

## NEW SOLUTIONS FOR FAST DATA TRANSFERS IN HBDP NETWORKS

Dan SCHRAGER<sup>1</sup>

*A broad selection of methods and techniques for data transfers is presented, starting with the TCP/IP protocol, its main enhancements, new non TCP protocols, and the class of applications and libraries able to sustain massive bulk transfers in HBDP networks. Distinctive features and comparison based on performance, effectiveness, efficiency and fairness are emphasized in each case. As a result, a new research direction has been identified, regarding generalization of the pipeline inter-process communication mechanism at networking level, by making fast interconnections between pairs of data producer/consumer processes distributed over the Internet. Preliminary experimental results are included as a proof of concept.*

**Keywords:** pipeline, parallel data streaming

**MSC2000:** 53C05.

### 1. Introduction

In the context of grid and cloud computing, the ability to transfer large amounts of data over wide area networks is a requirement for many actual research fields interested in HPC over geographically distributed data. For example, the Atlas experiment [2] is transferring 10 PB of data per year. The VIRGO collaboration [31] has storage needs in the order of many hundreds of TB, and very low latency requirements. In other scientific areas, like real-time data visualization of remote instrumentation or very-high-definition video streaming/conferencing, research optical networks (e.g. GLIF [12] or OptIPuter [26] ) are used for data streaming at high velocity.

TCP is the protocol most widely used in the Internet as it provides an important service, namely reliable data transfers as continuous byte streams. TCP's implementation is not restricted, so that it works both in normal Ethernet networks as well as in high bandwidth delay product networks having also a high bit error rate. Optimizing transfers of large data volumes in order to achieve an increased efficiency has led to new variations of TCP which address some of its initial limitations.

Another research direction is the design of new transport protocols. Such approaches usually involve changes in network routers, and therefore are mainly

---

<sup>1</sup>PhD Student, University *Politehnica* of Bucharest, Romania, e-mail: [danschrager@gmail.com](mailto:danschrager@gmail.com)

used in dedicated scientific networks that justify the high costs associated with implementation and dissemination.

There are numerous application-level performance improvement solutions. Among these, some use UDP to transfer actual data and TCP for the control channel. Others rely on network parallelism implementation, by using multiple TCP connections simultaneously.

A comparative analysis of the advantages and disadvantages of the above mentioned techniques and methods, without being exhaustive because of the existing huge volume, has served for the identification of new research directions. In essence it is about creating a *pipeline* mechanism of high capacity and speed which interconnects pairs of data producer/consumer processes, running on different machines in the network. Disentanglement between the new data transfer algorithms from data access details will lead to an architecture able to handle both regular file transfers and data streaming of size a priori unknown. This kind of interconnection will also make good use of existing storage and data access systems (e.g. Castor [5] at Cern [6]) and/or local system commands (e.g. cat, dd, tar, pmr [28]).

The next section presents a related work analysis of the networking transfers research domain, with accent on different variations of TCP, non-TCP protocols, and the large class of transfer applications and libraries. In Section 3 are presented preliminary experimental results using the new model of parallel data streaming and its main advantages over conventional solutions. Finally, conclusions are included in Section 4.

## 2. Related Work

This section presents a comparative analysis of existing research in the field of network transfers. Given the huge volume of existing information, only a selection of representative work has been included. Note also that distributed systems, P2P and overlay, have been deliberately omitted, as having little relevance to my current research interests.

### 2.1. Transmission Control Protocol

TCP [17], [21] is the protocol that standardizes how reliable transmission of data streams in the Internet are performed, as a constituent part of the TCP/IP model. TCP provides a bidirectional connection, allowing efficient exchange of data regardless of the characteristics of individual transport networks. Because TCP is typically implemented at the operating system level, it can be used by any application with a uniform access interface.

**2.1.1. TCP Performance.** Although able to adapt to a wide range of network types, TCP reaches a maximum transfer rate determined by its congestion avoidance algorithms. According to study in [24], the TCP bandwidth (BW) is approximated by equation:

$$BW = \frac{MSS}{RTT} \frac{C}{\sqrt{p}} \quad (1)$$

where  $p$  is the packet loss ratio, equal to the number of retransmitted packets divided by the total number of packets transmitted, and  $C$  is a constant depending on TCP characteristics.

