How to Improve System Performance Without Adding Latency

Understanding System Performance

In modern vision and AI applications, system performance isn't only about raw speed.
It's a balance between **throughput, latency, accuracy, power efficiency, and time-to-market**.

Higher resolutions, faster frame rates, and multi-camera configurations significantly increase data loads.
However, improving one parameter (such as throughput) can often compromise another, like latency or power.
The true challenge lies in **boosting overall performance while maintaining responsiveness and stability**.

For engineers looking to optimize throughput and responsiveness across high-speed vision systems, exploring Gidel's Frame Grabbers is a great starting point. These FPGA-based solutions form the foundation for many of the concepts discussed below.

What System Performance Looks Like in the Real World

In algorithmic trading, computers compete to "buy" and "sell" in microseconds. A system that reacts even one millisecond faster can secure a transaction worth millions, illustrating how performance and latency directly impact real-world outcomes.

In algorithmic trading, a millisecond can make or break a deal - a clear reminder that system performance is about much more than speed.

The Performance Dilemma

Every vision system faces a fundamental engineering trade-off: the more data you process, the harder it becomes to maintain real-time responsiveness.

This balance between **computational power and latency** defines the limits of many imaging systems, from autonomous vehicles to medical devices.
Instead of simply adding more processing power, the real breakthrough comes from **rethinking the system architecture** - designing a pipeline where each component contributes to higher performance **without slowing the response time**.

That's where **hybrid computing** comes into play.

Hybrid Computing: Low Latency FPGA Acceleration Where It Counts

Gidel's hybrid computing architecture combines **FPGA, CPU, GPU, and AI engines**, allowing each to do what it does best.

Task Type	Optimal Processor	Key Advantage
Random logic, decision branches, control loops	**CPU / GPU / AI Engine**	Flexibility for adaptive algorithms
Vector or repetitive data processing (histograms, gamma, compression)	**FPGA**	No speed limit, ultra-low latency, low power

This division enables the system to process more data in parallel, **increasing throughput without extending latency.**

In edge deployments that combine FPGA acceleration with embedded AI, solutions like the **Mini Jetson Frame Grabber family** deliver full hybrid performance within compact, power-efficient units.

Example: Hybrid Histogram Processing

A practical example of this hybrid approach is histogram-based image analysis.

- The **FPGA** performs fast number-crunching — incrementing pixel counts and building the histogram table.

- The **CPU** analyzes the completed table, detecting patterns, peaks, or applying corrections.

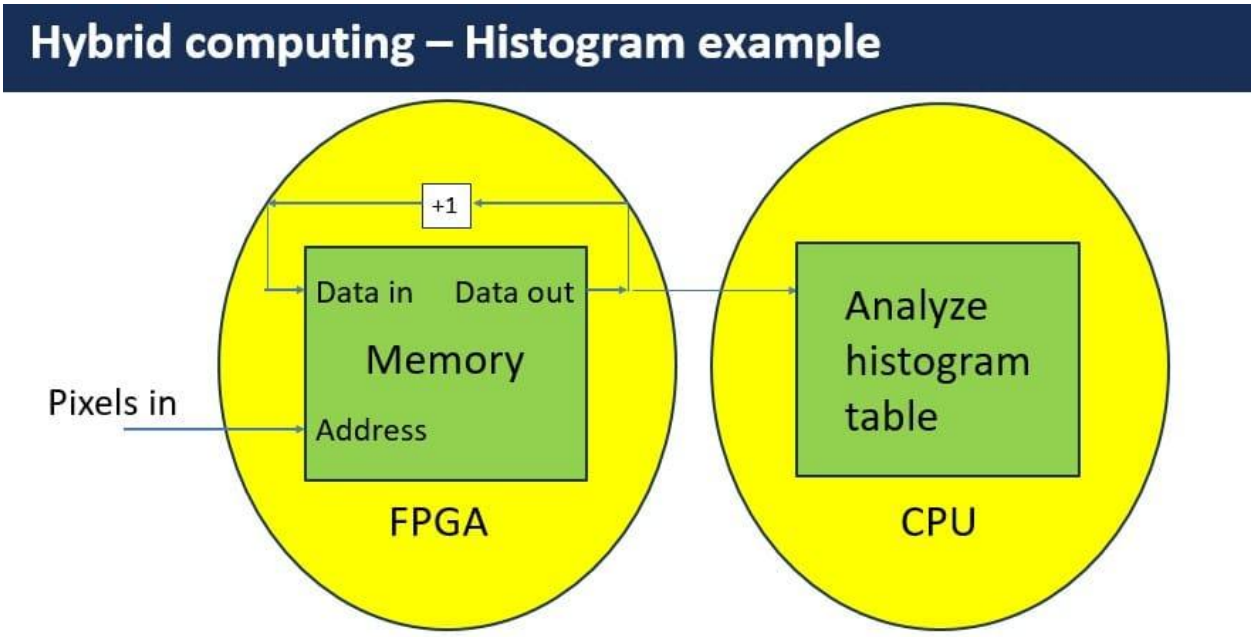

In a hybrid architecture, the FPGA handles repetitive number-crunching - incrementing pixel data - while the CPU analyzes the resulting histogram. Each processor does what it does best

Go Beyond CPU limits with Gidel Imaging Libraries

Offload heavy imaging compute to FPGA for faster processing, lower latency, and ultra-efficient real-time vision pipelines using **GIL - Gidel Imaging Libraries.**

One of the strongest demonstrations of this concept is Gidel's **HDR IP (High Dynamic Range Correction) image processing**.
Traditionally, HDR requires multiple exposures, reducing frame rate and increasing latency. Gidel's FPGA-based processing pipeline eliminates that bottleneck by performing multiple steps simultaneously:

- Single-frame HDR output
- Gamma correction and white balance
- Optional edge enhancement for better object detection
- On-FPGA JPEG compression - Explore Gidel's **Image Compression IPs** to learn more about real-time FPGA-based compression and data reduction technologies.

