
User Applications on a High-Latency Network

Scott Burleigh

24 January 2014

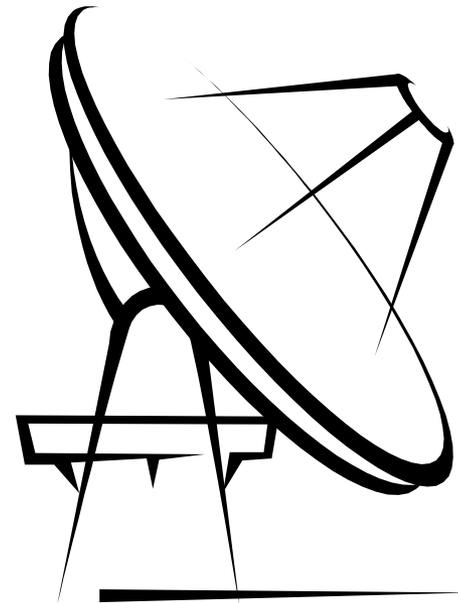
This research was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration. (c) 2014 California Institute of Technology. Government sponsorship acknowledged.

Outline

- An infrastructure proposal: a constellation of nanosatellites using delay-tolerant networking to provide low-cost access
- An illustration
- Some details: capacity, costs
- Application latency in this network
- Some applications that would work despite the latency
- A perspective on using a network
- Caveats and outlook

Satellites for Universal Network Access

- Earth-orbiting satellites can relay radio communications among sites on Earth.
- Can be visible from all points on Earth's surface, removing geographic and political obstacles.
- Not a new idea:
 - Geostationary (GEO): Exede (ViaSat), HughesNet (EchoStar), WildBlue, StarBand, Intelsat, Inmarsat, Thuraya
 - Low-Earth Orbiting (LEO): Globalstar, Iridium, Orbcomm, Teledesic



So, Problem Solved?

- Maintaining Internet connections with satellites isn't easy.
- GEO satellites do this by ensuring continuous radio contact with ground stations and customer equipment. But:
 - They are costly, on the order of \$300 million (manufacture & launch).
 - Each one provides communication to a limited part of Earth's surface.
 - Each one is a single point of failure.
 - While data rates are high, round-trip latencies are also high.
- LEO constellations do this by constantly switching connections among moving satellites.
 - Broad coverage areas, low latencies.
 - But data rates are lower than for GEO, more satellites are needed, and they're still expensive: \$150-\$200 million (manufacture and launch).

An Emerging Answer

- A capable nanosatellite, such as a CubeSat, could be fabricated and launched into low-earth orbit for .1% of the cost of a LEO communication satellite.
- An inexpensive nanosatellite constellation in LEO would have a broad coverage area, no single point of failure.
- There's a catch, though....

