# Research Statement

## *Vamsi Addanki*

My research bridges the domains of networked systems design and applied theory, combining practical, hands-on implementation with rigorous formal methods to tackle today's most pressing network challenges, particularly in datacenter networks. As a *networking* researcher, my goal is to build performant, reliable, scalable and resource-efficient networks. This is a surprisingly difficult task hindered by rigid fixed-function network devices, static physical interconnects and uncertainty of the underlying network traffic patterns. Over the past decades, these inefficiencies in network design were masked by the free scaling of compute, storage and bandwidth, governed by Moore's law. However, operators and hardware manufacturers witnessed a slowdown in Moore's law scaling in the recent years i.e., the network bandwidth and buffers are unable to keep with the explosive traffic demands generated by modern datacenter applications. The combination of strong systems and theory background allows me to investigate emerging technologies that are paving the way for a paradigm shift in how networks are designed and operated.

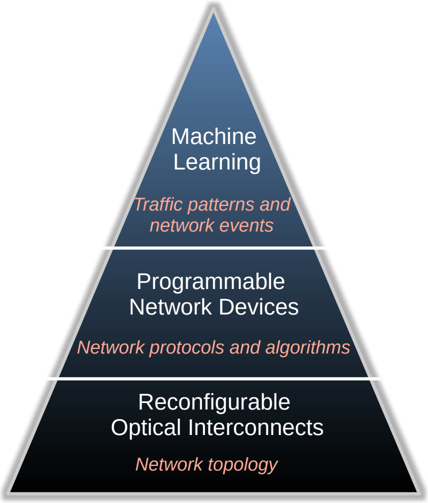

Figure 1: Three key enablers for a self-adjusting network design paradigm to support the rapid evolution of datacenter applications and to tackle the consequences of Moore's law slowdown.

In recent years, three emerging technologies have unlocked a wealth of opportunities to advance network design, performance, and scalability in datacenter networks. First, optical circuit-switching technologies alleviate the constraints of static network topologies, enabling more adaptive and responsive configurations. With reconfiguration possible at nanosecond scales, optical interconnects introduce new possibilities for dynamically aligning network topology with evolving traffic patterns. Second, programmable network devices overcome the limitations of fixed-function dataplanes, allowing for fine-tuned protocol deployment at scale, tailored to the unique demands of datacenter applications — all without the expense of additional hardware. Third, predictive models, especially transformers, have revolutionized the landscape of artificial intelligence. When combined with the programmability of network hardware, these prediction models open up a range of possibilities to build prediction-augmented systems that can automatically adapt to evolving traffic patterns in a fine-grained manner, without requiring manual parameter tuning by network operators.

To effectively harness the potential of these technologies, my research leverages theoretical concepts to build practical systems that can operate at scale while offering provable performance guarantees. By combining insights from multiple disciplines — networking (e.g., protocols, topologies), operating systems (e.g., Linux Kernel, eBPF), theory (e.g., graph theory, online algorithms), and other domains (e.g., optics, energy markets) — I adopt an interdisciplinary approach that addresses the multifaceted challenges of modern network infrastructure. This approach is essential for tackling the diverse workloads, stringent performance requirements, and need for energy efficiency that define today's networks. What sets my work apart is a strong focus on the co-design of network protocols and algorithms alongside the underlying infrastructure, ensuring that systems are not only theoretically sound but also practically deployable.

In my PhD dissertation, I addressed a range of challenges in modern datacenters, including congestion control, load-balancing, buffer-sharing, and optical interconnects. These contributions were published *(as a first author)* at top-tier networking conferences, including **SIGCOMM [1], NSDI [6, 4, 9], SIGMETRICS [3], and INFOCOM [8]**.

§1 I developed foundational techniques and introduced the first formal framework for analyzing the bandwidth of optical interconnects [3]. Building on these insights, I designed a *practical* high-bandwidth optical interconnect that adapts dynamically to network traffic patterns [2] — marking the first self-adjusting optical interconnect with formal bandwidth guarantees that account for hardware limitations, such as reconfiguration delays [7].

§2 In response to shrinking buffer sizes in datacenter network devices, I developed buffer-sharing algorithms that adapt to traffic patterns to improve throughput while providing *provable* guarantees [1, 4]. Further, I introduced the first learning-augmented buffer-sharing algorithm with formal performance bounds, independent of prediction accuracy [9] — advancing learning-augmented networked systems with explainable performance.

