



IS FPGA ACCELERATION OF FINANCIAL ANALYTICS VIABLE?

STAC Summit - Fall 2018

Agenda

Historical Challenges of FPGA Development

Ingredients Needed to Move FPGAs into the Computing “Mainstream”

What Financial Market Needs and What FPGAs Do Well

Two Case Studies:

- Partial Differential Equation (PDE) Solver
- In-line Analytics Processing

Make, Buy, Abstraction, Performance

Historical Challenges of FPGA Development

```
library ieee;
use ieee.std_logic_1164.all;
use ieee.numeric_std.all;

entity count is
  port
  (
    clk      : in std_logic;
    reset    : in std_logic;
    enable   : in std_logic;
    q        : out integer range 0 to 255
  );
end entity;

architecture rtl of count is
begin
  process (clk)
    variable cnt : integer range 0 to 255;
  begin
    if (rising_edge(clk)) then
      if reset = '1' then
        -- Reset the counter to 0
        cnt := 0;
      elsif enable = '1' then
        -- Increment the counter if counting is enabled
        cnt := cnt + 1;
      end if;
    end if;

    -- Output the current count
    q <= cnt;
  end process;
end rtl;
```

RTL Development

Behavioural Simulation

Synthesis


Timing Constraints

Layout


Place & Route

Verification

In System Debug



Ingredients needed to make FPGA in computing “mainstream”



Ecosystem Partners & Integrators

OEM Qualification

User Application

Data Framework (e.g. Apache Spark*)

Scalable Functions (e.g. PDE Solver)

Library Primitives (Mathematics, Statistics)

Developer SDKs - e.g. OpenCL*

Acceleration Stack, Drivers, BSPs & Interface IP

Boards and Platforms

FPGA Silicon & FPGA Design Tools

*Other names and brands may be claimed as property of others
OpenCL and the OpenCL logo are trademarks of Apple Inc. used by permission by Khronos.

Financial Market Needs

6 Key requirements for an accelerator:

- Equivalence of results
- High-Performance
- Scalability
- Parameterizable
- Easy to integrate and maintain
- Resilience



*Environment of regulatory and compliance workloads increase demand for HPC in finance
FSI customers evaluating Acceleration & Cloud*

What FPGA Does Well

- Custom Processing Pipelines
- Variable Precision Arithmetic
- Heterogenous Dataflows
- Diverse Memory Hierarchy
- Multiple Workloads
- High-bandwidth data caching
- Look-aside AND In-line Acceleration
- Parallel processing

