Streaming Large-Scale Microscopy Data to a Supercomputing Facility

Samuel S Welborn, Chris Harris, Stephanie M Ribet, Georgios Varnavides, Colin Ophus, Bjoern Enders, Peter Ercius

Microscopy and Microanalysis, 2024, **00**, 1–9
https://doi.org/10.1093/mam/ozae109

Original Article

Microscopy AND Microanalysis

Streaming Large-Scale Microscopy Data to a Supercomputing Facility

Samuel S. Welborn[1], Chris Harris[1], Stephanie M. Ribet[2], Georgios Varnavides[2], Colin Ophus[2], Bjoern Enders[1], and Peter Ercius[2,*]

[1]*National Energy Research Scientific Computing Center (NERSC), Lawrence Berkeley National Laboratory, Berkeley, CA 94720, US*
[2]*National Center for Electron Microscopy (NCEM), The Molecular Foundry, Lawrence Berkeley National Laboratory, Berkeley, CA 94720, US*

***Corresponding author:** Peter Ercius, E-mail: percius@lbl.gov

Abstract

Data management is a critical component of modern experimental workflows. As data generation rates increase, transferring data from acquisition servers to processing servers via conventional file-based methods is becoming increasingly impractical. The 4D Camera at the National Center for Electron Microscopy generates data at a nominal rate of 480 Gbit s^{-1} (87,000 frames s^{-1}), producing a 700 GB dataset in 15 s. To address the challenges associated with storing and processing such quantities of data, we developed a streaming workflow that utilizes a high-speed network to connect the 4D Camera's data acquisition system to supercomputing nodes at the National Energy Research Scientific Computing Center, bypassing intermediate file storage entirely. In this work, we demonstrate the effectiveness of our streaming pipeline in a production setting through an hour-long experiment that generated over 10 TB of raw data, yielding high-quality datasets suitable for advanced analyses. Additionally, we compare the efficacy of this streaming workflow against the conventional file-transfer workflow by conducting a postmortem analysis on historical data from experiments performed by real users. Our findings show that the streaming workflow significantly improves data turnaround time, enables real-time decision-making, and minimizes the potential for human error by eliminating manual user interactions.

Key words: streaming, 4D-STEM, high-performance computing, real-time processing, zmq

Introduction

In the era of big data, the scientific community faces significant challenges in data management (Rao, 2020; Spurgeon et al., 2021). This is especially evident at experimental user and core facilities, where advancements in instrumentation, such as faster detectors and increased light source brightness, have led to an exponential increase in data generation rates. The traditional methods of data storage and movement (e.g., personal flash drives) are becoming increasingly untenable.

In 2019, a new detector called the 4D Camera was installed on the TEAM 0.5 microscope at the National Center for Electron Microscopy (NCEM) facility of The Molecular Foundry at Lawrence Berkeley National Laboratory (LBNL) (Ercius et al., 2024). This detector produces data at a rate of 480 Gbit s^{-1} (equivalent to 87,000 frames s^{-1}), yielding datasets of up to 700 GB for a 15-s acquisition. Other microscopy facilities are installing similar high frame rate detectors with the ability to routinely generate >100 GB datasets (Chatterjee et al., 2021; Zambon et al., 2023). While these technological advancements provide new avenues for scientific exploration, they also pose significant challenges in data management, analysis, and acquisition. New opportunities for development include on-the-fly processing for quick feedback on an experimental approach and implementation of complex experimental pipelines, such as focal series or tomography (Pelz et al., 2021*b*, 2023) that leverage the capabilities of these advanced detectors. Given that microscope time is a limited and valuable resource, rapid data analysis that provides feedback on the quality of large data sets during a microscope session is crucial for improving throughput.

To mitigate these challenges, a collaborative effort involving high performance computing (HPC) experts at the National Energy Research Scientific Computing Center (NERSC), electron microscopy experts at NCEM, and software development experts at Kitware, Inc. led to the utilization of NERSC for data reduction and the development of a web frontend called *Distiller* to facilitate data management (Harris & Genova, 2023). HPC systems are typically accessed through command line interfaces, which are often unfamiliar to microscopists. *Distiller*, on the other hand, allows users to transfer and process data at NERSC through simple web-based interactions. This effort, which was part of a broader initiative at LBNL called The Superfacility Project, greatly improved the workflow for the 4D Camera (Enders et al., 2020; Harris & Genova, 2023; Welborn et al., 2024).

Despite its utility, data analysis at NERSC was constrained by file-based input/output (I/O) steps that created bottlenecks at two stages: (1) writing data from random-access memory (RAM) to local disk storage at NCEM and (2) file transfer from NCEM to NERSC before computation. We note that file-based data movement is the common workflow across most detector systems. The dependence on file-based I/O operations slows down data processing, constrains the scope of feasible experiments, and relies on file systems possibly shared by multiple users. At NERSC, for example, there is a greater amount of file system contention between users on both the

Received: June 28, 2024. **Revised:** August 22, 2024. **Accepted:** October 14, 2024

community file system and Perlmutter scratch, causing variability in I/O throughput. Conversely, NERSC's total network border bandwidth is 1.2 Tbit s^{-1} and, so far, we have observed no significant network contention. In our recent work, we showed that, by circumventing file-based operations through streaming data from detector buffer memory directly to NERSC compute node memory over the network, we improved throughput by 5- to 14-fold (Welborn et al., 2024). In the present work, we showcase the advantages of a streaming workflow for microscopy experiments using the 4D Camera as a case study.

This manuscript is organized as follows. In the Background section, we provide an overview of 4D scanning transmission electron microscopy (4D-STEM) and discuss difficulties in managing the substantial datasets generated by the 4D Camera. Then, we briefly outline the components of the streaming pipeline. Next, we describe enhancements to *Distiller* that obviate the need for an in-depth understanding of HPC. Finally, we demonstrate the practical benefits of streaming through a comparative analysis of real user experiments employing both workflows.

Background

Transmission Electron Microscopy and 4D-STEM

Transmission electron microscopy (TEM) provides insights into the atomic and molecular structure of materials, making it a cornerstone characterization technique across scientific disciplines from materials science to biology. Scanning TEM (STEM) operates in a mode where an electron probe is focused onto the sample and rastered over a two-dimensional set of probe positions. Postspecimen detectors register electron events in diffraction space that can be mapped to specific probe positions. The versatility of STEM extends its utility beyond conventional imaging, facilitating advanced analytical methods such as spectroscopy, electron tomography, ptychography, and holography (Ercius et al., 2015; Miao et al., 2016; Yasin et al., 2016, 2018; Stevens et al., 2018; Ophus, 2019, 2023; Ben-Moshe et al., 2021; Varnavides et al., 2023; Ribet et al., 2024).

Recent advancements in detector technology have ushered in a new era for STEM. Specifically, the introduction of direct electron detectors (DEDs) has dramatically accelerated data acquisition rates and opened new experimental possibilities (Levin, 2021; Ercius et al., 2024). DEDs can acquire data with a temporal resolution ranging from milliseconds to microseconds enabling a technique generally called 4D-STEM because two-dimensional (2D) diffraction patterns are acquired at a series of 2D probe positions (Ophus, 2019). The resulting 4D dataset contains a wealth of both structural and compositional information about the sample. Analysis of the diffraction patterns can reveal the sample's overall crystal orientation, strain, and material phase, enabling a detailed mapping of these properties to provide a comprehensive characterization of the material (Ophus, 2019). One of the applications of 4D-STEM is phase-contrast imaging—while detectors record only the intensity of the exit wave after interaction with the sample, it is possible to reconstruct the phase, leading to dose-efficient characterization of weakly scattering signals. Phase retrieval STEM methods, such as differential phase-contrast (DPC) (Dekkers & De Lang, 1974; Waddell & Chapman, 1979; Shibata et al., 2012; Cao et al., 2018), which measures the change in the center of mass of diffraction patterns, and advanced algorithms such as ptychography, offer enhanced contrast and resolution (Nellist et al., 1995; Enders & Thibault, 2016; Varnavides et al., 2023).

