

Orbital Optimization of a Distributed Heliocentric Relay Network for Mars-Earth Communications

Jules Pénot
 Massachusetts Institute of Technology
 77 Massachusetts Ave.
 Cambridge, MA 02139
 jpenot@mit.edu

Hamsa Balakrishnan
 Massachusetts Institute of Technology
 77 Massachusetts Ave.
 Cambridge, MA 02139
 hamsa@mit.edu

Abstract—Deep-space missions have long depended on Direct-to-Earth communications via NASA’s Deep Space Network, but the network’s limited capacity and the growing data downlink demands of modern missions pose major challenges for future robotic and crewed space exploration. To supplement the Deep Space Network, we propose a Heliocentric Relay Network (HRN): a distributed constellation of Ka-band relay satellites in solar orbit, designed for continuous high-data-rate Mars–Earth communications. We develop a modular hidden-genes genetic algorithm to determine HRN satellite orbits that maximize the average Mars–Earth downlink capacity over a 6.4-year simulation period. We show that with 25–100 Starship launches, average Mars–Earth data rates through the HRN range from 146 Mbps to 852 Mbps, more than 100 times the throughput of current Direct-to-Earth systems. These results demonstrate that a large-scale HRN using technologically-mature communication systems could enable high-data-rate, continuous Mars–Earth communications, supporting future exploration across the inner Solar System.

incurred while communicating directly across interplanetary distances also limits the achievable throughput of DTE communications. Meanwhile, space missions are increasingly demanding in data uplink and downlink volume: in response, NASA has a stated objective of attaining a 200 Mbps average downlink data rate from Mars to Earth, which corresponds to a 100-fold increase over existing capabilities [3], [4]. NASA has also put out a request for commercial proposals for a “Mars End-to-End Communication Service Architecture” capable of continuous, duplex, 100 Mbps Earth-Mars communications [5].

While new technologies such as optical communication systems promise significantly higher data rates, they do not inherently address other challenges in interplanetary communications such as interruptions during solar conjunctions [4]. Such up-time and reliability issues reduce the scientific productivity of autonomous missions through the Solar System, and pose a significant risk to the safety of astronauts during crewed missions beyond Earth orbit. Communication outages due to conjunctions and eclipses can be circumvented by using a single relay spacecraft in heliocentric orbit, but this approach does not increase end-to-end throughput compared to DTE communications [6].

TABLE OF CONTENTS

1. INTRODUCTION.....	1
2. SYSTEMS ENGINEERING AND ASSUMPTIONS....	2
3. HELIOCENTRIC SATELLITE DEPLOYMENT.....	3
4. NETWORK PERFORMANCE ANALYSIS.....	5
5. MODULAR GENETIC ALGORITHM.....	8
6. RESULTS AND DISCUSSION.....	11
7. FUTURE WORK.....	16
8. CONCLUSIONS.....	16
ACKNOWLEDGEMENTS.....	16
REFERENCES.....	16
BIOGRAPHY.....	18

1. INTRODUCTION

From the Voyager probes to the Perseverance rover, nearly all space exploration missions beyond cislunar space in the last 50 years have relied on Direct-to-Earth (DTE) communications for receiving commands, transmitting telemetry, and down-linking scientific data. In the DTE paradigm, spacecraft communicate directly (or occasionally through a nearby orbiter) with an Earth-based network of antennas, generally NASA JPL’s Deep Space Network (DSN). However, the DSN has now become a critical bottleneck for deep-space communications. With a limited number of antennas, the DSN is oversubscribed, with missions from both NASA and partner agencies having to negotiate for time using the network’s Earth-based antennas. Furthermore, the upcoming Artemis missions are expected to further strain the DSN’s resources in coming years [1], [2]. The large free space path loss

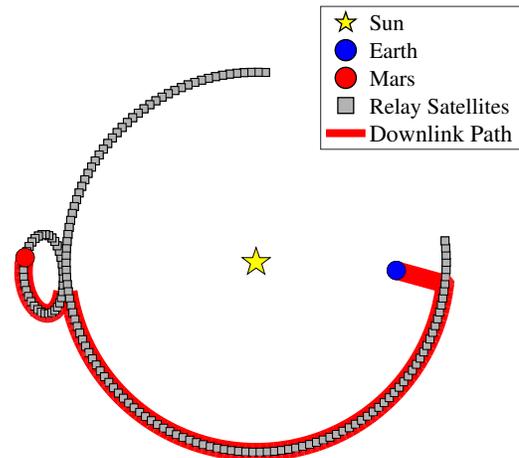


Figure 1: An Earth-Mars Heliocentric Relay Network with 165 relay satellites, deployed with 33 Starship vehicles.

As an alternative to the DTE communications paradigm, we propose the creation of a hyper-distributed network of radio-frequency (RF) relay satellites in Solar orbit for communications between Earth and Mars. In this work, we shall refer to this network as a **Heliocentric Relay Network (HRN)**. The HRN consists of a backbone of Ka-band relay satellites in heliocentric orbits, optimized for network throughput

and guaranteeing continuous communications during Mars-Sun-Earth conjunctions. Such a multi-hop, distributed interplanetary network presents several advantages over DTE RF communications, including reduced free-space path loss over communication links, built-in redundancy, lower size, weight and power (SWaP) of future Martian communications systems, and freeing DSN capacity for missions elsewhere in the Solar System [7].

This paper presents three main areas of investigation. First, we develop an approach for computing the number of relay satellites deliverable to a particular formation with a single launch based on parameters such as relay antenna diameter, the launch vehicle’s payload mass and volume constraints, and the satellites’ orbital insertion ΔV s. We then develop an efficient algorithm for evaluating the performance of a given satellite constellation, relying on a warm-started binary search algorithm to find the maximum-capacity path between Earth and Mars at each simulation timestep. To evaluate the performance of a network, we propagate the orbits of Mars, Earth, and all the satellites, over a 6.4-year simulation period, and take the average data rate from Mars to Earth over that simulation period as our measure of performance. Finally, we implement a modular genetic algorithm to find the optimal Mars–Earth relay network subject to a fixed number of satellite launches. We identify six categories of modular satellite formations (e.g., elliptical rings of satellites and planar Mutually Orbiting Groups (MOGs)) that maintain consistent inter-satellite spacing, and encode them into the design space under a unified genome. We assume the continuous use of one DSN 34 m antenna on Earth, the presence of relay satellites in Martian orbit, and the use of SpaceX’s Starship vehicle for launching and deploying the relay satellites.

The constellations identified have an average Mars–Earth data rate that increases substantially with the number of launches. Our results show that it is possible to achieve Mars–Earth data rates of over 146 Mbps with 25 Starship launches, 294 Mbps with 50 launches, 568 Mbps with 75 launches, and 852 Mbps with 100 launches. These data rates represent an over 100-fold increase in throughput compared to the current DTE system between the DSN and the Mars Reconnaissance Orbiter. The identified HRNs require substantial resources in the form of launch vehicles and spacecraft, but are able to surpass NASA’s objectives for Mars–Earth downlink throughput while using only technologically-mature Ka-band communications systems. In future work, these HRNs could integrate different communication technologies, and be expanded to provide simultaneous communication services to several points in the inner Solar System.

2. SYSTEMS ENGINEERING AND ASSUMPTIONS

Launch Vehicle and Satellite Assumptions

To quantitatively assess the performance of a given Mars–Earth relay network, we assume the use of a specific launch vehicle to deploy relay satellites into heliocentric orbit. We select SpaceX’s Starship Block III upper stage paired with the Super Heavy booster, due to its large payload capacity to Low Earth Orbit (LEO), large payload fairing volume, and increasing adoption in mission concepts and proposals [8]. We also used Starship as a representative launch vehicle in our previous work on the fuel-optimal deployment of satellites into planar heliocentric Mutually Orbiting Groups (MOGs) [9].

Table 1 summarizes key performance metrics for three iterations of Starship, compiled by Ars Technica from both SpaceX’s website and public announcements made by SpaceX officials [10], [11].

Table 1: Starship Performance Metrics [10], [11]

Starship Metrics	Block I	Block II	Block III
Payload to LEO (t)	N/A	100	200
Ship Propellant Load (t)	1200	1500	2300
Ship Dry Mass (t)	100	116	160

Additionally, for all of Starship’s propulsive maneuvers outside of the Earth’s atmosphere we use the specific impulse (Isp) of Starship’s Raptor vacuum thruster: 382 s in vacuum [12]. For the satellites’ onboard propulsion systems, we choose a specific impulse of 300 s, which corresponds to a typical specific impulse of the main engines of interplanetary spacecraft, such as Galileo’s 400 N monomethylhydrazine and nitrogen tetroxide main engine [13].

For this work, we assume that each relay satellite has two gimballed parabolic antennas, 8 m in diameter. We choose the largest antenna diameter that can fit inside Starship’s payload fairing, which has a cylindrical usable internal volume of 8 m in diameter and an extended payload bay height of up to 22 m, of which approximately 17.5 m are at the full 8 m diameter [14]. We assume an f/D ratio of 0.5, and use the focal length formula for a parabolic antenna, provided in Equation (1), to find the depth of the antenna body, $c = D/(16 \cdot f/D) = 1$ m. See Figure 2 for the naming conventions of the parabolic antenna dimensions used in this work. We use the calculated antenna depth c as an estimate of the stowed height of the antenna itself when stacked into the launch vehicle payload fairing, under the assumption that the receiver at the focal point of the parabolic antenna can be stowed for launch [15].

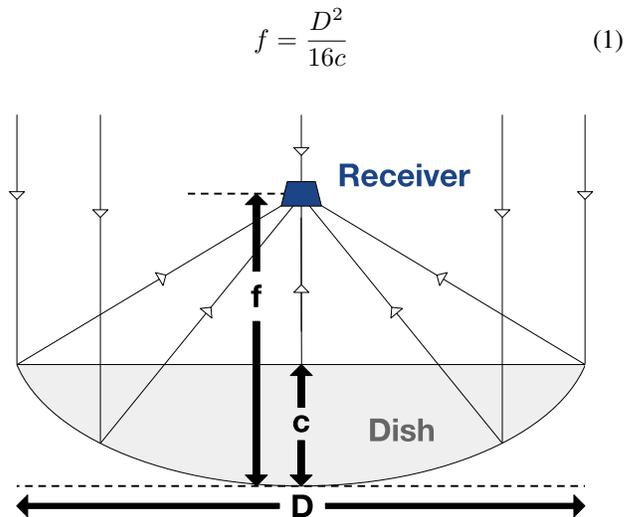


Figure 2: Diagram representing the major components and dimensions of a parabolic antenna. The receiver is located at the focal point of the parabolic dish.

We assume that each satellite’s payload bus has dimensions of approximately 2x2x1 m, similar to those of existing commercial communication satellites such as Starlink. This puts the

total height of one HRN satellite at under 3 m when stowed in the launch vehicle payload bay. Under these assumptions, Starship’s payload bay can fit 5 vertically-stacked HRN satellites.

We estimate the dry mass of the satellite bus to be 1000 kg excluding antennas, similar to the specifications of the Psyche spacecraft [16]. We assume a uniform areal mass density for the antenna body of 5 kg/m², which is that of the Mars Reconnaissance Orbiter (MRO) High Gain Antenna (HGA)—including the motorized gimbals and support structure—with a margin [17]. We then obtain the mass of each parabolic antenna body in Equation (2).

$$\begin{aligned}
 M_{\text{antenna}} &= \rho_a A_{\text{antenna}} = \rho_a \cdot \frac{\pi R}{6c^2} \cdot ((R^2 + 4c^2)^{\frac{3}{2}} - R^3) \\
 &= \rho_a \cdot \frac{64\pi}{3} D^2 (f/D)^2 \left(\frac{1}{4} + \frac{1}{64(f/D)^2} \right)^{\frac{3}{2}} - \frac{1}{8} \\
 &\approx 266.44 \text{ kg}
 \end{aligned}
 \tag{2}$$

The antenna mass estimate detailed in Equation (2) makes the total estimated dry mass per HRN satellite 1,533 kg. In our optimization algorithm, when we compute the achievable dry mass per satellite for the desired final formation (elliptical ring, circular ring, MOG, etc..) we compare the achievable dry mass to the nominal satellite dry mass. We then determine if the HRN satellite deployment to that formation is volume-constrained (we can launch all 5 satellites that can fit in the payload bay) or mass-constrained, in which case the achievable number of HRN satellites per launch is in the 1-4 range. If the total dry mass deployable in a Starship launch is less than the nominal dry mass of a single HRN satellite, then we conclude that the deployment procedure in question is incapable of deploying any satellites to the formation under analysis.

