the destination star and the probe does not slow down or change its motion in any way while releasing the replica probe. This allows us to use the same equations as in the single probe case, as the release of a new probe makes no change to the parent probe's motion in any way. We also assume that this replica probe is released with the same velocity and momentum as its parent probe.

This means, in the powered case, the parent probe slows as always on approach of the new star and then releases the replica with the same velocity. Both probes then choose a new destination star and accelerate away towards their chosen destinations.

In the slingshot case, the parent probe reaches the new destination star, and before it slingshots around the star it releases the replica probe. Both the parent and replica use the slingshot to boost their velocity. As the velocity boost from a slingshot trajectory depends on the angle between the stars, the parent and the replica will achieve different velocity boosts as they will have different destination stars.

Information transfer

In a situation with more than one probe exploring a Galaxy there needs to be a way for the probes to communicate which stars have already been visited. If the probes do not

communicate then we can easily end up in a situation where two or more probes land on the same star, and as they all have the exact same algorithm for selecting their next destination star, they may end up following the same path for the rest of the exploration. This could result in a fleet of probes following the same path around the Galaxy, which will not only increase the total exploration time, but could also cause chaos if said fleet ended up in a star system with life. One probe visiting may be of enormous consequence to civilization (Shostak & Almar 2002; Almár 2011), but if thousands of probes all descend at once or within a short time of each other it could be interpreted as hostile, causing panic (cf. Cantril 1940).

We prevent this from happening in our simulations by giving each probe perfect information at all points in time. That is, all probes know exactly where every other probe has been and is going to. When a probe selects its next destination star, it will select an unvisited star that no other probe is currently travelling to. Once a probe selects a destination star, all other probes know not to visit that star.

This situation is obviously unrealistic. While the model might be accurate at the start of the journey when the probes are relatively close together, as they get further and further apart, the information transfer between them should take longer. The rate of transfer of information is limited to the speed of light, so with greater distances there is a greater time delay between a probe making a decision, and the rest of the probes being made aware of this decision. However, we do not take account of this problem and just use the simple case of all probes knowing everyone's actions at all times. In later sections, we will describe some problems that arise using this assumption.

Results

Comparing single probes to self-replicating probes

In general, using self-replicating probes significantly reduces exploration time for all three scenarios when compared with a single non-replicating probe. The top plot in Fig. 2 shows the average cumulative travel time for single non-replicating probes exploring a box of 200 000 stars in the three scenarios, and the bottom plot in Fig. 2 shows the same quantity for the self-replicating probes. It is immediately clear that adding self-replication – simply by virtue of dividing the task of exploration between many agents – significantly reduces the exploration timescale. For the proper motion and powered scenarios, the exploration time is reduced by a factor of 1000, from $\sim 10^{10}$ years to $\sim 5 \times 10^6 - 10^7$ years, whereas the exploration time for the slingshot scenario is reduced by a factor of around 100, from 10^8 years to 10^6 years. Note that the ordering of the three scenarios remains the same – slingshot is the most effective algorithm, with powered now a close second to the maxspeed scenario. As the slingshot scenario is already quite efficient, the factor by which adding self-replication reduces exploration timescales is quite low, but in general the effect of self-replication is much more important than the effect of probe dynamics. We discuss the dependence of this efficiency increase on star number in the following section.

