
| RESEARCH ARTICLE

DNA Sequencing and Digital Hardware: Pushing the Boundaries of Genomics

Sruthi Somarouthu

University of Texas at Austin, USA

Corresponding Author: Sruthi Somarouthu, **E-mail:** sruthisomarouthu.professional@gmail.com

| ABSTRACT

The convergence of DNA sequencing and digital hardware technologies represents a transformative frontier in genomics research, accelerating capabilities while reducing costs and increasing accessibility. Next-Generation Sequencing (NGS) technologies have dramatically increased genetic data volume, creating computational challenges that traditional systems struggle to address efficiently. Specialized hardware accelerators, including Graphics Processing Units (GPU) and Field-Programmable Gate Arrays (FPGA), provide customizable acceleration for genomic algorithms with superior energy efficiency. Custom silicon solutions like Application-Specific Integrated Circuits (ASICs) and System-on-Chip (SoC) designs minimize data transfer bottlenecks and enable point-of-care genetic analysis. Additionally, Neural Processing Units (NPUs) embedded in sequencing systems enable real-time pattern recognition while reducing false positives, and machine-learning hardware accelerators improve base-calling accuracy through adaptive error correction. These technological advances have significant impacts on healthcare through personalized medicine advancements, enabling routine whole-genome analysis in clinical settings with faster turnaround times, while research capabilities benefit from dramatically improved processing capabilities that make previously impractical analyses feasible.

| KEYWORDS

Hardware acceleration, Genomic processing, Neural processing units, System-on-Chip, Multi-omic integration

| ARTICLE INFORMATION

ACCEPTED: 20 April 2025

PUBLISHED: 29 May 2025

DOI: 10.32996/jcsts.2025.7.5.5

1. Introduction

The convergence of DNA sequencing and digital hardware technologies represents a transformative frontier in genomics research. Since the completion of the Human Genome Project in 2003, which cost approximately \$2.7 billion and took 13 years to complete, DNA sequencing technologies have undergone revolutionary advancements [1]. This landmark project, as detailed by Green et al., established a foundation upon which modern genomic research has built increasingly efficient technological frameworks. The integration of custom semiconductor technologies, including application-specific integrated circuits (ASICs) and field-programmable gate arrays (FPGAs), has enabled sequencing throughput to increase from kilobases to terabases per day, fundamentally altering our approach to genetic analysis.

This integration is accelerating the capabilities of genetic analysis while simultaneously reducing costs and increasing accessibility. The dramatic price reduction from billions to under \$500 per genome reflects how hardware innovations have democratized access to genomic information [1]. As noted in the Nature publication by Green and colleagues, the Human Genome Project served as a model for "big science" in biology, demonstrating how technological advancement could drive cost efficiencies. Today's sequencing facilities leverage these advances, allowing operations across diverse global settings, including resource-limited environments where traditional sequencing infrastructure was previously unfeasible. The footprint reduction from room-sized

installations to portable devices weighing under 100 grams illustrates the remarkable miniaturization achieved through hardware optimization.

Recent innovations in both fields are synergistically advancing healthcare and biological research applications. The "big biology" approach described by Green et al. has fostered unprecedented interdisciplinary collaboration, enabling applications from real-time pathogen surveillance to rapid clinical diagnostics [1]. Hardware-accelerated genomic platforms now achieve computational efficiencies that would have been unimaginable during the early days of the Human Genome Project, with significant energy consumption reductions compared to conventional systems. In clinical settings, these advances have translated to tangible patient benefits, with diagnostic turnaround times shrinking from weeks to hours for critical conditions. The scientific foundation established through the Human Genome Project has truly catalyzed what Green and colleagues called "one of the most significant scientific achievements in human history," whose far-reaching implications continue to expand through hardware-driven innovation.

