Energy-Aware Data Routing for Disruption Tolerant Networks in Planetary Cave Exploration

Tiago Vaquero and **Martina Troesch** and **Marc Sanchez Net** and **Jay Gao** and **Steve Chien**
Jet Propulsion Laboratory
California Institute of Technology
4800 Oak Grove Drive
Pasadena, CA 91109

Abstract

Exploring subsurface structures with autonomous robots is of growing interest in the context of planetary caves studies. Communication in these environments can change rapidly as assets move around, which can complicate coordination among multiple assets. Limited lifetime must also be accounted for when exploring these subsurface structures, because it is likely that recharging the batteries of the robots will not be possible. The combination of uncertain communication and limited mission duration suggests that accounting for energy when transmitting data out of cave-like structures would be beneficial to mission success. Therefore, in this paper we investigate different energy-aware data routing strategies for multi-robot scenarios where asset lifetime is limited and benchmark their performance in a simulation environment.

1 Introduction

Planetary caves have increasingly raised interest from the planetary science and robotics communities for their enviromental and structural potential to host human habitats (Boston et al. 2003). Before settling humans in Moon or Mars caves, such unknown subterranean structures will need to be well studied and mapped, potentially by teams of autonomous robots (Husain et al. 2013). However, robotic exploration in such underground environments brings several challenges. In this paper, we focus on the challenges associated with transmitting science data out of a cave to a base station through networked robotic explorers that are constrained by intermittent communication and limited vehicle lifetime (likely to last only a few days due to the lack of sunlight to recharge batteries).

Communication in a cave environment has a high level of uncertainty in the reliability, capacity, and availability of the links between nodes (robots) in a network (Walsh and Gao 2018; Belov, Ellison, and Fraeman 2017). Even small distance variations between nodes can shift a strong signal to completely nonexistent.

A promising communication framework to address these challenges is Delay/Disruption Tolerant Networking (DTN), a technology developed at NASA to enable planetary internetworking. Although DTN was originally conceived for time-varying deterministic contact opportunities (e.g., between a spacecraft and ground station), it is currently being upgraded to also handle opportunistic contacts that cannot be predicted ahead of time. Furthermore, the literature offers several techniques to route data in opportunistic networks, traditionally using control flooding techniques (i.e. spraying/broadcasting data to every node) in the hope that a copy of the data will eventually reach the destination. However, this flooding approach assumes that energy is not a restrictive resource, an assumption not necessarily true for planetary cave exploration (Vaquero, Troesch, and Chien 2018). In fact, we believe that deliberately planning when and where to send science data becomes essential to optimally coordinate an energy-limited, multi-robot, cave exploration network. Indeed, balancing the use of energy becomes quite important when the objective is to cover a long distance into the cave.

In this work, we focus on energy-aware, smart, distributed routing capabilities in a multi-rover exploration scenario, in which science data is acquired by multiple science vehicles and routed to a target base station in a changing network environment. This is particularly relevant to multi-vehicle surface missions and subsurface missions in which 1) the environment (e.g., uneven terrains or cave structures) imposes inconsistent communication and significant bandwidth variation between the vehicles and 2) vehicle energy plays an important role. Under those conditions, science data products have to find their destination efficiently through the network to be processed. Efficiency becomes even more important when such data transportation is in the critical path to make crucial decisions in a timely manner or when limited vehicle lifetime is a factor.

In our ongoing project, we propose a set of energy-aware routing algorithms that are currently being investigated in the context of simulated cave exploration. They are broadly divided in two categories: 1) opportunistic protocols with a preferred data flow direction, and 2) schedule-based protocols in which some knowledge about the communication time windows and the bandwidths between vehicles are known or estimated. Note that the opportunistic protocols require minimal state information to perform the routing decisions and are simpler to implement. However, they might also lead to suboptimal routing decisions that reduce the total data volume that can be extracted out of the cave. Therefore, this paper is intended to be a preliminary exploratory analysis in which we define and compare these routing procedures under a simulated scenario.

The rest of this paper is organized as follows. First we introduce existing work on cave exploration and data routing in DTN. We then present the energy-aware data routing algorithms that are being studied in the context of this project. We then present preliminary results from a simulated multi-robot cave exploration scenario using the set of data routing algorithms to illustrate their advantages and disadvantages. Finally, we conclude by discussing the simulation results and outlining our next steps in the project.

2 Background

2.1 Cave Exploration

Exploring subsurface structures with robots is a hard problem due to the unknown environment, expected difficult terrain for robots to traverse, and anticipated communication challenges. Research suggests that a successful mission design for a cave scenario might utilize multiple robots to provide redundancy and carry a mix of instruments with different capabilities (Husain et al. 2013). Additionally, Dubowsky, Plante, and Boston (2006) have investigated the use of a large numbers of small, cheap robots to be deployed over a large area. They suggest these smaller robots may move better in a cave environment and mission success might be more robust to robot loss/failure, since there would be so many of them.

