Cycler Orbits and the Solar System Pony Express

Marc Sanchez Net
Jet Propulsion Laboratory
California Institute of Technology,
Pasadena, California, 91109
marc.sanchez.net@jpl.nasa.gov

Etienne Pellegrini
Jet Propulsion Laboratory
California Institute of Technology,
Pasadena, California, 91109
pellegrini@jpl.nasa.gov

Joshua Vander Hook
Jet Propulsion Laboratory
California Institute of Technology,
Pasadena, California, 91109
hook@jpl.nasa.gov

Abstract—In this work, we explore the concept of a secondary “data mule” consisting of a small satellite used to ferry data from a Mars mission to Earth for downlink. The concept exploits the fact that two nearby optical communicators can achieve extremely high data rates, and that a class of trajectories called “cyclers” can carry a satellite between Mars and Earth regularly. By exploiting cycler orbits, the courier needs minimal onboard propulsion.

However, cycler orbits have long periodicity, as it can take years for the satellite, Mars, and Earth to repeat their relative geometry. Therefore, we propose the use of a network of such cycler “couriers” on phase-shifted trajectories to achieve a regular cadence of downlink trips. We design a series of search and optimization steps that can output a set of trajectories that at first approximation have low onboard propulsion requirements and can be used for any regular logistics network to and from Mars, then derive the link budget for proximity optical communications to show that this network can ferry large amounts of data.

TABLE OF CONTENTS

1. INTRODUCTION.....	1
2. CONCEPT OVERVIEW	2
3. CYCLER ORBIT SELECTION	2
4. CROSS-LINK ANALYSIS	4
5. RESULTS AND DISCUSSION.....	7
6. CONCLUSION	8
REFERENCES	9

1. INTRODUCTION

With the recent resurgence of interest in manned missions, solar system logistics problems are catching the spotlight again. The oldest and best-studied logistics problem is actually that of data management, and the Deep Space Network (DSN) has had enormous responsibility and incredible successes in shepherding data from the remote reaches of the solar system.

The DSN relies on a terrestrial network of antennas to receive transmissions from deep-space transmitters. Through a series of technological upgrades and continuing development, the DSN is expected to grow in capacity over time.

However, the capability of the DSN to receive data from remote spacecraft is limited by a fundamental physical law that scales data rate by $\frac{1}{d^2}$ as distance, d , increases. This well-known relationship is the second tyrannical equation (after the rocket equation [1]), which limits our ability to explore the solar system. When compared to the immense size of data

gathering possible with modern instruments, we will always be able to gather more data than we can return to Earth. Scientific discovery, therefore, is constantly data-starved.

However, the converse of the $\frac{1}{d^2}$ law is that if a transmitter is capable of achieving even a modestly good data rate to a terrestrial receiver (large d), then it *could* be able to achieve an enormous data rate to a *nearby* receiver (as $d \rightarrow 0$). If that nearby asset could receive a massive data transfer and then transit to Earth, perhaps more data could be carried by ballistics than could be carried by electromagnetic waves—over time. This echoes the quote by Andrew Taunnenbaum, “Never underestimate the bandwidth of a station-wagon full of tapes hurtling down the highway” [2] and is reminiscent of the times when film canisters were literally dropped from orbital observers for pickup [3]. There was a time not long ago that if you downlinked data to a remote ground station (e.g., McMurdo in Antarctica), they would record it and then ship it back to the United States (using an actual ship). A more evocative analogy is suggested in the title, the Pony Express [4].

In this paper, we take the first steps to explore the concept of a solar-system-wide data-focused *courier* network. To limit the scope, we focus first on Mars, which has arguably the highest data transfer demands to date.

We consider a network of satellites on *cycler* orbits, that are periodic orbits which regularly visit two celestial bodies (in our case, Earth and Mars). Cycler orbits became of interest as a means to transport humans between Earth and Mars for low propellant costs [5], [6]. However, the complex requirement of performing hyperbolic rendezvous has made them impractical for human travel. Our proposed concept alleviates this problem since the data can be transferred to and from the couriers during Mars and Earth flybys, eliminating the need for rendezvous. Cycler orbits have also been studied in the context of planetary moon exploration, such as for Jovian and Saturn moon tours [7]. In the context of cyclers, a “visit” of flyby means an approach within a few thousand to tens of thousands of kilometers. A possible single-cycler conops is shown in Figure 1, and three example cycler orbits are shown in Figure 6. We propose that each courier receives a massive *cross-link* from a relay near Mars and then transits to Earth where it can be downlinked quickly.

We trade off the number of couriers against the total amortized data rate per year. Our data rates take into account mainly the ephemeris approximations and the communication technologies we anticipate will be available in the next decade. Other considerations are mentioned and discussed in Section 6 but are left to future work.

Figure 1: Possible concept of operations for a single courier and observer launched concurrently. The courier is inserted into a cycler orbit, which passes near Earth and Mars with regular cadence, but only once per 2-4 years. Thus, multiple couriers are required to achieve yearly or near-yearly downlink. In this paper, we discuss a multi-courier network and resulting data volumes transferred.

