GPU ACCELERATED HIGH PERFORMANCE DATA ANALYTICS

NVIDIA

Larry Brown Ph.D.
Federal Solutions Architecture Team Lead

NASA Goddard Workshop on Artificial Intelligence November 2018

“It’s time to start planning for the end of Moore’s Law, and it’s worth pondering how it will end, not just when.”

Robert Colwell
Retired Director, Microsystems Technology Office, DARPA

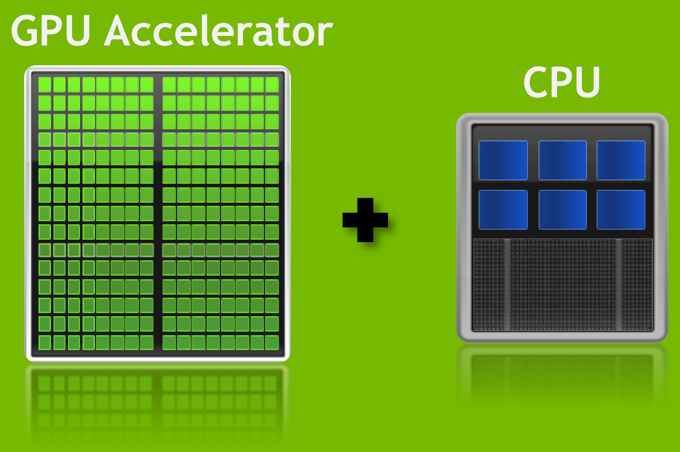

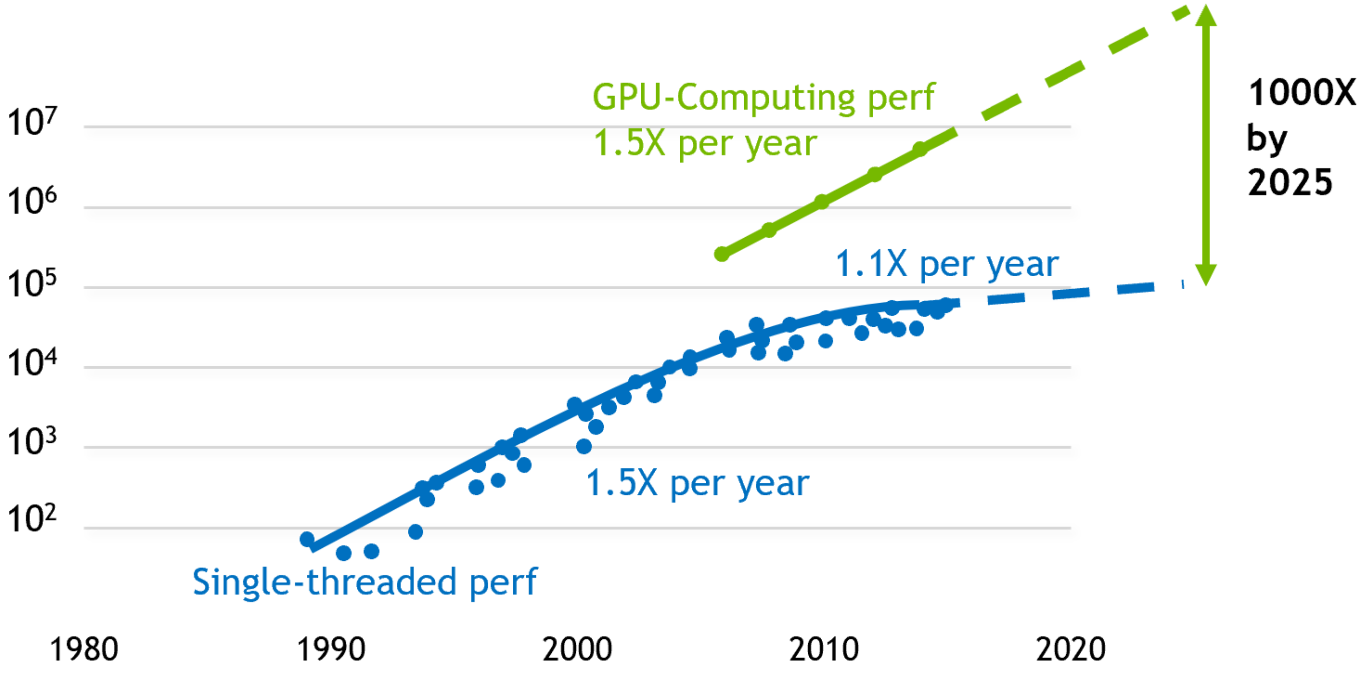

Original data up to the year 2010 collected and plotted by M. Horowitz, F. Labonte, O. Shacham, K. Olukotun, L. Hammond, and C. Batten New plot and data collected for 2010-2015 by K. Rupp

NVIDIA POWERS WORLD'S FASTEST SUPERCOMPUTERS

48% More Systems | 22 of Top 25 Greenest

ORNL Summit
World's Fastest
27,648 GPUs | 144 PF

LLNL Sierra
World's 2nd Fastest
17,280 GPUs | 95 PF

Piz Daint
Europe's Fastest
5,704 GPUs | 21 PF

ABCI
Japan's Fastest
4,352 GPUs| 20 PF

ENI HPC4
Fastest Industrial
3,200 GPUs| 12 PF

Half of Top500 compute power comes from accelerated systems.