The size of 4D-STEM data introduce significant challenges in data management. An illustrative case is the 4D Camera, which can accumulate 2D diffraction patterns at a rate of 87,000 Hz (nominally 200 TB h^{-1}) highlighting the need for informed data treatment beyond current capabilities at most electron microscopy laboratories. Solving the challenges that come with storing and processing large datasets in a timely manner necessitates an examination of the pathway data takes within the data acquisition (DAQ) system and processing workflow.

High Data Rate Acquisition and its Challenges

The DAQ system for the 4D Camera at NCEM, developed in-house at LBNL, integrates both software and hardware elements to achieve such high data rates (Fig. 1). The 4D Camera sensor (Fig. 1a, bottom) is partitioned into four sectors, each of which is connected to a dedicated receiving server via twelve 10 Gbit s^{-1} connections through field-programmable gate arrays (FPGAs). As the electron beam rasters across the sample (Fig. 1a, top), a 576 × 576 pixel frame is acquired at each scan position. Each 144 × 576 pixel sector is processed by an FPGA and transmitted to its corresponding data receiving server (Figs. 1a–1b). Upon completion of a scan, the data are written as binary data files (Fig. 1c) to flash storage. For a more comprehensive description of this DAQ system, the reader is referred to Ercius et al. (2024) and Welborn et al. (2024).

With a data rate of 480 Gbits s^{-1}, a single 15-s acquisition using the 4D Camera generates approximately 700 GB of data (Ercius et al., 2024). The large data volume manifests three distinct but interrelated challenges: (1) limited local disk storage capacity, where the available eight TB of flash storage can only accommodate eleven full scans; (2) the computational burden associated with processing large datasets, which overwhelms local dedicated resources; and (3) the time-intensive nature of writing large files to disk, which blocks the system from further data acquisition. Collectively, these challenges substantially reduce user productivity and waste precious beam time. It is important to note that challenges in data management and computational limitations extend beyond NCEM to other Experimental and Observational Science (EOS) facilities, and these problems will intensify in the future (Rao, 2020; Spurgeon et al., 2021). A notable unscalable example is the Event Horizon Telescope data transfer protocol, which involved physically transporting hard disk drives from the telescope to a processing facility to produce the now-famous black hole image (Doeleman et al., 2023).

Initial Mitigation Strategies

The 4D Camera is designed to acquire frames containing a small number of electrons, leading to a sparse data set. Thus, we can simultaneously mitigate the first challenge (storage capacity) and remove detector noise from our data through compression. The software package *stempy* (Avery et al., 2023) efficiently transforms the raw data into a more manageable sparse format by finding and keeping only the locations of single electron hits, a process called "electron counting" in this manuscript (Battaglia et al., 2009; Pelz et al., 2021*a*; Ercius et al., 2024). This transformation (represented

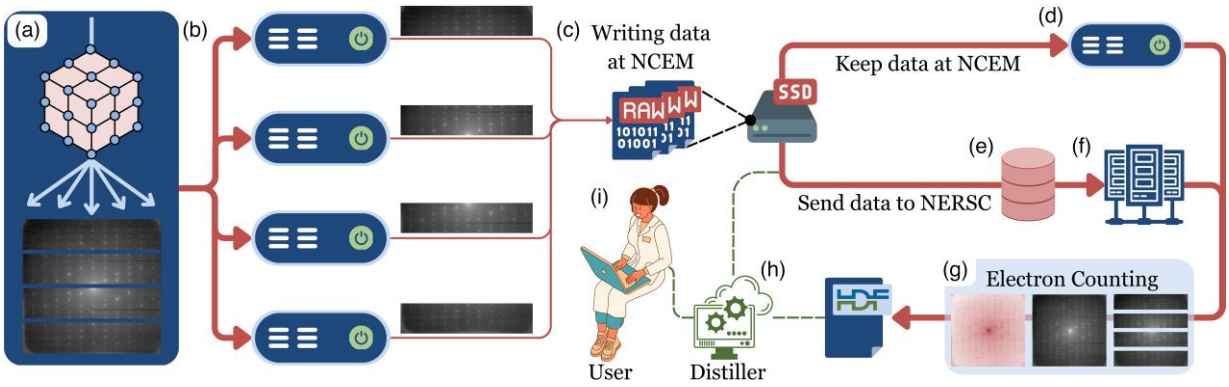

Fig. 1. Schematic illustrating both the DAQ system and the initial mitigation strategies for managing large-scale 4D-STEM datasets generated at NCEM. A user begins an experiment using the TEAM 0.5 microscope software for the four-sector 4D Camera (**a**). The camera is connected to data receiving servers through FPGAs (**b**). Each server ingests all data into RAM and subsequently writes it to an eight TB flash storage system (**c**), which takes around 150 s for a 700 GB dataset. The data are either processed locally at NCEM on a single server with ten CPU cores (**d**), or transferred to NERSC's filesystems (**e**) and processed with more robust compute resources (**f**). Data processing is illustrated in (**g**), showing the assembly of disconnected sectors into coherent frames and subsequent electron counting of these frames. This processed data are saved in a single HDF5 file. The *Distiller* web application (**h**) enables the user (**i**) to initiate file transfers to NERSC's file systems, perform electron counting, and launch analysis notebooks in NERSC's *Jupyter* environment. © Vectorslab, MedicaLineArt, Rolas Design, Juicy Fish, Icon54, Muhammad Wakas, shmaiinc, pixabay via Canva.com.

graphically in Fig. 1g) results in an order of magnitude data size reduction (alleviating some storage pressure) and the reduction of detector noise. The raw detector data are typically deleted after they have been electron counted.

While *stempy* significantly reduces storage requirements, it introduces the second challenge: computational demands for quickly processing large datasets. The processing time for this operation on local resources, represented in Figure 1d, is considerable. At NCEM, the computational resources are limited to ten CPU cores, which makes the electron counting of a 700 GB dataset a time-consuming task (10–12 min). During this time, the detector cannot be used because the same computational resources are shared for both data acquisition and processing. In contrast, each CPU node on NERSC's newest supercomputer, Perlmutter, is equipped with 128 CPU cores (Fig. 1f) and 512 GB DDR4 memory, and multiple compute nodes can be allocated to parallelize the electron counting process. Upgrades to NERSC's computational infrastructure (which have occurred since the 4D Camera was installed) translate into immediate improvements in both the NCEM processing pipeline and for the broader NERSC user community, thereby optimizing resource utilization. Absent this integration, any dedicated compute nodes installed at NCEM require local maintenance and remain underutilized, particularly in periods between experiments. Moreover, by integrating NCEM's workflow with NERSC's infrastructure, NCEM users gain access to NERSC's rich computing and data ecosystem, which is particularly advantageous for processing their data during and after an experiment. This integration not only streamlines NCEM's operations but also provides a blueprint for the efficient deployment of compute resources beyond a single detector or EOS facility (Enders et al., 2020; Bard et al., 2022).