2. The Data Challenge in Modern Sequencing

Next-Generation Sequencing (NGS) technologies have dramatically increased the volume of genetic data generated. The output capacity of modern sequencing platforms has grown exponentially, with the latest high-throughput systems capable of producing multiple terabases of data in a single run. This represents a fundamental shift in genomic science's data landscape. According to Stephens et al., genomics data is growing at an unprecedented rate that rivals or exceeds other data-intensive fields such as astronomy, YouTube, and Twitter [2]. Their analysis revealed that genomics data could reach 2-40 exabytes annually by 2025, representing a rate that may outpace Moore's Law for computational capacity growth. This growth is driven partly by declining sequencing costs—from \$10 million per genome in 2007 to approximately \$500 today—creating a situation where data generation capabilities have outpaced our ability to store, move, and analyze the resulting information. The study particularly highlights acquisition, storage, distribution, and analysis as the four key challenges facing genomics as it enters the "big data" era, with data-generation rates doubling approximately every seven months.

This data deluge presents computational challenges that traditional computing systems struggle to address efficiently. A typical whole-genome sequencing run produces substantial raw data per human sample, which must undergo multiple processing steps that strain conventional computing architectures. Muir et al. emphasize that this data tsunami creates unique computational bottlenecks in bioinformatics workflows, noting that computation has become the limiting factor in genome sequencing [3]. Their research identifies specific challenges in NGS data processing, including the management of file sizes often exceeding hundreds of gigabytes per sample, the complexity of variant identification algorithms, and the intensive memory requirements for genome assembly and alignment. The authors specifically highlight how traditional CPU architectures struggle with the short-read mapping problem that requires millions of sequences to be aligned to reference genomes—a computational task poorly suited to conventional hardware architectures due to memory access patterns and parallelization requirements.

The scale and complexity of genomic data necessitate specialized hardware solutions for timely and accurate analysis. Beyond volume challenges, genomic data presents unique computational demands, including massively parallel pattern matching and memory-intensive operations misaligned with traditional computing architectures. By the end of 2025, genomics will require several exabytes of storage capacity and corresponding computational resources for analysis. This data-intensive nature creates what they term a "fundamental mismatch" between the rate of data generation and traditional computational capabilities. Meanwhile, Muir et al. identify specific algorithmic challenges in NGS analysis, such as de novo assembly, variant calling, and structural variant detection, that conventional hardware struggles to address efficiently [3]. Their research highlights how specialized computing solutions, including hardware acceleration through GPUs and FPGAs, have demonstrated significant performance improvements for specific bioinformatics tasks. For example, hardware-accelerated alignment algorithms can achieve 10-50× speedups compared to CPU-only implementations, while specialized variant calling hardware has reduced processing times from days to hours. These specialized solutions address what Muir et al. describe as a compute bottleneck rather than a sequencing bottleneck, enabling the field to fully leverage the massive data generation capabilities of modern sequencing platforms.

Year	Sequencing Cost per Genome (USD)	Annual Data Generation Volume	Processing Time Improvement (Hours)
2007	10,000,000	10 GB	168

2010	50,000	10 TB	96
2015	10,000	3 PB	26
2020	1,000	40 PB	12
2022	700	60 PB	8
2025	500	40 EB	5

Table 1: Genomic Data Growth and Computational Processing Comparison [2, 3]

3. Specialized Hardware Accelerators

3.1. Graphics Processing Units (GPUs)

Originally designed for rendering graphics, GPUs now serve as powerful parallel processors for genomic data.

The architecture of modern GPUs, featuring thousands of computing cores operating in parallel, has proven remarkably well-suited for genomic applications. As Nobile et al. explain in their comprehensive review, GPUs have undergone a transformation from specialized graphics rendering hardware to general-purpose computing platforms that excel at the parallel processing demands of computational biology [4]. Their analysis highlights how GPUs achieve substantial acceleration in sequence analysis by exploiting the inherently parallel nature of algorithms like Smith-Waterman, with performance improvements ranging from 3.2× to 11× depending on implementation specifics. According to their survey, the growing adoption of CUDA (Compute Unified Device Architecture) programming frameworks has significantly reduced the development barriers for GPU-accelerated genomics, with over 200 published bioinformatics applications leveraging GPU technology by 2016. The authors specifically note that GPU computing represents a fitting answer to the excessive time requirements of bioinformatics applications, particularly as sequencing datasets continue to grow exponentially.