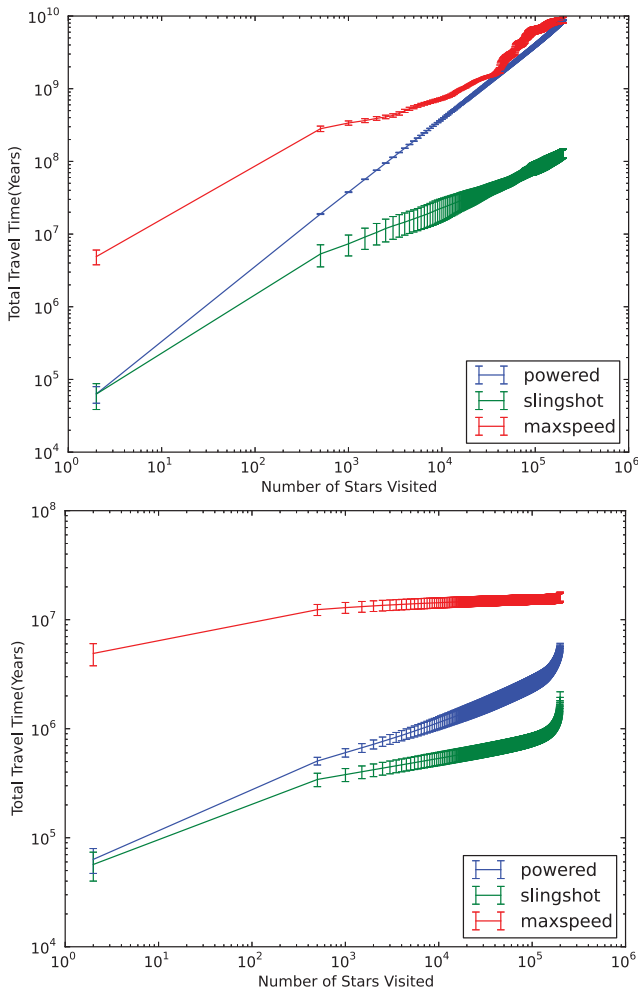


Fig. 2. The mean cumulative travel time for the three scenarios powered, slingshot and maxspeed as described previously, in a box containing 200 000 stars. The top plot shows the travel times for a single, non-replicating probe, and the bottom plot shows the travel time in the case of self-replicating probes. Every 500th data point was plotted to prevent overcrowding of the graph. The thickness of the lines indicates sample standard deviations as error estimates. The maxspeed case is represented in red, powered in blue and slingshot in green.

Note that the slingshot and powered scenarios in the self-replicating case show a significant upturn in travel time as the number of stars explored approaches 100%. A similar, but weaker feature is present in the single probe case for powered flight (see Fig. 2 of Forgan *et al.* 2012), which occurs when there are few unvisited stars left in the domain, with generally larger spatial separation between them than the mean stellar separation. We might naively think that this problem should no longer exist for the case of self-replicating probes. In this late stage of exploration, with many hundreds of thousands of probes, each unvisited star should be surrounded by a large number of potential visitors. The fact that this phenomenon persists is indicative of a deeper problem of modelling inter-probe communication (see the Discussion).

We can assess this phenomenon by investigating the time taken to explore 98 and 99% of the domain for the three

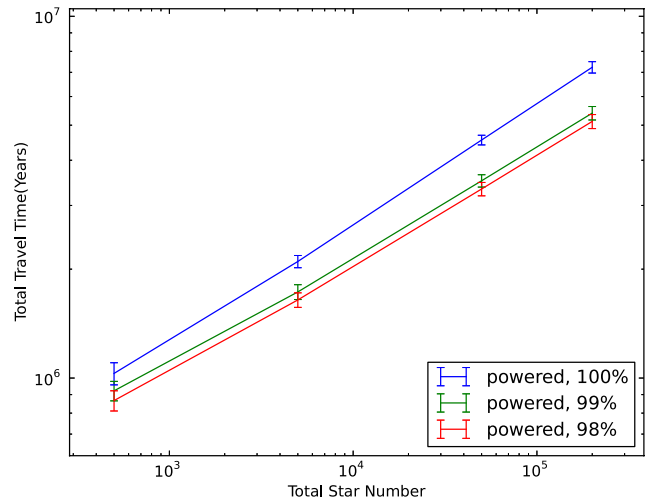


Fig. 3. Comparison of the exploration time for self-replicating probes to explore 100, 99 and 98% of the stars in the powered case.

