Modeling Interplanetary Communications after Telecommunication Networks, With Layering and Dynamic Satellite Management

Examiner Jeffrey Nickerson, USPTO

Abstract: Interplanetary communications can be improved by modeling satellite management after current telecommunication networks. Long-term, continuous bidirectional communication sessions can be established between planets by allowing satellite layers to perform location, trajectory, and timing based handoffs. Satellites may be dedicated to specific layers in order to minimize chokepoints and maximize the duration of point-to-point communications. Bandwidth for surface-to-layer, layer-to-layer, and planet-to-planet communications can be drastically improved by dynamically adjusting satellite positioning and functionality.

Introduction

Current interplanetary communication systems acknowledge that satellites orbiting a planet can be utilized as an endpoint in a point-to-point communication session. Typically, these point-to-point sessions happen between a fixed ground station (FGS) on a first planet and a fixed ground station on a second planet. The FGSs either directly transmit toward their target endpoint, or target an orbiting relay satellite, which then directs the transmission to the target endpoint. However, the duration of the communication sessions is limited by the overlapping rotation periods for either the FGSs (when transmitting directly from a source FGS to a target FGS) or the relay satellites (when transmitting from a source relay satellite orbiting a first planet to a target relay satellite orbiting a second planet). The overlapping periods are usually very short, on the order of a handful of hours of an earth day. As a result of short time windows for transmissions, outgoing data must be queued for transmission, and large data stores and backlogs are required at the source FGS or source relay satellite in order to store the communication information for transmission during these short windows.

Further consideration must be made regarding the bandwidth of links in the network, as some links are inherently choke points for communications. Direct communication between Satellite of planet A to satellite of planet B, and direct communication from FGS of planet A to FGS of planet B are typically on the magnitude of kilobits per second. Communication links between an FGS of planet A to a satellite of planet A tend to be on the magnitude of megabits per second. Therefore, the choke-point of communications in the overall network topology is between the satellite relays of differing planets, further amplifying the need for large data stores and backlogs at the source relay satellite.

Various schemes and considerations will be discussed to allow for both continuous transmissions between planets and to maximize bandwidth at all chokepoints.

Notations

Satellite_A = Set of satellites orbiting planet A at a first layer.

Satellite_B = Set of satellite orbiting planet B at a first layer.

Satellite_AA = Set of satellites orbiting planet A at a second layer.

Satellite_BB = Set of satellites orbiting planet B at a second layer.

Satellite_An = The nth satellite in the set Satellite_A. Example: Satellite_A1

FGS_A = Set of fixed ground stations on planet A.

FGS_B = Set of fixed ground stations on planet B.

FGS_An = The nth fixed ground station in the set FGS_A. Example: FGS_A1

Background

Each satellite in a layer (a layer such as Satellite_A) are treated as nodes in a mesh network, and maintain all required information in order to communicate to adjacent nodes. This may include routing tables, trajectory tables, and location tables of neighbors, so that a satellite may know which neighboring satellite to pass the communication to in order for it to reach its ultimate destination, the exact trajectory (relative speed, relative direction, absolute speed, absolute direction,) and location (relative location, absolute location) so that a communication may be effectuated. This can be handled by, for instance, calculating where satellite will be based on a time-associated past known location and a time-associated known trajectory. Standard routing protocol techniques, such as heartbeat messages, may be utilized to maintain up-to-date routing, trajectory, and location tables. Time values may be associated with all trajectory and location information to identify its age and validity. Considerations to reduce route flapping can also be realized and implemented. Unique to satellites, however, is that they can change their trajectory and relative physical location to both other nodes in the network and the planet, by using thrusters, etc.

