Internetworking with satellite constellations

PhD viva presentation

Centre for Communication Systems Research, University of Surrey

11am, Wednesday, 21 February 2001.

Lloyd Wood

http://www.ee.surrey.ac.uk/Personal/L.Wood/

work done in the Networks Group,
Centre for Communication Systems Research,
University of Surrey.

Overview of this talk and Lloyd’s PhD work:

- Geometry; orbital constraints, mesh network topology, handover and routing give path delays across constellation.

- TCP across LEO/MEO mesh networks – performance influences the approach to routing design.

- Multicast – a simple core-based approach; place the core (and its network state), shape the resulting tree.

- Controlled handover in rosette geometries for classes of service with smaller path delay variation.

- Architecture – support IP QoS and multicast using MPLS for a single control plane.

Without intersatellite links

Satellites *are* just 'last hop'.
Users need to be in same satellite footprint as gateway station for connectivity.
Coverage over remote areas (oceans) where there is no local infrastructure not possible.

Physical delay within space segment highly predictable and small; depends entirely on terminal-satellite-gateway hops during passes; likely to be symmetrical.

Delay in ground segment **unknown**.
Resulting total delay **unknown**.

Onboard processing, switching are *optional*. No onboard routing.

example constellations:
Globalstar, ICO, Skybridge

With intersatellite links (ISLs)

Satellites *aren't* just 'last hop'.
Users don't need to be in same satellite footprint as gateway station for connectivity.
Coverage over remote areas (oceans) where there is no local infrastructure possible.

Physical delay within space segment highly variable but deterministic; depends on where terminals and gateways are, path between them. May not be symmetrical.

Can neglect ground segment.
Can model space segment, find delay.

Onboard processing, switching, routing *are necessary*. It's a network.

example constellations:
Iridium, Teledesic, Spaceway

Concentrated on the satellite constellation with ISLs

- It's a single network. An autonomous system; single point of control and management (not a military viewpoint).

- Constraints of orbital motion and coverage mean that network can be defined and simulated with accuracy.

- Speed of light is a constraint. ISL propagation delay is largest factor, subsuming all others.

- What does traffic across the constellation experience? How do applications see it?

Focus on *delay* perspective. Latency indicates performance. (Congestion is multivariable; too many starting assumptions required to simulate congestion accurately.)

Constellation geometry - star *vs* rosette (LEO/MEO ISL designs)

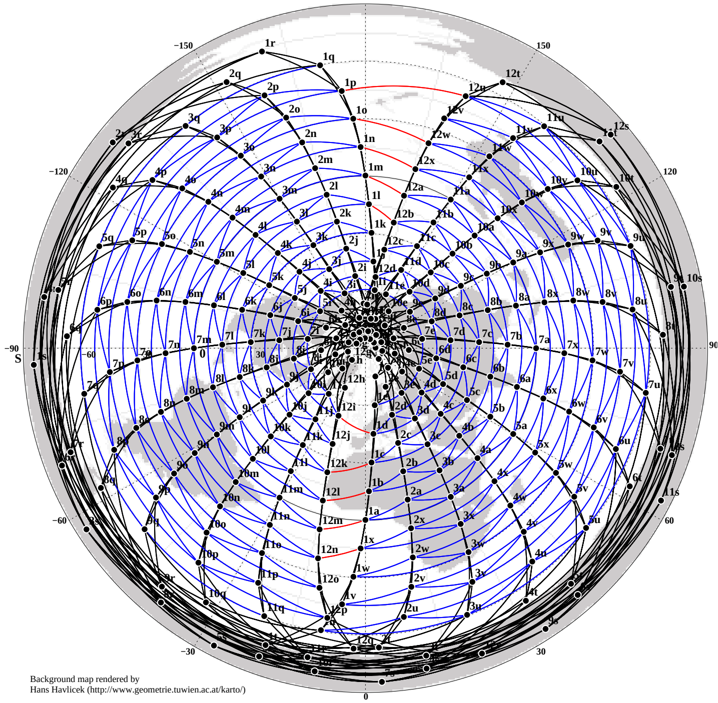

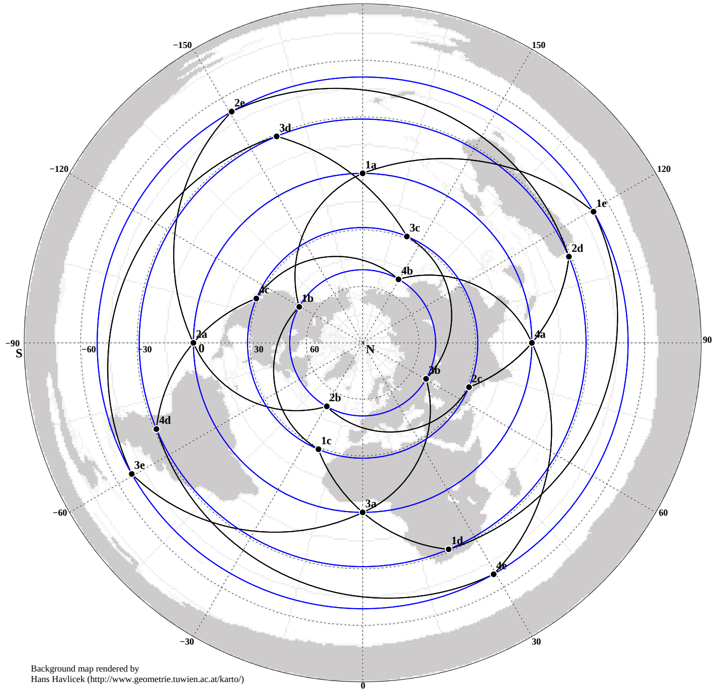

Teledesic **(Boeing design - 288 active satellites)**

Hughes *Spaceway NGSO* (20 active satellites)

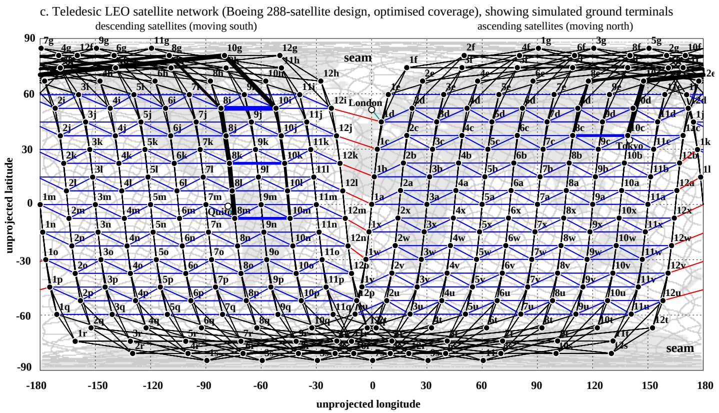

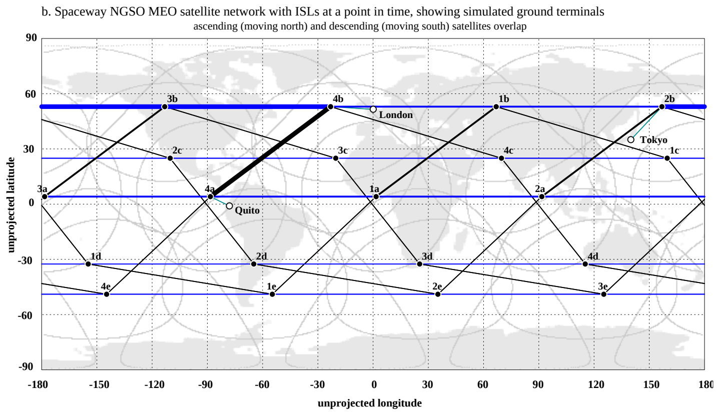

Intersatellite links give us mesh networks...