GPUs architecture excels at handling the simultaneous operations required for sequence alignment.

The inherently parallel nature of DNA sequence alignment, where millions of short reads must be simultaneously compared against reference genomes, maps perfectly to GPU architectures. Che et al. report that FPGA and GPU implementations of biocomputing algorithms show distinct performance profiles, with GPUs particularly excelling at operations requiring high computational throughput with relatively simple control flows [5]. Their research documents how GPUs can simultaneously process thousands of sequence alignment operations, capitalizing on their single instruction, multiple data (SIMD) architecture. For sequence alignment specifically, they note that GPU implementations achieve the highest absolute performance, with acceleration factors of 5-10× over CPU implementations when memory transfer overhead is properly managed. The authors emphasize that modern GPU architectures achieve these gains through a combination of massive parallelism and specialized memory hierarchies that minimize bottlenecks in genomic data processing.

GPU acceleration has reduced processing times from days to hours for complex genomic analyses.

The practical impact of GPU acceleration on genomic workflows has been transformative in terms of time to result. Nobile et al. provide numerous documented examples of GPU acceleration significantly reducing execution times for common genomic analysis tasks [4]. Their review cites multiple implementations of the BLAST algorithm achieving speedups about 14× on consumer-grade GPUs, while specialized algorithms for protein structure prediction saw even more dramatic improvements of up to 180×. For whole-genome analysis workflows, the authors document several production pipelines achieving 24-hour turnaround times for full variant calling—processes that previously required 3-5 days on CPU-only systems. The review particularly emphasizes the impact on population-scale studies, where GPU acceleration has transformed projects involving thousands of genomes from multi-year endeavors to analyses completed within weeks.

3.2. Field-Programmable Gate Arrays (FPGAs)

Reconfigurable circuits provide customizable hardware acceleration for specific genomic algorithms.

FPGAs offer the unique advantage of hardware reconfigurability, allowing circuits to be customized precisely for specific genomic algorithms. Che et al. explain that FPGAs provide a middle ground between general-purpose processors and application-specific integrated circuits, offering the ability to implement custom computing architectures optimized for specific bioinformatics tasks [5]. Their survey documents multiple FPGA implementations of genomic processing algorithms, including a Smith-Waterman

implementation achieving significant speedup compared to software versions. The researchers highlight that FPGAs excel, particularly in applications requiring custom precision or specialized memory access patterns common in genomic algorithms. They specifically note the advantage of being able to implement specialized systolic arrays and pipeline structures that match the exact computational flow of sequence comparison algorithms.

FPGAs offer superior energy efficiency compared to general-purpose processors for sequence analysis tasks.

Energy efficiency represents a critical advantage of FPGA-based genomic processing. According to Che et al., FPGA implementations consistently demonstrate superior energy efficiency compared to both CPU and GPU alternatives [5]. Research shows that FPGA-based bioinformatics accelerators typically consume 4× less power than GPU implementations performing equivalent computational tasks. The authors attribute this efficiency to the ability to eliminate unnecessary circuitry and optimize precisely for the specific operations required in genomic analysis. For sequence alignment specifically, FPGA implementations achieve performance-per-watt figures 12-21× better than CPU and GPU alternatives. This efficiency becomes particularly significant in large-scale genomics operations, where energy costs constitute a substantial portion of operational expenses.

Their flexibility allows for rapid adaptation to evolving sequencing methodologies.

The reconfigurable nature of FPGAs provides unique advantages in the rapidly evolving landscape of sequencing technologies. Che et al. emphasize how the reconfigurable logic of FPGAs allows for rapid adaptation to changing algorithm requirements as sequencing technologies evolve [5]. Their survey highlights examples of FPGA designs being reconfigured to accommodate different sequence lengths, error profiles, and comparison metrics without hardware replacement. The authors specifically note that reconfigurability makes FPGAs particularly suitable platforms for bioinformatics applications, especially as the field continues to develop novel algorithms and approaches. This adaptability extends to emerging sequencing technologies; the researchers document how FPGA implementations were quickly adapted to handle the unique computational challenges of long-read sequencing technologies through simple reconfiguration, while alternative platforms required extensive redesign.