In an HBDP network, with  $BW = 1$  Gbps,  $MSS = 1500$  bytes,  $RTT = 100$  ms,  $C = 1.2$  (standard TCP), substituting in equation (1) produces a  $p$  of a  $2 * 10^{-8}$  order of magnitude, which corresponds to a BER of a  $10^{-12}$  order of magnitude, challenging even for fiber optics. In addition, the AIMD algorithm proves to be too slow in its *slow-start* phase which affects its ability to probe and utilize fully the path's available bandwidth.

A solution to the performance issues TCP has in HBDP networks is the use of multiple simultaneous TCP connections, approach used by many transfer applications described in subsection 2.4 below. Collectively,  $N$  concurrent TCP connections behave similarly to a single connection with an  $MSS N$  times larger than that of one connection, at least in the additive increase (slow-start) phase. Also, noncongestive packet losses, misinterpreted as a signal for halving the congestion window, do not happen to all  $N$  connections simultaneously, so less bandwidth is released and the global transmission rate is less affected.

The following sections will present other solutions to improve TCP performance and other effective transfer techniques.

## 2.2. Various TCP implementations

Research in this area has focused on improving mechanisms for retransmission and congestion control in classical TCP. Usually changes are concerning the data sender in order to maintain compatibility with TCP. Implementations are performed in the operating system and therefore have some difficulty to disseminate.

**2.2.1. TCP SACK.** It is described in its final form in RFC 2018 [23] in 1996, but it was proposed for the first time since 1988 in RFC 1072 [33]. The system of cumulative packet confirmations characteristic to TCP (Reno) causes desynchronizations and reduction in the transmission rate in the case of multiple packet errors in the same data window. SACK's mechanism of selective acknowledgments is a strategy which corrects TCP's behavior when multiple data segment losses occur. Thus, the data receiver can inform the sender about each correctly received data segment, so that the transmitter will resend only the lost segments.

From a technical standpoint, TCP SACK uses two new TCP options, namely *Sack-Permitted*, which negotiates the SACK use when connection is established, and the *Sack* option itself, which may be used throughout the connection lifetime. TCP SACK does neither strengthen nor weaken TCP's security properties. In terms of eliminating unnecessary retransmissions, TCP SACK efficiency mechanism is strictly superior to TCP Reno conventional implementation.

TCP SACK behavior in response to congestion is discussed in [30], where it is shown that indeed significantly increases data transfer performance, especially for highly loaded networks or even in the absence of congestion, in case of multiple clustered errors.

**2.2.2. Multipath TCP.** It is an extension of TCP able to simultaneously use different paths in the network for interconnecting computers with multiple network interfaces and different IP network addresses. Currently (2012) is presented in an IETF advanced document [1], being developed since 2010.

The Multipath TCP technique is transparent to applications as the system is implemented in a new sockets library at the operating system level. Multipath TCP is compatible with TCP Reno at the network system function calls, hence there is no need to rewrite existing applications. When connecting to hosts where it is not yet implemented, it is able to use standard TCP by default. As it uses several TCP connections for each main interconnection, congestion management is amended to avoid aggressivity against normal TCP connections and it also sends data over paths with the highest available bandwidth. The Multipath protocol is implemented using new TCP options (e.g. MP\_CAPABLE, MP\_JOIN, MP\_PRIO, MP\_FASTCLOSE, MP\_FAIL, DSS, ADD\_ADDR, REMOVE\_ADDR) to ensure unhampered crossing of diverse intermediate middleboxes like NAT, PEP, firewalls, IDS, etc. Protocol security is ensured using cryptographic keys of type SHA-1. Implementation minimizes memory requirements and processing by applying opportunistic retransmission strategies, penalizing slower paths and autoconfiguring network buffer sizes.

Multipath TCP performance has been studied in the laboratory in wireless environments with applications using the HTTP protocol, good results being obtained comparatively to TCP Reno in [8].

