

## 6.033 2020 Design Project Frequently Asked Questions (FAQs) Updated on: March 11, 2020

1. Connectivity among the various satellites, relays, and antennas:
  - a. *Q: is it the case that there is a 90% chance that one LEO satellite has connectivity to a GEO satellite, or that every LEO satellite has connectivity with some GEO satellite?* A: The latter – each LEO satellite has a 90% probability of connectivity with some GEO satellite.
  - b. *Q: Can a Mars (or Moon) based antenna communicate directly with a GEO satellite, or must it communicate through a relay?* A: They can communicate directly, but only if there is a direct communication path. If the antenna is facing away from Earth, the relays will be necessary. Notice that the relays may also provide some resiliency and flexibility in transmission that may not be possible with the antenna that is on a rotating surface.
  - c. *Q: From where is the 30-minute update delivery time measured?* A: It doesn't matter. It could be from either the control computers on Earth or the antenna. The time it takes to deliver the software traffic through the terrestrial network to the antenna is negligible.
  - d. *Q: Can a node send and receive with other nodes simultaneously?* A: Each node has one laser and one receiver, so it can only be sending one bundle at a time and receiving one bundle at a time, but it can do those in parallel, in other words what is called "full duplex". The bandwidths given in the table are in each direction. This means that at any instant, a laser will be pointing at one receiver only, so for that period of time, that sending node will only be in contact with that receiving node. But, the "nearest neighbor" information provided by the routing table indicates which other nodes are valid potential links. Also, since the communication is duplex, A might have B and C as nearest neighbors and be sending to B while receiving from C.
  - e. *Q: Is there any flexibility with respect to orbits? Can they be changed?* A: Only NASA has control over them. Remember that the LEO satellites are in Polar orbits and quite low, and in 24 hours they will "cycle" back to the same spot. The GEO satellites are essentially in an Equatorial orbit and much higher, so they have a much larger range of reachability. The orbits of these satellites is determined by their altitude (remember your orbital mechanics).
  - f. *Q: Do all GEO satellites collect telemetry data? How is this managed?* A: Yes, they all do. Notice that there may be a lot of overlap in what they "see" of the Sun, except when they are blocked by the Earth. What they "see" of the Earth will be much more limited and generally quite stable, since they are in geocentric orbits.
2. Communication between Bundle Nodes:
  - a. *Q: Can a node send more than one bundle simultaneous (same for receiving)?* A: No, each bundle comes into or goes out on a single queue which can send/receive one bundle at a time, but input interfaces (antennas) and output interfaces (lasers) can operate independently of each other, so a bundle node can be sending and receiving at the same time
  - b. *Q: Is the bandwidth of a node the amount indicated in each direction (so, is a 1GB/s device that much in each direction or half that in each direction)?* A: One can get the full capacity in each direction, so a 1GB/s device can both send and receive at 1GB/s.
  - c. *Q: If two Bundle nodes are communicating and one has a limit of 622MB/s and the other has a limit of 1.2GB/s, how fast can they communicate?* A: They are limited by the slower device.
  - d. *Q: Can a Bundle node split its bandwidth and receiving from multiple other Bundle Nodes in parallel at lower rates?* Yes, sort of. A Bundle node will need to receive each whole bundle essentially atomically (all as a unit), but there is no reason that it can't receive a bundle from

one Bundle Node and then a bundle from the next as long as they are in communication range.

- e. *Q: Why do LEO satellites have less than perfect connectivity at the Poles?* A: Because there is more interfering activity, so there will be more noise.
- f. *Q: Where does custody information go?* A: First, see the update to the document and Errata for more detail. Second, in order to transfer custody a “receiving” custody node must notify its predecessor that it will take over custody. Optionally (if the Bit 15 flag is set), it may also notify the report-to node about progress.

