Ptychography under jittery XFEL beams with MPI

Delight Nyanhete

Mentors: Stefano Marchesini, Yuan Ni

PTYCHOGRAPHY

Ptychography is a high-resolution imaging technique that combines coherent diffractive imaging with scanning probe microscopy. A localized illumination probe scans a sample in overlapping steps, capturing a far-field diffraction pattern at each position.

Using the redundancy in overlapping regions, iterative phase retrieval algorithms reconstruct both the amplitude and phase of the sample at nanoscale resolution

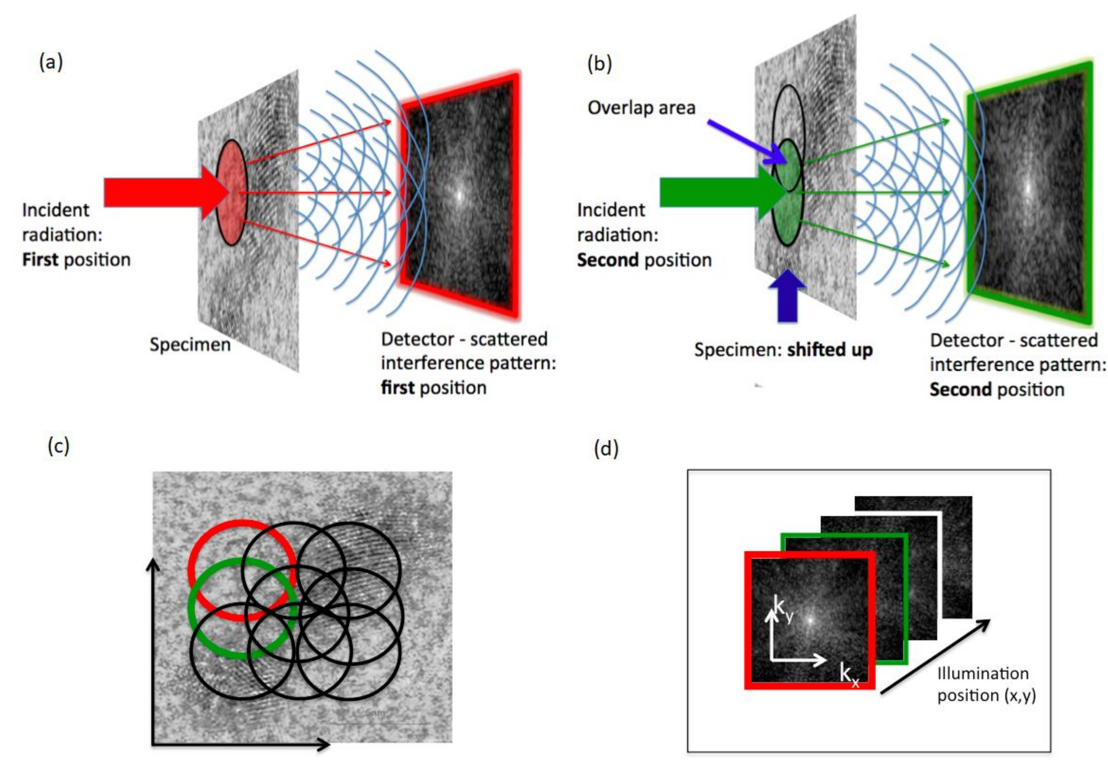

XFEL JITTER PROBLEM

XFEL delivers femtosecond pulses at 1 MHz, but every shot arrives with a slightly different tilt, so the illumination on the sample changes shot to shot. Assuming one fixed probe blurs the reconstruction and introduces artifacts.

