

Efficient Image Reconstruction Architecture for Neutral Atom Quantum Computing

Jonas Winklmann^{*,1}, Yian Yu^{*,2}, Xiaorang Guo¹, Korbinian Staudacher², and Martin Schulz¹

¹Chair of Computer Architecture and Parallel Systems, Technical University of Munich, Garching, Germany

²MNM-Team, Ludwig-Maximilians-Universität München, Munich, Germany

Email: {jonas.winklmann, xiaorang.guo, martin.w.j.schulz}@tum.de, yian.yu@campus.lmu.de, staudacher@nm.ifi.lmu.de

Abstract—In recent years, neutral atom quantum computers (NAQCs) have attracted a lot of attention, primarily due to their long coherence times and good scalability. One of their main drawbacks is their comparatively time-consuming control overhead, with one of the main contributing procedures being the detection of individual atoms and measurement of their states, each occurring at least once per compute cycle and requiring fluorescence imaging and subsequent image analysis.

To reduce the required time budget, we propose a highly-parallel atom-detection accelerator for tweezer-based NAQCs. Building on an existing solution, our design combines algorithm-level optimization with a field-programmable gate array (FPGA) implementation to maximize parallelism and reduce the run time of the image analysis process. Our design can analyze a 256×256 -pixel image representing a 10×10 atom array in just $115 \mu\text{s}$ on a Xilinx UltraScale+ FPGA. Compared to the original CPU baseline and our optimized CPU version, we achieve about $34.9\times$ and $6.3\times$ speedup of the reconstruction time, respectively. Moreover, this work also contributes to the ongoing efforts toward fully integrated FPGA-based control systems for NAQCs.

Index Terms—Quantum Computing, Neutral Atoms, Atom Detection, Image Reconstruction, FPGA

I. INTRODUCTION

Typical neutral atom quantum computers (NAQCs) require at least two atom detection steps, one for initial detection and one for the final readout, with possibly more to detect atom losses during sorting or, especially for fault-tolerant computing [1], for mid-circuit measurements. Most other required steps, such as gate execution and atom loading, are comparatively fast. As a result, the detection process ought to be performed rapidly, making parallel algorithms highly desirable.

While the capability to resolve single atoms through microscopic imaging is well-explained and influential in many adjacent fields [2, 3, 4, 5, 6], the ongoing transition of NAQC from lab experiment to generally useful computational platform imposes new challenges in terms of robustness and independence. Field-programmable gate array (FPGA)-based control architectures emerge as a promising solution that fulfills all these requirements, delivering a low-latency, programmable hardware required for real-time atom detection tasks.

Therefore, in this work, we propose a highly parallel image deconvolution solution for atom detection based on the

* These authors contributed equally to this work.

This work was funded by the German Federal Ministry of Education and Research (BMBF) under the funding program *Quantum Technologies - From Basic Research to Market* under contract numbers 13N16077 and 13N16087, as well as from the Munich Quantum Valley (MQV), which is supported by the Bavarian State Government with funds from the Hightech Agenda Bayern.

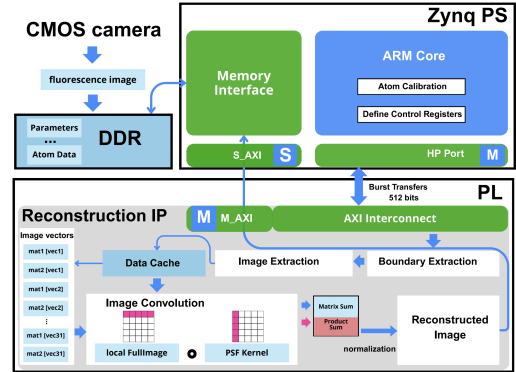


Fig. 1. Overall Architecture of the Reconstruction Accelerator. In this figure, *mat1* and *mat2* represent the PSF kernel and the atom’s image, respectively.

forementioned hardware devices. Our solution utilizes the hardware-software co-design strategy to optimize an existing atom detection method, namely the projection-based state-reconstruction algorithm presented in work [7]. Through this approach, we obtain a CPU-optimized version with high inherent parallelism and streamlined logic. Building on this, we develop an FPGA-based reconstruction accelerator that, once connected to a camera, can directly bridge the gap between image generation and FPGA-based rearrangement [8, 9] or readout procedures. Overall, with many microwave pulse generation platforms already being based on FPGAs [10, 11, 12], the vision of a fully integrated sorting and readout device is nearing completion.

II. SYSTEM ARCHITECTURE

Fig. 1 shows the overall system architecture of our FPGA-based accelerator, where the programmable logic (PL) and processing system (PS) work together to process atom positions and apply the state-reconstruction algorithm. The PS side is mainly responsible for software-driven primitive atom calibration, control register configuration, and high-level data communication. Meanwhile, the PL side implements dedicated hardware development, which provides substantial benefits in the parallel execution of the state-reconstruction process.

