Frequently Asked Questions (FAQ) v1 SolarNet, March 7, 2025

Note: This document includes answers to questions posed by students in their DP Prep assignments as well as on piazza. The questions were all over the place, hence the length of this document. If you don't find your question answered in here, ask it on piazza. The updated material as of March 11 is the section on Specialized (Telemetry) Satellites.

Software updates

1. *Q: Are nodes completely inactive during an update?* A: If you think about your computer or phone, when an update is happening, there is some initial process that involves things like figuring out if there is enough space in the file system, which files will be replaced, etc. That part of the process is interruptible but will need to be repeated, if the update is restarted. Then there is a period when the code is being installed, etc. That is not interruptible. An update must either happen completely or not at all and during the crucial "update" part of it, the device is inactive for anything else, like transmitting and receiving traffic.
2. *Q: Does a satellite need to store a complete update in order to update its software?* A: Yes.
3. *Q: What is the 30 minute time constraint on updates?* A: Once an update is being sent out, it is expected that any device receiving it will begin the update process within 30 minutes of that. In other words, the sender of the update is assuming that all updates will start within 30 minutes of when the "send update" is initiated. By knowing how long such an update takes, the management can then know an upper bound on the time by which they can expect all updates to be completed. As with your own computer, some updates are very fast, a few seconds, and other can take up to 5 minutes.
4. *Q: What happens if an update fails to complete?* A: If an update fails to complete, it is possible that the device will need to be taken out of the normal path of operation until it can be updated, but this is extremely rare. Most updates will occur when needed and if they don't the device will be able to continue operation, perhaps with some slightly reduced capability. These updates can be to internal operations of the node or possibly to the protocols among the nodes. If the protocol is being revised and the update fails, that node will not support the update. For this reason it is always a good idea to make such updates backwards compatible, so at least the original model of operation will continue to work.
5. *Q: Are updates scheduled or unscheduled?* A: Both. Table 4 provides a schedule for routine updates. In addition, if the management finds something critical that needs updating sooner, they can send that out. They do not have to wait for the next scheduled update. The expectation is that these will be extremely rare, but you can expect that there may be 1-2 a year for each kind of device.
6. *Q: Will an update bring down the whole system?* A: This is a design choice for you. Clearly, a simple approach is to have every node begin updating as soon as possible, but that might mean that the whole network would be down for some period of time. Another approach would be some form of coordination, so they do not all update at the same time. This is unspecified at this time and up to you.
7. *Q: What about confirmation of completion of a correct update?* A: The update process, as with your computer, has a verification phase to be confident that the update completed correctly. Once that is done, you will need to arrange for a message to be sent back to the originator to

report that it is complete. Until the originator receives that message it must assume that the node has not completed its update.

8. *Q: How demanding are the updates on bandwidth?* A: The spec tells you how large each file is and how many there are, as well as the bandwidth capacity of each link. It would be a good idea to figure out the minimum amount of time (if the satellites are lined up perfectly) it would take to deliver an update to a satellite, and then come up with a model of what could disrupt that, as well as how you will deal with imperfect communication.

Flexibility and extensibility

1. *Q: What kinds of extensions should we expect? More devices, different types of devices, longer range devices?* A: These are all possibilities. Starlink is at something like 6,000 or so satellites now and is expecting to be at over 40,000 in five years. NASA is not planning on SolarNet growing that large because it servers a different purpose. In addition, NASA definitely assumes there will be upgrades in devices and possibly new types of devices. They already have a probe (Parker Solar Probe) that has gotten very close to the Sun. They also have probes reaching the outer parts of the Solar system. But in SolarNet they are not designing for massive populations on other planets with commercial sorts of communications requirements. We suggest that you not design for these kinds of challenges from the beginning. Rather, design for the immediate requirements and constraints and then consider extensibility.
2. *Q: Does SolarNet integrate in some way with the deep-space network?* A: The deep-space network consists of probably 2 Earth-based antennas, using optical communications, so they need direct line-of-sight. That is actually what is being used now for the Parker Solar Probe. What we are proposing here for the solar telemetry satellites is the same approach except that if the satellite is not within range of one of the Earth-based satellites and has determined that there is an urgency to delivering data (based on the preliminary analysis that a Solar flare may be predicted), the SolarNet can be used for more expedited delivery of the data for more complete analysis.

