



High-Performance Trading with FPGA Accelerators, Low Latency NICs, and Server- Class Processors


Vasudevan Visvanathan
Product Marketing

STAC-T0 benchmark preliminary results¹ (AMD + Exegy)

Why implement the benchmark?

- AMD performance claims quantified and validated
 - Specified application
 - Specified measurement method
 - Independent verification

Actionable Latency	UL3524 (GTF)	Notes
Minimum (Message A)	13.9ns	Includes Transceiver and MAC/PCS round trip, algorithm and Ethernet protocol.



1: Performance analysis performed by Exegy and audited by STAC Research. Results are preliminary and released with permission by STAC Research. Claims have not been independently verified by AMD. (Performance and latency are impacted by a variety of variables. Results are specific to Exegy IP and may not be typical. GD-181)

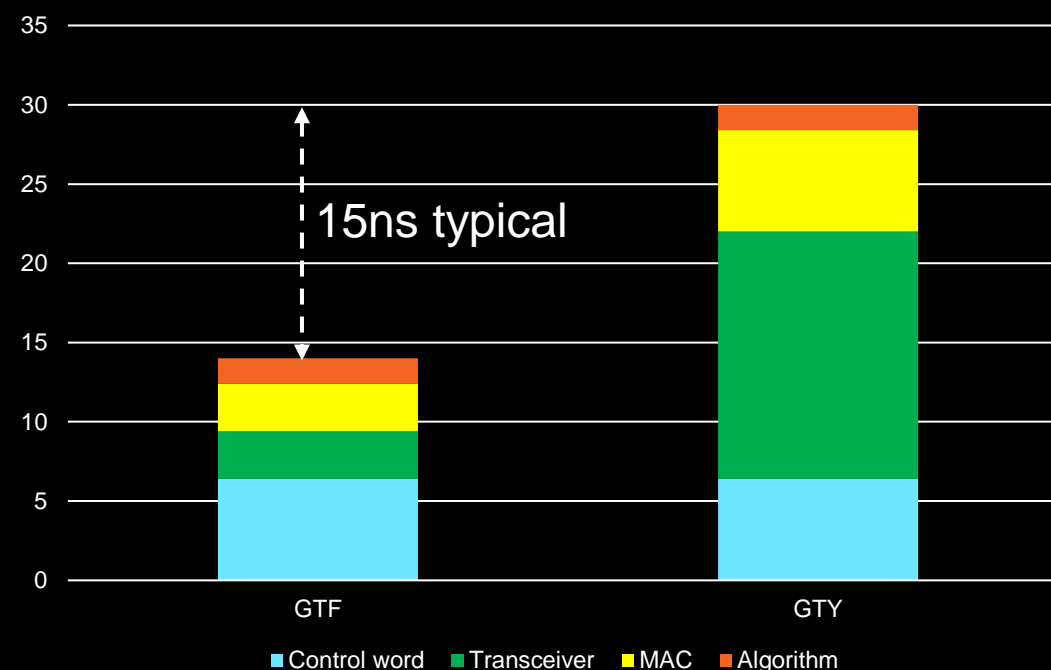
GTF versus GTY in T2T application (NOT STAC BENCHMARK)

Hypothetical analysis to illustrate latencies in T2T system
(Major latency contributors only)

Component	GTF	GTY	Assumptions
Ethernet Protocol (blue)	6.4ns	6.4ns	Identical
Transceiver latency (green)	2.34ns	15.6ns	AMD qual numbers
MAC (Yellow)	3.1ns	6.4ns	4 clks 644M soft MAC
ALGO (Red)	1.5ns	1.5ns	Identical fast Algo
Total	13.44ns	30ns	

*Not STAC
Benchmark*

Application latency delta GTF and GTY



How much lower T2T latency can be achieved?

Note: Hypothetical comparison based on common assumptions of latency of standard functions. Claims have not been independently verified by AMD. (Performance and latency are impacted by a variety of variables. GD-181)

Introducing the Alveo™ V80 for Big Data Compute

- **First Broad-Market Versal™ Accelerator Card**
Data analytics, Fintech, HPC, compute, storage, networking
- **2X Performance vs. Previous Gen¹**
2X Memory Bandwidth, 2X Logic Density
- **“Vivado™ First” Design Flow**
New example design (“AVED”) for ease of bring-up