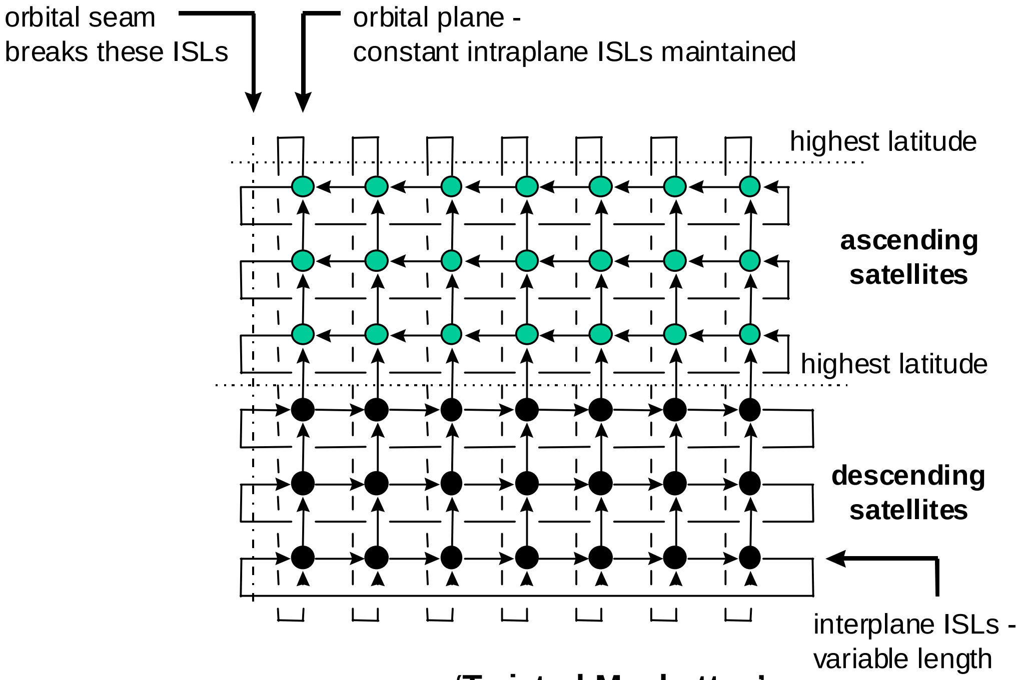

'Twisted Manhattan' satellite network variant

path delay (milliseconds) over a day

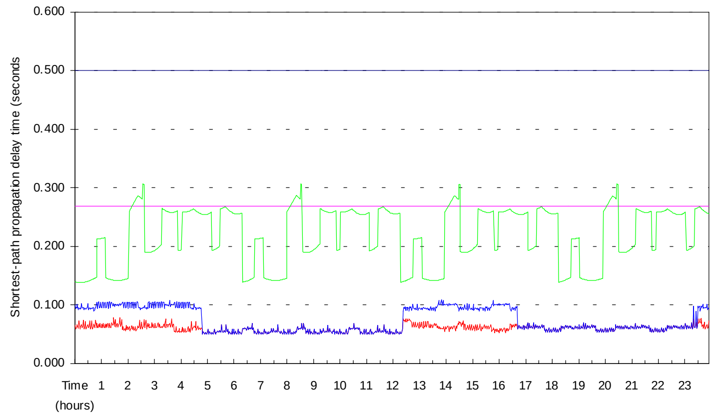

and as number of hops through the mesh:

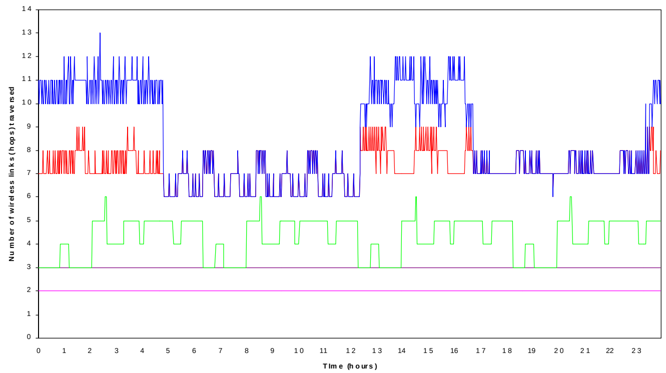

Path delay from Quito to London (by LEO/MEO/GEO)

...and orbital motion gives us handover, path delay variation.

So, taking path delay further…

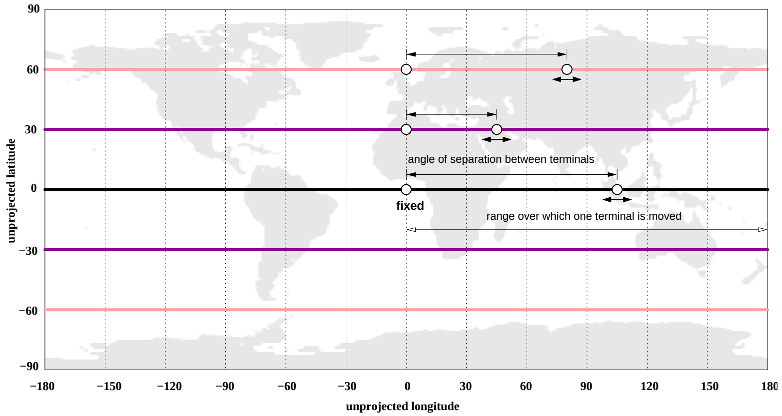

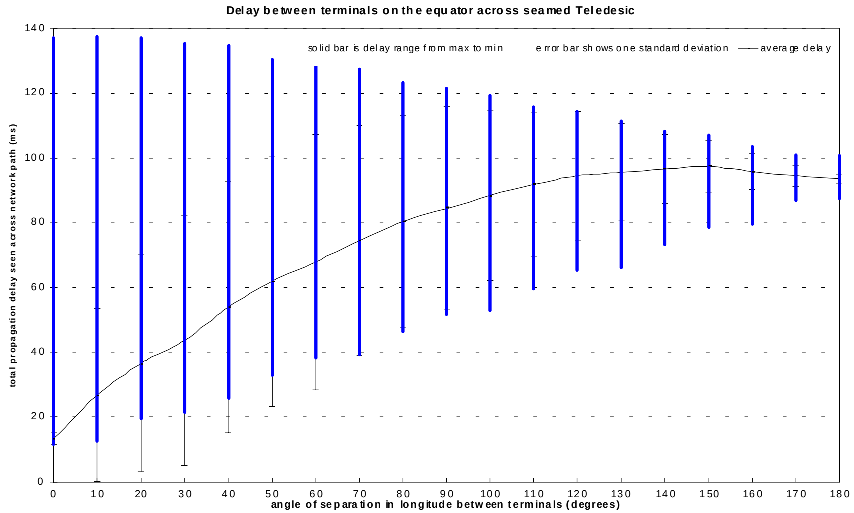

Teledesic *with* and *without* cross-seam links; every line shows a 24-hour simulation at a given latitude separation between terminals.

Smaller delays and delay variation *with* cross-seam links.

Note difference at 0º separation.

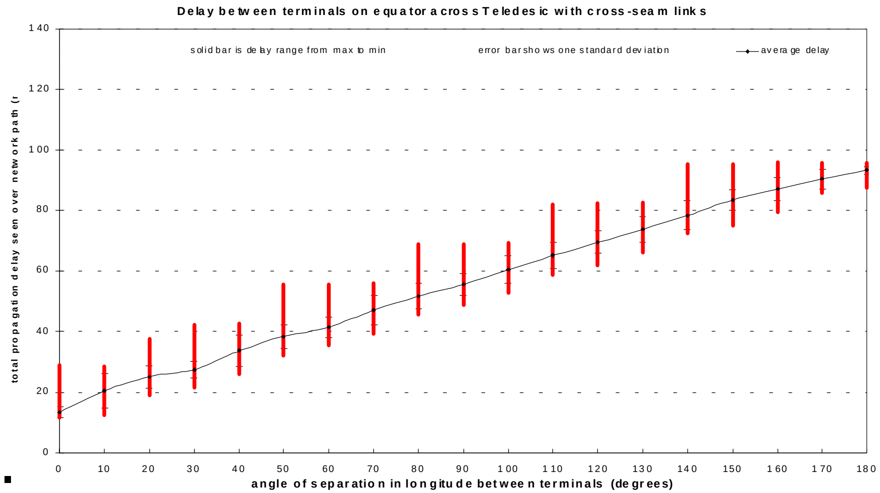

Handover and packets in flight

Star seam is a worst case, but there are transient spikes in path delay for packets in flight (subject to what your implementation does; could just drop packets with very hard handover…) *every* time your terminal undergoes a handover.

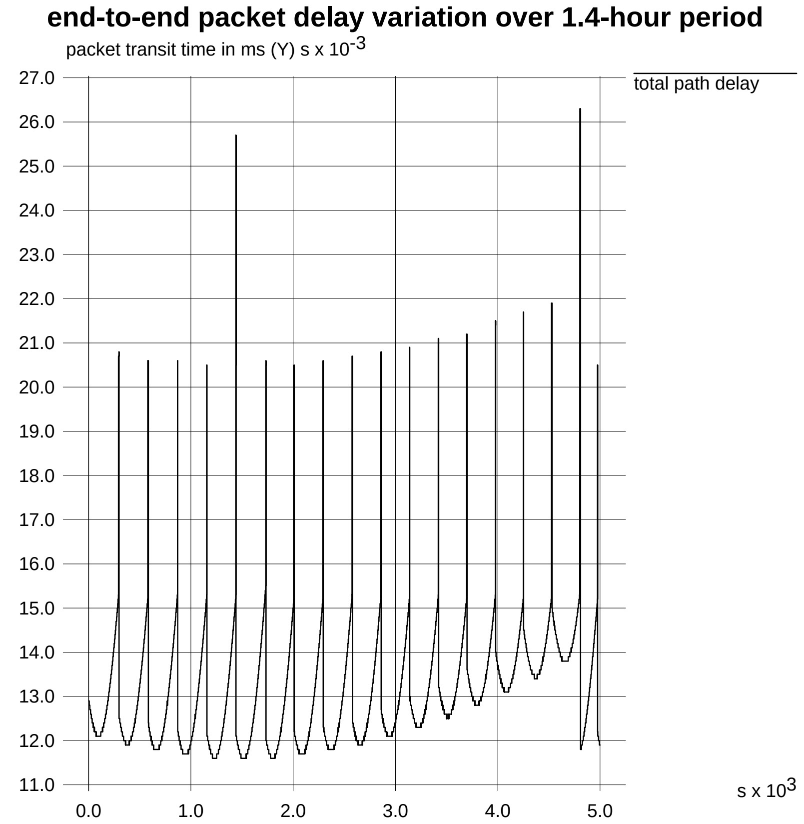

This has implications for moving network state between satellites. If you don't want to disrupt high-rate traffic, you want small spikes of no more than 1 ISL hop.