Planetary Network Assumptions

This work focuses on the orbital optimization of the HRN itself, and not on the design, operation, or tasking of the communications infrastructure at the Earth or Mars ends of the interplanetary network. We therefore make several assumptions about the Earth and Mars interfaces of the network to compromise between fidelity and simplicity of analysis.

For the Earth interface of the HRN, we assume the use of one DSN 34 m-diameter Beam Waveguide (BWG) antenna for the entire duration of the Mars–Earth distributed network’s operation. Monopolizing any of the DSN’s 70 m antennas for long durations is less realistic since those larger antennas are fewer in number, are particularly oversubscribed, and prioritize communications with outer Solar System missions such as New Horizons, the Voyager probes, and Juno [1], [18].

The particular 34 m antenna used at a given time may belong to any of the three DSN complexes (Goldstone, Madrid, or Canberra) so long as the HRN satellite down-linking data to Earth at that moment is in the visibility cone of the antenna in question. Since the visibility cones of the DSN sites cover the entire sky past GEO, and the Earth’s radius of 6,378 km is small compared to the distances between HRN satellites, we model the Earth’s interface with the broader HRN as a single *node* at the Earth’s position that has the link budget parameters of a DSN 34 m BWG antenna.

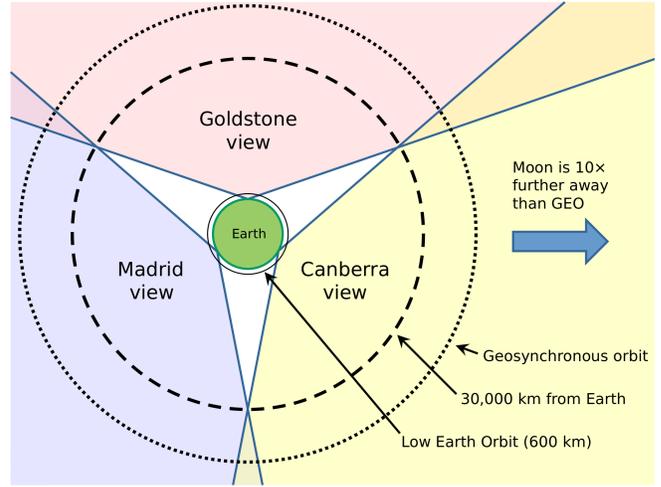


Figure 3: Field of view of the Deep Space Network antennas, looking down from the North Pole. Spacecraft more than 30,000 km from Earth are always in view of at least one station, and hence at least three 34 m BWG antennas [19].

In this work, we also assume that there already exists a constellation of relay satellites in Mars orbit, with the following characteristics:

- Each Mars relay satellite has identical communications subsystems to the HRN satellites, including 8 m Ka-band parabolic antennas.
- There is a sufficient number of Mars relay satellites to guarantee that at least one is visible to any HRN satellite at any time.
- The Mars network is responsible for relaying data and telemetry to and from other spacecraft, landers, and rovers in orbit and on the surface of Mars.

Similarly to our treatment of the Earth–HRN interface, we thus model the Mars–HRN interface as a single node that can always communicate with its nearest neighbors in the broader relay network. When evaluating the performance of a candidate HRN, Mars is therefore treated similarly to any other HRN satellite, except that the Mars node follows the exact ephemerides of Mars’ orbit. Future work may focus on the design, orbital optimization, and detailed deployment, of a Mars-orbiting communications constellation for interfacing with a heliocentric relay network, but such high-fidelity modeling is outside the scope of this paper.

3. HELIOCENTRIC SATELLITE DEPLOYMENT

To quantitatively estimate the number of HRN satellites that can be deployed with a single launch, it is necessary to compute the ΔV required for the orbital insertion of the HRN satellites. We develop several deployment procedures to deploy HRN satellites into the three main types of formations incorporated into our modular HRN: Mutually Orbiting Groups (MOGs), circular rings of satellites, and elliptical rings of satellites.

For a given satellite formation, we then select the most fuel-efficient deployment procedure, and obtain the corresponding launch vehicle and satellite ΔV s. We compute the total HRN satellite dry mass deliverable using the selected procedure’s ΔV s and the Tsiolkovsky rocket equation, presented in Equa-

tion (3). The final delivered satellite dry mass dictates the number of satellites deliverable to the MOG per launch N_{pl} . The total number of satellites delivered to the MOG with N_1 launches is thus $N_{\text{sats}} = N_{\text{pl}} \cdot N_1$.

$$\Delta V = I_{\text{sp}} \cdot g_0 \cdot \ln\left(\frac{m_0}{m_f}\right) \quad (3)$$

Mutually Orbiting Group Deployment

We have previously published work on the fuel-optimal deployment of planar, heliocentric, MOGs [9]. We reuse that work here for estimating the deployment ΔV of MOGs in the context of our broader HRN optimization framework.

To deploy MOGs, we consider two distinct MOG deployment strategies developed in this previous work [9]:

1. The Hohmann Insertion strategy: the launch vehicle enters a circular heliocentric deployment orbit of radius a_{MOG} (the semi-major axis of the MOG satellite orbits). The satellites then one-by-one execute a Hohmann transfer from the deployment orbit to their final elliptical MOG orbit.
2. The Spread-out Insertion strategy: the launch vehicle directly enters the MOG formation (an elliptical deployment orbit). MOG satellites execute MOG phasing maneuvers to insert into their final MOG orbits.

Circular Ring Deployment

We consider two different satellite insertion strategies for the deployment of circular, heliocentric rings of satellites. The first involves an offset circular deployment orbit, which we will refer to as the ‘‘Offset Insertion’’ (OI) strategy. The second involves the launch vehicle inserting directly into the circular ring and the HRN satellites executing phasing maneuvers within the ring. We will refer to this second deployment paradigm as the ‘‘Phasing Insertion’’ (PI) deployment strategy.

Offset Insertion—We portray the Offset Insertion deployment strategy in Figure 4 and summarize its major steps below:

1. Satellites launch aboard some launch vehicle, entering a LEO parking orbit. The launch vehicle upper stage is refueled in orbit.
2. The upper stage containing the satellites carries out burns to exit LEO, carrying out a Hohmann transfer to a circular deployment orbit of radius R_D .
3. Over the course of some max number O_{max} of deployment orbit periods, the satellites separate and execute a Hohmann transfer to their final position in the circular ring

The ΔV of a satellite in this deployment scenario is that of a Hohmann transfer between circular orbits of radius R_D and R , for which we provide a formula in Equation (4) [20].

$$\Delta V_{\text{sat}} = \sqrt{\frac{\mu}{R_D}} \left(\sqrt{\frac{2R}{R+R_D}} - 1 \right) + \sqrt{\frac{\mu}{R}} \left(1 - \sqrt{\frac{2R_D}{R+R_D}} \right) \quad (4)$$

where μ is the Sun’s gravitational parameter.

Similarly, the launch vehicle ΔV is that of a Hohmann transfer, this time between the Earth’s orbit (approximated as a 1 AU circular orbit) and the circular deployment orbit or

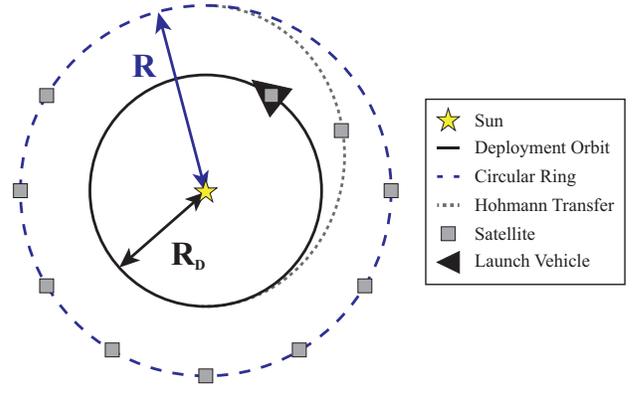


Figure 4: Illustration of step 3 of the Offset Insertion deployment strategy.

radius R_D . We however also account for the first burn of the Hohmann transfer coinciding with the burn to attain exit the Earth’s gravitational well, as shown in Equation (5).

$$\Delta V_{\text{LV}} = \sqrt{\Delta V_1^2 + \frac{2\mu_E}{R_E + a}} - \sqrt{\frac{\mu_E}{R_E + a}} + \|\Delta V_2\| \quad (5)$$

where μ_E is the Earth’s gravitational parameter, R_E is the Earth’s radius, $a = 200$ km is the altitude of the parking orbit, and ΔV_1 and ΔV_2 are the ΔV s associated with the first and second burn of the launch vehicle’s Hohmann transfer, respectively. ΔV_1 and ΔV_2 are presented in Equations (6) and (7).

$$\Delta V_1 = \sqrt{\frac{\mu}{1 \text{ AU}}} \left(\sqrt{\frac{2R_D}{R_D + 1 \text{ AU}}} - 1 \right) \quad (6)$$

$$\Delta V_2 = \sqrt{\frac{\mu}{R_D}} \left(1 - \sqrt{\frac{2 \text{ AU}}{R_D + 1 \text{ AU}}} \right) \quad (7)$$

The satellite ΔV , ΔV_{sat} , goes to 0 as the deployment radius R_D gets close to the circular ring radius R . However, as the difference between R_D and R gets small, the deployment time using the OI strategy gets infinitely long. Indeed, this deployment strategy relies on the launch vehicle orbiting the Sun at a different period than the satellites in the circular orbit, such that the satellites are able to insert at different true anomalies and create the desired circular ring of satellites. To avoid unrealistically-long deployment times, we set an upper limit on the number of periods of the deployment orbit, O_{max} , before which all satellites should have entered their insertion Hohmann transfer. For a given O_{max} , we have two ideal deployment radii R_D , one larger than R (R_D^{over}) and the other smaller (R_D^{under}). We present analytical expressions for these radii in Equations (8) and (9).

$$R_D^{\text{under}} = R \cdot \left(\frac{O_{\text{max}}}{O_{\text{max}} + 1} \right)^{\frac{2}{3}} \quad (8)$$

$$R_D^{\text{over}} = R \cdot \left(\frac{O_{\text{max}}}{O_{\text{max}} - 1} \right)^{\frac{2}{3}} \quad (9)$$

For our Earth-Mars HRN, we selected $O_{\max} = 3$, which keeps the maximum deployment time of circular rings using the OI strategy on the order of $O_{\max} \cdot T_{R^{\text{over}}} \leq O_{\max} \cdot T_{\text{Ma}} \cdot 1.5 \approx 8.46$ years.

We evaluate both the R_D^{under} and R_D^{over} options when choosing the deployment strategy to use for some specific circular ring of satellites.

Phasing Insertion—We also consider a second circular ring deployment strategy: Phasing Insertion (PI). Its major steps are as follows:

1. Satellites launch aboard some launch vehicle, entering a LEO parking orbit. The launch vehicle upper stage is refueled in orbit.
2. The upper stage containing the satellites carries out burns to exit LEO and carry out a Hohmann transfer to the final circular orbit, of radius R
3. The satellites enter elliptical phasing orbits, as portrayed in Figure 5 each reaching their final position in the circular ring after some maximum number of phasing orbit periods O_{\max}

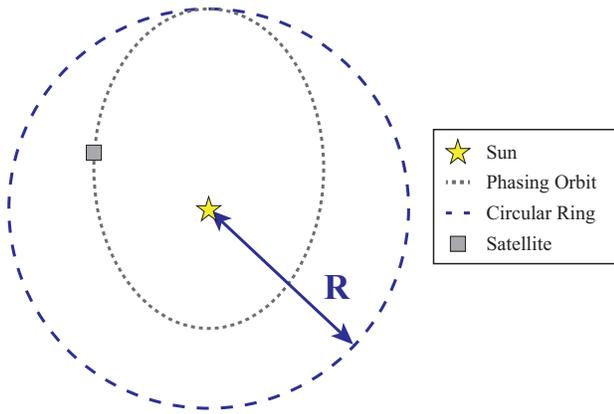


Figure 5: Phasing maneuver within a circular orbit. The spacecraft enters an elliptical phasing orbit, then later reinserts into the circular orbit with a different true anomaly.