LOW LATENCY

INHERENTLY PARALLEL

HIGH PERFORMANCE

VARIABLE PRECISION

REPROGRAMMABLE

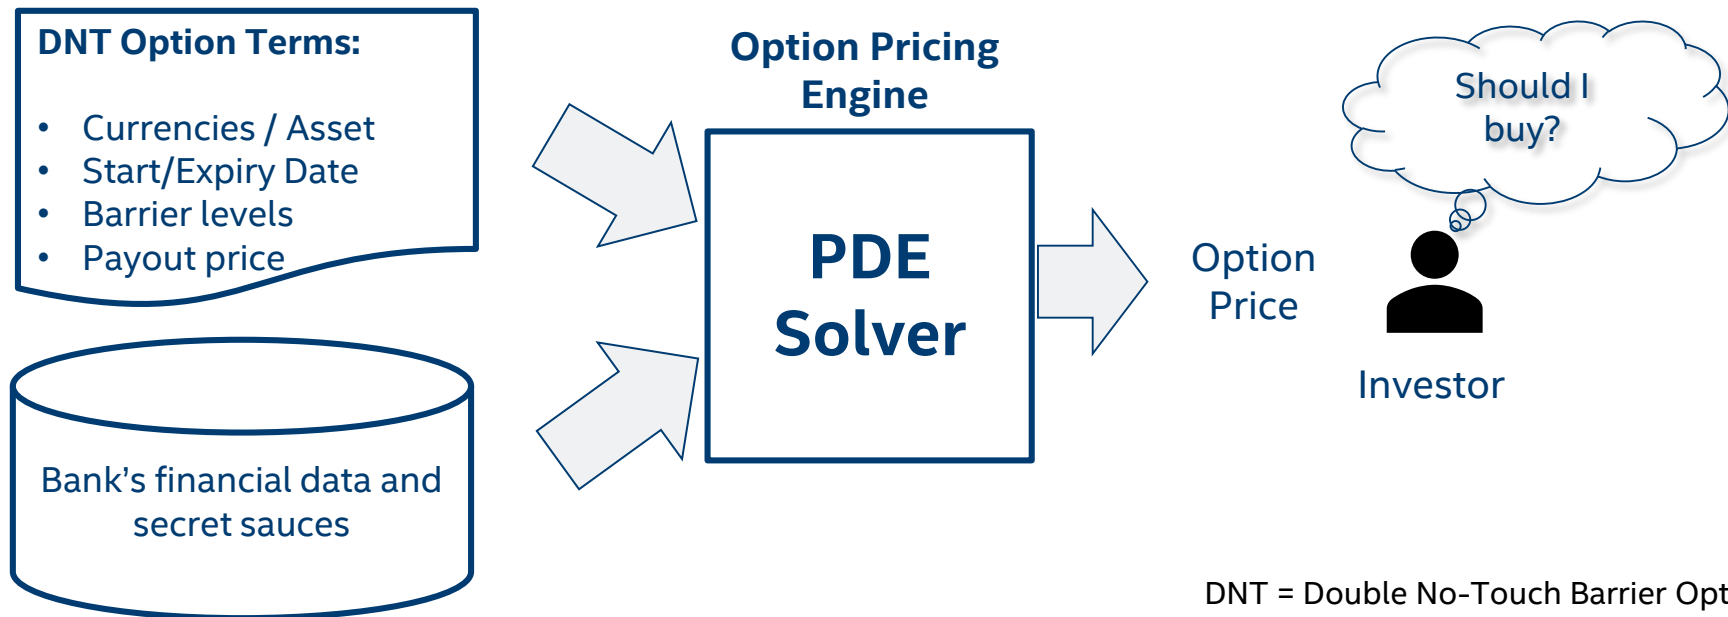
ENERGY EFFICIENT



CASE STUDIES

PDE Solver – Double No-Touch (DNT) Option Pricing Engine

How much should the bank charge the Investor as **Option Price**?



DNT = Double No-Touch Barrier Option
PDE = Partial Differential Equation

The Math(s)

Intuitively, algorithm requires knowing how **option values change w.r.t. time**

This is modeled by parabolic PDE:

$$\frac{\partial V}{\partial t} = a \frac{\partial^2 V}{\partial x^2} + b \frac{\partial V}{\partial x} + c \frac{\partial^2 V}{\partial y^2} + d \frac{\partial V}{\partial y} + e \frac{\partial^2 V}{\partial x \partial y} - rV$$

Where:

- x is spot level, y is volatility, t is time
- Coefficients a, b, c, d, e, r are **dependent** on x, y, t, and in the form:

$$w = w_t(t) + w_x(x, t)w_y(y, t)$$

Creating a PDE Solver in FPGA

FPGA used to provide a solver for a particularly computationally challenging workload

- **Intent:** Improve time to results
 - More results (Present Value of Options (PVs)) in a given amount of time or compute resource
- **Starting point:** C-model implementation of PDE Solver created (880 lines of C code)
- **End point:** Optimised OpenCL* implementation (920 lines of OpenCL)

Task	Dev. Time	Result
Convert C model to OpenCL	2 weeks	142 PV/s**
Optimise pipeline	1 week	174 PV/s**
New C Code + Open CL optimisations	2 weeks	387 PV/s**
Scale infrastructure (4 x FPGA Cards)	1 week	1511 PV/s**


** See appendix for server configuration

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors.

Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more complete information visit www.intel.com/benchmarks.

Performance results are based on testing as of 01 Sep 2018 and may not reflect all publicly available security updates. See configuration disclosure for details. No product can be absolutely secure.

Throughput Averaged Over 50 Consecutive Batches (1940 PDEs/batch)