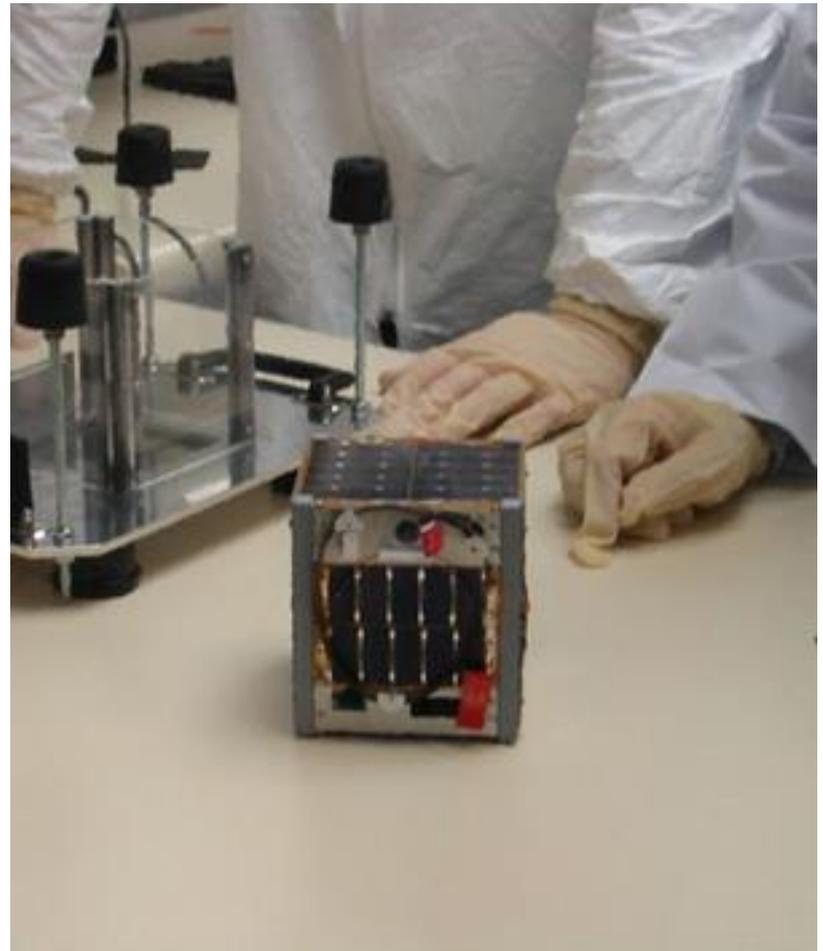


Photo used by permission: <http://www.cubesat.org>

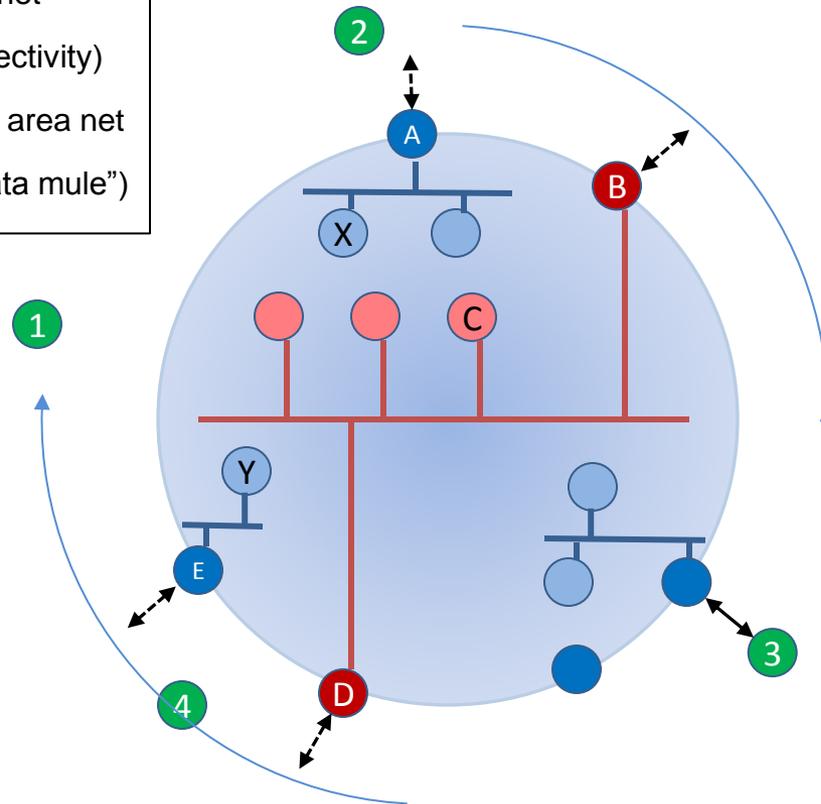
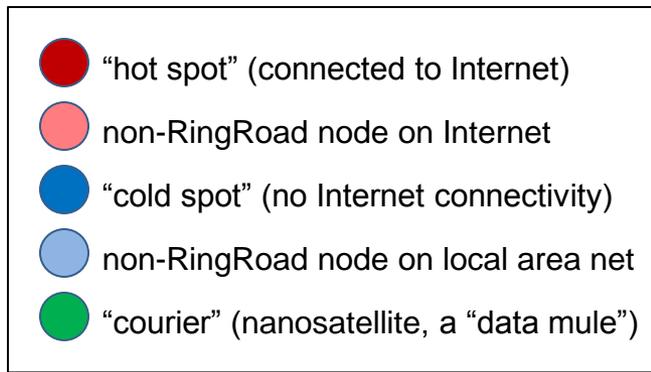
“Nano” Communication Satellites

- Nanosatellites are inexpensive in part because they are simple.
 - Maneuverability is limited.
 - Normally only a single radio (UHF or S-band).
- This means the constant connection switching among cross-linked satellites that normally enables LEO constellations to maintain Internet connections is not feasible.
- So we can't expect a constellation of low-cost nanosatellites to be able to sustain thousands or millions of continuous end-to-end Internet connections.

Solving that “Catch”

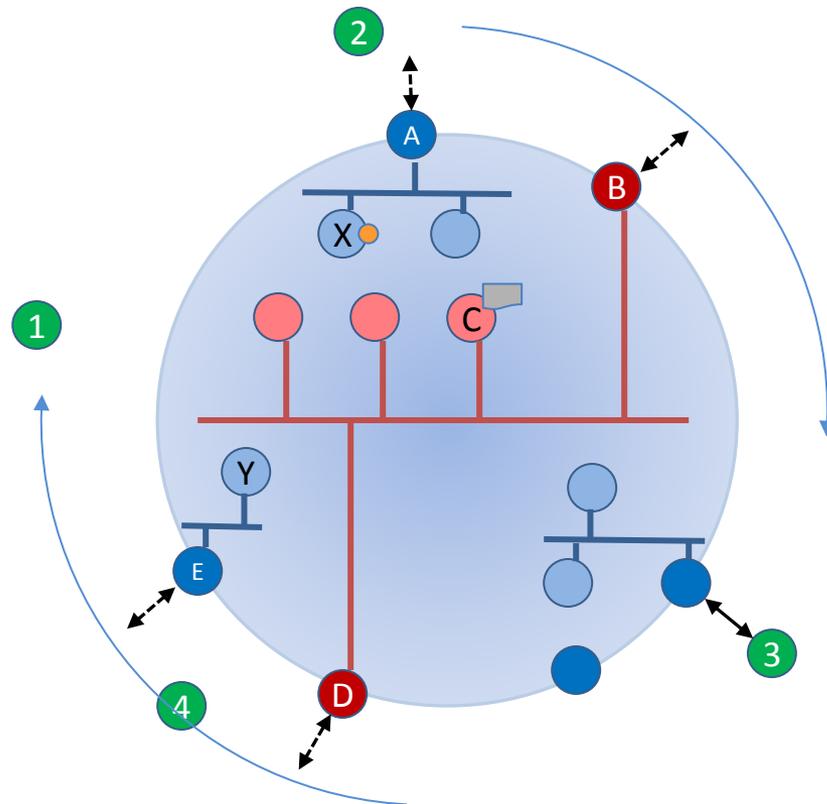
- But maintaining continuous end-to-end Internet connections is not the only possible network service.
- ***Delay-Tolerant Networking (DTN)*** technology, emerging over the past decade, would enable a different model:
 - At each moment, each satellite’s radio would be in contact only with the ground station directly below it (when there is one).
 - While the satellite and ground station are in contact, they would exchange network traffic. The satellite would function as a router.
 - When contact is broken, *the router satellite would retain outbound messages in local storage while it continues on its orbital track.*
 - When contact with the next ground station begins, network traffic exchange would resume.