Every satellite in a layer is aware of the physical topology (and overlay topology, if one should exist) and physical relativity to the planet of every other satellite in the layer at any given point in time. Trajectory update information may be broadcast periodically to neighbors from a single node. At which point, a receiving neighbor may broadcast to his neighbors. Since the trajectory information contains a unique identifier for the satellite that originated the broadcast, and each node knows the physical layout of the entire layer, the receiving neighbor knows to which neighbors to broadcast the message (the direction away from the originator). Every receiving node checks to see if its distance in the 3D orbital layer from the originating node is greater than half the circumference of the cross-section of the layer (assuming the layer is the same elevation). If so, the node stops propagation, as the nodes in the direction away from the reception direction will be heading back towards the originator and these nodes will receive their propagation message from the other side. If the satellites making up the layer are not the same elevation, then satellites may need to only broadcast updates once to each physical neighbor (since each satellite knows the entire physical layout of the layer). This may be realized with regard to each planet (or moon or space object) and to each satellite layer. When changing trajectory, satellites should broadcast their new trajectory and a time and location pair with that trajectory, shortly after taking it.

Source Planet to Layer 1 Communications

Static coverage areas exist when the satellites orbiting a planet are in geostationary orbit. That is, a FGS_A is always talking to the same subset of Satellite_A. For example, FGS_A1 may be capable of talking to Satellite_A1, Satellite_A2, and Satellite_A3, because these satellites are always within a certain angle range with respect to FGS_A1's transceiver and in geostationary orbit.

Coverage areas on planet A are broken up via hexagonal areas (any shape may be recognized) on the surface of the planet. Inside a first hexagonal coverage area is a subset of FGS_A, for example FGS_A1 and FGS_A2. These coverage areas may cover hundreds of miles on the surface, or be very concentrated (a single military base, several miles) depending on the elevation of the Satellite_A layer. A subset of Satellite_A is in geostationary orbit above the first hexagonal coverage area, for example Satellite_A1, Satellite_A2, and Satellite_A3. Cross coverage may be possible. That is, Satellite_A3 may be capable of communication with two adjacent coverage areas. The entire planet could be broken up into coverage areas, with multiple satellites in geostationary orbit above each area.

The satellites in Satellite_A monitor communications and bandwidth use for their respective coverage area. This may be done in a decentralized fashion, or a single satellite for each coverage area may be designated as a master management satellite. This designation may be dynamic and be maintained similar to the Bluetooth promotion/demotion for master/slave statuses. Assume a centralized master satellite for each coverage area. One may recognize that one coverage area may have extremely heavy traffic while another coverage area may have infrequent traffic. The centralized master satellites for each coverage area may communicate their network statistics usage histories with one another and map a general network usage topology of the entire Satellite_A. Each of the master satellites, or one designated master satellite, may then generate a social network model of the Satellite_A layer, assigning things such as centrality, connectedeness, etc, to each satellite or coverage area, based on the collected network statistics. By doing so, the Satellite_A can realize which coverage areas are being utilized the most and which are not, usually using some scale and a function of weighted social networking attributes. The social network modeling make be done periodically and a history can be kept. Trends may be analyzed in the models to determine if certain coverage areas are generating heavy traffic during certain times (which may be days, years, decades, etc).

Since satellites may alter their trajectory and ultimately their geostationary orbit location, satellites are capable of moving to the entire other side of a planet. This could be user commanded from a FGS to move a satellite to cover the area of a newly constructed FGS on the other side of the planet. Or, the satellites may manage their location by themselves using the social networking modeling and physical topology awareness. Satellites can be implemented with artificial intelligence to ensure collisions do not occur with other satellites in Satellite_A, since each node knows the physical location and trajectory of all other nodes in the network. Other detection mechanisms may be employed, similar to current autonomous car technology, to ensure no collisions occur with satellites external to Satellite_A. The entire Satellite_A network may therefore converge towards "hot spot" locations that have heavy traffic, to enable higher bandwidth capabilities to these coverage areas. Since history trends of traffic use may be analyzed, satellites may move towards a specific coverage area prior to a known timeframe that receives heavy traffic. This is especially useful when it may take extended periods of time for a satellite to alter its position to the correct location. When this converging occurs, the coverage areas on the surface area may be contracted or expanded. If more satellites are in a specific coverage area, the coverage area may be split into multiple, smaller coverage areas. Alternatively, coverage areas that receive very low traffic during a certain timeframe may be monitored by only one satellite during these times. And, if several adjacent low traffic coverage areas exist, the same single satellite may monitor the adjacent coverage areas, essentially expanding the coverage areas to a single coverage area.