Throughput scales linearly to the number of FPGAs

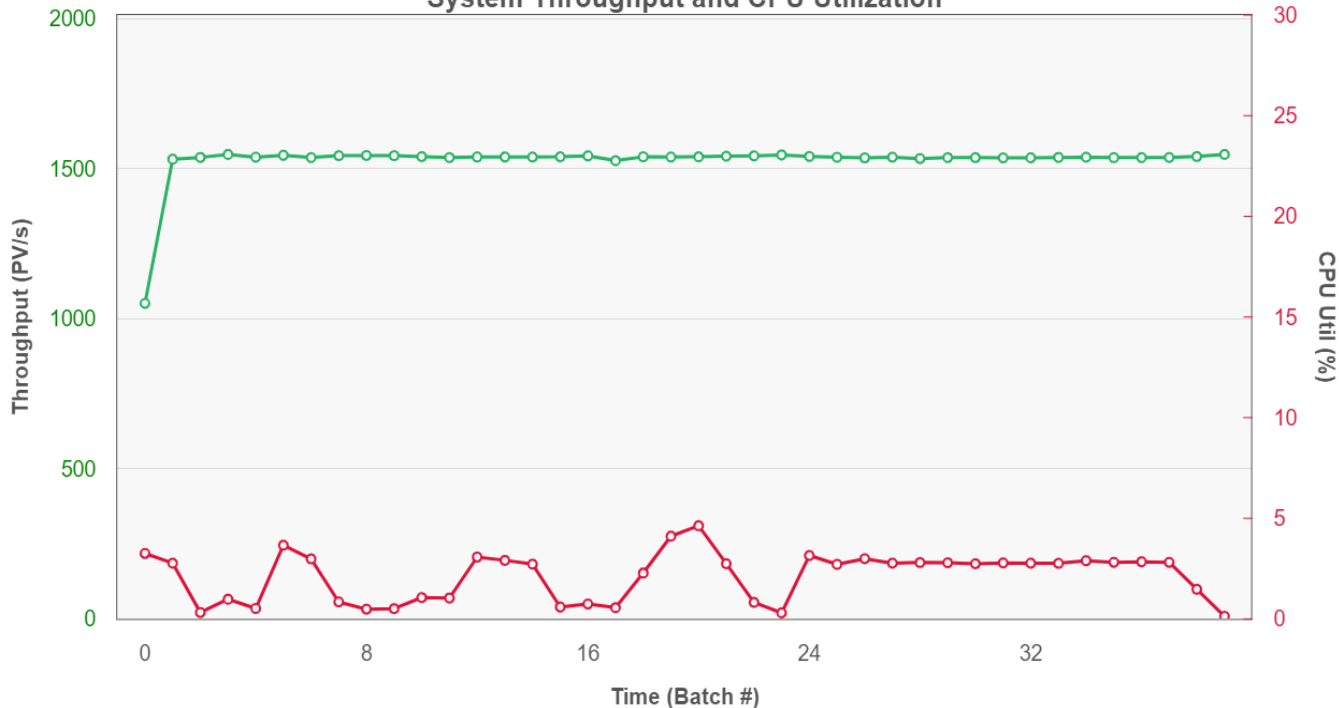
** See appendix for server configuration

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors.

Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more complete information visit www.intel.com/benchmarks.

Performance results are based on testing as of 01 Sep 2018 and may not reflect all publicly available security updates. See configuration disclosure for details. No product can be absolutely secure.

System Throughput and CPU Utilization



Number of FPGAs:
4

Avg Throughput:
1526 PV/s

Max CPU Util:
4.63 %

Not STAC Benchmarks

Peak throughput is sustainable and can be reached quickly

** See appendix for server configuration
Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more complete information visit www.intel.com/benchmarks.
Performance results are based on testing as of 01 Sep 2018 and may not reflect all publicly available security updates. See configuration disclosure for details. No product can be absolutely secure.



Flexible Allocation of FPGA Devices

Same Total Throughput



Better Fault Tolerance

Smaller Hardware Footprint

In-Line Data Processing Example

FPGA offers the possibility to process ingress market data before sending to the CPU

TCP offload in the FPGA reduces host software load

Integrating pricing / risk / compliance calculations offers:

- Near real-time calculations
- Implement custom functions in C/C++ with high-level design tools
- Run SW applications and frameworks on CPU cores