The “Ring Road” Concept



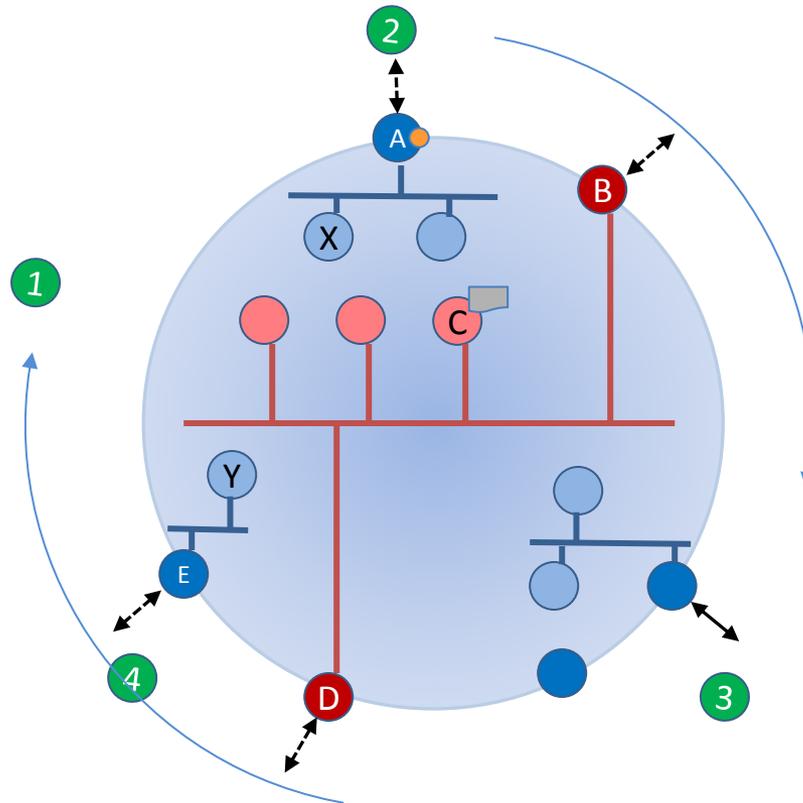
Downloading a File

User at X requests a file from Internet host C.



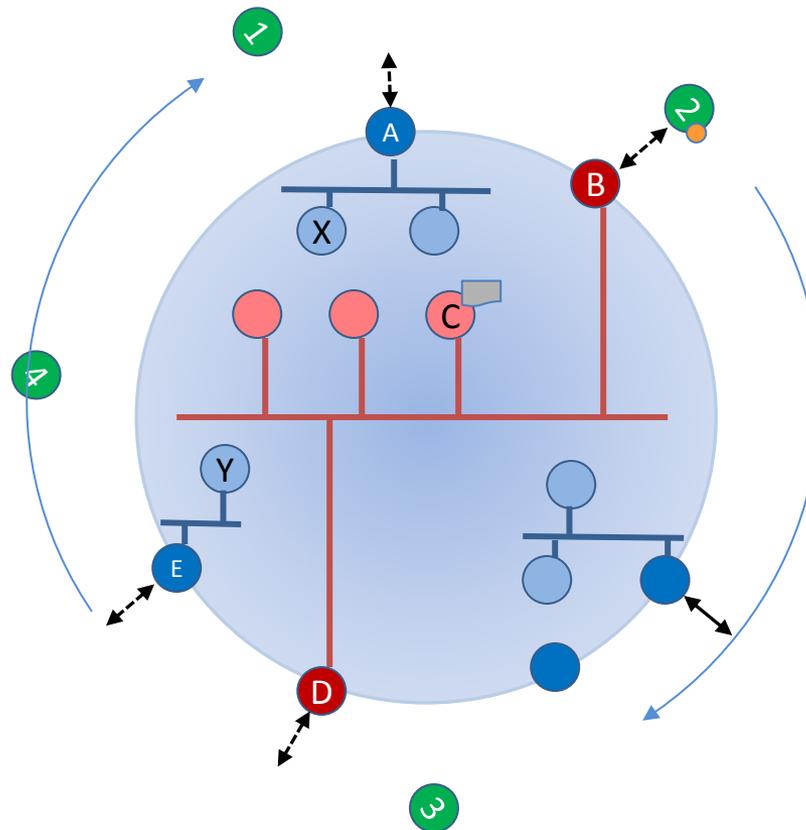
Step 1: request is routed over LAN to “cold spot” A, where it waits for a satellite overflight.

Downloading a File



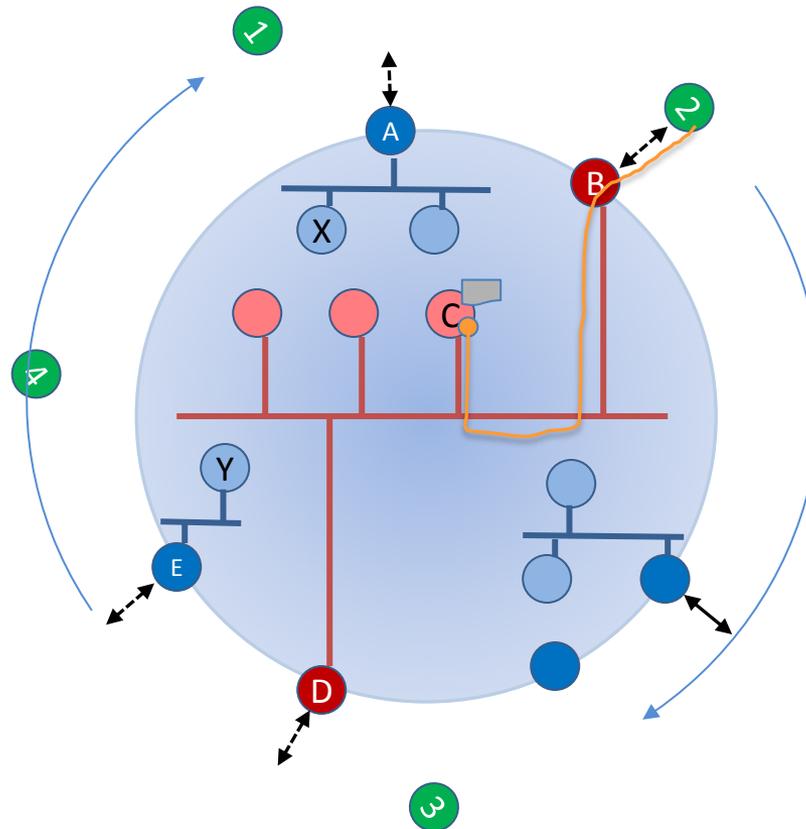
Step 2: request is uploaded to courier 2, where it waits for contact with a “hot spot”.