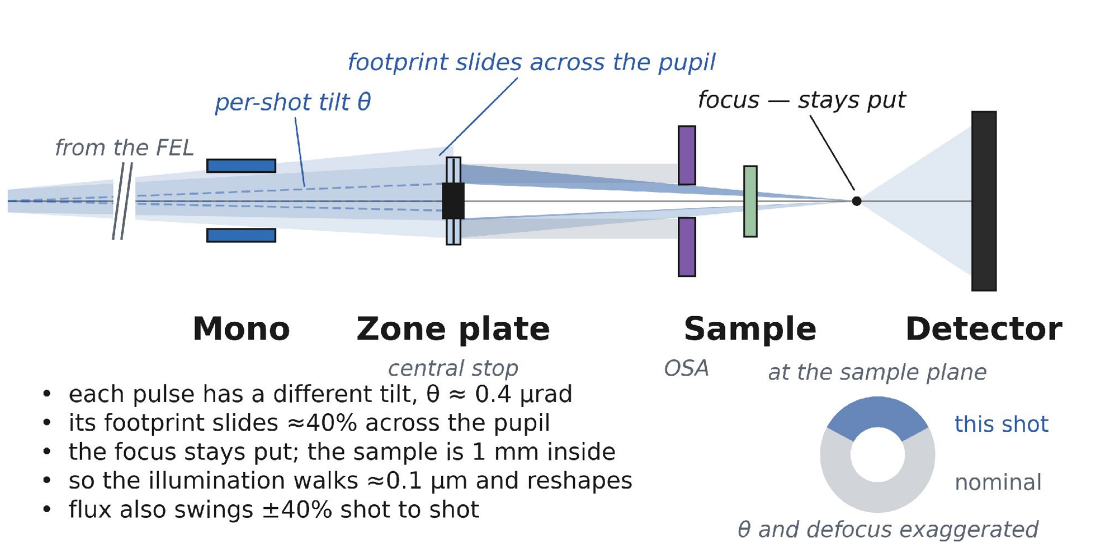

Beamline modelling and the jittery-beam simulations: Himanshu Goel, using SRW (Synchrotron Radiation Workshop; Chubar & Elleaume, EPAC-98 1998), developed at Brookhaven National Laboratory, collaborators on this project. Using Genesis FEL pulses simulated by J Tang (SLAC)

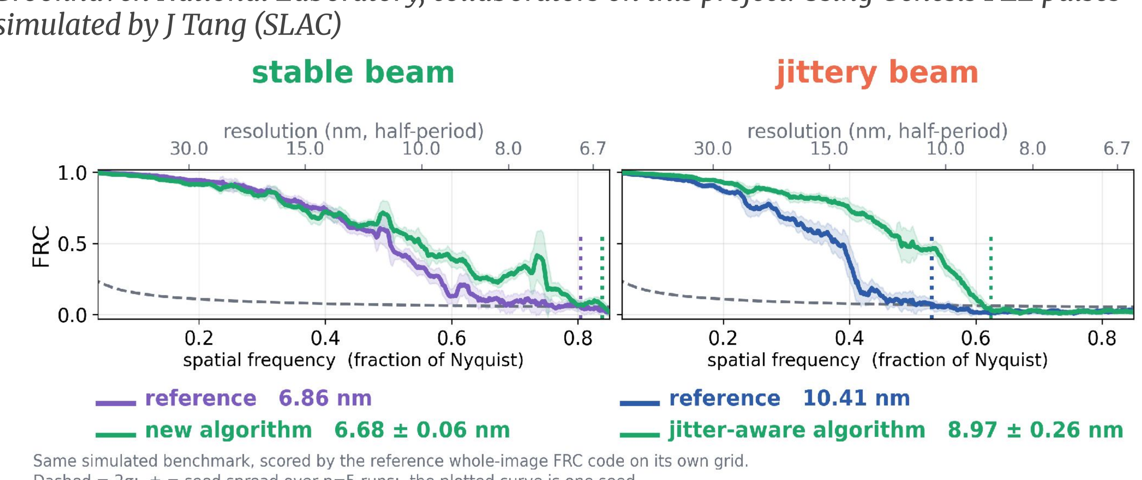

Modeling per-shot beam jitter introduces new degrees of freedom that significantly increase computation over a standard shared probe. Benchmarking on a single H200 GPU reveals a 5.6× per-iteration overhead, jumping from 0.195 s (shared probe) to 1.094 s (rank-2 per-shot probes) (See https://github.com/slac-lcls/sharpy for software)