Operations for MSL, a surface rover on Mars that can recharge its battery, currently involves communicating the rover state, relevant observations and scientific data from Mars to operators on Earth, who then create a plan that is uplinked to the rover to be executed every sol or martian day (Chattopadhyay et al. 2014; Gaines et al. 2016). With the limited lifetime of rovers in caves, this operating paradigm becomes infeasible as the time to wait for plans would take up energy that could be used for collecting and transmitting valuable science data. This highlights that autonomous operations is important to make such missions possible.

To address these challenges in a multi-rover mission scenario in a Martian cave, we presented the *Dynamic Zonal Relay with Sneakernet Relay Algorithm* in our previous work (Vaquero, Troesch, and Chien 2018). In this algorithm, the rovers spread out into the cave, each one following its neighbor, while collecting and sending data towards a base station at the mouth/entrance of the cave. They incrementally extend further into the cave as data is transferred out. The routing decision is based purely on the existence of a connection with a rover that is closer to the base station. At some point, the rovers spread out enough that the communication connection between rovers is lost, and they need to drive back and forth (sneakernet) to transfer data out of the cave. The preliminary experiments that we performed with this algorithm indicated that transferring data is a major source of energy usage during the mission and that more data could potentially be sent out if data was being dynamically routed (instead of having rovers drive back and forth).

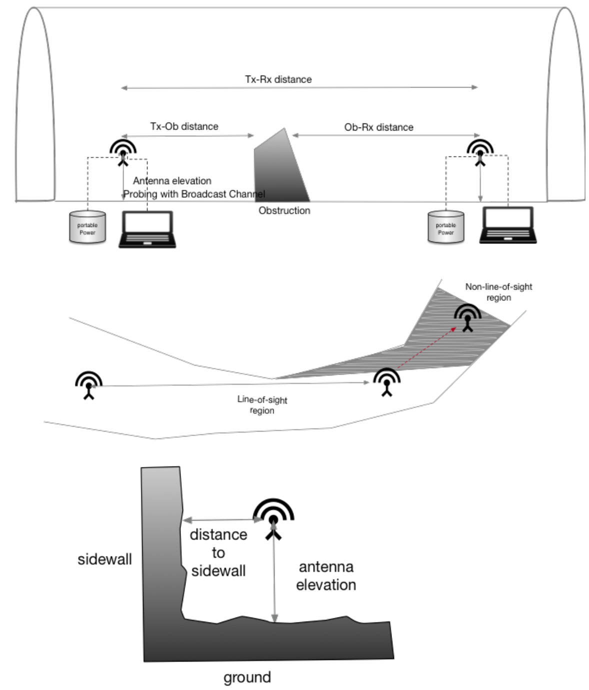

Figure 1: Diagram of important factors affecting communication inside a cave.

Communication Environment in Planetary Caves As previously mentioned, a major performance driver for planetary cave exploration is the communication environment experienced by the rovers. In (Walsh and Gao 2018), we presented a communication model for a 100 meter long lava tube (tunnel) with a 1 meter diameter opening on both ends and a 4 meter diameter in the center with large obstacles for radio wave propagation. This model predicts that a cave's geometry can cause large *constructive* and *destructive* fading effects (see Figure 1). For the former, multiple signal reflections arrive at the receiving antenna aligned in phase, therefore significantly increasing the signal strength. For the latter, these reflected waves arrive at the receiver out-of-phase, thus canceling each other out and causing a deep fade of 10dB or more. Furthermore, the transition between constructive and destructive combining can occur with a movement as short as the wavelength of the signal, which for the case of 2.4GHz WiFi, is just 12.5 cm.

Field experiments in various tunnel configurations confirms such behavior (see Figure 2). Additionally, data collected during these field activities has allowed us to calibrate average and standard deviation estimates for both large and small scale fading effects (see Figure 3). Together with an 802.11b waveform link budget, this data allows us to derive the probability distribution of link outage and achievable data rates as a function of the distance between two communicating rovers (see Figure 4), which we then apply to our network simulations. In particular, the channel model assumed for planetary cave communication maps distance d between transmitter and receiver to signal to noise ratio SNR using the following function:

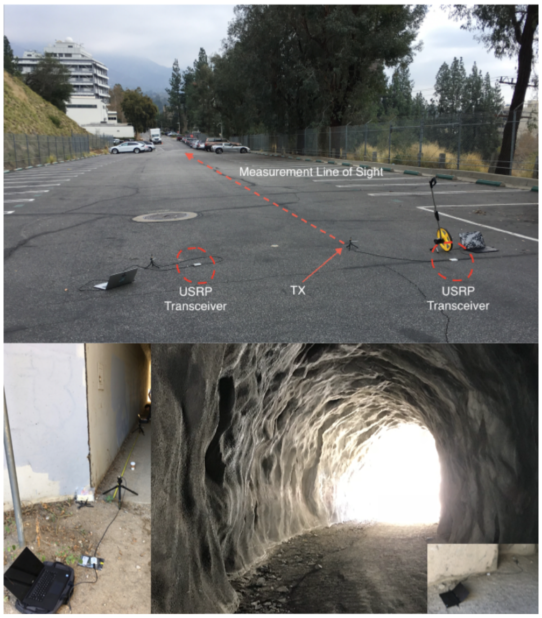

Figure 2: Pictures from the field tests.