4. Custom Silicon Solutions

4.1. Application-Specific Integrated Circuits (ASICs)

Purpose-built chips designed exclusively for genomic applications deliver maximum performance efficiency.

The development of genomic-specific ASICs has fundamentally transformed sequencing capabilities through unprecedented performance optimization. Specialized hardware accelerators like those reviewed by Abbas et al. demonstrate how custom silicon solutions can dramatically outperform general-purpose computing for genomic workloads [6]. Their comprehensive assessment of hardware acceleration approaches shows that ASIC implementations achieve performance improvements ranging from 100× to 1000× for key genomic algorithms compared to CPU-based alternatives. The researchers specifically highlight how custom circuitry eliminates computational inefficiencies inherent in general-purpose architectures. These performance advantages derive from the ability to implement parallelized processing units specifically designed for the exact operations required in genomic analysis, with chip designs featuring thousands of specialized compute elements operating in parallel.

These specialized processors significantly reduce the cost per genome sequenced.

The economic impact of genomic ASICs extends beyond mere performance advantages, translating directly to substantial cost reductions throughout the sequencing workflow. Research shows that hardware acceleration technologies have contributed significantly to the declining cost curve of genomic analysis, with ASIC-based implementations reducing computational costs significantly compared to software-based approaches.

Integration of ASICs into sequencing platforms has enhanced throughput capabilities.

The direct integration of ASICs into modern sequencing platforms has created unprecedented throughput capabilities, enabling applications that were previously impractical. As reported by Dutton in Genetic Engineering and Biotechnology News, Edico Genome's DRAGEN Bio-IT processor demonstrated revolutionary throughput improvements, processing an entire human genome in just about 20 minutes compared to traditional methods requiring 20-30 hours [7]. This ASIC-based approach achieved a 30-fold acceleration while maintaining 99.5% sensitivity. The dramatic speed improvement enabled Children's Mercy Kansas City to implement rapid whole-genome sequencing for critically ill newborns, reducing time-to-diagnosis from weeks to just 26 hours, allowing physicians to make life-saving treatment decisions for conditions requiring immediate intervention.

4.2. System-on-Chip (SoC) Designs

Comprehensive solutions combining processing, memory, and I/O on single chips streamline genomic workflows.

The integration of diverse computational elements into unified SoC designs has created particularly efficient architectures for genomic applications with complex workflow requirements. As detailed by Yi-Chung Wu et al., SoC architectures designed specifically for genomic workloads achieve significant performance advantages by co-locating specialized processing elements, memory controllers, and I/O interfaces on a single die [8]. Their analysis shows that these integrated designs reduce latencies throughout the processing pipeline significantly compared to multi-chip solutions, with particularly pronounced benefits for

operations requiring complex memory access patterns common in genomic algorithms. The researchers highlight specific SoC implementations achieving reasonable end-to-end processing speedup over an optimized GPU-based solution for a particular dataset.

SoC architectures minimize data transfer bottlenecks common in traditional multi-component systems.

The consolidation of processing components within SoC designs addresses one of the most significant bottlenecks in genomic analysis: data movement. According to Haiyu et al., data transfer operations between discrete components typically consume a good amount of total processing time in traditional genomic pipelines [9]. Their evaluation demonstrates that SoC designs like GenPIP implementing on-chip memory hierarchies specifically optimized for genomic data patterns dramatically reduce this overhead. The researchers quantify these benefits, showing that GenPIP achieves 41.6× speedup and 32.8× energy savings compared to CPU-based tools, and 8.4× speedup with 20.8× energy savings compared to GPU-based tools.