### 2.3. Non-TCP protocols

Such protocols control congestion using feedback from intermediate routers. Implementations, although efficient, require changes in routers, sender and receiver and therefore dissemination is difficult in terms of the high prices involved.

**2.3.1. XCP** [9]. It is a new protocol that controls better than TCP the congestion and ensures efficiency, fairness and stability of transmission in HBDP networks. The protocol generalizes the ECN [19] protocol extension and introduces a new concept of decoupling utilization control from fairness control. This increases its flexibility in implementing policies for the allocation of available bandwidth, allows to distinguish between losses due to congestion and those caused by transmission errors, and increases security by quickly detecting deviations from the protocol.

ECN is extended at router level by including a congestion header in each transmitted packet. Field  $H_{feedback}$  is initialized by sender, modified by each router along the path, and the whole header is copied back by receiver in acknowledgement packets returned to sender, so that it can adapt properly its congestion window,  $cwnd$ , according to constrains imposed by the most loaded router. Routers compute

efficiency based on a MIMD formula with (two) parameters invariant to number of participating flows, delays or bottleneck bandwidth, determined on the basis of ensuring stability, and with the purpose to increase the traffic rate proportional to the spare bandwidth. Fairness control is based on an AIMD algorithm which splits equally the excess of available throughput between sources, while the throughput decrease of a flow is proportional to its current throughput. Additionally, a bandwidth shuffling technique is used to ensure rapid convergence when efficiency is around optimal. Therefore, at least 10% of the traffic is redistributed among individual flows.

Compared to TCP, XCP is stable and efficient regardless of the relative throughput of individual flows. Simulation results show that XCP is performing as well as TCP in normal networks but outperforms it largely in HBDP networks. XCP ensure fairness, efficiency, smaller size of router queues and packet loss almost zero (less than one per million), both in steady traffic as well as variable. Although it can be implemented in routers requiring negligible computational power (several additions and three multiplications), XCP is likely to disseminate only in networks dedicated to scientific research and less on the Internet in general.

#### 2.4. Fast transfer applications and libraries

At network application level there are many available implementations, some included in libraries, based on different techniques such as using multiple parallel TCP connections, or UDP for data transfers and TCP only for control. Dissemination of applications is easy because we deal with software located outside the operating system.

**2.4.1. XFTP** [20]. It is a client/server application which extends the FTP protocol described in RFC 959 [16] with a multi-socket implementation.

The protocol extension applies only to the active mode of an FTP connection (commands MULT and MPRT) and is implemented by using  $n$  simultaneous TCP connections, with buffers to avoid deadlocks and multiplexed asynchronously in a single process. Transferred files are divided into  $m$  equal parts (of size 8K, with  $m >> n$ ), each sent via one of the available sockets, together with an identification field that tells the receiver its position relative to the beginning of file, and then reordered upon arrival.

XFTP managed to achieve a throughput as high as 90% across a satellite T1 link, compared to only 24% in the case of regular FTP.

**2.4.2. PSockets** [13]. It is a library of functions written in C++ that make applications obtain optimum use of the network without having to adjust the transmission window described in RFC 1323 [32].

The technique used, called network striping, is based on dividing data equally in a number of parts determined by the number of parallel connections created, and sending them asynchronously over the open sockets. In this way the inconvenience of reordering data upon arrival is eliminated, and thus very little overhead is encountered at reception, as compared to the method used by XFTP.

PSockets is a library easy to include in applications interested in moving efficiently bulk data without extra manual network tuning, because it uses functions with same interface as that provided by regular sockets.

**2.4.3. *Globus GridFTP*** [35]. It is a client/server file transfer application designed specifically for Grid environments, being part of GT4 [15].

The GridFTP protocol [34] is based on FTP and its extensions standardized by IETF, on RFC 2228 [22] (security) and on RFC 2389 [27] (negotiation), plus a number of new commands (e.g. SPAS, SPOR, ERET, ESTO, SBUF) and related options. Thus the protocol is capable of simultaneous transfers between multiple servers, use of parallel connections, partial file transfers (regions), automatic negotiation of transfer parameters (window, buffer) and restart and recovery of transfers disrupted by errors.