2. CONCEPT OVERVIEW

As shown in Figure 1, we propose a set of satellites (the couriers) forming the backbone of a downlink network between Earth and a planetary body (Mars, in our working example) by exploiting nearly ballistic cycler orbits [5] and high-data-rate optical communications.

A cycler orbit is a trajectory that passes close to two planets regularly with periodicity proportional to their synodic periods. The courier satellites could be launched together with a primary mission (the observer)². In this case, the couriers detach after launch, possibly during Earth-Mars transfer.

The observer continues on to Mars to conduct surveys of the planet using high data rate instruments. The couriers use solar-electric propulsion to get into Mars-Earth cycler orbits. The orbits are phase-shifted such that once every 6-8 months a courier passes by Mars and then to Earth.

By routinely passing near enough for high data rate communications, the courier can retrieve from the observer, and shepherd to Earth, at least a petabit per pass from Mars to Earth. By placing couriers in phase-shifted cycler orbits, the network can achieve an average of multiple-petabit-per-year data dumps.

To show the feasibility of this concept, we will present various cycler networks and discuss the trade-offs of each. We proceed as follows. First, we enumerate a large corpus of nearly ballistic cycler orbits that begin in 2030. Then we formulate an orbit selection problem as a discrete *Set Cover* instance so the required downlink windows can be serviced by at least one cycler. Then we evaluate the cross-link capabilities of likely optical communications technologies in the 2030-2040 time frame. Taken together, we can project the number of visits per year, and the total data carried by each

² If a single-launch stretches credibility, then it may be reasonable to assume that a capable transmitter exists at Mars, as is required to reach the DSN growth projections for the coming decades.

visiting satellite, to show the total amortized data rate and data volume per year for comparison to the DSN. Since cycler orbits repeat indefinitely (ideally), the network can exist in perpetuity, up to the limits of our ability to engineer robust spacecraft.

3. CYCLER ORBIT SELECTION

In this section we formulate the network of cycler trajectories needed to achieve a desirable downlink cadence. A large number of potentially-ballistic cycler trajectories can be computed in a patched-conics, two-body problem using the Star software available at JPL. A preliminary enumeration of cycler orbits yielded 1000s of candidates with periods between 2 and 6 years. Those trajectories can later be optimized in a more realistic model in order to determine the real ballistic cycler trajectories, using for instance the Cosmic optimal control software within the Monte suite [8], or using methods such as the ones described in Reference [9]. The final step consists in selecting the required number of cycler orbits to use for the couriers. The cadence of Mars visits and latency before Earth return are both optimized against the minimal number of independent spacecraft required.

Enumeration of Orbits

The preliminary enumeration of potential cycler orbits is done using the Star software available at JPL. Star is a broad-search tool that produces database of trajectories in the patched-conics model. In order to limit the search to a manageable trade space, a number of constraints have to be enforced. The current enumeration uses the following constraints:

- Launch date: 2030 - 2034. The launch date in Star corresponds to the insertion into the cycler orbit. The couriers can still potentially be launched at the beginning of the period (2030), and wait in the system until the appropriate time for injection.
- Time-of-flight: 15 years.
- Number of cycles: 3. Each of the trajectories returned

Figure 2: Full dataset returned by Star with Mars-Earth transit duration constrained to be shorter than 12 months

by the broad-search has exactly 3 Mars visits. Most cyclers found have a period equal to 2 Earth-Mars synodic periods, or 4.3 years, and can therefore have 3 visits of Mars during the assumed time-of-flight. However, some cyclers found have shorter periods and therefore shorter time-of-flights.

- Sequence of flybys in a cycle: A single cycle is either an Earth-Mars(-Earth) (EM) trajectory or an Earth-Earth-Mars(-Earth) (EEM) trajectory. Different combinations are left to explore in order to exhaust the search space.
- Maximum allowable ΔV at a flyby: 10 m/s. This preliminary enumeration is focused on near ballistic cyclers. Limiting the ΔV in Star allows only near ballistic solutions in the patched-conics model. The optimization of the trajectories in a more realistic model is necessary to obtain the actual ΔV cost.

Over 7900 trajectories were generated by Star and separated into two batches. The first batch is constrained to have every Mars-Earth (ME) leg shorter than 12 months in duration in order to reduce the risk to the data while flying in deep space. The second, a superset of the first, is constrained to have ME transits shorter than 18 months. The 12 months dataset is shown in Figure 2. Three examples of different cycler orbits (with different sequences of flybys) are shown in Figure 6. Note that Star can return “hybrid” cyclers, which use an EM cycle followed by an EEM cycle, and viceversa.

With the current search parameters for the enumeration, all returned cyclers need at most 60 m/s once injected into their trajectories. Considering our current strategy of having a single launch into a low-energy escape orbit, the necessary ΔV for injection into the cycler trajectories ranges from 0.5 to 4 km/s and varies according to the departure V_∞ at Earth returned by Star. Consequently, in a first-order approximation, the cyclers are considered to be ballistic, and the departure V_∞ could be used as a proxy for ΔV in the selection process.

Network Design Using Set Cover

Given the large number of cyclers found, we wish to select a small number that can be phased to provide a regular but well-spaced cadence of Earth downlinks. It is not considered a good solution if multiple cyclers visit Mars then Earth in quick succession, followed by a decade-long wait, for example.