Security and Integrity

1. *Q: Should we be designing for security?* A: First, you should not be thinking about that now. We won't get to talking about security until later in the term. This is not a complete security solution, but it is a first step. Second, the Bundle Protocol already provides encryption. We didn't raise that because we won't get to security until later in the term.
2. *Q: What about data integrity?* A: All data carries checksums. These address problems of individual bit errors. Larger data integrity problems are not addressed, but you can set aside at least for now.
3. *Q: Who has access to the network?* A: For our purposes now, you can assume that only authorized users have access.

Routing and Forwarding

1. *Q: Can we assume that routing tables are up to date?* A: Remember that this is a relatively small network, so routing information and updates can be quite efficient. So, routing information is extremely likely to be quite up to date (within one minute). It is only if a satellite loses all communication that its routing information may be out of date, but it also can't communicate at that point anyway.

2. *Q: How much simultaneity is there is sending and receiving on satellites?* A: As discussed on p. 6, each satellite can simultaneously send one bundle and receive one bundle. It cannot send two simultaneously or receive two simultaneously. It has one sending laser and one receiving antenna and each can do at most one thing at a time.
3. *Q: What do the routing tables tell us?* A: First, for each destination address there will be a next hop that is the "best" choice and will include the number of bytes that can be guaranteed to be transmissible to that next hop. This provides an indication of how long the routing algorithm believes this link will be available. It says nothing about future predictions beyond that. In addition, the routing protocol assumes that a complete path must exist from a next hop for that next hop to be a valid next hop. Notice, that that does not guarantee that that path will continue to exist. Failures and congestion can happen. In addition, the routing table will provide the set of neighbors that are available for finding an alternate or additional path. Again for each of those there will be a number of bytes that can be guaranteed to be transmissible to that neighbor. Note this does not tell you how much capacity that neighbor has to store or forward bundles. That can only be determined by querying the node itself. So, you cannot redesign the "best next hop" algorithm but you can enhance the system by allowing for side-stepping it or extending it. In addition, there is no guarantee that the next neighbors will have an alternate path, although it is highly likely that somewhere along the way, there will be an alternate path.
4. *Q: Can the routing protocol be extended to include storage capacity?* A: As stated in a number of places, any changes to protocols will require an extraordinary case to be made to NASA. Much of this is already specified, has gone through rigorous vetting and is implemented. In addition, since the routing protocol sends updates once a minute, I would question whether a potential receiving node could have given an accurate prediction of storage available a minute ago in a dynamic transmission system such as this.
5. *Q: How are the routing table updates transmitted?* A: Using the Bundle Protocol.
6. *Q: Can routing information be out of date, so forwarding to a failing node might be a problem?* A: Yes, that is always possible. Generally, failures are unpredictable. When a satellite tries to forward to another one as a next hop, it is always possible that that receiving satellite might either already be down or that it might go down during the communication. In addition, of course, the forwarding satellite could go down during a communication and fail to send the complete bundle. These failures are rare and random. Notice that that is one of the reasons that bundle delivery is only best effort. Think of the postal system. One puts a letter into the mail. In general, you don't know whether it got there, or whether possibly the envelop split open and part of it got there. Only with reporting can you know for sure that it got there and notice that reporting is also done by mail.
7. *Q: It is stated that "90% of the time a LEOCom can reach one of the GEOComs" (pg. 6). Does this refer to communication between a given LEOCom and GEOCom or does this mean that 90% of the time a LEOCom is able to reach some GEOCom?* A: The latter, a given LEOCom has a 90% probability of reaching some GEOCom.
8. *Q: What is the MTU (Maximum Transmission Unit)?* A: Consider the LEOComs. They are all in equal polar orbits. Their LEOComs neighbors will always be the same set of satellites, so if they are communicating with each other there is no required limit on their MTUs. That said, there is good reason to not "allow" that – sharing of resources. As long as satellite A is forwarding to satellite B, it can't be also forwarding to satellite C, so setting limits provides improved sharing of resources.
9. *Q: Is SolarNet isolated or does it interoperate with other networks?* A: In this project it is isolated. There is no defined interface between SolarNet and other networks (e.g. Starlink or the broader terrestrial Internet).