§3 My work on low-latency datacenter networks rearchitects datacenter transport [6, 5], challenging conventional approaches and demonstrating that inherent traffic patterns can be leveraged to optimize the network performance.

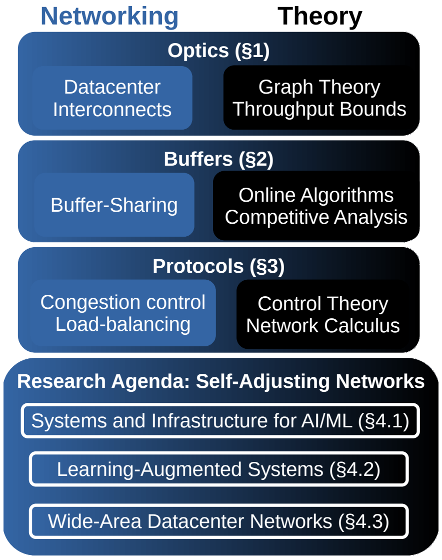

Figure 2: Overview of my prior work (§1, §2, §3) and future research agenda (§4).

In addition to developing innovative solutions for networked systems, I am deeply committed to ensuring that my research is reproducible and extensible. I believe that making research outputs accessible fosters future advancements and collaboration. To this end, all my publications are accompanied by open-source code. Notably, the GitHub repository hosting the source code for several of my works on datacenter networks has been utilized by numerous research groups and has garnered over a hundred stars from GitHub users, reflecting its impact in the research community.

§4 **My vision is to develop a self-adjusting datacenter network design.** Reconfigurable optical interconnects, programmable network dataplanes, and prediction models form a powerful combination for creating networks where both the protocols and the underlying infrastructure can dynamically adjust to the evolving demands of diverse applications. This vision is especially relevant in the context of emerging GPU clusters for AI/ML workloads, where traffic patterns, while structured and predictable, can vary significantly across different training jobs. Such variability highlights the need for self-adjusting networks that can optimize performance and efficiency on-demand, in response to changing workload requirements.

A self-adjusting network enhances the scalability and performance of datacenters while laying the groundwork for new paradigms in network management. By dynamically aligning resources to workload demands, such networks reduce operational overhead and promote sustainable resource use. Realizing this vision will empower datacenters to meet the demands of tomorrow's complex workloads — from AI/ML to next-generation cloud applications — all while advancing the imperative for energy-efficient design.

The rest of the statement is organized as follows: §1, §2, §3 summarize my prior work and §4 outlines my research agenda.

## 1 High-Bandwidth Optical Interconnects for Datacenters

The network infrastructure is no longer able to keep up with the exponential growth in the volume of data generated by modern applications. The traditional network design based on electrical packet-switched interconnects are reaching their limits in terms of bandwidth, latency, and energy efficiency. Optical circuit-switched interconnects have emerged as a promising solution to address these challenges. However, unlike static networks, a multitude of factors such as the reconfiguration delay, the granularity of switching, and the topology of the network play a crucial role in determining the performance of optical interconnects. Three fundamental questions remained unanswered in the literature:

***(Q1)*** *What is the limit on the throughput achievable by optical interconnects?*

***(Q2)*** *How can we design high-throughput optical interconnects that can adapt dynamically to traffic patterns while providing formal throughput guarantees?*

***(Q3)*** *Can optical interconnects provide high throughput even under demand uncertainty?*

These questions are surprisingly difficult and require non-trivial methodologies to answer them. Multiple decades of research has focused on understanding the throughput of *static* networks. e.g., multicommodity flow problems and its variants. However, optical interconnects are typically dynamic in nature. Throughput of such networks remained largely unexplored in the literature and lacks the necessary analytical tools. In my dissertation, I answered all the above questions.

**Mars** [3] is an optical interconnect design motivated by the pursuit of finding the throughput bounds of such networks. I developed fundamental techniques to analyze the throughput of optical interconnects. I established a relation between the throughput of a periodic circuit-switched optical interconnect and that of a directed static graph, allowing the analysis of the throughput of optical interconnects using standard graph-theoretic techniques. This relationship and subsequent analysis revealed the inherent trade-offs between throughput, latency, and buffer requirements in networks built with optical interconnects. Given these tradeoffs, I designed Mars, a periodic circuit-switched optical interconnect that achieves near-optimal throughput while minimizing latency and buffer requirements.