We downselected the cyclers for membership in the network

by formulating a Set Cover problem [10], [11]. Set Cover is a well-known discrete optimization problem that can be solved efficiently using Integer Linear Programming [11]. A solution takes the form of a minimum-weight selection of sets, such that the union of all items in any set is contained in the union of the selected sets. In our case, we used cyclers (which are sets of ME transits over time) to cover the required time periods.

Specifically, we divided the total time period of 18 years (15 years of operation plus offset in launch time) into n *downlink windows* of a fixed span of days to match the frequency of visits desired per year f^* . For example, $f^* = 1$ implies $n = 18$, and $f^* = 2$ nets 36 (assuming even division—rounding can cause an extra bin).

No downlink was possible in the first three years and so all windows in this time were discarded, leaving 15 years. To ensure a solution was always produced, we further removed any window that was not feasible for any cycler in our dataset. Thus, the network can only achieve an actual recurring frequency $f = \#visits/15$. A more exhaustive set of cyclers to select from would fill these gaps. Each window was considered “covered” if a selected courier visited Earth during that date range *and had just come from Mars*. Anecdotaly, we had hoped to achieve a yearly downlink ($f \approx 1.0$) with a small number of cyclers (e.g., 4-8).

The GLPK solver was used to solve an Integer Linear Program as follows[12]. GLPK can use the dual-simplex method to efficiently solve this ILP, as was mentioned in[11]. Let there be n downlink windows and k cyclers to select from.

$$\begin{aligned} &\text{Minimize}_x c^T x \\ &\text{s.t., } Ax \geq b \end{aligned} \quad (1)$$

Where,

- x, c are size $1 \times k$
- b is size $1 \times n$
- A is size $n \times k$
- $x_i = 1$ if cycler i was selected, 0 otherwise
- $A_{i,j} = 1$ if j^{th} cycler visits Earth during window i
- $b_i = 1$ if *any* cycler visits Earth during window i
- $c_i = 1$ for all i

In our problem, n ranged from 10s to 100s, k was around 8000, and GLPK returned solutions nearly instantaneously on a modern laptop computer. Note that we have ignored the cost weighting c to just select the minimum *number* of cyclers, but could easily use other metrics such as ΔV .

In Section 5 we show the full solution with discussion. Some interesting trade points are shown in Table 1. Note that as few as 6 couriers are required to average a yearly downlink from the initial dataset we explored, and 3 can achieve bi-annual. Conveniently, ESPA rings can lift 6 small satellites as secondary payloads [13], but a major direction for future work could be to explore secondary payload options so that 3-5 couriers could be added to the network with primary missions. The solutions are very sensitive to n , and the gaps become much more pronounced at higher numbers of cyclers, so it is challenging to show the scaling of frequency against number of cyclers. However, by varying f from .5 to 24, we achieved the following results, which show a rough scaling law, and decent results for low numbers of cyclers.

Table 1: Example Downlink Cadence

Max 18-Month ME Transit		
min # cyclers	Achievable n	Achievable f
03	9	0.5
04	11	0.7
05	13	0.9
06	15	1.0
07	20	1.3
10	28	1.9
13	34	2.3
14	38	2.5
16	43	2.9
20	53	3.5
29	74	4.9
35	93	6.2
40	102	6.8
50	130	8.7

We assume each courier receives its data payload when it flies by Mars on the way to Earth. The observer (or a relay) at Mars uses an optical link to offload data, and we bound the data that can be transferred to a cycler next.

4. CROSS-LINK ANALYSIS

For the purposes of this paper, we focus solely on the proximity link at Mars between the courier and the observer. We assume that both of them carry the same optimal terminal, which is built as an evolution of the NASA's Lunar Laser Communication Demonstration (LLCD) terminal, and utilizes DPSK for data transmission.

Our preliminary analysis is based on a few randomly-selected but typical ballistic cyclers (c.f. S1L1 [14]). For each of them, we propagate the portion of the cycler orbit close to the Mars flyby (from approximately -50000 km to 50000 km) and obtain range and velocity profiles as a function of time. Figure 3 shows the obtained results for the 10 randomly-selected cyclers. We observe that notable differences are confined to the period of time near to closest approach with Mars.

To limit the amount of time the observer is tied communicating with a courier, we assume that it never transmits for more than 2.5 consecutive hours. Therefore, and given that

the encounter between the observer and the courier lasts three days approximately (assuming the ± 50000 km definition), we constrain the optimal corss-link to only operate for about 3.5% of the encounter. During the remaining 96.5% of the time, the observer is free to resume normal operations (i.e., science gathering or communicating with the DSN). An alternative formulation could transmit for a day or two to reduce the peak data rate requirement in the next section, and allow regular DSN operation on all other days in the year. Either way, the impact to DSN is constrained to be very small.