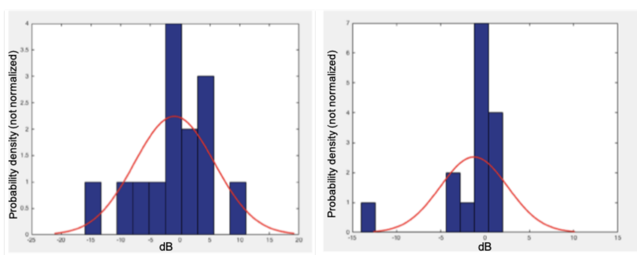

(a) (b)

Figure 3: Mueller tunnel small scale (fast) and large scale (flat) fading density functions derived from the field tests. (a) Small scale fading $\mathcal{N}(-0.97, 6.72)$. (b) Large scale fading $\mathcal{N}(-1.25, 3.79)$.

$$SNR(d) = SNR_o - 10n\log_{10}(d/d_o) - X - Y. \quad (1)$$

SNR_o is the SNR needed to achieve a bandwidth of 1 Mbps, and has a value of 45 dB. d_o is the distance where SNR_o is achieved, and has a value of 20 m. $n = 2.4$ is the exponent for SNR degradation per inverse d^n. X represents the large scale fading $\mathcal{N}(-1.25, 3.79)$, which is re-evaluated at changes of 4 m, and Y represents the small scale fading $\mathcal{N}(-0.97, 6.72)$, which is re-evaluated at changes of 0.125 m. Finally, the SNR is converted to a bandwidth BW estimate (in Mbps) using the following step function:

$$BW = \begin{cases} 0 & \text{if } 0 \leq SNR \leq 37 \\ 1 & \text{if } 37 < SNR \leq 40 \\ 2 & \text{if } 40 < SNR \leq 44 \\ 5.5 & \text{if } 44 < SNR \leq 47 \\ 11 & \text{if } SNR > 47 \end{cases} \quad (2)$$

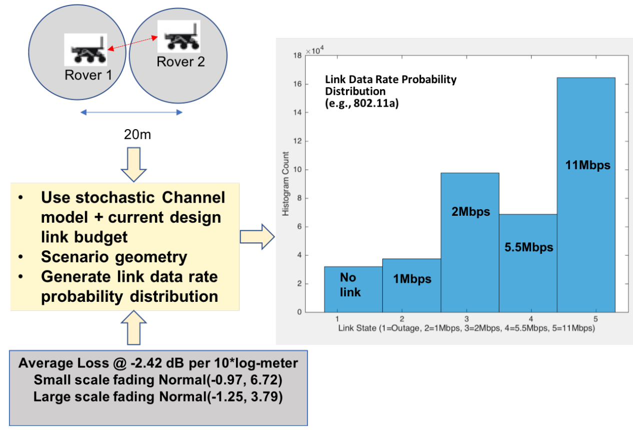

Figure 4: Process for generating the link data rate probability distribution.

2.2 Disruptive Tolerant Network

Delay/Disruption Tolerant Networking (DTN) is a networking paradigm originally conceived and developed for deep space communications characterized by long propagation delay and episodic connectivities arising from planetary and orbital motion. At its core, DTN replaces parts of the traditional TCP/IP protocol stack and avoids using handshaking procedures, which cannot be performed due to long delays and/or lack of end-to-end data path from source to destination. Therefore, DTN operates in a hop-to-hop basis instead of the Internet's end-to-end design.

Several technology demonstrations with DTN have been conducted in the past. For instance, DTN was initially demonstrated on the EPOXI spacecraft in 2009 (Wyatt et al. 2009), and has since been deployed on the ISS, supporting science commanding and data collection between ground investigators and on-board instrumentations. Additionally, DTN has also become an active area of research for applications in related domains such as MANET, sensor-webs, ad hoc vehicular networks, planetary cave exploration (Wyatt et al. 2018) and space-based interferometry (Belov et al. 2018).

Data Routing in DTN A wide body of literature has considered the problem of routing data in DTNs. In general, all proposed approaches seek to strike an acceptable trade-off between efficiency (i.e., utilization of bandwidth and buffer) and performance (delay and capacity). However, they vary in the degree of network state knowledge assumed for making routing decisions (Araniti et al. 2015). In that sense, at the lower end of the scale of network knowledge, a variety of "flooding" techniques has been studied. Having no network knowledge except its immediate neighbors, each node replicates received data units to multiple neighbors in order to ensure a high likelihood of eventual delivery to the destination. Recognizing the substantial overhead cost of pure flooding, a number of similar techniques seek to improve efficiency by limiting the degree of replication, introducing directionality to the flooding based on a spanning tree structure (additional knowledge), and imposing a waiting period for receipt confirmation (assumes acceptable delay and bi-directionality).

