Wenfei Huang, Meike Hofmann, Peter Amthor, and Boris Koldehofe. "Energy-Aware and Flexible Optical Data Stream Processing: A Study of the Averaging Operator in Free Space", to appear in Proceedings of the 15th International Symposium on Communication Systems, Networks and Digital Signal Processing 2026 (CSNDSP 2026), July 15-17, 2026, Edinburgh, Scotland UK.

Energy-Aware and Flexible Optical Data Stream Processing: A Study of the Averaging Operator in Free Space

Wenfei Huang, Meike Hofmann, Peter Amthor, Boris Koldehofe
Technische Universität Ilmenau
98693 Ilmenau, Germany
[first.last]@tu-ilmenau.de

***Abstract*—Infinite data streams in modern applications place great demands on data computation and movement resources in conventional digital infrastructure, resulting in a trend of steeply increasing energy consumption. Optical computing provides an alternative by performing computation during light propagation, potentially alleviating the energy overhead inevitable with digital electronics. In this study, we make a first step towards this goal by implementing the averaging operator, a core primitive of stream processing, based on free-space optical computing. Our results demonstrate a high level of accuracy (relative error of approximately 2%) across various input types. Regarding energy consumption, we found that the optical implementation scales extremely well and is more efficient than digital computation for processing high-dimensional data. This hints at a practical methodology for mapping various stream aggregation operations onto optical hardware.**

***Index Terms*—optical computing, stream processing, aggregation, average operation, free-space optical computing**

I. Introduction

Current data-driven applications, such as real-time analytics and online recommendation platforms, are generating a sheer number of continuous data streams [1]–[3]. As data volumes keep growing, the resources required to process data – including separate stages like transmission, computation and storage – are constantly expanding, imposing tremendous pressure on data centers. To cope with the low-latency and high-throughput requirements of such workloads, stream processing has evolved as a fundamental computing model for real-time data processing. It also facilitates better data management, storage and downstream analysis [4]. Further, in-network computing has been adopted to relocate parts of the computations from data centers into the networking fabrics, reducing the need for data movement and mitigating bottlenecks associated with centralized data processing [5].

Despite these advances, the underlying infrastructure remains dominated by digital electronic devices. Therefore, supporting the sheer volume of data necessitates a larger number of data centers. Globally, the number of active data centers has increased from about $8\,000$ in 2021 [6] to $12\,000$ in 2025 [7]. Together with the impact of more powerful hardware in data centers, it accounts for approximately $1.5\,\%$ of worldwide electricity consumption, which is expected to more than double by 2030 [8]. Also, the cooling costs of data centers are far from negligible [9].

To counter these energy efficiency issues, optical computing has emerged as a promising alternative to digital electronics. Besides the energy-efficient computation that optics promises, research in optical physics indicates that photons can naturally perform computations whilst propagating at the speed of light [10]. This property co-locates data movement and computation within the physical medium itself, thereby significantly reducing the burden of explicit data transfer between separate memory and processing units. By contrast, in current data centers, data must be frequently transferred between computing units (e.g., central processing units (CPUs), graphics processing units (GPUs), neural processing units (NPUs), etc.) and communication fabrics (e.g., switches and routers), resulting in substantial overhead for data movement and excessive energy consumption [11].

Given the above promising capability, existing research on optical computing mostly concentrates on neural network (NN)-centric primitives [12] such as matrix multiplication and nonlinear activation functions, while other general-purpose data operations, such as filtering and aggregation in stream processing, have not received much attention. Specifically, it remains an open question whether stream operations can be accurately and efficiently achieved on a physical optical system. Moreover, most studies rely on simulation software to model optical systems, failing to capture realistic hardware imperfections, environmental noise, and system-level overheads, making the accuracy of computation and the measurement of energy consumption unreliable.

This work intends to investigate how optical computing can be applied to data processing. For this, we evaluate the feasibility and energy efficiency of a widespread stream processing operation, namely aggregation. To optically implement the aggregation and characterize energy consumption, the choice of the optical system is critical. Waveguide-based optics offers a compact system and precise outputs, but it relies on circuits fabricated in a fine-grained manner, resulting in limited reconfigurability after fabrication, analogous to application-specific integrated circuits (ASICs). Free-space

This work has been partially funded by the Carl Zeiss Foundation in the "Ilmenau School of Green Electronics" (CZS: P2022-00-135, P10: "Optical Computing in Data Processing Applications").