Recognizing the advantages of centralized compute/storage resources for managing large datasets, the *Distiller* (Fig. 1h) application was developed to facilitate user interactions with the detector and NERSC. During data acquisition, *Distiller* presents status and metadata using a user-friendly web-based frontend, allowing users (Fig. 1i) to initiate data transfers to NERSC (Fig. 1e). Then, the data are electron counted using *stempy* on Perlmutter (Figs. 1f–1g) (Enders et al., 2020; Harris & Genova, 2023). After counting, the end result is a single sparse HDF5 file ready for further analysis. NERSC can then provide access restrictions based on user credentials, compute for further analysis, and file transfer to other sites. We provide a screencast of this workflow in Supplementary Video 1. It is also important to recognize that by collaborating with software development experts at Kitware and HPC specialists at NERSC, we avoided the technical debt often associated with *ad hoc* scripts developed by microscopists, who do not typically have the bandwidth to develop seamlessly integrated tools like *Distiller*.

Despite these advancements, writing/reading large files to/from disk remains an unresolved bottleneck, leading to the third challenge that impedes the efficient transfer of high-volume data.

I/O Bottlenecks in Data Transmission

Four critical I/O operations slow down the transmission of data from NCEM to NERSC:

1. Writing the data to a local drive at NCEM.
2. Reading the data from the local drive and transferring it to NERSC over a fiber network.
3. Writing the data to NERSC's file systems.
4. Reading the data into NERSC compute node memory for electron counting.

These file I/O bottlenecks present a dual challenge: they slow down data transfer and analysis and also restrict the types of experiments that can be conducted. For instance, they preclude the possibility of running automated experiments over extended periods (Pattison et al., 2023), because human intervention is required to manage data transfer and counting once the local eight TB file system is full.

Streaming Data from NCEM to NERSC

To overcome the I/O bottlenecks outlined above, we have developed a streaming service that facilitates the transmission of microscope data from NCEM servers to NERSC compute nodes without using file storage. The foundation of our solution is a socket-based network that facilitates RAM-to-RAM data transfer for real-time processing (Fig. 2). Sockets serve as integral components in networked systems, facilitating the

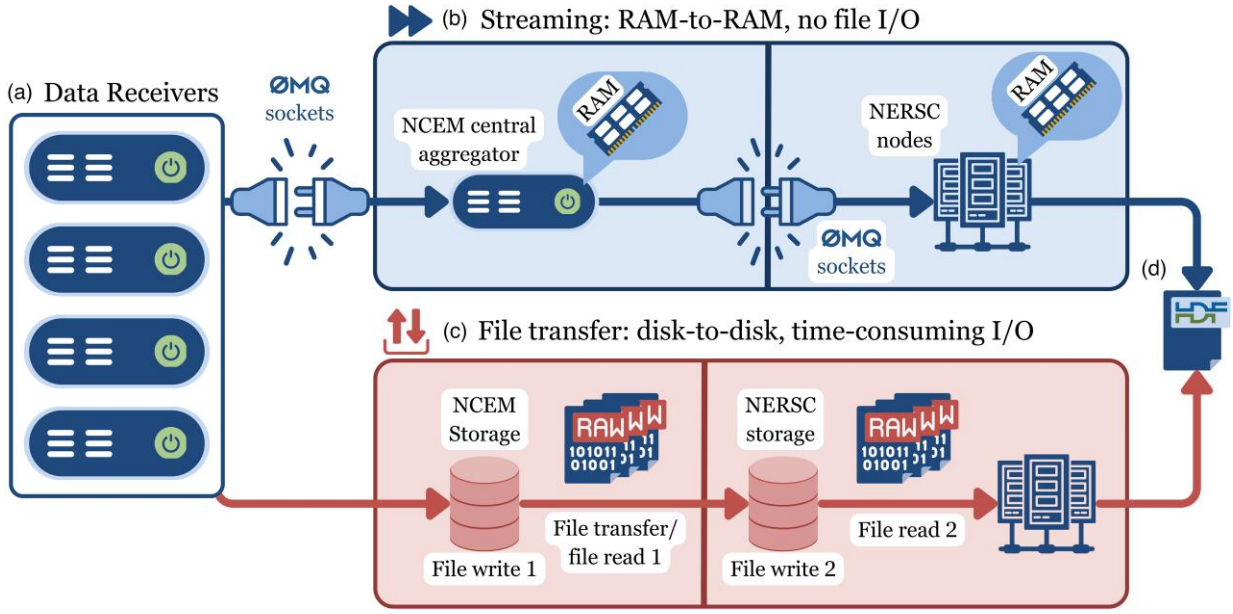

Fig. 2. Schematic comparison of the data streaming pipeline (blue pathway, **a**–**b**–**d**) with the file transfer pipeline (red pathway, **a**–**c**–**d**). Starting from the data receivers (**a**), the streaming approach employs *ZeroMQ* sockets to bypass raw file disk storage at NCEM, enabling direct RAM-to-RAM transfer. Sockets are created on the data receivers, a centralized aggregator server at NCEM, and NERSC compute nodes to facilitate this transmission. Conversely, the file transfer approach requires several intermediate file storage operations to move the data from NCEM to NERSC. In both pathways, the thick vertical line indicates the network border between NCEM and NERSC. Using *stempy*, the data are electron counted and saved in a single HDF5 file for further processing (**d**). © Icon Jam, Rolas Design, Juicy Fish, Prosymbol, Muhammad Wakas via Canva.com.

exchange of data packets between interconnected devices; by using sockets, we are taking advantage of the progress made in commercial internet infrastructure to improve scientific computing. Our architecture utilizes Zero Message Queue (*ZeroMQ*), a network socket library, to establish communication between the key elements of our pipeline: the data receiving servers at NCEM, a centralized aggregator server at NCEM, and the compute nodes at NERSC. We should note here that other user facilities are starting to take the streaming approach to data movement. A notable recent example from the Advanced Photon Source streams data to the Argonne Leadership Computing Facility from Experimental Physics and Industrial Control System (EPICS)-based beamline detectors (Veseli et al., 2023). In our pipeline, we were unable to make use of this package as the 4D Camera does not have an EPICS areaDetector driver. This section's content serves as a high-level synopsis of our approach. For a more comprehensive overview of the methods and system architecture, the reader is directed to our recent technical work (Welborn et al., 2024).

Intercepting File Write at NCEM

The data receiving servers at NCEM (Fig. 2a) handle detector data retrieval, data formatting, and disk storage of raw data files (see Background section). Traditionally, each server accumulates data for one sector of the detector in memory during a scan and writes it to disk as files (Fig. 2c) after acquisition is complete. We replaced this disk write operation with our *ZeroMQ* streaming operation represented by the outlet socket attached to Figure 2a. These sockets transmit the data from the server's RAM to a central aggregator server.

Routing the Data to NERSC

The aggregator server routes data to NERSC for frame reassembly and processing. Its sockets are graphically represented by the inlet socket attached to Figure 2b. The routing strategy on the central aggregator uses sector metadata (the frame number) to forward data (outlet socket attached to the aggregator in Fig. 2b) to its corresponding node at NERSC (inlet socket attached to NERSC nodes in Fig. 2b). This data routing ensures equitable distribution of frames across the NERSC compute nodes, maintaining a consistent computational load across them. Further, it guarantees that all sectors of a given frame are routed to the same NERSC compute node—sector data are initially dispersed among the receiving servers (see Fig. 1b), and they must be assembled on the same NERSC node before processing (see Fig. 1g).

Live Electron Counting at NERSC

On the NERSC nodes, processes with *ZeroMQ* sockets initiate outbound connections to the *ZeroMQ* sockets on the aggregator server, which ingest data into the nodes' RAM (Fig. 2b). Full frames are automatically processed using the electron counting algorithm in the *stempy* package (Avery et al., 2023). After all frames have been received, the sparse, electron-counted data are saved in a single HDF5 file (Fig. 2d). The entire system is now ready for another acquisition. It is important to note that NERSC security policy only permits outbound connections from NERSC compte nodes to servers outside the NERSC network—servers at NCEM cannot initiate the socket connections.