The most prominent example of controlled routing is Epidemic Routing (Vahdat, Becker, and others 2000), which controls the overhead by comparison of the data buffer between neighbors in order to limit data exchange only on the differences. In the end, the flooding technique is most suitable for networks that have an abundance of bandwidth and by necessity or preference will operate with low knowledge of network/link states. A second category of routing algorithms operates based on statistical knowledge of the network, which is either supplied a priori or derived while in operation from encounter history, location information, and other metrics. These algorithms utilize decision processes similar to traditional link state or distance vector (Perkins, Belding-Royer, and Das 2003) to estimate destination dependent cost and delivery probability. The most prominent example of statistical routing for DTN is ProPHET (Grasic et al. 2011).

A completely different approach to routing in DTNs assumes a high degree of network state knowledge, typically supplied to nodes a priori. For instance, Contact Graph Routing (CGR) (Araniti et al. 2015) uses a contact plan as a list of scheduled communications capabilities that is preloaded in spacecraft or regularly updated to automate the relay operation between the avionics system and the communications system. To make routing decisions, CGR first turns this contact plan into a contact graph and then performs a Dijkstra search over the network's time-variant topology. Additionally, the amount of data being sent over each contact is recorded so as to ensure that no links are overbooked.

While routing in DTN has been studied in the past, its applicability to the problem of planetary cave exploration in an energy-constrained operational environment has not yet been considered. Therefore, in this paper we consider both of them as central factors for selecting routing protocols that maximize science return, and study their performances using a simulated cave exploration environment.

3 Approach

This section describes the set of routing protocols to benchmark for multi-vehicle cave exploration. We focus on exploration scenarios in which a set of rovers collect science data and transmit them towards a base station located at the cave entrance. In particular, we propose two energy-aware opportunistic routing protocols that build upon our previous work on Martian cave exploration (Vaquero, Troesch, and Chien 2018). Additionally, we describe Energy-Aware Contact Graph Routing, a proposed extension to CGR that selects paths of minimal energy in a time-varying network topology.

3.1 Opportunistic Protocols

In traditional opportunistic DTNs, epidemic strategies are generally used. However, those strategies might compromise the lifetime of the vehicles since transmission of replicated data can consume a significant amount of energy. We believe that in a cave environment a more cautious data transfer approach needs to be developed. Therefore, we consider exploration scenarios in which vehicles' tactical navigation through the environment allows for a preferable flow of data through the network of robots. For example, in (Vaquero, Troesch, and Chien 2018) rovers $r1$ through $r4$ space themselves deep into the cave (rover $r4$ being the deepest into the cave), forming a communication chain connected to the base station at the cave's mouth. In this case, a rover deeper into the cave should send data to its immediate neighbors up the communication chain while taking into consideration its energy state and that of its neighbors. Therefore, we propose two opportunistic protocols that make routing decisions based on both this preferable data path (out of the cave), and energy considerations.

Energy Estimate Energy to transfer data is a natural metric to support the routing process in a cave environment. Herein, a vehicle i selects its immediate neighbor vehicle j that 1) is connected directly to the source vehicle i, 2) is up in the communication chain ($i > j$), 3) requires the lowest amount of energy to transfer E_{ij}, and, in case of a tie in criterion (3), 4) is the closest to the mouth of the cave. To compute the energy consumed while transmitting a packet of size S from node i to node j, we assume that both the transmitter and receiver must be turned on, drawing P_{tx_i} and P_{rx_j} Watts of power. Therefore, this energy consumption in Joules can be simply estimated as $E_{ij} = (P_{tx_i} + P_{rx_j}) \frac{S}{BW_{ij}}$, where BW_{ij} is the data rate in bits per second obtained from Equation 2. Note that for this routing procedure to work, all neighboring nodes must be aware of each other's radio power consumptions. The mechanisms by which this information is shared across the system is assumed ideal and beyond the scope of this paper, albeit reasonable to implement since it only needs to be shared locally.

Energy Left Due to the fact that a vehicle's battery in a cave would decrease monotonically (no recharging), it might be wise to balance out the data transfer through nodes which have more battery slack. Thus, in this protocol, vehicle i selects the immediate neighbor vehicle j ($i > j$) that 1) has a direct connection with the source vehicle, 2) has the greatest amount of energy left after transferring the data L_j, and, in case of a tie in criterion (2), 3) is the closest to the mouth of the cave. In this work, we compute the amount of energy left after receiving a packet as $L_j = B_j - P_{rx_j} \frac{S}{BW_{ij}}$, where B_j is the battery level of node j and the rest of the parameters are defined as before. Note that for this routing procedure to work, all neighboring nodes must be aware of each other's battery levels. As in the Energy Estimate protocol, the mechanism by which this information is shared is beyond the scope of this paper.

3.2 Schedule-based Protocols

Autonomous cave exploration strategies cannot only deliberately create periodic, scheduled communication contacts between certain vehicles (tactical contacts), but also use that knowledge to reason about routing science data products out of the cave. In these strategies, one could try to maintain communication at all times (e.g. in (Mukhija, Krishna, and Krishna 2010; Pei and Mutka 2012)), which will provide a static contact plan where all the vehicles are connected to the base station until a certain estimated horizon (e.g., a fixed conservative horizon or the vehicles' lifetime). Alternately, loss of communication can happen intentionally or expectedly as part of the vehicle coordination process. In the aforementioned sneakernet strategy for example, it assigns relay vehicles that would serve as data mules with contacts known a priori. Similarly, in a cave exploration scenario, operators could dynamically assign certain vehicles to either perform data muling at a certain frequency between two locations, or to keep the connectivity between two or more groups of explorers. Those examples illustrate cases in which a schedule of communication contacts between vehicles can be estimated or reinforced by an operator or an autonomous system.