Wenfei Huang, Meike Hofmann, Peter Amthor, and Boris Koldehofe. “Energy-Aware and Flexible Optical Data Stream Processing: A Study of the Averaging Operator in Free Space”, to appear in Proceedings of the 15th International Symposium on Communication Systems, Networks and Digital Signal Processing 2026 (CSNDSP 2026), July 15-17, 2026, Edinburgh, Scotland UK.

optical computing, resembling field-programmable gate array (FPGA), allows computations to be flexibly reconfigured at runtime through programmable spatial modulators and an open physical layout, which precisely addresses the need for operator reconfiguration in stream processing—adapting to dynamic data processing.

In this paper, we present an in-depth study of a stream aggregation operator as a potential candidate for optical computing, taking the averaging operation as a concrete and representative case study. We experiment with a free-space optical system featuring an spatial light modulator (SLM) display. We analyze the averaging results and measure the end-to-end energy usage of the entire optical system. The digital implementation of an equivalent computation serves as a baseline to evaluate the accuracy and energy cost of optical averaging. Throughout the comparative analysis, our work aims to deepen the understanding of the energy benefits that optics can bring and paves the way for the development of future optical stream processing applications.

The rest of the paper is organized as follows. Section II introduces the background of stream processing and optical computing. Section III presents our approach to conceptually mapping the stream processing pipeline to an optical implementation channel of the averaging operator and energy models for both the optical and digital averaging. In section IV, we demonstrate our system, data, and evaluation. Section V reviews related work, and section VI concludes the paper.

II. Background

We first present the background and terminology of stream processing, aggregation operations, and optical computing.

A. Stream Processing and Aggregation

Stream processing systems continuously handle infinite inputs of data tuples or events by applying a set of operators. A stream processing application is modeled as a directed dataflow graph, where nodes denote operators and edges are data streams between the nodes, connecting upstream data sources (producers) to downstream data sinks (consumers). In Fig. 1a, we illustrate an elementary, linear dataflow graph, a stream processing pipeline with sequentially chained operators[1], where p and c stand for producer and consumer. Each node ω in the operators is an abstract operation that transforms an input stream into an output stream, which can be classified as an aggregation, a pattern, or a filter. Aggregation computes a summary over the input stream, such as sums, averages, and counts. An operator can be executed by different computing substrates, such as CPUs, GPUs, and optics. In our work, we select a single operator ω_1 as an instance of aggregation to study whether it is feasible to replace digital processors with optics while preserving the averaging semantics.

[1]While not covering parallel processing, the pipeline model already allows us to address a variety of interesting real-world applications [13], [14].

B. Optical Computing Systems

Optical computing exploits the physical properties of light, such as superposition, interference, and diffraction, to accomplish computation during propagation. Optical computing systems, which are specifically designed for carrying out computational tasks, can be categorized into two physical platforms. The first system is free-space optical system, where the optical components, like lenses, SLMs, and polarizers, are arranged to manipulate and process the light field. The light field can be obtained by displaying an image on the SLM screen and then illuminating it with a laser. When light passes through multiple components, e.g., two lenses forming a *4f* stage [15] where the lenses are spaced by twice the focal length, various coherent optical transformations (computations) can be achieved, such as Fourier transforms, spatial filtering, and summation. The SLM modulates grayscale values to phase delays, and the polarizer transforms phase into light intensity. Through empirical calibration on the SLM, a monotonic gray-to-intensity correlation function can be established—an initial step in our averaging experiments.

The second system is waveguide-based optics, which utilizes components like integrated modulators and interferometers to direct optical signals within substrates such as integrated circuits on-chip or fiber optic cables.

III. Approach: Optical averaging

This section introduces our approach to optical averaging in free space. We first present a conceptual channel for averaging, followed by a formulation of optical averaging. We then describe energy models for end-to-end energy consumption of our optical averaging channel and the digital equivalent.

A. Optical Averaging Channel

We map an averaging operator w_1 (cf. Fig. 1a) to an optical averaging channel (Fig. 1b), with three stages: encoding, optical averaging w_{avg}, and decoding.

Encoding: The data originate in the digital domain, so to interface with the optical system, an initial encoding step is essential. In a free-space optical system, this step is carried out by digitally loading the original data onto a modulator, in our case, the SLM, which converts it into an optical representation when illuminated.

Optical averaging: The encoded data propagate in the form of light through the optical system composed of polarizers and cascaded *4f* stages. These stages spatially integrate the light field, thereby enabling physical accumulation. By fixing the number of pixels and using an equivalent pixel size during encoding, normalization is realized implicitly.