**Vermilion** [2] is a Self-Adjusting optical interconnect that dynamically adapts to the underlying traffic patterns. Importantly, I proved the first formal throughput guarantees for a self-adjusting (demand-aware) optical interconnect that breaks the throughput barrier of periodic circuit-switching. To prove this result, I developed a novel martix-nearness measure by leveraging insights from matrix rounding — a technique developed in the 80's at the intersection of mathematics and social sciences for aggregating large survey data while retaining statistical properties — highlighting my interdisciplinary approach to solving networked systems problems.

**Vermilion, Pt. 2** [7] demonstrates that a self-adjusting optical interconnect can provide high throughput not only under perfect knowledge of the communication patterns (consistency) but also provides formal guarantees even under demand uncertainty (robustness). Such algorithmic design principles are increasingly more critical in the context of learning-augmented systems that achieve *explainable* performance.

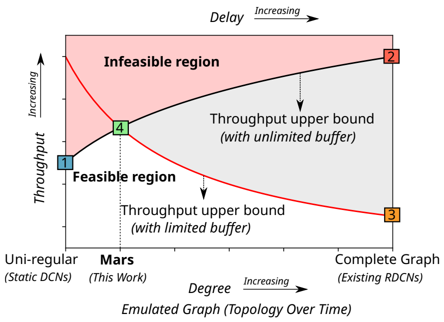

**ACM SIGMETRICS 2023**

[3] Vamsi Addanki, Chen Avin, and Stefan Schmid. *"Mars: Near-Optimal Throughput with Shallow Buffers in Reconfigurable Datacenter Networks"*

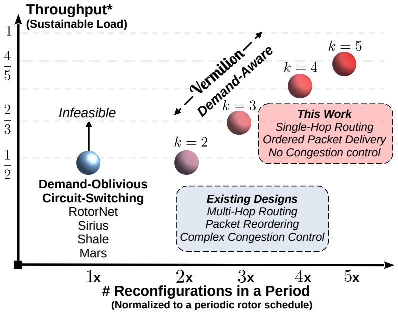

**Under Submission**

[2] Vamsi Addanki, Chen Avin, Manya Ghobadi, Goran Dario Knabe, and Stefan Schmid. *"Vermilion: Breaking the Throughput Barrier of Periodic Circuit-Switching (A Simple Demand-Aware Optical Interconnect)"*

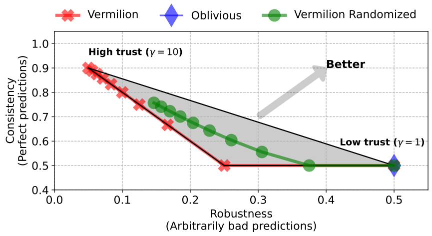

**Under Submission**

[7] Vamsi Addanki, Maciej Pacut, Leon Kellerhals, Goran Dario Knabe, and Stefan Schmid. *"Vermilion, Pt. 2: Augmenting Reconfigurable Optical Network Design with Machine-Learned Predictions"*

## ② Buffer Sharing in Post Moore's Law Era

Buffers are central to the performance of networked systems, absorbing transient bursts in traffic and minimizing packet losses. Network devices enjoyed the free scaling from Moore's law for decades, allowing the use of large buffers to absorb bursts. However, Moore's law scaling has been progressively slowing down in the recent years and switch vendors are unable to scale buffer sizes proportional to bandwidth. The shrinking buffer sizes in modern network devices poses a significant challenge to network performance. This trend posed several fundamental questions to the research community:

***(Q1)*** *Can shallow buffers achieve high burst absorption while maintaining high throughput?*

***(Q2)*** *Can different traffic types be isolated in the network?*

***(Q3)*** *To what extent can buffer-sharing algorithms be improved within the hardware capabilities, while providing formal throughput guarantees?*

My dissertation addressed these challenges using novel techniques, providing practical solutions along with formal guarantees.

**ABM** [1] represents a significant breakthrough in buffer-sharing by fundamentally rethinking how buffers accommodate the diverse demands of network traffic. It introduces an innovative buffer-sharing technique that effectively combines the strengths of conventional Buffer Management and Active Queue Management approaches. In designing ABM, I conducted a thorough analysis of state-of-the-art algorithms using network calculus, yielding valuable insights into their limitations. I established formal guarantees for ABM, such as high burst absorption and bounded latency, which position it as a robust solution for modern networks. Notably, ABM is the first algorithm capable of effectively isolating different traffic types, offering a practical solution to buffer sharing in the post-Moore's law era.