Workflow from the User's Perspective

Many users, particularly those without experience in HPC, may find the prospect of initiating a streaming job on a supercomputing cluster to be daunting. To address this, we extended the functionalities of *Distiller* (Harris & Genova, 2023). Prior to this work, *Distiller* served as a web portal primarily for cataloging data sets, tracking metadata, and

initiating processing jobs at NERSC. Our enhancements allow users to initiate a streaming compute job through the *Distiller* web interface. Supplementary Video 2 demonstrates starting a session using the *Distiller* interface and subsequently collecting several acquisitions.

With a single mouse click in *Distiller*, the necessary connections between NCEM and NERSC are automatically established. NCEM has access to NERSC's "realtime" queue, so the nodes that make these connections are provided to the user within 1 min. This enables users to focus on their experiments while the data are seamlessly streamed to NERSC. As a result, datasets are rapidly available for further analysis, eliminating user distraction and delays associated with manually starting a separate job for each dataset.

The user monitors the progress of their streaming session and initiates data analysis notebooks using NERSC's *Jupyter* ecosystem directly from the *Distiller* web interface (see Supplementary Video 3), which is enabled by NERSC's Superfacility API (Thomas et al., 2017; Enders et al., 2020; Henderson et al., 2020; Parkinson et al, 2020; Thomas & Cholia, 2021). Integration of data acquisition, transfer, and analysis into a unified workflow enhances user productivity and enables more complex, data-intensive experiments. The streaming capability has been utilized on the TEAM 0.5 microscope for about 8 months providing streamlined data transfer and analysis for real user experiments. The code for *Distiller* is publicly accessible and can be found in Harris & Genova (2023).

From the perspective of user authentication, *Distiller* uses a NERSC collaboration account to initiate compute jobs at NERSC through the Superfacility API (Enders et al., 2020; Bard et al., 2022; Welborn et al., 2024). Electron-counted data are saved by this account into the NCEM project directory on NERSC's community file system. Each user of the 4D Camera has access to this directory through their own NERSC account. In the case that the user would like to protect their data, they can elect to move it into their own account. For more information on NERSC's authentication and Unix File Permissions policies, see the documentation at https://docs.nersc.gov/.

Microscope Stability Experiment and Workflow Comparison

Stability Experiment

In order to explore the capabilities enabled by the streaming approach, we conducted a real experiment that mimics a typical high-throughput microscopy workflow—the collection of data at regular time intervals for an extended period, hereafter referred to as a multi-scan experiment. The goals were threefold: first, to generate a large volume of data that would challenge the streaming system's capabilities; second, to quantify the microscope's stability over time; and third, to show that the system can produce many high-quality 4D-STEM datasets amenable to advanced analytical techniques, such as ptychography.

The experiment was performed on the aberration-corrected TEAM 0.5 outfitted with the all piezo-electric TEAM Stage. This stage offers exceptional stability, with a nominal drift rate of 2 pm s^{-1}, and allows for tilting up to ± 180° within the 2.5 mm pole piece gap (Ercius et al., 2012). The microscope was operated at an accelerating voltage of 300 keV, a convergence angle of 17.1 mrad, a sample tilt of 0°, a probe current of 20 pA, and a probe step size of 0.36 Å. A standard sample made of gold nanoparticles with approximate diameter of 5–10 nm was prepared by chemical vapor deposition (CVD) of gold onto an ultrathin carbon substrate. Using an automated data collection script, we acquired 60 4D-STEM datasets at 55-s intervals (total duration of 55 min), each with dimensions of 512 × 512 × 576 × 576. Each dataset consists of 173 GB of raw data (262,144 diffraction space frames), culminating in a total data volume of 10.4 TB streamed to NERSC. Each dataset was successfully acquired, transmitted to NERSC, reduced by electron counting, and stored for further analysis. The total data volume was reduced from 10.4 TB down to a more manageable size of 125 GB through counting. At NERSC, we use four Perlmutter CPU nodes for real-time processing for a total of four node hours for this experiment. For *in situ* experiments requiring a faster cadence, the user can reduce the real-space dimension so that processing occurs within seconds of acquisition, as we showed in our technical work (Welborn et al., 2024).

The large number of high-quality 4D-STEM scans acquired during this hour-long experiment provides an opportunity to measure changes in the microscope using advanced techniques such as parallax or tilt-corrected bright field and ptychography (Varnavides et al., 2023; Yu et al., 2024). Here, we used the `ParallaxReconstruction` and `SingleslicePtychographicReconstruction` classes in py4DSTEM version 0.14.3 (Savitzky et al., 2021; Varnavides et al., 2023) to perform the reconstructions on each dataset in the time series. The reconstructions were executed on eight 40 GB NVIDIA A100 GPUs (two Perlmutter GPU nodes), with the workload evenly distributed among the GPUs, allowing for simultaneous reconstruction of eight datasets at a time (see our accompanying data repository in https://github.com/swelborn/welborn-microscopy-streaming-paper-with-code for the reconstruction settings along with an example *Jupyter* notebook for one of the acquisitions). Representative (a) parallax and (b) ptychography reconstructions at the beginning (i), middle (ii), and end (iii) of the series are shown in Figure 3. These reconstructions indicate that atomic resolution is maintained throughout the experiment, owing in part to the exceptional stability of the TEAM 0.5 microscope and stage—we made no adjustments to the microscope during the experiment.

We quantify the microscope's stability by inspecting the estimated microscope parameters from the parallax and ptychography results. In a parallax reconstruction, virtual images from different positions in the bright field disk are aligned through cross-correlation. The image shifts are imparted on these virtual images based on the gradient of the aberration surface of the incoming beam and the rotation between real and reciprocal space in the microscope setup. By fitting the aberration profile of these shifts and rotations, we can estimate changes in low-order aberrations during the course of the experiments. Figure 4a shows the estimated change in defocus over the full hour of data acquisition, amounting to a drift rate of 0.5 pm s^{-1}, which is either due to stage or lens drift. The defocus drift value is not typically measured due to the projection nature of the STEM.

We also expect other aberrations to change during the course of the experiment due to lens drift (Schramm et al., 2012), and we can estimate the probe astigmatism (A1) in X and Y for all 60 datasets (Fig. 4b) based on the $A1_X$ and $A1_Y$ determined from the probe estimate over time. Both astigmatism directions have a drift rate of approximately 0.2

