

Improving patient outcomes with NVIDIA CUDA[®] -powered healthcare solutions



Yogeshwar Misal

Technical Lead, CitiusTech

Summary

CUDA technology accelerates healthcare by speeding up medical imaging and drug discovery, enhancing patient outcomes. It optimizes real-time data processing for diagnostics and personalized treatments. However, implementing CUDA requires overcoming hardware and software challenges. To address these challenges, CitiusTech provides strategic integration and expertise to maximize its benefits in healthcare.



Introduction

In an emergency medical situation, every second counts. Take a stroke unit, for instance. Rapid imaging helps doctors to make quicker decisions and even stop a stroke in its tracks. But traditional CPU processing means vital 3D brain scans could take longer to render—precious time during which brain cells die, potentially leading to irreversible damage. Or consider a pandemic situation, where developing a new drug quickly is a must. But the process of molecular simulation in drug discovery with traditional processing can take months or even years—a delay that can result in countless lives lost.

Healthcare is always in a race against time, and it needs powerful processing for performance-intensive tasks. And that's where NVIDIA's CUDA⁽¹⁾ (Compute Unified Device Architecture) can be a game changer. This platform is designed with a software framework to expand the power of GPU (Graphic Processing Unit) acceleration. It allows developers to process large datasets, images, and complex computations faster than a traditional CPU. In Healthcare, it accelerates image processing for MRI, CT, and PET scans, allowing for real-time analysis and faster diagnostics. It drastically reduces genome sequencing⁽²⁾ time, facilitating rapid insights for personalized medicine, speeding up molecular simulations, and aiding faster virtual screening and drug discovery⁽³⁾ processes.



Understanding CUDA's architecture for healthcare applications

CUDA transforms a GPU from a graphics processor into a powerful computing device capable of handling thousands of tasks simultaneously. Unlike a regular CPU, which processes tasks sequentially (one after the other), a GPU with CUDA can execute many tasks at the same time. This is like having thousands of tiny workers (cores) who can perform small tasks simultaneously, rather than waiting for one worker to finish before the next one starts. This capability is what makes GPUs with CUDA incredibly fast and efficient for complex calculations.

- **Scientific computing:** Take climate modeling or astrophysical phenomena. Researchers need to simulate years of processes and CUDA helps them accomplish it in days or even hours, leading to faster breakthroughs and more iterative scientific exploration.
- **Machine learning:** Artificial intelligence is evolving exponentially and has its footprint across industries and functions. To continuously evolve, these AI and Machine Learning models need to be trained. With CUDA's rapid processing AI models can learn from enormous datasets in a fraction of the usual time.
- **Data Analytics:** There is a flood of Big data, that is expected to reach 181 Zettabytes⁽⁴⁾ by 2025. While data is the real power, processing and extracting insights from these large repositories is still challenging. CUDA enables real-time processing and analysis of large volumes of data on the fly.

What happens when healthcare moves at GPU speed?

- **Accelerated diagnostics with medical imaging:** Radiologists inundated with countless images can accelerate image processing⁽⁵⁾ tasks, such as segmentation, registration, and reconstruction, transforming hours of work into mere minutes, ultimately improving patient outcomes. In a stroke situation discussed earlier, this drastic reduction can mean the difference between full recovery and permanent disability for stroke patients.
- **Personalized treatments with Genomics and bioinformatics:** Genetic diseases require rapid identification to begin appropriate treatments, for instance, in newborns showing symptoms of genetic disorders. CUDA speeds up genome sequencing, alignment, and variant analysis. This acceleration allows for immediate intervention and significantly improving outcomes by enabling timely treatment plans tailored to the genetic profiles of patients⁽⁶⁾.