10. *Q: What is the probability that a specific GEOCom can communicate with other GEOComs*? A: 100% except during failures (and possibly things like Solar Storms). The GEOComs are all in the same orbit as each other, around the equator and because they are geosynchronous each stays above the same point on the Earth. They appear to just sit in place by orbiting at the same rate the Earth is rotating.
11. *Q: Does SolarNet have a fixed or dynamic topology?* A: Topology is defining which nodes can communicate directly with which others. Within the set of LEOComs neighbors are persistent. This is also true among the GEOComs. But, when you consider the relative orbits of LEOComs and GEOComs, adjacent satellites will change with time. This also means that the number of LEOComs and GEOComs in direct communication with each other may vary with time.
12. *Q: How often are satellites and relays moved?* A: Probably extremely rarely. Moves might occur to address failures of other nodes. Generally, this is expected to be a very stable situation.
13. *Q: What flexibility exists in the Bundle Protocol to move lower priority bundles (or others) to neighboring satellites?* A: This is part of your challenge. You absolutely can allow bundles that previously would have had a single path designated to take alternate or duplicate paths. To do this, you will need to move bundles (fragments) to alternate nodes.
14. *Q: What does "overhead may be affected by lateral movement among satellites at the same altitude in order to achieve some performance tradeoff" mean?* A: The routing table recommends the best next hop for a bundle. If you choose to "divert" a bundle to a neighbor instead of that next best hop, that by definition is less than the next best hop, so it will have cost you something (perhaps in terms of delay). There may be many good reasons to do this – that is the tradeoff.
15. *Q: What is the relationship between LEOComs having only a 90% probability of reaching GEOComs and LEOComs having lower probability of communicating when at the poles?* A: They are closely related. When the LEOComs are at the poles they are the farthest point from the GEOComs orbits and the curvature of the Earth may cause problems. You can expect that the routing protocol and routing algorithm will understand this and it is unlikely that the best next hop for traffic on a LEOCom at a pole will be a GEOCom. Also the same in reverse for traffic going from GEOComs to LEOComs.
16. *Q: Tell me more about expiration times.* A: They can mean whatever the application chooses them to mean. For some purposes, the expiration time is critical. Whether it is landslide or solar loop data, if it arrives hours late it is of no use for predicting landslides and Solar storms, but the data is not then useless. The same data is used for all sort of long-term studies, etc. So, the expiration time is a hint to the forwarding system. For this particular use at this particular time, delivery has this level of urgency. If the expiration time arrives and there is other extremely urgent data, then it would be a good idea to drop the expired data and move on. But, if there is nothing urgent, then continuing to delivery might be valuable.
17. *Q: Do we know the expected delivery time of a bundle?* A: You can figure it out. The bits are traveling at the speed of light. The amount of time it takes can be measured from the time the first bit is send until the time the last bit arrives.
18. *Q: Can different fragments of a single bundle take different paths?* A: Absolutely, and that may be a particularly useful approach, because fragmentation implies that only part of the bundle could be forwarded along the "best" current path.
19. *Q: What happens if a bundle reaches its maximum hopcount?* A: It is a candidate to be dropped, but if resources are available to continue forwarding, that may be the better option. The hopcount is only a hint.