**Reverie** [4] is inspired by the intriguing paradoxes that arise in buffer-sharing scenarios within datacenters. As a systems researcher with theory background, I took these paradoxes as a challenge and conducted a first formal analysis of SONiC, the state-of-the-art buffer sharing technique in datacenters. My analysis revealed concrete answers, pin-pointing the root causes to these issues. With valuable feedback from major switch vendors and operators, I designed Reverie, a buffer-sharing mechanism that isolates RDMA and TCP traffic within the network buffers, while improving throughput and burst absorption for both traffic types.

**Credence** [9] is a learning-augmented buffer-sharing system with formal throughput guarantees, offering explainable performance even while treating predictions as a blackbox. After reaching the limits of improving throughput guarantees with conventional drop-tail buffer-sharing algorithms, I shifted focus toward prediction-based methods. Unlike traditional approaches that rely on perfect prediction accuracy, I designed Credence to provide formal guarantees even under highly inaccurate predictions. Practical within today's hardware constraints, Credence surpasses the fundamental throughput limitations of drop-tail algorithms, marking a substantial advancement in buffer-sharing techniques.

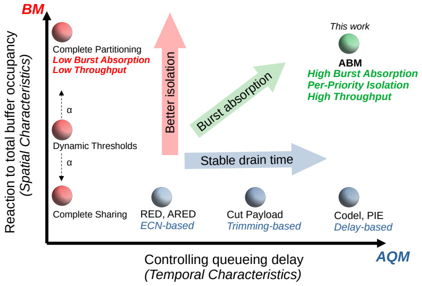

**ACM SIGCOMM 2022**

[1] Vamsi Addanki, Maria Apostolaki, Manya Ghobadi, Stefan Schmid, and Laurent Vanbever. *"ABM: Active Buffer Management in Datacenters"*

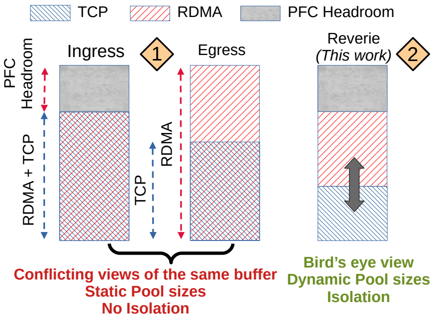

**USENIX NSDI 2024**

[4] Vamsi Addanki, Wei Bai, Stefan Schmid, and Maria Apostolaki. *"Reverie: Low Pass Filter-Based Switch Buffer Sharing for Datacenters with RDMA and TCP Traffic"*

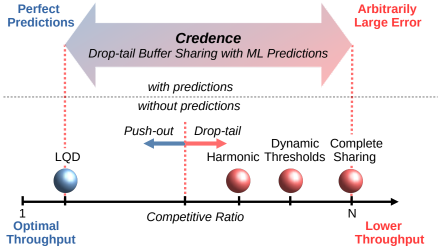

**USENIX NSDI 2024**

[9] Vamsi Addanki, Maciej Pacut, and Stefan Schmid. *"Credence: Augmenting Datacenter Switch Buffer Sharing with ML Predictions"*

## 3 Low Latency Datacenter Networks

Latency and completion times for the transmissions in a datacenter are of paramount importance, that impacts the performance of various applications critically. Decades of research has focused on achieving this goal using various techniques such as congestion control and load-balancing algorithms. Yet, it remains a surprisingly difficult task to achieve low latency while maintaining high throughput, especially in emerging datacenters e.g., RDMA-based GPU clusters with AI/ML workloads.

***(Q1)*** *Does rapid reaction to congestion fundamentally entail throughput loss?*

***(Q2)*** *Can optimal load-balancing be achieved using single-path transport?*

My research addresses these core challenges in networking, advancing beyond conventional approaches to congestion control and load balancing. I introduce novel techniques that are both *simple* and *practical*, backed by *provable* performance guarantees.

