A Multi-path Transport Protocol to Exploit Network Diversity in Airborne Networks

V. Sharma,
ECSE Dept, RPI,
sharmv@rpi.edu

V. Subramanian,
ECSE Dept, RPI,
subrav@rpi.edu

K. Kar,
ECSE Dept, RPI,
koushik@ecse.rpi.edu

S. Kalyanaraman,
IBM Research Labs, Bangalore, India
shivkumar-k@in.ibm.com

K.K. Ramakrishnan,
AT&T Labs Research,
kkrama@research.att.com

B. Ganguly,
MIT Lincoln Labs,
ganguly@mit.edu

ABSTRACT

Airborne links are playing an increasingly important role in defense and military scenarios. This makes it important to investigate the performance of existing transport protocols over such links. Airborne links experience high variation in quality due to mobility, weather and other effects such as blockage.This translates to high loss rate environments for which current protocols are not designed. Consequently, the standard protocols fail to deliver high data rates. In prior work, we presented Loss-Tolerant TCP (LT-TCP) which is designed for such airborne links.We have now extended this protocol to take advantage of the availability of multiple paths in the network.We call this Multi-Path LOss-Tolerant TCP (MPLOT). In this paper, we investigate the statistical characteristics of airborne links using data gathered from actual experiments and develop mathematical models for such links. We then proceed to use these models to test MPLOT in the ns-2 network simulator. Our results show that while standard protocols such as TCP-SACK are unable to perform well, MPLOT can deliver higher goodput and lower latency than conventional transport protocols like TCP-SACK.MPLOT can leverage path diversity in the network to deliver even higher goodput as number of paths increase even while the total bandwidth remains fixed.

I. Introduction

The concept of Network Centric Operations (NCO) has been a central theme for Department of Defense (DoD) for over a decade. NCO proposes networking well informed systems across geographical locations so that information can be shared, processed and directed to the system(s) that need it with minimum delay. This would improve mission efficiency by responding to events quickly. Advances in communications technology, especially wireless communications have allowed this concept to be extended to the battle-space.

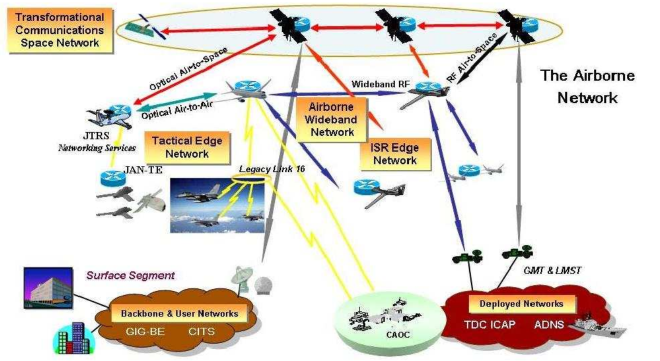

Fig. 1. Airborne networks are crucial for connecting geographically dispersed systems in battle-space.

Airborne networks play a key role in implementing NCO by connecting geographically dispersed locations that may not be able to connect due to terrain or range limitations. Figure 1 depicts the crucial role played by airborne networks in connecting different systems in battle-space. Such networks would use internet like transport protocols to transfer data. However, unlike commercial wireless networks, airborne networks are formed rapidly in battle conditions. We can expect them to face considerable uncertainty in operating conditions (*e.g.* jamming, interference, high mobility, channel impairments and capture effects). As a result, conventional transport protocols would see variable capacity, unpredictable packet erasures and volatile delays from such paths. Hence, it is paramount that transport protocols tolerate such volatility of while maintaining stable goodput (data rate) and latency for the systems using them.

One way to accomplish this objective is to use the diversity offered by wireless airborne networks. The main source of diversity is the existence of multiple paths in wireless networks. Nodes in a wireless network do not need an explicit physical connection to exchange information. Hence, wireless networks offer an increased

978-1-4244-2677-5/08/$25.00 2008 ©IEEE

opportunity to form multiple paths. The emergence of multi-homing and directional transmission has made possible the use of multiple paths in parallel with negligible inter-path interference. Each path experiences delay and capacity variations and losses independent of other paths. We denote this existence of independent paths in the network as network diversity. Transport protocols can potentially use network diversity to counter the volatility of a single path by transmitting data across different paths intelligently. Such transport protocols should ideally aggregate capacities of multiple, lossy paths and leverage diversity among paths to yield stable and high goodput *e.g.* such a protocol can chose to transmit critical data on paths experiencing "better" conditions and use other paths for transporting data not required urgently, thus improving the goodput compared to a single path.