The entire reconstruction accelerator, a customized Intellectual Property (IP) core, adopts the dataflow design, as shown in the bottom part of Fig. 1, enabling task-level parallelism across different modules. Before execution, the DDR memory is initialized with the atom position grid,




Fig. 2. Result of the detection algorithm with a 30×30 atom array as an example. (a) A raw image of the atom array captured by a camera (simulated). (b) The output of our reconstruction accelerator. Darker pixels denote higher reconstructed brightness. (c) Thresholded boolean result using a calibrated threshold. Black denotes a detected atom.

a 31×31 PSF kernel, and the fluorescence image. The *Boundary Extraction* module first identifies the local region of interest for each atom. Following this, the *Image Extraction* module fetches the corresponding pixel data and PSF kernel through a 512-bit AXI interface. This high-bandwidth data is then decoded to 32-bit precision, and the PSF kernel is pre-processed into a projector to facilitate the subsequent *Image Convolution* stage.

Within the *image convolution* module, two parallel computational paths are executed: (1) an element-wise matrix multiplication between the local image detail and the projector kernel to derive the *product sum* (indicated by pink arrows in Fig. 1), and (2) a summation of all elements within the projector matrix to obtain the *matrix sum* (indicated by blue arrows in Fig. 1). To achieve high throughput, the matrix operations are decomposed into 31 concurrent vector processing units, each featuring internal parallelization. Furthermore, a logarithmic reduction algorithm is implemented to compute vector sums via a four-stage adder tree. This architecture reduces the computational complexity from $\mathcal{O}(n)$ to $\mathcal{O}(\log n)$, enabling the summation of 31 elements to be completed within only five clock cycles.

Ultimately, the *output aggregation* module calculates a normalized value for each convolution result from the previous step, thereby providing the normalized brightness at each atomic site.

III. EXPERIMENTS AND EVALUATION

We evaluate this work based on the quality of the reconstructed image, reconstruction run time, and FPGA resource utilization. For run time comparison, we benchmark the accelerator against the original CPU baseline [7] and our optimized CPU implementation (CPU-opt). The FPGA design is implemented on a Xilinx ZCU216 board [13] at 100 MHz.

A. Reconstructed Image

Fig. 2 shows the result of applying the reconstruction algorithm to a 30×30 atom array. After the reconstruction procedure, we obtain a so-called *emission matrix*, where its values of emissions indicate the brightness of atoms. Denoting higher emission values by darker pixels, we generate the reconstructed image of atoms shown in Fig. 2(b). The algorithm’s precision has been shown previously [14], so ensuring that our adaptation




Fig. 3. Run time comparison of the reconstruction process for various sizes of the atom array (10×10 to 40×40) among CPU-baseline, CPU-opt, and FPGA. Error bars represent the standard deviations of the run time.

performs equivalently to the original is sufficient. This is the case up to insignificant differences that can be attributed to rounding errors.

B. Run Time of Reconstruction

As illustrated in Fig. 3, we compare the reconstruction run time for various array sizes of the original algorithm (CPU-baseline), the CPU-optimized version (CPU-opt), and the proposed FPGA accelerator. The Y-axis is plotted on a logarithmic scale, with error bars representing the standard deviation of run time across 50 independent trials for each data point. The results demonstrate that our FPGA accelerator not only consistently outperforms the CPU counterparts in terms of speed but also exhibits superior stability with minimal run-time variance. Specifically, for a 10×10 array, the FPGA achieves a speedup of $34.9 \times$ ($115 \mu s$ vs. $4012 \mu s$) over the CPU-baseline and $6.3 \times$ ($115 \mu s$ vs. $730 \mu s$) over the CPU-opt.

C. Resource Utilization

Resource utilization is a critical metric for evaluating scalability. Thanks to our fixed parallelization and time-multiplexing design, resource consumption remains constant, independent of the atom array size. Based on the report of the Vivado implementation, the accelerator is highly efficient, utilizing only one-quarter of lookup tables (LUTs) and $\sim 15\%$ of flip-flops (FFs), with negligible digital signal processing (DSP) and block RAM (BRAM) usage. This compact footprint reserves significant capacity for additional logic, making the design well-suited for integration into unified quantum control systems and future HPC-quantum platforms (HPCQC) [15, 16, 17].

IV. CONCLUSION

This work presented a highly parallel atom-detection accelerator for atom detection in NAQCs. By employing a software-hardware co-design strategy, we leverage algorithm-level optimization and an FPGA-based parallel architecture to minimize reconstruction latency. Our implementation achieves an ultra-low latency of $115 \mu s$ for a 10×10 array and $1.825 \mu s$ for a 40×40 array, representing a $34.9 \times$ speedup over the CPU baseline and $6.3 \times$ over an optimized CPU version. The design is characterized by its low hardware cost and excellent scalability, demonstrating significant potential for integration into NAQCs control systems.