**PowerTCP** [6] is a congestion control algorithm specifically designed to meet the ultra-low latency requirements of modern datacenter networks. PowerTCP's design emerged from analyzing the current landscape of congestion control algorithms, which typically operate along a single dimension: either "voltage" (inflight data) or "current" (transmission rate). Leveraging this observation and the classic bottleneck link model, I established a relationship between the sender's transmission window size and the product of "voltage" and "current" — a metric we define as "power" in the networking context, revealing an interesting connection to electrical systems from a control-theoretic perspective. Building on these insights, PowerTCP is designed to react quickly to congestion while maintaining high throughput during convergence. I proved several theoretical properties that underscore PowerTCP's relevance to modern datacenters and implemented a working prototype in the Linux Kernel, demonstrating its practicality.

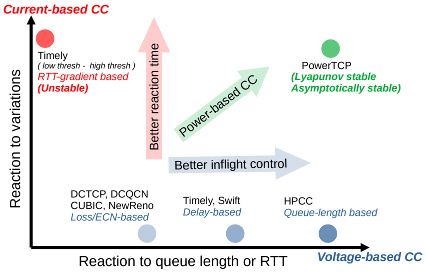

**USENIX NSDI 2022**

[6] Vamsi Addanki, Oliver Michel, and Stefan Schmid. *"PowerTCP: Pushing the Performance Limits of Datacenter Networks"*

**Ethereal** [5] is a transport protocol designed specifically for large-scale distributed training in GPU clusters. Previous approaches, including efforts from hyperscale datacenter operators, attempted to improve load-balancing using single-path transport semantics within existing hardware. However, most eventually abandoned this approach, with major vendors joining forces to develop hardware that supports "packet spraying" — a form of multipath transport. Challenging the prevailing assumption in both academia and industry, I formally proved that optimal load-balancing can still be achieved with single-path transport in the context of AI/ML workloads, presenting a practical alternative to load-balancing that requires no new hardware. Even in environments with multipath-capable hardware, Ethereal provides a more scalable and resource-efficient solution. Designing Ethereal required a deep understanding of the underlying network topology, combined with the characteristic of AI/ML workloads, which revealed interesting properties that can be leveraged to achieve optimal load-balancing. This work underscores my passion for tackling real-world challenges in networked systems and reflects my approach to problem-solving, which is firmly rooted in theoretical foundations while ensuring practicality.

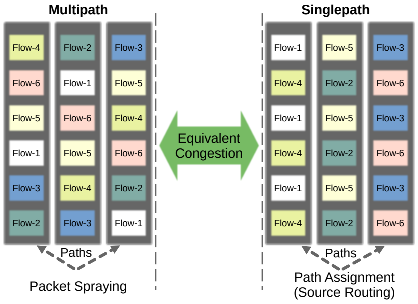

**Under Submission**

[5] Vamsi Addanki, Prateesh Goyal, and Ilias Marinos. *"Challenging the Need for Packet Spraying in Large-Scale Distributed Training"*

# ④ Research Agenda: Self-Adjusting Networks

The massive scale of investment and the economic value of the datacenter as well as the wide-area-network infrastructure (even comparable to many developed nation GDPs) demands a paradigm shift in the design of protocols and algorithms. The traditional best-effort paradigm is no longer sufficient to meet the stringent requirements of the trillion-dollar network infrastructure, especially in the case of dedicated GPU clusters. Traditional datacenters with general-purpose processors served a variety of workloads often in a best-effort manner as the traffic patterns largely remained unpredictable. Much of the research efforts, ranging from congestion control, topologies, load balancing, routing, and so on, have been focused on designing protocols and algorithms that can handle this unpredictability, thus compromising on performance. Similarly, traditional wide-area-networks involve costly high-bandwidth long-distance cables that are often over-provisioned and traffic-engineering techniques are employed to handle the unpredictability of traffic patterns. However, first, the emerging GPU clusters are built for specific purposes, such as AI/ML training workloads that generate *predictable traffic patterns*, and demand a different approach; second, the emerging technologies such as optical circuit switches allow not just traffic-engineering but also topology-engineering that can be optimized for specific traffic patterns. The network infrastructure is evolving into a collection of specialized clusters, each optimized for specific workloads.

My research agenda involves architecting a self-adjusting network paradigm that tailors protocols and algorithms to the specific needs of the evolving network demands. In the following, I outline the key research directions that will drive my future work.