The ΔV expended by the launch vehicle is similar to that expended in the OI deployment scenario, except that the bus containing the satellites inserts directly into the final circular orbit of radius R instead of some other deployment orbit. Therefore we have the launch vehicle upper stage ΔV presented in Equation (10).

$$\Delta V_{LV} = \sqrt{\Delta V_1^2 + \frac{2\mu_E}{R_E + a}} - \sqrt{\frac{\mu_E}{R_E + a}} + \|\Delta V_2\| \quad (10)$$

where ΔV_1 and ΔV_2 are as defined in Equations (11) and (12).

$$\Delta V_1 = \sqrt{\frac{\mu}{1 \text{ AU}}} \left(\sqrt{\frac{2R}{R + 1 \text{ AU}}} - 1 \right) \quad (11)$$

$$\Delta V_2 = \sqrt{\frac{\mu}{R}} \left(1 - \sqrt{\frac{2 \text{ AU}}{R + 1 \text{ AU}}} \right) \quad (12)$$

The satellite ΔV s in this deployment scenario are no longer all equal—since satellites need to end up at different true anomalies in the circular ring, they all have different true anomaly changes ΔP_i and therefore, different phasing orbit semi-major axes. Suppose one launch vehicle is deploying N_{sats} satellites into a fraction f_R of a circular ring (for instance, if we were deploying a full circular ring of $\frac{1}{f_R} N_{\text{sats}}$ satellites using $N_1 = \frac{1}{f_R}$ launches). Then the phase change assignment of satellite i is given by Equation (13).

$$\Delta P_i = 2\pi f_R \cdot \left[\left(i - \frac{1}{2} \right) \cdot \frac{1}{N_{\text{sats}}} - \frac{1}{2} \right] \quad (13)$$

We then obtain the phasing ΔV of satellite i from Equation (14) [20]. We set $O_{\max} = 3$ so as to obtain similar worst-case deployment times as with the Offset Insertion deployment strategy.

$$\Delta V_{\text{sat}}^i = 2 \left\| \sqrt{\frac{\mu}{R}} \left(\sqrt{2 - \left(1 + \frac{\Delta P_i}{2\pi \cdot O_{\max}} \right)^{-2/3}} - 1 \right) \right\| \quad (14)$$

Elliptical Ring Deployment

The Offset Insertion strategy isn't feasible for deploying *elliptical* rings of satellites because there is a loss of rotational symmetry in the shape of the satellites' final orbits. Consequently, we only consider the Phasing Insertion deployment strategy when computing the deployable satellite dry mass to elliptical satellite rings. The Phasing Insertion deployment procedure for elliptical orbits is in all ways identical to the procedure for circular orbits, except that the final ring is elliptical instead of circular. This difference adds some nuance to the ΔV evaluation procedure, namely the need to compare the execution of phasing maneuvers at the final ring's perihelion and at its aphelion.

4. NETWORK PERFORMANCE ANALYSIS

The existence of any distributed Mars–Earth relay network accomplishes most of the objectives of the Mars–Earth HRN, including avoiding communication outages during solar conjunctions, and reducing the number of Mars missions needing dedicated DSN antenna time. Therefore, the primary goal of our HRN orbital optimization is to maximize the data throughput from Mars to Earth, maximizing the scientific returns and productivity of current and future Mars missions. Our performance metric used to evaluate the fitness of a candidate HRN is thus the **average Mars–Earth downlink data rate** over a simulation duration (6.39 years) of three Mars–Earth synodic periods. We chose an integer multiple of the Earth–Mars synodic period as the simulation length in order to not bias the HRN topology towards a particular Earth–Mars arrangement, and 6.39 years is a reasonable mission duration for spacecraft subjected to the harsh environment of interplanetary space.

In this section, we detail several aspects of our network performance evaluation methodology, including how we compute the achievable data rate between nodes, and how we take advantage of the predictable evolution of the network's structure under Keplerian dynamics to efficiently identify

the maximum-capacity path between Mars and Earth in the HRN’s time-varying network graph.

Link Budget

We present the specifications and parameters of the proposed communication scheme, including assumed modulation schemes and forward error correction (FEC) codes, in Table 2. We use the parameters presented in Table 3 as the link performance parameters of the satellite communications systems, both at the Mars node of the network and at the heliocentric satellite nodes. We use the performance parameters presented in Table 4, taken from the DSN 34 m BWG antenna’s documentation, as the link performance parameters for the DSN antennas at the Earth node of the network. In order to be conservative in our link performance estimates, we use a DSN 34 m antenna system noise equivalent temperature (T_{sys}) of 60 K, which is larger than that measured during all nominal Ka-band operations reported during calibration in 2018 [21].

Table 2: Communication Link Parameters

Parameter	Symbol	Value	Units
Carrier band	-	Ka	-
Carrier frequency	f	32	GHz
Modulation type	-	QPSK	-
Modul. spectral efficiency	η	2	bps/Hz
FEC type	-	LDPC	-
FEC code rate	r	1/2	NDU
System losses	L_s	3	dB
Link Margin	L_m	3	dB

Table 3: Satellite Link Budget Parameters

Parameter	Symbol	Value	Units
Antenna diameter	D_{antenna}	8	m
Antenna efficiency	e_a	0.6	-
System noise temperature	T_{sys}	150	K
Transmit Power	P_{tx}	100	W
Max. Data Rate	R_{max}	1	Gbps

Table 4: DSN 34 m Antenna Link Budget Parameters [21]

Parameter	Symbol	Value	Units
Antenna diameter	D_{antenna}	34	m
Antenna efficiency	e_a	0.66	-
System noise temperature	T_{sys}	60	K

For any given link in the network, at some timestep, we can evaluate the received power P_{rx} at the receiving antenna using the RF link budget equation, described in Equation (15). In the link budget equation, P_{tx} is the transmit power of the transmitting antenna in W, and $G_{\text{tx}}, G_{\text{rx}}$ are the antenna gains over isotropic antennas, in dimensionless units, of the transmitting and the receiving antenna respectively [22].

$$P_{\text{rx}} = \frac{P_{\text{tx}} \cdot G_{\text{tx}} \cdot G_{\text{rx}}}{L_s \cdot L_m \cdot \text{FSPL}} \quad (15)$$

The gain of a parabolic antenna over an isotropic radiator, in non-dimensional units, is given by Equation (16), where λ is the wavelength corresponding to the carrier frequency [22].

$$G_{\text{antenna}} = \left(\frac{\pi D_{\text{antenna}}}{\lambda} \right)^2 \cdot e_a \quad (16)$$

Finally, the free space path loss (FSPL) accounts for the power lost to the propagation of the RF signal over large distances. A large FSPL is the main factor limiting the data rate of communications over interplanetary distances. We use the formula presented in Equation (17) to calculate the FSPL associated with a given link distance d [22].

$$\text{FSPL} = \left(\frac{4\pi d}{\lambda} \right)^2 \quad (17)$$

The noise spectral density for microwave signals is given by $N_0 = k \cdot T_{\text{sys}}$, where k is Boltzmann’s constant and T_{sys} is the receiving system’s equivalent noise temperature. N_0 is generally expressed in W/Hz, and is related to the signal to noise ratio (SNR), the bandwidth B and the received power P_{rx} by Equation (18) [22].

$$B = \frac{P_{\text{rx}}}{N_0 \cdot \text{SNR}} \quad (18)$$

For a given link, we obtain the bandwidth B from Equation (18) by assuming that the SNR is some desired value, in our case $\text{SNR}_{\text{desired}} = 10$ dB. Such an SNR, when combined with our usage of QPSK and the LDPC forward error correction scheme, should lead to bit error rates (BER) of less than 10^{-8} , which is reasonable for the downlink of Martian scientific data and imagery [22].

We finally obtain the data rate R_b of the link by incorporating the QPSK modulation spectral efficiency η and the LDPC code rate r : $R_b = \eta \cdot r \cdot B = 2 \cdot 1/2 \cdot B = B$ in our case.

To validate our link budget methodology, we compared our estimated data rate against the measured data rate record of MRO at the Earth-Mars closest approach distance of $d = 0.373$ AU. For this comparison, we used the link budget parameters for MRO’s 3 m High Gain Antenna (HGA), communicating using X-band (8 GHz) with the DSN’s 70 m antenna, as well as the correct “bits per symbol” of the used BPSK modulation scheme and other relevant parameters relevant to the MRO’s turbocode error correction [17].

We obtain an estimated data rate of 4.98 Mbps using the methodology detailed in this section, which is within an acceptable range of the maximal data rate of 6.0 Mbps achieved by MRO during its X-band communications with the Earth [17]. Our data rate estimates throughout the rest of this work, which use the conservative DSN antenna parameters presented in Table 4, are therefore consistent with the performance of existing RF communications systems.

Frequency Band Restrictions

For space to Earth Ka-band links, the ITU has allocated the 31.8–32.3 GHz frequency band. In practice, this spectrum

is also subdivided into multiple channels and shared among several users [23].

The QPSK modulation and rate-1/2 LDPC coding listed in Table 2 would therefore exceed the allocated Ka-band spectrum at the maximal data rate of 1000 Mbps listed in Table 3. Accordingly, the HRN-DSN link is assumed to employ more spectrally efficient modulation and forward error correction schemes (e.g., GMSK and higher-rate LDPC codes) to remain compliant with ITU frequency allocations. This assumption applies specifically to the HRN-DSN link and therefore does not affect inter-satellite links within the HRN, or the achievable end-to-end Mars–Earth data rates obtained in this work.

Maximum Capacity Path

The network graph at some timestep has the Earth, Mars, and HRN satellites as nodes, and the achievable data rate between any two nodes as edge weights. Evaluating the Mars–Earth downlink data throughput of some HRN at some timestep t is equivalent to identifying the maximum data rate transmissible along any single path between Mars and the Earth in the network graph at that timestep. With a large number of nodes and combinatorially-many edges, identifying the path that enables the largest capacity feasible is a non-trivial integer optimization problem with no known ideal formulation. This problem is the “Maximum Capacity Path” (MCP) problem, of which we plot an example in Figure 6. The MCP problem is also referred to as the “maximum bandwidth path problem” and the “widest path problem”. It can also be mapped to the Bottleneck Shortest Path problem if the edge weights in the network represent a cost to minimize instead of a capacity to maximize as in the MCP [24].

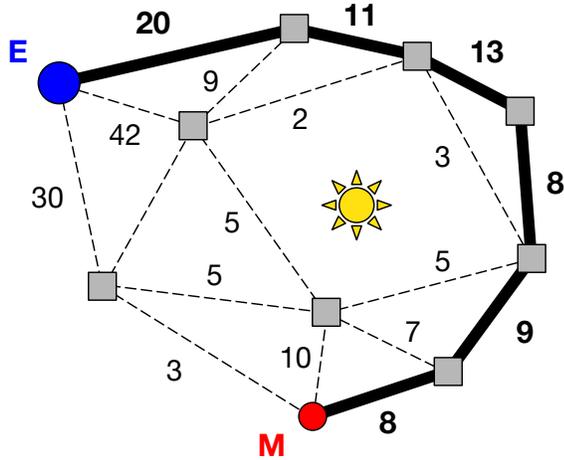


Figure 6: Illustration of a HRN at some timestep t , with internode achievable data rates as edge weights. A Maximum Capacity Path (MCP) is shown in bold. The *limiting capacity* of the network represented is 8, because no path exists between E and M with all edge capacities greater than 8.