20. *Q: If the transmission of a bundle fails midstream, can the sender pick up from where it left off and continue sending?* A: Extremely unlikely because the sender can't know how many bits actually got there.
21. *Q: What does it mean for bundles to be interleaved and how does that impact reconstruction?* A: These are two different issues. When a node has queues of bundles to be forwarded let's say at a single priority, it could choose to forward all the bundles of one queue before all the bundles of another, or it could decide to do something that takes bundles from each queue in succession. For example, this could be a round-robin kind of scheme, where it takes one bundle from each queue before going to the next. This is unrelated to reconstruction which occurs only when a bundle has been fragmented. See the discussion of fragmentation elsewhere in other questions.

Custody nodes and custody bundles

1. *Q: What are custody bundles?* A: See p. 8. Custody bundles are bundles for which the sender wants a reasonable degree of reliable delivery. In TCP, we have end-to-end acknowledgement of each packet. In the Bundle Protocol, end-to-end acknowledgement often may be feasible, if it takes 20 min. for each acknowledgement to arrive. So, instead the Bundle Protocol supports a store-and-forward approach in which each custody node remains responsible for a custody node until it has been notified by a succeeding custody node that it now is prepared to take custody. This step by step assurance avoids the long round trips between, for example, Earth and Mars, while providing assurance of delivery.
2. *Q: Who assigns custody nodes*? A: NASA (See p. 8). They will assign some number of LEOComs and some number of GEOComs to do this extra work. In addition, all the relays, the specialized telemetry nodes the antennas on Earth, the Moon, and Mars will be custody nodes.
3. *Q: Can a custody node do other things while waiting for fragments of a bundle for which it has received some but not all fragments?* A: Yes. Each fragment delivery is a separate bundle. The custody node awaiting further fragment bundles for one particular full custody bundle can receive other bundles while waiting. As with all nodes, they can be receiving on their one antenna and sending with their laser simultaneously.
4. *Q: Can custody nodes also handle non-custody bundles?* A: Yes, custody nodes are "super nodes" in that they are regular nodes that also happen to support custody.
5. *Q: How is fragment and bundle handling related to priority in a custody node?* A: They are handled in exactly the same way as with non-custody nodes. Highest priority handled first, etc.
6. *Q: How many custody nodes are there?* A: You can assume that at least 1/3 of the nodes of each type (LEOComs and GEOComs) are custody nodes. NASA is prepared to add a few more (up to ½ of the nodes) if you can make a case for that and explain the tradeoffs.
7. *Q: Do custody bundles have to pass through only custody nodes?* A: No, that would be impossible. See the description on p. 8.
8. *Q: Are custody nodes special in any other way, like more reliable and less prone to failures?* A: No. They are just souped up nodes.
9. *Q: How should custody nodes balance long-term storage with shorter term transmission requirements?* A: In whatever way you think best.

Priority levels

1. *Q: Who decides on priority levels?* A: The application that is generating the ADU. This is not something that you have any control over. This is true for ALL applications, whether it is telemetry data or other kinds of transmissions from other locations such as ground control and

transmissions originating from the Moon or Mars. The applications always decide on priority when a bundle is being originated.

2. *Q: Can a priority level change inside the network?* A: It is not an option I had considered. If you think it is important and can make the case for it then suggest it. Remember that it is a suggestion to NASA, because they hadn't designed for it.
3. *Q: How is the priority level indicated?* A: See bits 7-8 in Table 3. Note that the space for all these flags is part of the primary block of a bundle, whether or not any particular bundle has those flags and values set or not.
4. *Q: Does priority level have any impact on how storage is handled, especially in case the storage is full?* A: It could but it doesn't have to. This is design choice you can consider.
5. *Q: When traffic arrives at the same priority level, but between different source-destination pairs, what happens?* A: See pp. 8-9. There is no predefined determination of what happens here, but your design could make choices about how this is handled or leave the decisions to the underlying system. It only guarantees what is specified in this part of the text.
6. *Q: Are there ways to differentiate the priority of traffic within a single priority between a single source/destination pair?* A: Nothing has been provided to you. If you think this is important make the case for both the fact that a solution is needed and how you propose to handle it.