In this paper, we present *MPLOT*, the *Multi Path LOss-tolerant Transport* protocol to attain the above mentioned objectives. MPLOT uses block Forward Error Correction (FEC) coding at the transport layer to counter high packet erasures. MPLOT counters path volatility by constantly monitoring and adapting to the changing conditions of paths. In particular, MPLOT transmits urgently required packets on paths that have shorter delays and lower losses while other packets are transmitted on longer paths. MPLOT overcomes the out of order arrival problem associated with multi-path protocols by using sequence agnostic properties of FEC coding and intelligent use of paths. We have already shown in [1] that MPLOT is able to effectively use multiple paths even in presence of high packet losses to deliver higher goodput.

In the next section, we discuss some prior work to address the problem of high packet losses in wireless networks, use of multiple paths and modeling packet erasures in wireless networks. We follow it by providing a brief description of MPLOT. We then use the measurements from actual experiments performed on airborne links to generate a model for packet losses on such links and then use the model in ns-2 simulations to show that MPLOT uses multi-path diversity efficiently to deliver high goodput in presence of losses that are similar to that observed in airborne networks.

II. RELATED WORK

Baldatoni *et al.* [2] proposed a version of TCP with FEC (but without adaptivity) that works for small error rates. Rizzo showed the feasibility of transport-layer high-speed FEC computation in [3]. Although [3] mentions the idea of FEC in TCP, a specific scheme has not been studied. Recent attempts at FEC with TCP have met with limited success ([4] for less than 10% erasure rates). Success with higher erasure rates have not been reported to the best of our knowledge. TCP Westwood [5] uses an estimate of output rate to guide congestion control, and has been effective for low erasure rates (under 5%). Overall, despite growing interest, there has been no clear baseline proposal that offers a significant increase in TCP performance over a wide range of erasure rates. Powerful and efficient error correction techniques have been proposed recently (see [6]) that enable such operations to be done efficiently. In this work, we assume the use of Reed-Solomon codes as the FEC mechanism.

We proposed a robust transport protocol called Loss-Tolerant TCP (LT-TCP) (see [7] and references therein) It estimates end-to-end packet losses to provision FEC adaptively to match the changing path conditions. We showed that LT-TCP consistently delivers higher goodput than standard transport protocols like TCP-SACK. However, LT-TCP uses only a single path and it's performance is dependent on the volatility of the path itself. Our work builds upon LT-TCP and extends its principles over multiple paths to counter the limitations of a single path.

There have been some attempts to use multiple-paths in the last decade. Lee *et al.* propose simple modifications to TCP like increasing the fast re-transmission threshold and delayed ACKs to address re-ordering issues in multi-path transport in [8]. Lim *et al.* [9] propose a multi-path transport framework for lossy networks, where they transmit multiple copies of a single packet on different paths. The performance of this scheme degrades rapidly as packet errors increase beyond $15-20\%$.

Several multi-path transport protocols ([10], [11], [12]) have been proposed for lossy wireless networks. However, such schemes provide limited packet redundancy, due to which they cannot handle multiple highly lossy paths effectively. Zhang *et.al.* propose mTCP in [10] to provide connection redundancy. However, mTCP only allows a single reverse path and no packet redundancy. As a result, it is inadequate for lossy networks. pTCP proposed by Hseih *et al.* in [11] aggregated bandwidths of different paths and separates reliability and congestion control functions but its packet redundancy is limited and only used in case of timeouts. RCP, proposed in [12] relies solely on retransmission of lost packets for recovery, which seriously limits its performance in lossy environments.

The bit-error rate characteristics of the wireless links differ from wired links significantly. Over the years, many models have been proposed to model the bit and packet errors exhibited by wireless channels. The popular 2-state Gilbert-Elliot model for bit-errors was proposed in [13] and [14]. In this model, each state corresponds to a specific channel quality which is either noiseless or completely noisy. An extension to the 2-state Markov model was suggested by Wang *et al.* in [15]. The authors in [15] suggested the use of multiple error states to reflect the changing network conditions. This model is appropriate for Rayleigh fading channels with slowly varying conditions. The Markov models are simple for analysis but cannot explain the heavy tailed behavior of packet bursts observed for wireless networks.