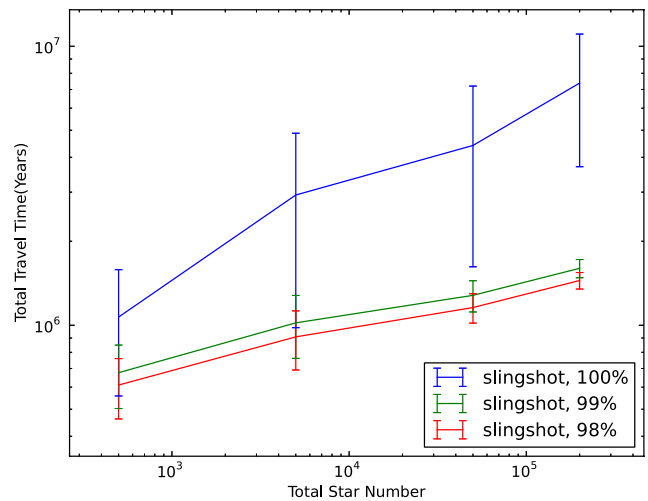


Fig. 4. Comparison of the exploration time for self-replicating probes to explore 100, 99 and 98% of the stars in the slingshot case.

scenarios as a function of total star number (Figs. 3–5). In the powered case (Fig. 3), there is a clear power-law relationship between the number of stars in the domain and the total travel time. The time taken to explore 98% of the system is the same as the time taken to explore 99% of the system (to within the sample standard deviation). The time taken to explore the final 1% of the domain increases the final travel time by around 40%!

In the slingshot case (Fig. 4), we can see that the travel time increases more slowly with extra star number – the velocity boosts due to the slingshots allow the probes to explore at greater speed than in the powered case. Again, the 98 and 99% curves are tightly correlated with overlapping errors, but the 100% case takes nearly three times longer to explore the domain for 200 000 stars, with much larger errors than the powered case due to the larger range of speeds that the probe population can adopt.

The maxspeed case (Fig. 5) shows more complicated behaviour. In low star number runs, the time penalty for

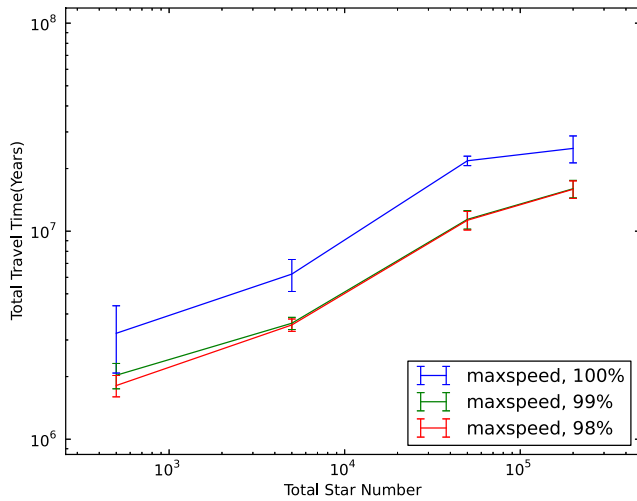


Fig. 5. Comparison of the exploration time for self-replicating probes to explore 100, 99 and 98% of the stars in the maxspeed case.

travelling large distances to achieve a high velocity boost is weaker, as the box itself is smaller. As the box size increases, the penalty for finding the maximum boost increases, and begins to level off once the star number reaches 100000, where the time penalty is approximately balanced by achieving high velocity quickly. Once more, we can see that the last 1% of the stellar population requires half the total travel time to explore.