Downloading a File



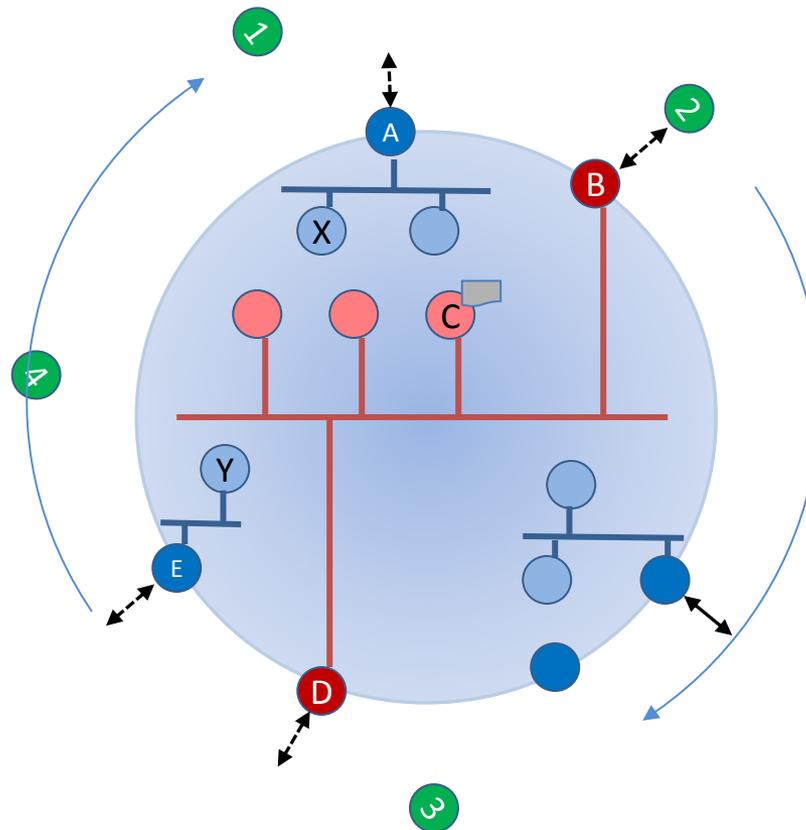
Step 3: request is downloaded to “hot spot” B, which immediately forwards it to server C.

Downloading a File



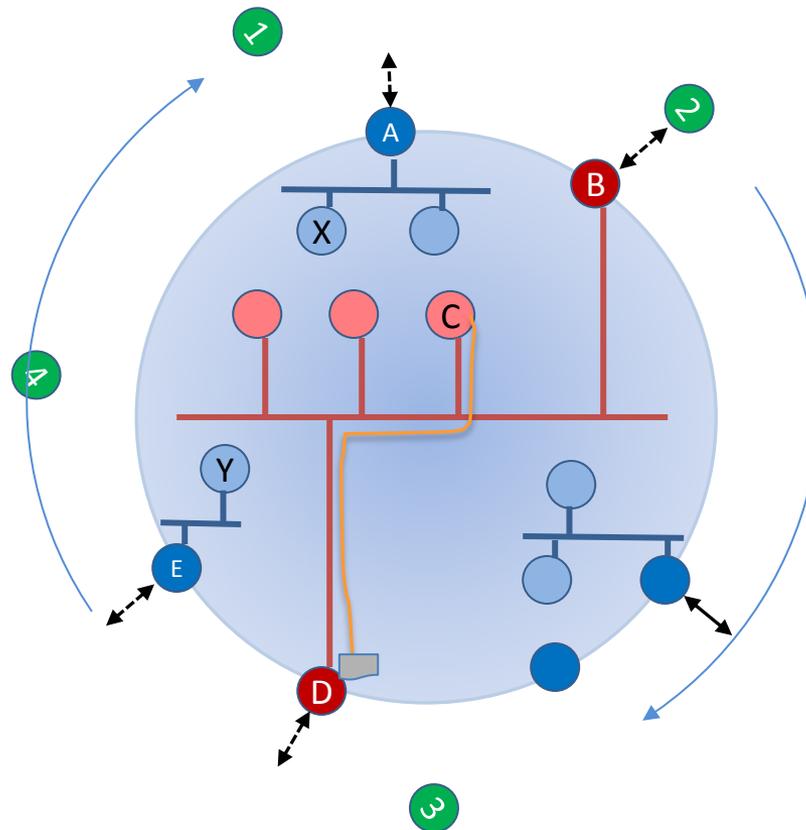
Step 3: request is downloaded to “hot spot” B, which immediately forwards it to server C.

Downloading a File



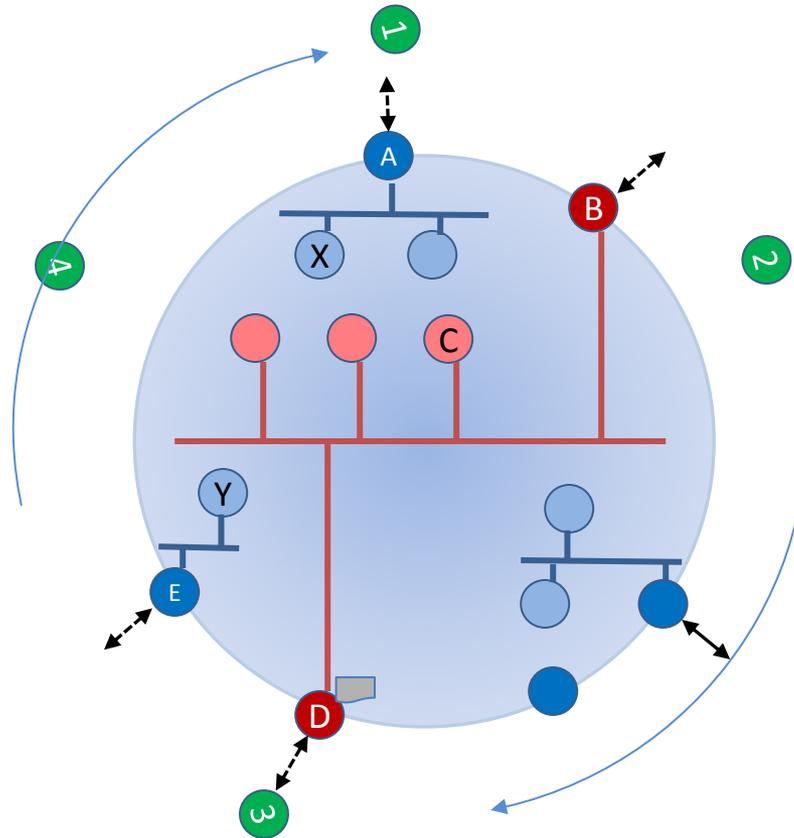
Step 4: server C determines that the earliest transit to X is by courier 3 via “hot spot” D.