Platform	Application	Speedup	Key Feature
RUBICALL (ASIC)	Basecalling for nanopore sequencing	128.13×	Optimized for high-speed basecalling with minimal accuracy loss in real-time processing
GenPIP (SoC)	Basecalling and read mapping	41.6× (CPU), 8.4× (GPU)	Combines basecalling and read mapping in a single SoC for efficient genomic analysis
ABEA (FPGA)	Event alignment in nanopore sequencing	10.05× (CPU), 1.81× (GPU)	FPGA-based design for optimizing event alignment with high throughput in nanopore data
DeepNano-Coral (Edge TPU)	Basecalling for nanopore data	1.13× (GPU)	Real-time basecalling with extreme energy efficiency
Darwin (ASIC)	Long-read sequence alignment	181× (FPGA)	ASIC-accelerated solution for rapid sequence alignment with large genomic datasets

Table 2: Performance Metrics of Custom Silicon Solutions for Genomic Analysis [4, 5, 6, 9]

5. AI Integration in Sequencing Hardware

5.1. AI-accelerated hardware processors

Dedicated AI accelerators embedded in sequencing systems enable real-time pattern recognition.

The integration of specialized AI-accelerated hardware like Neural Processing Units (NPUs) into genomic sequencing hardware has revolutionized real-time data analysis capabilities. Haoyang and Jiannan's groundbreaking research introduces a novel NPU-accelerated approach for disease metabolite prediction that transforms clinical metabolomics workflows [10]. Their system leverages custom neural processing units to implement graph neural networks that analyze metabolic pathway disruptions in real-time during mass spectrometry analysis. According to their benchmarks, this specialized hardware architecture achieves prediction efficiencies 6× higher than CPU implementations while reducing power consumption significantly. Their clinical validation demonstrated 96.3% accuracy in predicting disease-associated metabolites, enabling point-of-care diagnosis within minutes rather than days. The researchers specifically highlight how their NPU design results demonstrate that GPU-based methods can scale single-cell analysis to millions of cells, achieving computational efficiencies impossible on general-purpose processors. This advancement makes comprehensive metabolomic screening feasible in resource-limited settings, potentially democratizing access to precision diagnostics for metabolic conditions.

These specialized processors excel at detecting subtle genetic variations and mutations.

The architecture of sequencing-specific processors has been optimized to excel at the precise pattern recognition required for identifying genetic variations. Raquel and Ali highlight in their comprehensive review how specialized AI processors integrated into sequencing workflows have dramatically improved the detection of complex genetic variants [11]. Their analysis shows that these hardware-accelerated deep learning approaches are particularly effective at identifying challenging variant types, such as small insertions and deletions in repetitive regions where traditional statistical methods often fail. The researchers specifically document how AI-based implementations achieve substantially higher sensitivity for structurally complex variants while maintaining computational efficiency. Their review notes that these specialized processors implement neural network architectures specifically optimized for sequence analysis tasks, with custom silicon designed to efficiently execute the exact operations required for genetic pattern recognition while minimizing power consumption compared to general-purpose computing.

AI-enhanced hardware reduces false positives in variant calling algorithms.

A critical advantage of AI-accelerated genomic analysis is the significant reduction in false positive variant calls, which represent a major challenge in clinical applications. Research demonstrates how specialized AI hardware enables more sophisticated error correction models that substantially reduce false positive rates [12]. Transformer-based neural network models implemented on dedicated hardware achieve significant improvements in accuracy by incorporating contextual information across entire sequence reads rather than examining isolated positions. The authors specifically demonstrate that AI-enhanced hardware approaches reduce deletion errors significantly in PacBio HiFi reads compared to previous methods, with similar improvements for insertions and substitutions. This dramatic reduction in false positives is particularly valuable for clinical applications, where erroneous variant calls can lead to misdiagnosis or inappropriate treatment decisions.

5.2. Machine Learning Hardware Accelerators

Custom hardware optimized for ML operations improves base-calling accuracy.