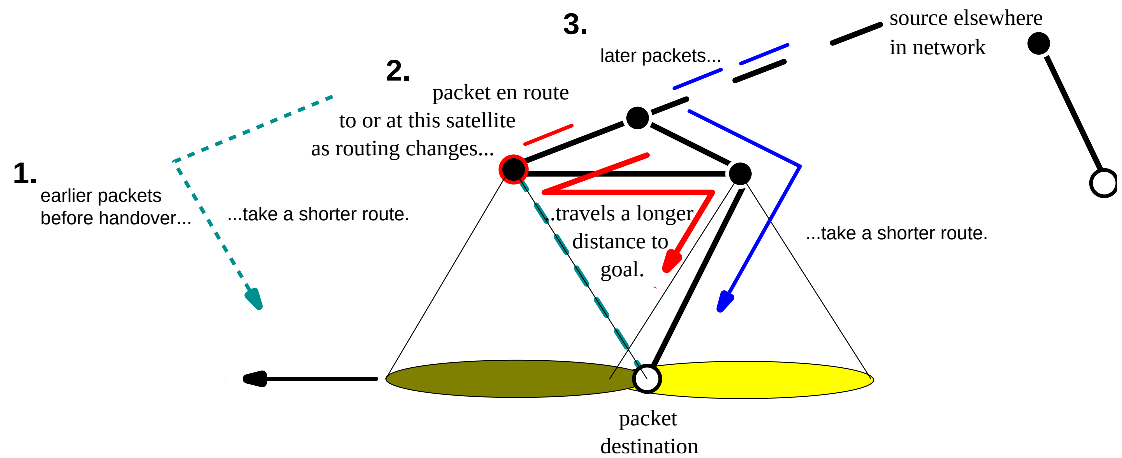

TCP over the constellation mesh

Examined interaction of fast recovery and delayed acks with multipath routing to simulate load balancing.

TCP’s fixed three-dupack threshold presumes congestion, can’t cope with large amounts of reordering. Raise that threshold and performance improves.

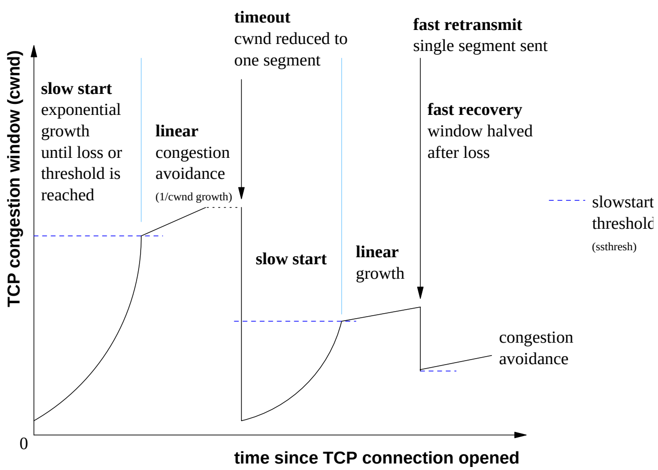

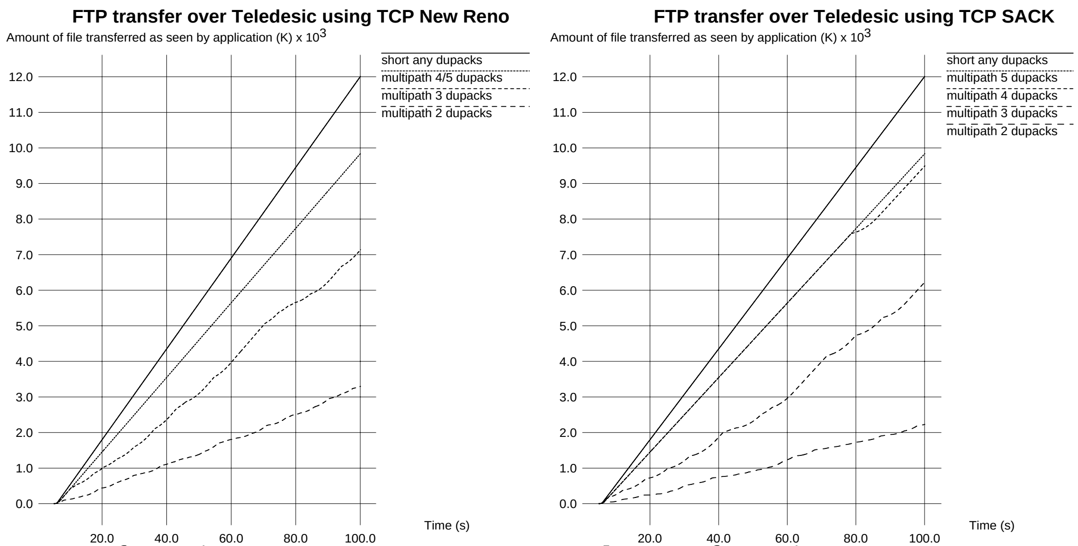

a. Transfers using New Reno TCP **b. Transfers using SACK TCP**

Progress of FTP transfer between terminals at Quito and Tokyo using *Teledesic*.

effect of dupack threshold on throughput over multiple LEO paths

Delayed acknowledgements

Increased ack delay decreases throughput; window growth in recovery is slowed… for both SACK and New Reno.

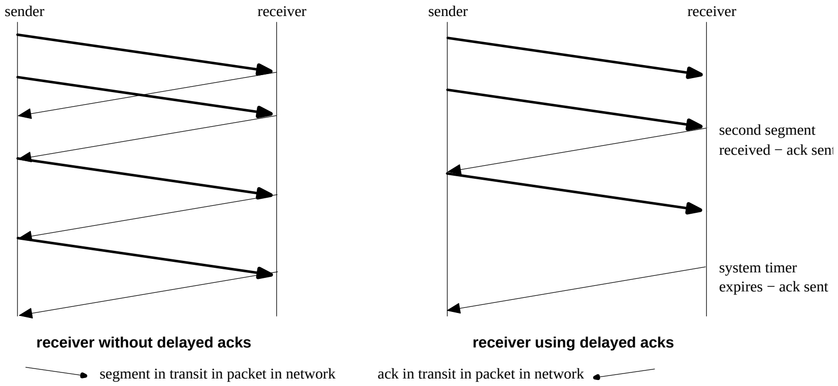

Similar results were surprising; experience tells us SACK performs better in most scenarios. What happened?

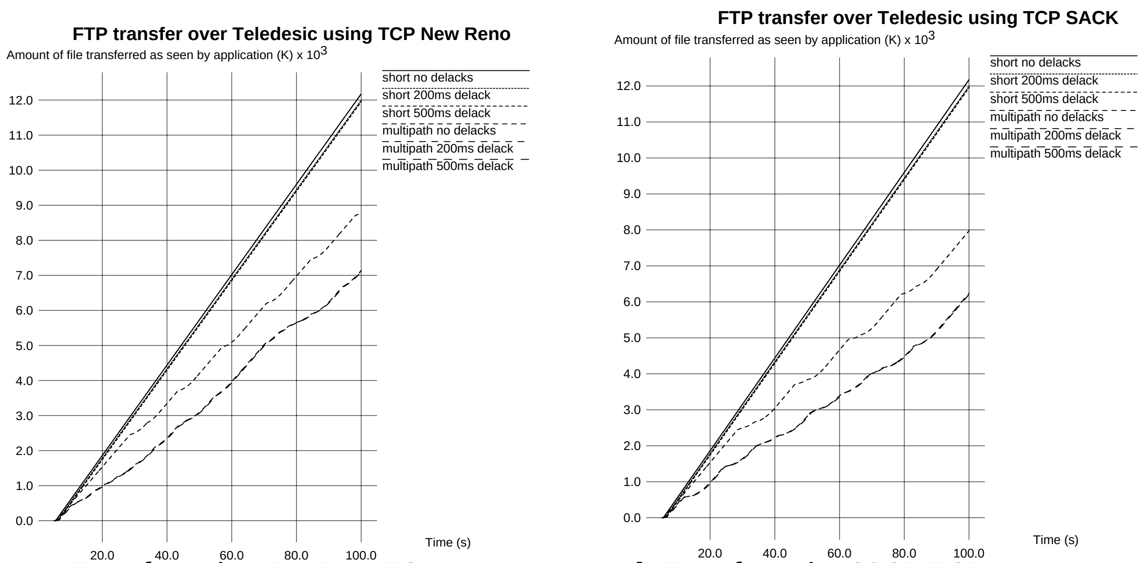

a. Transfers using New Reno TCP **b. Transfers using SACK TCP**

Progress of FTP transfer between terminals at Quito and Tokyo using *Teledesic*

delayed acks degrading file transfer over multiple LEO paths

This is due to a delack implementation decision

Should only delay in-order acks at the right edge of the window when nothing is outstanding. SACK can then grow its congestion window faster in recovery than New Reno. This *should* increases flow performance - needed for multiple paths.