In this section we propose a DTN enhancement that considers the temporal constraints available in the contact plan and the energy requirements. In particular, we consider an energy-aware version of contact graph routing in which paths are selected so as to minimize the energy spent while transmitting and storing a packet.

Energy-Aware Contact Graph Routing Energy-Aware Contact Graph Routing (ECGR) extends traditional CGR by finding paths of minimal energy over a time-varying topology of pre-scheduled contacts. In particular, our proposed algorithm is based on the well-known Dijkstra shortest path search procedure (Golden 1976) and, consequently, can only find *the single best path* between an origin-destination pair.

ECGR can be notionally broken down in three parts: time-varying topology construction; computation of energy metric; and selection of best path. To construct the network topology, we assume that ECGR takes a contact plan as input. This plan specifies contact opportunities between nodes as a six element tuple: start time, end time, origin node, destination node, average data rate, and range in light seconds. Additionally, ECGR also knows the hotel load (in units of Watts) expected in all network nodes, as well as the power consumed by the 802.11b radio to transmit and receive data.

To construct the graph topology, ECGR essentially inherits the procedures from time-based CGR and works on a contact graph where vertices identify contacts between two robots and edges symbolize periods of time where data is being stored somewhere in the network (Fraire et al. 2017). Prior to building the graph, a few initialization step are required: First, an initial contact C_{ini} from the transmitter to itself is added to the contact plan. It is valid from time 0 to time infinity and has an infinite data rate (similarly, a final contact C_{end} from the receiver to itself is also appended). Then, each contact in the contact plan is initialized with two state variables, Early Transmission Time (ETT) and Early Arrival Time (EAT), which are both set to infinity. ETT indicates the earliest possible time that a bit can depart a given node through the contact. On the other hand, the EAT indicates the earliest time at which a bit can reach a node through a given contact.

To build the contact graph from the initialized contact plan, we define two contacts C_i and C_j as neighbors if (1) the destination of C_i is equal to the origin of C_j and (2) the EAT to C_j is less than the time at which C_j ends. Let EAT_j and ETT_j denote the EAT and ETT of C_j, respectively. Then, EAT_j is computed as

$$ETT_j = \max(t_{s_j}, EAT_i) \tag{3}$$

$$EAT_j = ETT_j + OWLT, \tag{4}$$

where t_{s_j} is the start time of C_j and $OWLT$ is the one-way light-time delay in units of seconds. In other words, data can only be transmitted through C_j after C_j has started, and data has arrived from the previous contact C_i. Similarly, data can only arrive at the destination of C_j after you start sending it through C_j and $OWLT$ seconds of propagation delay have elapsed.

With these definitions, the contact graph can now be formalized as a set of vectices C_i (recall they represent contacts) connected through a set of edges of arbitrary but equal weight such that an edge between vertices C_i and C_j is present if and only if they are neighboring contacts. Furthermore, paths through this time varying network topology can be obtained using traditional shortest-path algorithms and minimizing the EAT of data to the destination. Therefore, CGR traditionally searches paths that optimize the time at which data is delivered to the destination.

ECGR, instead, seeks to find paths that minimize total energy consumption for a given packet. To that end, we first define the concept of total stored time. Let C_i and C_j be two neighboring contacts as previously defined. Then, the total time data needs to be stored in a node between its arrival from contact C_i and transmission using contact C_j can simply be estimated as $TST_j = ETT_j - EAT_i$. Therefore, multiplying this value by a robot's hotel load gives us a proxy for the energy spent while storing the packet in memory. Next, we consider the energy spent per transmission. In this case, we pessimistically assume that the time it takes to transmit a packet is approximately 1 second (the packet size is 1Mbit and the minimum 802.11b data rate is 1Mbps), and therefore the energy spent per transmission is simply the power consumed by the radio times this unitary transmission time. Finally, at each iteration of the Dijkstra algorithm we select as candidate for the next vertex to explore, the contact that, up to that point, has used less energy to transmit the information.

4 Experimental Results

4.1 Setup

Our preliminary studies include a set of simulated robotic cave exploration scenarios and the aforementioned data routing protocols. In what follows we describe the simulation and routing setup.

Motion Model for the Simulated Cave Exploration Inspired by our previous work (Vaquero, Troesch, and Chien 2018), we assume that the planetary cave under exploration can be modeled as a long, slender tunnel. A static base station located at the cave entry acts as a sink for all data collected by the exploration robots. These are spread across the

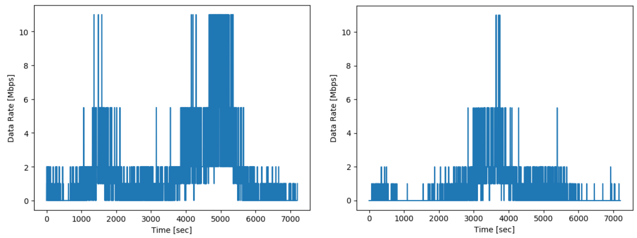

(a) (b)

Figure 5: (a) Data rate between base station and node 1. (b) Data rate between base station and node 2.