Decoding: The resulting optical signals are captured by a photodiode and decoded into readable digital data by an optical power meter, which integrates analog-to-digital conversion. The readout values present the light intensity (μW).

Since light modulation is a physical and analog phenomenon by nature, we omit the optical-to-analog conversion in our channel. However, we still specify the analog domain when we measure the cost in the decoding step.

Wenfei Huang, Meike Hofmann, Peter Amthor, and Boris Koldehofe. “Energy-Aware and Flexible Optical Data Stream Processing: A Study of the Averaging Operator in Free Space”, to appear in Proceedings of the 15th International Symposium on Communication Systems, Networks and Digital Signal Processing 2026 (CSNDSP 2026), July 15-17, 2026, Edinburgh, Scotland UK.

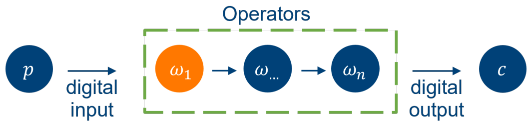

(a) Stream processing pipeline

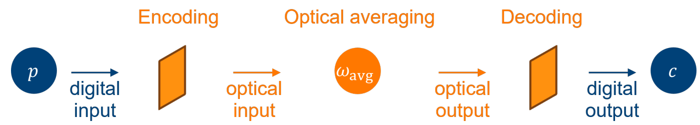

(b) Optical averaging channel

Figure 1: Mapping a stream processing operator within a pipeline to an optical averaging channel: orange parts are enabled by the optics

B. Optical Averaging Formulation

Before performing averaging, the correlation between 8-bit grayscale values $x \in [0, 255]$ and the output intensity is established. A single grayscale value is encoded across $W \times H$ pixels utilized on SLM (W is the number of pixels in the width direction and H is the height), propagating through the averaging channel and outputting the intensity I via the correlation function $I = f(x)$. The correlation is performed only once and is then used for all subsequent averaging computations. We illustrate our optical averaging formulation in Algorithm 1 below.

Conceptually, the tuple in Algorithm 1 is encoded into $(f(x_1), f(x_2), ..., f(x_n))$, and the output is accumulated by the spatial contribution of each element, which arithmetically calculates the average $\frac{f(x_1)}{n} + \frac{f(x_2)}{n} + \cdots + \frac{f(x_n)}{n} = \frac{I_1+I_2+\cdots+I_n}{n}$. For instance, given a digital input $(240, 10, 10, 240)$ for 2×2 pixels SLM display. Fig. 2a shows the encoding of a single value of 10 and outputs intensity $I_{10} = f(10)$, and Fig. 2b is 240 and outputs $I_{240} = f(240)$. With $n = 4$, the data loaded on the SLM screen consists of 4 equivalent blocks, as shown in Fig. 2c, each block referring to one pixel. The final readout value is $\frac{I_{240}+I_{10}+I_{10}+I_{240}}{4}$.

In practical stream processing systems, aggregation is typically applied over sliding or tumbling windows, where the number of data points per window can vary. In the present implementation, we leave out windowing in the optical domain.

Algorithm 1 Optical Averaging Formulation

Input: Digital input tuple $(x_1, x_2, \ldots, x_n)$, with $n \leq W \times H, x_i \in [0, 255](i \in [1, n]), (W \times H) \bmod n = 0$

Procedure:

1: Given n data elements in the tuple as n equal-sized, non-overlapping blocks, each block containing $\frac{W \times H}{n}$ pixels

2: **for** block $i = 1$ to n **do**

3: All pixels in block i are loaded by a single value x_i

4: **end for**

5: Load n blocks ($W \times H$ pixels) onto the SLM

Output: Readout the intensity value I

Notes:

The pixels on a block can be placed anywhere within SLM

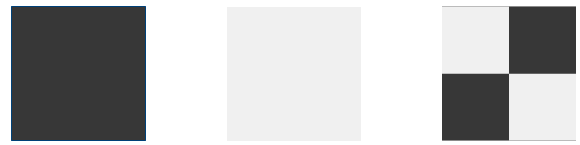

(a) $\begin{bmatrix} 10 & 10 \\ 10 & 10 \end{bmatrix}$ (b) $\begin{bmatrix} 240 & 240 \\ 240 & 240 \end{bmatrix}$ (c) $\begin{bmatrix} 240 & 10 \\ 10 & 240 \end{bmatrix}$