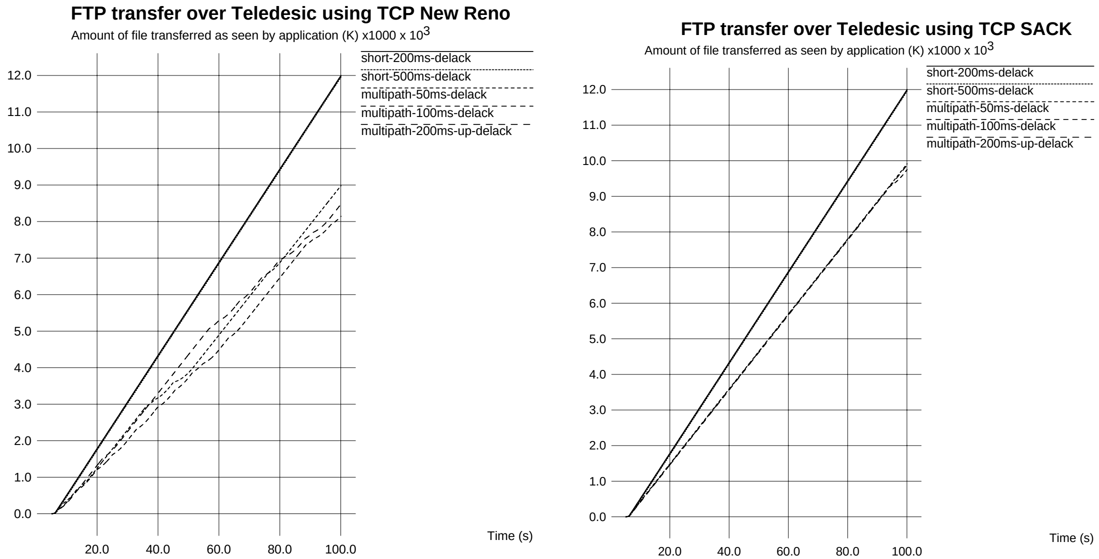

a. Transfers using New Reno TCP **b. Transfers using SACK TCP**

Progress of FTP transfer between terminals at Quito and Tokyo using *Teledesic*

delacks only take place on packets received at the right edge of the window (RFC2581 SHOULD)

delacks degrading rate of file transfer over LEO obeying RFC2581

TCP performance at its best with ordered flows; this encourages flow-based approach to routing/traffic engineering.

Multicast in the constellation mesh

Took a core-based approach; assume that the constellation is likely to be dealing with small number of users in a group.

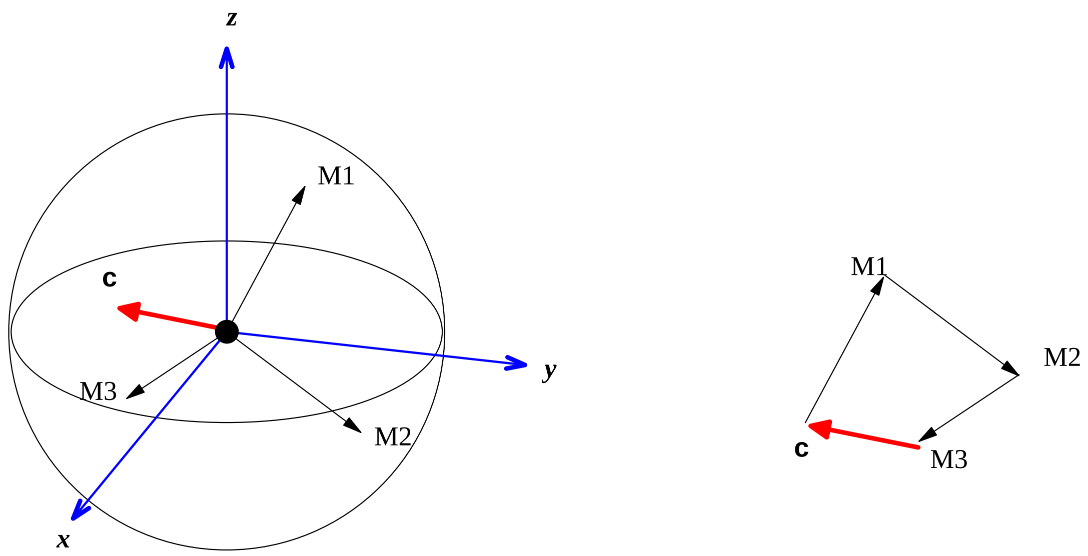

Satellites move, but terminals are fixed; place a fixed core using linear algebra (vector summation) and have its state inherited by the nearest satellite and passed on over time.

Hierarchy of terminals - satellites - core maps nicely to the core-based tree (CBT) hierarchy…

Core is a satellite; duplicate in mesh, conserve air interface.

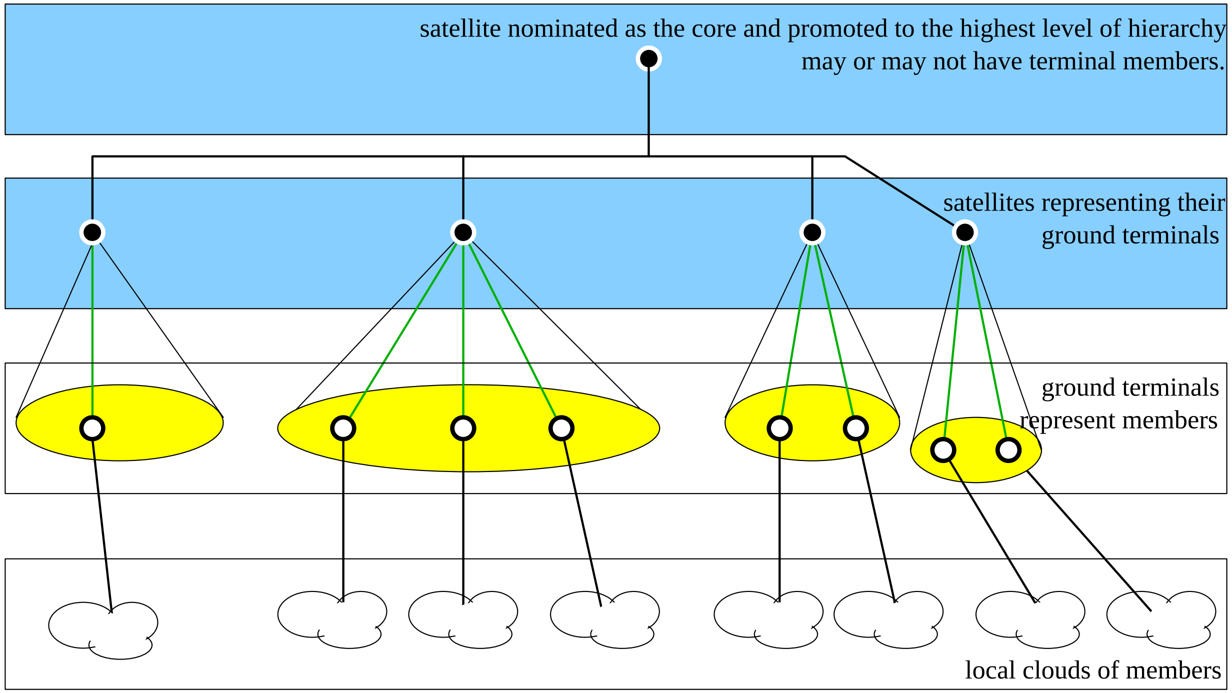

…but how does multicast perform?

CBT in ISL mesh

Simulated movement of constellation over the course of 24 hours; measured path delays between all terminals via the core, plotted means. Delay and variation in delay are worse than unicast shortest path, but not far worse…

Benefit of capacity saving in links due to duplicating packets at satellites.

