

Abhigyan Gandhi

+44 7423024085 | +91 9589959545 | [LinkedIn](#) | [GitHub](#) | [Portfolio](#) | [Email](#) | **Work authorisation:** UK Graduate visa valid to Dec 2027 - no sponsorship required to start. Authorised to work in India. **Availability:** Immediate

SUMMARY

High-Performance C++ Engineer (MSc, University of Leeds) specializing in low-latency systems and GPU programming. Experienced across trading infrastructure, CUDA compute, and real-time Vulkan rendering, with a habit of profiling and benchmarking before optimising. Recipient of the Game Republic Red Kite Games Technology Award 2025.

SKILLS

Languages: C++(17/20), C#, CUDA, Python, SQL, TypeScript.

Systems and Performance: Multithreading, SIMD, cache-aware data structure design, memory management and custom allocators, parallel computing, low-level optimisation.

Profiling and Debugging: perf, Tracy, Nsight, RenderDoc, gdb, Valgrind.

GPU and Parallel Compute: CUDA, Vulkan, WebGPU, OpenGL, SPIR-V, DirectX12, GLSL/HLSL/ WGSL.

Build and Infrastructure: CMake, GoogleTest, Git, Linux, Docker, GitLab CI/CD, Azure DevOps, UE5, Unity, Node.js, MATLAB.

Machine Learning: TensorFlow, Scikit-learn, Pandas, CNNs, GANs, computer vision.

PROJECTS

Limit Order Book Matching Engine - C++ 17, CMake, GoogleTest, 2026

- Built a price-time-priority limit order book and NASDAQ TotalView-ITCH 5.0 feed handler; validated on a full real trading day(**368M messages, 11.2 GB**), reconstructing APPL's book to its exact closing quote and cross-checked against a std::map reference to an identical checksum.
- Engineered for the hot path: flat price-level arrays with O(1) top-of-book, intrusive FIFO queues, an object pool allocator(no malloc), and O(1) cancel via an intrusive hash map, sustained **~2.1M book-ops/sec** at **p50 281 ns/p99 1.7 µs**.
- Built an RDTSCP-timed benchmark harness (latency percentiles vs a std::map baseline, **~1.2x throughput** on real data); it surfaced a tail-latency spike from a linear top-of-book scan, which a hierarchical occupancy bitmap then removed.
- 25 GoogleTest cases** across book, parser and bitmap. [[GitHub Repo](#)]

CUDA Compute Kernels - C++, CUDA, 2026

- Implemented benchmarked GPU kernels on an RTX 4070; achieved **1.37x speedup** on tiled shared-memory matrix multiplication over a naive implementation through coalesced global memory access.
- Profiled occupancy, memory throughput and warp divergence under Nsight Compute to attribute performance to specific hardware limits. [[GitHub Repo](#)]

Read-Time Fluid Simulation - WebGPU, WGSL, 2026

- Implemented a Navier-Stokes (Stable Fluids) solver executing **50+ compute dispatches** per frame; sustains **60+ FPS** advecting over **1M texels** in a single-file browser application. [[Live Demo](#)] [[GitHub Repo](#)]. Also available as a wallpaper for windows through Lively Wallpaper. [[Wallpaper Download](#)]

Hardware-Accelerated Ray Tracer - C++, Vulkan, 2026

- Built a real-time ray tracer using **VK_KHR_ray_tracing_pipeline**, implemented BLAS/TLAS acceleration structures, Shader Binding Table setup and recursive traversal for shadows and reflections.
- Managed explicit **GPU synchronisation**, descriptor lifetimes and device memory allocation. [[GitHub Repo](#)]

Multi-Pass Vulkan Game Engine - C++, Vulkan, 2025

- Developed a multi-pass renderer with a PBR pipeline, GPU-driven culling, proprietary ECS architecture and vehicle physics simulation; sustained **100+ FPS** on RTX hardware. [[GitHub Repo](#)]
 - Won the Game Republic Red Kite Games **Technology Award** at the student showcase 2025 (group project). [[Award](#)]
-

EXPERIENCE

Tech Designer (Freelance) | Creative AF Games | Mar 2026 - Jun 2026 | Remote, UK

- Engineered character ability systems and gameplay mechanics in Unreal Engine 5 using **C++ and Blueprints**.
- Profiled and resolved **CPU/GPU bottlenecks** to hold frame budget targets on target hardware.
- Delivered features on time alongside design and engineering teams.

AI/ML Intern | Stella Stays | Aug 2021 - Jan 2022 | Dubai, UAE

- Built a predictive pricing model using LSTM and regression, improving base price accuracy to **10% MAPE**.
- Developed a room allocation algorithm using minimum slack scheduling and clustering, improving allocation **efficiency by 15%**.
- Implemented an NLTK sentiment analysis pipeline for review data at **80% classification accuracy**.

Game Development Intern | Sentient Labs | Jun 2020 - Aug 2020 | Dubai, UAE

- Build a WebGL game in Unity; increased **user engagement by 25%** and **average session duration by 18%**.
 - Implemented modular UI systems, reducing **UI load time by 40%**; integrated **50+ art and audio assets** at a sustained **60+ FPS**.
-

EDUCATION

MSc High-Performance Graphics and Games Engineering, Merit | University of Leeds | Sept 2024 - Dec 2025

- Dissertation: comparative performance and image-quality benchmarking of Intel Open Image Denoise against NVIDIA Real-Time Denoisers, concluding that a hybrid real-time/offline pipeline gives the best performance-fidelity tradeoff.

B.E. (Hons) Computer Science, Minor in Data Science, 8.72/10 | BITS Pilani | Sept 2018 - Sept 2022

- Final Year Project: Mobile robot path planning using deep learning and CNNs.
-

ADDITIONAL

Ticket Sales Advisor (Part-time) | Leeds United | Oct 2025 - Aug 2026 | Leeds, UK

- Resolved **time sensitive problems efficiently** in a fast-paced, high pressure environment, ensuring a positive experience.
- Delivered **calm, professional customer service**, assisting supporters with ticketing issues, collections and enquiries, showcasing **teamwork** and **communication skills**.