Discussion

Race conditions, information and the final 1% of the stars

In assuming the probes are perfectly informed of the swarm’s exploration of the box, we unwittingly created a problem for the final 1% of the unexplored box. When probes select their next target, this is done on a first-come first-served basis. The first probe to lay claim to a star will visit it, and all other probes will know immediately not to also choose that star. When there are plenty of stars to choose from, this does not cause any visible issues. This is only a problem if there are few stars to choose from, where we can end up with the scenario as depicted in Fig. 6. In this simplified situation, there are three stars and two probes. Each probe is assigned a star: the probe to reach their star first will go on to visit the final star. However, that probe is not necessarily the closest to the final star, yet as long as it reaches its destination star first, the second probe will not consider a visit to the last star. In this case, we have artificially extended the exploration time of the probes by prohibiting the second probe from going onto the last star, even though it is much closer and would get there in less time. This is an instance of a *race condition* or *race hazard*, which occurs in electronic systems where incorrect sequencing of events can result in unwanted delays (Karam & Buhr 1990).

This underlines a logistical issue in stellar exploration. We have allowed probes a great deal of autonomy in their exploration; the fleet is not broken up into groups with well-defined areas of exploration, as is the case with the

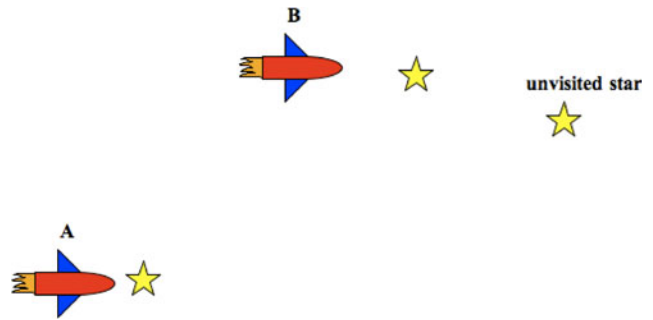


Fig. 6. An example of a race condition in the simulation. Although probe B is much closer to the unvisited star, as probe A will reach its destination first, it will tell all other probes it is going on to the unvisited star, thus forcing B to chose another more distant star (or causing B to stop if we are nearing the end of the exploration) which can have the effect of increasing the overall exploration time.

non-replicating studies of Bjørk (2007) and Cotta & Morales (2009). There is no chain of command in the fleet, the fleet continues to grow, and the probes will act according to simple programming. Also, the assumption of perfect information is not well motivated. For this assumption to be valid, the travel time for communication between probes must be significantly lower than the travel time for probes to move towards each other. In the simulation, the travel time for communication is zero; in reality, communication between the fleet will proceed at the speed of light, c . As the probes achieve speeds of around $0.01 c$ in the course of exploring the box, and the stellar density is sufficiently high, perfect information is not strictly possible.

Other limitations of the model

We share some of the limitations of the model described in Forgan *et al.* (2012), e.g. not evolving the positions of the stars despite their possessing non-zero velocity vectors. As velocity vectors will change during the motion of the stars in a gravitational potential, the ability of the maxspeed algorithm to select a star that will have a favourable velocity vector when a probe arrives at it depends significantly on the travel time – stars are much less likely to maintain a similar velocity vector for increasing lengths of time. We again ignore relativistic effects in this analysis, as they are not significant – for speeds of $0.01c$, the Lorentz factor $\gamma = 1.00005$, showing that classical physics is an acceptable approximation. This limiting speed of $0.01c$ is a consequence of the maximum Δv being linked to the maximum value of the rotation angle δ that the velocity vector moves through during the slingshot:

$$\Delta v_{\max} = \frac{u_{\text{esc}}^2}{(u_{\text{esc}}^2/2u_i) + u_i}, \tag{4}$$

where

$$u_{\text{esc}} = \sqrt{\frac{2GM_*}{R_*}} \tag{5}$$

and we assume solar values for the star mass M_* and radius R_* . Compact objects such as neutron stars and black holes may

provide slightly better Δv_{\max} (Dyson 1963), although it is unclear how such a craft may deal with the extremely strong tidal forces upon its hull. Besides, any analysis of these types of slingshots must of course consider general relativistic effects, which we do not.

Are advanced orbital dynamics necessary?