Implementation is modularized based on criteria of efficiency. The main components are the protocol interpreter (PI) at both server and client level and the data transfer process (DTP). In turn, DTP consists of three modules, data access, data processing and data channel protocol module - with direct access to the network. Security operations are based on GSI [14] and Kerberos [4], on both the control channel and optionally on the data channel.

Compared with other FTP servers, GridFTP gets better performances both in terms of achieved transfer speed as well as scalability by supporting many concurrent clients without excessive load.

**2.4.4. *RBUDP*** [11]. It is a data transfer application for high bandwidth networks, dedicated or capable of QoS (e.g. optical networks). Actual data transfers are done via UDP while the control traffic is signaled with TCP.

The RBUDP algorithm does not use congestion control or the slow-start mechanism of TCP and so is able to send data at user specified rates. Each UDP transmission phase is followed by retransmission of lost packets only, as signaled via TCP, until all errors, caused by UDP unreliability, are eliminated.

Experimental results demonstrate that RBUDP can perform large bulk data transfers at rates close to the full bandwidth of a link.

### 3. The New Data Streaming Model

It has become apparent from the above analysis that a new approach could generalize the task of high throughput data transfers in HBDP networks. The novel idea is based on combining pipelines existing at the operating system level with multiple concurrent TCP connections. In this way the pipe programming paradigm which interconnects data producing and consuming processes locally is extended between hosts distributed in the Internet, while transfer rates are in the same time maximized.

It is worth emphasizing the generality and elegance of the envisioned approach which although using a client/server networking application, strictly specialized in high velocity data streaming, manages to exploit, through encapsulation, the pipe

programming paradigm - specific to modern operating systems (Unix) - based on reuse of existing system components, as opposed to building new monolithic applications. Needless to say, the advantage of easy dissemination of applications will be preserved too.

Disentanglement between fast data transfer algorithms from data access details leads to an architecture able to handle both regular file transfers and data streaming of size a priori unknown. This kind of interconnection will also make good use of existing storage and data access systems, local system commands and of local applications at times uniquely specialized in storage access and retrieval (e.g. rftar [29]). Furthermore, multi-path implementation at application level will also become possible.

### 3.1. Preliminary experimental results

I was encouraged in the research and development of the new algorithms because my technology had early adopters. For example, the particle physics department of the Weizmann Institute of Science [36], the VIRGO experiment, the MWT2 grid center [25], all attracted by its flexibility and/or performance. Figure 1 illustrates the performance obtained by the latter, achieved during a disk-to-disk inter-cluster transfer between University of Chicago and Indiana University (members of MWT2) over a Starlight 10 Gbps link.

To evaluate the performance of the new transfer algorithms I created an application (bbftpPRO [3]) that implements the fast data streaming model (using pipes forked/joined over parallel TCP connections), alongside the simple file striping approach (where a file is divided in equal parts sent concurrently over existing sockets), used as a term of comparison. Integrating both algorithms in the same application has allowed accurate configuration of networking parameters of the TCP protocol, in

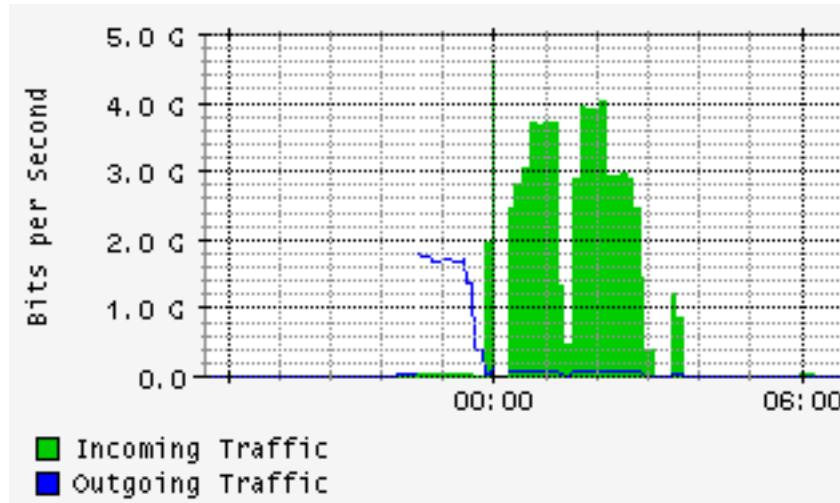


FIG. 1. Disk-to-disk streaming over a 10 Gbps Starlight link

the same way for each of the experimented transfer models. In addition, instrumentation was also common, which gave confidence in the results obtained measuring the transfer rate for each of the cases considered.