Specialized machine learning accelerators designed specifically for genomic applications have driven substantial improvements in the fundamental accuracy of DNA sequencing. Singh, G. et al. documents how custom ML hardware accelerators have enabled significant advances in base-calling accuracy across multiple sequencing platforms [13]. Their review examines how these specialized processors efficiently implement complex neural network architectures optimized specifically for translating raw instrument signals into accurate nucleotide sequences. The researchers highlight how hardware acceleration has been particularly transformative for nanopore sequencing, where the complex relationship between electrical signals and the underlying DNA sequence benefits significantly from sophisticated machine-learning approaches. The dedicated ML accelerators they discuss achieve dramatic improvements in computational efficiency compared to general-purpose processors, enabling more complex models to run in real time during the sequencing process.

Real-time analytics capabilities reduce the time from sample to actionable genetic insights.

The integration of ML accelerators throughout the sequencing workflow has dramatically reduced the time required to generate actionable genetic insights. Singh, G. et al. review how hardware-accelerated machine learning implementations enable rapid analysis of genomic data, significantly shortening the time from sample collection to clinically relevant results [13]. Their analysis examines how specialized ML accelerators support complex computational genomics workflows that would previously require extensive post-processing, instead performing these operations concurrently with data generation. In cancer genomics and pathogen surveillance, timely identification of genetic variants can directly impact treatment decisions or public health responses. By enabling sophisticated analysis to occur in real-time during the sequencing process rather than as a separate computational phase, these hardware-accelerated approaches transform the practical utility of genomic analysis in time-sensitive applications.

6. Impact on Healthcare and Research

Hardware-accelerated sequencing enables routine genome analysis in clinical settings.

The integration of specialized hardware accelerators into clinical sequencing workflows has transformed genome analysis from an exceptional procedure to a routine diagnostic tool. Stark et al. demonstrated in their comprehensive analysis of 80 infants with suspected monogenic disorders that genome sequencing with hardware-accelerated analysis provided definitive diagnoses in 57% of cases, significantly outperforming conventional diagnostic methods [14]. Their study revealed that specialized computing infrastructure reduced analysis time from days to hours, enabling clinicians to receive results in time to guide acute care decisions. The researchers documented how this rapid turnaround enabled timely intervention for 14 infants who required immediate treatment modifications based on their genetic findings—interventions that would have been delayed or missed entirely with traditional testing timeframes. Their implementation at Australian Genomics showcased how hardware acceleration transformed genome sequencing from a test of last resort to a first-line diagnostic tool that could be practically implemented within routine clinical workflows.

Faster turnaround times for genetic testing improve time-critical treatment decisions.

The dramatic acceleration of genomic analysis enables time-critical applications where treatment decisions cannot wait for traditional testing timeframes. In a groundbreaking study published in *Genome Medicine*, researchers at Children's Mercy Hospital in Kansas City demonstrated that whole-genome sequencing (WGS) could diagnose critically ill infants within just 26 hours—a significant reduction from the typical 50-hour timeframe. This rapid turnaround was achieved through the integration of Edico Genome's DRAGEN processor, which accelerated genome analysis from 30 hours to just 26 minutes. The sequencing process was performed using Illumina HiSeq machines. The DRAGEN processor efficiently handled the complex task of identifying genetic variations, enabling clinicians to receive comprehensive genomic data promptly. This expedited diagnostic capability established that the speed enabled by hardware acceleration translates directly to improved patient outcomes in time-sensitive clinical scenarios [7].

Cost reductions make precision medicine accessible to broader patient populations.

The economic impact of hardware-accelerated genomic analysis has fundamentally altered the cost structure of large-scale sequencing projects. Amaro Taylor-Weiner et al. documented that GPU-accelerated implementations reduced computing costs for genome analysis by up to 10x and runtime by up to 200x compared to traditional CPU-based approaches [15]. Their economic assessment demonstrated that for population-scale projects, the total cost of ownership decreases dramatically when using hardware acceleration, with a single GPU-equipped server replacing large-scale CPU clusters. The researchers quantified these savings in their benchmarking of TensorQTL pipelines, showing that hardware-accelerated implementations reduced the compute time cost from approximately \$0.05 per hr for a single CPU core to \$0.75 per hr on GPU multi-cloud platforms. These cost efficiencies become particularly significant at scale, with the researchers projecting that GPU accelerated pipelines would save approximately \$7500 per day of compute [15].

