ATLAS: AI-Assisted Transport for AI Workloads in Data Center Networks

Leonardo Alberro*†, Pedro Casas†, Matías Richart*, Eduardo Grampin*
*Universidad de la República, Montevideo, Uruguay †Austrian Institute of Technology, Vienna, Austria
{lalberro, grampin, mrichart}@fing.edu.uy, pedro.casas@ait.ac.at

***Abstract*—Distributed AI training workloads are reshaping data center network (DCN) traffic patterns, generating structured and predictable communication behavior that differs fundamentally from traditional cloud applications. Existing transport protocols such as DCTCP and flow-aware variants rely on reactive flow classification, where transport decisions are made without early knowledge of flow characteristics, limiting their ability to optimize latency-sensitive communication. We propose ATLAS, an AI-assisted transport framework for distributed AI workloads in DCNs. ATLAS exploits the regularity of AI communication patterns to enable prediction-driven flow-size estimation and proactive transport-layer decision making. By exposing flow characteristics early, ATLAS aims to improve congestion control and flow prioritization beyond conventional reactive mechanisms. We present the overall ATLAS framework, including workload-aware traffic characterization, learning-based flow-size prediction, and integration with DCTCP-style transport protocols. Preliminary results show that early flow-size awareness can substantially improve latency-sensitive communication, reducing short-flow completion times by up to 75% under perfect early flow classification, while preserving long-flow throughput.**

***Index Terms*—Data center networks, transport protocols, congestion control, flow-size prediction, AI workloads.**

I. Introduction

Modern data center networks must simultaneously support low latency for short transfers and high throughput for bandwidth-intensive flows. This objective is particularly important in AI data centers, i.e., data center infrastructures specifically designed to support large-scale distributed learning and large language model (LLM) workloads, where training jobs repeatedly exchange gradients, parameters, partitions, and synchronization messages across thousands of accelerators and servers. Communication delays in these exchanges can propagate to training iteration time, making transport-layer inefficiencies increasingly costly as clusters scale.

AI workloads differ from traditional web and microservice traffic. They often exhibit recurring computation-communication phases, synchronized exchanges, and persistent worker-to-worker or worker-to-server communication patterns. This regularity creates an opportunity: rather than treating flow size as hidden information that can only be inferred after observing transmitted bytes, the transport layer can be assisted by early predictions derived from workload and system context.

Existing data center network (DCN) transport protocols [1] already exploit the distinction between short and long flows. DCTCP [2] uses ECN feedback to maintain shallow queues, while flow-aware variants such as FA-DCTCP [3] apply different congestion-window reductions depending on whether a flow is considered short or long. However, this classification remains reactive: flows are initially treated as short by default, and are only reclassified as long once the amount of transmitted data exceeds a predefined byte threshold. This delay is problematic in high-speed DCNs, where many short flows complete within a few RTTs and where incorrectly treating early long-flow bytes as short-flow traffic can increase queueing pressure.

To advance the design of transport protocols for AI data centers, we propose *ATLAS*, an AI-assisted transport framework for AI workloads in DCNs that leverages workload-aware flow-size prediction to enable proactive and flow-aware transport-layer decisions. ATLAS is based on a simple hypothesis: if AI workloads generate predictable communication behavior, then transport protocols should use predictions as first-class inputs to congestion-control and flow-prioritization decisions. ATLAS does not replace congestion control with a black-box learner. Instead, it keeps the transport mechanism interpretable and deployable, while using learning to expose early flow characteristics that conventional mechanisms obtain too late.

II. ATLAS: Prediction-Assisted Transport

Figure 1 shows the ATLAS closed control loop. The framework monitors AI workload communication, builds per-flow feature representations, predicts transport-relevant flow characteristics – in particular flow size, and feeds these predictions into a flow-aware transport mechanism. The resulting behavior is then observed again, enabling future refinement of the prediction pipeline. Next, we describe each of these stages.

A. Workload-aware traffic visibility

ATLAS starts from the observation that useful flow-size signals may exist before or at flow initiation. These signals include temporal context from recent flows, host-level state, memory and disk activity, network counters, machine identity, and communication calls associated with the workload runtime. In practical deployments, such signals may come from a combination of host telemetry, in-network telemetry, and programmable devices such as SmartNICs [4]. The key requirement is not exhaustive packet capture, but per-flow visibility rich enough to support early transport decisions.