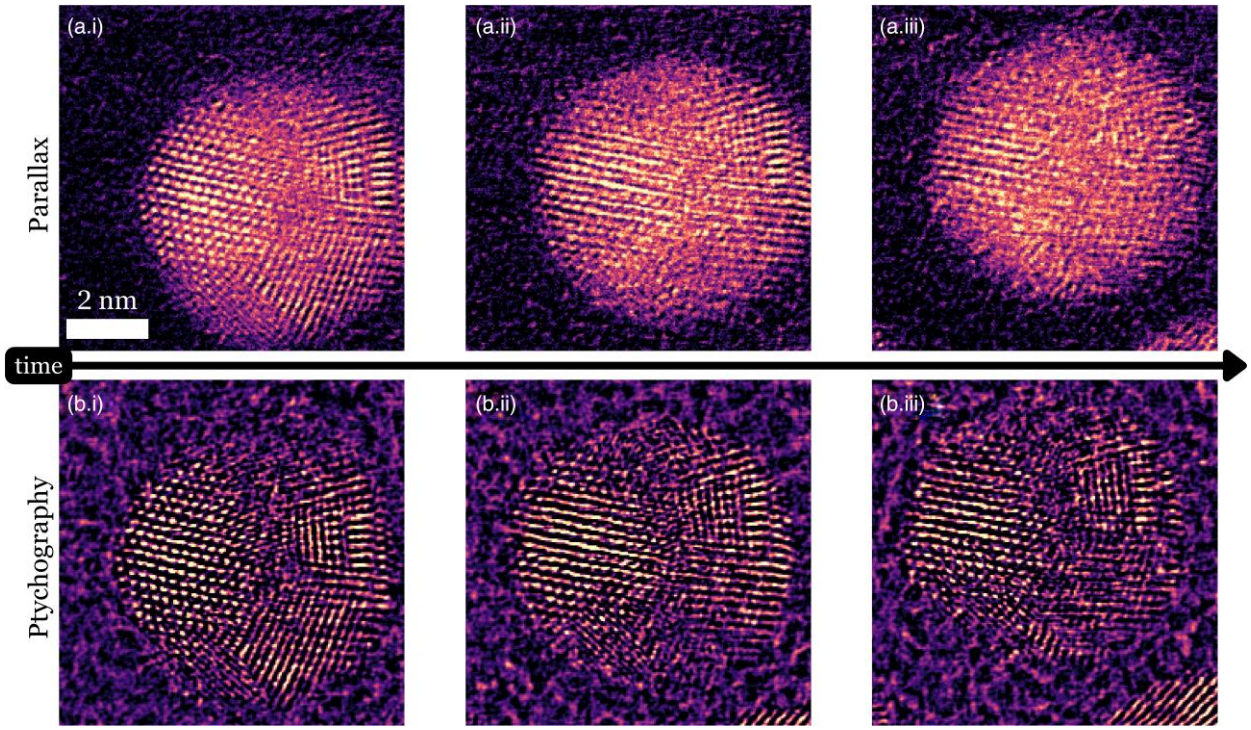

Fig. 3. Reconstructions of the same gold nanoparticle using (**a**) parallax and (**b**) ptychography over the course of the nearly hour-long experiment. (i), (ii), and (iii) display reconstructions from the experiment's start, middle, and end.

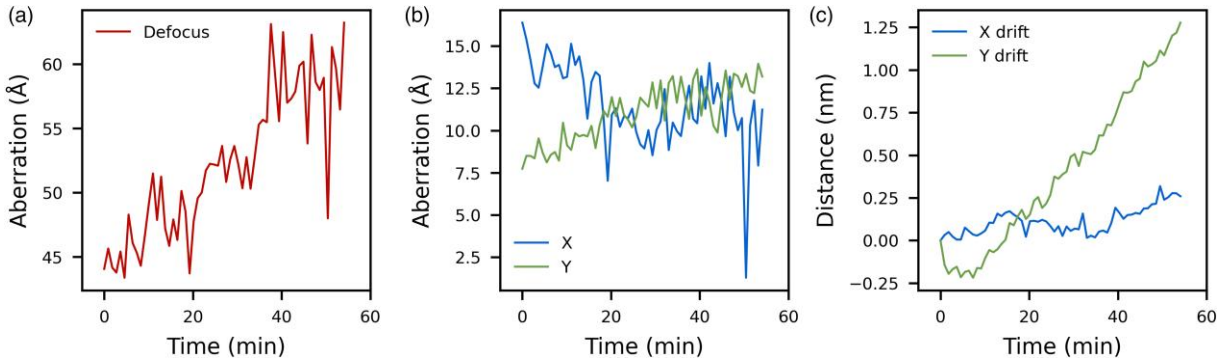

Fig. 4. Fitted parameters from reconstructions in Figure 3. (**a**) Defocus (C_1) drift of the TEAM 0.5 during the experiment, 0.3 pm s^{-1}. (**b**) Probe astigmatism (X and Y) drift throughout the experiment, both drifting at about 0.2 pm s^{-1}. (**a**) and (**b**) were both fit using the data in Figure 3a. (**c**) Lateral drift of the sample, fit with cross-correlation using the first ptychography reconstruction as the basis ($x = 0$, $y = 0$).

pm s^{-1}. Iterative electron ptychography can be used to solve for the object as well as the probe from a 4D-STEM dataset and produces a high-resolution and high signal-to-noise reconstruction (Varnavides et al., 2023). We further quantified the lateral drift of the sample by employing cross-correlation techniques on the high-resolution ptychographic reconstructions. The sample exhibited a movement of approximately 1.3 nm in the positive Y direction and around 0.3 nm in the positive X direction (Fig. 4c). These shifts are well within the published stability limits of the microscope stage, which allows for a maximum drift of 6.6 nm over the 55-min experiment duration, as calculated from the stage's drift rate of 2 pm s^{-1}.

Workflow Comparison

In our recent work (Welborn et al., 2024), we established that streaming a dataset from NCEM to NERSC is 5- to 14-fold faster than the conventional file transfer workflow in terms of raw throughput of raw detector data without counting electron events (i.e., the electron beam was off). Here, we will compare these two workflows through an analysis of historical data from four real user experiments with electron events, as illustrated by the timelines in Figure 5. File transfer experiments exhibit extended durations due to concurrent dataset transmissions and manual user interactions with *Distiller*, which introduce delays. Conversely, streaming maintains consistent and reliable transfer times, making data immediately accessible at NERSC postacquisition.

To construct these plots, we determined the last-modified timestamps for two key files created on the NERSC file system for each acquisition. The "start time" is marked by the timestamp of the simultaneously acquired HAADF-STEM image file uploaded to NERSC immediately at the end of each acquisition, and the "end time" by the timestamp of the electron-counted data file. This pair forms one of the

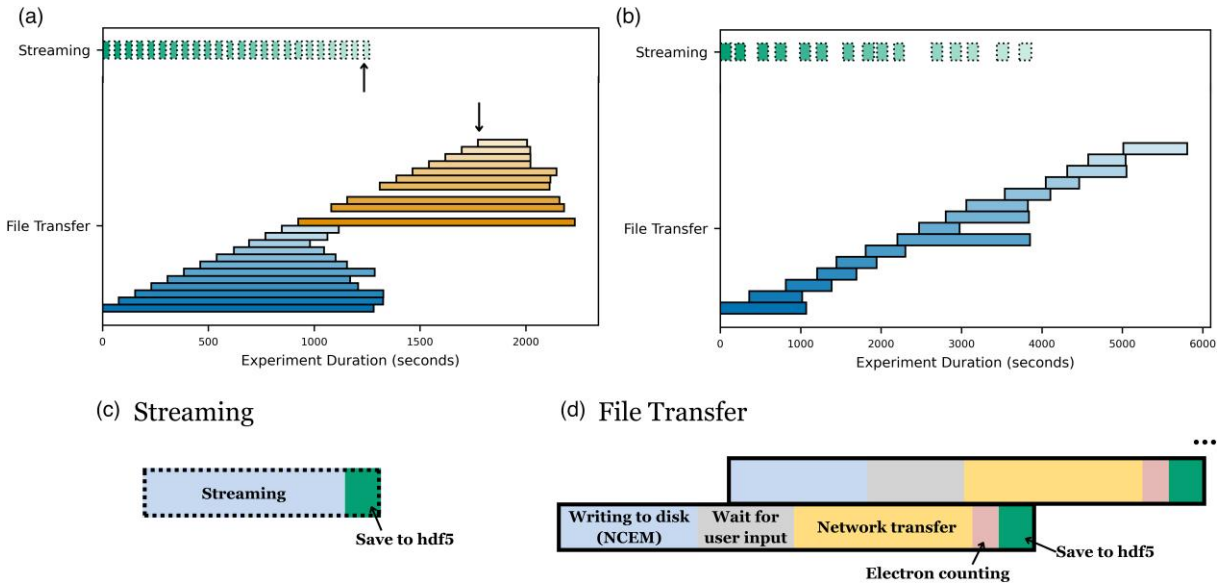

Fig. 5. Timeline diagrams of the streaming workflow compared to the file transfer workflow for two different types of experiments: (**a**) multi-scan, where data are automatically acquired at regular intervals similar to the stability experiment; and (**b**) 4D-STEM tomography, where data are collected at semi-regular intervals, but adjustments must be made to the microscope between acquisitions. The left side of each horizontal bar represents the acquisition start time, and the right side indicates the time when the electron-counted data are available at NERSC. (**c**) and (**d**) qualitatively indicate the serial steps taken within each of these bars for streaming and file transfer, respectively. The arrows in (**a**) represent the 24th acquisition in both workflows.