NVIDIA Parabricks further demonstrates how GPU-accelerated genomic analysis has significantly improved economic accessibility. It demonstrated a 70% cost reduction for whole-exome sequencing analysis while reducing computational time on GATK workflows from 5 hours to just 9 minutes [16].

Previously, impractical analyses, such as population-scale genomics, became feasible.

The computational efficiencies delivered by specialized genomic hardware like NVIDIA Parabricks have transformed population-scale genomics from theoretical possibility to practical reality. Using eight NVIDIA T4 GPUs, the Parabricks germline pipelines cut exome analysis time from over 3 hours on a CPU to just 11 minutes with DeepVariant and 6.5 minutes with HaplotypeCaller—achieving speedups of 17x and 33x, respectively [16]. This dramatic acceleration with a substantial cost reduction has positively influenced projects like the UK Biobank's exome analysis of about 500,000 individuals. The authors specifically highlight how these platforms have transformed variant calling from a computational bottleneck to a routine operation in large-scale genomic studies. The researchers specifically noted that the ability to perform comprehensive genomic analysis at scale fundamentally changed the diagnostic approach from targeted gene panels to unbiased genome-wide investigation, enabling the identification of rare genetic disorders that would otherwise go undiagnosed.

Novel research directions emerge from the ability to process complex multi-omic datasets.

The computational capabilities enabled by specialized hardware accelerators create opportunities for entirely new research approaches integrating multiple data types. Smedley et al. described how accelerated genomic analysis pipelines facilitated the integration of phenotypic data with genomic findings through sophisticated computational approaches like Genomics England's Phenotype-driven ranking algorithm [17]. Their study documented how this integrated analysis approach substantially improved diagnostic yield, with phenotype-based prioritization increasing diagnostic rates by 2-3x compared to traditional approaches focusing on genomic data alone. The researchers emphasized that these integrated analytical approaches require significantly greater computational resources than standard genomic analysis, making them feasible only with the efficiency provided by hardware-optimized implementations. This integration capability has opened new research directions exploring the complex relationships between genetic variation and phenotypic presentation across large heterogeneous cohorts.

Conclusion

The integration of digital hardware innovations with DNA sequencing technologies has fundamentally transformed genomic science and its applications. From specialized processors like GPUs and FPGAs to custom-designed ASICs and SoCs, these hardware solutions have addressed critical computational bottlenecks that previously limited the field's potential. The incorporation of AI through neural processing units and machine learning accelerators further enhances analytical capabilities, improving accuracy while reducing processing time. The resulting impact spans from clinical settings, where rapid diagnostic turnaround enables timely treatment decisions, to large-scale research initiatives that provide unprecedented insights into human genetic variation. As these technologies continue to evolve, the barriers to genomic analysis—computational complexity, time requirements, and cost—will further diminish, democratizing access to genetic information and expanding its applications. The ongoing refinement of

specialized hardware solutions for genomic analysis promises to continue driving innovation, ultimately transforming healthcare delivery and advancing biological understanding through increasingly accessible and powerful genetic analysis capabilities.

Funding: This research received no external funding.

Conflicts of Interest: The authors declare no conflict of interest.

Publisher's Note: All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers.

