GENERATION AND VALIDATION OF EMPIRICALLY-DERIVED TCP APPLICATION WORKLOADS

Félix Hernández-Campos

A dissertation submitted to the faculty of the University of North Carolina at Chapel Hill in partial fulfillment of the requirements for the degree of Doctor of Philosophy in the Department of Computer Science.

Chapel Hill
2006

Approved by:
Advisor: Kevin Jeffay
Reader: F. Donelson Smith
Reader: Ketan Mayer-Patel
Reader: J. Steve Marron
Reader: Andrew Nobel
Reader: Jan Prins

ABSTRACT

FÉLIX HERNÁNDEZ-CAMPOS: Generation and Validation of
Empirically-Derived TCP Application Workloads.
(Under the direction of Kevin Jeffay)

This dissertation proposes and evaluates a new approach for generating realistic traffic in networking experiments. The main problem solved by our approach is generating closed-loop traffic consistent with the behavior of the entire set of applications in modern traffic mixes. Unlike earlier approaches, which described individual applications in terms of the specific semantics of each application, we describe the source behavior driving each connection in a generic manner using the a-b-t model. This model provides an intuitive but detailed way of describing source behavior in terms of connection vectors that capture the sizes and ordering of application data units, the quiet times between them, and whether data exchange is sequential or concurrent. This is consistent with the view of traffic from TCP, which does not concern itself with application semantics.

The a-b-t model also satisfies a crucial property: given a packet header trace collected from an arbitrary Internet link, we can algorithmically infer the source-level behavior driving each connection, and cast it into the notation of the model. The result of packet header processing is a collection of a-b-t connection vectors, which can then be replayed in software simulators and testbed experiments to drive network stacks. Such a replay generates synthetic traffic that fully preserves the feedback loop between the TCP endpoints and the state of the network, which is essential in experiments where network congestion can occur. By construction, this type of traffic generation is fully reproducible, providing a solid foundation for comparative empirical studies.

Our experimental work demonstrates the high quality of the generated traffic, by directly comparing traces from real Internet links and their source-level trace replays for a rich set of metrics. Such comparison requires the careful measurement of network parameters for each connection, and their reproduction together with the corresponding source behavior. Our final contribution consists of two resampling methods for introducing controlled variability in network experiments and for generating closed-loop traffic that accurately matches a target offered load.

ACKNOWLEDGMENTS

First of all, I must thank Kevin Jeffay and Don Smith for their guidance and encouragement throughout my doctoral program. Their patience and friendship have been invaluable all these years. I also thank them, together with other faculty and student members of the Distributed and Real-Time Systems group (DiRT), for building a phenomenal infrastructure for Internet measurement and experimental networking research. DiRT students have greatly contributed to my doctoral experience, most especially Jay Aikat and David Ott.

My committee members and other collaborators have contributed tremendously to my efforts. I am specially in debt with Steve Marron and Andrew Nobel, who have greatly enriched the statistical side of my work. In this regard, being part of SAMSI's "Network Modeling for the Internet" program and of the inter-disciplinary Internet study group at UNC gave me superb opportunities to widen my understanding of Internet research. I must also thank UNC's Department of Computer Science as whole, including faculty, students and staff, for creating an outstanding research and teaching environment. Overall, my years at UNC were an incredibly positive experience.

I thank the National Science Foundation, IBM, Cisco, Intel, Sun Microsystems and others for supporting this work. I am specially grateful to the Computer Measurement Group (CMG) for their doctoral fellowship.

Finally, I thank my family for their support. Their constant example of hard-work, and their respect for intellectual endeavors has motivated me during my entire life. My wife's help with the editing of this manuscript was invaluable, as was her constant encouragement during my graduate studies. More than anybody else, my parents gave me my passion for knowledge, and it is to them that I dedicate this doctoral dissertation.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

LIST OF ABBREVIATIONS

ACK	Positive acknowledgment TCP segment
ADU	Application Data Unit
API	Application Programming Interface
AQM	Active Queue Management
BGP	Border Gateway Protocol
BPF	Berkeley Packet Filter
C.I.	Confidence Interval
CCDF	Complementary Cumulative Distribution Function
CDF	Cumulative Distribution Function
DAG	Data Acquisition and Generation
FIFO	First-In First-Out
FIN	TCP control flag indicating “no more data from sender”
FTP	File Transfer Protocol
GB	Gigabyte
GPS	Global Positioning System
HTML	HyperText Markup Language
HTTP	HyperText Transfer Protocol
I/O	Input/Output
ICMP	Internet Control Message Protocol
IP	Internet Protocol
IRC	Internet Relay Chat
ISP	Internet Service Provider
K-S	Kolmogorov-Smirnov test
KB	Kilobyte
Kpps	Kilo packet per second
LRD	Long-Range Dependence
MB	Megabyte
MIME	Multipurpose Internet Mail Extensions
MSS	Maximum Segment Size
MTU	Maximum Transmission Unit

Mbps	Megabit per second
NNTP	Network News Transfer Protocol
OSTT	One-Side Transit Time
PMA	Passive Measurement and Analysis
Q-Q	Quantile-Quantile
RED	Random Early Detection
RFC	Request For Comments
RST	TCP control flag indicating “connection reset”
RTT	Round-Trip Time
SMTP	Simple Mail Transfer Protocol
SSH	Secure Shell
SYN	Synchronize TCP control segment
SYN-ACK	Positive acknowledgement of SYN segment
TCP	Transport Control Protocol
UDP	User Datagram Protocol
UNC	University of North Carolina at Chapel Hill
URL	Universal Resource Locator