Optical Technology

Since the launch date of the proposed system is uncertain and well in the future, we need to make reasonable assumptions about optical technology progression over the coming decades. To that end we first surveyed similar studies done in the past (see [15]) and, based on their insights, defined the following three baseline technological states:

- (T1) Both the courier and observer are equipped with an LLCD-like terminal with a 0.1cm aperture and a 0.5W laser. The system electronics can support rates up to tens of Gbps.
- (T2) The observer is equipped with an optical terminal that is heritage from the Deep Space Optical Communications (DSOC) system. It has a 0.22 cm aperture and a 0.5 W laser³. The courier is still equipped with an LLCD-like terminal to save mass. Finally, the system electronics can support up to hundreds of Gbps.
- (T3) Both the courier and the observer carry a 0.22 cm optical terminal heritage from DSOC. Additionally, the observer is able to output 1 W of laser power, and its electronics can support up to 1Tbps.

Next, we consider the ability of the optical communication system to adapt its data rate as the distance between the courier and the observer decreases. In that sense, two limiting cases are conceivable: Fixed rate link vs. perfectly-adaptable data rate. Using the former will result in a lower bound for the total returnable data rate over a flyby, while the latter will result in an upper bound. For this paper, we will assume that the optical terminal implements a near-optimal adaptive coding and modulation scheme, and thus we will report only the upper bound values.

Details of Link Budget

We consider a DPSK optical terminal with a soft-decision decoder using the 1/2 LDPC code defined in the DVB-S2 standard. To approximate its performance over an inter-satellite link, we first estimate the DPSK channel capacity and then assume pessimistically a 2 dB coding gap. Moreover, we also assume that a further 3 dB loss is incurred in the optical fiber couplings at the receiver.

For the optical heads at both the transmitter and receiver, we first assume a 7% relative depointing. This value is then used to estimate the pointing losses assuming that the telescope has no obscuration and that the filter illuminates the entire aperture with uniform plane-wave illumination [16]. Additionally, we also aslso assume, somewhat pessimistically, a 35% optical transmission efficiency (values up to 50% have been reported in the literature [17]).

Table 2 presents the optical link budget assuming technology scenario T1 and a nominal distance of 10000 km. Similar link budgets have also been computed for the other two tech-

³The DSOC laser is currently developed to output 4 W of power. However, this is considered excessive burden for a close-range link.

Figure 3: Flyby Characteristics

nology scenarios (T2 and T3) at the same reference distance. Then, data rates as a function of range between the observer and courier are scaled quadratically as follows:

$$R(t) = R_{\text{ref}} \left(\frac{d_{\text{ref}}}{d(t)} \right)^2, \quad (2)$$

where $d_{\text{ref}} = 10000$ km, $d(t)$ is the range at a given point in time, and R_{ref} is the data rate obtained from the link budget calculation and takes values 16 Gbps for T1, 150 Gbps for T2 and 740 Gbps for T3.

Data Rate and Data Volume Profiles

Figure 4 shows the data rates achievable between the courier and the observer as a function of the technology assumed and the specific flyby trajectory. We can see that results are significantly more sensitive to our technological assumptions than to the specific flyby trajectory, and lead to approximately one order of magnitude differences. Furthermore, we also note that in the most aggressive case the link would need to support up to 1 Tbps, a rate that imposes stringent requirements on the electronics of the spacecraft.

Figure 5 shows the resulting data volume returnable over a 2.5h encounter between the courier and the observer as a function of technology available and flyby trajectory. Once again we note that the effect of the technology assumptions is significantly larger and can be summarized as follows: With T1, approximately 100 Tbit/flyby can be returned; with T2, up to 1 Pbit/flyby can be sent; and finally, with T3 up to 5 Pbit/flyby can be exchanged.

Given that our proposed concept enhances the yearly data volume returnable from Mars, it is informative to compare this increase with current and future DSN capabilities. In that sense, Table 3 estimates the returnable yearly data volume as a function of the ground station (GS) support and flight terminal (FLT) capabilities. It can be observed that with just Ka-band support, our system can approximately double the amount of data that can be returned from Mars with just one courier flyby per year. Similar results are also true if the DSN incorporates optical deep space links in its operations, except for the extreme case where the flight terminal can deliver 100W of laser power using 4 Wavelength Division Multiplexing. Note that the numbers in Table 3 are based

Figure 4: Data Rate for Different Flybys and Technology Assumptions. Color indicates technology assumed and line style indicates flyby.