978-3-903176-83-6 ©2026 IFIP

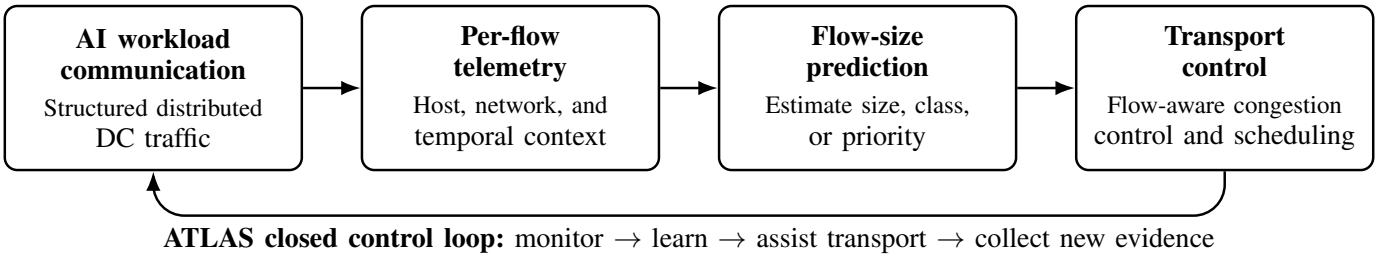

Fig. 1: ATLAS closed control loop. The framework monitors AI workload communication, builds per-flow feature representations, predicts flow sizes, and feeds these predictions into flow-aware transport control. The resulting behavior is observed again, enabling future refinement of the prediction pipeline.

This also motivates a DCN-specific definition of a flow. Distributed AI frameworks often reuse persistent TCP connections, so application messages are not always aligned with TCP connection lifetimes. ATLAS therefore treats a flow as a burst of packets between hosts, separated from other bursts by an inter-arrival gap. This definition better captures the units whose completion time matters for transport optimization.

B. Cross-workload flow-size prediction

Recent work has explored the use of machine learning for flow-size estimation in data center networks [5], [6]. However, these efforts largely formulate the problem from a per-workload perspective, training and deploying a separate predictor for each individual application or distributed framework. While effective in controlled settings, workload-specific predictors present important deployment limitations for operational AI data centers. Production AI clusters execute heterogeneous and continuously evolving workloads, often combining multiple distributed frameworks with partially overlapping feature spaces. Maintaining, retraining, and integrating one predictor per workload therefore becomes operationally complex and difficult to scale. ATLAS instead favors a unified cross-workload predictor capable of learning transport-relevant communication patterns shared across related AI workloads.

We formulate flow-size prediction as a supervised regression problem over current-flow telemetry and a temporal window of previous flows. The feature space combines host-level, network-level, and temporal information describing the execution context in which flows are generated. Since workloads may expose only partially overlapping features, the model relies on feature alignment, explicit handling of missing features, and shared representations across heterogeneous datasets. Gradient-boosted decision trees are particularly attractive in this setting because they combine strong predictive performance on heterogeneous tabular telemetry with low inference latency and implementation simplicity.

To evaluate this approach, we use the public dataset introduced in [5], which contains traces from four distributed AI applications: Spark PageRank, Spark K-Means, Spark SGD, and distributed TensorFlow training. The traces include detailed system- and flow-level instrumentation, enabling the study of flow-size predictability from host-side observations. Figure 2 illustrates the flow-size distributions for representative workloads. Although all workloads exhibit the heavy-tailed

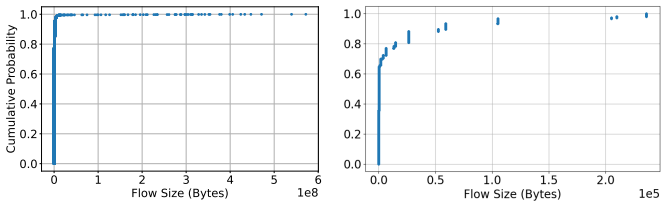

(a) Spark Page Rank. (b) TensorFlow.

Fig. 2: Flow size CDF for distributed AI workloads.

structure characteristic of data center traffic, their distributions differ substantially in scale and communication behavior. For example, in Spark PageRank, approximately 80% of flows are smaller than 100 KB, whereas in TensorFlow the same fraction falls below 10 KB. These differences highlight the challenge of designing workload-specific predictors and motivate the need for unified models capable of generalizing across heterogeneous AI communication patterns.