Table 1: Simulation Setup Parameters

Parameter	Value	Units
Cave dimensions	100×440	meters
Num. of nodes	10	-
Node speed	0.05	m/s
Radio type	802.11b	-
Packet size	1	Mbit
Packet generation rate	100	kbps
Packet TTL	∞	sec
Simulation duration	7200	seconds
Hotel power	5	W
Radio power	1	W
Robot battery	21.6	kJ
Num. random observations	20	-

cave in overlapping sections of 40 meters (i.e., robot 1 explores up to a 40 meter depth, robot 2 moves from 20 to 60 meters deep, and so on) and are allowed to move freely within their assigned exploration area. Furthermore, we simulate motion in 1 second increments through a random waypoint model (Tracy et al. 2002) in which robots advance at 0.05 m/s towards randomly pre-selected targets. Once they reach them, they wait in that position for up to 1 minute collecting samples, choose another location, and start moving towards it. This process is repeated until the end of the simulation.

As robots move, we collect information about their relative location to each other and feed it to the cave communication model from Equations 1 and 2. This allows us to create discrete data rate timelines that indicate the bandwidth available between any two given nodes at any point in time. Note that this data rate is, by construction, the net information rate that the 802.11b protocol offers to the upper layers of the communication stack. Therefore, overhead such as coding bits or multiple access collisions are already accounted for.

Figure 5 shows the simulated data rate between the base station and nodes 1 and 2. Since node 1 is exploring an area limited to a depth of 40 meters, its communication channel is in good state for large periods of time and the robot can transmit up to 1.4GB of data to the base station in 2 hours. In contrast, node 2 is exploring a region deeper into the cave and therefore its communication channel with the base station is impaired during certain periods of time. This results in robot 2 only being able to send 0.5GB to the base station over the same 2 hour period. Similar trends would be observed for robots further into the cave, which would necessarily need to route most of their data through peers rather than directly contacting the base station.

Exploration Scenarios Table 1 summarizes the main parameters used in our simulation environment to generate the experimental results described in this section. We simulate the aforementioned long, slender cave with 10 nodes placed in a quasi-linear topology. All of them generate packets of 1Mbit at a constant average rate of 100kbps. Each packet is destined to the base station located at the cave entry and has infinite Time-To-Live. Furthermore, robots utilize 5W of power to move and conduct their operations (including science data collection), and require 1W of power to transmit any packet – values inherited from our previous analysis (Vaquero, Troesch, and Chien 2018). Finally, each scenario is run 20 times randomizing both the rovers' motion and the underlying communication channels to ensure statistical significance of the results.

Several simplifications were made for the initial set of results reported in this paper. For instance, we consider that robots are pre-positioned in their desired area of exploration without any energy penalty and they do not need to exit the cave upon battery exhaustion (they just need to send the data back). Similarly, power consumption of all rovers is assumed nearly constant over time, with no significant differences between a robot in stand-by mode and in data collection mode. Also, we assume that no robots fail due to unexpected circumstances and they operate according to the predefined plan captured in the contact plan (i.e., they are not adaptable to environment changes). Therefore, the results reported in the following sections must be interpreted as an upper bound on system performance, and will be revised once experimental data from hardware-in-the-loop testing is conducted.

Routing Approaches Four routing procedures are tested: Opportunistic with energy estimate minimization; Opportunistic with energy left maximization; Time-based Contact Graph Routing (CGR); and Energy-Aware Contact Graph Routing (ECGR). For the opportunistic protocols, the only state information provided to the routing procedures is a preferred path sequence that specifies the line topology of the cave (e.g., node i will only be able to send data to nodes $i-1$, $i-2$, and so on, as well as the base station). This will ensure that data is always routed towards the cave's mouth, albeit the exact next hop for a given packet will be decided in real-time depending on the state of the network. On the other hand, both CGR and ECGR require a contact plan as state information to route data. This contact plan is constructed together with the robot motion and communication channel using the following steps. First, the position of all robots as a function of time is computed. Then, we estimate the expected bandwidth for all origin-destination pairs in the network in 1 second increments. Next, we discretize the bandwidth of each pair into 5 minute intervals and compute the average data rate over them. Finally, the contact plan is built using these 5 minute averages, where a contact is assumed to be present if the average data rate is greater than 500kbps.

Metrics We simulate the system for 2 hours and record the following information: 1) *percentage of packets* that arrive to the base station (note that the maximum possible value for this percentage is not necessarily 100% as some environments are so challenging that you cannot guarantee packet delivery); 2) *latency*, measured as the time a packet is delivered to the base station minus its creation time in units of seconds; 3) *routing overhead*, measured as the number of routing calls performed per packet delivered; and 4) *robot death time*, measured in seconds from the start of the simulation, due to battery exhaustion.