Using the aforementioned techniques, it can be realized that the layer 1 satellite network may be dynamically adjusted to provide optimal bandwidth and coverage to a planet, allowing for efficient transfer of communications between the surface and the layer 1 satellite network. Each planet may have a layer 1 network accordingly. Furthermore, when dynamically adjusting, handoffs may be made between the layer 1 satellites to accommodate ongoing communications, similar to that of telecommunication networks.

Layer 1 to Layer 2 Communications

A second layer of satellites orbiting a planet may be realized. These satellites may or may not be in geostationary orbit. However, like the layer 1 satellites, all the satellites in this layer maintain physical topology layouts of the entire layer, including trajectory and routing information.

The layer 2 satellites are preferably orbiting at a far slower speed than the planet's rotation (an orbital period of the planet's months or years), dynamic role switching can be utilized between satellites in preferably both layers 1 and 2. Select satellites in each layer may to be dynamically selected as gateways to the other layers. Both layer 1 and layer 2 satellite networks may have gateways, or only one layer or the other may have a gateway. If both layers have gateways, the layers may be connected in a point-to-point fashion (same number of gateways) or multipoint-to-point (layer 1 has more gateways than layer 2), or point-to-multipoint (layer 2 has more gateways).

The gateways are dedicated to communicating to another layer's gateway. The gateways maintain trajectory and location information of every gateway within their layer and with every gateway in layers with which they communicate. This allows gateways to choose a best fit target gateway (based on relative location) for the transmissions.

These gateways maintain traffic usage information between their layer and the destination layer. Again, each set of gateways in a layer may have a master gateway, who can manage the traffic and network usage information and generate models accordingly. This master gateway may alter the number of gateways in its layer based on demand and usage. Also, the master gateway may direct other gateways to move to the areas around the planet that are in high usage. Layers may exchange traffic history information and therefore the layer 2 master gateway may be able to base the number and location of layer 2 gateways on its own, without monitoring layer 1 to layer 2 communication trends.

The purpose of gateways within layer 1 is to receive communications from the coverage satellites in layer 1 (coverage satellites may also be gateway satellites, dynamic role switching, etc), and then feed the communication to a layer 2 gateway. This means that non-gateway satellites in layer 1 should route their communications to a best-fit layer 1 gateway (relative location), based on their routing and trajectory and location tables. Once layer 2 gateways receive the communication from a layer 1 gateway, they relay the communication to a selected layer 2 satellite.

All communications may be bidirectional and the opposite may occur for receiving an interplanetary communication, as one would realize.

Source Layer 2 to Destination Layer 2 Communications

Layer 2 satellites are tasked with sending communications over an extremely far distance to a destination planet's layer 2 satellite network. Therefore, layer 2 satellites maintain not only all trajectory and location information of their own layer, but also all trajectory and location information of the destination planet's layer 2 satellites.

Terminology:

Line of Sight entrance point – The earliest point in time of a satellite's orbit that the satellite has a clear line of sight with a target satellite.

Line of Sight exit point – The latest point in time of the satellite's orbit that the satellite has a clear line of sight with the target satellite.

Interplanetary transmission delay – Time of delay for a communication to travel from source satellite orbiting a first planet to a target satellite orbiting a second planet. This dynamically changes based on where the satellites are in their orbit, and the relative orbits of the planets.

Valid reception phase – The location of a receiving satellite's orbit in which it can receive a communication from a transmitting satellite orbiting another planet. Based on interplanetary transmission delay and line of sight entrance/exit points.

Valid transmission phase – The location of a transmitting satellite's orbit in which it can transmit a communication to a receiving satellite orbiting another planet. Based on interplanetary transmission delay and line of sight entrance/exit points.