Downloading a File



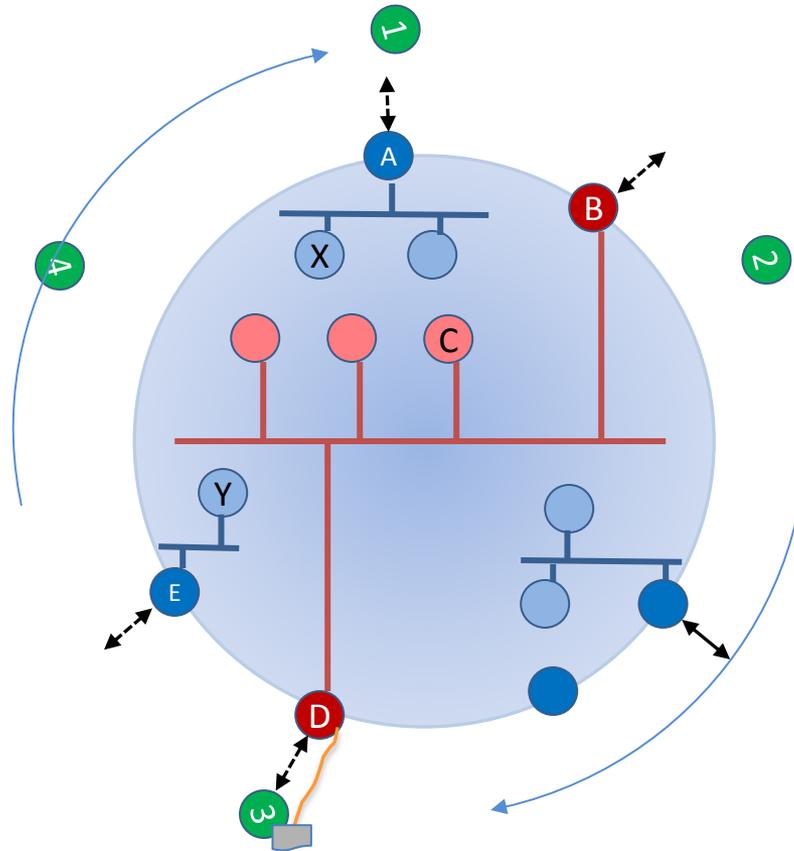
Step 4: server C determines that the earliest transit to X is by courier 3 via “hot spot” D.

Downloading a File



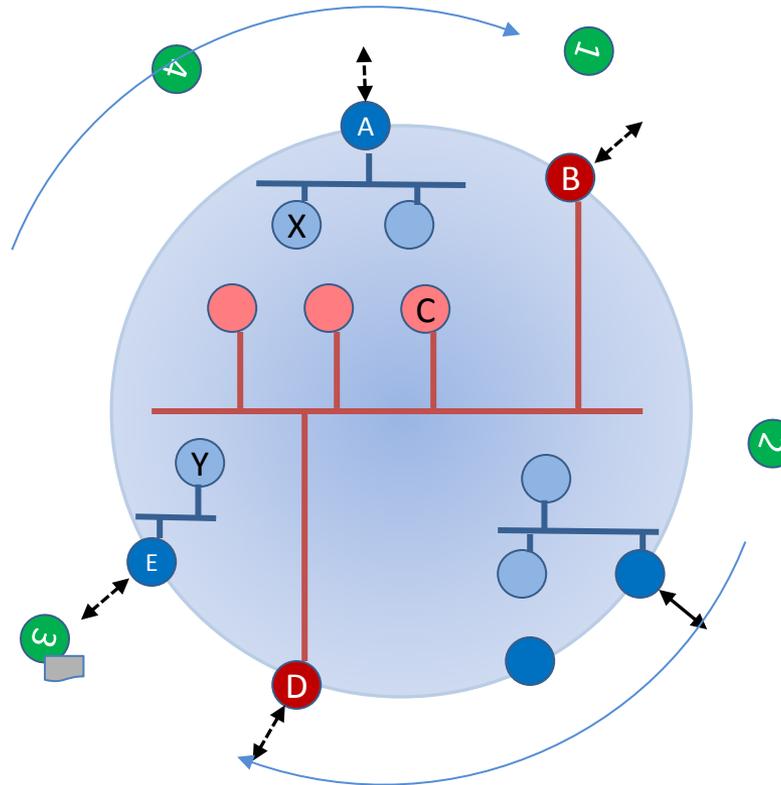
Step 5: “hot spot” D uploads the file to courier 3.

Downloading a File



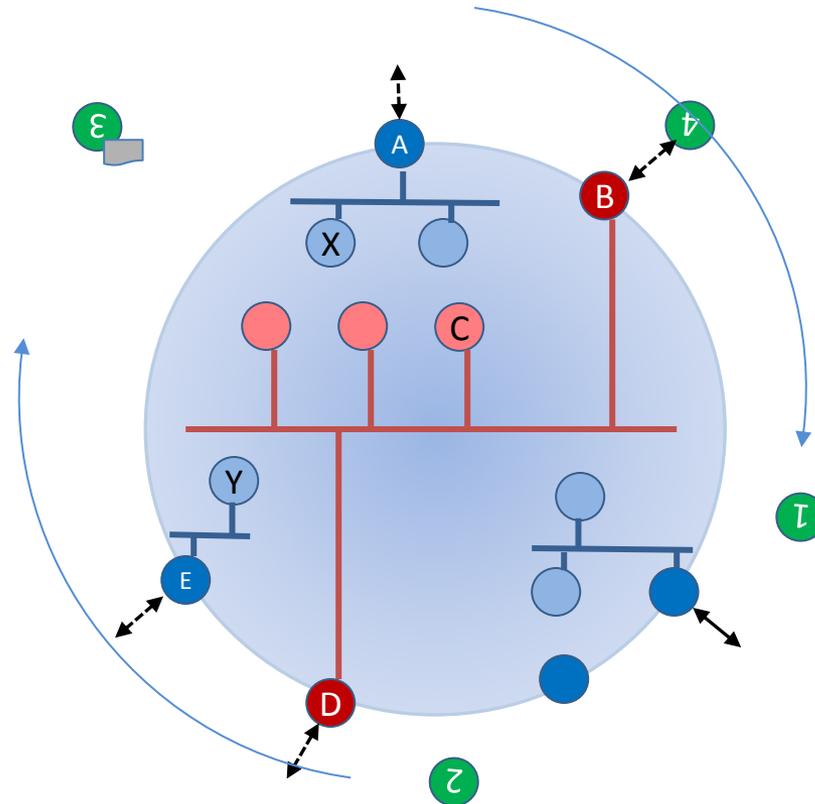
Step 5: “hot spot” D uploads the file to courier 3.