Relays

1. *Q: Where are the relays located?* A: They are in the vicinity of the Moon and Mars. There is one stably located (orbiting) over the ground antenna on each of the Moon and Mars and another orbiting farther out in such a way that generally, even if the relay over the antenna is in the shadow of that body with respect to Earth that second relay is reachable both by that direct antenna relay and Earth. Notice that sometimes the antenna focused relay will also have direct contact with Earth, when not in the shadow.
2. *Q: What about GEOCom connectivity directly to a Moon antenna?* A: Notice that the same side of the Moon is always facing the Earth. It rotates at the same rate that it revolves round the Earth. So, if an antenna on the Moon is unreachable directly by the GEOComs, that will always be true and vice versa if it can communicate directly with Earth that will always be true. Note that even if a Moon antenna is facing the Earth, that does not mean that there will always be a GEOCom available for communication. This is a challenge for you.

Failure and Loss

1. *Q: Is there any relationship between failures?* A: Many failures are independent. If a piece of hardware on a satellite fails, there is no particular reason to believe that that would cause some cascade of failures among other satellites. That said, if there is a design flaw in one satellite, it is likely that that exists in others as well, although each nodes experiences will differ from its neighbors because they see different traffic. In addition, there are some sorts of failures that may be seen in multiple satellites simultaneous, in particular the effects of Solar Storms. If it happens to be a short Solar Storm, then it is possible that all the satellites in direct line of site with the Sun at the arrival of the storm will be impacted and others will not.
2. *Q: What happens when there is a satellite failure?* A: That totally depends. Sometimes rebooting solves the problem, so all satellites will have auxiliary capability to attempt to reboot. Sometimes that doesn't solve the problem. At that point, NASA has a couple of options. It may move other satellites slightly to improve coverage in the face of one fewer satellites. At a slower pace, they may launch a replacement.

3. *Q: What happens to persistent storage when a satellite fails?* A: Exactly what happens on your computer. If the persistent storage didn't fail, then, if the device can be rebooted, the persistent storage will still persist. If it can't be rebooted, it won't help much to have it stored persistently. Unlike your computer, there is no service desk to take it to, to extract what is on the storage.
4. *Q: What are the most common failure modes?* A: We don't know, but possibly a more important question for you to consider is what are the most important failure modes and how frequently might they occur.
5. *Q: How much retransmission should occur when bundles are lost?* For non-custody bundles, this is completely an application specific decision. For custody bundles, the local custody node decides this. It may depend on a number of factors such as life-time or hopcount indicators as well as local storage capacity. If it is a low-priority custody bundle, the local decision may be simply to not resend it more than once or twice before reporting back that transmission failed.
6. *Q: What fail-safe methods exist for continuous operation if satellites and relays unexpectedly fail?* A: It isn't clear there is one. If the critical satellites or relays fail, then communications through them has failed and there is nothing to be done other than try again later, if possible.
7. *Q: Should the system degrade gracefully?* A: Absolutely, that is the best. If some of the high priority data can be delivered that is better than none. If there are humans on the Moon or Mars, some communication with Earth is better than none.

Orbits

1. *Q: Is a LEO orbit 1.5 hours pole to pole or all the way around?* A: An orbit is all the way around, not halfway.
2. *Q: How accurate are orbit predictions?* A: Extremely accurate. And NASA can tweak the satellites as needed to keep them in orbit. (Part of the management functionality.)

Fragmentation

1. *Q: How is fragmentation handled?* A: Using the `fragment_bundle` call. It will produce two fragments, the original one of the specified size and the remainder. Every fragment is given an ID that will cause reassembly to work. The Bundle Protocol handles reassembly for you.
2. *Q: Can all bundles be fragmented?* A: No. As stated on p. 11, administrative bundles can have no flags set. By definition, that includes the fragmentation flag.
3. *Q: Is there a minimum fragment size?* A: No, but the smaller the fragments, the more overhead there is in sending because each fragment requires a full header block. You could choose to decide on a minimum size and explain why you have chosen that.
4. *Q: How does reassembly work?* A: The node receiving the bundle will need to collect the fragments as they arrive, possibly out of order. Each will be labeled with its place in the whole bundle. Notice reassembly only needs to occur at a custody node, for a custody bundle or at the final destination for non-custody bundles.