Another stochastic model was proposed by Carvalho *et al.* in [16] but the model was tested only for indoor IEEE 802.11g networks with links operating over a small distance. It's applicability to a wider range of networks/links was not investigated. Chaotic maps were proposed in [17] as another means of modeling wireless links errors. However, the estimation of model parameters was too complex and it ignored the longest runs/error burst observed in the actual traces. Jiao *et al.* proposed a new approach to simulate/model the packet errors in [13]. They proposed deriving "gap" and "run" distribution from the channel traces. A gap is the burst of corrupted packets observed between packets correctly transmitted and a run is the length of packets correctly received between two corrupted/lost packets. The cumulative distribution for such gaps and runs can be easily computed from the traces observed in actual networks. The gap/run length can be zero also, which implies a burst of corrupted packets or a length of packets correctly received. This approach has the advantage of being independent of any underlying assumption on the link or the network. The distribution also ensures that model emulates the error burst statistics observed in the channel itself. However, the gap model generally results in a non-parametric error distribution which cannot be analyzed easily.

III. MPLOT DESCRIPTION

In this section, we present a brief description of the Multi-Path Loss-Tolerant TCP (MPLOT) and describe its major components. We provide a more detailed description in [1]. MPLOT *constructs a block* of data and FEC packets that will be transmitted across the available paths. The block size B is determined by the path-windows such that the average transmission time is median Round Trip Time (RTT) of paths and is given by the following expression:-

$$B = \sum_{i=1}^{M} w_i \frac{RTT_{med}}{RTT_i} \quad (1)$$

Here, M is the number of available paths, w_i and RTT_i are the window size (in packets) and RTT of path i respectively. RTT_{med} is the median RTT across all paths.

Reliability is organized at the *aggregate flow manager*, across paths. We perform hybrid FEC/ARQ functions at the aggregate level across individual paths. Provisioning FEC packets based on aggregate parameters helps in averaging out the volatility of individual paths leading to a smoother, more stable performance.

Congestion control is done on a per-path basis. The per-path congestion window determines when a path can accept packets from the aggregate flow manager. Explicit congestion notification (ECN) on a path is used to distinguish congestion losses from those due to lossy links. Latest Aggregate reliability status is fed back on all paths. Thus the information about a packet received on a long path can reach the source through a shorter reverse path, shortening the effective round trip times. Moreover, if any single reverse path is subject to heavy loss or disruption, the feedback to update reliability status and advance windows for that path can arrive at the source through other paths. The source can thus advance the window for any path based upon the feedback received on a shorter or error-free reverse path. Per-path disruptions in the forward direction will lead only to per-path timeouts like in TCP, but will not affect the congestion window dynamics of other paths.

MPLOT uses *intelligent packet mapping* strategies to map packets to individual paths. When a path's congestion window advances and offers a transmission opportunity, an appropriate data or FEC packet (possibly out-of-order from a future block) is mapped to that path. Path parameters (loss rate, RTT, window) are combined into a *rank* function that is used to decide which packet is picked for a given transmit opportunity. In particular, higher ranked paths have shorter RTTs, lower loss rates and higher window sizes, and data and FEC packets from the earliest un-recovered block are mapped to these paths.

IV. MODELING AIRBORNE LINKS

In this section, we use the data from experiments conducted by Lincoln Laboratory on the wireless link between a Boeing 707 aircraft and 2 ground stations. The data used to generate these models were obtained from five test flights conducted by engineers in August 2006. We used these experiments to generate a record of IP packets sent/received and lost on the airborne links. We then used this data to develop a model for packet losses occurring in such networks and used it to run simulations in ns-2 to test MPLOT in conditions that are statistically similar to channels used in the experiments. We provide the details of these experiments in [18].

In section II, we discussed several models that have been proposed for bit/packet errors on wireless networks. The links tested in the experiment are different from the type of links studied for significant number of models discussed in section II. We cannot use models mentioned in [16], [17] because they have been only tested for paths operating over short distances in IEEE 802.11 wireless networks. The paths in such models are generally modeled as Rayleigh faded paths without a direct line-of-sight component. In contrast, the links in our experiments of airborne networks operate over larger distances than in wireless LANs and unlike Rayleigh faded channels, have a strong line-of-sight component.