We have seen that the act of adding self-replication to a probe reduces the difference in total travel time between the powered and slingshot cases significantly (Fig. 2). It is therefore important to ask, 'is it worthwhile bothering to program probes to carry out complicated orbital manoeuvres at all?'

We would argue the following points: firstly, while there is now only a factor of a few differences in travel time, this is still significant if the exploration is expanded to the entire Galaxy. We have limited our study to 200 000 stars – increasing that number by a factor of a million (i.e. 10^{11} stars) would mean that a factor of a few becomes much more important. Also, if we consider the trend of total exploration time with star number (Figs. 3 and 4), then (if we take 99% of the stars to be an acceptable goal), we can see that the slingshot algorithm scales much better with star number than the powered algorithm, and it seems the difference in travel time will continue to increase as the total number of stars increases. Assuming power laws for how both cases scale with star number, a rough extrapolation suggests that a fleet using the powered algorithm would take around 10^8 years to explore the Galaxy, compared with around 10^7 for a fleet using the slingshot algorithm.

Even if this is not the case, and there is not a strong difference in travel time on Galactic scales, there is a strong economic argument for the slingshot algorithm. The massive savings in fuel across the entire fleet by avoiding using powered flight cannot be underestimated (although it is possible that alternative renewable fuel sources such as solar energy may be used). Being able to achieve a higher velocity may help the process of collecting raw materials from the interstellar medium during transit to the next star. While the total amount of material collected depends only on the column density of material and path length, the ability to compress the material more efficiently, to conserve space, would benefit from the higher velocities that slingshot trajectories can offer.

Improvements to the model

Future work should include resolution of the race condition problem described above. This could be achieved in a number of ways, the simplest of which could be to include a travel time limit on the probes. If the nearest unvisited star to them would take longer to travel to than this preprogrammed limit, the probe will simply cease its exploration and will stay on its current star and not produce a replica. This could have the effect of reducing the overall travel time, by preventing probes choosing a destination star at a great distance to themselves, instead of leaving the stars for other probes that are nearby. Another possible way would be for the probes to monitor the position of all other probes and to implement a priority system where the probe closest to the unvisited star is assigned.

A related improvement to these simulations would be to make the probes' information transfer realistic. When a star gets marked as visited, there would be a delay in this information reaching the probes, proportional to the distance of the star from members of the probe fleet. This could be simulated by assuming each probe leaves a beacon that emits a signal showing it has been visited, where the signal would propagate throughout space at the speed of light. Probes would then update their own local catalogues of visited stars when they come within the range of the signal. Alternatively, probes could leave behind their travel history on all visited stars in some form of data box. This box would also contain information on where the probe is next travelling to. Any probe that then later visits an already visited star would sync with the data box, with both entities merging their catalogue data as to what stars have been visited. This way the probes exchange information when they encounter already visited stars. This could help reduce travel time, as probes will avoid stars that have already been visited by other probes with which they have crossed paths.

Conclusions

We have presented results from MCR simulations of a fleet of self-replicating probes exploring a region of the Galaxy, investigating the effect of slingshot orbital dynamics on the total exploration timescale. These are the first numerical simulations that explicitly model self-replicating probes undergoing realistic orbital trajectories in a stellar population reminiscent of the Milky Way. We build on the single probe results of Forgan *et al.* (2012), and we explore the same three scenarios for probe flight: moving to the nearest neighbour under powered flight, moving to the nearest neighbour using slingshot trajectories, and moving to the star that provides the greatest velocity boost via slingshots.

We find that nearest-neighbour, slingshot-assisted flight is the most time effective exploration method for self-replicating probes (as is the case with a single non-replicating probe). Owing to our assumption of perfect information exchange amongst the fleet, the last 1% of stars take much longer to explore due to 'race conditions' being set up within the probe network, which require more sophisticated decision-making algorithms to avoid. Despite this, our results confirm that a fleet of self-replicating probes can explore the Galaxy in a timescale commensurate with those normally assumed when posing the Fermi Paradox (Hart 1975), with powered flight at the upper limits of the timescale and slingshot flights at the lower end. Both are still orders of magnitude less than the age of the Earth, proving that the question underlying the Fermi Paradox is well posed.

