Networking over Next-Generation Satellite Systems

by

Thomas Ross Henderson

B.S. (Stanford University) 1990
M.S. (Stanford University) 1991

A dissertation submitted in partial satisfaction of the
requirements for the degree of
Doctor of Philosophy

in

Engineering– Electrical Engineering
and Computer Sciences

in the

GRADUATE DIVISION
of the
UNIVERSITY of CALIFORNIA at BERKELEY

Committee in charge:

Professor Randy H. Katz, Chair
Professor Dorit S. Hochbaum
Professor Steven R. McCanne

Fall 1999

The dissertation of *Thomas Ross Henderson* is approved:

Chair Date

Date

Date

University of California at Berkeley

Fall 1999

Networking over Next-Generation Satellite Systems

Abstract

Networking over Next-Generation Satellite Systems

by

Thomas Ross Henderson
Doctor of Philosophy in Engineering– Electrical Engineering and Computer Sciences

University of California at Berkeley

Professor Randy H. Katz, Chair

Thanks to both the rapid deployment of the Internet and advances in satellite technology, the market for broadband satellite services is poised for substantial growth in the coming decade. Current communications satellite systems have generally been designed to provide either voice or data transaction (low data rate) services through small terminals, or trunking (high data rate, or broadband) services through large terminals. However, technological advances are enabling new systems that combine broadband data rates with small terminals, thereby providing more affordable "last-mile" network access to home and small business users worldwide. In particular, two types of broadband satellite systems are under development: high-power satellites deployed at traditional geostationary (GEO) orbits, and large constellations of satellites deployed at much lower (LEO) orbits.

In this thesis, we explore research problems that have arisen from this shift in satellite network architectures. When using GEO satellites to provide Internet access service, the performance of the Internet's Transmission Control Protocol (TCP) is degraded by the high latency and high degree of bandwidth asymmetry present in such systems. We therefore undertook a comprehensive study of TCP performance in the context of broadband satellite systems used for network access. We first studied whether TCP's congestion avoidance algorithm can be adjusted to provide better fairness when satellite connections are forced to share bottleneck links with other (shorter delay) connections. Our data suggests that adjustments of the policy used in that algorithm may yield substantial fairness benefits without compromizing utilization. We next demonstrated how minor variations in TCP implementations can have drastic performance implications when used over satellite links (such as a reduction in file transfer throughput by over half), and from our observations constructed a satellite-optimized TCP implementation using standardized options. We explored the performance of TCP for short data transfers such as Web traffic, and found that two experimental options relating to how TCP starts a connection, when used together, could reduce the user-perceived latency by a factor of two to three. However, because not all of these options are likely to be deployed on a wide scale, and because even the best satellite-optimized TCP implementation is vulnerable to the fairness problems identified above, we explored the performance benefits of splitting a TCP connection at a protocol gateway within the satellite network, and found that such an approach can allow the performance of the satellite connection to approach that of a non-satellite connection. Carrying this work one step further, we construct a satellite-optimized transport protocol that can be used in such a split-connection environment, and demonstrate how it outperforms

TCP in a satellite environment characterized by large amounts of bandwidth asymmetry. In particular, our protocol, which we have dubbed the “Satellite Transport Protocol,” uses up to an order of magnitude less traffic on the bandwidth-constrained reverse channel than is needed by TCP.

In contrast to research on GEO systems, research on LEO systems is still in its infancy. While LEO systems are being designed specifically to avoid the high latencies found in GEO systems, their design is challenging from a packet routing perspective due to the highly time-varying nature of the LEO network topology. Moreover, even the most basic system properties of such constellations is not well documented in the literature. We constructed a detailed packet-level LEO network simulator and identified some fundamental delay performance results of commercially proposed constellations. We then explored whether or not geographic-based network addresses can be constructively used in the design of distributed or centralized packet routing systems. We found that the construction of a distributed packet routing algorithm based on geographic-based packet forwarding is fundamentally challenging due to subtle topological properties of LEO networks. However, we demonstrated that geographic-based addresses are useful in centralized routing systems, enabling significant reductions in routing traffic and on-board routing tables. Specifically, we constructed routing strategies based on geographic-based addresses that potentially reduce by an order of magnitude the amount of routing traffic exchanged between satellite nodes and a centralized routing center on the Earth’s surface, and the number of forwarding table entries required for non-local destinations.