Standard Markovian models like the Gilbert-Elliot model do not explain the error burst distribution observed from the experiments. We compare the complimentary cumulative distribution function (ccdf) of the length of packet error bursts from the traces and the Gilbert-Elliot model (configured from the traces observed from experiments) in Figure 2. It is clear that the Gilbert-Elliot model is a poor model for the link because the Gilbert model only accounts for the average statistics but does not explain the higher order statistics of error burst length.

Due to the limitations of the models mentioned above, we decided to use the gap model proposed in [13] to model the airborne wireless links in consideration. The gap model uses the cumulative distribution of the "gap" and "run" actually observed in the network. As a result, it will follow the error burst distribution observed in the networks exactly, unlike the Gilbert-Elliot model and is simple enough to be used in our simulations.

The distribution of gaps and runs is derived from traces of IP packets obtained from the experiments. We calculated the frequency of different gap and run lengths observed to derive the distributions. We then used these distributions to determine the gap/run length stochastically in our simulations.

In the next section, we use the gap/run distributions to test MPLOT's performance in airborne networks and show that it delivers higher and more stable goodput than compared to conventional transport protocols like TCP-SACK.

V. Performance Evaluation

In this section, we present the results from the simulations performed to analyze MPLOT performance over networks that exhibit packet loss characteristics similar to airborne networks.

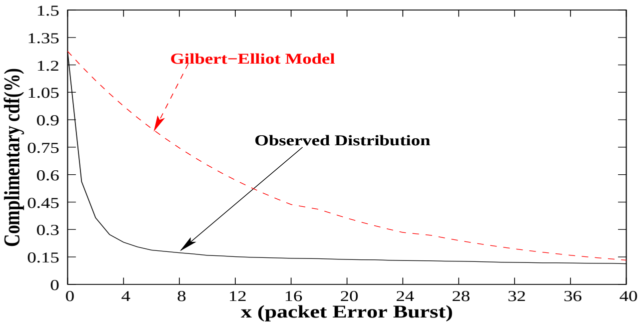

Fig. 2. The distribution of length of error burst from the Gilbert-Elliot model does not follow the actual distribution observed from the traces. Hence, the Gilbert-Elliot model would be a poor choice for modeling and simulation.

We used the standard ns-2 network simulator for our simulations. We consider a network topology that has multiple paths between a source and a destination. Figure 3 shows the topology used in simulations when 5 sources transport data over 3 paths to their respective destinations. This topology provides an abstraction of the physical routes (paths), where the different parallel paths in the topology correspond to different, possibly overlapping routes in the underlying network. In our simulations, we vary the number of paths and delays of paths but keep the total bottleneck bandwidth fixed at $10Mb/s$ textite.g. when we simulate 2 paths, then the bandwidth of each path is $5Mb/s$ and when we simulate 4 paths, the bandwidth of each path is $2.5Mb/s$. This allows us to study MPLOT's ability to leverage diversity independent of bandwidth aggregation effects. Each path, unless stated otherwise has a Round Trip Time(RTT) of $80ms$, which is similar to the scale of delays that would be experienced on paths in airborne networks. We use ECN capable Random Early Detect (RED) queues with a minimum and maximum thresholds of 2500 and 5000 bytes respectively for our simulations.

A total of 5 MPLOT sources transmit on the paths with 5 UDP sources operating on each path. The transmission rate of UDP sources is fixed such that the combined UDP sources on each path share approximately half the path bandwidth. The erroneous links exhibit the same distribution for packet loss bursts that was calculated from experiments. The packet losses on each path are independent of losses on any other paths. In the subsequent sections, we report the average values and their 95% confidence interval obtained after running 8 simulations for each set of parameters. We ran each simulation for 300 seconds to obtain the values for computation.