## 4.1 Systems and Infrastructure for AI/ML Workloads

### *Leveraging Predictability (and Mutability) of Traffic Patterns to Design Network Protocols*

In contrast to traditional datacenters, large-scale GPU clusters exhibit predictable traffic patterns. Workloads are pre-determined, and the traffic patterns are explicit due to the nature of collective communication in distributed training. The traffic matrices generated by common collectives are mutable. For example, an allReduce operation results in a uniform traffic matrix when using an all-to-all algorithm, or a permutation matrix with a ring algorithm. This mutability can be exploited to optimize network performance. Recent research has proposed alternative collective communication algorithms that adapt the traffic matrix based on various objectives, often without considering the underlying topology or assuming ideal network conditions.

However, the state-of-the-art protocol suite, designed with a best-effort philosophy for commodity hardware, does not provide ideal network conditions; network performance is influenced by topology, routing, and congestion control. For instance, a decade of research on optimizing congestion control, load-balancing, scheduling, buffer-sharing, still leave surprising and unique problems under collective communication in GPU clusters: while some communication patterns (e.g., point-to-point) perform well under specific protocols, others suffer (e.g., all-to-all ) i.e., the best-effort philosophy of designing "one ring to rule them all" does not hold anymore.

My research aims to bridge this gap by co-optimizing collective communication algorithms with underlying network protocols, particularly routing, load-balancing, and congestion control, to maximize performance. My recent works are a first step in this direction. Credence [9] exploits the predictability of traffic patterns and improves throughput when multiple congestion hotspots contend for commonly shared buffer space. Ethereal [5] challenges the need for packet spraying in GPU clusters and shows that the properties of the underlying traffic matrices can be exploited to optimize network performance. In the future, I plan to pursue fundamental questions such as: How can we design network protocols that can exploit the mutability of traffic matrices? How can we co-optimize collective communication algorithms with underlying network protocols to maximize performance?

### *Self-Adjusting Optical Interconnects for GPU Clusters*

Emerging technologies using optical circuit switches enable network infrastructures to adapt their topology and capacity provisioning based on observed traffic patterns, marking a shift from traditional datacenter network designs. Conventional topologies like the fat-tree are built to offer high bisection bandwidth and low latency for unpredictable traffic patterns. In contrast, today's GPU clusters generate predictable and often structured traffic, which opens opportunities to optimize network topology and capacity provisioning for performance and cost-efficiency. In the context of the post-Moore's law era, developing high-performance optical interconnects is essential, supporting a transition toward optical computing.

My recent work [3, 2, 7] on reconfigurable networks is a foundational step in understanding the trade-offs between reconfigurability and performance. My research agenda now seeks to address key questions: How can we co-optimize dynamic network topologies and collective communication algorithms to maximize performance? What are the strategies for building failure-resilient optical interconnects? How can we achieve fully automated control plane management for optical interconnects, minimizing manual intervention?

### *Transport Protocol for GPU clusters*

My research agenda includes a special focus on the design of a high-performance, reliable, and practical transport protocol tailored specifically for GPU clusters, addressing the unique demands of AI/ML workloads. Congestion control and load balancing are increasingly critical in this domain, as underscored by the focus from hyperscale datacenter operators, including the Ultra Ethernet Consortium.

Building upon my prior work [5, 6], I aim to develop a transport protocol that can adapt dynamically to AI/ML traffic patterns while offering formal guarantees on both throughput and latency. I plan to collaborate with industry to gain insights into the real-world challenges and to validate the performance of the proposed transport protocol. To this end, I am already in discussions with collaborators from Microsoft Research and plan to foster strong, ongoing engagement between academia and industry. My vision is to design an optimal transport protocol for AI/ML workloads without necessitating costly new hardware, enabling high-performance solutions for both hyperscale datacenters and emerging HPC clusters at universities.

### *An Open Source Massively Parallel Network Simulator*

A network simulator is indispensable for identifying performance bottlenecks in large-scale networks and designing new protocols and algorithms. However, current open-source simulators like NS3 and OMNET++, though widely used, struggle with the scale required for today's GPU clusters. While these simulators technically support parallelization, achieving efficient multi-core simulations remains a challenging and time-intensive process for researchers.

In my previous work, I relied extensively on NS3 to evaluate datacenter network performance. Even for a modest datacenter topology with 256 servers running at 25 gigabit per second link speeds, an NS3 simulation could take days to complete. Given that today's GPU clusters operate at link speeds of up to 800 gigabit per second with thousands of GPUs, scaling these network simulations effectively is both impractical and largely inaccessible.