- Processing speeds exceeding **1 gigapixel per second**

The result is **superior image quality at real-time speed - without increasing latency.**

To see how this works in practice, explore the High Dynamic Range (HDR) IP, which performs single-frame HDR with real-time gamma, white-balance, and enhancement directly on the FPGA.

Original Image >100MP - Real-Time Processing

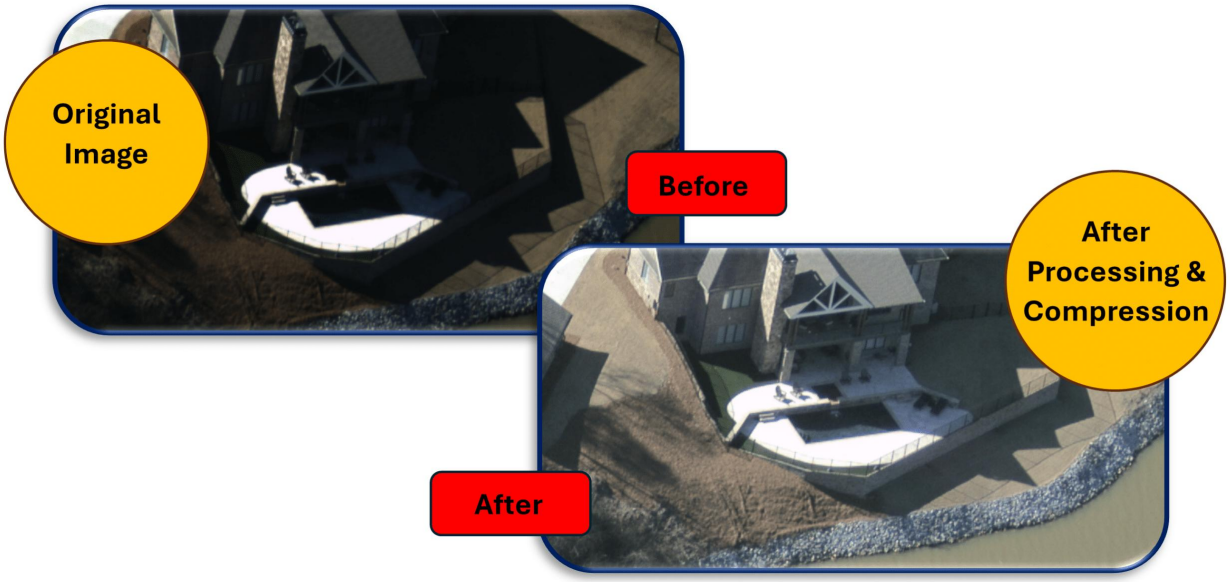

Beyond Raw Speed - The Real Drivers of System Performance

Reducing Development Cycles

True system performance isn't limited to runtime metrics.
Long development and validation loops can slow innovation just as much as inefficient code.

Gidel's modular FPGA environment, combined with tools such as the **Camera Simulator**, allows engineers to accelerate every stage of development - from design to validation.
Running live video streams through FPGA hardware shortens test cycles, reduces rework, and helps teams reach proof-of-concept and production faster.
Shorter development time means products reach the market sooner, an often overlooked but critical part of overall performance.

Power Efficiency as a Competitive Advantage

As data rates and AI workloads grow, power efficiency becomes a defining factor in performance.
Unlike traditional CPUs or GPUs, FPGAs execute logic in parallel hardware paths, minimizing wasted cycles and heat generation.

This results in lower energy consumption, higher stability, and consistent real-time operation - even under intensive workloads.
Power efficiency doesn't just reduce cost; it ensures reliability and predictability, especially in mission-critical or embedded environments.

Scalability for Future Demands

Performance requirements rarely stay constant.
Systems that can scale - in resolution, frame rate, or algorithmic complexity - maintain their value over time.

Gidel's hybrid architecture allows incremental scaling: new FPGA modules can be integrated alongside existing CPU or GPU systems without redesigning the entire pipeline.
This flexibility helps companies keep up with evolving demands in vision, robotics, and AI without compromising latency or stability.

The Bottom Line - Faster, Smarter, and Ready for Tomorrow

Boosting performance without increasing latency isn't about adding more power — it's about using the right power in the right place.
By combining FPGA acceleration, modular design, and hybrid computing, Gidel helps engineers achieve measurable improvements across all key metrics:

- **Higher throughput** at real-time speeds
- **Deterministic low latency** under any workload
- **Power efficiency and scalability**
- **Reduced time-to-market**

Organizations that adopt this architectural approach gain more than performance - they gain agility and long-term competitive advantage.

Gidel's solutions, including **Frame Grabbers**, **FPGA Modules**, **FantoVision Edge Computers**, and **Camera Simulators** providing the building blocks to create faster, smarter, and more efficient imaging and vision systems.
Ready to see what your system can do when latency is no longer the limit?
Request a demo or **contact our team** to explore how Gidel's FPGA-powered solutions can accelerate your next-generation vision platform.

Sky Blue Microsystems GmbH
Geisenhausenerstr. 18
81379 Munich, Germany
+49 89 780 2970, info@skyblue.de
www.skyblue.de