**3.1.1. File Striping vs. File Streaming Performance Evaluation.** In a first series of experiments I was interested to determine whether the use of pipes has an impact on transfer performance. For this purpose I performed transfers between pair of hosts linked by an ISP with a bottleneck bandwidth of 100 Mbps. Every time I transferred a disk file large enough to attain TCP equilibrium ( $\sim 300$  MB), in the file striping mode and via two *cat* processes interconnected by remote pipes in the fast parallel streaming mode. I varied the number of parallel connections and I recorded the transfer rate in both styles, as reported by bbftpPRO. The average speed measurements are presented in Fig. 2.

I have extended same type of experiments in the 1 Gbps domain too. This time, to avoid drag on performance caused by disk access delays, I transferred repeatedly a 400 MB in memory file, both striped and streamed with pipes. The average throughput as a function of the number of connections is presented in Fig. 3.

I also measured the performance of the two models over an IPoIB Infiniband network, capable of rates up to 10 Gbps, in a grid cluster, with nodes having each: 12 cores with hyper-threading enabled (total 24 processors), running at 3 GHz, and having 32 GB of installed RAM. To level the comparison field, a 10 GB in memory file was used during the striped and streamed transfers. The average throughput measurements are shown in Fig. 4, where the number of simultaneous connections varies between 1 and 24.

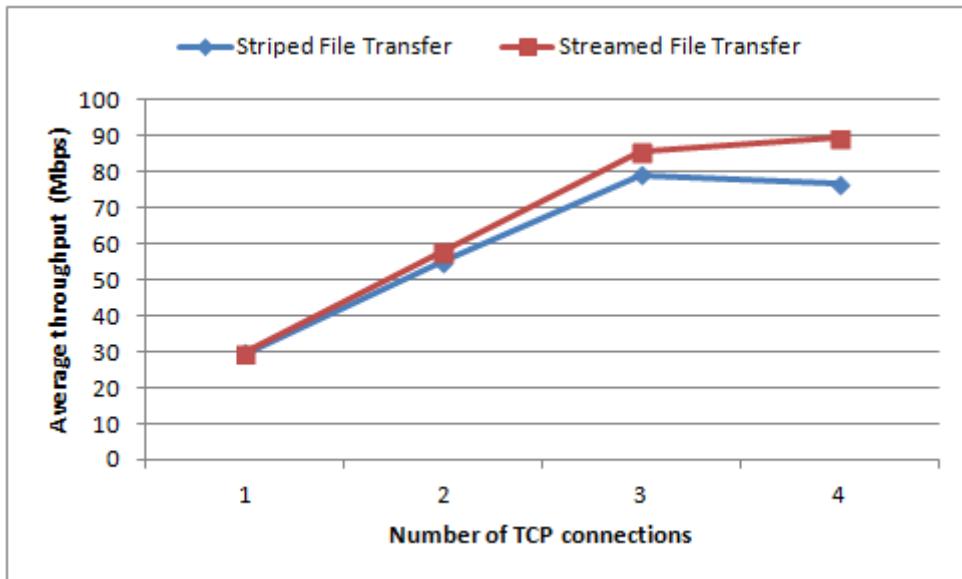


FIG. 2. File striping vs. file streaming, 100 Mbps network

**3.1.2. Analysis of Experimental Results.** From a quantitative point of view, the fast streaming algorithms are as rapid as those based on parallel file striping. This results from data shown in Fig. 2, 3 and 4, where the average total throughput increases approximately linearly with the number of parallel connections until reaching a saturation value of about 70 to 90% of the available bandwidth, for both methods as well.