MPLOT exploits the diversity across paths to gain additional goodput over what can be obtained from a single path with the same aggregate capacity. The diversity can be due to differences in losses (loss diversity) on paths or

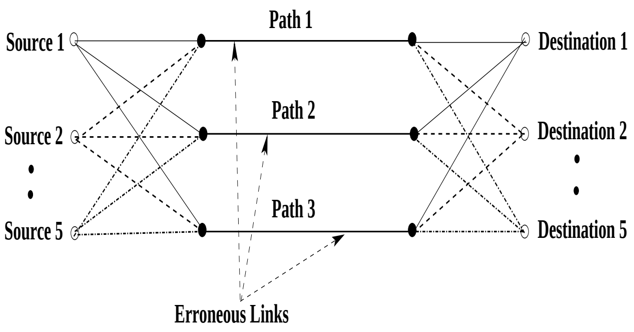

Fig. 3. The Network Topology Simulated for 3 paths with 5 sources. The packet losses on each path are independent.The total bandwidth is $10Mb/s$ and each path has a Round Trip Time (RTT) of $80ms$.

the difference in delays of paths (delay diversity). In the following sections, we use simulations to show MPLOT's ability to exploit loss and delay diversity and compare it's performance under different conditions to the performance of standard single-path TCP-SACK protocol and a multi-path protocol pTCP([11]).

We select pTCP for comparison with MPLOT because it also organizes reliability at an aggregate level and performs congestion control of each path independently. It also, with a limited capability provides packet redundancy by mapping the same packet to more than one path in case of a timeout.

A. Loss Diversity

In this section, we study MPLOT's ability to exploit diversity across paths because of differences in losses on each path. For this purpose, we simulate MPLOT for different number of paths while keeping delays and bandwidths of the paths identical. We also look at the performance of TCP-SACK(which can only use one path) and pTCP (which can use multiple paths) in the same conditions and with the same bottleneck capacity. Figure 4

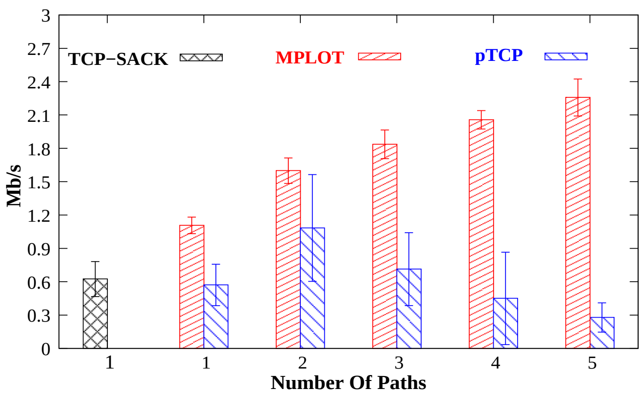

Fig. 4. Goodput achieved by MPLOT and pTCP for 1 through 5 paths. pTCP delivers a significantly low goodput and is unable to use diversity across paths to gain goodput, unlike MPLOT

compares the aggregate average goodput achieved (with

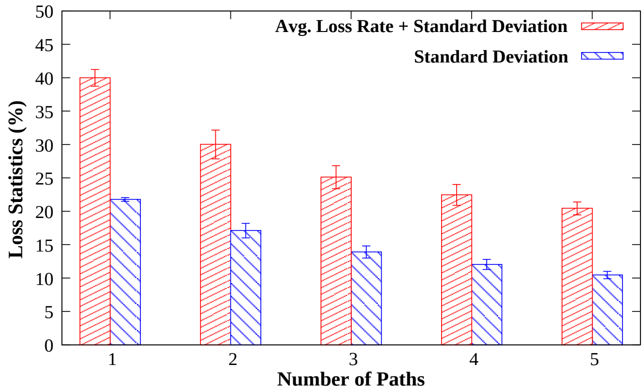

Fig. 5. Standard Deviation of packet losses measured by a MPLOT source. As number of paths increase from 1 to 5, the standard deviation reduces by half from 22% to 10%.

a 95% confidence interval) by 5 TCP-SACK, pTCP and MPLOT sources when operating over different number of paths. We observe that TCP-SACK sources only manages a total goodput of about $0.6Mb/s$ from a path with bandwidth of $10Mb/s$ while MPLOT on a single path itself achieves $1.1Mb/s$. We also note that the goodput of MPLOT sources increases with the number of paths (even though the total bandwidth is fixed at $10Mb/s$) and achieves a value of $2.25Mb/s$ for 5 paths. This is about a 100% improvement over the single path goodput. We also note that pTCP's performance, though better than TCP-SACK is poor compared to MPLOT. In fact, pTCP's goodput declines as number of paths increases beyond 2. This decline in pTCP's goodput is due to the fact that unlike MPLOT, pTCP maps packets to paths in order. This mapping scheme is unable to fully use the diversity across paths and results in a loss of goodput as number of paths increases.