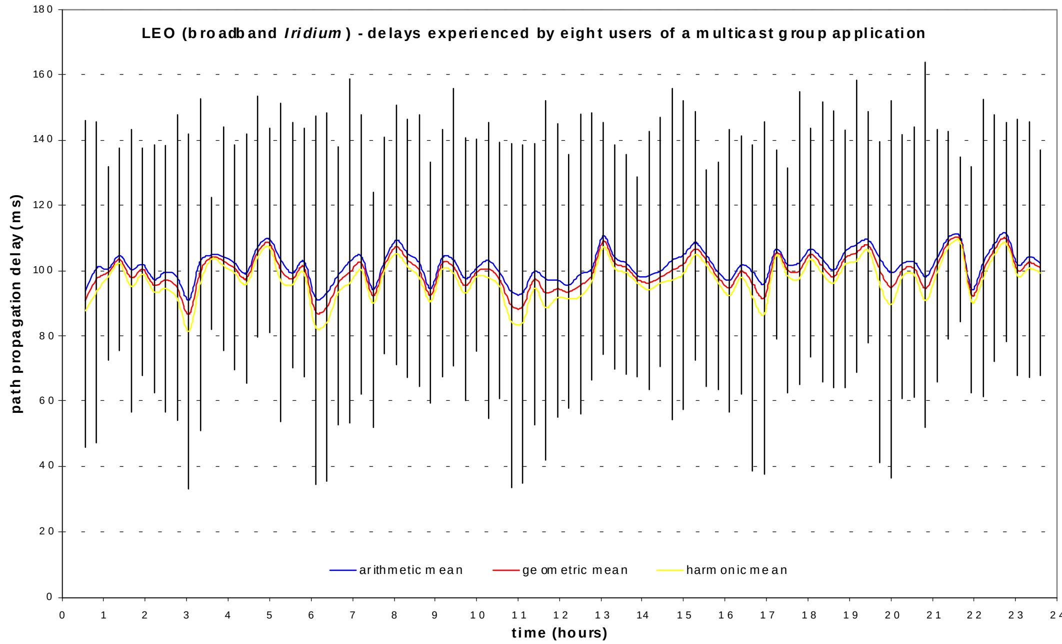

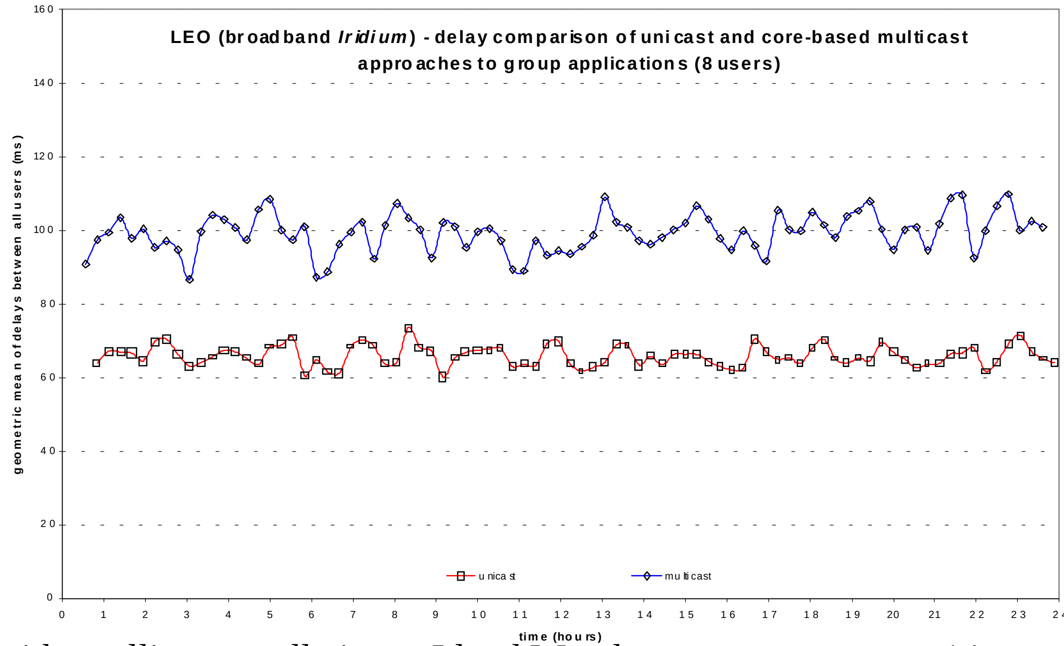

Capacity saving metrics

Chang-Sirbu scaling law says that ratio of multicast/unicast hops used is proportional to (group size)$^{0.8}$

And it holds for the constellation network, too.

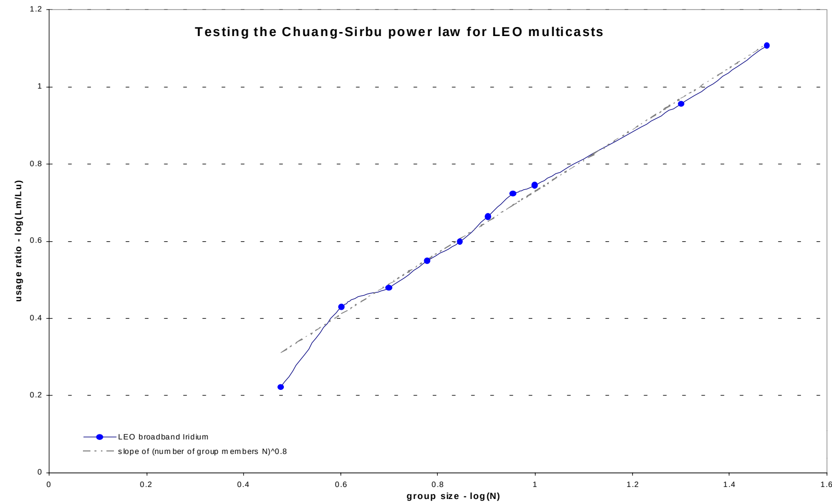

Multicast/unicast hop savings at LEO

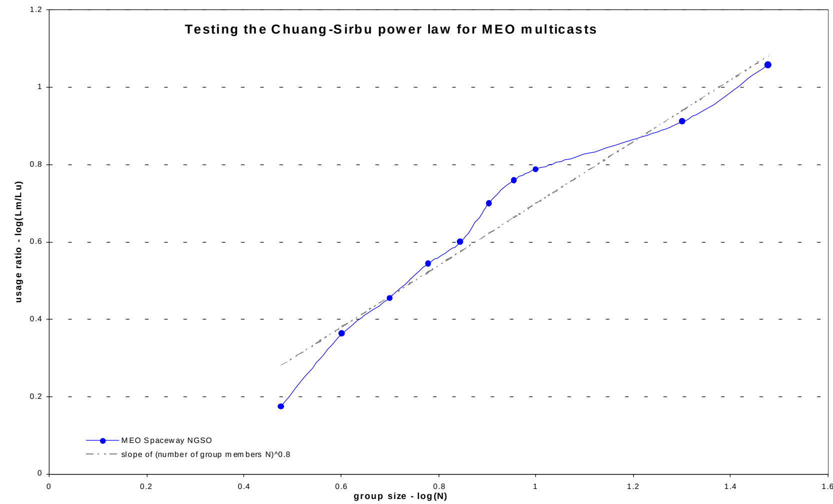

Multicast/unicast hop savings at MEO

Also invented a variation on the vector algorithm for star constellations without cross-seam links (not that you'd want to do that - core state moves around).

Looked at multicast across rosette (*Spaceway NGSO*):

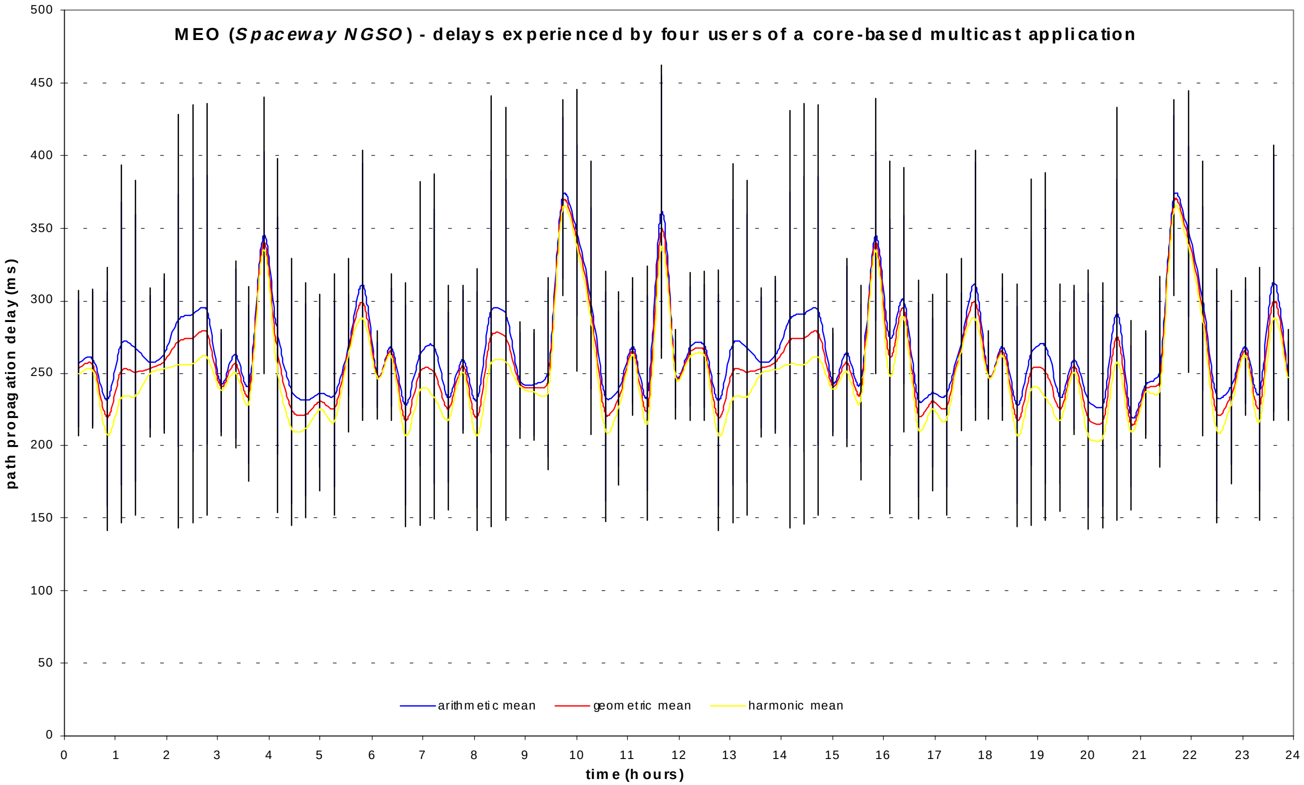

Delay variation is all over the place, due to handover…

Handover in rosettes

ascending and descending satellites mean two parts of the ISL mesh can be seen overhead; by multihoming ground terminals we increase reliability, and can also introduce two sets of path delays…

Current proposals (e.g. *Globalstar*) just use this for diversity in the air interface.