horizontal bars in Figure 5. It is important to note that these timestamps are synchronized using the same clock, avoiding timing discrepancies across different devices in the distributed workflow environment. For clarity, we do not detail the intermediate steps between the start and finish in Figures 5a–5b, instead displaying representative timeline snippets in Figures 5c–5d. For the file transfer workflow (Fig. 5d), there are five serial steps: writing the data to the local drive at NCEM (light blue), waiting for user input to initiate network transfer (grey), the network transfer to NERSC file system (yellow), followed by electron counting (pink), and finally saving to hdf5 (green). Conversely, there are only two steps in the streaming workflow: streaming (light blue), followed by saving to hdf5 (green).

In Figure 5a, we compare the file transfer and streaming workflows for multi-scan experiments similar to the stability experiment described above. Each acquisition amounted to 173 GB of raw data, with data dimensions of $512 \times 512 \times 576 \times 576$. During the experiment using the file transfer workflow, the user allowed a batch of acquisitions to accumulate on the NCEM file system and then initiated many NERSC transfer jobs using *Distiller* in reverse-acquisition order. This results in a pyramid-shaped timeline for each batch, since the most recent acquisition was transferred to NERSC first. In Figure 5a, two batches are shown: the first starting with the dark blue bar and ending with the light blue bar, and the second starting with the dark orange bar and ending with the light orange bar. Notably, the second batch exhibits gaps indicating missing acquisitions. These omissions could either be deliberate, perhaps due to adjustments in the microscope setup causing concerns with these acquisitions, or unintentional due to transfer failures. In either case, this underscores the disadvantages of having a human in the loop for repetitive file transfer tasks, as real-time decision-making distracts from the ongoing experiment.

The delay between acquisition and data availability is significantly longer in the file transfer workflow compared to streaming. For instance, in the first batch, the user waited 20 min for the initial acquisition to be available (dark blue bar in Fig. 5a). Even the last acquisition in the first batch, one of the shortest timeline bars in the series, required about 270 s to become accessible—almost an order of magnitude slower than the consistent 30-s time to processed data observed in the streaming workflow. Our previous work showed that the average file transfer duration for similar sized data sets is approximately 139 s (refer to the Results section of reference Welborn et al., 2024). However, that analysis did not account for additional overhead found in real experiments, such as simultaneous data transfer and acquisition, Perlmutter queue times, and user interactions needed to initiate transfers in *Distiller*. Together, these delays amounted to doubling the waiting period for the microscope user. Conversely, the streaming acquisitions, represented in teal, were available approximately 30 s after each acquisition as no human interaction is required between acquisitions, and simultaneous data transfer and acquisition do not occur. The arrows in Figure 5a point to the last (24th) acquisition in both series, including the omitted file transfer acquisitions, indicating the streaming workflow enables collection of data at a faster rate.

In Figure 5b, we compare the workflows for a 4D-STEM tomography experiment, where a user spends time between acquisitions to align the sample and microscope at a set of rotation angles. Each acquisition amounted to 695 GB of raw data, with data dimensions of $1{,}024 \times 1{,}024 \times 576 \times 576$. Here, the user was able to tilt, center, and focus the object in the field of view faster than the file transfer pipeline was able to produce processed data. The user thus had to wait for the NERSC process to complete before acquiring a new scan. Further, the user was required to initiate file transfers (disrupting their focus on the experiment) and monitor the file transfer process during the experiment to avoid overtaxing the system. Conversely, the streaming workflow (initiated with one interaction at the start of the experiment) produced finalized data before the next scan was initiated indicating processing time was less than microscope operation time.

In both cases, it is clear that more data can be acquired in a shorter amount of time with better consistency. There is also additional benefit in removing several extra steps from the experimental workflow, especially the need for users to initiate processing jobs.

Conclusions

In this work, we demonstrate the advantages of a streaming-based data transfer workflow over traditional file-based workflows, which often suffer from performance bottlenecks due to disk I/O operations. By bypassing local and remote disk I/O and transferring data directly over the network to a HPC center, our pipeline enables on-the-fly processing on remote hardware with better capabilities. This streaming pipeline seamlessly connects a high frame rate direct electron detector (the 4D Camera) to an HPC center (NERSC).

The pipeline's capabilities were demonstrated through an hour-long experiment, where 60 4D-STEM datasets totaling over 10 TB of raw data were acquired, streamed, and electron-counted in real time at NERSC, resulting in a compressed data size of 125 GB. This experiment not only tested the streaming workflow's ability to handle large volumes of data but also evaluated the entire system's capacity to produce large numbers of high-quality datasets suitable for advanced analyses such as parallax and ptychography.

A key benefit of our streaming approach, beyond the already-established increase in raw throughput (Welborn et al., 2024), is the significant reduction in human interaction required. Our comparative analysis of historical data from real user experiments reveals that automating the data transfer process increases throughput, minimizes the potential for human error, and eliminates the overhead associated with manual interactions.

Furthermore, the user focused design of our solution abstracts away the complexities of HPC, allowing researchers to focus on scientific inquiry rather than the intricacies of computation. This streaming system is integrated into the *Distiller* web frontend, simplifying the workflow. The system is in daily use on the TEAM 0.5 microscope at NCEM.

This work represents an important step forward in the integration of HPC resources with EOS facilities, addressing critical challenges in data management and computational efficiency. It serves as a model for similar integrations in other data-intensive scientific domains, having implications that extend beyond the immediate context of one electron microscopy detector. Future work will focus on expanding its applicability to other experimental setups and analytical techniques beyond electron counting.

Availability of Data and Materials

The authors have declared that no datasets apply for this piece.

Supplementary Material

To view supplementary material for this article, please visit http://academic.oup.com/mam/article-lookup/doi/10.1093/mam/ozae109#supplementary-data.

Acknowledgments

Work at the Molecular Foundry was supported by the Office of Science, Office of Basic Energy Sciences, of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231. This research used resources of the National Energy Research Scientific Computing Center (NERSC), a U.S. Department of Energy Office of Science User Facility located at Lawrence Berkeley National Laboratory, operated under Contract No. DE-AC02-05CH11231 using NERSC awards BES-ERCAP0024753 and BES-ERCAP0024754. This work was partially funded by the US Department of Energy in the program "Electron Distillery 2.0: Massive Electron Microscopy Data to Useful Information with AI/ML." We would like to thank A Saha and A Bhalla-Levine for compiling experimental information for use in Figure 5.

Conflict of Interest

No competing interest is declared.

References

Avery P, Harris C, Ercius P, Genova A, Hanwell MD & Zhao Z (2023). Openchemistry/stempy: Stempy 3.3.3.