Downloading a File

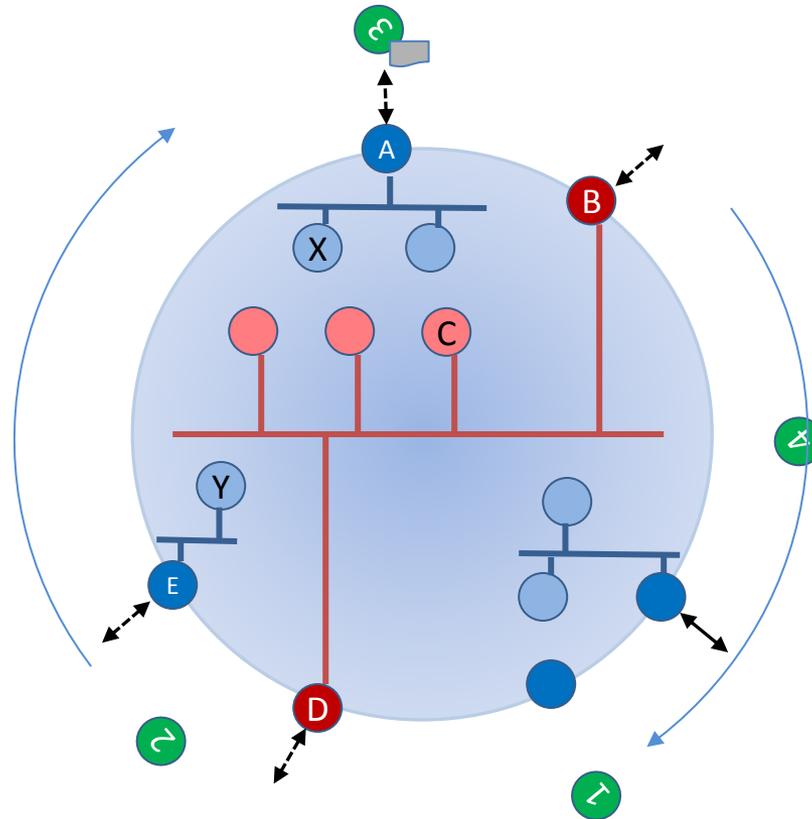


(Contact with “cold spot” E can’t contribute to this data exchange.)

Downloading a File



Downloading a File



Step 6: courier 3 downloads the file to “cold spot”, which immediately routes it to client X.

Proposed Constellation Size

- Couriers would be in near-circular low-Earth orbits at inclination of about 50 degrees (so the coverage area would be from 50 degrees South latitude to 50 degrees North); altitude about 500 km.
- Radius of satellite visibility would be somewhat greater than 1000 km.
- Maximum separation of satellites would be at the equator, which is about 40000 km, so 10 orbital planes would be needed: each satellite would cross the equator twice per orbit, covering a circle of 2000 km diameter at each point.
- Coverage per satellite at the equator would be 18 degrees of longitude; 12 satellites per orbital plane would allocate 18 degrees of latitude coverage (out of 200) to each satellite, but say 15 to ensure some overlap. So 150 spacecraft in all.

Network Capacity

- Satellites at 500 km altitude travel at about 7.8 km per second, an orbital period of 90 minutes; 16 orbits per day. Any single satellite would be in view of any single ground station for about 256 seconds per orbit, the maximum contact length.
- S-band transceivers can transmit at 230 kbps, could upload up to 7.4 MB during a single satellite contact of maximum length.
- Up to 20 contacts of max. length per orbit, so almost 150 MB could be inserted into the network per satellite per orbit, 2.4 GB per satellite per day.
- For 150 satellites, 360 GB per day. This is about 4 MB per second, 32 Mbps. But maximum upload would be impossible because each satellite would often be over open sea; figure 16 Mbps.

Deployment Cost

- Estimated assembly cost per satellite: \$100,000.
 - Total satellite assembly cost for network: \$15 million.
- Estimated launch cost per group of three CubeSats: \$200,000.
 - Total launch cost for network: \$10 million.
- 20 contacts per orbit (40,000 km Earth circumference, 2000 km diameter per contact) for 150 satellites would require 3000 base stations; 1500 if one-half of the covered surface is unoccupied. Estimated cost per base station, including S-band radio modem: \$2000.
 - Total base station cost for network: \$3 million.

Service Cost Estimate

- Suppose network lifetime of 5 years and operating cost of \$1 million per year. Total lifetime cost of the network would be $\$15 + \$10 + \$3 + \$5 = \$33$ million.
- At 16 Mbps, total lifetime traffic would be about 300 TB. Mean cost of transmission would be less than *\$0.11 per MB*.
- Commercial satellite services (as of September 2013):
 - Iridium Pilot: \$7.41 per MB
 - Inmarsat BGAN shared service: \$6.49 per MB
 - Thuraya GmPRS NOVA: \$5.00 per MB

Latency

- Total network capacity and data rate to subscribers could be increased 100-fold or more by using C-band or Ku-band radios, not currently available for CubeSats (but coming).
- But higher data rate can't solve the end-to-end round-trip time problem: latency in communication between satellites and ground stations would be negligible, but round-trip latency in the network would be extremely variable and could be very high.
 - Suppose a subscriber is on Seram Island in Indonesia. The preceding hot spot on the Ring Road track is, say, Manado (800 km distant) and the next one is perhaps Darwin, 1200 km distant. Total round-trip time for an Internet database query is then $1200 / 7.8 = 154$ seconds plus $800 / 7.8 = 103$ seconds, a total of 4.3 minutes.

Potential Applications

- Can a network be usable when round-trip time is 5 minutes?
 - DTN was originally devised to enable automated network communications in interplanetary space flight operations, where the round-trip time between Earth and Mars varies between 8 and 40 minutes. (Which is why it was called “delay-tolerant networking”.)
- Some network applications are useless over very long round-trip latencies: Skype, teleconferencing, telnet, ssh.
- But other applications are innately delay-tolerant:
 - Email. We send email to a distant server; addressee picks it up later.
 - News feeds, weather advisories, crop reports – anything that is a unidirectional data flow.
 - Archiving, backups, documentation of events – any communication where reliability and security are more important than timeliness.