The reason that MPLOT is able to get such a significant improvement in goodput as number of paths increases while total bandwidth remains fixed is evident by observing the aggregate loss statistics, measured by an MPLOT source, in Figure 5. We note that the average standard deviation of packet loss measured across paths reduces from 22% for 1 path to 10% for 5 paths. This is a reduction of more than 50%. This reduction is due to the fact that MPLOT constantly monitors losses on each path and maps packets to paths accordingly. As a result of this reduction, MPLOT allocates proportionally less number of PFEC packets to a block of same size which is evident from the reduction in the sum of aggregate packet loss rate and standard deviation of packet loss rate in Figure 5.

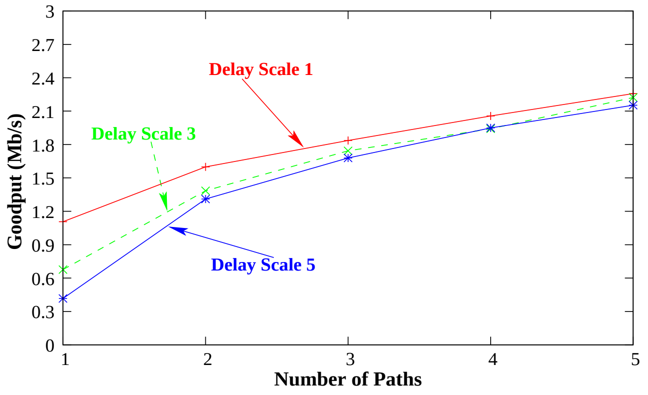

Fig. 6. Goodput obtained by MPLOT for different number of paths as delay scale changes to 1,3 and 5. As number of paths increases, the difference in goodput obtained for different delay scales reduces to almost zero.

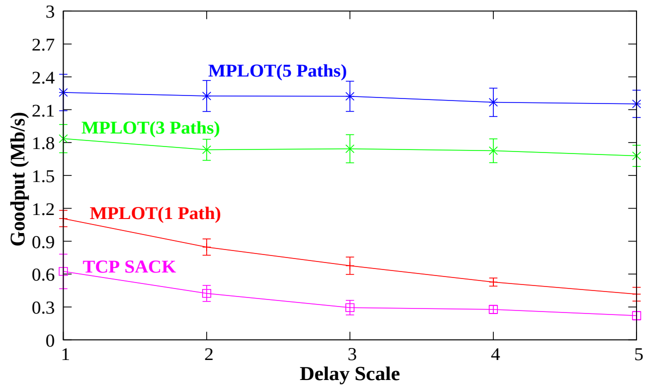

Fig. 7. Goodput obtained by MPLOT on 1,3 & 5 paths and TCP SACK for different delay scales. The goodput for TCP SACK and single path MPLOT reduces with increasing delay scale but the goodput of MPLOT over 3 and 5 paths remains almost constant with delay scale.

B. Delay Diversity

Different paths in a network may have different delays. MPLOT uses the shorter reverse paths to send feedback for packets sent on paths with longer delays. As a result, the apparent round trip time of the longer paths is reduced. *e.g* if we use 2 paths- one with a one way delay of 50 ms and second with a one way delay of 5ms, then the feedback for a packet sent on the longer path can potentially be received in $50 + 5 = 55ms$ instead of $100ms$. This is a 45% reduction in RTT which has far reaching implications.

MPLOT updates the Selective ACK(SACK) scoreboard for packets received on all paths and sends it back on different paths. The SACK information received from the shortest path will be the latest and can be used to update parameters for all the paths. This would result in a more gradual degradation in goodput with increasing delay than single path protocols and multi-path transport protocols that do not use the shorter paths to compensate for higher delays on some paths.

In order to study the ability of MPLOT to exploit this "delay diversity", we scale the delay of one the paths with a "delay scale" factor. A delay scale of 3 implies that the delay of one of the paths is thrice the delay on other paths *i.e* the ration of maximum to minimum RTT is 3. The paths are independent and have equal bandwidths. We then observe the difference in goodput achieved for different delay scales.