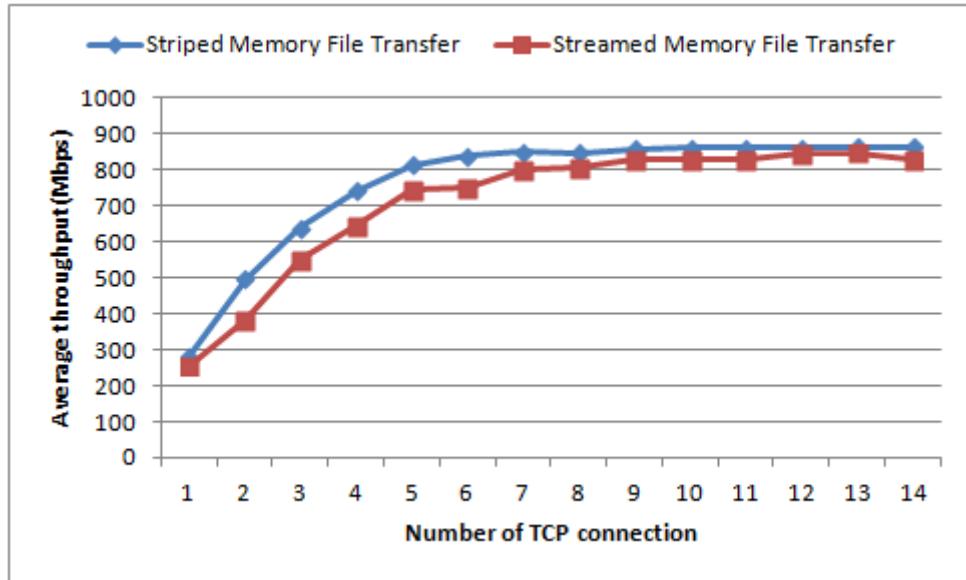


FIG. 3. File striping vs. file streaming, 1 Gbps network

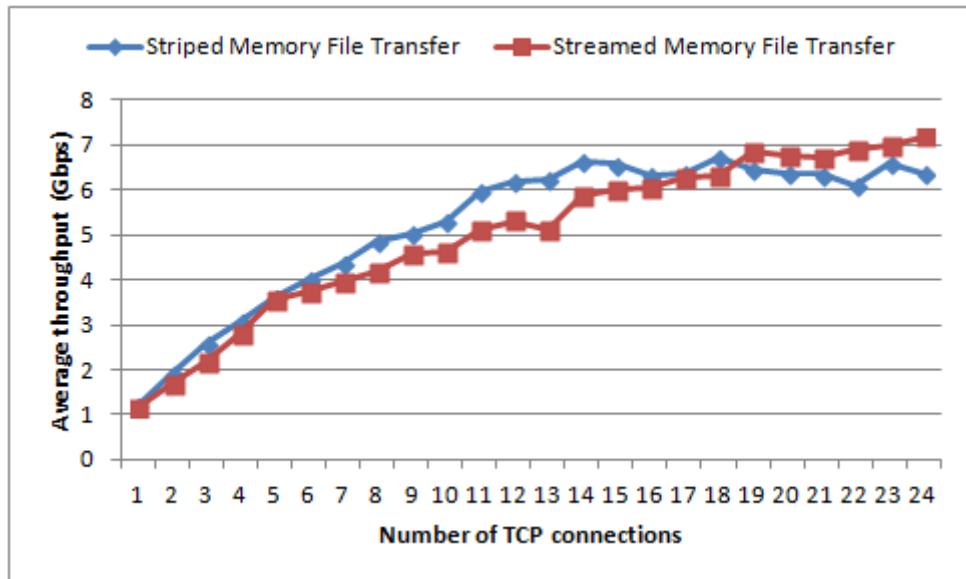


FIG. 4. File striping vs. file streaming, 10 Gbps network

While the parallel streaming rate is slightly superior to the regular striping rate in the 100 Mbps domain, in the 1 and 10 Gbps networks, where every microsecond counts, striping appears to have an early advantage. This happens because of the synchronization requirements of streaming which add a small overhead, absent in the case of (independent) striping. However, when transmission parameters are well tuned (e.g. increased buffer sizes) and for a larger number of parallel connections, the average throughput becomes mostly equal, as pictured in the same figures.