Figure 5: Data Volume for Different Flybys and Technology Assumptions

Table 2: DPSK Link Budget for T1

LINK PARAMETERS			LINK BUDGET		
Link distance	10000	km	Tx Avg Power	26.99	dBm
Inter-satellite Link	Yes	-	Tx Photons/Pulse	2.47E+08	-
TRANSMITTER			Tx Antenna Gain	106.14	dB
			Tx Array Gain	0.00	dB
Modulation Type	DPSK		Tx Transmission Loss	-4.56	dB
Modulation Order	2		Tx Pointing Loss	-0.09	dB
			EIRP	128.48	dBm
Tx Wavelength	1.55	μm	Isotropic Space Loss	-278.18	dB
Photon Energy	1.28e-19	J/photon	Atmospheric Loss	0.00	dB
Tx Avg. Power	0.5	W	Path Loss	-278.18	dB
Tx Data Rate	1.58e+10	bps			
Tx Data Rate	15.76	Gbps			
			Irradiance at Rx Aperture	5.61e-6	W/m ²
Tx Aperture Diameter	0.10	m	Rx Antenna Gain	106.14	dB
Tx FOV	4.07	arcsec	Rx Array Gain	0.00	dB
Tx Array Size	1	Apertures	Rx Transmission Loss	-4.56	dB
Tx Aperture Efficiency	0.35		Rx Pointing Loss	-0.09	dB
Relative Rx Depointing	0.07		Rx Fiber Coupling Loss	-3.00	dB
Tx Depointing	0.28	arcsec	Total Optical Path Loss	-78.19	dB
Code Rate	1/2		Avg Power at Rx Detector	-51.20	dBm
Coded Symbol Rate	3.15e+10	sps	Photons/Pulse at Rx Detector	3.75	-
Modulated Symbol Rate	3.15e+10	baud			
RECEIVER			Required Photons/Pulse	1.88	-
			Link Margin	3.00	dB
Rx Aperture Diameter	0.10	m			
Rx FOV	4.07	arcsec			
Relative Rx Depointing	0.07				
Rx Depointing	0.28	arcsec			
Rx Optical Transmission	0.35				
Rx Array Size	1	Apertures			
Fiber Coupling Loss	-3.00	dB			
Modulator Type	Soft-Decision				
Capacity at Rate 1/2	2.37	photons/bit			
Capacity at Rate 1/2	1.19	photons/pulse			
Code gap (no fading)	2.00	dB			
Required Photons/Pulse	1.88	-			

on reasonable but simplistic assumptions that do not take into account operational constraints at neither ends of the communication channel. Therefore, they should be considered a first-order of magnitude anchoring point.

Onboard Storage

Recent consumer electronics using NAND Flash cells (specifically 3D NAND configurations) have achieved 0.2 Petabit storage on a small form factor hard drive requiring 5W power during operation (30TB SSDs are available off the shelf at the time of this writing [18]). The largest mass storage unit being used on a JPL mission is the EMIT Mercury SSDR, which uses SLC NAND to provide 440 GB of storage. MRAM, FRAM, and ReRAM technologies have also made significant advancements in recent years [18]. The primary focus of this study, NAND Flash technology, offers the most promising balance of speed, density, and power. Reliability, however, poses several issues. TLC (and MLC to a lesser extent) 3D-NAND devices have low Program-Erase (P/E) limits, on the order of 3000, but this is not an issue for infrequent data transfers like we propose. TLC NAND is especially sensitive to Total Ionizing Dose, radiation-induced soft errors (Single bit upsets, multi-bit upsets, and functional interrupts), as well as some radiation-induced hard errors (stuck bits, dielectric breakdown) [19], [20], [21]. Thank-

fully, for the purpose of this study, all of these effects can be mitigated with increased redundancy (i.e., increased weight and power).

If we assume a memory density similar to that of 2019 commercial cutting edgwhich is a fair assumption given that NAND devices used in space lag commercial cutting edge by about 10 years⁴ and we assume a redundancy factor to account for the aforementioned reliability issues, we may now get a rough picture of the expected weight for our Pb-scale Mass Storage Unit. If we assume a gram/bit of 6.5x10-16 for our F-NAND (based on Samsung 30TB V-NAND SSD), our 1Pb MSU comes to 0.7kg. Even with a reliability factor of 40%, this still only comes out to 1kg. This is a very workable weight for the proposed system, and allows for a great deal of flexibility for additional redundancy and error correction systems.

However, read/write speed must be high, given the optical link speeds of 100s of Gbps. The UltraScale+ SoC is currently able to provide up to 32.75 Gb/s for serial transmission [22], so interface bottlenecks are not expected to be too

⁴From EMIT Instruments SSDR and MSL NVCMAM, which use 2009 Micron SLC NAND and 2002 Micron SLC NAND technologies respectively.

Table 3: Current and Future DSN Capabilities

Reference year	2020	2025-2030	>2030	>2030	2020	2030	>2030
Frequency band	X-band	Ka-band	Ka-band	Ka-band	Optical	Optical	Optical
GS support	1x34m	2x34m	3x34m	7x34m	5m	12m	12m
GS scheduling	90%	90%	30%	15%	30%	90%	90%
FLT size	3m	3m	3m	3m	22cm	40cm	50cm
FLT power	1x100W	1x180W	1x500W	1x500W	1x4W	1x10W	4x20W
FLT capability	1x MRO	x5 MRO	x14 MRO	x14 MRO	x1 DSOC	x8 DSOC	x100 DSOC
Data Rate @ 1.5AU	500kbps	20Mbps	85Mbps	200Mbps	850kbps	36Mbps	490Mbps
Yearly Data Volume	30Tbit	1.3Pbit	1.8Pbit	2.1Pbit	20Tbit	2.3Pbit	30.4Pbit

significant of a hurdle. The remainder of the study would be concerned with refining requirements to bus and memory architectures in terms of reliability, speed, and data volume against size, weight, and power (SWaP). If required, and as discussed in 3.3.2.3, different communication regimes can reduce the R/W speed requirement by 20-30x. However, total data volume and reliability seems, at first approximation, to be feasible.

5. RESULTS AND DISCUSSION

Combining the results presented, we can automatically enumerate a large batch of cycler orbits, then select the set of orbits required to achieve a given “quality of service” in terms of number of visits per decade and cadence of such visits.