To address this gap, my goal is to develop a new, open-source, massively parallel network simulator that makes multi-core simulations easily accessible and user-friendly. By engaging both PhD and undergraduate students in its development, I aim to create a tool that empowers researchers to conduct large-scale network simulations, fostering collaboration and advancing reproducibility across the community.

### *Minimizing Energy Cost of GPU Clusters*

Large-scale training jobs in dedicated GPU clusters often exhibit a repetitive pattern of computation and communication. Notably, the training process can be paused, checkpointed, saved to storage, and resumed later without affecting accuracy. This flexibility in GPU workloads opens exciting opportunities to design algorithms that strategically pause and resume jobs, thereby optimizing energy costs and minimizing carbon emissions. This approach is inspired by the fact that both energy costs and carbon emissions fluctuate over time and location, often dropping when renewable energy sources are plentiful.

My ongoing work [10] draws an intriguing connection between energy cost minimization and the classic online $k$-min search algorithm, setting the stage for algorithmic trading in energy procurement. Looking ahead, I plan to collaborate with industry and economists to gain insights into the practices and challenges of real-world energy procurement in datacenters. My research agenda aims to design resilient, practical algorithms for sustainable datacenter operations that can adapt to market volatility and fluctuating availability of renewable energy sources.

## 4.2 Learning-Augmented Systems

The self-adjusting network paradigm demands a radical shift in the design of network protocols and algorithms, particularly those running directly on hardware, to be optimized for workload-specific needs. Traditional network devices were not built to support this level of adaptability, thus underscoring the importance of programmable network hardware. Fortunately, programmable switches have become a reality, supported by mature programming languages such as P4, as well as programmable optical switches that directly manipulate light pathways based on port mapping. This programmability, combined with predictable traffic patterns, unlocks a range of research directions that remain largely unexplored.

Learning-augmented algorithms are a hot-topic within the theory community, demonstrating promising results across various fields. These algorithms offer a hybrid approach: they provide near-optimal performance when predictions are accurate, while still ensuring robust performance even when predictions are far from perfect. This duality is particularly appealing for networked systems, where reliability and explainable performance are essential. However, integrating learning-augmented algorithms into networked systems remains nascent, with current models often requiring high prediction accuracy to deliver strong performance.

My research agenda seeks to advance this area by designing learning-augmented networked systems capable of delivering formal guarantees, not only under high prediction accuracy but also when faced with inaccuracies. My prior work on Credence [9] represents a foundational step, showing that learning-augmented buffer-sharing algorithms can maintain throughput guarantees despite imperfect predictions. Building on this, I aim to explore learning-augmented approaches across congestion control, buffer-sharing, packet classification, scheduling, and load balancing. My vision is to pioneer a new class of learning-augmented network algorithms that can be deployed within programmable network devices, creating self-adjusting networks that automatically adapt to evolving traffic patterns, thus eliminating the need for manual tuning and oversight.

## 4.3 Wide-Area Datacenter Networks

Recent advancements in deep learning have led to model sizes that exceed the capacity of a single datacenter, requiring multiple datacenters to work collaboratively to train a single model. Consequently, wide-area datacenter networks are becoming critical, demanding the development of new protocols and algorithms to optimize performance across geographically dispersed datacenters. The predictable traffic patterns of AI/ML workloads present an opportunity to design tailored protocols that adapt to their specific requirements. However, unlike intra-datacenter networks, wide-area datacenter networks are characterized by significant propagation delays, introducing a unique set of challenges that call for innovative solutions.

To address these challenges, my research agenda focuses on creating adaptive, latency-aware protocols that can overcome the high propagation delays inherent in wide-area datacenter networks. This involves leveraging traffic predictability in AI/ML workloads to design algorithms that: *(i)* optimize bandwidth allocation using reconfigurable network devices such as ROADMs and optical switches *(ii)* minimize latency by tailoring congestion control and load-balancing, and *(iii)* desynchronizing the training job to a large extent, reducing the frequency of inter-datacenter communication. By closely collaborating with industry, I aim to develop these solutions within real-world constraints, ensuring they are both practical and robust. Ultimately, this work will contribute to a cohesive, high-performance framework for wide-area datacenter networks, enabling seamless scaling of AI/ML applications and supporting the next generation of distributed, data-intensive workloads.