- **Expedited drug discovery and development:** The road from research to market can be lengthy and fraught with challenges. However, with CUDA, organizations can optimize molecular simulations and drug design processes^[7], identifying potential drug candidates with remarkable efficiency, expediting the discovery of life-saving medications, and reducing the time and cost associated with bringing these innovations to market.
- **Better clinical decisions with real-time insights:** By integrating CUDA, CDSS can analyze real-time data from multiple sources instantly, providing critical insights^[8] at the point of care and in high-stakes environments like the ICU.

Healthcare organizations are successfully integrating CUDA to enhance operational efficiency and patient outcomes. For instance, United Imaging Healthcare^[9] used NVIDIA's CUDA-based solutions to accelerate medical imaging workflows, significantly improving image reconstruction and deep learning model performance. This has led to faster image processing times in radiology applications, ensuring more accurate diagnostics and patient outcomes.

Similarly, Mount Sinai Health System^[10] adopted CUDA for genomic analysis, with GPUs accelerating one genomic analysis by more than 6x. This rapid data processing has enabled the identification of therapeutic targets for personalized cancer treatment while lowering operational costs.

The benefits make a compelling case for implementing CUDA in Healthcare, but it is easier said than done. Healthcare technology is a complex ecosystem burdened with legacy systems, budget constraints, regulatory constraints, and a complex web of edge solutions. Successfully implementing and integrating CUDA needs careful planning.

Overcoming technical barriers with strategic solutions

CitiusTech specializes in healthcare technology and understands the industry's complexities. We help clients incorporate CUDA-enabled GPU acceleration into their systems, enhancing workflow efficiency. Our expertise in high-performance computing can significantly improve patient care by maximizing the benefits of GPU technology.

01. Hardware compatibility and infrastructure limitations

Many Healthcare institutions still operate on legacy systems that are incompatible with GPUs and upgrading to GPU-capable infrastructure requires substantial capital investment. A phased implementation strategy offers a way forward. By gradually integrating GPU-enabled hardware into existing systems, organizations can spread out costs and minimize disruptions.

Solution

Most healthcare providers have outdated systems that are not immediately compatible with GPU technology. CitiusTech assists clients in phased implementation strategies that align with their budget and operational needs. For instance, we recently set up a private GPU farm integrated with AWS Sagemaker for a client. This hybrid approach allowed the client to leverage existing infrastructure while accessing cloud-based GPU resources which enabled them to use the power of CUDA, reduce upfront costs, and ensure scalability.

02. Software integration and development challenges

Developing CUDA applications demands specialized expertise in parallel programming and a deep understanding of GPU architecture. Training programs focused on CUDA development can help build in-house expertise, positioning teams to thrive in a GPU-driven future. Alternatively, collaborating with experts can provide the additional support needed to navigate complex CUDA development. High-level libraries like CUDA-X AI or cuDNN, simplify the application development process.

Solution

Developing applications that use the full power of CUDA requires specialized skills in parallel programming and GPU architecture/ CitiusTech has equipped their in-house teams with the required training and necessary skills to develop and optimize CUDA applications.

03. Data privacy and security concerns

Ensuring that CUDA applications comply with regulations like HIPAA or GDPR is non-negotiable. Additionally, sensitive patient data must be handled with the utmost care, especially when utilizing cloud-based GPU resources.

Solution

Encrypting data at rest and in transit ensures that sensitive information remains protected and regular compliance audits ensure all processes are up to the highest standards. Layering in access control mechanisms adds another level of defense, making sure only authorized personnel have access to critical data.

- CitiusTech understands the importance of and is always compliant with regulations like HIPAA and GDPR
- Our teams place a strong emphasis on data security
- Regular compliance audits and stringent access control mechanisms
- All this when utilizing cloud based-CUDA enabled-GPU resources



04. Scalability and performance optimization

Managing GPU resources efficiently and adapting algorithms to fully exploit the parallelism of CUDA can be complex, often requiring a rethink of existing approaches.

Solution

Dynamic scaling tools allow real-time resource allocation and ensure that systems operate optimally under fluctuating loads. Performance profiling tools identify bottlenecks, helping developers fine-tune their applications for peak performance, and data/task parallelism unlocks the full power of CUDA's parallel processing capabilities.