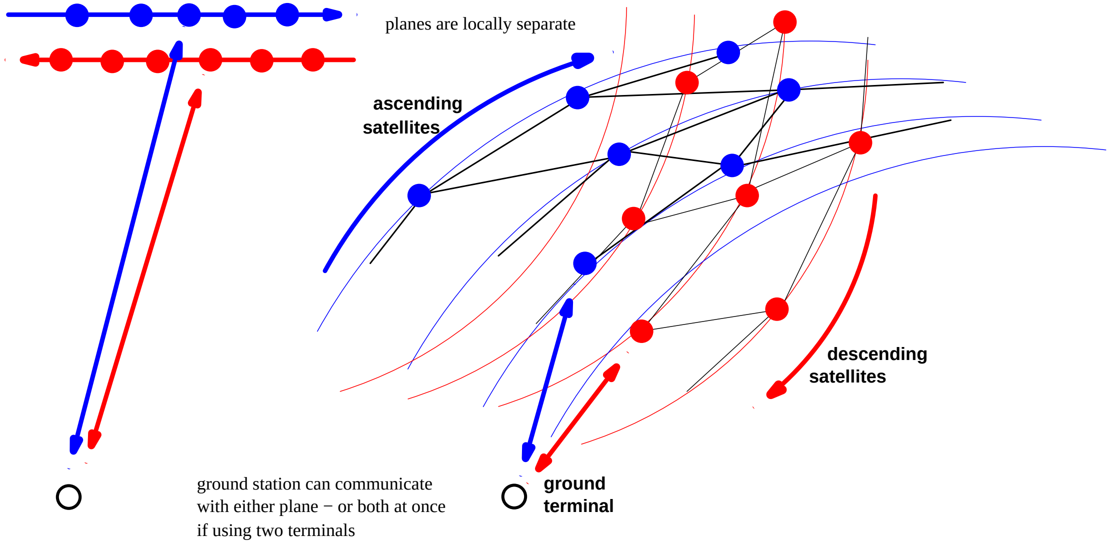

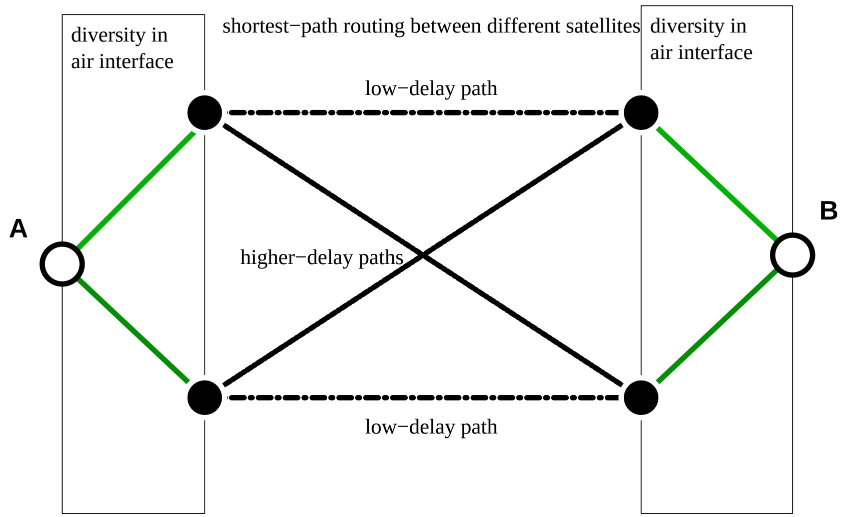

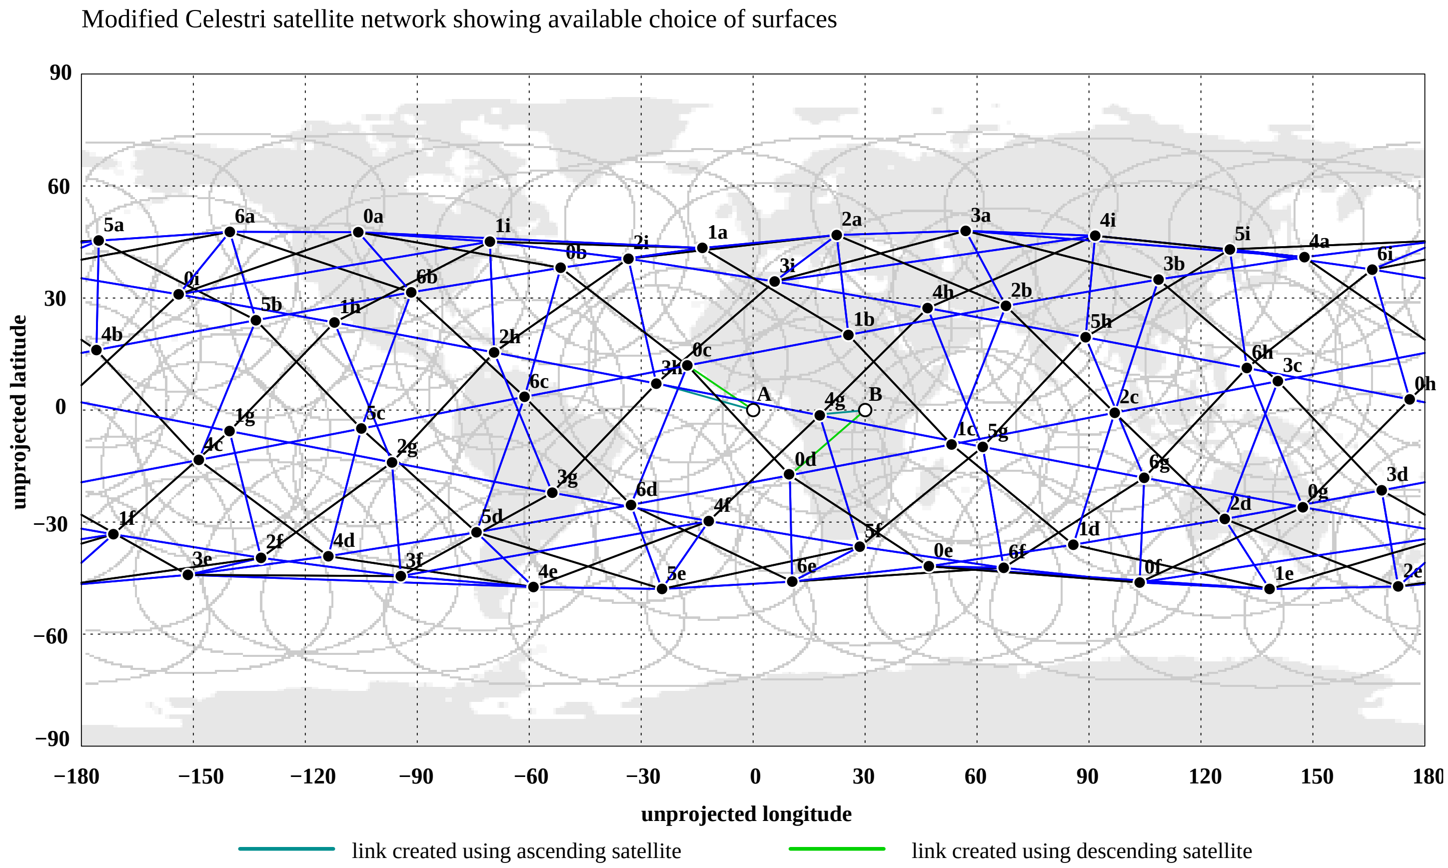

Examined Motorola’s LEO *Celestri* proposal.

But we have to modify the existing geometry
of this proposal slightly to get the delay separation we want.