The goodputs obtained for different paths with delay scales of 1,3 and 5 are shown in Figure 6. We note that as the number of paths increases, the difference in goodput for delay scales 1 and 5 reduces from $0.7Mb/s$ ($1.1Mb/s$ vs $0.4Mb/s$) for 1 path to less than $0.1Mb/s$ ($2.25Mb/s$ vs $2.15Mb/s$) for 5 paths. We observe the same behavior for goodput obtained for delay scale of 3.

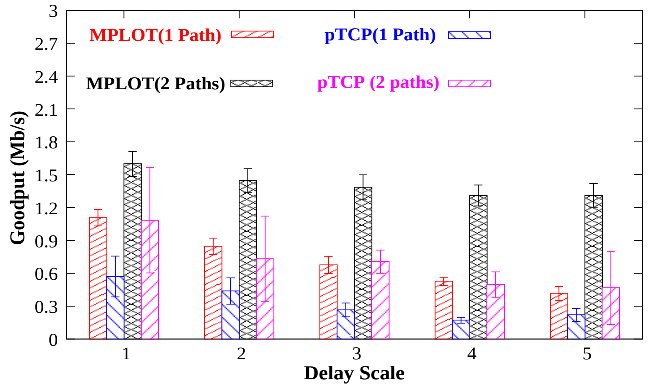

Fig. 8. Goodput obtained by MPLOT and pTCP on 1 & 2 paths and for different delay scales. The goodput for MPLOT and pTCP reduces with increasing delay scale.For 2 paths, the rate of goodput reduction for MPLOT is considerably lower than pTCP.

For comparison with TCP SACK, we show the goodput achieved by MPLOT over 1,3 and 5 paths for delay scales of 1 through 5 along with goodput of TCP SACK with a 95% confidence interval in Figure 7. We observe from the Figure that the goodputs of single path MPLOT and TCP SACK reduces significantly with increasing delay scale. However, the increasing delay scale does not have an adverse effect on the goodput achieved by MPLOT over 3 and 5 paths. This shows that MPLOT is using the shorter reverse path effectively to compensate for longer delays in the forward direction.

A comparison with pTCP in shown in Figure 8 for the same conditions. We observe that even for 2 paths, the reduction in goodput of MPLOT is less severe with delay scale than pTCP. pTCP operating over 2 paths sees it's goodput reduce from $1Mb/s$ for delay scale 1 to $0.5Mb/s$ for a delay scale of 5, which is a reduction of 50% while MPLOT's goodput only reduces by 20%. This shows that MPLOT is able to effectively use shorter reverse paths to counter delay variations among paths.

VI. Summary and Conclusions

In this paper, we looked at airborne networks which are playing an ever-increasing role in defense and military scenarios. Airborne networks are exposed to challenging environments due to the harsh conditions under which they operate; namely, weather effects, wing and body blockage, constant mobility and changing orientation of the airborne nodes. In earlier work, we had shown how current transport protocols such as TCP-SACK are unable to perform well on airborne networks. We had proposed a protocol called Loss-Tolerant TCP that was designed to overcome the limitations of TCP-SACK when dealing with highly lossy links.

In this paper, we extend our work and develop a protocol called Multi-Path Loss-Tolerant (MPLOT) that can realize significant bandwidth gains through effective use of any available multiple heterogeneous paths. While MPLOT is conceptually similar to LT-TCP, several additional challenges include the need to perform efficiently over multiple paths, congestion control over multiple paths and others. Our techniques include providing reliability on a multi-path basis ,congestion-control on a per-path basis, a packet mapping scheme etc.. To test our mechanisms, we used data gathered from actual test flights and initially developed a model for the airborne links. This model was then used to drive ns-2 simulations where we compared the performance of TCP-SACK, pTCP and MPLOT. We show how MPLOT outperforms TCP-SACK and pTCP by exhibiting better loss and delay-scale diversity.

In summary, in this paper, we presented the challenges faced by airborne networks and the need for protocol solutions to mitigate these challenges. We developed a model for the wireless links using actual data gathered from test flights. We then used these models to show how our transport protocol MPLOT, which exploits path and loss diversity and out-performs current solutions.