4.2 Results

The performance of the four routing approaches are shown in Figure 6. Surprisingly, results indicate that there is little difference between using opportunistic vs. contact-based routing schemes in a cave environment for traditional networking metrics such as percentage of packets delivered to destination or latency. Similarly, there is also little performance difference with respect to the robot's lifetime, as energy spent in packet transmission is significantly lower than energy used for other robot operations and therefore the latter dominates battery depletion. Note that this conclusion is only applicable to the scenario reported in this paper, where data generation per node was constrained to very conservative numbers. We expect that differences in performance between the considered routing algorithms will be observed when more data-intense scenarios are run (e.g., gigabytes of data transferred, as opposed to megabytes). However, we observe an almost two order of magnitude increase in routing overhead when using the opportunistic methods as compared to CGR and ECGR. Unfortunately, properly characterizing the impact of this overhead with respect to energy is not possible in this study since no CPU models were available while collecting results. Intuitively, we know that the per call computational load of an opportunistic algorithm is minimal, whereas CGR and ECGR require a Dijkstra search over a potentially large contact graph. However, since we cannot estimate how many CPU cycles are required for either of them, it is impossible to translate this overhead to time and energy effects that can then be measured during simulation time.

5 Conclusion

This paper considered the problem of routing data in a planetary cave environment. In particular, we analyzed different routing strategies assuming that non-rechargeable mobile robots are placed inside a long, slender cave to conduct science activities, and information generated needs to be sent back to a base station located at the cave entrance. Four routing strategies were considered: opportunistic routing minimizing energy or maximizing remaining battery; contact-based scheduling using CGR; and energy-based scheduling using ECGR. Of those, ECGR was developed as a novel extension to CGR that considers energy instead of time as the primary routing metric.

Each routing alternative was benchmarked against a quasi-linear topology in which 10 robots are exploring a 440m long cave and only the two closest rovers to the entrance can communicate directly with the base station. Additionally, we modeled energy constraints in the system by assuming that each robot has an initial battery charge that decays progressively due to both robot movement and data transmission. Results show that the performance of the four tested algorithms is very similar in terms of percentage of packets delivered to destination and their latency. However, we observe that the opportunistic-based routing schemes perform up to two orders of magnitude more routing calls than their CGR-based counterparts. This is partly due to the fact that CGR-based routers can proactively store packets that are not currently routable and wait for a future contact. In contrast, opportunistic approaches must re-try routing them constantly until they succeed.

Several areas of future work have been identified while conducting this study. First and foremost, more test runs must be performed in order to properly compare these routing procedures. While our current work has included up to 20 randomized observations per run, the simulation setup has remained mostly constant. Moreover, we will study different configurations and strategies for the multi-robot cave exploration to analyze where each routing strategy thrives. In future work, we will consider several other factors of interest such as the computational load imposed by opportunistic vs. CGR-based routing algorithms and their impact on the lifespan of these battery-powered rovers. Similarly, we will also characterize the scalability of these algorithms

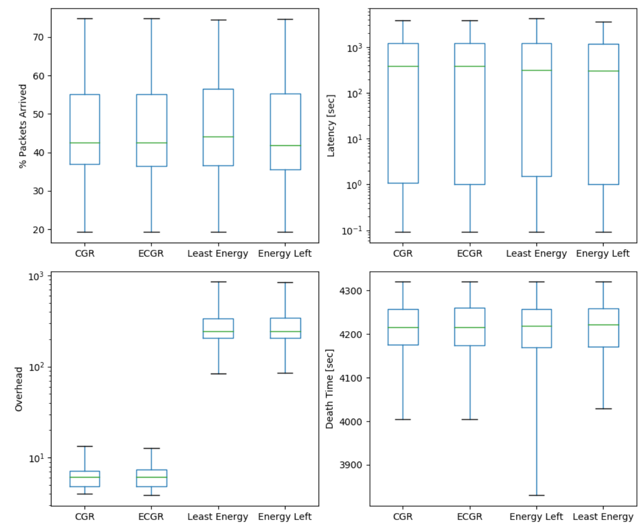

Figure 6: (a) Percentage of packets delivered to destination. (b) Latency of packets arrived in seconds. (c) Routing overhead. (d) Node death time.

with respect to traffic generation rate and number of nodes. Finally, in this paper it was considered that all robots have exactly the same capabilities and starting configuration. Scenarios with heterogeneous rovers and/or different initial conditions should also be tested (e.g. (Bechon et al. 2018)).

6 Acknowledgments

The research was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under contract with the National Aeronautics and Space Administration.

References

Araniti, G.; Bezirgiannidis, N.; Birrane, E.; Bisio, I.; Burleigh, S.; Caini, C.; Feldmann, M.; Marchese, M.; Segui, J.; and Suzuki, K. 2015. Contact graph routing in dtn space networks: overview, enhancements and performance. *IEEE Communications Magazine* 53(3):38–46.

Bechon, P.; Barbier, M.; Grand, C.; Lacroix, S.; Lesire, C.; and Pralet, C. 2018. Integrating planning and execution for a team of heterogeneous robots with time and communication constraints. In *IEEE International Conference on Robotics and Automation, ICRA 2018, Brisbane, Australia, May 21-25, 2018*, 1091–1097.