Figure 2: Examples of 2×2-block data input

C. Modeling Energy Consumption

To reason about end-to-end energy consumption, we consider the difference between Fig. 1a and Fig. 1b for a semantically identical averaging operation: The cost $E_{digital}$ of the digital implementation equals the energy used by a digital processor to compute the average (Eq. 1). In contrast, the cost of the optical implementation E_{test} includes the energy spent on input encoding, noted as E_{enc}, optical averaging E_{opt_avg}, output decoding E_{dec}, and the added cost of additional power supplies for the required optical hardware $E_{hardware}$ (SLM and laser in our system), as given in Eq. 2. To obtain the numerical average as an output, the correlation step has to be carried out with an identically structured energy cost $E_{correlation}$ (Eq. 3). The sum of both forms the total energy cost E_{opt} (Eq. 4).

$$E_{digital} = E_{dit\_avg} \tag{1}$$

$$\begin{aligned} E_{test} = E_{enc} + E_{opt\_avg} + E_{dec} \\ + E_{hardware} \end{aligned} \tag{2}$$

$$\begin{aligned} E_{correlation} = E'_{enc} + E'_{opt\_avg} + E'_{dec} \\ + E'_{hardware} \end{aligned} \tag{3}$$

$$E_{opt} = E_{correlation} + E_{test} \tag{4}$$

In line with the consensus on common assumptions in optical computing—light propagation through a passive free-space optical system does not incur additional energy consumption—we assume that the cost of averaging during propagation is dissipation-less, $E_{opt\_avg} = E'_{opt\_avg} = 0$.

IV. Prototypical Evaluation

This section experimentally validates the proposed averaging channel. We set up a prototype of a free-space optical system that features three stages described in section III. The experiment aims to evaluate computational accuracy and end-to-end energy costs by comparing the equivalent digital version.

Wenfei Huang, Meike Hofmann, Peter Amthor, and Boris Koldehofe. "Energy-Aware and Flexible Optical Data Stream Processing: A Study of the Averaging Operator in Free Space", to appear in Proceedings of the 15th International Symposium on Communication Systems, Networks and Digital Signal Processing 2026 (CSNDSP 2026), July 15-17, 2026, Edinburgh, Scotland UK.

A. System Overview

Our free-space optical computing system implements a full pipeline for computing the average, as shown in Fig. 3. The laser (HeNe, $543\,\mathrm{nm}$) is expanded, polarized ($45°$) and guided onto the phase-only SLM (HOLOEYE GAEA 2.1, 3840×2160 pixels, $60\,\mathrm{Hz}$). The reflected light is analyzed by a second polarizer to convert the phase into intensity values and is imaged by a first *4f* stage into an intermediate image plane. A square aperture is installed as a field stop to limit the extent of the data. A second *4f* stage images the pattern onto the photodiode and the optical power meter (PM100USB) measures the optical power. The SLM is connected to a laptop running Windows 11 operating system (OS) via an HDMI cable. The optical power meter is the bridge between the photodiode and the laptop. It translates the output power into readable digits, which are transmitted to the laptop through a USB-A 2.0 interface. In the end, we read the values from the software paired with the converter.

B. Correlation and Test Data

We select a centered square area on the SLM containing 1740×1740 pixels as our input image for both correlation data and test data. To avoid complicating the computation, we calibrate the SLM to expect a linear correlation between the input grayscale value and the output intensity.

During the correlation, we encode the single grayscale value iteratively for every 5-value increment, starting from 0 (thus 52 data in total). The SLM displays each input for $0.1\,\mathrm{s}$, including the subsequent test data averaging. After recording the response intensity values, we formulate the correlation function demonstrated in Fig. 4, which is nearly linear.

To understand how different input sizes affect energy consumption when compared to digital averaging, we generate test data in three different dimensions: 3-dimensional, 174-dimensional, and 1740-dimensional data (n-dimensional data refer to $n \times n$ data elements in the input tuple, thus $n \times n$ blocks). To investigate any impact of geometrical patterns, we further encode the data in three different patterns for each dimension: normal, vertical, and horizontal. The grayscale values in one block are drawn randomly. Fig. 5 shows examples of 3-dimensional data in the three patterns, which is exactly what the SLM screen is displaying in our setting. There are nine different types of data inputs combining dimensions and patterns. For each combination, five are generated, so the total number of test data is 45 and thus the number of averaging computations.