We formulate the Max-Capacity Path integer optimization problem at some time t . Let $G_t = (V, E)$ be the undirected graph representing the interplanetary network at time t . Let C_{ij} be the edge weight (capacity) between nodes i and j . Let s be the start node and t be the target (arrival) node – in this interplanetary application, Mars and Earth respectively.

We have the decision variable:

$$x_{ij} = \begin{cases} 1 & \text{if edge } (i,j) \text{ is selected to be in the path} \\ 0 & \text{otherwise} \end{cases}$$

and the auxiliary variable C_{\min} that represents the capacity of the minimum capacity link along the selected path. We obtain the integer optimization problem presented in Equation (19)

$$\begin{aligned} & \max(C_{\min}) \\ \text{s.t.} & \\ & C_{ij} \cdot x_{ij} \geq C_{\min} \cdot x_{ij} \quad \forall (i,j) \in E \\ & \sum_{i \in V} x_{it} - \sum_{j \in V} x_{tj} = 1 \\ & \sum_{i \in V} x_{is} - \sum_{j \in V} x_{sj} = -1 \\ & \sum_{i \in V} x_{ik} - \sum_{j \in V} x_{kj} = 0 \quad \forall k \in V \setminus \{s, t\} \\ & x_{ij} \in \{0, 1\} \quad \forall (i,j) \in E \\ & C_{\min} \in \mathbb{R}_+ \end{aligned} \tag{19}$$

To solve this integer optimization problem, we employ a binary search algorithm adapted from a thesis on bottleneck problems in graphs [25], which we outline in Algorithm 1.

Algorithm 1 Find Max-Capacity Path (MCP) and C_{\min}

- Input:** $G = (V, E)$
Output: MCP, C_{\min}
- 1: $F \leftarrow E$
 - 2: **while** $|F| > 1$ **do**
 - 3: $\lambda \leftarrow$ mean edge weight in F
 - 4: Delete all edges $e = (i, j)$ with $C_{ij} < \lambda$;
let G' be the resulting graph
 - 5: **if** there is a path from s to t in G' **then**
 - 6: $F \leftarrow \{e \in F : C_{ij} \geq \lambda\}$
 - 7: **else**
 - 8: $F \leftarrow \{e \in F : C_{ij} < \lambda\}$
 - 9: **end if**
 - 10: **end while**
 - 11: $C_{\min} \leftarrow$ weight of remaining edge
 - 12: Construct $G_D = (V_D, E_D)$, the same graph but with internode distances as edge weights instead of available data rate. Delete all edges in E_D with indices (i, j) such that $C_{ij} < C_{\min}$ in E ; let G_{final}^D be the resulting graph
 - 13: $MCP \leftarrow$ shortest path (Dijkstra’s) from s to t in G_{final}^D
-

Finding the shortest path at the last step of Algorithm 1 returns the shortest path in the heliocentric network at some timestep t that is still a MCP. Thus, we preserve our network performance metric – the Mars–Earth data rate at that timestep – all the while minimizing the total communication delay along the final selected path by minimizing the total distance traveled along the path.

Network Analysis in a Time-Varying Graph

The structure of the heliocentric network changes between timesteps as the Earth, Mars, and the HRN satellites orbit the Sun. To use the average data rate over the simulation period as the objective function to minimize in our genetic algorithm, we thus need to find the data rate from Mars to Earth at each individual timestep. Finding each data rate by running Algorithm 1 independently at each timestep proves to be excessively computationally-expensive for a function that will be called millions of times in a single run of the genetic algorithm.

Instead, we can warm-start each timestep (except the first) with information from the previous timestep. This approach is effective because the structure of the network changes very little between consecutive timesteps at the chosen simulation temporal resolution of 72 h intervals. When the networks at t and $t - 1$ are similar, computing the data rate along the edges of MCP_{t-1} in G_t provides an excellent lower bound for C_t^{\min} , accelerating the convergence of Algorithm 1. We summarize this approach in Algorithm 2.

Algorithm 2 Warm-started MCP in a Time-Varying Graph

Input: Timesteps $t = 1$ to T , graphs $\{G_t\}_{t=1}^T$

Output: $\{C_t^{\min}, MCP_t\}_{t=1}^T$

- 1: Precompute link capacity $C_{ij}(t)$ for all t
 - 2: Prune G_1 by removing edges with $C_{ij} \leq C_{DTE}$; Let G_1^p be the resulting graph
 - 3: Run Algorithm 1 on G_1^p to obtain C_1^{\min}, MCP_1
 - 4: **for** $t = 2 \rightarrow T$ **do**
 - 5: $C_t^* \leftarrow$ minimum edge weight along MCP_{t-1} in G_t
 - 6: Prune G_t to keep edges with $C_{ij} \geq C_t^*$; let G_t^p be the resulting graph
 - 7: Run Algorithm 1 on G_t^p to obtain C_t^{\min}, MCP_t
 - 8: **end for**
-

Simulation

To model Earth and Mars’ orbits over our simulation period, we use planetary ephemerides in the J2000 ecliptic frame generated using NASA JPL NAIF’s SPICE toolkit [26], [27].

To propagate the positions of HRN satellites throughout the simulation period in an accurate, numerically stable manner, we propagate the satellite orbits in Keplerian coordinates. Keplerian orbital propagation is symplectic: varying the true anomaly of a Keplerian orbit has no impact on the orbit’s total energy. A drawback, however, is that this approach cannot account for 3rd body or solar perturbations, since we don’t explicitly model the forces acting on an object in an elliptical orbit.

An object orbiting some central body in an elliptical orbit can be described by 6 Keplerian coordinates, summarized in Table 5.

The only Keplerian orbital element that varies over time is the true anomaly, which represents the angle between the orbiting object’s position and the orbit’s periapsis. There is no known analytical expression for the true anomaly, which we obtain from the mean anomaly M by solving Kepler’s equation (Equation (20)). The mean anomaly M corresponds to what

Element Name	Element Symbol
Semi-major axis	a
Eccentricity	e
Inclination	i
RAAN	Ω
Argument of Periapsis	ω
True Anomaly	θ

Table 5: Keplerian Orbital Elements

the true anomaly θ would be if the orbit were circular, and is itself linearly propagated forward in time between timesteps.

$$M = \text{atan2}[\sqrt{1 - e^2} \cdot \sin(\theta), e + \cos(\theta)] - \frac{e \cdot \sqrt{1 - e^2} \cdot \sin(\theta)}{1 + e \cdot \cos(\theta)} \quad (20)$$

When we need to obtain a satellite’s position at some number of discrete time steps, we propagate its position within its Keplerian elliptical orbit, obtaining the true anomaly at each timestep. We then convert from Keplerian coordinates to Cartesian coordinates, obtaining the satellite’s position and velocity in the x-y-z plane at each timestep of the simulation.

5. MODULAR GENETIC ALGORITHM

Genetic Algorithms

A genetic algorithm (GA) is a population-based optimization algorithm that emulates some of the processes of natural selection. The genome of some individual in the GA represents the variables of the optimization problem. A genetic algorithm curates a population of individuals over several generations, as outlined in Figure 7. At each generation, the fitness of each individual is evaluated using the GA’s objective function; the obtained fitness values guide the use of genetic operations such as selection, mutation, and crossover. Over successive generations, the population’s average fitness increases, and converges towards some locally optimal solution.

Modular Genome

In order to design a HRN composed of several distinct satellite formations and constellations, we have developed a modular genetic algorithm genome. Each formation or constellation, which we refer to as a “feature”, is described by N_v variables. We can therefore encode a design space of N_f distinct features in a genome of $N_v \cdot N_f$ variables, as we have represented in Figure 8.

Meanwhile, each feature in our modular genome can correspond to a number of different types of formations and constellations, namely different variants of planar heliocentric MOGs, and different types of elliptical and circular satellite rings. We encoded these different options as an integer feature index, and used other variables in that feature’s genome only when relevant to that category of formation. By selectively using some subset of the parameters in a genome depending on the feature index, we have essentially implemented a modular version of the Hidden Genes Genetic Algorithm (HGGA) [28].

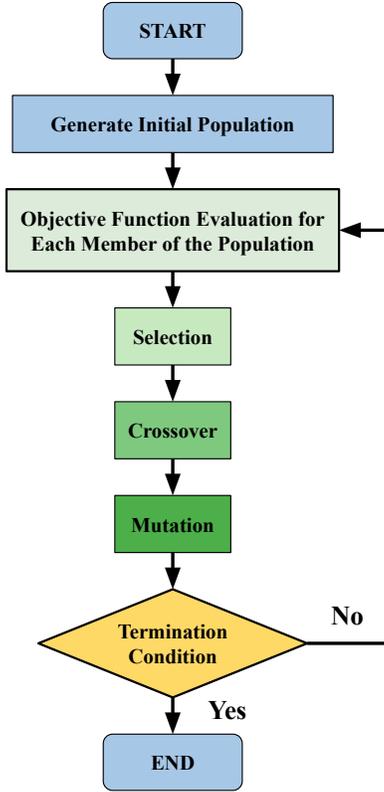


Figure 7: Block diagram outlining the steps of a genetic algorithm.

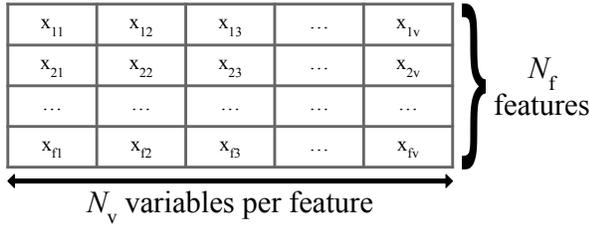


Figure 8: Modular GA genome encoding N_f features of N_v variables each.

We provide the details of the parameter assignments of our modular GA implementation in Table 6.

To constrain the optimization problem, we set a maximal total number of launches for a constellation, N_1 . To enforce this constraint across multiple features, we impose the following linear constraint on the number of launches N_1^i assigned to some i th feature: $\sum_{i=1}^{N_f} N_1^i \leq N_1$. Additionally, upper and lower bounds are placed on each of the GA variables, as we detail in Table 6.

Formations Overview

We plot representative examples of the six distinct formation options in Figure 9. Each plot also relates some of each formation’s parameters, as listed in Table 6, to the dimensions

and layout of the plotted satellite formations. Within a modular GA genome, each feature is attributed a number of launches N_{launches} that translates to a number of satellites N_{sats} once ΔV and fairing volume constraints are taken into account.

The first type of feature is a generic, planar, Mutually Orbiting Group (MOG), as we defined in prior work on heliocentric MOG deployment [9]. All MOG satellites are in planar elliptical orbits of semi-major axis a_{MOG} and eccentricity e_{MOG} . As we have represented in Figure 9a, a_{MOG} is therefore the distance between the center of the MOG and the solar system barycenter, while e_{MOG} independently determines how spread out the formation’s satellites are around the MOG center. GA variable 6, referred to as ω_0 for this formation type, serves as a phasing variable, defining the angle with respect to the positive x axis that the formation center lies at at the start of the HRN network performance simulation. Within the context of this network optimization, these generic MOGs are useful because they maintain their structure over time, all the while bridging the gap between different features, notably circular and elliptical rings, that lie at different distances from the Sun.

The second type of formation we include in our modular genome is a MOG variation that we refer to as a MOG “at a planet”. The integer variable PlanetID determines near which planet (Earth or Mars) the MOG is placed, fully determining the MOG semi-major axis a_{MOG} and placement ω_0 such that the MOG edge remains in close proximity to the planet throughout the planet’s elliptical orbit. We plot an example MOG placed at Mars in Figure 9b. While this feature type is included in the design space of the generic MOG feature, we include this variation as a distinct option because it would otherwise be unlikely for a genetic algorithm to identify such a specific pair of $(a_{\text{MOG}}, \omega_0)$ variables for the MOG to remain near a specific planet for the entire simulation period.