C. From reactive to proactive flow awareness

ATLAS uses predictions as early transport-layer signals. In a DCTCP-style implementation, the predicted flow size can be mapped to a hard short/long flow class or to a soft confidence score. The former preserves the structure of existing flow-aware protocols, while the latter allows the congestion response to interpolate between standard DCTCP behavior and short-flow treatment according to prediction confidence. The learning component is therefore not a replacement for congestion control, but an input that makes existing flow-aware mechanisms more effective.

III. Preliminary Results

A. Prediction Results

Table I summarizes representative preliminary results comparing workload-specific predictors [5], [6] against the ATLAS unified cross-workload model. Overall, the unified predictor preserves most of the predictive quality achieved by specialized per-workload models. The reduction in predictive performance remains limited across workloads, and in the case of Spark KMeans the unified model slightly improves the reported R^2 score. These results suggest that a single predictor can successfully capture common transport-relevant structure across heterogeneous AI workloads while avoiding the deployment complexity of per application models.

TABLE I: Representative cross-workload flow-size prediction results.

Workload	Individual R^2	Unified R^2	ΔR^2
Spark KMeans	0.884	0.888	+0.004
Spark PageRank	0.855	0.795	-0.061
Spark SGD	0.600	0.560	-0.040
TensorFlow	0.970	0.963	-0.007

Importantly, predictive accuracy alone is not sufficient to assess the usefulness of flow-size estimation for AI-assisted transport. The key objective is not the prediction itself, but rather its impact on transport-layer behavior and overall network performance. In practice, even predictors with non-negligible byte-level error may still provide substantial benefits if they expose sufficiently accurate early flow-size information to guide congestion-control and scheduling decisions. This motivates evaluating prediction mechanisms through their transport-level effects on metrics such as flow completion time (FCT), queue occupancy, and throughput efficiency.

B. Transport-Level Gains

To quantify the transport-level value of early flow-size awareness, we evaluate prediction-guided flow-aware transport in ns-3 using real distributed AI training traces [5] and representative DCN traffic patterns. The evaluation compares three behaviors: standard DCTCP, which performs no flow differentiation; reactive flow-aware DCTCP (FA-DCTCP) [3], where flows are initially treated as short and only later reclassified after exceeding a byte threshold; and early size-informed variants, where flow classification is available before transport decisions are taken, by some prediction mechanism as the one used by ATLAS.

To isolate the intrinsic benefit of early flow-size knowledge, we consider an oracle-assisted variant in which the exact flow size is exposed at flow creation time – i.e., a perfect predictor. This oracle design is not intended as a deployable mechanism; instead, it defines an upper bound on the gains achievable by ATLAS and allows us to quantify the impact of delayed reactive classification.

Figure 3 reports representative short-flow average FCT (AFCT) reductions achieved by the oracle-assisted transport under different link speeds. The results show that flow classification has a major impact on transport performance. At 100 Mbps, early flow awareness reduces short-flow AFCT by 45% relative to FA-DCTCP and by 41.5% relative to standard DCTCP. At 1 Gbps, the corresponding gains are 34.3% and 16%, respectively. The largest improvements appear at 10 Gbps, where short flows complete particularly quickly and delayed classification becomes increasingly costly: early flow awareness reduces short-flow AFCT by 74.4% relative to FA-DCTCP and by 27.8% relative to DCTCP. While not reported here, results also show that the AFCT reductions achieved for short flows are not obtained by systematically sacrificing long-flow throughput. In fact, throughput remains largely stable across configurations, and in several cases even improves relative to both FA-DCTCP and standard DCTCP.

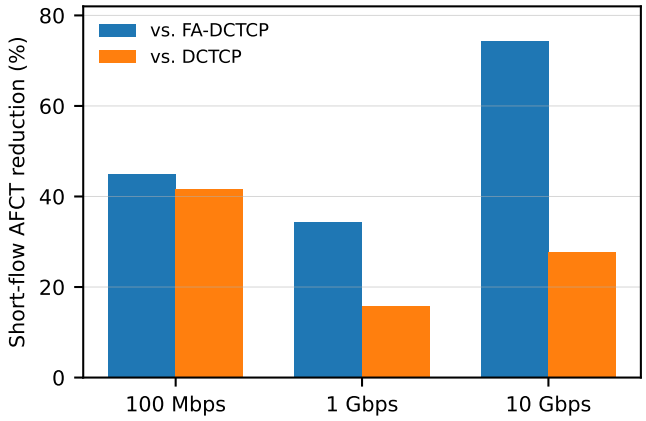

Fig. 3: Short-flow AFCT reduction obtained by exposing exact flow-size information early. The oracle result estimates the headroom available to ATLAS-style prediction-assisted transport.