In any case, piped streaming holds an overall qualitative edge over file oriented striping, and as analyzed in a related paper, [10], is able of notably improved performance in cases of application level multi-path.

### 3.2. Design advantages of the pipe/streamed model

The related work section has emphasized both strengths and weaknesses of the existing network transfer solutions. Thus, changes to the TCP protocol have a long duration of standardization (e.g. Sack, Multipath) and a difficult dissemination - involving kernel upgrades. Non-TCP transport level protocols (e.g. XCP) require changes regarding congestion handling in routers as well as modifications of both data senders and receivers. Such requirements are difficult to accomplish and expensive and are therefore an obstacle to dissemination, except for dedicated scientific networks. Although they have the best chance to dissemination, libraries of network functions need to be included in new applications, while existing network applications are sometimes too specialized and mainly oriented toward file transfers.

My design has none of these drawbacks and instead has a few remarkable advantages such as:

- ease of dissemination, since it is implemented as an ordinary client/server application.
- generality, it transfers at high rate standard output/input of data processes spread over HBDP networks.
- reuse of existing storage system applications or local commands, through encapsulation of the pipe programming paradigm.
- supports both arbitrary data streaming and regular file transfers, by separating its fast transfer algorithms from the data access details.
- multi-path support, at application level.

## 4. Conclusions

This paper presented a comparative analysis of research in the field of network data transfers. Representative works in the area of TCP enhancements, non-TCP protocols, and applications or libraries dedicated to networking transfers have been included. The advantages and limitations of existing solutions have been set out. This led to a new research direction in the field of bulk transfers in HBDP networks, based on generalization of the pipe programming paradigm from its local (Unix)

meaning toward the Internet at large, while ensuring in the same time high transfer rates at parallel speed.

Preliminary experiments and results described in the previous section demonstrate the effectiveness of the new solutions envisaged. Thus the performance of the streamed transfers, in terms of achieved throughput is not worse than in the case of simple striping of files, under synchronization requirements, specific to data streaming. In a related paper [10] as well and planned ones, I will demonstrate in detail that the proposed model and the new fast transfer algorithms meet the modern requirements of todays research domains interested in transmitting massive bulk data elegantly, efficiently, effectively and at parallel speeds.

## REFERENCES

- [1] *A. Ford, C. Raiciu, M. Handley and O. Bonaventure*, TCP Extensions for Multipath Operation with Multiple Addresses draft-ietf-mptcp-multiaddressed-09, IETF draft, Jun. 2012.
- [2] *ATLAS Experiment*, URL: <http://www.atlas.ch/>
- [3] *bbftpPRO*, URL: <http://bbftppro.myftp.org/>
- [4] *B. Neuman and T. Ts'o*, Kerberos: An Authentication Service for Computer Networks, IEEE Communications Magazine, 32 (9). 33-88. 1994.
- [5] *CASTOR*, URL: <http://castor.web.cern.ch/>
- [6] *CERN*, URL: <http://public.web.cern.ch/public/>
- [7] *C. Jin, D. Wei and S. Low*, FAST TCP: Motivation, Architecture, Algorithms, Performance, Proceedings of IEEE Infocom, Hong Kong, Mar. 2004.
- [8] *C. Raiciu, Ch. Paasch, S. Barre, A. Ford, M. Honda, F. Duchene, O. Bonaventure and M. Handley*, How Hard Can It Be? Designing and Implementing a Deployable Multipath TCP, USENIX Symposium of Networked Systems Design and Implementation (NSDI'12), Apr. 2012.
- [9] *D. Katabi*, Decoupling Congestion Control and Bandwidth Allocation Policy with Application to High Bandwidth-Delay Product Networks, PhD thesis, Massachusetts Institute of Technology, Mar. 2003.
- [10] *D. Schrager and F. Radulescu*, Efficient Algorithms for Fast Data Transfers Using Long and Large Pipes in WAN Networks, 19th International Conference on Control Systems and Computer Science (CSCS-19), May 2013
- [11] *E. He, J. Leigh, O. Yu and T. DeFanti*, Reliable Blast UDP : Predictable High Performance Bulk Data Transfer, Proceedings of IEEE International Conference on Cluster Computing, 2002.
- [12] *GLIF Facility*, URL: <http://www.gliif.is/>
- [13] *H. Sivakumar, S. Bailey and R. L. Grossman*, PSockets: The case for application-level network striping for data intensive applications using high speed wide area networks, Proceedings of the IEEE/ACM SC2000 Conference, 2000.
- [14] *I. Foster, C. Kesselman, G. Tsudik and S. Tuecke*, A Security Architecture for Computational Grids, 5th ACM Conference on Computer and Communications Security, 1998.
- [15] *I. Foster*, Globus Toolkit version 4: Software for service-oriented systems, IFIP International Conference on Network and Parallel Computing, pp. 213, 2005.
- [16] *J. Postel and J. Reynolds*, FILE TRANSFER PROTOCOL (FTP), RFC 959, Oct. 1985.
- [17] *J. Postel*, Transmission Control Protocol, RFC 793, Sep. 1981.