Enyx* nxTCP function integrated into FPGA

Enyx* nx_tcp SW libraries plug into Acceleration Stack

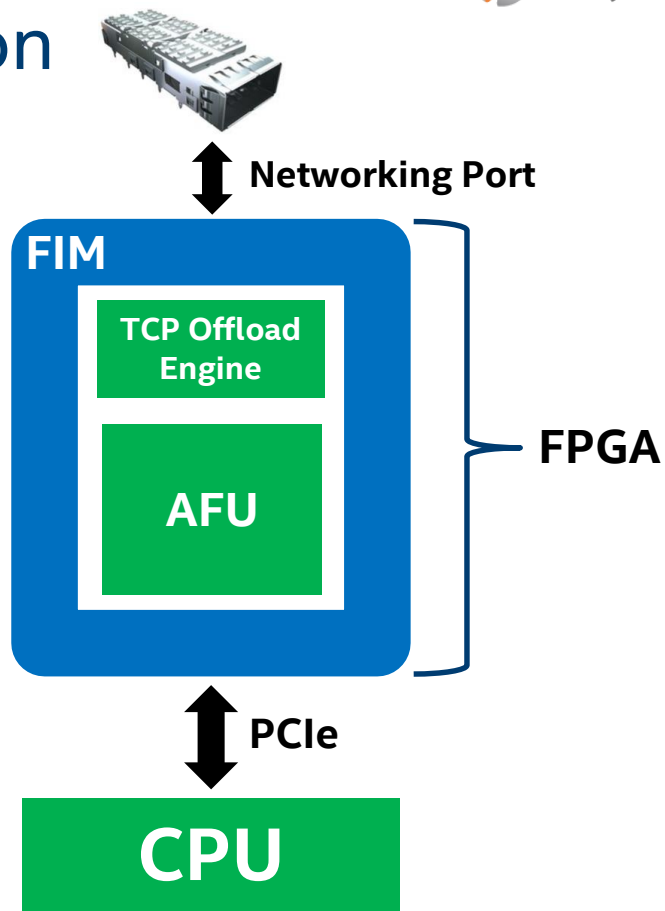
- UDP and Market Feed Decoder are available

*Other names and brands may be claimed as property of others

In-Line Accelerator Implementation

Hardware integration:

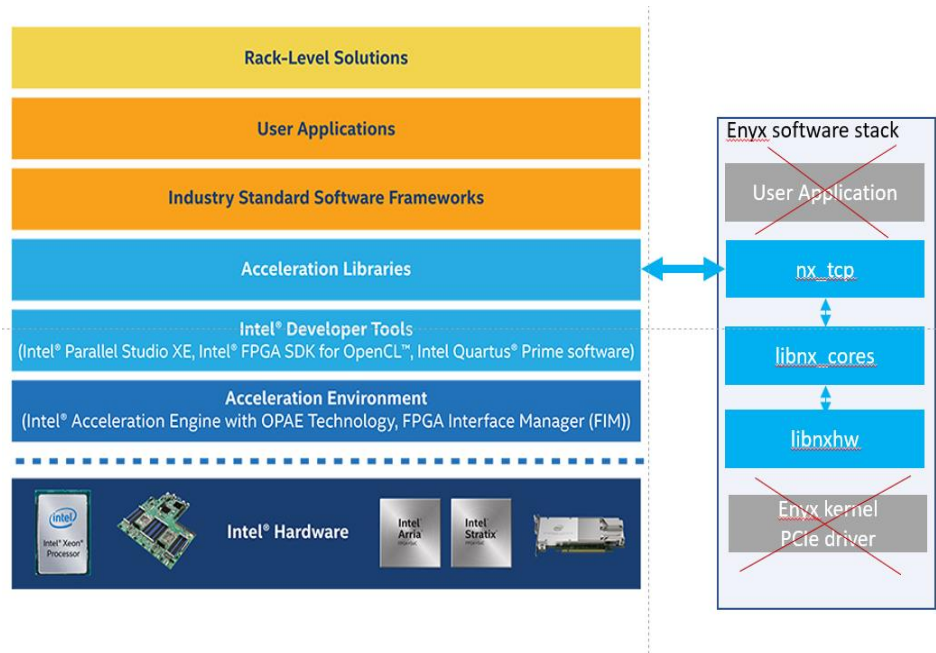
- FPGA Interface Manager (FIM) is part of Intel® Acceleration Stack for Intel® Xeon® CPU with FPGAs
- TCP Offload Engine from Enyx* is provided ready for Accelerator Functional Unit (AFU) attachment (loopback mode)
- Design AFUs with HLS



*Other names and brands may be claimed as property of others

Enyx* TCP offload plugs into Intel® Acceleration Stack for Intel® Xeon® CPU with FPGAs