3. Resources on a Bundle Node:

- a. *Q: Is the storage specified per device entirely disk storage? What is the read/write timing for the disk?* A: Yes, the storage is non-volatile disk storage. Each system will also have volatile memory, but that is much smaller. The volatile memory is large enough to hold a small number of bundles in addition to whatever the system on each device requires for normal operation. The read/write times are not specified. They are at least an order of magnitude shorter than the average time of direct connection between communicating nodes. If you need some more concrete value for these numbers, you will need to provide a realistic number and strong justification for it.
- b. *Q: Does the storage in this system need to support applications as well as the bundle protocol’s need for network storage?* A: You can assume that the antennas on Earth, the Moon and Mars are not the machines on which the applications are running, but are connected to them. So, an antenna needs to deliver traffic to some other device that will have its own storage. But...that is not true for the telemetry data collected on the GEO satellites. Their data will be being collected and stored in the only storage available until it is sent or deleted. You will need to consider this in your design.
- c. *Q: Is the storage on the antennas part of the distributed network storage system?* A: Yes.
- d. *Q: Do all the devices run the same operating system?* A: No, all the antennas run one OS. All the LEOs run a different one. All the GEOs yet another one. And the relays a fourth one. Note this means that they may not be running a Unix file system, but rather a specialized system to meet their particular needs. You need to be designing the storage system to meet those needs.

4. The Bundle Protocol:

- a. *Q: How reliable is the Bundle Protocol?* A: The basic Bundle Protocol is “best-effort”. We do not know with what probability a bundle arrived correctly at its destination. Additional reliability can be achieved through the use of the various administrative reporting bundles, at the cost of more traffic. But, there are two “unreliability” aspects to these options. First, as a bundle moves from one Bundle node to another, each one can decide unilaterally whether or not it will actually do any of this reporting. Its decision is independent of any other Bundle Node’s decision. In addition, the administrative bundles are only sent “best-effort” and hence may not arrive. Because the Bundle Protocol is “best-effort” you must assume that there will be some circumstances in which normal operation of the Bundle Protocol will be pretty unreliable, and possibly others in which most of the traffic is delivered. There are no known numbers for this. Providing estimates of reliability will increase the effectiveness of your argumentation for your design choices.
- b. *Q: Will a bundle continue to be transmitted after it has reached its lifetime or hopcount limit?* A: It might, but that would be a complete waste of resources. These choices about time and hopcount limit come directly from the application using the Bundle Protocol, so any bundle that has reached such a limit is no longer of interest to the application. So, it could be forwarded, but it is better dropped.

- c. *Q: Who determines fragmentation?* A: By default, the Bundle Protocol will determine the maximum size for a bundle to be sent in real time. This may be limited by connection time or possibly other factors. This automatic fragmentation will be done for you, leaving perhaps a fragment bundle in the Bundle Node for later delivery. In addition, you can decide that you want to require fragmentation of a bundle (or set of bundle) and request that. You will need to specify the size of the first fragment. (The second will be determined by the remainder of the original bundle.) If you want to do this, you will need to explain and justify your decision.
  - d. *Q: What happens if a bundle is sent on multiple paths and fragmented differently on these paths?* A: Every fragment is clearly labeled with exactly where it fits into the whole. As long as all the bits of the whole have arrived at the destination at least once, the whole bundle can be reassembled. Any duplicate bits are just that and can be ignored. If any set of bits does not arrive or is corrupted beyond repair, the bundle will not be received by the application at the destination.
  - e. *Q: When and where does fragment reassembly happen?* A: It can happen anywhere when all the fragments have been collected in a single location. This could be in a Bundle Node inside the network, or it could be only at the final destination. Note that in order for a custody node to accept custody of a custody bundle, it must have the whole, original bundle. It cannot accept custody of only some of the fragments of a bundle.
  - f. *What about custody transfer for bundles?* A: This had been updated in both the design document and the errata document. See those for more details.
  - g. *Q: How does the bundle protocol deal with dropped or rejected bundles? Does it resend them?* A: Remember that this is “best effort”. The Bundle Protocol itself does not provide increased reliability. In fact, it is the application that is using the bundle protocol that needs to decide what happens if a bundle times out, is dropped due to lack of storage space, or is dropped because it has been corrupted to the point that error correction can’t fix it. The Bundle Protocol cannot know what is best done when a bundle is not delivered.
  - h. *Q: Who sets the priority flags?* A: In practice the application in handing an ADU to the Bundle Protocol, will also indicate its priority. The Bundle Protocol itself actually composes the bundle. That said, if you think that the Bundle Protocol should be increasing priority for some reason, that is an option. It is not an option to decrease priority. Notice that it will also be the application that does or does not request custody handling for the bundle.
  - i. *Q: How is time kept track of, for lifetime expiration?* A: There are two parts to this. The ID of a bundle includes its creation time. In addition, there is a field in the primary block that includes a lifetime value. When a bundle is fragmented, each fragment has an “expanded” id which includes the original and where it fits into the whole. Thus, it is always possible to look at a bundle, whether it is a fragment or the whole, and know whether or not it has expired.
  - j. *Q: Can we create bundles to send to other nodes?* A: Yes, if, for example, you need a node to communicate directly with a neighbor or some other node, you can do that. Think of it as your “application”, which needs some communication.
5. Routing:
- a. *Q: Does the routing table figure out things like orbital trajectories, or do we need to figure that out?* A: The routing protocol will provide you with links and how long they will be available (from the time you ask). So, for each destination, the routing table will tell you the address of the next Bundle Node in the path toward the destination. It will also tell you your two closest neighbors within the same type of satellite that you are (e.g. two nearest LEO neighbors if you are a LEO satellite, assuming there are two within communication range). You do not need to do any such calculations. Note that this means that if you are a LEO