The third type of feature included in the modular genome is a MOG that “includes” a planet - it’s a MOG of same semi-major axis **and** eccentricity as a planet’s elliptical orbit. This is an important distinction from the second feature, since this inclusion into a MOG guarantees that the planet doesn’t oscillate around the edge of the MOG due to the eccentricity of its orbit, a second order effect which limits the average data rate between planets and a MOG or arbitrary eccentricity placed near them. This MOG is also slightly inclined ($i_{\text{MOG}} = i_{\text{planet}}$) and RAAN-based, such that the planet (Mars in Figure 9c) is truly integrated into a MOG as if it were a MOG satellite itself. This feature is therefore not perfectly in the ecliptic of the solar system, but ensures that the planet remains tied in to the HRN even when it dips out of the ecliptic during its inclined orbit. Since the inclination of the Earth is near-zero (by definition) and that of Mars is only 1.85° , we don’t include inclination change maneuvers in the MOG insertion ΔV computed for these formations.

The fourth type of formation employed in our modular GA is the planar circular ring of satellites, portrayed in Figure 9d. The radius R of the formation determines the radius of the satellites’ circular orbits. Meanwhile, the 8th variable, $frac$, determines the fraction of feature to fill with satellites. Finally, the 7th variable, θ_0 , determines the initial true anomaly of the first satellite of the ring, acting as a sort of phasing parameter for the gap in the ring created when $frac$ is non-zero.

Table 6: Variables used in the modular genome of the genetic algorithm. Dashes indicate that a given variable is inactive for that feature index.

Variable Number	1	2	3	4	5	6	7	8
Feature Name	Feature Index	Number of Launches	Planet Selected	Semi-major Axis	Eccentricity	Argument of Periapsis	True Anomaly of Satellite 1	Fraction of Feature Filled
MOG	1	N_{launches}	-	a_{MOG}	e_{MOG}	w_0	-	-
MOG @ Planet	2	N_{launches}	PlanetID	-	e_{MOG}	-	-	-
MOG Incl. Planet	3	N_{launches}	PlanetID	-	-	-	-	-
Circular Ring	4	N_{launches}	-	R	-	-	θ_0	frac
Elliptical Ring	5	N_{launches}	-	a_{ring}	e_{ring}	w_{ring}	θ_0	frac
Ring Along Planet's Orbit	6	N_{launches}	PlanetID	-	-	-	-	frac
Variable Lower Bound	0	0	1	a_{min}	0	0	0	0
Variable Upper Bound	6	N_1	2	a_{max}	1	2π	2π	1
Integer Variable	Yes	Yes	Yes	No	No	No	No	No

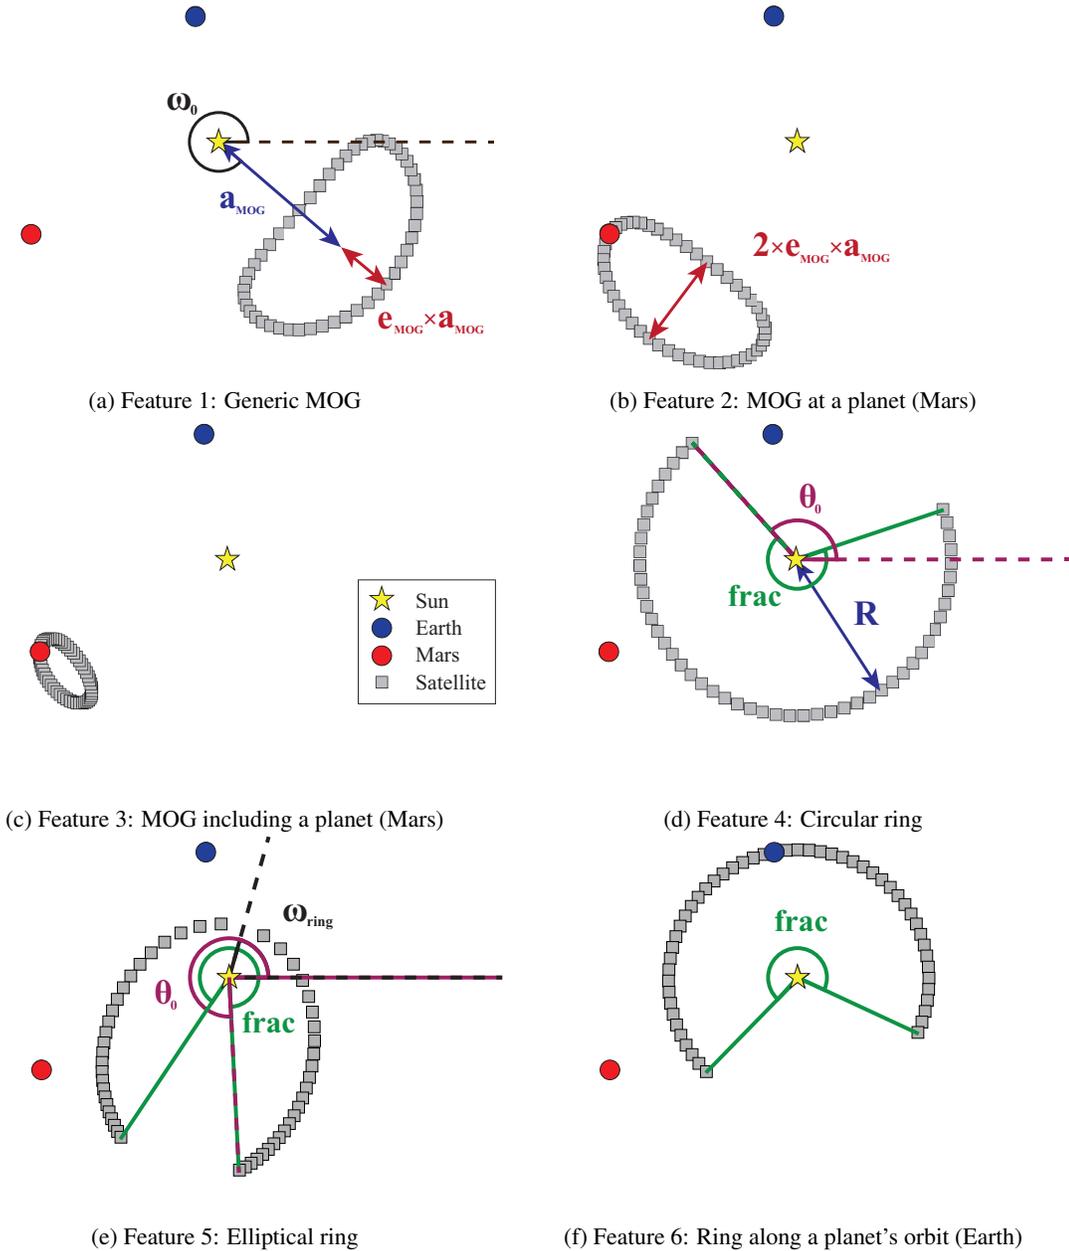


Figure 9: Representative examples of the 6 different modular GA features.

The fifth type of feature included in our design space is the planar elliptical ring of HRN satellites represented in Figure 9e. This formation is similar in all ways to the circular ring, except that the satellites lie on an elliptical orbit of non-zero eccentricity e_{ring} . Fully defining the elliptical orbit also requires a second additional variable, ω_{ring} , the argument of periapsis of the elliptical orbit.

The sixth and final formation included in the GA’s design space is the “Ring along a planet’s orbit”, exemplified by Figure 9f. This feature is constructed by inserting relay satellites at even mean anomaly intervals along a planet’s real (inclined, elliptical) orbit. As in other similar features, frac determines the fraction of the mean anomaly filled by satellites, while PlanetID specifies which planet’s orbit to use. Just like feature 3, this formation is slightly inclined, but we neglect the additional ΔV incurred by inclination change maneuvers since it is small compared to the overall insertion ΔV .

Objective Function

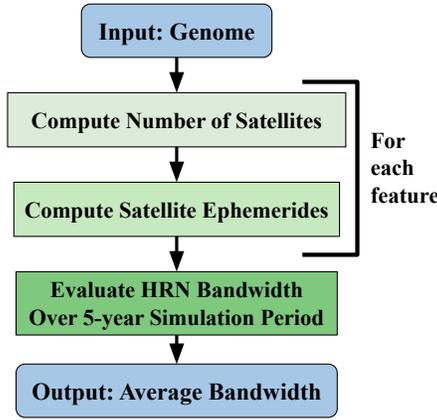


Figure 10: GA Objective function Outline

The objective function of our genetic algorithm is outlined in Figure 10. To evaluate the fitness of a given genome, we first compute the number of satellites N_{sats} deployable to each feature with the attributed number of launches N_{launches} . We then use that N_{sats} and the rest of the genome’s variables to compute the ephemerides of each feature’s satellites. Finally, we simulate the structure and performance of the resulting HRN over the full 6.4-year simulation period. The fitness of the genome being evaluated is then simply the Mars–Earth downlink capacity averaged over the simulation period.

Warm Starts

Consider the GA solution obtained for some pair of parameters (N_f^*, N_l^*) ; that solution is **valid** for any $N_f \geq N_f^*$ and $N_l \geq N_l^*$. We use this property to warm-start our modified genetic algorithm. When starting a GA, we take some number of warm-start individuals N_{ws} from neighboring, existing solutions, and integrate them into the initial population of the genetic algorithm.

We also carry out a version of this warm-start procedure to back-propagate solutions with larger N_l to smaller N_l . If some existing solution X_{large} has a larger N_l than sanctioned by the current GA parameters, we randomly remove launches from X_{large} ’s features until the resulting X_{large}^* has a total number of launches compliant with the N_l constraint.

Implementation Details

We implemented this modular genetic algorithm in MATLAB R2024a using the $\text{ga}()$ function from the Global Optimization Toolbox [29]. We summarize the parameters and settings used in Table 7. We scale the GA population size affinely with N_f since the genome size increases linearly (and the search space combinatorially) with N_f , as was illustrated in Figure 8.

Table 7: Summary of $\text{ga}()$ parameters and settings used in our implementation. Settings not listed in this table remain at their default values.

Parameter Name	Value
FunctionTolerance	1E-05
MaxStallGenerations	20
PopulationSize	$200 \cdot (N_f + 1)$
InitialPopulationMatrix	Warm-start individuals
N_{ws}	10% of PopulationSize

After the genetic algorithm completes, we use MATLAB’s $\text{fmincon}()$ as a hybrid function [30]. Using the same objective function as in the GA and holding the integer variables fixed, we provide the GA’s best solution as the initial guess for $\text{fmincon}()$. The solver then tweaks the continuous variables of that genome, converging towards the local optimum identified by the GA. We summarize this procedure in Figure 11.

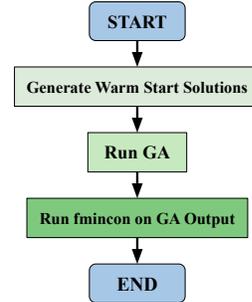


Figure 11: Diagram summarizing the full optimization procedure for some pair of parameters (N_f, N_l) .

Individual runs of the genetic algorithm for a given set of parameters (N_f, N_l) are independent of other GA runs. We therefore run these genetic algorithms in parallel batches of jobs on MIT Lincoln Lab’s SuperCloud HPC cluster [31].

We run these parallel batches of jobs over the parameter grid $(N_f, N_l) \in \{1, 2, 3\} \times \{1, 2, \dots, 125\}$. We run the GA over the entire grid of parameters several times, to enable the forward and back propagation of solutions across neighboring N_l values thanks to the warm-start procedure. This helps homogenize solutions, enforce monotonicity, and leads to the emergence of only a few distinct solutions structures over the (N_f, N_l) space.