This extra cost is driven by per-frame Gram solves and per-pixel mode updates rather than Fast Fourier Transforms (FFTs), since coherent rank-K modes are summed into a single exit wave prior to transformation. Consequently, scaling to higher ranks adds minimal marginal cost (rank-3 costs 6.1×; rank-8 costs 8.4×). Because jitter-aware fitting requires a converged initial reconstruction to start from, it operates as a second pass on top of a jitter-blind run. This pushes total end-to-end execution time closer to an order of magnitude and creating the clear computational bottleneck that demands multi-GPU parallelization

MESSAGE PASSING INTERFACE

The higher cost of jitter-aware fitting motivates a distributed implementation. AP is decomposed spatially across MPI ranks so that most computation remains local, with communication introduced only for synchronization and shared object boundaries.

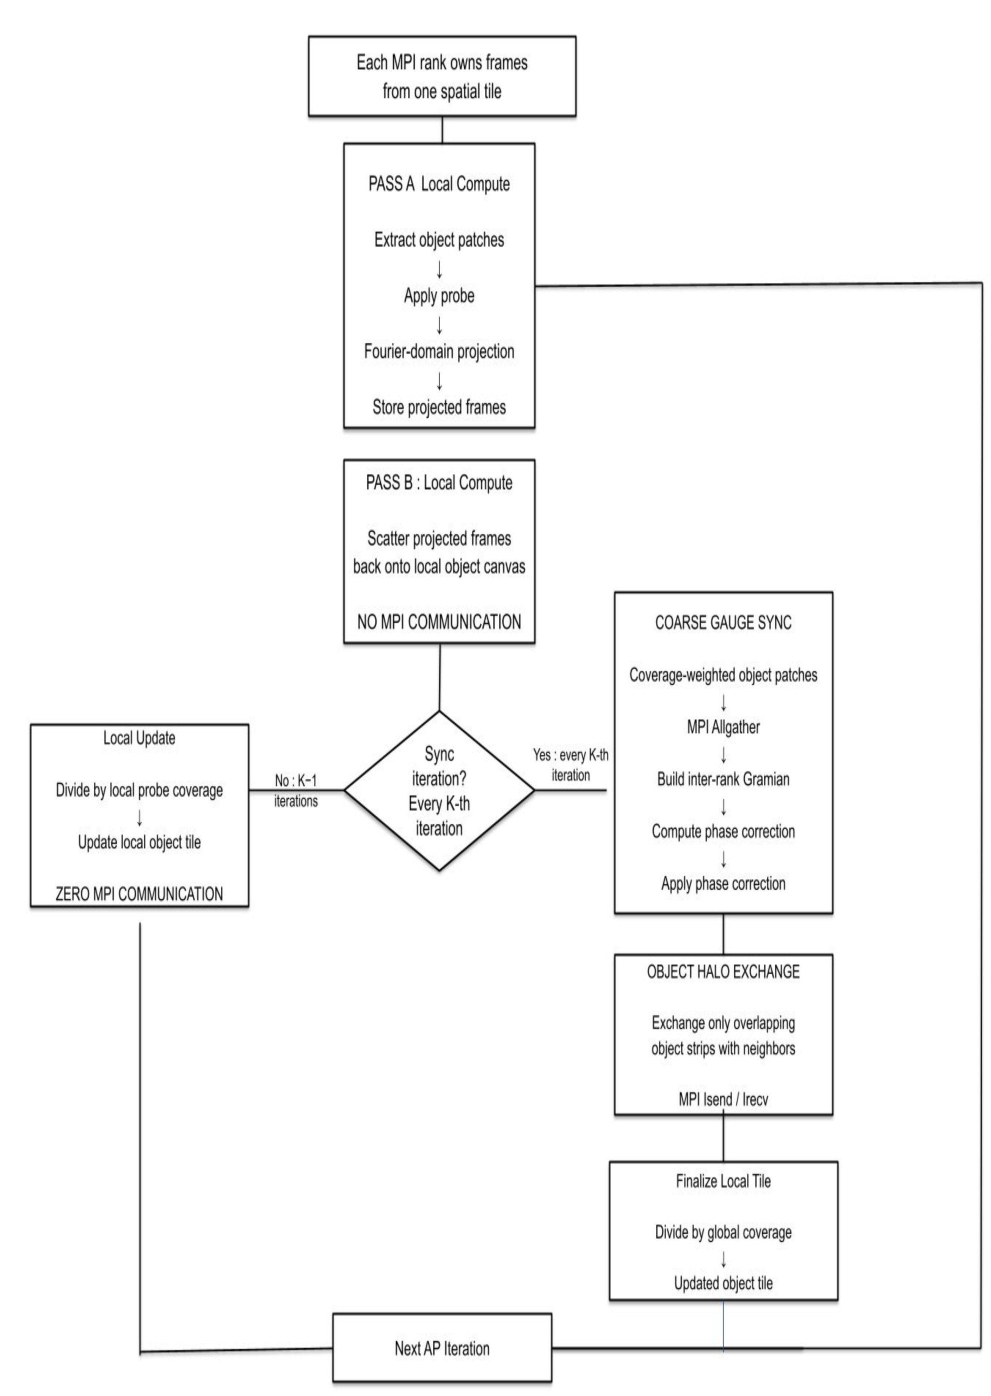

Distributing the reconstruction introduces a communication bottleneck. Full-canvas AllReduce communicates the global object, while object halo exchange restricts communication to overlapping regions between neighboring spatial tiles.

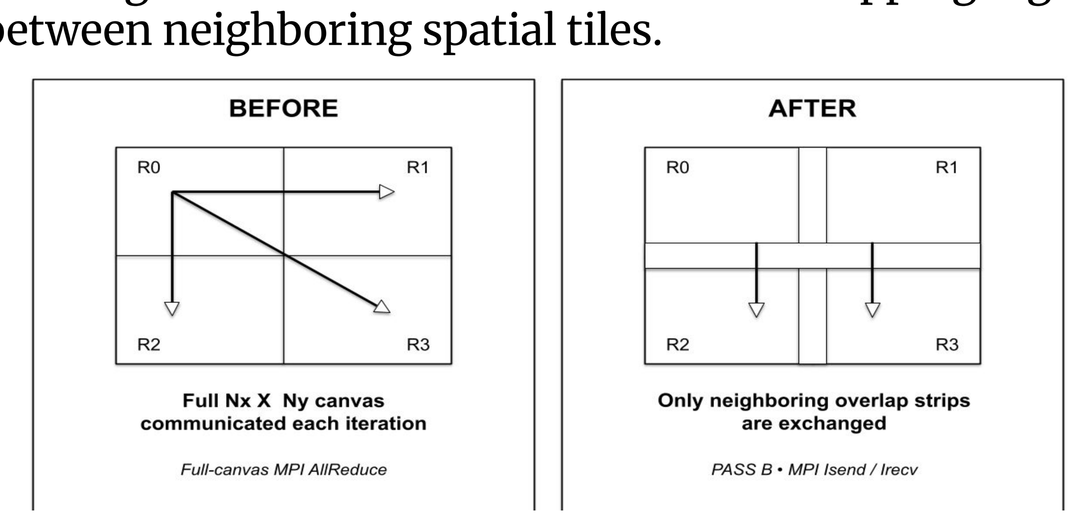

THROUGHPUT AND SCALING RESULTS

To address the 5.6x compute bottleneck of per-shot probe modeling, we parallelized the Alternating Projections algorithm across GPUs using MPI. Replacing naive global canvas transfers with localized spatial halo exchanges eliminates communication bottlenecks and enables efficient multi-node scaling