- Our team understands to maximize the benefits of CUDA, we require efficient management of GPU resources and optimization of algorithms for parallel processing
- Or experts utilize dynamic scaling and performance profiling tools to identify and eliminate bottlenecks, ensure applications run at peak efficiency
- Above helps our clients to unlock the full potential of CUDA's capabilities

05. Interoperability with existing software systems

Healthcare IT is a patchwork of different systems, from Electronic Health Records (EHR) to specialized diagnostic tools. Ensuring that CUDA-powered solutions can communicate effectively with these systems is essential to avoid creating more silos.

Solution

Creating robust APIs that facilitate communication between CUDA applications and existing healthcare systems enables seamless integration. Adopting industry-standard data formats and protocols further enhances compatibility across systems. Middleware solutions can also bridge gaps between disparate systems, ensuring smooth data flow and interoperability across the board.

- The heterogeneity of healthcare IT systems is always an integration challenge.
- Our teams specialize in creating middleware solutions and robust APIs that facilitate seamless communication between CUDA-powered applications and existing systems like EHRs and laboratory information systems
- We adhere to industry- standard data formats and protocols, we ensure that new technologies enhance existing workflows



06. Futureproofing and technology evolution

The pace of change in GPU technology is relentless. How should organizations ensure the longevity of their investments?

Solution

Building applications with modular architecture can enable easy updating of components, without a need for a complete overhaul. Favoring open standards and platforms reduces dependency on single vendors, ensuring agility.

- Citiustech applies a forward-thinking approach to system designs always.
- We promote modular application architectures, and the adoption of open standards and platforms ensuring our clients remain agile and can adapt to technological advancements with minimal friction.

CUDA-powered Breakthroughs: The Future of Accelerated Healthcare

The possibilities in healthcare grow exponentially with technology. We have reached a point where Generative AI can unlock unprecedented insights and innovations. However, without the right infrastructure, these advancements remain just potential. CUDA changes this. It can accelerate the healthcare journey towards an ideal of value-based, interoperable, and human-centric care. Every piece of medical data, from a rural clinic to biggest cities, seamlessly integrates. CUDA's power enables a real-time, universal patient profile accessible across continents, guiding doctors to make the most informed decisions. This level of care coordination and data fluidity enhances patient outcomes, streamlines operational efficiencies, and reduces overall healthcare costs. CUDA is the golden ticket to fully leverage AI's capabilities, on the way for a future where technology and human-centric care converge to improve lives effectively.



References

1. <https://developer.nvidia.com/cuda-zone>
2. <https://blogs.nvidia.com/blog/world-record-genome-sequencing-parabricks/>
3. <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC7524381/>
4. <https://www.statista.com/statistics/871513/worldwide-data-created/>
5. <https://developer.nvidia.com/blog/accelerating-medical-image-processing-with-dali/>
6. <https://www.nvidia.com/en-us/claragenomics/>
7. <https://www.nvidia.com/en-in/industries/healthcare-life-sciences/drug-discovery/>
8. <https://www.nvidia.com/en-in/deep-learning-ai/solutions/data-analytics/>
9. <https://www.nvidia.com/en-us/case-studies/united-imaging-healthcare/>
10. <https://blogs.nvidia.com/blog/mount-sinai-genomics-on-gpus/>





Shaping Healthcare Possibilities

CitiusTech is a global IT services, consulting, and business solutions enterprise 100% focused on the healthcare and life sciences industry. We enable 140+ enterprises to build a human-first ecosystem that is efficient, effective, and equitable with deep domain expertise and next-gen technology.

With over 8,500 healthcare technology professionals worldwide, CitiusTech powers healthcare digital innovation, business transformation and industry-wide convergence through next-generation technologies, solutions, and products.

www.citiustech.com

Shaping Healthcare Possibilities