---

- [18] *J. Schopf, M. DArcy, N. Miller, L. Pearlman and I. Kesselman*, Monitoring and discovery in a webservices framework: Functionality and performance of the globus toolkits mds4, Globus, Tech. Rep. ANL/MCS-P1248-04-5, 2005.
- [19] *K. Ramakrishnan, S. Floyd and D. Black*, The Addition of Explicit Congestion Notification (ECN) to IP, RFC 3168, Sep. 2001.
- [20] *M. Allman, S. Ostermann and H. Kruse*, Data Transfer Efficiency over Satellite Circuits Using A Multi-Socket Extension to the File Transfer Protocol (FTP), Proceedings of the ACTS Results Conference, NASA Lewis Research, 1995.
- [21] *M. Allman, V. Paxson and W. Stevens*, TCP Congestion Control, RFC 2581, Apr. 1999.
- [22] *M. Horowitz and S. Lunt*, FTP Security Extensions, RFC 2228, Oct. 1997.
- [23] *M. Mathis, J. Mahdavi, S. Floyd and A. Romanow*, TCP Selective Acknowledgment Options, RFC 2018, Oct. 1996.
- [24] *M. Mathis, J. Semke and J. Mahdavi*, The macroscopic behavior of the TCP congestion avoidance algorithm, Computer Communications Review, July 1997.
- [25] *MWT2*, URL: <http://mwt2.usatlasfacility.org/>
- [26] *OptIPuter*, URL: <http://www.optiputer.net/>
- [27] *P. Hethmon and R. Elz*, Feature negotiation mechanism for the File Transfer Protocol, RFC 2389, Aug. 1998.
- [28] *pmr*, URL: <http://zakalwe.fi/~shd/foss/pmr/>
- [29] *rftar*, URL: <http://castorold.web.cern.ch/castorold/DIST/CERN/savannah/CONTRIB.pkg/>  
Dan.Schrager@weizmann.ac.il/
- [30] *S. Floyd*, Issues of TCP with SACK, URL: [ftp://ftp.ee.lbl.gov/papers/issues\\_sa.ps.Z](ftp://ftp.ee.lbl.gov/papers/issues_sa.ps.Z), Jan. 1996.
- [31] *VIRGO Detector*, URL: <https://wwwcascina.virgo.infn.it/>
- [32] *V. Jacobson, R. Braden and D. Borman*, TCP Extensions for High Performance, RFC 1323, May 1992.
- [33] *V. Jacobson and R. Braden*, TCP Extensions for Long-Delay Paths, RFC 1072, Oct. 1988.
- [34] *W. Allcock, J Bester, J Bresnahan, A Chervenak, L Liming and S Tuecke*, GridFTP: Protocol Extensions to FTP for the Grid, Global Grid Forum Recommendation Document GFD.20, 2005.
- [35] *W. Allcock, J. Bresnahan, R. Kettimuthu, M. Link, C. Dumitrescu, I. Raicu and I. Foster*, The Globus Striped GridFTP Framework and Server, SC 05: Proceedings of the 2005 ACM/IEEE conference on Supercomputing, Nov. 2005.
- [36] *WIS*, URL: <http://www.weizmann.ac.il/particle/>