6. RESULTS AND DISCUSSION

Network Performance versus Number of Launches

We present the main results of our analysis in Figure 12. The first curve, in blue, is the performance of the HRN networks

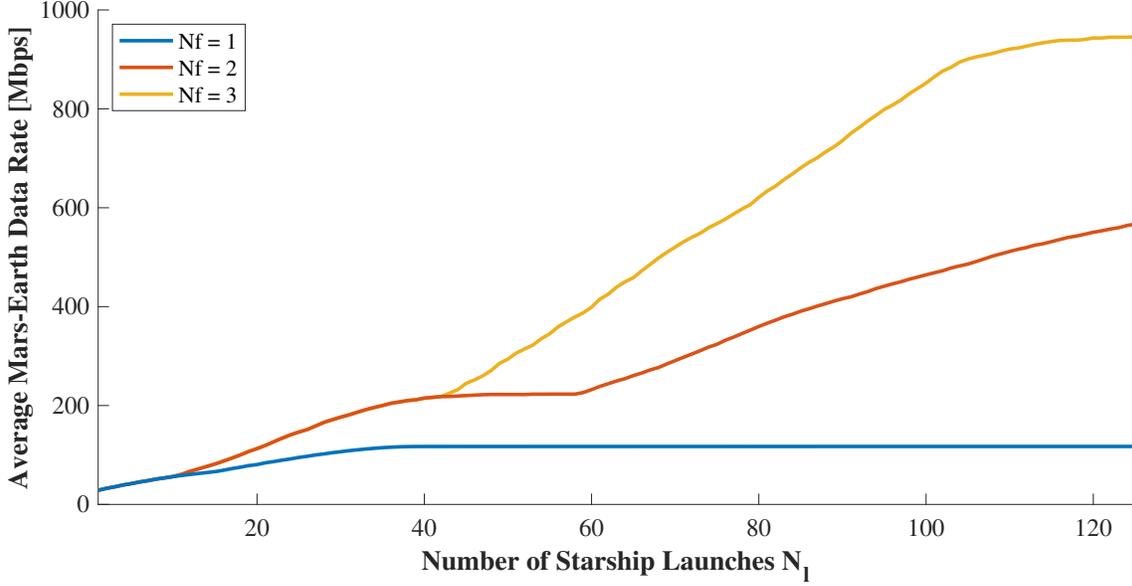


Figure 12: Average Mars–Earth data rate as a function of the number of Starship launches devoted to deploying the HRN.

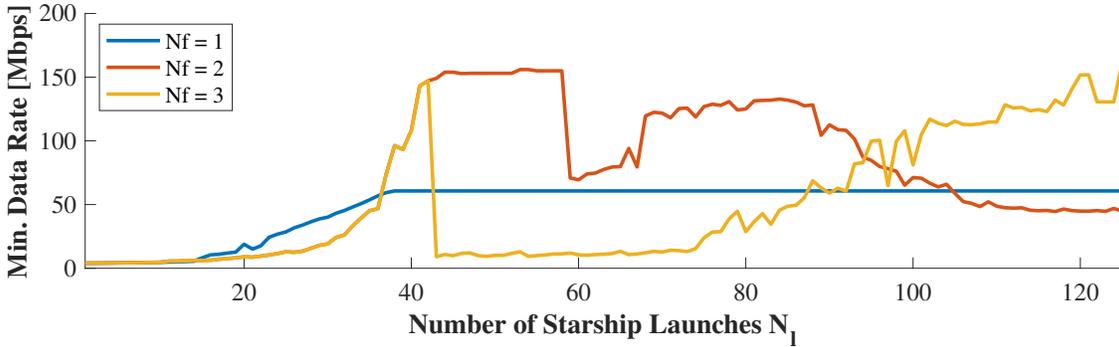


Figure 13: Minimum Mars–Earth data rate as a function of the number of Starship launches devoted to deploying the HRN.

that use only one distinct feature ($N_f = 1$). The second, in red, is that of the HRN networks with $N_f = 2$ features. The third, in yellow, is the performance curve of the HRN networks with $N_f = 3$ features.

For $N_1 \in \{1, \dots, 11\}$, all N_f values lead to the same optimal solutions since there are too few launches, and therefore too few HRN satellites, for the use of two or three distinct features to improve over the HRN solution obtained with $N_f = 1$.

For $N_1 > 11$, the $N_f = 2$ solutions start yielding higher average data rates than the $N_f = 1$ solutions. Starting at around $N_1 = 40$ launches, the $N_f = 1$ HRN networks plateau at an average data rate of 117.2 Mbps. Meanwhile, the $N_f = 2$ networks continue to increase along with N_1 , reaching a 566.8 Mbps average with 125 Starship launches.

Finally, the $N_f = 3$ solutions diverge from the $N_f = 2$ solutions for $N_1 \geq 43$. The $N_f = 3$ networks tend towards the maximal data rate $R_{\max} = 1000$ Mbps, surpassing an average of 800 Mbps by $N_1 = 96$ launches.

The objective of our optimization is to maximize the average Mars–Earth downlink data rate, and therefore the network throughput. Nonetheless, we are also interested in the minimum downlink data rate recorded during the evaluation of the HRNs returned by our GA. We plot the guaranteed minimum data rate over the 6.4-year simulation period of the GA solutions in Figure 13 for the same ranges of (N_f, N_1) parameters. There are large variations in minimum data rate in Figure 13 even when the throughput increases monotonically in Figure 12. These variations reveal that the GA solutions have different underlying network structures for different pairs of (N_f, N_1) parameters.

Network Architectures

These discrepancies in performance and minimum downlink data rate guarantees suggest important differences in the structure of the HRNs discovered by our genetic algorithm for different combinations of (N_f, N_1) parameters. We can summarize these qualitative differences in network structure by grouping similar solutions together into several different **architectures**. We represent the ranges of parameters, and average data rates, corresponding to each architecture in Figure 14. We plot representative members of each of the

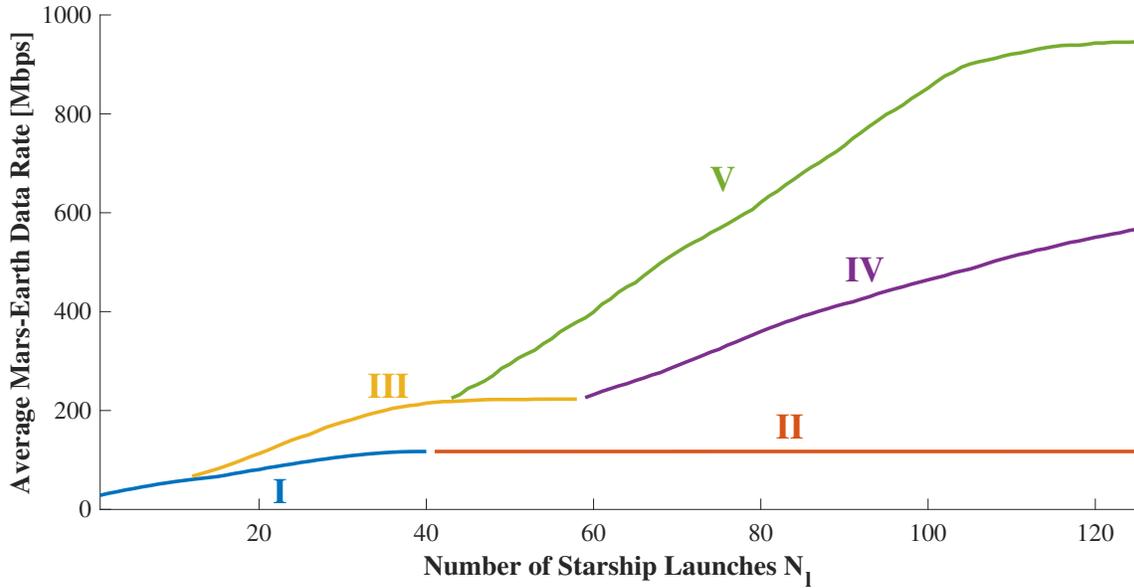


Figure 14: Average Mars–Earth data rate as a function of Starship launches, with colored curve segments indicating the five HRN architectures emerging from our optimization.

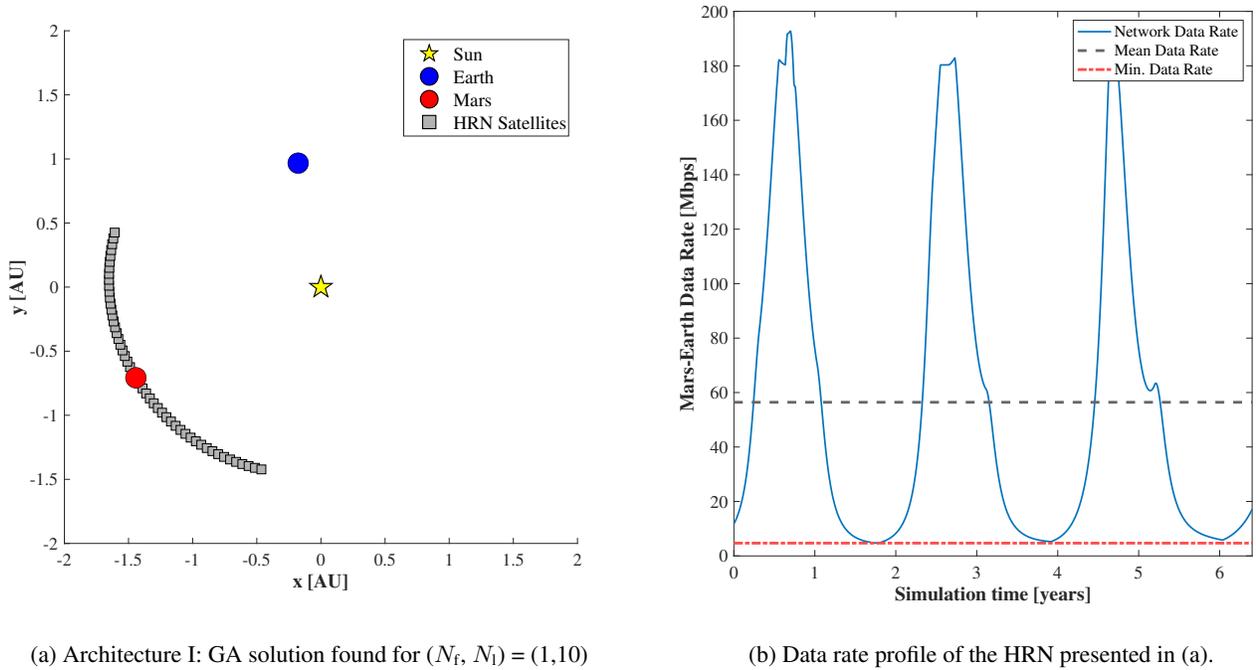
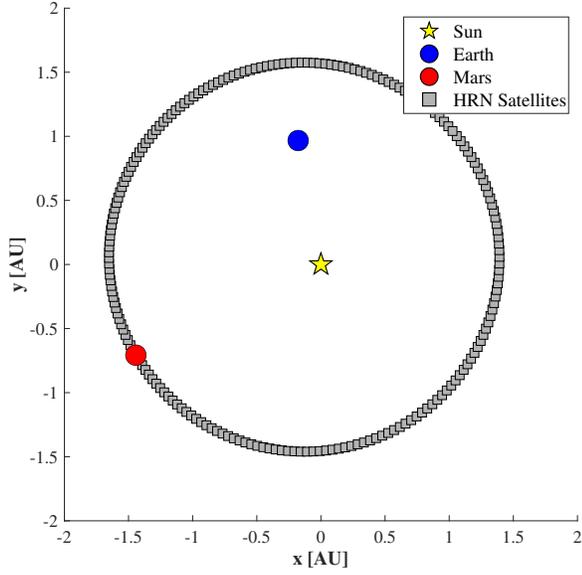


Figure 15: Representative example of HRN Architecture I.

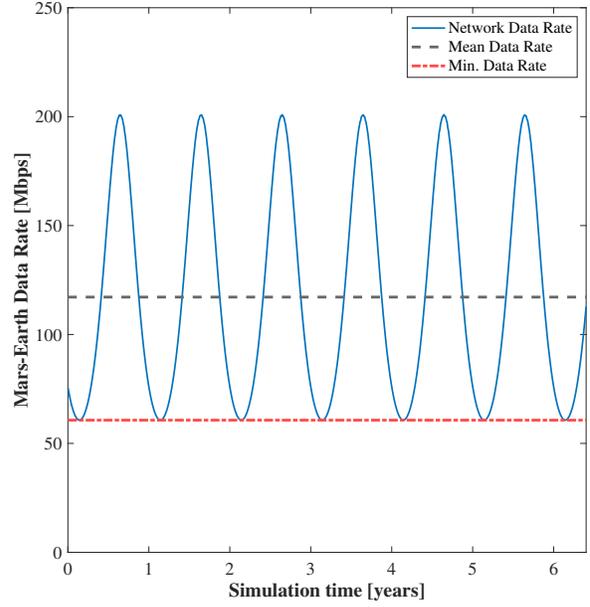
major uncovered HRN architectures in Figures 15-19, along with their throughput profiles.