REFERENCES

- [1] E. Deist, Y.-H. Lu, J. Ho, M. K. Pasha, J. Zeiher, Z. Yan, and D. M. Stamper-Kurn, “Mid-circuit cavity measurement in a neutral atom array,” *Phys. Rev. Lett.*, vol. 129, p. 203602, Nov 2022. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevLett.129.203602>
- [2] N. Schlosser, G. Reymond, I. Protsenko, and P. Grangier, “Sub-poissonian loading of single atoms in a microscopic dipole trap,” *Nature*, vol. 411, pp. 1024–7, July 2001.
- [3] W. S. Bakr, J. I. Gillen, A. Peng, S. Fölling, and M. Greiner, “A quantum gas microscope for detecting single atoms in a hubbard-regime optical lattice,” *Nature*, vol. 462, no. 7269, p. 74–77, Nov. 2009. [Online]. Available: <http://dx.doi.org/10.1038/nature08482>
- [4] E. Haller, J. Hudson, A. Kelly, D. A. Cotta, B. Peaudecerf, G. D. Bruce, and S. Kuhr, “Single-atom imaging of fermions in a quantum-gas microscope,” *Nature Physics*, vol. 11, no. 9, p. 738–742, Jul. 2015. [Online]. Available: <http://dx.doi.org/10.1038/nphys3403>
- [5] J. F. Sherson, C. Weitenberg, M. Endres, M. Cheneau, I. Bloch, and S. Kuhr, “Single-atom-resolved fluorescence imaging of an atomic mott insulator,” *Nature*, vol. 467, no. 7311, p. 68–72, Aug. 2010. [Online]. Available: <http://dx.doi.org/10.1038/nature09378>
- [6] M. Morgado and S. Whitlock, “Quantum simulation and computing with rydberg-interacting qubits,” *AVS Quantum Science*, vol. 3, p. 023501, June 2021.
- [7] D. Wei, “Microscopy of spin hydrodynamics and cooperative light scattering in atomic hubbard systems,” Ph.D. dissertation, Ludwig-Maximilians-Universität München, Sept. 2023.
- [8] X. Guo, J. Winklmann, D. Stober, A. Elsharkawy, and M. Schulz, “Design of an fpga-based neutral atom rearrangement accelerator for quantum computing,” in *2025 Design, Automation & Test in Europe Conference (DATE)*, 2025, pp. 1–6.
- [9] S. Wang, W. Zhang, T. Zhang, S. Mei, Y. Wang, J. Hu, and W. Chen, “Accelerating the assembly of defect-free atomic arrays with maximum parallelisms,” *Physical Review Applied*, vol. 19, no. 5, p. 054032, 2023.
- [10] L. Stefanazzi, K. Treptow, N. Wilcer, C. Stoughton, C. Bradford, S. Uemura, S. Zorzetti, S. Montella, G. Cancelo, S. Sussman *et al.*, “The qick (quantum instrumentation control kit): Readout and control for qubits and detectors,” *Review of Scientific Instruments*, vol. 93, no. 4, 2022.
- [11] J. Liu, Y. Lee, H. Deng, C. Clayton, G. Yang, and X. Wu, “Risc-q: A generator for real-time quantum control system-on-chips compatible with risc-v,” *arXiv preprint arXiv:2505.14902*, 2025.
- [12] Y. Xu, G. Huang, N. Fruitwala, A. Rajagopala, R. K. Naik, K. Nowrouzi, D. I. Santiago, and I. Siddiqi, “Qubic 2.0: An extensible open-source qubit control system capable of mid-circuit measurement and feed-forward,” *arXiv preprint arXiv:2309.10333*, 2023.
- [13] A. M. D. Inc., “Amd zynq™ ultrascale+™ rfsocs: The industry’s only single-chip adaptable radio platform,” 2025. [Online]. Available: <https://www.amd.com/en/products/adaptive-socs-and-fpgas/soc/zynq-ultrascale-plus-rfsoc.html>
- [14] J. Winklmann, A. Alberti, and M. Schulz, “Comparison of atom detection algorithms for neutral atom quantum computing,” in *2024 IEEE International Conference on Quantum Computing and Engineering (QCE)*. IEEE, Sep. 2024, p. 1048–1057. [Online]. Available: <http://dx.doi.org/10.1109/QCE60285.2024.00124>
- [15] A. Elsharkawy, X. Guo, and M. Schulz, “Integration of quantum accelerators into hpc: Toward a unified quantum platform,” in *2024 IEEE International Conference on Quantum Computing and Engineering (QCE)*, vol. 01, 2024, pp. 774–783.
- [16] P. Döbler and M. S. Jattana, “A survey on integrating quantum computers into high performance computing systems,” *arXiv preprint arXiv:2507.03540*, 2025.
- [17] R. Ramsauer and W. Mauerer, “Towards system-level quantum-accelerator integration,” *arXiv preprint arXiv:2507.19212*, 2025.