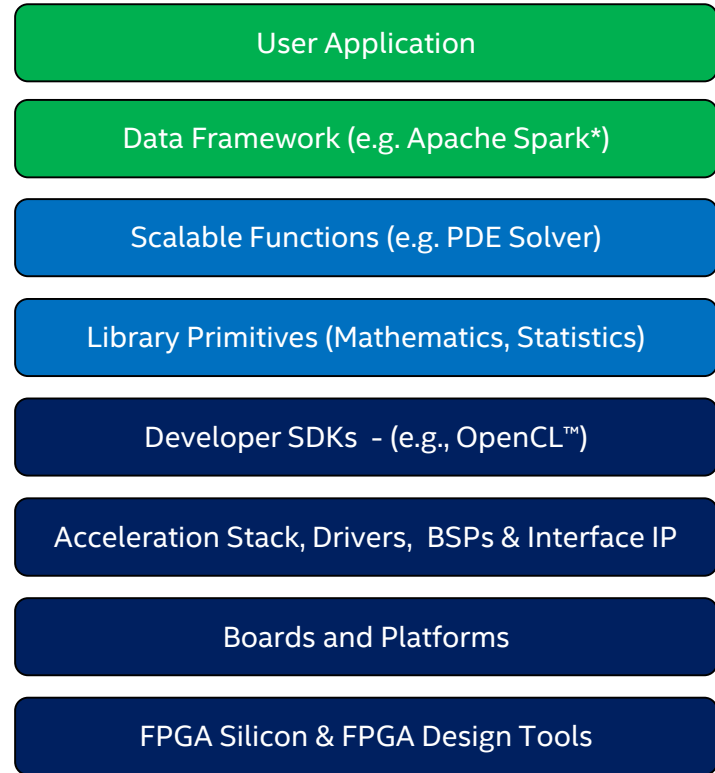

- Software integration:
 - TCP/OE management is linked to the OPAE software stack
 - External library plug in
 - OPAE is a robust, open software stack that includes:
 - Up-streamed Linux drivers
 - Tools to discover, enumerate, program, share, query accesses, manipulate FPGAs



*Other names and brands may be claimed as property of others

OPAE = Open Programmable Acceleration Engine

Final Thoughts on Accelerator Ease of Use



*Other names and brands may be claimed as property of others
OpenCL and the OpenCL logo are trademarks of Apple Inc. used by permission by Khronos.

**System configuration for performance testing

Server configuration:

Dell PowerEdge R740
2 x Intel® Xeon® Gold 6132 @ 2.6 GHz
192GB (12 x 16GB) RDIMM, 2666MT/s, Dual Rank

Operating System:

Red Hat Enterprise Linux: Release 7.5 with Linux kernel 3.10.0-862.el7.x86_64

FPGA:

Intel Programmable Acceleration Card with Intel Arria® 10 GX FPGA, Acceleration Stack version 1.0

Test performed during August 2018.

OpenCL code was developed within Intel® Programmable Solutions Group.

Functional correctness was verified by comparison with single-precision floating point results from CPU, using the “==” operator in C/C++

Tests were performed with pre-production, proof-of-concept code.

Not all capabilities are part of shipping products.

Legal Disclaimer

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Performance varies depending on system configuration. No computer system can be absolutely secure. Check with your system manufacturer or retailer or learn more at intel.com.

Some results have been estimated or simulated using internal Intel analysis or architecture simulation or modeling, and provided to you for informational purposes. Any differences in your system hardware, software or configuration may affect your actual performance.

Tests document performance of components on a particular test, in specific systems. Differences in hardware, software, or configuration will affect actual performance. Consult other sources of information to evaluate performance as you consider your purchase. For more complete information about performance and benchmark results, visit <http://www.intel.com/benchmarks>.

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more complete information visit <http://www.intel.com/benchmarks>.

Statements in this document that refer to Intel's plans and expectations for the quarter, the year, and the future, are forward-looking statements that involve a number of risks and uncertainties. A detailed discussion of the factors that could affect Intel's results and plans is included in Intel's SEC filings, including the annual report on Form 10-K.

All products, computer systems, dates and figures specified are preliminary based on current expectations, and are subject to change without notice. The products described may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

No license (express or implied, by estoppel or otherwise) to any intellectual property rights is granted by this document.

Intel does not control or audit third-party benchmark data or the web sites referenced in this document. You should visit the referenced web site and confirm whether referenced data are accurate.

© 2018 Intel Corporation. Intel, the Intel logo, Stratix, Arria, and Xeon are trademarks of Intel Corporation in the U.S. and/or other countries.

*Other names and brands may be claimed as property of others.