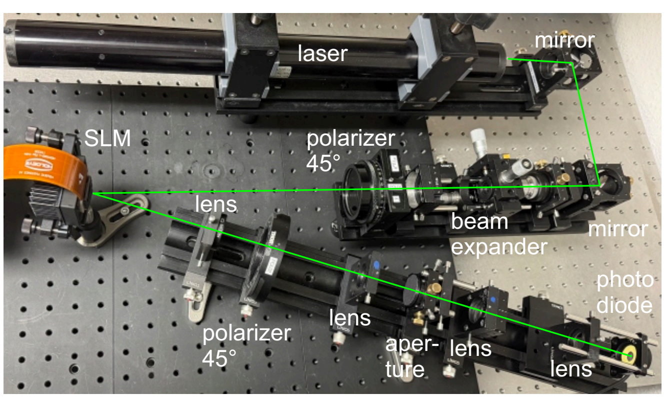

Figure 3: Photo of our free-space optical system.

C. Computation Accuracy vs. Digital Baseline

We load the test data into SLM Python application programming interface (API) in ascending order by dimension, with the pattern order of each dimension being normal, vertical, and then horizontal. As the correlation function is nearly linear, the digital average results are obtained by first averaging the grayscale values and then fitting them to this function. Two metrics are used to evaluate the accuracy against the digital baseline, relative error (RE) calculated in Eq. 5 and root mean squared error (RMSE) in Eq. 6. We categorize the accuracy results by different dimensions in Table I and different patterns in Table II. In the two tables, the average RE and RMSE reach 1.8% and 0.68, respectively, but deviations in accuracy are evident across different dimensions and patterns.

$$RE_i = \frac{|Optical\_result_i - Digital\_result_i|}{|Digital\_result_i|} \tag{5}$$

$$RMSE = \sqrt{\frac{1}{n}\sum_{i=1}^{n}(Optical\_result_i - Digital\_result_i)^2} \tag{6}$$

D. Energy Consumption vs. Digital Baseline

We use Intel Performance Counter Monitor (PCM) as our measurement tool for both digital and optical implementations. It monitors CPU energy usage across the entire machine. We measure the cost across the three dimensions separately to observe their impact on energy consumption.

We follow Eq. 4 to compute the energy produced in our optical implementation. To compute E_{test} in Eq. 2, we detail

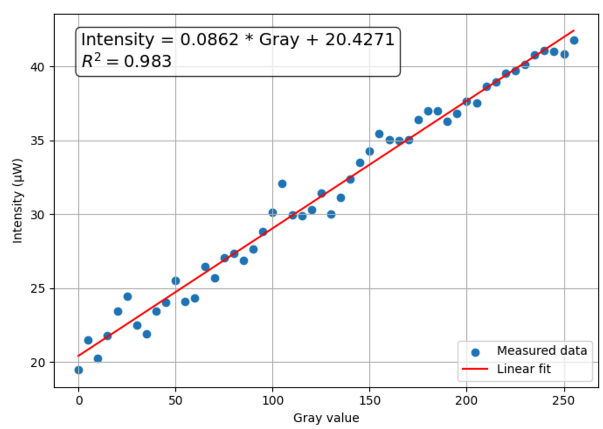

Figure 4: Correlation function between single gray-scale value and intensity (μW).

Wenfei Huang, Meike Hofmann, Peter Amthor, and Boris Koldehofe. “Energy-Aware and Flexible Optical Data Stream Processing: A Study of the Averaging Operator in Free Space”, to appear in Proceedings of the 15th International Symposium on Communication Systems, Networks and Digital Signal Processing 2026 (CSNDSP 2026), July 15-17, 2026, Edinburgh, Scotland UK.

(a) normal (b) vertical (c) horizontal

Figure 5: Examples of 3-dimensional data in different patterns

the cost from the encoding ($E_{enc} = E_{SLM\_driving} + E_{HDMI}$) since programming the input and loading it onto SLM is powered by the laptop. The same applies to the decoding ($E_{dec} = E_{Readout\_software} + E_{USB}$). $E_{SLM\_driving}$ and $E_{Readout\_software}$ are characterized by PCM, while E_{HDMI} and E_{USB} can only be estimated by their standard power specifications (USB-$2.5\,$W, HDMI-$1.5\,$W). Power supplies for the SLM and the laser in $E_{hardware}$ are gauged by an electric power meter (NETIO PowerBOX 4KF), where the real-time load of the SLM is $15\,$W, and the laser is $19\,$W. For 45 test data, each of them is displayed for $0.1\,$s. The cost, comprising E_{HDMI}, E_{USB}, and $E_{hardware}$, can then be computed by multiplying the total power load by the time required to display and obtain the results. The whole procedure for computing E_{test} is identical to that for correlation (in Eq. 3). Since the correlation step only needs to be executed once, the cost will be distributed evenly among all the test data.