References

- Almár, I. (2011). *Acta Astronaut.* (UK) **69**, 899.
- Benford, G., Benford, J. & Benford, D. (2010a). *Astrobiology* **10**, 491.
- Benford, J., Benford, G. & Benford, D. (2010b). *Astrobiology* **10**, 475.
- Björk, R. (2007). *Int. J. Astrobiol.* **6**, 89.

- Bracewell, R.N. (1960). *Nature* **186**, 670.
- Brin, G.D. (1983). *QJRAS* **24**, 283.
- Cantril, H. (1940). *The Invasion from Mars: A Study in the Psychology of Panic*. Transaction Publishers, New Brunswick, USA and London, UK.
- Cartin, D. (2013). *Int. J. Astrobiol.*, in press.
- Chyba, C.F. & Hand, K.P. (2005). *ARA&A* **43**, 31.
- Cirkovic, M.M. (2009). *Serb. Astron. J.* **178**, 1.
- Cotta, C. & Morales, A. (2009). *J. Br. Interplanet. Soc.* **62**, 82.
- Dyson, F.J. (1963). Gravitational machines, In *Interstellar Communication*, ed. Cameron, A.G.W., pp. 115–120. Benjamin Press, New York.
- Forgan, D. & Nichol, R. (2010). *Int. J. Astrobiol.* **10**, 77.
- Forgan, D.H., Papadogiannakis, S. & Kitching, T. (2012). *J. Br. Interplanet. Soc.*, in press.
- Freitas, R.A. (1980). *J. Br. Interplanet. Soc.* **33**, 251.
- Freitas, R.A. (1983a). *Br. Interplanet. Soc.* **36**, 496.
- Freitas, R.A. (1983b). *Br. Interplanet. Soc.* **36**, 501.
- Golden, B. & Assad, A. (eds) (1988). *Vehicle Routing: Methods and Studies, vol.16 of Studies in Management Science and Systems*. North-Holland, Amsterdam.
- Gurzadyan, G.A. (1996). *Theory of Interplanetary Flights*. Gordon and Breach, Amsterdam.
- Hart, M.H. (1975). *QJRAS* **16**, 128.
- Horvat, M., Nakić, A. & Otočan, I. (2011). *Int. J. Astrobiol.* **11**, 51.
- Karam, G. & Buhr, R. (1990). *IEEE Trans. Softw. Eng.* **16**, 829.
- Lineweaver, C. (2001). *Icarus* **151**, 307.
- Lineweaver, C.H., Fenner, Y. & Gibson, B.K. (2004). *Science* **303**, 59.
- Loeb, A. & Turner, E.L. (2012). *Astrobiology* **12**, 290.
- Sagan, C. & Newman, W.I. (1983). *QJRAS* **24**, 113.
- Shostak, S. & Almar, I. (2002). in 34th COSPAR Scientific Assembly, Second World Space Congress, pp. IAA-9-1-06.
- Surdin, V.G. (1986). *Astron. Vestn. (Russia)* **19**, 354.
- Tipler, F.J. (1980). *QJRAS* **21**, 267.
- Toth, P. & Vigo, D. (eds) (2001). *The Vehicle Routing Problem, Vol. 9*. Society for Industrial and Applied Mathematics, Philadelphia, PN.
- Webb, S. (2002). *If the Universe is Teeming with Aliens— where is Everybody? Fifty Solutions to the Fermi Paradox and the Problem of Extraterrestrial Life*, Springer, New York, p. 288, please see: <http://adsabs.harvard.edu/abs/2002iuta.book.....W>.
- Wiley, K.B. (2011). arXiv e-prints 1111.6131.