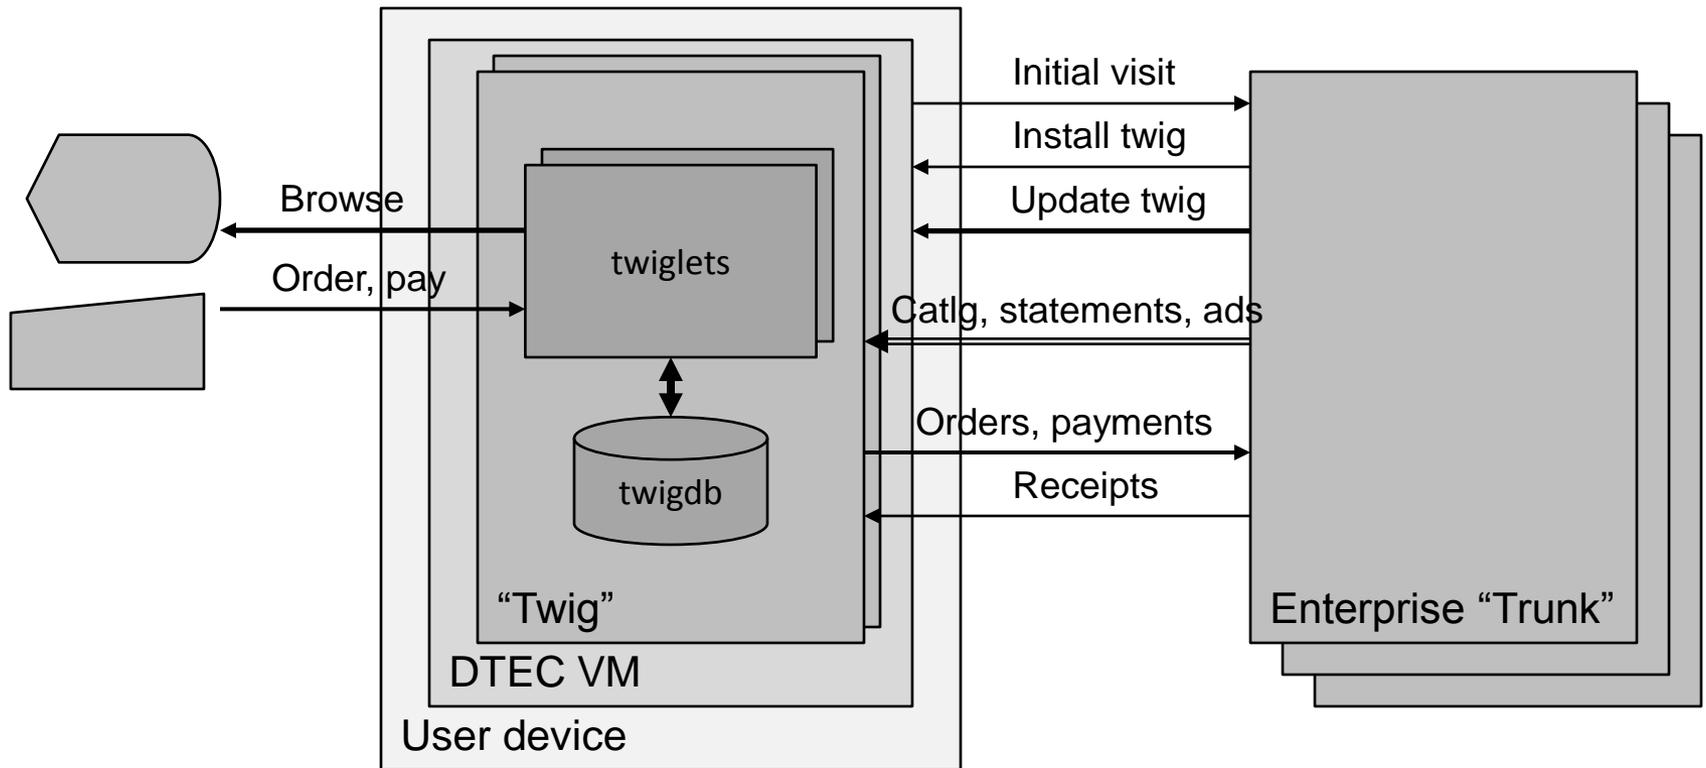
Non-Obvious Potential Applications

- Warning signals, distress signals, requests for disaster relief: you want them to arrive as soon as possible, but where latency is unavoidable, you at least want reliability.
- Most of today's heavily used applications are naturally delay-tolerant.
 - File transfer, including media download: you may want it to happen as soon as possible, but it's not conversational.
 - Social media: Facebook, Twitter, Instagram, Tumblr are multitasking-friendly – you don't have to respond instantly, and you don't expect anyone else to either.

Delay-Tolerant e-Commerce (1 of 2)

- Even much electronic commerce and finance could be accomplished over a high-latency network: ***anything that we used to do by mail could be done on a DTN-based network.***
 - Merchandise used to be purchased by mail order, from catalogues received in the mail.
 - Banking – deposits, payments, statements – used to be done by mail.
- In every case, applications on a high-latency network can provide the same functionality but ***much more quickly and securely, with graphics and automation.***

Delay-Tolerant e-Commerce (2 of 2)