From 1-11 launches, all N_f curves correspond to Architecture I, represented in Figure 15. Architecture I consists of a partial ring of HRN satellites distributed along Mars’ orbit. This architecture emerges from the low number of satellites available at small N_f , slightly increasing the HRN’s average throughput but providing poor minimum data rate guarantees when the Earth is far from the partial ring.

As N_f increases, the $N_f = 1$ network performance increases as the Mars satellite ring fills out. At $N_f = 41$, the ring is completely full of satellites, yielding Architecture II, represented in Figure 16. At that point, the link between the Earth and the Mars ring becomes the network bottleneck, causing the $N_f = 1$ solution performance to plateau at around 117.2 Mbps. Due to its symmetric layout around the Sun, however, the data rate over time, represented in Figure 16b, oscillates between minimum and maximum values that correspond to the Earth passing by Mars’ aphelion and perihelion. Archi-

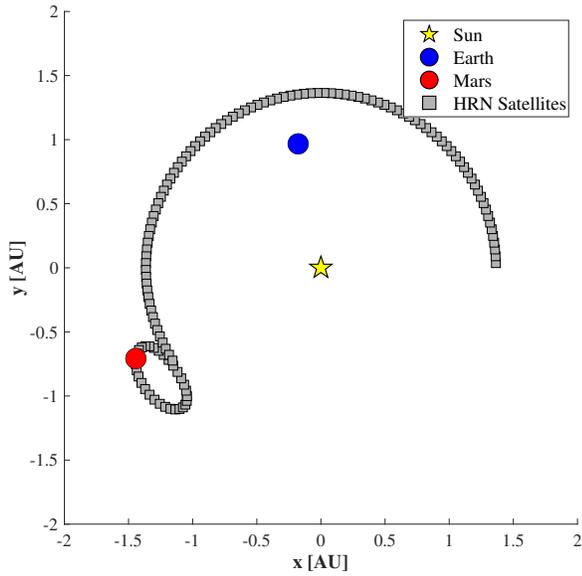


(a) Architecture II: GA solution found for $(N_f, N_i) = (1,40)$

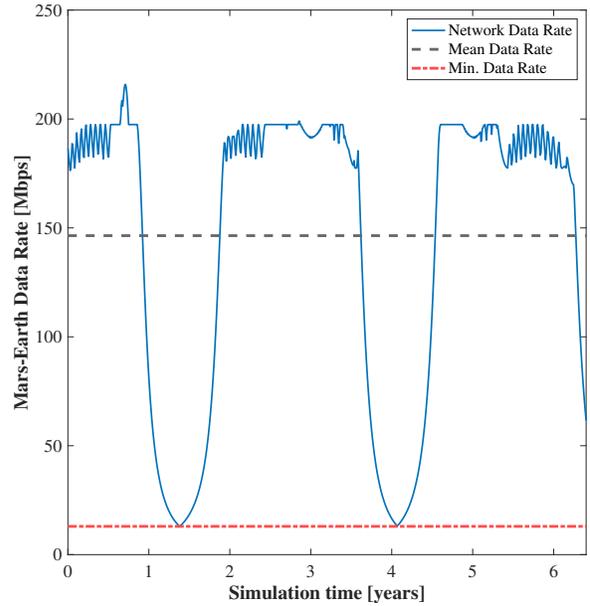


(b) Data rate profile of the HRN presented in (a).

Figure 16: Representative example of HRN Architecture II.



(a) Architecture III: GA solution found for $(N_f, N_i) = (2,25)$



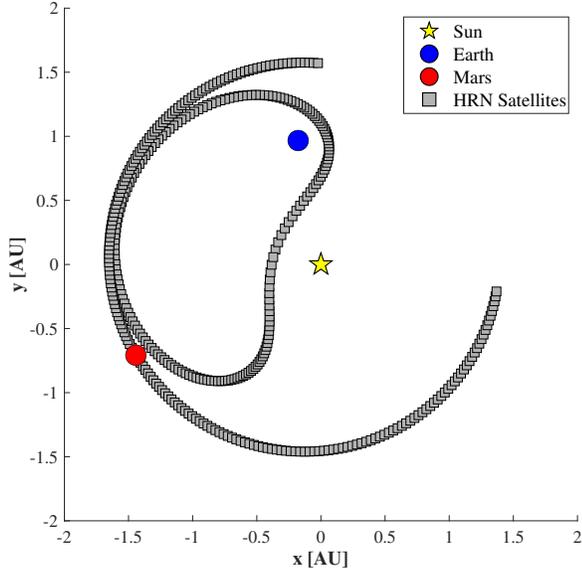
(b) Data rate profile of the HRN presented in (a).

Figure 17: Representative example of HRN Architecture III.

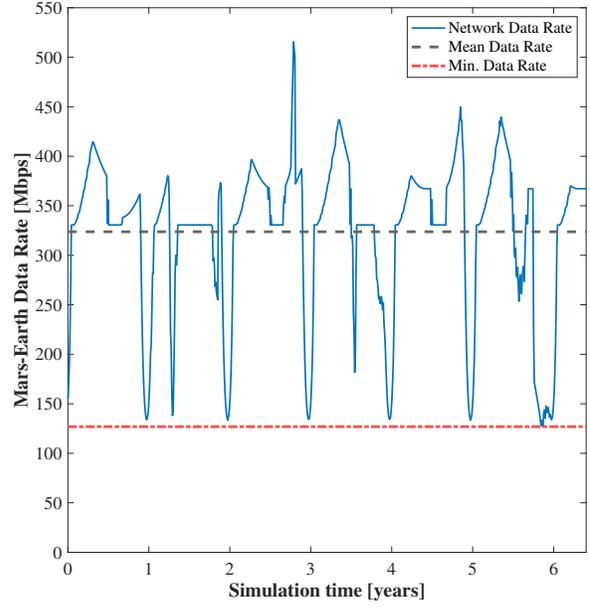
Architecture II provides reasonable minimum data rate guarantees over the full simulation period of 60.7 Mbps, as seen in Figure 13.

Meanwhile, the $N_f = 2$ and $N_f = 3$ solutions diverge from the $N_f = 1$ solutions for $N_i \geq 12$, yielding Architecture III. Architecture III consists of a MOG that includes Mars,

in close proximity to a partial circular ring. This architecture effectively increases the throughput between the Earth and the circular ring of Architectures I and II by reducing the ring's radius, all the while using the Mars MOG to keep the red planet connected to the HRN. Figure 17 contains an example Architecture III solution, along with its data rate profile.

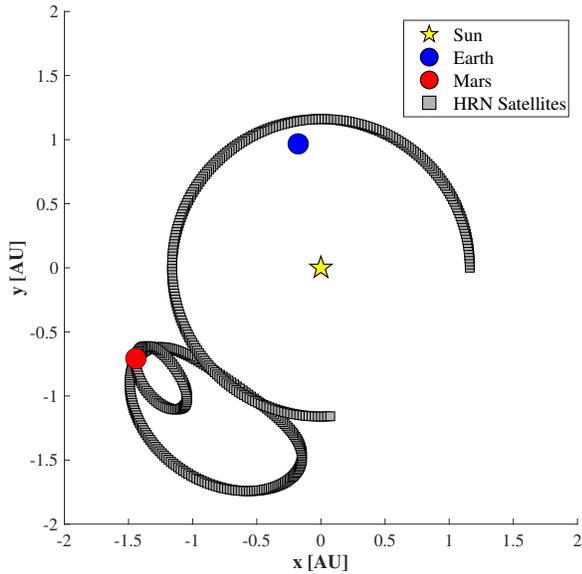


(a) Architecture IV: GA solution found for $(N_f, N_i) = (2, 75)$

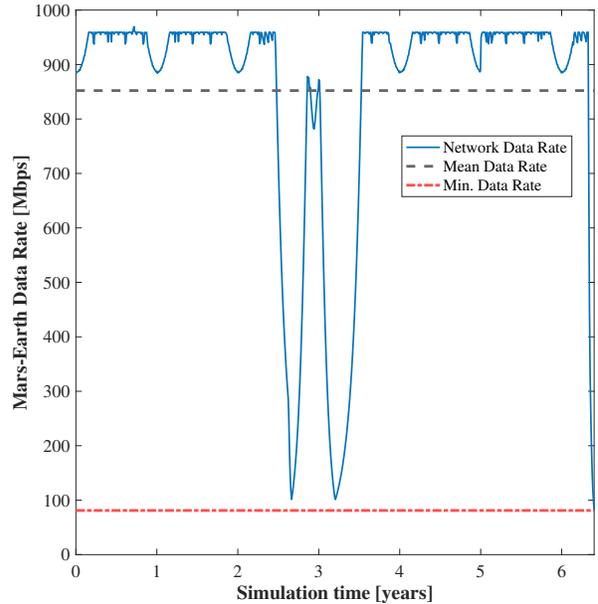


(b) Data rate profile of the HRN presented in (a).

Figure 18: Representative example of HRN Architecture IV.



(a) Architecture V: GA solution found for $(N_f, N_i) = (3, 100)$



(b) Data rate profile of the HRN presented in (a).

Figure 19: Representative example of HRN Architecture V.

As its partial circular ring is extended into a full ring, Architecture III's mean downlink throughput plateaus near 223.1 Mbps, additional Starship launches yielding diminishing returns past $N_i = 48$. Filling out the circular ring does however significantly improve the minimum data rate of the Architecture III solutions, which reaches 154.0 Mbps for $44 \leq N_i \leq 52$, as shown in Figure 13.

We plot the last two-feature architecture, Architecture IV, in Figure 18. Architecture IV consists of a partial ring of satellites along Mars' orbit, paired with a large MOG that remains close to the Earth. Architecture IV continues to increase the average Mars–Earth throughput compared to Architecture III's performance plateau, at the cost of returning to lower

minimum data rate guarantees due to the incompleteness of its satellite ring.

Finally, the only three-feature architecture selected by our GA, Architecture V, is represented in Figure 19. Architecture V consists of a circular ring of satellites in close proximity to Earth, a MOG that includes Mars, and a larger, planar, MOG near Mars that ties the Mars MOG and the circular ring together. The combination of the two MOG types allows for some satellites to be in constant proximity to Mars, all the while moving the radius of the circular ring closer to the Earth, further increasing the data rate between the Earth and the circular ring. The smaller, inclined, MOG also reduces the gap between Mars and the HRN when Mars is far from the ecliptic due to its inclined orbit. As we can see in Figure 19b, the gap in the circular ring leads to an occasional dip in network throughput, providing poor minimum throughput guarantees. As N_1 increases, the gap is progressively filled in, increasing Architecture V's minimum data rate to close to 150 Mbps (Figure 13). Finally, Architecture V's average downlink throughput curve starts to plateau as it approaches the enforced data rate limit of 1000 Mbps.

7. FUTURE WORK

Our methodology and the results presented in this paper rely on several assumptions that may be relaxed or otherwise modified in future work. These include the assumption that all satellites should have two antennas of diameter 8 m—future work could investigate decreasing the antenna diameter, allowing more satellites to be packed into a single launch vehicle but reducing antenna gain. Future work may also examine the sensitivity of our results to the use of different launch vehicles, higher-fidelity mass and power modeling of the HRN satellites, and the use of different DSN antennas at the Earth interface.

Another direction for future work is the inclusion of additional types of heliocentric formations in the design space of our modular genetic algorithm. Candidates for inclusion include 2D planar flower constellations, as well as constellations that rely on multi-body dynamics.