Broadband satellite networks are likely to become an important niche of the future Internet because of their unique ability to provide network access from almost any point on the globe, particularly those underserved by terrestrial infrastructure. However, because the design of Internet protocols is driven by the performance of the wired Internet, satellite network engineers must be vigilant in assisting in the design and deployment of satellite-friendly protocols and in considering how satellite networks interwork with the wired Internet. Broadband LEO networks are likely to be deployed later than their GEO counterparts, and the design of these networks is still in its infancy, particularly since such networks are significantly more ambitious technically than anything that has been previously attempted. It is our hope that the findings presented in this thesis may contribute to a better understanding of how to design protocols for these next-generation systems while stimulating further work in this area.

Professor Randy H. Katz
Dissertation Committee Chair

Contents

List of Figures

List of Tables

Acknowledgements

Pursuing a doctoral degree is a long journey that one cannot make alone, and in this space I'd like to thank the many people who have contributed to my particular journey.

First, I thank my parents for all of the love, support, and encouragement they have provided me throughout my life. I am forever grateful for the sacrifices they made to give me the wonderful opportunites that I have had.

Randy Katz has been a fantastic research advisor on many levels. I have been continually impressed with his ability and dedication in leading a very large group of students towards their degrees, while also teaching and serving as department chair. Randy gave me a lot of freedom to pursue my research, while also providing wise advice and helpful discussions along the way. I feel very fortunate to have had Randy as an advisor and role model during my time at Berkeley.

I am grateful to the many members of the BARWAN research group for their help, friendship, and support. Giao Nguyen helped me a great deal in programming related to the *ns* simulator. Venkat Padmanabhan contributed device drivers used in many of my network experiments. Keith Sklower helped out tremendously by porting satellite card drivers from the Windows platform to Unix. Hari Balakrishnan wrote the initial TCP SACK implementation that I used in experiments. I also especially thank Hari, Venkat, Giao, Elan Amir, and Tina Wong for several technical discussions and reviews of draft versions of this research. For everyone else associated with BARWAN, I thank you for your contributions to my work and to a fun and challenging research project.

I am grateful to the other two professors on my committee, Steve McCanne and Dorit Hochbaum, for their interest and concern along the way and for their help in writing this dissertation. Both provided me with technical assistance, and the classes that they taught on Computer Networks and Network Flows were especially helpful to my work. I would also like to thank Jean Walrand for serving on my quals committee, and Joseph Kahn and Avideh Zakhor for sponsoring my research in my first year at Berkeley.

I thank Sally Floyd for her help with my initial TCP research by providing unpublished research results as well as several discussions on different aspects of TCP. The TCP fairness work reported herein derives from a class project that I undertook with Emile Sahouria, who in particular wrote several components of the background HTTP traffic generators used herein. Bruce Mah and Steve Gribble provided HTTP traffic traces that drove that traffic generator.

The *ns* simulator has been invaluable for my research, and I would like to thank everyone in the VINT project who has made *ns* such a powerful tool for networking research.

I thank Son Dao at HRL Laboratories for supporting my work both through sponsoring BARWAN and employing me as an intern over two summers. I enjoyed working with Son and the other members of HRL, including Yongguang Zhang, Bo Ryu, Dennis Connors, and Dante DeLucia. Hughes DirecPC also graciously agreed to host a Berkeley computer at their network uplink center that was used for some of my satellite network experiments.

I would also like to thank the many people who have provided me with unpublished technical data or assistance, including Don Hoffman and Marie-Jose Montpetit at Teledesic, and Lloyd Wood.

Finally, I thank my wife Sue, whose love, support, and dedication throughout this whole journey has helped get me through it all. From putting up with long hours to helping me celebrate

little victories along the way, I couldn't have asked for better companionship.

Tom Henderson
December, 1999