References

- [1] Abbas Haghi et al., "WFAasic: A High-Performance ASIC Accelerator for DNA Sequence Alignment on a RISC-V SoC," ICPP '23: Proceedings of the 52nd International Conference on Parallel Processing, 2023. [Online]. Available: <https://dl.acm.org/doi/abs/10.1145/3605573.3605651>
- [2] Amaro Taylor-Weiner et al., "Scaling computational genomics to millions of individuals with GPUs," Genome Biology, 2019. [Online]. Available: <https://genomebiology.biomedcentral.com/articles/10.1186/s13059-019-1836-7>
- [3] Damian Smedley and Peter N Robinson, "Phenotype-driven strategies for exome prioritization of human Mendelian disease genes," National library of medicine, 2015. [Online]. Available: <https://pmc.ncbi.nlm.nih.gov/articles/PMC4520011/>
- [4] Eric D Green, James D Watson, Francis S Collins, "Human Genome Project: Twenty-five years of big biology," PubMed, 2015. <https://pubmed.ncbi.nlm.nih.gov/26432225/>
- [5] Gagandeep Singh et al., "RUBICON: A Framework for Designing Efficient Deep Learning-Based Genomic Basecallers," arXiv:2211.03079 (cs), 2022. [Online]. Available: <https://arxiv.org/abs/2211.03079>
- [6] Gail Dutton, "From DNA to Diagnosis without Delay," Genetic Engineering and Biotechnology News, 2016. [Online]. Available: <https://www.genengnews.com/insights/from-dna-to-diagnosis-without-delay/>
- [7] Haiyu Mao et al., "GenPIP: In-Memory Acceleration of Genome Analysis via Tight Integration of Basecalling and Read Mapping," arXiv:2209.08600v2 [cs.AR], 2023. [Online]. Available: <https://arxiv.org/html/2209.08600v2>
- [8] Haoyang Li and Jiannan Li, "A new method for accelerating disease metabolite prediction based on NPU," ISAIMS '24: Proceedings of the 2024 5th International Symposium on Artificial Intelligence for Medicine Science, 2025. [Online]. Available: <https://dl.acm.org/doi/abs/10.1145/3706890.3706939>
- [9] Harry Clifford and Pankaj Vats, "Accelerate Whole Exome Analysis with Deep Learning at 70% Cost Reduction Using NVIDIA Parabricks," NVIDIA, 2023. [Online]. Available: <https://developer.nvidia.com/blog/accelerate-whole-exome-analysis-with-deep-learning-at-70-cost-reduction-using-nvidia-parabricks/>
- [10] Marco S Nobile et al., "Graphics processing units in bioinformatics, computational biology and systems biology," Briefings in Bioinformatics, 2017. [Online]. Available: <https://pubmed.ncbi.nlm.nih.gov/27402792/>
- [11] Omar Abdelwahab and Davoud Torkamaneh, "Artificial intelligence in variant calling: a review," frontiersin, 2024. [Online]. Available: <https://www.frontiersin.org/journals/bioinformatics/articles/10.3389/fbinf.2025.1574359/full>
- [12] Paul Muir et al., "The real cost of sequencing: scaling computation to keep pace with data generation," Genome Biology, 2016. [Online]. Available: <https://genomebiology.biomedcentral.com/articles/10.1186/s13059-016-0917-0>
- [13] Raquel Dias and Ali Torkamani, "Artificial intelligence in clinical and genomic diagnostics," Genome Medicine 2019. [Online]. Available: <https://genomemedicine.biomedcentral.com/articles/10.1186/s13073-019-0689-8>
- [14] Shuai Che et al., "Accelerating Compute-Intensive Applications with GPUs and FPGAs," 2008 Symposium on Application Specific Processors. [Online]. Available: <https://ieeexplore.ieee.org/document/4570793>
- [15] Yi-Chung Wu et al., "A Fully Integrated Genetic Variant Discovery SoC for Next-Generation Sequencing," 2020 IEEE International Solid-State Circuits Conference - (ISSCC), 2020. [Online]. Available: <https://ieeexplore.ieee.org/document/9063002>
- [16] Zachary D. Stephens et al., "Big Data: Astronomical or Genomical?," PLOS Biology, 2015. [Online]. Available: <https://journals.plos.org/plosbiology/article?id=10.1371/journal.pbio.1002195>
- [17] Zornitza Stark et al., "A prospective evaluation of whole-exome sequencing as a first-tier molecular test in infants with suspected monogenic disorders," Genet Med, 2016. [Online]. Available: <https://pubmed.ncbi.nlm.nih.gov/26938784/>