We see in Table 4 three networks for 3, 5, and 6 cyclers. The cyclers used for the 3-cycler network are shown on Figure 6. We note some clustering of visit dates, likely caused by the launch constraint (see for example the 6 cycler network has 3 visits in 2044). Note the optimization problem in Section 3 specified *at least* one cycler per window, so duplicates are acceptable. Further work will have to find a network with less clustering. Still, the staggering of visits is not bad, with at least a few months to gather data at Mars before a downlink transit. We can further constrain the problem to eliminate duplicate visits in the same window by setting the inequality constraint in our optimization to be an equality constraint. This shows an expected scaling of 3 visits per cycler added to the network, with about 2 of them being *unique* per cycler.

Table 5 shows a first estimate of the overall network data volumes and amortized data rate (volume divided by mission time). We show how the assumptions on technology from Section 4 affect overall data rate and volume.

On the conservative side, a 100 Tb dump per year—about the minimum enabled with a single launch and most cautious technology—could bring back a Landsat-like survey of Mars with seasonal differencing without significantly affecting DSN. On the upper end, a 6 Pb dump and 2-3 downlinks per year *might* be enabled with secondary payloads on the next few launches to Mars, and would enable 180 Pb per decade with only 5% usage of the Mars transmitter. This corresponds to roughly half a gigabit per second amortized data rate with only 5% of the DSN transmitter used. Given the low overall transmitter usage, this could be considered a good “bang-for-the-buck” technology.

On the other side of the problem, the use of proximity links could mean much smaller transmitter and receivers at both Mars and Earth, as well as possibly lower power requirements

Table 4: Example Cycler Networks

3-Cycler Network / 9 downlinks		
downlink date (Y-M-D)	cyclers #	ME transit
2033-10-29	51	174.16 days
2035-11-25	1	233.62 days
2037-12-19	1141	165.09 days
2038-04-24	51	220.83 days
2040-05-29	1	332.28 days
2042-05-30	51	212.25 days
2044-06-09	1	278.25 days
2046-08-24	1141	295.64 days
2048-09-19	1141	143.25 days

5-Cycler Network / 15 downlinks		
downlink date (Y-M-D)	cyclers #	ME transit
2033-10-30	50	196.16 days
2035-12-21	84	297.57 days
2036-03-08	397	262.41 days
2037-12-19	1141	165.09 days
2038-01-28	1066	229.01 days
2039-12-28	84	364.59 days
2040-06-19	50	270.23 days
2042-04-11	397	265.89 days
2044-05-11	84	358.25 days
2044-07-27	1066	277.19 days
2044-09-20	50	293.25 days
2046-08-24	1141	295.64 days
2046-10-18	397	351.25 days
2048-07-19	1066	105.25 days
2048-09-19	1141	143.25 days

6-Cycler Network / 18 downlinks		
downlink date (Y-M-D)	cyclers #	ME transit
2033-10-11	362	126.20 days
2035-12-21	84	297.57 days
2036-04-19	53	272.33 days
2037-11-07	529	335.18 days
2038-05-07	1102	229.81 days
2039-12-28	84	364.59 days
2040-02-10	53	330.50 days
2040-07-08	362	346.19 days
2040-07-11	612	241.19 days
2042-07-04	529	251.72 days
2042-08-09	1102	257.64 days
2044-04-27	53	281.25 days
2044-05-11	84	358.25 days
2044-07-07	612	239.23 days
2046-06-20	529	132.25 days
2046-08-18	612	171.25 days
2046-09-06	362	179.25 days
2048-08-30	1102	179.25 days

Figure 6: The 3 cycler network. The Mars and Earth visit times are shown for each cycler.

for the same total data volume. The impact to latency would have to be carefully considered, however, as wait times are significant.

A manned mission to Mars would produce much more data than current DSN (even with optical communications) can handle. For example, a single camera on the ISS generates HD video feeds on the order of megabits per second continuously, which regularly has thousands of online viewers.

HD video feeds from landers and orbiters (and eventually manned outposts) could be played back in real time (after downlink) on Earth, exciting the global community about ongoing space exploration and manned mission priorities. We imagine a high-resolution video of the evolution of a Mars outpost would excite the public. A massive data set would allow academic and citizen scientists to view, zoom, filter, and work with terabytes of new imaging, weather, and hyperspectral data per year.

While Mars is somewhat unique in the solar system for the attention it has received from the space exploration community, there is nothing to suggest this concept would not work elsewhere.

6. CONCLUSION

We have proposed a network of small satellites, probably as secondary payloads, that can augment the DSN with large capacity downlink by trading latency for massive volume. Cyclers are often associated with Buzz Aldrin, who proposed using them to transfer goods and astronauts to Mars[6]. In effect, we have proposed that the first cargo these cyclers should carry is data. The Cycler Courier Network evokes images of the old Pony Express that raced across the western United States. With this network, missions may rely on recurring data transfers and gather much more data potentially reducing the need for large antennas on each mission by using the observer (perhaps now as a shared archive) to cache until pickup.

We have focused on the problem of enumerating the network of cyclers, and selecting appropriately spaced orbits. We used a thorough, but assumption-driven approach to calculating cross-link data rates, which bounded the total data transfer over the cycler network.