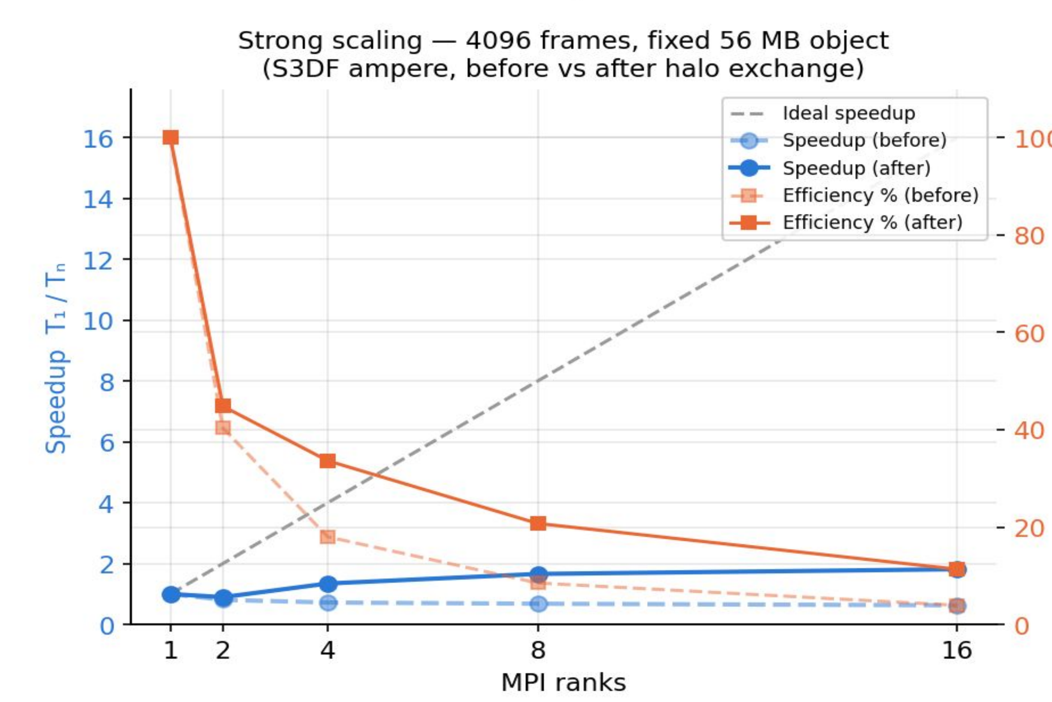

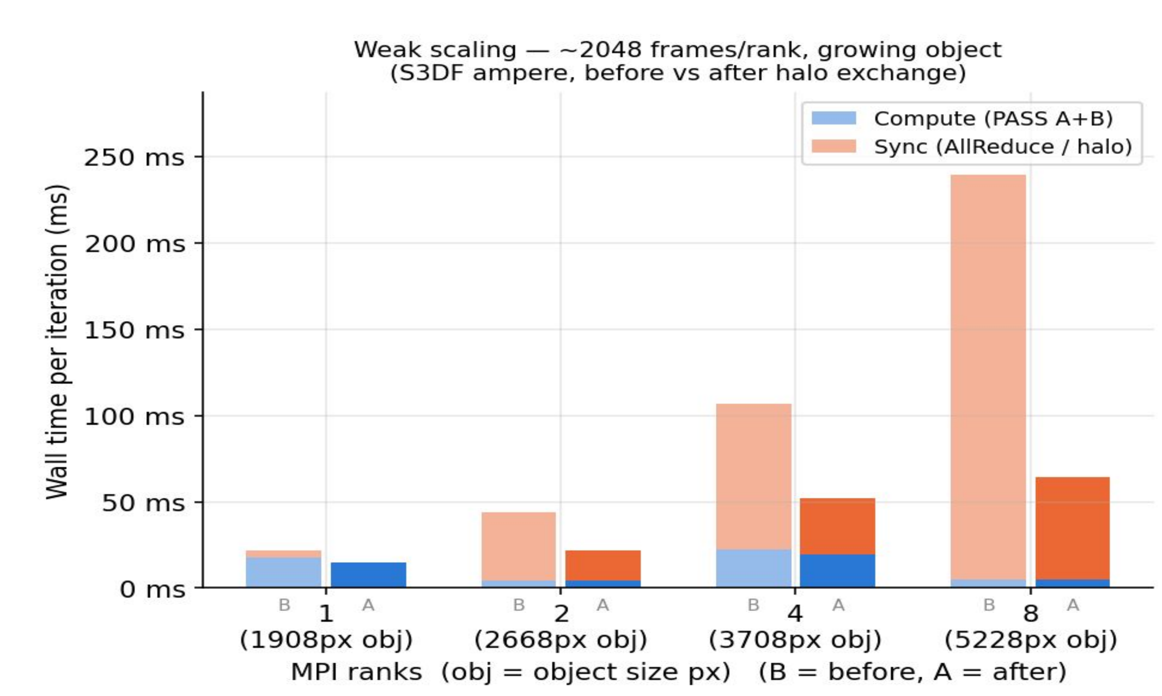

To verify algorithm stability and solver behavior using the jitter-aware formulation, reconstructions were evaluated on stable beam data under a fixed 60-second budget. Distributed execution across 16 MPI ranks significantly outperformed the single GPU, reaching 9,327 iterations and achieving lower error (NMSE = 1.77e-06) while preserving full spatial resolution

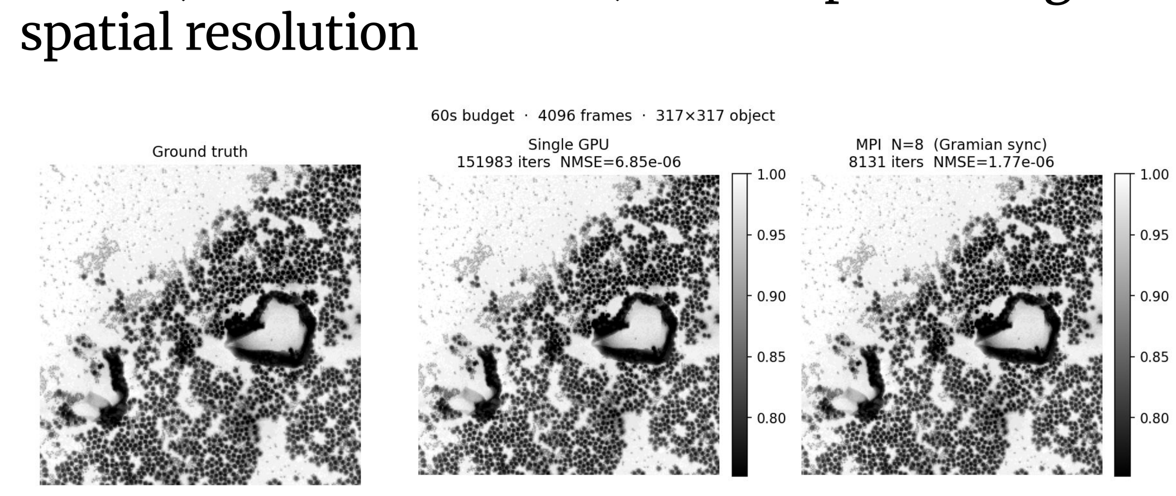

CONCLUSION

We resolved the computational bottleneck of jitter-aware ptychography by distributing the AP algorithm across multi-GPU nodes using MPI. Replacing global canvas transfers with localized spatial halo exchanges eliminated communication latency, restoring strong scaling up to 16 ranks while accelerating convergence and lowering reconstruction error. Future work includes deploying this solver on live LCLS-II

Use of the Linac Coherent Light Source (LCLS), SLAC National Accelerator Laboratory, is supported by the U.S. Department of Energy, Office of Science, Office of Basic Energy Sciences under Contract No. DE-AC02-76SF00515. Thanks to the LCLS Summer Internship Program and the SLAC ptychography group.

Many thanks to Himanshu Goel of BNL for the SRW simulations. NERSC for computing resources. J. Tang (SLAC) for GENESIS (Reiche, NIM A 429, 243 (1999) pulse simulations. See https://github.com/slac-lcls/sharpy for software