Naming or addressing

1. *Q: What naming hierarchy is used for the nodes in this network?* A: As specified in the spec each device is given an IPv6 address (128 bits). In addition, given that there is such a small number of nodes in the network, every device will have an internal list of all the devices including their type of device and IP address. It isn't clear that hierarchy plays any particular role in this.

Reporting

1. *Q: Is there any way to enforce reporting, if the reporting flags have been turned on?* A: No, a node makes that decision. In general, the nodes will try to honor reporting request with priority given to custody transfers, but the node is the final arbiter of whether or not it has the resources to do this and might choose other traffic to send instead. In addition, since all administrative bundles are only best effort (no custody permitted on administrative bundles), the system must be prepared for administrative bundles not to arrive.

Storage

1. *Q: How much storage is there on Earth?* A: More than enough and it is easy to add more as needed, since it is on Earth.
2. *Q: How will storage be allocated dynamically?* A: You are designing both the local storage management and any distributed storage coordination. That is part of your challenge.
3. *Q: How will the system determine which information can be deleted from the network storage?* A: That is part of your design.
4. *Q: What about data integrity?* A: In general, all systems today provide significant data integrity checking. If you feel you need more on top of that, that is part of your design.

Solar Storms

1. *Q: What are the effects of Solar Storms?* A: They are magnetic storms. They generally effect all types of radio communication. So, anything that uses radios, cellphone communication, GPS, etc. Also, presumably TV and radio communication. In addition, they can have a significant impact on electronics, in particular in satellites themselves – they're operation. They have also been known to have an impact on the power grid. Typically these storms last for somewhere between a few minutes and few hours, but the after effects can last much longer than that. They range significantly in intensity with the strongest ones being extremely rare.
2. *Q: What is the frequency distribution of Solar Storms?* A: The Sun goes through an 11-year cycle, with the frequency and strength of Solar Storms growing and then shrinking over that 11-year cycle.
3. *Q: What is the processing and notification time for Solar Storm alerts?* A: The solar loop analysis at its best will be able to give 5 hours of warning of an impending solar storm. Because no number is given you can assume that the analysis of the data on the ground takes a very small amount of time (think a few minutes) to decide whether a warning should be posted. In addition, posting can happen essentially instantaneously. So, it is advantageous to get that data to Earth as quickly as possible so that any warnings can be posted as quickly as possible.
4. *Q: What is the impact of false positives and false negatives in predicting Solar Storms?* A: If a Solar Storm occurs and no warning went out, then anyone who could have prepared for the situations discussed above will not have done so. On the other hand, too many false positives will lead to the population not believing those warnings and hence not taking precautions when they could have. Both are a problem. You can decide how important each is.
5. *Q: Is notification of an impending Solar Storm automated or manual?* A: I'm not quite sure why this matters, but in general, sending out mass alerts is best monitored by humans and then automated. So, if the system predicts some level of storm, it is probably wisest to have final human approval for an alert to be sent out, just to check for other mitigating circumstances.

Specialized (telemetry) satellites

(This is also reflected in the updated Specification.)

1. There are two specilialized types of satellites, land telescopes and solar probes. There are two of each type. The land telescopes are LEO satellites at a slightly different altitude than the LEOComs. During normal, non-critical periods, they transmit their data directly to one of two ground receivers, when they are over them, but in an emergency, they can use the SolarNet syste to communicate with either nearby LEOComs or GEOComs. The solar probes are in the same geosynchronous orbit as the GEOComs and use the SolarNet system for all communications. All software updates to these satellites use the SolarNet system. For simplicity, in this work, we assume that the land telescopes receive the same software updates as the LEOComs and the solar probes those of the GEOComs.