Testing controlled handover for multihomed terminals

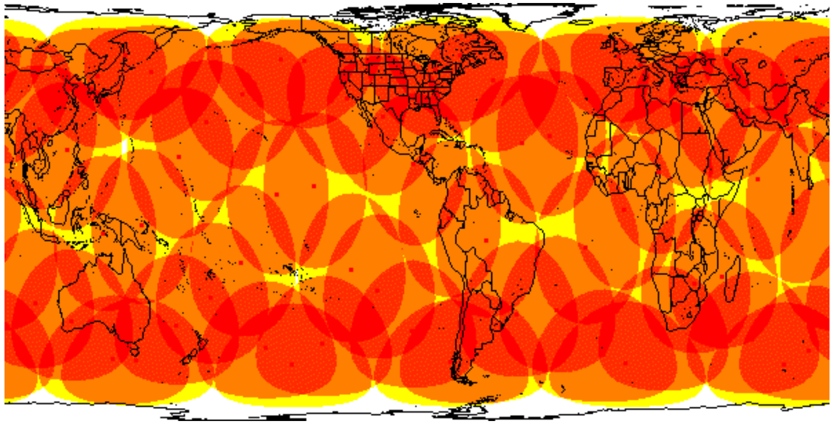

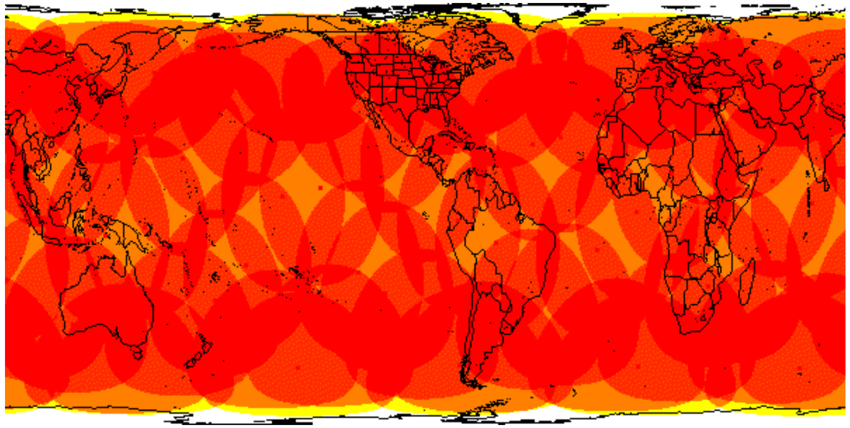

LEO *Celestri* coverage (single in parts – multihoming limited)

Modifying for double surface coverage (lower mask elevation angle by 6º)

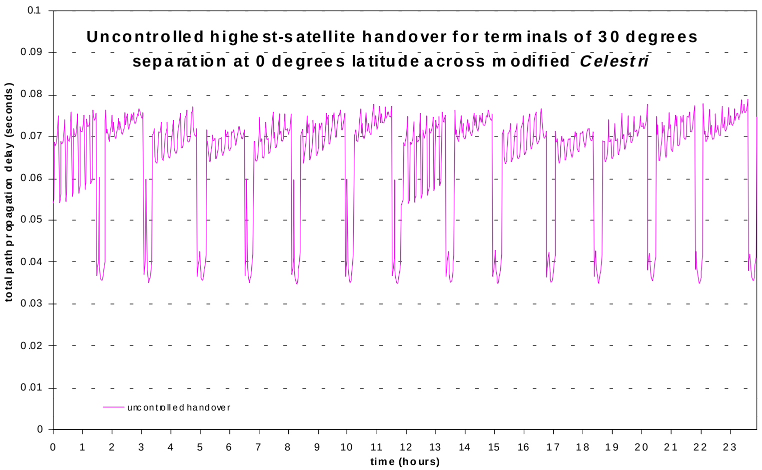

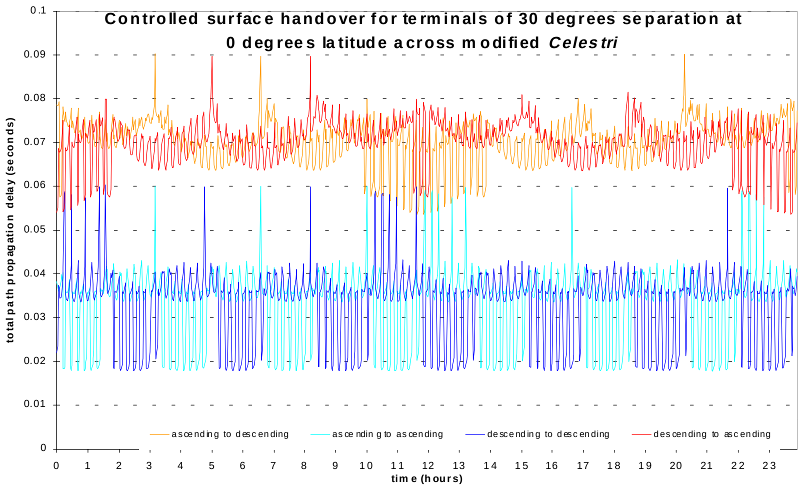

Single-homed terminals pick the highest visible satellite and flip between surfaces

Mujlti-homed terminals stick to a surface and give two sets of delay classes

Reusing our tried and trusted delay evaluation method

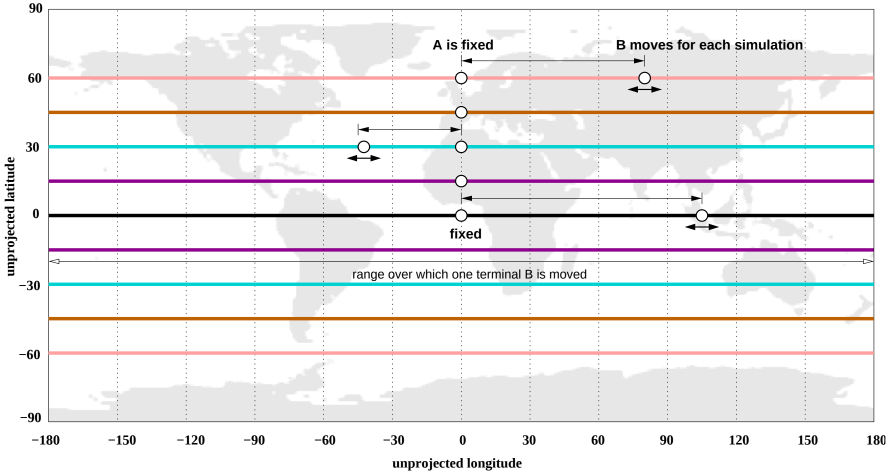

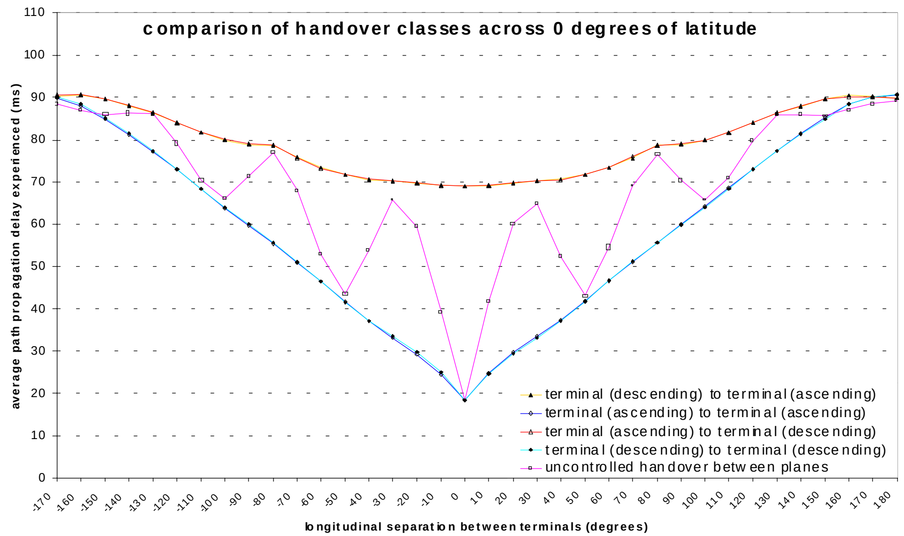

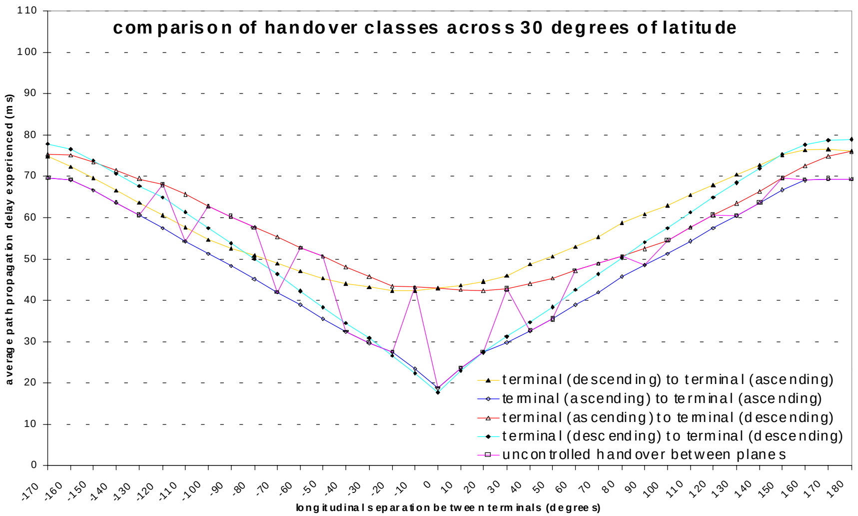

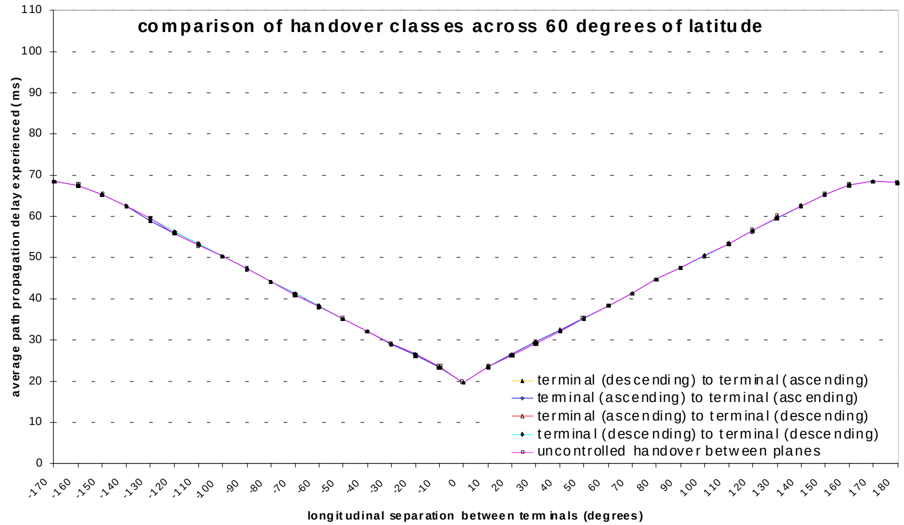

Uncontrolled handover falls between two limits. At limits of coverage (60º latitude) effect vanishes.