Several areas of future work are necessary to further refine these results.

For instance, the current analysis does not consider size, weight and power (SWaP) as a constraint in the system. Therefore, there is no penalty for increasing telescope size and laser power, and we have assumed that existing commercial solutions for radiation-hardened data storage and computing will be a good reference point for the courier's capabilities. We have suggested numerous times that the couriers could be secondary payloads. While ESPA rings can deploy up to six 180kg spacecraft [13], we have made no attempt to ensure that a courier would actually fit within that weight class. Related, the use of proximity links could mean much smaller transmitter and receivers at both Mars and Earth for comparable amortized data rate as DSN, as well as possibly lower power requirements. This is a good direction for future work as well.

We have not considered the use of Ka band for this problem. The use of RF vs optical could make the technology useful imminently, but may bring down the total data volume.

Similarly, we did not consider how data will be sent from the courier to Earth. This can be done using traditional RF communications, by establishing a link that runs not only during flyby but also potentially during part of the cruise. Alternatively, data can be returned over a deep space optical link using PPM modulation, a fact that further complicates the design of the optical terminal since now both a coherent and noncoherent modulation scheme need to be supported⁵. In either case, further analysis is needed to understand the SWaP limitations that a courier would impose on the optical terminal (and other communication payloads).

The corpus of cycler orbits generated in Section 3 could be expanded. As noted, there are gap years in the data set, which we believe is a product of our constraints used to guide the search, rather than celestial mechanics. A more liberal search, coupled with a fast selection subroutine would likely produce more favorable downlink cadences.

As a next step, the nearly ballistic patched-conics cycler orbits have to be optimized in the full model. Many of them will necessitate large amounts of Delta V once the influence of other celestial bodies, the oblateness of the Earth and Mars, and other perturbations are taken into account. Next, the trajectories must be constrained to allow insertion into the cycler orbit using solar electric propulsion (SEP), which will allow 2 or more data couriers in cycler orbits from a single launch. SEP allows great flexibility in the choice of orbital

⁵At the time of writing, coherent optical modulations cannot be used in space-to-ground links due to atmospheric effects.

Table 5: Expanded Downlink Volumes for various network sizes and technology assumptions (See Section 4)

# Cyclers	Tech (T1,T2,T3)	Total (Pb)	Ammortized Rate (Mbps)
3	T1	0.90	1.90
3	T2	9.00	19.03
3	T3	54.00	114.16
4	T1	1.10	2.33
4	T2	11.00	23.25
4	T3	66.00	139.52
5	T1	1.50	3.17
5	T2	15.00	31.71
5	T3	90.00	190.26
6	T1	1.80	3.81
6	T2	18.00	38.05
6	T3	108.00	228.31
7	T1	2.00	4.23
7	T2	20.00	42.28
7	T3	120.00	253.68
11	T1	2.80	5.92
11	T2	28.00	59.19
11	T3	168.00	355.15
12	T1	3.20	6.76
12	T2	32.00	67.65
12	T3	192.00	405.89
14	T1	3.80	8.03
14	T2	38.00	80.33
14	T3	228.00	481.99
16	T1	4.40	9.30
16	T2	44.00	93.02
16	T3	264.00	558.09
20	T1	5.30	11.20
20	T2	53.00	112.04
20	T3	318.00	672.25
23	T1	6.20	13.11
23	T2	62.00	131.07
23	T3	372.00	786.40
25	T1	6.80	14.38
25	T2	68.00	143.75
25	T3	408.00	862.51
28	T1	7.30	15.43
28	T2	73.00	154.32
28	T3	438.00	925.93
32	T1	8.80	18.60
32	T2	88.00	186.03
32	T3	528.00	1116.18
35	T1	9.30	19.66
35	T2	93.00	196.60
35	T3	558.00	1179.60
40	T1	10.20	21.56
40	T2	102.00	215.63
40	T3	612.00	1293.76
50	T1	13.00	27.48
50	T2	130.00	274.82
50	T3	780.00	1648.91

parameters (when given enough time). This flexibility can potentially be used to attain a first Earth flyby with the correct conditions for insertion into a cyclor orbit. Flexibility is greater for low-energy orbits, therefore a launch into low C3 escape orbits appears as an ideal candidate for this strategy. Higher energy orbits combined with lunar flybys could also enable this strategy.

Finally, a study of long-term, phrasal growth of a courier

network over time would better position this concept with other solar system logistics problems and better estimate the usefulness.

ACKNOWLEDGMENTS

The research was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration. The authors would like to thank Dr. Damon Landau of the Jet Propulsion Laboratory for his invaluable help in using his software, Star. Also, the authors would like to thank Wilson Parker, of the Jet Propulsion Laboratory, for his in-depth knowledge and help on spacecraft memory technologies.