Furthermore, in future work we may examine the use of different communication technologies, including but not limited to optical communications, phased array antennas, and deployable antennas. These different technologies vary in flight heritage, robustness, throughput, and spectral efficiency, meriting further investigation in the context of heliocentric relay networks for interplanetary communications.

A natural extension of the work presented in this paper is adapting our methodology to the design of HRNs for high-data-rate communications with other planets (Venus, Mercury) and loci of interest throughout the solar system such as Earth-Sun Lagrange points. The modular architecture of the HRN makes it scalable to extending communications services to several distinct points in the Solar System simultaneously. This opens avenues of investigation into the multi-objective optimization of HRNs for interplanetary communications, balancing equity and throughput in multi-user interplanetary networks.

8. CONCLUSIONS

We presented the concept of a Heliocentric Relay Network (HRN): a distributed, Sun-orbiting constellation of relay satellites designed for continuous high-data-rate Mars–Earth communications. We developed a methodology that incorporates satellite deployment considerations, link budget computation, and network throughput evaluation, to obtain an HRN's average Mars–Earth downlink throughput over a 6.4-year simulation period. With the network analysis as our objective function, we implemented a modular genetic algorithm to optimize the HRN satellites' orbits for maximal Mars–Earth data throughput. We obtained average Mars–Earth downlink data rates of over 146 Mbps, 294 Mbps, 568 Mbps, and 852 Mbps for HRNs deployed using 25, 50, 75, and 100 Starship launches respectively. These throughputs represent an up to 400-times increase over the data rates achievable with the current Direct-to-Earth communications paradigm, surpassing NASA's stated downlink capacity objectives for Earth-Mars communications. The HRN uses only flight-proven RF communications systems and offloads the oversubscribed Deep Space Network. This modular HRN design and optimization framework can be extended for communications with multiple planetary bodies and libration points simultaneously, laying the foundations for a future interplanetary internet.

ACKNOWLEDGEMENTS

J. Pénot was supported by the National Science Foundation Graduate Research Fellowship under Grant No. 2024377105. The authors acknowledge the MIT SuperCloud and Lincoln Laboratory Supercomputing Center for providing HPC resources that have contributed to the research results reported within this paper.

REFERENCES

- [1] M. D. Johnston, "Scheduling NASA's Deep Space Network: Priorities, Preferences, and Optimization," in *International Conference on Automated Planning and Scheduling*. Guangzhou, China: JPL Open Repository, 2020. [Online]. Available: <https://dataverse.jpl.nasa.gov/file.xhtml?fileId=66710&version=3.0>
- [2] C. Shouraboura, D. Tran, and M. D. Johnston, "Prioritization and Oversubscribed Scheduling for NASA's Deep Space Network," Jun. 2016, nTRS Author Affiliations: Jet Propulsion Laboratory (JPL), California Institute of Technology (CalTech) NTRS Meeting Information: ICAPS 2016 - SPARK Workshop; 2016-06-13 to 2016-06-15; undefined NTRS Report/Patent Number: JPL-CL-16-2214 NTRS Document ID: 20190033614 NTRS Research Center: Jet Propulsion Laboratory (JPL). [Online]. Available: <https://ntrs.nasa.gov/citations/20190033614>
- [3] W. J. Weber, R. J. Cesarone, D. S. Abraham, P. E. Doms, R. J. Doyle, C. D. Edwards, A. J. Hooke, J. R. Lesh, and R. B. Miller, "Transforming the deep space network into the Interplanetary Network," *Acta Astronautica*, vol. 58, no. 8, pp. 411–421, Apr. 2006. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0094576506000129>
- [4] J. Milton, "Scenario-based uncertainty quantification for deep space optical communications," Thesis, Massachusetts Institute of Technology, Sep. 2023, accepted: 2024-04-17T21:08:58Z. [Online]. Available:

<https://dspace.mit.edu/handle/1721.1/154177>

- [5] NASA GSFC, “CIS Capability Studies IV: Lunar Trunkline Communication Architectures and Mars End-to-End Communication Service Architectures (NextSTEP-3: Appendix C),” NASA GSFC, Request for Proposals (RFP), Jul. 2025. [Online]. Available: <https://sam.gov/workspace/contract/opp/833efc09ce41401d95f52069065082e6/view>
- [6] K.-M. Cheung, H. Xie, C. Lee, P. Carter, W. Jun, and G. Lightsey, “Deep Space Relay Architecture for Communications and Navigation,” in *2023 IEEE Aerospace Conference*, Mar. 2023, pp. 1–19, iSSN: 1095-323X. [Online]. Available: <https://ieeexplore.ieee.org/document/10115636/?arnumber=10115636>
- [7] J. Breidenthal, “The merits of multi-hop communication in deep space,” in *2000 IEEE Aerospace Conference. Proceedings (Cat. No.00TH8484)*, vol. 1, Mar. 2000, pp. 211–222 vol.1, iSSN: 1095-323X. [Online]. Available: <https://ieeexplore.ieee.org/abstract/document/879389>
- [8] M. Knapp, L. Paritsky, E. Kononov, and M. M. Kao, “Great Observatory for Long Wavelengths (GO-LoW) NIAC Phase I Final Report,” Apr. 2024, arXiv:2404.08432 [astro-ph]. [Online]. Available: <http://arxiv.org/abs/2404.08432>
- [9] J. Pénot and H. Balakrishnan, “Heliocentric Mutually Orbiting Groups: Parametrization, Phasing Maneuvers, and Orbital Insertion,” in *2025 AAS/AIAA Astrodynamics Specialist Conference*, Boston, MA, Aug. 2025, forthcoming in proceedings. [Online]. Available: https://www.researchgate.net/publication/394485873_Heliocentric_Mutually_Orbiting_Groups_Parametrization_Phasing_Maneuvers_and_Orbital_Insertion
- [10] E. Berger, “Elon musk just gave another mars speech—this time the vision seems tangible,” <https://arstechnica.com/space/2024/04/elon-musk-just-gave-another-mars-speech-this-time-the-vision-seems-tangible/>, 2024, accessed: March 26, 2025.
- [11] SpaceX, “Starship: Service to earth orbit, moon, mars and beyond,” 2025, accessed: March 26, 2025. [Online]. Available: <https://www.spacex.com/vehicles/starship/>
- [12] E. Musk, “Making humans a multi-planetary species,” *New Space*, vol. 5, 2017.
- [13] T. Barber, F. Krug, K. Renner, T. Barber, F. Krug, and K. Renner, “Final Galileo Propulsion System in-flight characterization,” in *33rd Joint Propulsion Conference and Exhibit*. Seattle, WA, U.S.A.: American Institute of Aeronautics and Astronautics, Jul. 1997. [Online]. Available: <https://arc.aiaa.org/doi/10.2514/6.1997-2946>
- [14] “Starship Users Guide,” 2020. [Online]. Available: https://newspaceconomy.ca/wp-content/uploads/2022/10/starship_users_guide_v1.pdf
- [15] A. J. Alqaraghuli, A. Singh, and J. M. Jornet, “Novel CubeSat Combined Antenna Deployment and Beam Steering Method Using Motorized Rods for Terahertz Space Networks,” in *2021 IEEE Conference on Antenna Measurements & Applications (CAMA)*. Antibes Juan-les-Pins, France: IEEE, Nov. 2021, pp. 560–561. [Online]. Available: <https://ieeexplore.ieee.org/document/9703480/>
- [16] “Spacecraft,” 2024. [Online]. Available: <https://www.jpl.nasa.gov/press-kits/psyche/mission/spacecraft/>
- [17] J. Taylor, D. K. Lee, and S. Shambayati, “Mars Reconnaissance Orbiter,” in *Deep Space Communications*. John Wiley & Sons, Ltd, 2016, pp. 193–250. [Online]. Available: <https://onlinelibrary.wiley.com/doi/abs/10.1002/9781119169079.ch6>
- [18] W. A. Imbriale, *Large Antennas of the Deep Space Network*, 1st ed. Wiley, Jan. 2003. [Online]. Available: <https://onlinelibrary.wiley.com/doi/book/10.1002/0471728497>
- [19] “NASA Deep Space Network,” Jun. 2025, page Version ID: 1297646155. [Online]. Available: https://en.wikipedia.org/w/index.php?title=NASA_Deep_Space_Network&oldid=1297646155
- [20] D. Vallado, *Fundamentals Of Astrodynamics And Applications*, 4th ed. Microcosm Press, 2013. [Online]. Available: <http://archive.org/details/FundamentalsOfAstrodynamicsAndApplications>
- [21] R. C. LaBelle and D. J. Rochblatt, “Calibration and performance measurements for the nasa deep space network aperture enhancement project (daep),” *Acta Astronautica*, vol. 147, pp. 37–47, Jun. 2018. [Online]. Available: <https://linkinghub.elsevier.com/retrieve/pii/S009457651731679X>
- [22] J. H. Yuen, *Deep Space Telecommunications Systems Engineering*, ser. Applications of Communications Theory. NASA JPL, 1983. [Online]. Available: <https://descanso.jpl.nasa.gov/dstse/DSTSE.pdf>
- [23] Dong K Shin, “Frequency and Channel Assignments,” Jet Propulsion Laboratory, Tech. Rep. DSN No 810-005, 201, Rev. C, Dec. 2014. [Online]. Available: <https://deepspace.jpl.nasa.gov/dsndocs/810-005/201/201C.pdf>
- [24] R. Duan and S. Pettie, “Fast Algorithms for (max, min)-Matrix Multiplication and Bottleneck Shortest Paths,” in *Proceedings of the 2009 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, ser. Proceedings. Society for Industrial and Applied Mathematics, Jan. 2009, pp. 384–391. [Online]. Available: <https://epubs.siam.org/doi/abs/10.1137/1.9781611973068.43>
- [25] L. Buddrus, “Bottleneck Problems in Graphs,” 2022, master’s Thesis. [Online]. Available: <https://www.mi.fu-berlin.de/inf/groups/ag-ti/theses/download/Buddrus22.pdf>
- [26] C. H. Acton, “Ancillary data services of NASA’s Navigation and Ancillary Information Facility,” *Planetary and Space Science*, vol. 44, no. 1, pp. 65–70, Jan. 1996. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/0032063395001077>
- [27] C. Acton, N. Bachman, B. Semenov, and E. Wright, “A look towards the future in the handling of space science mission geometry,” *Planetary and Space Science*, vol. 150, pp. 9–12, Jan. 2018. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0032063316303129>
- [28] O. Abdelkhalik, “Hidden Genes Genetic Optimization for Variable-Size Design Space Problems,” *Journal of Optimization Theory and Applications*, vol. 156, no. 2, pp. 450–468, Feb. 2013. [Online]. Available: <https://doi.org/10.1007/s10957-012-0122-6>

- [29] “ga - Find minimum of function using genetic algorithm - MATLAB,” 2024. [Online]. Available: <https://www.mathworks.com/help/gads/ga.html>
- [30] “fmincon - Find minimum of constrained nonlinear multivariable function - MATLAB,” 2024. [Online]. Available: <https://www.mathworks.com/help/optim/ug/fmincon.html>
- [31] A. Reuther, J. Kepner, C. Byun, S. Samsi, W. Arcand, D. Bestor, B. Bergeron, V. Gadepally, M. Houle, M. Hubbell, M. Jones, A. Klein, L. Milechin, J. Mullen, A. Prout, A. Rosa, C. Yee, and P. Michaleas, “Interactive supercomputing on 40,000 cores for machine learning and data analysis,” in *2018 IEEE High Performance Extreme Computing Conference (HPEC)*. IEEE, 2018.

BIOGRAPHY



Jules Pénot received his B.S. in Mechanical Engineering with an Aerospace minor from Caltech in 2024. He is currently a graduate student pursuing an S.M. and Ph.D. in MIT’s Department of Aeronautics and Astronautics. His doctoral research focuses on the optimal design of distributed satellite relay networks for interplanetary and cislunar communications.



Hamsa Balakrishnan is the William E. Leonhard (1940) Professor of Aeronautics and Astronautics at the Massachusetts Institute of Technology. Her research is in the design, analysis, and implementation of control and optimization algorithms for cyber-physical infrastructures, with an emphasis on aerospace systems.