The energy measurements per averaging operation are recorded in Table III. We can observe that in lower dimension, digital averaging is more energy-efficient than the optical version, with a factor of 10^2. However, when it comes to 1740-dimensional data, the optical system outperforms the digital implementation.

E. Evaluation Summary and Discussion

The averaging implementation prototyped on the free-space optical system achieves high computational accuracy, with an average RE of around 2%, while the unexpected deviations need further investigation in our future work. In terms of energy efficiency, the amount of energy consumed by digital computing is significantly determined by the data scales, while optical computing is akin to a system that incurs a heavy but one-time expense at the start and will not suffer dramatically from dimensional surges.

Table I: Accuracy across dimensions

Metric	Dimensions(-dim)			Entire	Standard deviation
	3	174	1740		
average RE (%)	2.10	1.13	2.17	1.80	0.47
RMSE	0.92	0.41	0.72	0.68	0.20

Table II: Accuracy across patterns

Metric	Patterns			Entire	Standard deviation
	normal	vertical	horizontal		
average RE (%)	2.56	1.40	1.42	1.79	0.54
RMSE	0.90	0.49	0.52	0.63	0.18

Table III: Energy measurements of optical and digital averaging per operation

Per averaging	Optical			Digital		
	3-dim	174-dim	1740-dim	3-dim	174-dim	1740-dim
Energy cost (J)	22.91	23.05	26.65	0.15	0.43	41.99

System latency, as another factor, is determined by the encoding time, the refresh rate of SLMs and the sensor readout time. In experiments with 1740-dimensional data, we observe benefits in optical system ($9.98\,$s) compared to digital averaging ($34.51\,$s). While SLMs enables fast and flexible prototyping, their cost and bulky size constrain miniaturization, making practical uses suitable at the system level, e.g., server-side plugged accelerators. This work focuses on mapping a single stream operator, but conceptually, multiplexing of optical paths can support the concurrent processing of distinct operators, offering parallelism analogous to that in stream processing.

V. Related Work

Optical computations: Recent studies largely explore optical computing as a hardware accelerator for NN primitives, which refer to the optical implementations of matrix multiplication using free-space or integrated photonic systems [16]–[19], also nonlinear activation functions [20], and convolution and pooling operations [21], [22]. Some architectures of NNs adopt global average pooling [23], which is mathematically equivalent to the averaging operation.

Averaging in data processing applications: As highly generalizable as matrix multiplication, averaging can also be positioned in a variety of diverse domains, including image processing [24], in-network traffic characterization [25], digital signal processing [26], and remote sensing [27], not only within NNs. Under real-time stream processing scenarios, averaging is typically implemented as a windowed aggregation function and is practically used as a building block in, e.g., monitoring, anomaly detection, and statistical analysis [28], [29].

Free-space optical system: In the literature [30]–[32], two modern optical computing paradigms: free-space optical computing and waveguide-based optics, along with their applications are reviewed. In a free-space optical system, the *4f* stage is versatile, enabling spatial operations such as optical Fourier transformation. The SLM and the digital micro-mirror device (DMD) are two types of modulators commonly used in free-space optical systems to encode digital data into the optical domain. Among the two, the SLM is the standard choice when high refresh rates are not necessary for experimental systems [33].

Relevance: The studies mentioned above leave a clear gap at the intersection of optical computing and stream processing. Existing work typically simulates the optical system using software or does not conduct a qualitative and quantitative energy analysis. Our work addresses these gaps by integrating 1) the investigation of an optically feasible operator that is

Wenfei Huang, Meike Hofmann, Peter Amthor, and Boris Koldehofe. "Energy-Aware and Flexible Optical Data Stream Processing: A Study of the Averaging Operator in Free Space", to appear in Proceedings of the 15th International Symposium on Communication Systems, Networks and Digital Signal Processing 2026 (CSNDSP 2026), July 15-17, 2026, Edinburgh, Scotland UK.

missing but essential in the context of real-world big data and stream processing applications, 2) the establishment of a practical free-space optical system rather than a software simulation, 3) the end-to-end energy measurement of the optical implementation, and 4) the evaluation of the energy consumption and accuracy of the optical implementation against an identical digital baseline.