REFERENCES

- [1] V. Dvornychenko, "The generalized tsiolkovsky equation," 1990.
- [2] A. S. Tanenbaum *et al.*, "Computer networks, 4-th edition," *ed: Prentice Hall*, 2003.
- [3] B. Crothers, M. J. Lanphear, M. B. Garino, M. P. P. Konyha III, and M. E. P. Byrne, "Us space-based intelligence, surveillance, and reconnaissance," 2009.
- [4] G. D. Bradley, *The Story of the Pony Express: An Account of the Most Remarkable Mail Service Ever in Existence, and Its Place in History.* AC McClurg, 1914.
- [5] W. M. Hollister, "Periodic orbits for interplanetary flight," *Journal of Spacecraft and Rockets*, vol. 6, no. 4, pp. 366–369, 1969.
- [6] D. V. Byrnes, J. M. Longuski, and B. Aldrin, "Cyclor orbit between earth and mars," *Journal of Spacecraft and Rockets*, vol. 30, no. 3, pp. 334–336, 1993.
- [7] R. P. Russell and N. J. Strange, "Cyclor trajectories in planetary moon systems," *Journal of Guidance, Control, and Dynamics*, vol. 32, no. 1, pp. 143–157, 2009.
- [8] S. Evans, W. Taber, T. Drain, J. Smith, H.-C. Wu, M. Guevara, R. Sunseri, and J. Evans, "Monte: The next generation of mission design and navigation software," *CEAS Space Journal*, vol. 10, no. 1, pp. 79–86, 2018.
- [9] R. P. Russell and C. A. Ocampo, "Optimization of a broad class of ephemeris model earth-mars cyclers," *Journal of Guidance, Control, and Dynamics*, vol. 29, no. 2, pp. 354–367, 2006.
- [10] G. Andrew, T. Hoffman, and C. Krabek, "On the generalized set covering problem," *CDC, Data Centers Division, Minneapolis*, 1968.
- [11] R. D. Koncal and H. M. Salkin, "Set covering by an all integer algorithm: Computational experience," *Journal of the ACM (JACM)*, vol. 20, no. 2, pp. 189–193, 1973.
- [12] A. Makhorin, "Glpk (gnu linear programming kit), 2000," *B B*, 2014.
- [13] J. R. Maly, S. A. Haskett, P. S. Wilke, E. Fowler, D. Sciulli, and T. E. Meink, "Espa: Eelv secondary payload adapter with whole-spacecraft isolation for primary and secondary payloads," in *Smart Structures and Materials 2000: Damping and Isolation*, vol. 3989. International Society for Optics and Photonics, 2000, pp. 430–439.
- [14] T. T. McConaghy, J. L. Longuski, and D. V. Byrnes,

“Analysis of a class of earth-mars cycler trajectories,” *Journal of spacecraft and rockets*, vol. 41, no. 4, pp. 622–628, 2004.

- [15] J. Breidenthal, H. Xie, C.-W. Lau, and B. MacNeal, “Space and earth terminal sizing for future mars missions,” in *2018 SpaceOps Conference*, 2018, p. 2426.
- [16] W. K. Marshall, “Transmitter pointing loss calculation for free-space optical communications link analyses,” *Applied optics*, vol. 26, no. 11, pp. 2055–2057, 1987.
- [17] J. Breidenthal, H. Xie, and L. Clare, “Optical ground segment performance summary,” *The Interplanetary Network Progress Report*, vol. 42, p. 205, 2016.
- [18] Simone Bertolazzi, “MRAM technology and market trends technology and market analyst,” in *Flash Memory Summit Proceedings 2019*, 2009.
- [19] Farokh Irom and Gregory R. Allen, “Radiation tests of highly scaled, high-density, commercial, nonvolatile nand flash memories,” Jet Propulsion Laboratory, Tech. Rep., December 2012.
- [20] F. Irom and D. N. Nguyen, “Single event effect characterization of high density commercial nand and nor nonvolatile flash memories,” *IEEE Transactions on Nuclear Science*, vol. 54, no. 6, pp. 2547–2553, 2007.
- [21] G. Cellere, A. Paccagnella, A. Visconti, M. Bonanomi, and S. Beltrami, “Single event effects in nand flash memory arrays,” *IEEE transactions on nuclear science*, vol. 53, no. 4, pp. 1813–1818, 2006.
- [22] Xilinx, “Ultrascale architecture gty transceivers user guide,” Xilinx, Tech. Rep., September 2007.

Joshua Vander Hook received a Ph.D. in Computer Science in 2015 from the University of Minnesota. He is currently a distributed systems technologist in the Autonomous Systems Division of NASA Jet Propulsion Laboratory. He is most interested in navigation, experiment design, and activity planning problems for both space flight and military applications.

Marc Sanchez is a telecommunications engineer in the Communication Architectures and Research Section at JPL. His research interests include delay tolerant networking and its impact on distributed applications such as computational task sharing, spacecraft constellation management, as well as design of space communication systems in challenged environments such as the surface of the Moon or Mars. Marc received his PhD in 2017 from MIT, and also holds degrees in both telecommunications engineering and industrial engineering from Universitat Politècnica de Catalunya, Barcelona.

Etienne Pellegrini is a Mission Design Engineer in the Mission Design and Navigation Section of the Jet Propulsion Laboratory, where he works on tour design for Europa Clipper. His research interests include astrodynamics, nonlinear optimization, and high-performance computing. Etienne received his Ph.D. in 2017 from the University of Texas at Austin, where he worked on the development of a Differential Dynamic Programming trajectory optimization software.