Bard D, Snavely C, Gerhardt L, Lee J, Totzke R, Antypas K, Arndt W, Blaschke J, Byna S, Cheema R, Cholia S, Day M, Enders B, Gaur A, Greiner A, Groves T, Kiran M, Koziol Q, Lehman T, Rowland K, Samuel C, Selvarajan A, Sim A, Skinner D, Stephey L, Thomas R & Torok G (2022). LBNL superfacility project report, Research Org.: Lawrence Berkeley National Lab. (LBNL), Berkeley, CA (United States).

Battaglia M, Contarato D, Denes P & Giubilato P (2009). Cluster imaging with a direct detection CMOS pixel sensor in transmission electron microscopy. *Nucl Instrum Methods Phys Res Sec A: Accel Spectrom Detect Assoc Equip* **608**(2), 363–365. https://doi.org/10.1016/j.nima.2009.07.017

Ben-Moshe A, Da Silva A, Müller A, Abu-Odeh A, Harrison P, Waelder J, Niroui F, Ophus C, Minor AM, Asta M, Theis W, Ercius P & Alivisatos AP (2021). The chain of chirality transfer in tellurium nanocrystals. *Science* **372**(6543), 729–733. https://doi.org/10.1126/science.abf9645

Cao MC, Han Y, Chen Z, Jiang Y, Nguyen KX, Turgut E, Fuchs GD & Muller DA (2018). Theory and practice of electron diffraction from single atoms and extended objects using an EMPAD. *Microscopy* **67**(suppl_1), i150–i161. https://doi.org/10.1093/jmicro/dfx123

Chatterjee D, Wei J, kvit A, Bammes B, Levin B, Bilhorn R & Voyles P (2021). An ultrafast direct electron camera for 4D stem. *Microsc Microanal* **27**(S1), 1004–1006. https://doi.org/10.1017/S1431927621003809

Dekkers N & De Lang H (1974). Differential phase contrast in a stem. *Optik* **41**(4), 452–456. Available at https://pascal-francis.inist.fr/vibad/index.phpaction=getRecordDetail&idt=PASCAL7530074894

Doeleman SS, Barrett J, Blackburn L, Bouman KL, Broderick AE, Chaves R, Fish VL, Fitzpatrick G, Freeman M, Fuentes A, Gómez JL, Haworth K, Houston J, Issaoun S, Johnson MD, Kettenis M, Loinard L, Nagar N, Narayanan G, Oppenheimer A, Palumbo DCM, Patel N, Pesce DW, Raymond AW, Roelofs F, Srinivasan R, Tiede P, Weintroub J & Wielgus M (2023). Reference array and design consideration for the next-generation event horizon telescope. *Galaxies* **11**(5), 107. https://www.mdpi.com/2075-4434/11/5/107. https://doi.org/10.3390/galaxies11050107

Enders B, Bard D, Snavely C, Gerhardt L, Lee J, Totzke B, Antypas K, Byna S, Cheema R, Cholia S, Day M, Gaur A, Greiner A, Groves T, Kiran M, Koziol Q, Rowland K, Samuel C, Selvarajan A, Sim A, Skinner D, Thomas R & Torok G (2020). Cross-facility science with the superfacility project at LBNL. In *2020 IEEE/ACM 2nd Annual Workshop on Extreme-scale Experiment-in-the-Loop Computing (XLOOP)*, pp. 1–7. IEEE.

Enders B & Thibault P (2016). A computational framework for ptychographic reconstructions. *Proc R Soc A Math Phys Eng Sci* **472**(2196), 20160640. https://doi.org/10.1098/rspa.2016.0640

Ercius P, Alaidi O, Rames MJ & Ren G (2015). Electron tomography: A three-dimensional analytic tool for hard and soft materials research. *Adv Mater* **27**(38), 5638–5663. https://doi.org/10.1002/adma.v27.38

Ercius P, Boese M, Duden T & Dahmen U (2012). Operation of TEAM I in a user environment at NCEM. *Microsc Microanal* **18**(4), 676–683. https://doi.org/10.1017/S1431927612001225

Ercius P, Johnson IJ, Pelz P, Savitzky BH, Hughes L, Brown HG, Zeltmann SE, Hsu S-L, Pedroso CCS, Cohen BE, Ramesh R, Paul D, Joseph JM, Stezelberger T, Czarnik C, Lent M, Fong E, Ciston J, Scott MC, Ophus C, Minor AM & Denes P (2024). The 4D camera: An 87 kHz direct electron detector for scanning/transmission electron microscopy. *Microsc Microanal*. In press. https://doi.org/10.1093/mam/ozae086

Harris C & Genova A (2023). Distiller. https://github.com/OpenChemistry/distiller.

Henderson ML, Krinsman W, Cholia S, Thomas R & Slaton T (2020). Accelerating experimental science using Jupyter and NERSC HPC. In *Tools and Techniques for High Performance Computing: Selected Workshops, HUST, SE-HER and WIHPC, Held in Conjunction with SC 2019, Denver, CO, USA, November 17–18, 2019, Revised Selected Papers 6*, pp. 145–163. Springer.

Levin BDA (2021). Direct detectors and their applications in electron microscopy for materials science. *J Phys Mater* **4**(4), 042005. https://doi.org/10.1088/2515-7639/ac0ff9

Miao J, Ercius P & Billinge SJ (2016). Atomic electron tomography: 3D structures without crystals. *Science* **353**(6306), aaf2157. https://doi.org/10.1126/science.aaf2157

Nellist P, McCallum B & Rodenburg JM (1995). Resolution beyond the 'information limit' in transmission electron microscopy. *Nature* **374**(6523), 630–632. https://doi.org/10.1038/374630a0

Ophus C (2019). Four-dimensional scanning transmission electron microscopy (4D-STEM): From scanning nanodiffraction to ptychography and beyond. *Microsc Microanal* **25**(3), 563–582. https://doi.org/10.1017/S1431927619000497

Ophus C (2023). Quantitative scanning transmission electron microscopy for materials science: Imaging, diffraction, spectroscopy, and tomography. *Annu Rev Mater Res* **53**(1), 105–141. https://doi.org/10.1146/matsci.2023.53.issue-1

Parkinson DY, Krishnan H, Ushizima D, Henderson M & Cholia S (2020). Interactive parallel workflows for synchrotron tomography. In *2020 IEEE/ACM 2nd Annual Workshop on Extreme-scale Experiment-in-the-Loop Computing (XLOOP)*, pp. 29–34. IEEE.