These preliminary results suggest that delayed threshold-based classification is not merely an implementation detail, but a fundamental limitation of reactive flow-aware transport. In practice, realistic predictors such as the one used by ATLAS would operate between the two extremes represented by standard DCTCP and oracle-assisted transport. The oracle results therefore establish the performance headroom that prediction-guided mechanisms such as ATLAS can potentially recover through early workload-aware flow-size estimation, while preserving the throughput behavior of long flows.

IV. Discussion and Concluding Remarks

ATLAS proposes a broader direction for AI-native DCN transport: instead of replacing congestion control policies, machine learning can provide early workload-aware signals that assist existing transport mechanisms. In this setting, learning is used in a structured and narrowly scoped manner, preserving the robustness and deployability of established DCN transport protocols.

A central challenge moving forward is systems integration. Prediction mechanisms must operate within tight DCN timescales, requiring low-latency inference and scalable telemetry collection. Tree-based predictors are particularly attractive in this regard, since they can be efficiently implemented inside the kernel [7] and potentially offloaded into programmable data planes and SmartNICs using P4-based frameworks [8], [9].

Another key challenge is understanding how prediction errors affect transport behavior. In particular, misclassifying long flows as short can have a disproportionate network-wide impact, increasing queue occupancy and harming competing short flows. This suggests that future predictors should optimize not only aggregate prediction accuracy, but also the transport-level cost of different error types.

More broadly, AI workloads evolve rapidly over time, communication patterns may shift across training phases, and telemetry infrastructure must remain scalable and deployable in production environments.

ACKNOWLEDGMENTS

This work has been supported by the FWF Austrian Science Fund, Project reference I-6653, as part of the EU CHIST-ERA-2022-SPiDDS-02 project GRAPHS4SEC. Leonardo Alberro is also supported by the Comisión Académica de Posgrado (CAP) of the Universidad de la República (Udelar) through a Graduate Fellowship 2023.

REFERENCES

[1] L. Alberro and E. Grampin, "Transport Protocols in the Data Center: Understanding Proposals and Research Challenges," *Computer Networks*, vol. 274, p. 111793, 2026.

[2] M. Alizadeh *et al.*, "Data Center TCP (DCTCP)," in *Proceedings of the ACM SIGCOMM 2010 Conference*, ser. SIGCOMM '10. New York, NY, USA: Association for Computing Machinery, 2010, p. 63–74.

[3] S. Joy and A. Nayak, "Improving Flow Completion Time for Short Flows in Datacenter Networks," in *2015 IFIP/IEEE International Symposium on Integrated Network Management (IM)*, 2015, pp. 700–705.

[4] L. Alberro *et al.*, "INTFusion: Unifying Network and Host Telemetry in Data Center Networks," in *Proceedings of IFIP Networking 2026*, Lugano, Switzerland, May 2026, To appear.

[5] V. Dukic *et al.*, "Is Advance Knowledge of Flow Sizes a Plausible Assumption?" in *Proceedings of the 16th USENIX Conference on Networked Systems Design and Implementation*, ser. NSDI'19. USA: USENIX Association, 2019, p. 565–580.

[6] S. M. Hosseini, S. Sadrhaghighi, and M. Ghaderi, "Flow Size Prediction with Short Time Gaps," in *IEEE INFOCOM 2024 (INFOCOM WKSHPS)*, 2024, pp. 1–7.

[7] DMLC Contributors, "Treelite: Universal Model Exchange and Serialization Format for Decision Tree Forests," https://github.com/dmlc/treelite, 2026, accessed: 2026-04-28.

[8] C. Zheng, M. Zang, X. Hong, L. Perreault, R. Bensoussane, S. Vargaftik, Y. Ben-Itzhak, and N. Zilberman, "Planter: Rapid Prototyping of In-Network Machine Learning Inference," *SIGCOMM Comput. Commun. Rev.*, vol. 54, no. 1, p. 2–21, Aug. 2024. [Online]. Available: https://doi.org/10.1145/3687230.3687232

[9] B. Brandino, P. Casas, and E. Grampín, "Detecting Attacks at Switching Speed: Ai/Ml and Active Learning for in-Network Monitoring in Data Planes," in *2024 IEEE 32nd International Conference on Network Protocols (ICNP)*, 2024, pp. 1–6.