Telemetry in general, communcations , analysis, and storage

1. *Q: is it possible to have conflicting data?* A: Absolutely. There are unlikely to be absolutes in either scenario. The conclusions are always probabilities that an event will happe.
2. *Q: Where do the analyses happen?* A: In both cases, preliminary analysis is done on the satellites to make enough of a determination to decide whether or not the data sets should be sent with high priority to the ground and hence to the relevant organizations for more complete analysis and a determination of whether or not to recommend an alert.
3. Q: How automated are the alerts for either Solar Storms or land/mud slide predictions. As stated in the Solar Storms section, for all of these notification is quite automated through the cellular phone system, as well as email and any other means. Consider a similar situation at MIT and in the City of Cambridge. In cases of urgent messaging, messages are sent out appropriately blanketing whatever range is needed. These systems exist. You do not need to invent them. For cellular coverage for example in Cambridge, there are some sorts of alerts that are sent to all phones with a particular geographic area or for particular types of events. As a resident of Cambridge I am notified whenever there is a snow emergency. I can sign up to be notified when yard waste pickup and street cleaning start and stop. If there is a hazardous waste spill, any cell phone within the danger area is notified. One person makes a decision and all such notifications go out as appropriate.
4. *Q: How does information get to and from the ground-based antennas?* A: The Internet. That is outside the scope of what you are designing here and already exists.
5. *Q: How much additional storage is required for the data collection satellites?* A: As specified in Table 1, each of these has 10TB. It is estimated that with their normal transmission models, that will provide a stable state, although the usage may fluctuate as transmissions capacities are or are not available. Once data has been transmitted from one of these satellites it can be deleted.
6. *Q: Should high priority data be stored on the telemetry satellites differently than non-high priority data?* A: I'm not sure why it would be, but if you think so and can make a case for it, you may do that, explaining what the differences are and why.
7. *Q: What if there is both a land/mud slide and solar storm prediction simultaneously? Which takes priority?* A: The system as it stands does not know which messages are which. Remember priority 3 flows between A-B and A-C are not given an ordering inside the system as it stands. First, if you think it is necessary to make a choice of one over the other, you will need to make the case for that. In addition, you will need to propose how you will make that happen. Think about other alternatives as well.
8. *Q: How does the system prioritize high-priority bundles when capacity is limited?* A: This is closely related to the question above. It is your task to consider how you might do this. Is it

better to divert one bundle to another path or to fragment bundles and send them piecemeal and as a result more slowly but overall giving a fairer share to each? Are there other options?

Other Use Cases

1. *Q: How strict is the 10-minute delivery time for normal messages (Section 4.2a))?* A: Clearly, this is for non-critical communications. If it is 10 min, 5 sec. it is not a big deal. If it is 20 min. rather than 10, it is a bigger deal. If it is an hour, there's potentially a big problem.

Broad questions

1. *Q: Are there plans to make this international?* A: Good long term question. NASA certainly collaborates with other space programs. Clearly, the land/mud slide question in Kyrgystan is an example of international cooperation.
2. *Q: Are there other regulatory and policy questions?* A: Absolutely, but that is pretty far outside the bounds of what we can do in a DP project in 6.180.
3. *Q: Is space debris a consideration?* A: In general, absolutely. For your design, there are much more important things to consider. Set this one aside.
4. *Q: Can data be backed up on Earth?* A: You can assume that all data is stored and backed up persistently on Earth. The telemetry satellites will not keep archival data. They do not have enough storage for that, nor is that a reasonable place to keep it.
5. *Q: Could TCP be modified or improved to eliminate the need for the Bundle Protocol?* A: There is some discussion of this now in the IETF. It is in very early stages and will have to overcome the challenge of reliability with extremely long roundtrip times.
6. *Q: Is the Bundle Protocol useful beyond a scenario such as this?* A: Absolutely. It is actually a "delay tolerant protocol" and delays can happen in many other terrestrial types of situations as well, where communications are both intermittent and slow. Think of extremely rural areas. I believe that there are places in Africa and probably significant parts of the region including Kyrgystan and many of its neighbors, as well as regions such as Nepal, Tibet and mountainous India, etc., where communication is very fragile and slow. For all of those sorts of situations, the same problems arise and the solutions are equally applicable.