Belov, K.; Branch, A.; Broschart, S.; Castillo-Rogez, J.; Chien, S.; Clare, L.; Dengler, R.; Gao, J.; Garza, D.; Hegedus, A.; Hernandez, S.; Herzig, S.; Imken, T.; Kim, H.; Mandutianu, S.; Romero-Wolf, A.; Schaffer, S.; Troesch, M.; Wyatt, E. J.; and Lazio, J. 2018. A space-based decametric wavelength radio telescope concept. *Experimental Astronomy*.

Belov, K.; Ellison, D.; and Fraeman, A. A. 2017. Pisgah lava cave communication test. Technical Report JPL D-100993, Jet Propulsion Laboratory.

Boston, P.; Frederick, R.; Welch, S.; Werker, J.; Meyer, T.; Sprungman, B.; Hildreth-Werker, V.; Thompson, L.; and Murphy, D. 2003. Human utilization of subsurface extraterrestrial environments. *Gravitational and Space Biology Bulletin* 26(2).

Chattopadhyay, D.; Mishkin, A.; Allbaugh, A.; Cox, Z.; Lee, S.; Tan-Wang, G.; and Pyrzak, G. 2014. The mars science laboratory supratactical process. In *Proceedings of the 13th International Conference on Space Operations (SpaceOps 2014), Pasadena, CA*, 3977–3991. AIAA.

Dubowsky, S.; Plante, J.-S.; and Boston, P. 2006. Low cost micro exploration robots for search and rescue in rough terrain. In *IEEE International Workshop on Safety, Security and Rescue Robotics, Gaithersburg, MD, USA*.

Fraire, J. A.; Madoery, P.; Burleigh, S.; Feldmann, M.; Finochietto, J.; Charif, A.; Zergainoh, N.; and Velazco, R. 2017. Assessing contact graph routing performance and reliability in distributed satellite constellations. *Journal of Computer Networks and Communications* 2017.

Gaines, D.; Anderson, R.; Doran, G.; Huffman, W.; Justice, H.; Mackey, R.; Rabideau, G.; Vasavada, A.; Verma, V.; Estlin, T.; et al. 2016. Productivity challenges for mars rover operations. In *Proceedings of 4th Workshop on Planning and Robotics (PlanRob)*, 115–125. London, UK.

Golden, B. 1976. Shortest-path algorithms: A comparison. *Operations Research* 24(6):1164–1168.

Grasic, S.; Davies, E.; Lindgren, A.; and Doria, A. 2011. The evolution of a dtn routing protocol-prophetv2. In *Proceedings of the 6th ACM workshop on Challenged networks*, 27–30. ACM.

Husain, A.; Jones, H.; Kannan, B.; Wong, U.; Pimentel, T.; Tang, S.; Daftry, S.; Huber, S.; and Whittaker, W. L. 2013. Mapping planetary caves with an autonomous, heterogeneous robot team. In *IEEE Aerospace Conference*, 1–13.

Mukhija, P.; Krishna, K. M.; and Krishna, V. 2010. A two phase recursive tree propagation based multi-robotic exploration framework with fixed base station constraint. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 4806–4811.

Pei, Y., and Mutka, M. W. 2012. Steiner traveler: Relay deployment for remote sensing in heterogeneous multi-robot exploration. In *IEEE International Conference on Robotics and Automation*, 1551–1556.

Perkins, C.; Belding-Royer, E.; and Das, S. 2003. Ad hoc on-demand distance vector (AODV) routing. RFC Editor, United States. https://tools.ietf.org/html/rfc3561.

Tracy, C.; Jeff, B.; Vanessa, D.; et al. 2002. A survey of mobility models for ad hoc network research. *Wireless Comm. & Mobile Computing (WCMC)* 2(5):483–502.

Vahdat, A.; Becker, D.; et al. 2000. Epidemic routing for partially connected ad hoc networks. Technical report, CS-200006, Duke University.

Vaquero, T.; Troesch, M.; and Chien, S. 2018. An approach for autonomous multi-rover collaboration for mars cave exploration: Preliminary results. In *International Symposium on Artificial Intelligence, Robotics, and Automation in Space (i-SAIRAS 2018)*. Also appears at the ICAPS PlanRob 2018.

Walsh, W., and Gao, J. 2018. Communications in a cave environment. In *Proceedings of the 2018 IEEE Aerospace Conference*, 1–8. IEEE.

Wyatt, J.; Burleigh, S.; Jones, R.; Wissler, S.; and Torgerson, L. 2009. Disruption tolerant networking flight validation experiment on nasa's epoxi mission. In *Proceedings of the 1st International Conference on Advances in Satellite and Space Communications*, volume 1, 187–196.

Wyatt, E. J.; Belov, K.; Castillo-Rogez, J.; Chien, S.; Fraeman, A.; Gao, J.; Herzig, S.; Lazio, T. J. W.; Troesch, M.; and Vaquero, T. 2018. Autonomous networking for robotic deep space exploration. In *International Symposium on Artificial Intelligence, Robotics, and Automation for Space (ISAIRAS 2018)*.