# 5 Conclusion

In summary, my research agenda is to architect a "Self-Adjusting" paradigm for network infrastructure and protocol design. Three key enablers for this vision are: *(i)* reconfigurable optical interconnects *(ii)* programmable network hardware and *(iii)* prediction models like transformers. In pursuit of this vision, my priority lies in designing specialized protocols and algorithms that are both *performant as well as implementable.*

## References

[1] **Vamsi Addanki**, Maria Apostolaki, Manya Ghobadi, Stefan Schmid, and Laurent Vanbever. “ABM: Active Buffer Management in Datacenters”. In: *Proceedings of the ACM SIGCOMM 2022 Conference*. SIGCOMM ’22. Amsterdam, Netherlands: Association for Computing Machinery, 2022, pp. 36–52. ISBN: 9781450394208. URL: `https://doi.org/10.1145/3544216.3544252`.

[2] **Vamsi Addanki**, Chen Avin, Manya Ghobadi, Goran Dario Knabe, and Stefan Schmid. “Vermilion: Breaking the Throughput Barrier of Periodic Circuit-Switching (A Simple Demand-Aware Optical Interconnect)”. In: *Under Submission* (2024). URL: `https://nextcloud.inet.tu-berlin.de/s/DanmSk4os8nB5pJ`.

[3] **Vamsi Addanki**, Chen Avin, and Stefan Schmid. “Mars: Near-Optimal Throughput with Shallow Buffers in Reconfigurable Datacenter Networks”. In: *Proc. ACM Meas. Anal. Comput. Syst.* 7.1 (Mar. 2023). URL: `https://doi.org/10.1145/3579312`.

[4] **Vamsi Addanki**, Wei Bai, Stefan Schmid, and Maria Apostolaki. “Reverie: Low Pass Filter-Based Switch Buffer Sharing for Datacenters with RDMA and TCP Traffic”. In: *21st USENIX Symposium on Networked Systems Design and Implementation (NSDI 24)*. Santa Clara, CA: USENIX Association, Apr. 2024, pp. 651–668. ISBN: 978-1-939133-39-7. URL: `https://www.usenix.org/conference/nsdi24/presentation/addanki-reverie`.

[5] **Vamsi Addanki**, Prateesh Goyal, and Ilias Marinos. “Challenging the Need for Packet Spraying in Large-Scale Distributed Training”. In: *CoRR* abs/2407.00550 (2024). arXiv: `2407.00550`. URL: `https://arxiv.org/abs/2407.00550`.

[6] **Vamsi Addanki**, Oliver Michel, and Stefan Schmid. “PowerTCP: Pushing the Performance Limits of Datacenter Networks”. In: *19th USENIX Symposium on Networked Systems Design and Implementation (NSDI 22)*. Renton, WA: USENIX Association, Apr. 2022, pp. 51–70. ISBN: 978-1-939133-27-4. URL: `https://www.usenix.org/conference/nsdi22/presentation/addanki`.

[7] **Vamsi Addanki**, Maciej Pacut, Leon Kellerhals, Goran Dario Knabe, and Stefan Schmid. “Vermilion, Pt. 2: Augmenting Reconfigurable Optical Network Design with Machine-Learned Predictions”. In: *Under Submission* (2024). URL: `https://nextcloud.inet.tu-berlin.de/s/HfMQJEm7XXmAcfb`.

[8] **Vamsi Addanki**, Maciej Pacut, Arash Pourdamghani, Gabor Rétvári, Stefan Schmid, and Juan Vanerio. “Self-Adjusting Partially Ordered Lists”. In: *IEEE INFOCOM 2023 - IEEE Conference on Computer Communications*. 2023, pp. 1–10. DOI: `10.1109/INFOCOM53939.2023.10228937`.

[9] **Vamsi Addanki**, Maciej Pacut, and Stefan Schmid. “Credence: Augmenting Datacenter Switch Buffer Sharing with ML Predictions”. In: *21st USENIX Symposium on Networked Systems Design and Implementation (NSDI 24)*. Santa Clara, CA: USENIX Association, Apr. 2024, pp. 613–634. ISBN: 978-1-939133-39-7. URL: `https://www.usenix.org/conference/nsdi24/presentation/addanki-credence`.

[10] Maciej Pacut, Prajna Bhat, **Vamsi Addanki**, Anshul Gandhi, and Stefan Schmid. “Algorithmic Trading in the Energy Derivatives Market: Applications for Energy Procurement in Large-Scale GPU Clusters”. In: *Work in Progress* (2024).