Acknowledgment

We thank our sponsors MIT Lincoln Labs (grant letter PO # 7000013888), AT&T Labs Research for their support and our collaborators from MIT Lincoln Labs, Stephen McGarry and Linda Zeger for insightful discussions.

References

[1] V. Sharma, S.Kalyanaraman, K.Kar, K. Ramakrishnan, and V. Subramanian, "MPLOT: A transport protocol exploiting multipath diversity using erasure codes," in *IEEE INFOCOM*, 2008.

[2] L. Baldantoni, H. Lundqvist and G. Karlsson, "Adaptive End-to-End FEC for Improving TCP Performance over Wireless Links." in *ICC*, June 2004.

[3] L. Rizzo, "On the feasibility of software fec," dEIT Technical Report LR-970131. [Online]. Available: citeseer.ist.psu.edu/rizzo97feasibility.html

[4] T. Anker, R. Cohen, and D. Dolev, " Transport Layer End-to-End Error Correcting," The School of Computer Science and Engineering , Hebrew University, Tech. Rep., 2004.

[5] S. Mascolo, C. Casetti, M. Gerla, M. Y. Sanadidi, and R. Wang, "TCP Westwood: Bandwidth Estimation for Enhanced Transport over Wireless Links," in *Mobile Computing and Networking*, 2001, pp. 287–297.

[6] M. M. J. W. Byers, M. Luby and A. Rege, "A Digital Fountain Approach to Reliable Distribution of Bulk Data," in *SIGCOMM*, Aug-Sep 1998.

[7] V. Subramanian, S. Kalyanaraman, and K.K.Ramakrishnan, "An end-to-end transport protocol for extreme wireless network environments," in *IEEE MILCOM*, 2006.

[8] Y. Lee, I. Park, and Y. Choi, "Improving tcp performance in multipath packet forwarding schemes," *Journal of Communications and Networks*, vol. 4, no. 2, pp. 148–157, june 2002.

[9] J. Chen, K. Xu, and M. Gerla, "Multipath tcp in lossy wireless environment," in *Proc. IFIP Third Annual Mediterranean Ad Hoc Networking Workshop (Med-Hoc-Net '04)*, 2004.

[10] M. Zhang, "Understanding internet routing anomalies and building robust transport layer protocols," Ph.D. dissertation, Department of Computer Science, Princeton University, sep 2005.

[11] H.-Y. Hsieh and R. Sivakumar, "A transport layer approach for achieving aggregate bandwidths on multi-homed mobile hosts," in *Proc. ACM Mobicom*, 2002.

[12] H. Hsieh, Y. Z. K.-H. Kim, and R. Sivakumar, "A receiver-centric transport protocol for mobile hosts with heterogeneous wireless interfaces," in *Proc. ACM Mobicom*, 2003.

[13] C. Jiao, L. Schwiebert, and B. Xu, "On modeling the packet error statistics in bursty channels," in *LCN '02: Proceedings of the 27th Annual IEEE Conference on Local Computer Networks*. Washington, DC, USA: IEEE Computer Society, 2002.

[14] X. Yu, J. W. Modestino, and X. Tian, "The accuracy of gilbert models in predicting packet-loss statistics for a single-multiplexer network model," in *INFOCOM 2005. 24th Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings IEEE*, vol. 4, Mar. 2005, pp. 2602–2612.

[15] H. S. Wang and N. Moayeri, "Finite-state markov channel-a useful model for radio communicationchannels," *IEEE Transactions on Vehicular Technology*, vol. 44, no. 1, Feb. 1995.

[16] L. Carvalho, J. Angeja, and A. Navarro, "A new packet loss model of the IEEE 802.11g wireless network for multimedia communications," *IEEE Transactions on Consumer Electronics*, vol. 51, no. 3, pp. 809–814, Aug. 2005.

[17] A. Kopke, A. Willig, and H. Karl, "Chaotic maps as parsimonious bit error models of wireless channels," in *INFOCOM 2003. Twenty-Second Annual Joint Conference of the IEEE Computer and Communications Societies. IEEE*, vol. 1, 2003.

[18] B. Ganguly, V. Subramanian, S. Kalyanaraman, and K.K.Ramakrishnan, "Performance of disruption-tolerant network mechanisms applied to airborne networks," in *IEEE MILCOM*, 2007.