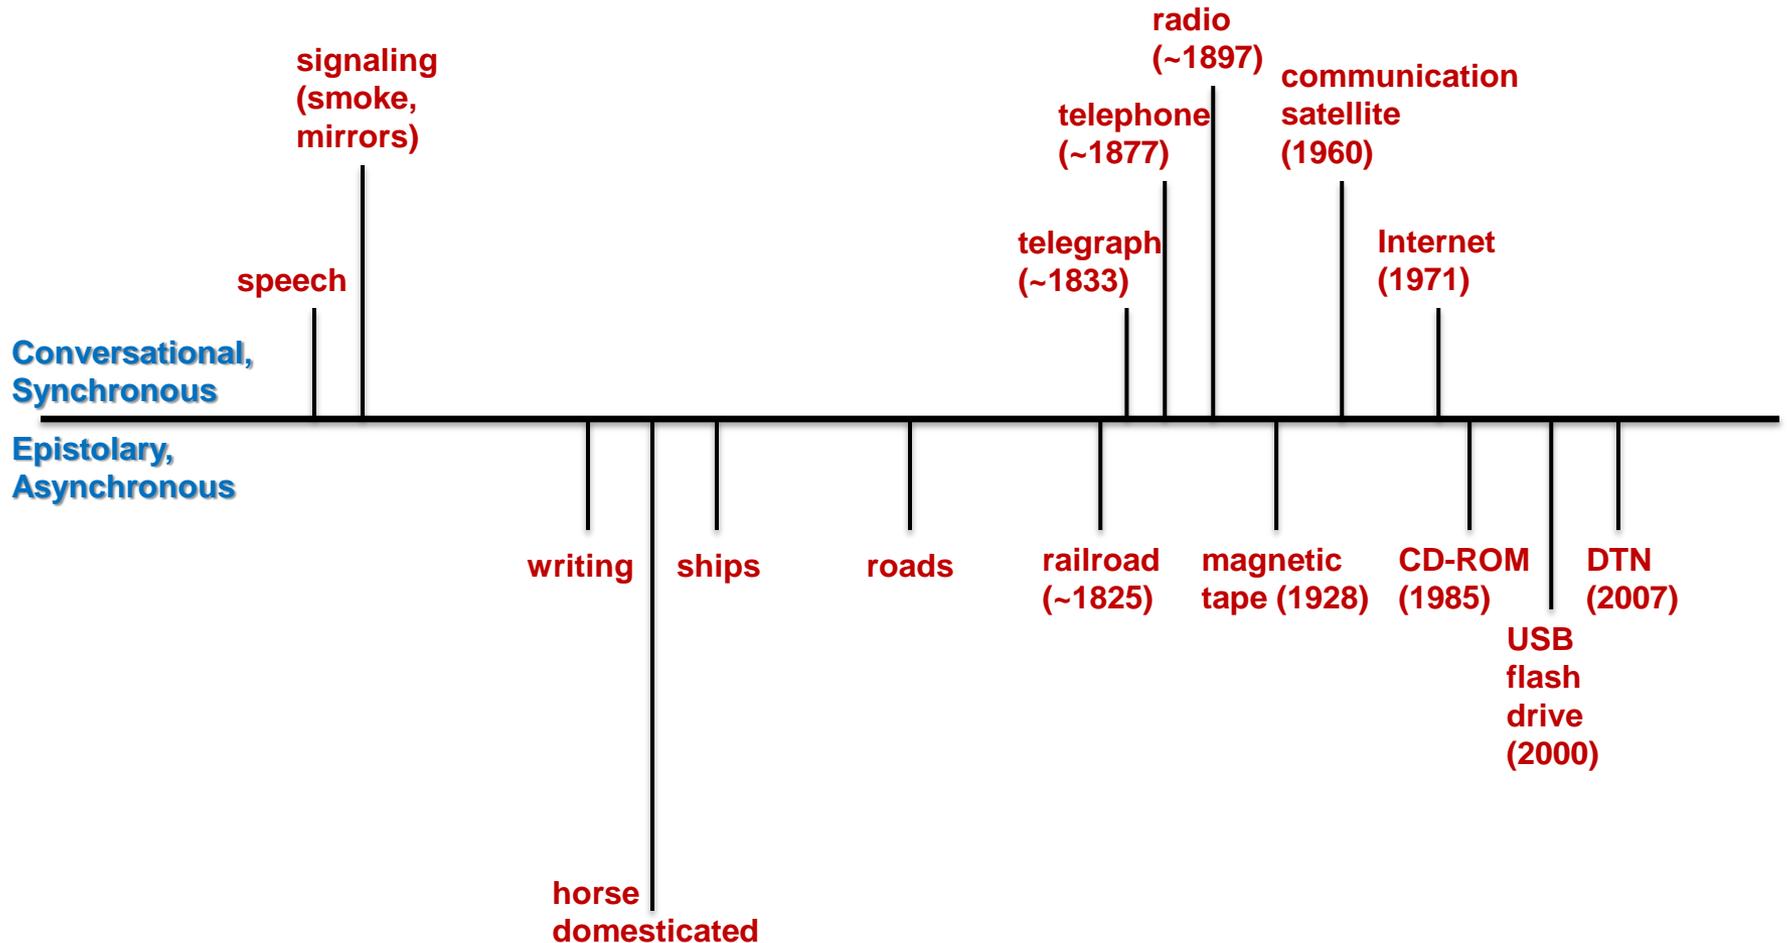


“DTEC VM” hosts a local, private Web server, remotely managed by the merchant, bank, etc. The user sees Web pages, exactly as in today’s e-Commerce, but with much higher speed and reliability.

A Perspective on Using a Network

- Our concept of how to use a network derives from the origin of the Internet: modulation and demodulation of data over telephone lines. We expect a network to operate like a telephone, only more so.
- A high-latency network breaks this expectation: the network no longer operates like a telephone.
- But communication that is non-telephonic is not necessarily inconvenient or inefficient. For many communication applications, telephony is overkill.
- In some cases – e.g., signing up with a health care exchange – non-telephonic electronic access might actually be *better*.

Communications Technology Timeline



Caveats

- Operations: the core DTN protocol implementations are mature, but supporting software is needed.
 - Route computation could be based on Contact Graph Routing.
 - Scalable network management protocols and tools would be needed.
- Security:
 - Basic network security measures (authentication, integrity, confidentiality) are supported by Bundle Security Protocol.
 - A scalable key distribution system would be needed.
- The elephant in the room: where will the funding come from?

Outlook

- DTN on satellites in low-Earth orbit is already a well-proven concept: Change Request 013799, authorizing deployment of two permanent DTN gateway nodes on the International Space Station, was approved on 10 September 2013.
- The Ring Road network concept has some clear advantages:
 - Would surmount geographic obstacles.
 - Not disabled by earthquake, storm, flood, fire.
 - Difficult to disable intentionally.
 - Low barrier to entry.
 - Low operating cost.
- It would be a high-latency network, unsuitable for some applications. But for many applications it could work as well as the Internet. Perhaps even better.

Thanks!

Questions?