satellite and are moving a bundle toward the Moon, it will give you the next GEO satellite on the path, you can opt for 1) to send it to that GEO satellite, 2) to send it to one or both of the neighbor LEO satellites, or 3) some combination of all of the above. You get to decide and discuss/justify your choice.

- b. *Q: Can each device know the whole topology of the network?* A: it is unlikely and you can't depend on it. Some of the trajectories are completely predictable, but some, especially relays may move around. Satellites could probably be moved to different orbits under direct command as well, or may be offline.
- c. *Q: What about custody routing?* A: Again, see the updates to the design document and the Errata. The routing table will also provide you with the next custody node address, in case the bundle requires custody. It will also provide the next hop address toward that custody node and the amount of time it will be available.
- d. *Q: How much does each node know about all the other nodes?* A: It does not know full connectivity at all times. With respect to connectivity it only knows what the Routing Protocol has installed in its routing table. But, it does know some additional static information. It knows the IP address of every node in the system. It also knows the type of that device (antenna, LEO, GEO, or relay).

6. Use cases:

- a. *Q: Is the video possibly realtime streaming (two-way)?* A: Not necessarily, especially to Mars where the round-trip time will be a minimum of between 8 and 40 minutes.
- b. *Q: The software update use case requires that all messages be received within 30 minutes, but it can take up to 20 minutes for a message to be sent to the mars. If the message isn't successfully received, how do we meet the 3- minute deadline?* A: If a message fails to be delivered to Mars, that does not necessarily mean that it must be resent from Earth. If the network storage is available, the bundles that make up the update could be stored along the way in various nodes. This may be a case where a few more relays along the way would help to solve the problem, because of the delays. If you believe that this is important to your design, justify that choice.
- c. *Q: In the routine communications use case, if a bundle is past its deadline, should it be delivered?* A: Most likely not. As mentioned above, if the application has determined that it is not useful after a specific time, it probably should not be delivered. If you can think of situations in which it would still be useful, explain and justify your choice.

7. General questions:

- a. *Q: What can you tell us about future extension plans?* A: Nothing is known at this time other than that something is likely to come along. More satellites or relays, larger capacity satellites or relays, more applications, other more distant planets, etc. Everything is possible.
- b. *Q: Can we make some assumptions about things that are not specified?* A: Yes, if there are facts you believe are necessary, you can make and justify assumptions about what you need to know. They should be reasonable.
- c. *Q: At what layer does the distributed storage system sit?* A: it is orthogonal to the protocol layers. It is a service that is provided by the OS to meet the needs of the protocol (or any protocol layer that chooses to use it). Notice that although that is true, because the Bundle Protocol is the only game in town, for the distributed storage system to be distributed it must use the Bundle Protocol, despite that the Bundle Protocol is also using the storage.

- d. *Q: What can you tell me about Solar Flares and predicting them?* A: The increases in measurements at best allow for a 24-48 hour prediction of a Solar Flare. In addition, a Solar Flare can last for 24-48 hours. That said, there have been times when up to five were seen in a 48-hour period.