Launched May 8, 2024

AMD Solarflare™ Low-latency Ethernet Adapters

Field-proven for Fintech networking



1x 100GbE



2x 10/25GbE



4x 10/25GbE

Proven Technology & Robust Roadmap

Deployed worldwide for 15+ years to 2500+ customers

Continuous improvements

Worldwide Support for Critical Applications

24/7 Support to meet compliance requirements

Long Term Support available

Custom Designed Performance Adapters

Targeted silicon design

Low-latency and high-performance hardware

Open Source, Robust Software Stack

Open Onload BSD API-compliant kernel bypass

TCPDirect optimized BSD API kernel bypass

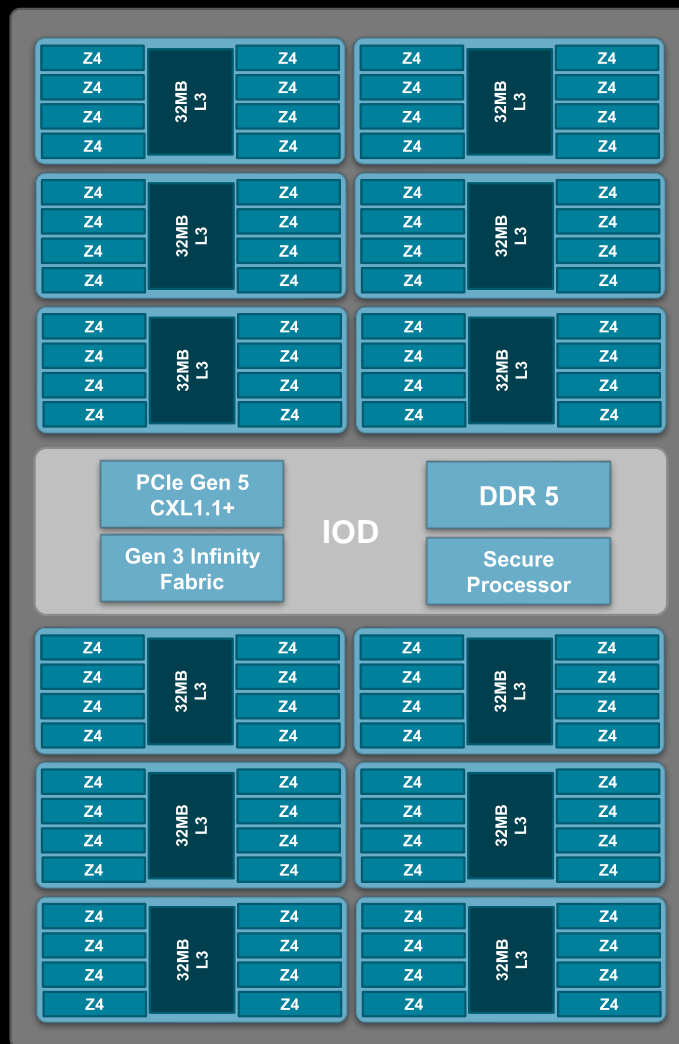

ef_vi link layer kernel bypass

4th Gen AMD EPYC CPUs for FinTech

AT A GLANCE

Trading Simulation & Research

- **Unprecedented core density**
To maximize throughput for embarrassingly parallel workloads (e.g., *Monte Carlo methods*)
- **Large L3 caches**
To improve performance for cache-bound applications (e.g., *multifactor PDEs*)
- **Updated IOD & AMD Gen3 Infinity Fabric™**
To reduce latency for thread-distributed calculations & to enable higher CPU↔GPU bandwidth for heterogeneous applications (e.g., *deep learning*)



Machine Learning & Deep Learning

- **New instruction set extensions**
BFLOAT16, VNNI, AVX-512 for improved AI capability
- **Open-source software ecosystem**
Optimized AMD libraries for classical machine learning & inference

Cloud Availability

- 4th Gen AMD EPYC deployed across major CSPs
- Well-suited for burst workloads (e.g., *regulatory scenarios*)



AMD 

Disclaimer & Attribution

Timelines, roadmaps, and/or product release dates shown in these slides are plans only and subject to change.

The information contained herein is for informational purposes only and is subject to change without notice. While every precaution has been taken in the preparation of this document, it may contain technical inaccuracies, omissions and typographical errors, and AMD is under no obligation to update or otherwise correct this. Advanced Micro Devices, Inc. makes no representations or warranties with respect to the accuracy or completeness of the contents of this document, and assumes no liability of any kind, including the implied warranties of noninfringement, merchantability or fitness for particular purposes, with respect to the operation or use of AMD hardware, software or other products described herein. No license, including implied or arising by estoppel, to any intellectual property rights is granted by this document. Terms and limitations applicable to the purchase or use of AMD's products are as set forth in a signed agreement between the parties or in AMD's Standard Terms and Conditions of Sale.

©2024 Advanced Micro Devices, Inc. All rights reserved. AMD, the AMD Arrow logo, Alveo, EPYC, Solarflare, and combinations thereof are trademarks of Advanced Micro Devices, Inc. Other product names used in this publication are for identification purposes only and may be trademarks of their respective owners.