Pattison AJ, Pedroso CCS, Cohen BE, Ondry JC, Alivisatos AP, Theis W & Ercius P (2023). Advanced techniques in automated high-resolution scanning transmission electron microscopy. *Nanotechnology* **35**(1), 015710. https://doi.org/10.1088/1361-6528/acf938

Pelz PM, Brown HG, Stonemeyer S, Findlay SD, Zettl A, Ercius P, Zhang Y, Ciston J, Scott MC & Ophus C (2021*b*). Phase-contrast imaging of multiply-scattering extended objects at atomic resolution by reconstruction of the scattering matrix. *Phys Rev Res* **3**(2), 023159. https://doi.org/10.1103/PhysRevResearch.3.023159

Pelz P, Ercius P, Ophus C, Johnson I & Scott M (2021*a*). Real-time interactive ptychography from electron event representation data. *Microsc Microanal* **27**(S1), 188–189. https://doi.org/10.1017/S1431927621001288

Pelz PM, Griffin S, Stonemeyer S, Popple D, Devyldere H, Ercius P, Zettl A, Scott MC & Ophus C (2023). Solving complex nanostructures with ptychographic atomic electron tomography. *Nat Commun* **14**(1), 7906. https://doi.org/10.1038/s41467-023-43634-z

Rao R (2020). Synchrotrons face a data deluge. *Phys Today* **2020**(2), 0925a. https://doi.org/10.1063/PT.6.2.20200925a

Ribet SM, Varnavides G, Pedroso CCS, Cohen BE, Ercius P, Scott MC & Ophus C (2024). Uncovering the three-dimensional structure of upconverting core-shell nanoparticles with multislice electron ptychography. *Appl Phys Lett* **124**(24), 240601. https://doi.org/10.1063/5.0206814

Savitzky BH, Zeltmann SE, Hughes LA, Brown HG, Zhao S, Pelz PM, Pekin TC, Barnard ES, Donohue J, Rangel DaCosta L, Kennedy E, Xie Y, Janish MT, Schneider MM, Herring P, Gopal C, Anapolsky A, Dhall R, Bustillo KC, Ercius P, Scott MC, Ciston J, Minor AM & Ophus C (2021). py4dstem: A software package for four-dimensional scanning transmission electron microscopy data analysis. *Microsc Microanal* **27**(4), 712–743. https://doi.org/10.1017/S1431927621000477

Schramm SM, van der Molen SJ & Tromp RM (2012). Intrinsic instability of aberration-corrected electron microscopes. *Phys Rev Lett***109**(16), 163901. https://doi.org/10.1103/PhysRevLett.109.163901

Shibata N, Findlay SD, Kohno Y, Sawada H, Kondo Y & Ikuhara Y (2012). Differential phase-contrast microscopy at atomic resolution. *Nat Phys* **8**(8), 611–615. https://doi.org/10.1038/nphys2337

Spurgeon SR, Ophus C, Jones L, Petford-Long A, Kalinin SV, Olszta MJ, Dunin-Borkowski RE, Salmon N, Hattar K, Yang W-CD, Sharma R, Du Y, Chiaramonti A, Zheng H, Buck EC, Kovarik L, Penn RL, Li D, Zhang X, Murayama M & Taheri ML (2021). Towards data-driven next-generation transmission electron microscopy. *Nat Mater* **20**(3), 274–279. https://doi.org/10.1038/s41563-020-00833-z

Stevens A, Yang H, Hao W, Jones L, Ophus C, Nellist PD & Browning ND (2018). Subsampled stem-ptychography. *Appl Phys Lett* **113**(3), 033104. https://doi.org/10.1063/1.5040496

Thomas R, Canon S, Cholia S, Gerhardt L & Racah E (2017). Toward interactive supercomputing at NERSC with Jupyter. In *Cray User Group (CUG) Conference Proceedings*. Cray User Group (CUG).

Thomas R & Cholia S (2021). Interactive supercomputing with Jupyter. *Comput Sci Eng* **23**(2), 93–98. https://doi.org/10.1109/MCSE.2021.3059037

Varnavides G, Ribet SM, Zeltmann SE, Yu Y, Savitzky BH, Dravid VP, Scott MC & Ophus C (2023). 'Iterative phase retrieval algorithms for scanning transmission electron microscopy.' arXiv, arXiv:2309.05250, preprint: not peer reviewed.

Veseli S, Hammonds J, Henke S, Parraga H & Schwarz N (2023). Streaming data from experimental facilities to supercomputers for real-time data processing. In *Proceedings of the SC '23 Workshops of The International Conference on High Performance Computing, Network, Storage, and Analysis*, SC-W '23, pp. 2110–2117. New York, NY, USA: Association for Computing Machinery.

Waddell E & Chapman J (1979). Linear imaging of strong phase objects using asymmetrical detectors in stem. *Optik* **54**, 83–96. Available at https://pascal-francis.inist.fr/vibad/index.php?action=getRecordDetail&idt=PASCAL8030339586

Welborn SS, Enders B, Harris C, Ercius P & Bard DJ (2024). Third combined workshop on interactive and urgent supercomputing. In *High Performance Computing: ISC High Performance 2024 International Workshops*, Hamburg, Germany, May 12–16, 2024. Lecture Notes in Computer Science.

Yasin FS, Harvey TR, Chess JJ, Pierce JS & McMorran BJ (2016). Development of stem-holography. *Microsc Microanal* **22**(S3), 506–507. https://doi.org/10.1017/S143192761600338X

Yasin FS, Harvey TR, Chess JJ, Pierce JS, Ophus C, Ercius P & McMorran BJ (2018). Probing light atoms at subnanometer resolution: Realization of scanning transmission electron microscope holography. *Nano Lett* **18**(11), 7118–7123. https://doi.org/10.1021/acs.nanolett.8b03166

Yu Y, Spoth KA, Colletta M, Nguyen KX, Zeltmann SE, Zhang XS, Paraan M, Kopylov M, Dubbeldam C, Serwas D, Siems H, Muller DA & Kourkoutis LF (2024). 'Dose-efficient cryo-electron microscopy for thick samples using tilt-corrected scanning transmission electron microscopy, demonstrated on cells and single particles.' bioRxiv, https://doi.org/10.1101/2024.04.22.590491, 04 2024, preprint: not peer reviewed.

Zambon P, Vávra J, Montemurro G, Bottinelli S, Dudina A, Schnyder R, Hörmann C, Meffert M, Schulze-Briese C, Stroppa D, Lehmann N & Piazza L (2023). High-frame rate and high-count rate hybrid pixel detector for 4D stem applications. *Front Phys* **11**, 1308321. https://doi.org/10.3389/fphy.2023.1308321

TESCAN
PERFORMANCE IN NANOSPACE

TESCAN AMBER X 2

PLASMA FIB-SEM REDEFINED

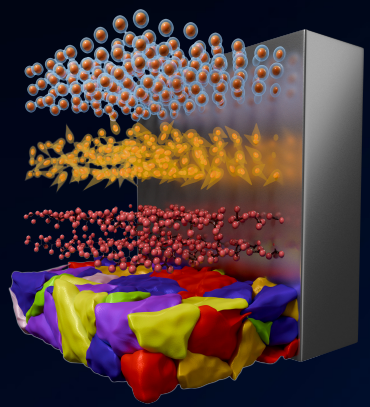

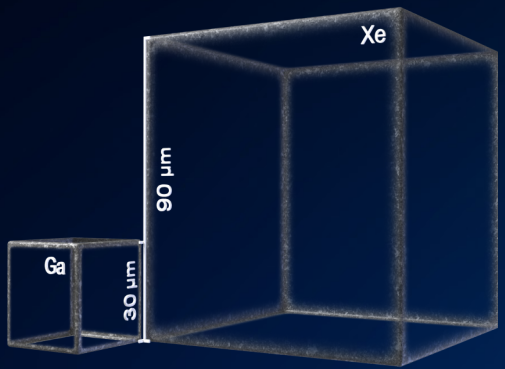

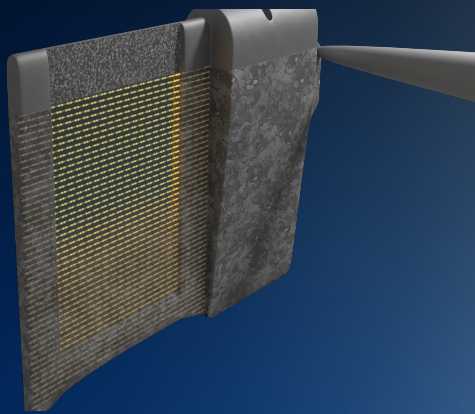

UTILITY REDEFINED

SPEED REDEFINED

PRECISION REDEFINED

info.tescan.com