At MEO, with *Spaceway NGSO*...

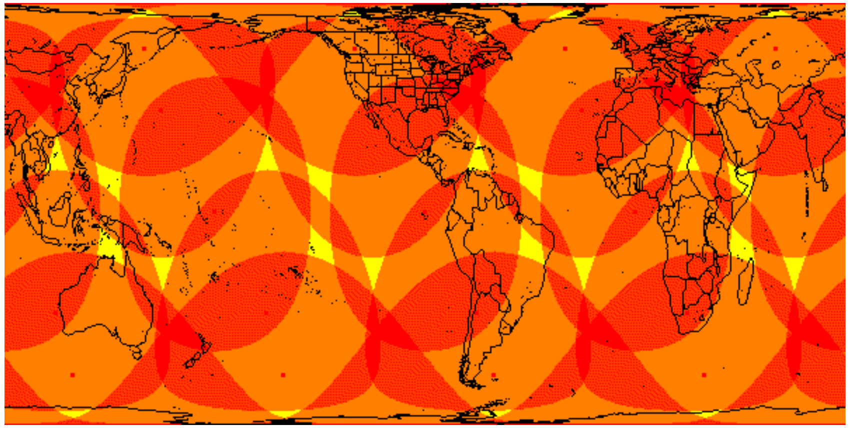

Some gaps in coverage are visible

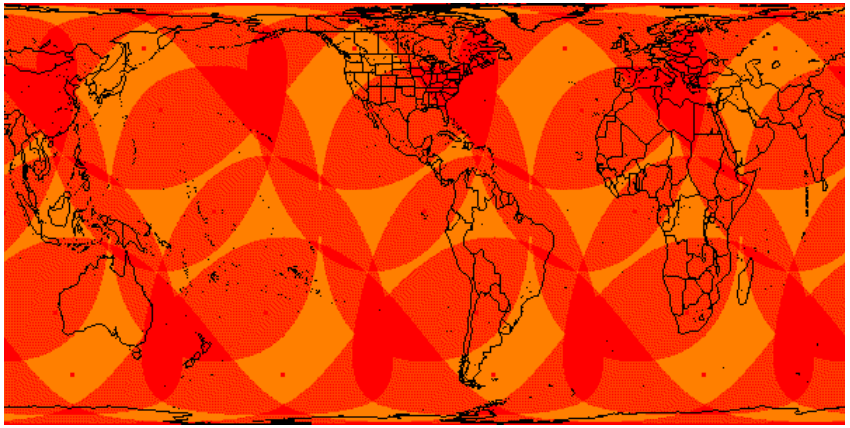

Lower elevation angle gives dual coverage (but not double surface coverage)

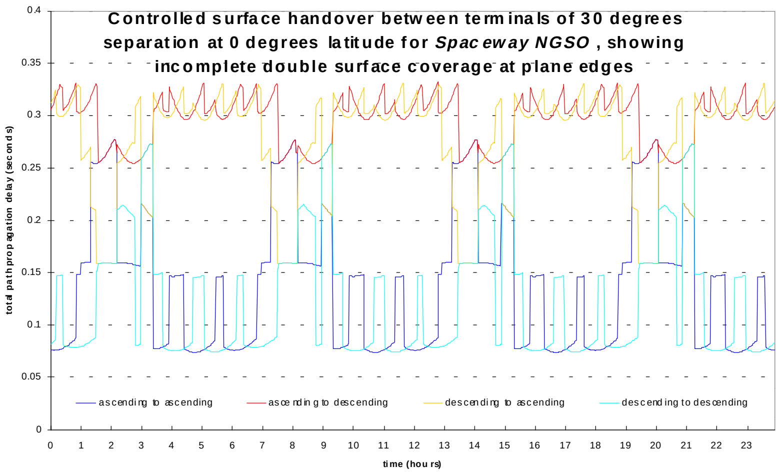

Orbital planes are too far apart; multihomed terminals still forced to hand over to other surface

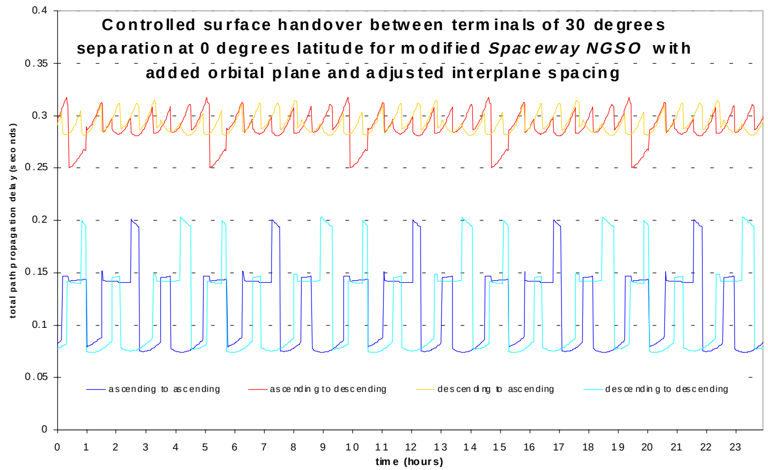

Insert an extra plane – classes of delay and service are introduced

Architectural considerations

Support for IP multicast and IP QoS in the constellation are desirable. For both of those, we need support for IP routing.

Much work has been done on wireless ATM on the air and ISL interfaces – typically two cells packed in a MAC-layer frame with checksum.

MPLS is the logical way to combine IP routing with an ATM infrastructure.

Constellation is an autonomous system; this influences routing. Combine BGP, IGP, MPLS, tunnelling and even twice-NAT so that satellites can inherit *minimal* ‘virtual node’ routing state as they move, while terminals have fixed IP addresses.

Achievements #1

- Showed the desirability, from viewpoint of delay and delay variation, of cross-seam links in star constellation networks.

- Demonstrated interesting transient effects on traffic in ISLs due to handover, with implications for moving state.

- Showed that TCP multipath performance is sensitive to delack implementation, and that a flow/traffic engineering approach to routing in the constellation mesh benefits TCP.

- Proposed a simple, robust approach to multicast in the constellation by computing core position; resulting capacity savings obey the Chuang-Sirbu scaling law.

Achievements #2

- Proposed an architecture for the constellation using MPLS.
- Introduced a way of managing delay with handover in the rosette constellation with ISLs, leading to classes of service. No previous proposals (*M-Star*, *Celestri*, *Spaceway NGSO*) permit this, but geometries can be modified to do so.
- Showed importance of handover in the constellation for its effect on end-to-end traffic delays.
- Publications – journal and conference papers.
- Some software, too.

Related publications

Two peer-reviewed journal papers

Wood, Pavlou, Evans, ‘Effects on TCP of routing strategies in satellite constellations’, *IEEE Communications Magazine*, March 2001.

Wood, Clerget, Andrikopoulos, Pavlou, Dabbous, ‘IP routing issues in satellite constellation networks’, *International Journal of Satellite Communications*, January/February 2001.

Two related conference papers

Wood, Pavlou, Evans, ‘Managing diversity with handover to provide classes of service in satellite constellation networks’, 19th AIAA ICSSC, Toulouse, April 2001.

Wood, Cruickshank, Sun, ‘Supporting group applications via satellite constellations with multicast’, IEE ICT ’98, Edinburgh.

Also ICC ’00 (Andrikopoulos, diffserv), an internet-draft on ARQ, COST contributions...

Related software contributions

Satellite Visualisation software *SaVi*
Wrote most of the scripts in existence, simulating commercial and idealised constellations based on public descriptions. Several packaged with *SaVi* 1.0 release.

Network simulator *ns*
Wrote enhancements to Tom Henderson's code, including graphical constellation network plot perl scripts; sundry speedups and bugfixes. All for *ns* 2.1b7.

Webpages
Widely-referenced teaching resource; various awards.
http://www.ee.surrey.ac.uk/Personal/L.Wood/