Select signaling satellites within layer 2 are dedicated to continually receiving trajectory and location information from the destination layer 2, and select satellites are dedicated to continually transmitting trajectory and location information. Due to a very high delay in transmission time between planets (for example, speed of light between earth to mars is in the magnitude of minutes), the transmission of trajectory and location information of satellites entering their communication phases should occur a certain amount of time before they actually enter the communication phase. For example, Satellite_AA1 wants to send a communication to Satellite_BB1. Satellite_AA1 can only communicate when both Satellite_AA1 and Satellite_BB1 have entered a valid communication phase of their orbits. Trajectory and location information of Satellite_AA1 is passed via the selected signaling satellite of Satellite_AA to a signaling satellite of Satellite_BB, and vice versa, at a predetermined time period prior to Satellite_AA1 and Satellite_BB1 entering their valid communication phases. Once Satellite_AA1 and Satellite_BB1 enter their communication phases they *should not* alter their trajectory or location, as the transmission delay does allow for real-time trajectory and location adjustments between planets.

A master layer 2 satellite manages its planet's layer 2 signaling satellites. As the role of a signaling satellite must be assigned to a satellite in a valid communication phase, the signaling role should be handed off from a satellite just prior to exiting its valid communication phase to a satellite just entering its valid communication phase. The master layer 2 satellite can direct this handoff to occur smoothly on its end and should instruct the current signaling satellite to transmit timing information (a future time), trajectory, and location information of the next signaling satellite to the destination planet's signaling satellite. The destination signaling satellite of the target planet receives this handoff information and knows to start transmitting and receiving signaling information with a new signaling satellite, of notified trajectory and location, at the point in time the timing information designates. There may be more than one master layer 2 satellite managing more than one concurrent signaling satellite, for one or multiple interplanetary links. In no event should a signaling link be ended once started (This paper does not address how to initiate a continuous signaling link), as restarting continuous signaling links is likely an extremely difficult task that will take manual involvement. A master layer 2 satellite may therefore direct layer 2 satellites to move into position so that they can be assigned the role of signaling satellite without losing connection of the signaling link.

As with the handoff of signaling satellite roles and notifying destination planet of the handoff, the same must occur with all layer 2 satellites prior to entering and exiting their valid communication phases. The signaling satellites of each respective planet inform each other of timing, trajectory, and location information of layer 2 satellites at a predetermined time prior to them entering their communication phases.

Like all other intercommunications, layer 2 master satellites monitor traffic usage between planets and dynamically adjust satellite locations and trajectories accordingly. Preferably, there are tens or hundreds of satellites, each with a very long orbital period and in a point to point communication with a destination layer 2 satellite, which is also accompanied by tens or hundreds of destination layer 2 satellites, each with a very long orbital period and in point-to-point communications. Once a layer 2 satellite exits a valid communication phase, it may accelerate its orbit and adjust its trajectory towards another valid communication phase. At which point, it will slow down and re-adjust as self-calculated or specified by a master layer 2 satellite. The second valid communication phase may or may not be the same communication phase the satellite just left. If the layer 2 satellites are in communication with multiple different planets on different sides of the planet they are orbiting, the satellite may be accelerated towards a valid communication phase associated with a different destination planet. If certain interplanetary links are predicted to receive heavy traffic (based on current traffic analysis, modeling based on social networking theory, or historical network usage trends), then the multiple layer 2 satellites may converge towards their specific valid communication phase with that planet, in order to maximize bandwidth provided for the interplanetary link.

All routing of layer 2 satellite communications takes into consideration where the communication is ultimately headed (such as a specific planet) and routes the communication data structure (which may be packets, etc) to a layer 2 satellite in a valid communication phase with the correct destination planet. A master layer 2 satellite may act as a distributer to make sure that load distribution among all satellites in valid communication phases are evenly distributed. Alternatively, the entire layer 2 satellite network may be decentralized, where each satellite maintains information about which satellites are in which valid communication phases and load information, and distributes and routes communications accordingly. Furthermore, each satellites may recognize transmission delays within their layer by managing histories that map delays to relative locations of the planet they're orbiting. This allows, for instance, a layer 2 satellite attempting to send a communication to a target planet, to identify a satellite will be leaving a valid communication phase shortly, and not route packets to it if the transmission delay would be such that the satellite would have left the valid communication phase.