VI. Conclusions and future work

In this work we investigate the feasibility of practically implementing stream aggregation – by the example of averaging – through a free-space optical system. Our prototype demonstrates a promisingly high accuracy with an average RE of 2%, and improved energy efficiency when processing high-density data, despite not always being more efficient in absolute numbers. Beyond the experimental implementation, we introduce a conceptual mapping from a linear stream processing pipeline to a three-stage averaging channel, where we elucidate how averaging occurs in optics and provide a basis for end-to-end energy consumption models. Future work will extend this study to understand the deviation in the accuracy by experimenting with more data, and to enhance energy efficiency by utilizing the full resolution of the SLM, and investigating more complex operations such as dynamic windowing.

References

[1] M. A. Alam, A. Nabil, A. A. Mintoo, and A. Islam, "Real-Time Analytics In Streaming Big Data: Techniques And Applications," *Non human journal*, vol. 1, pp. 104–122, Oct. 2024.

[2] S. Chang, Y. Zhang, J. Tang, D. Yin, Y. Chang, M. A. Hasegawa-Johnson, and T. S. Huang, "Streaming Recommender Systems," in *Proceedings of the 26th International Conference on World Wide Web*, WWW '17, pp. 381–389, International World Wide Web Conferences Steering Committee, Apr. 2017.

[3] T. Kolajo, O. Daramola, and A. Adebiyi, "Big data stream analysis: A systematic literature review," *Journal of Big Data*, vol. 6, p. 47, June 2019.

[4] O.-C. Marcu and P. Bouvry, "Big data stream processing," tech. rep., University of Luxembourg, Sept. 2024.

[5] A. Sapio, I. Abdelaziz, A. Aldilaijan, M. Canini, and P. Kalnis, "In-Network Computation is a Dumb Idea Whose Time Has Come," in *Proceedings of the 16th ACM Workshop on Hot Topics in Networks*, HotNets '17, pp. 150–156, Association for Computing Machinery, Nov. 2017.

[6] "Powering intelligence: Analyzing artificial intelligence and data center energy consumption," White Paper 3002028905, Electric Power Research Institute (EPRI), May 2024.

[7] R. Leichter, "Number of Data Centers by Country (November 2025) | Cargoson." https://www.cargoson.com/en/blog/number-of-data-centers-by-country, Nov. 2025.

[8] "Energy and ai," tech. rep., International Energy Agency, Apr. 2025.

[9] Q. Zhang, Z. Meng, X. Hong, Y. Zhan, J. Liu, J. Dong, T. Bai, J. Niu, and M. J. Deen, "A survey on data center cooling systems: Technology, power consumption modeling and control strategy optimization," *Journal of Systems Architecture*, vol. 119, p. 102253, Oct. 2021.

[10] P. L. McMahon, "The physics of optical computing," *Nature Reviews Physics*, vol. 5, pp. 717–734, Dec. 2023.

[11] T. Hoßfeld, "Energy use in data centers: Current figures and trends," *VDE ITG News*, 2025.

[12] T. Fu, J. Zhang, R. Sun, Y. Huang, W. Xu, S. Yang, Z. Zhu, and H. Chen, "Optical neural networks: Progress and challenges," *Light: Science & Applications*, vol. 13, p. 263, Sept. 2024.

[13] M. Salehe, Z. Hu, S. H. Mortazavi, I. Mohomed, and T. Capes, "VideoPipe: Building Video Stream Processing Pipelines at the Edge," in *Proceedings of the 20th International Middleware Conference Industrial Track*, Middleware '19, pp. 43–49, Association for Computing Machinery, Dec. 2019.

[14] P. Wiener, P. Zehnder, and D. Riemer, "Managing geo-distributed stream processing pipelines for the IIoT with StreamPipes edge extensions," in *Proceedings of the 14th ACM International Conference on Distributed and Event-based Systems*, DEBS '20, pp. 165–176, Association for Computing Machinery, July 2020.

[15] A. Lugt, "Signal detection by complex spatial filtering," *IEEE Transactions on Information Theory*, vol. 10, pp. 139–145, Apr. 1964.

[16] R. Hamerly, L. Bernstein, A. Sludds, M. Soljačić, and D. Englund, "Large-Scale Optical Neural Networks Based on Photoelectric Multiplication," *Physical Review X*, vol. 9, p. 021032, May 2019.

[17] Y. Shen, N. C. Harris, S. Skirlo, M. Prabhu, T. Baehr-Jones, M. Hochberg, X. Sun, S. Zhao, H. Larochelle, D. Englund, and M. Soljačić, "Deep learning with coherent nanophotonic circuits," *Nature Photonics*, vol. 11, pp. 441–446, July 2017.

[18] Z. Zhong, W. Wang, M. Ghobadi, A. Sludds, R. Hamerly, L. Bernstein, and D. Englund, "IOI: In-network Optical Inference," in *Proceedings of the ACM SIGCOMM 2021 Workshop on Optical Systems*, OptSys '21, pp. 18–22, Association for Computing Machinery, Aug. 2021.

[19] H. Zhou, J. Dong, J. Cheng, W. Dong, C. Huang, Y. Shen, Q. Zhang, M. Gu, C. Qian, H. Chen, Z. Ruan, and X. Zhang, "Photonic matrix multiplication lights up photonic accelerator and beyond," *Light: Science & Applications*, vol. 11, p. 30, Feb. 2022.

[20] M. Miscuglio, A. Mehrabian, Z. Hu, S. I. Azzam, J. George, A. V. Kildishev, M. Pelton, and V. J. Sorger, "All-optical nonlinear activation function for photonic neural networks [Invited]," *Optical Materials Express*, vol. 8, p. 3851, Dec. 2018.

[21] S. Colburn, Y. Chu, E. Shilzerman, and A. Majumdar, "Optical frontend for a convolutional neural network," *Applied Optics*, vol. 58, p. 3179, Apr. 2019.

[22] H. Sadeghzadeh, S. Koohi, and A. Paranj, "Free-Space Optical Neural Network Based on Optical Nonlinearity and Pooling Operations," *IEEE Access*, vol. PP, pp. 1–1, Oct. 2021.

[23] M. Lin, Q. Chen, and S. Yan, "Network in network," in *2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14-16, 2014, Conference Track Proceedings* (Y. Bengio and Y. LeCun, eds.), 2014.

[24] C. Shao, P. Kaur, and R. Kumar, "An Improved Adaptive Weighted Mean Filtering Approach for Metallographic Image Processing**," *Journal of Intelligent Systems*, vol. 30, pp. 470–478, Jan. 2021.

[25] C. Brülhart, N. S. Bülbül, N. Ole Tippenhauer, and M. Fischer, "Transparent TSN for Agnostic End-hosts via P4-based Traffic Characterization at Switches," in *2024 IEEE 49th Conference on Local Computer Networks (LCN)*, pp. 1–8, Oct. 2024.

[26] D. Purnamasari, K. Aji, and H. Sukri, "Digital Moving Average Filter Application for Echo Signals and Temperature," *E3S Web of Conferences*, vol. 328, p. 02007, Dec. 2021.

[27] J. Liu, Z. Gao, Y. Li, S. Lv, J. Liu, C. Yang, J. Liu, Z. Gao, Y. Li, S. Lv, J. Liu, and C. Yang, "Ranging Offset Calibration and Moving Average Filter Enhanced Reliable UWB Positioning in Classic User Environments," *Remote Sensing*, vol. 16, July 2024.

[28] J. Verwiebe, P. M. Grulich, J. Traub, and V. Markl, "Survey of window types for aggregation in stream processing systems," *The VLDB Journal*, vol. 32, pp. 985–1011, Sept. 2023.

[29] D. J. Hill and B. S. Minsker, "Anomaly detection in streaming environmental sensor data: A data-driven modeling approach," *Environmental Modelling & Software*, vol. 25, pp. 1014–1022, Sept. 2010.

[30] J. Hu, D. Mengu, D. C. Tzarouchis, B. Edwards, N. Engheta, and A. Ozcan, "Diffractive optical computing in free space," *Nature Communications*, vol. 15, p. 1525, Feb. 2024.

[31] R. Chen and W. Gao, "Free-space Optical Computing Systems," *Annalen der Physik*, vol. n/a, no. n/a, p. 2400390, 2025.

[32] N. L. Kazanskiy, M. A. Butt, S. N. Khonina, N. L. Kazanskiy, M. A. Butt, and S. N. Khonina, "Optical Computing: Status and Perspectives," *Nanomaterials*, vol. 12, June 2022.

[33] A. B. Ayoub and D. Psaltis, "High speed, complex wavefront shaping using the digital micro-mirror device," *Scientific Reports*, vol. 11, p. 18837, Sept. 2021.