THE NEW HPC MARKET

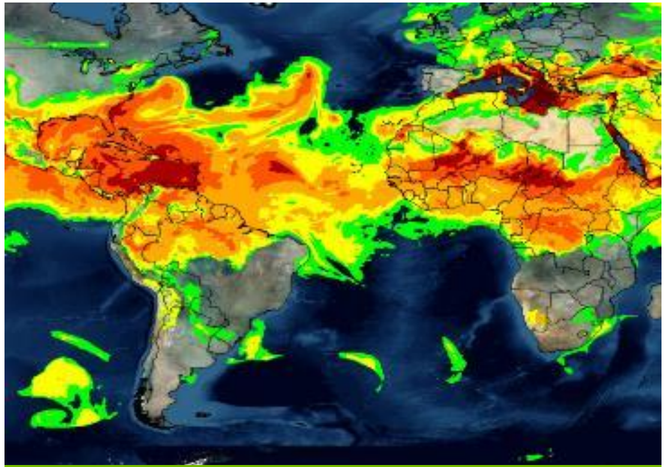

SIMULATION

MACHINE LEARNING

DEEP LEARNING

RAPIDS — OPEN GPU DATA SCIENCE

Software Stack Python

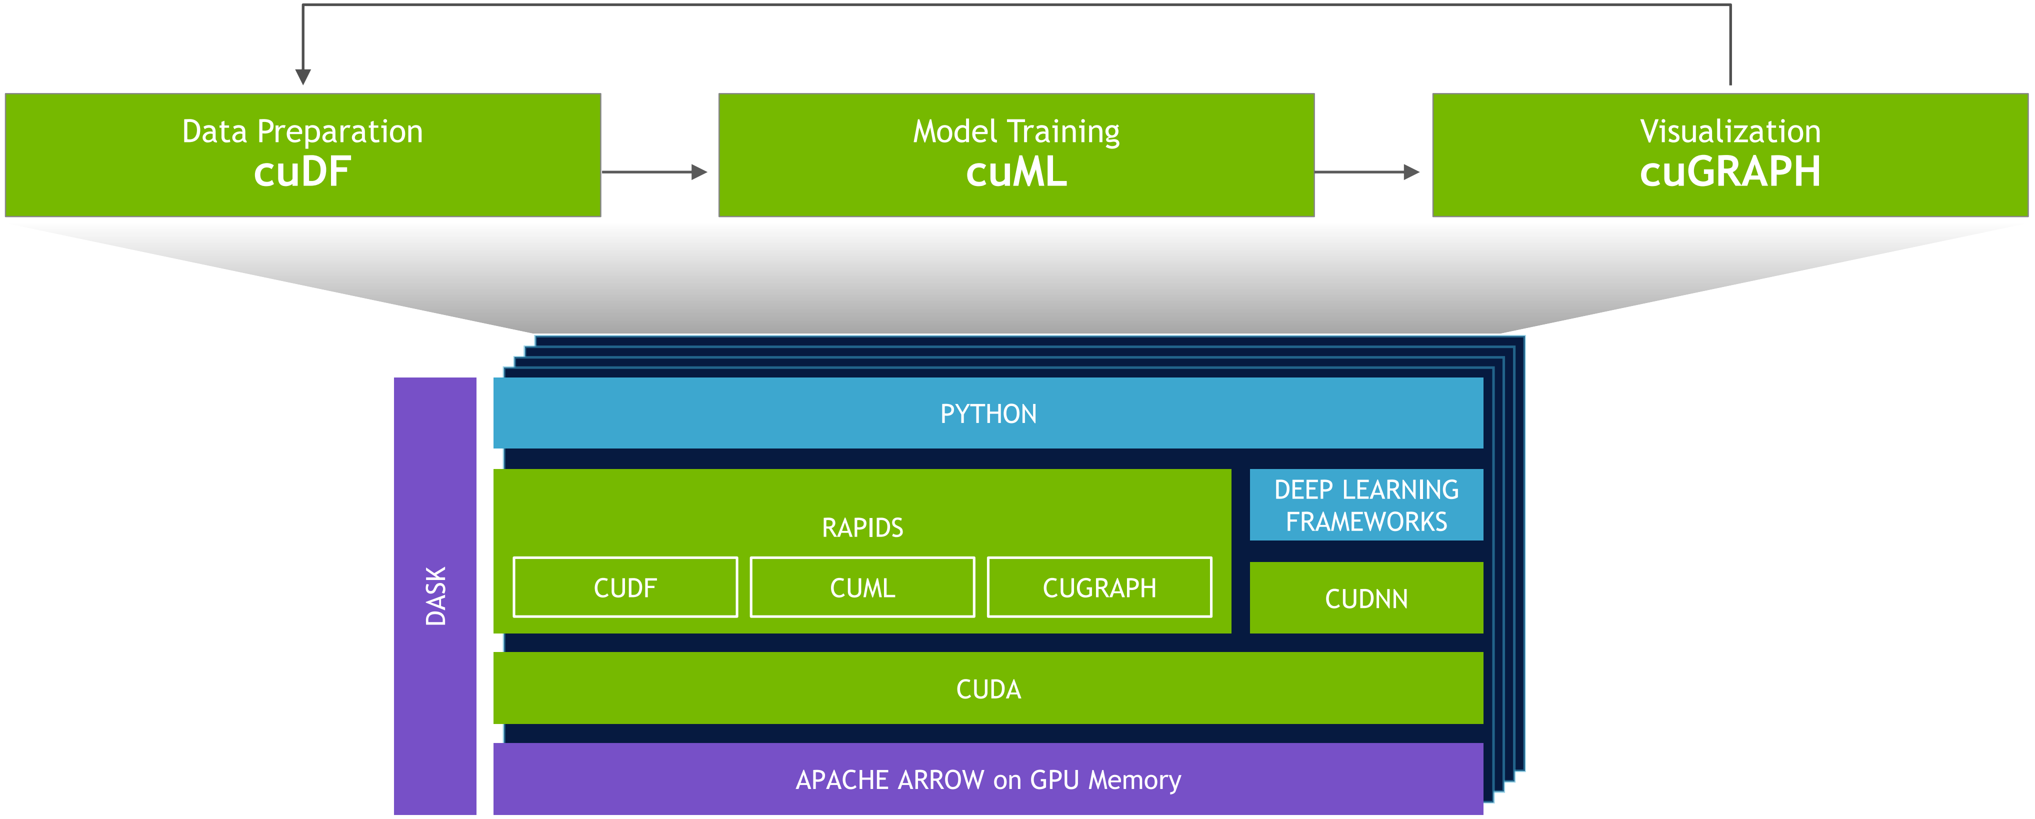

RE-IMAGINING DATA SCIENCE WORKFLOW

Open Source, End-to-end GPU-accelerated Workflow Built On CUDA

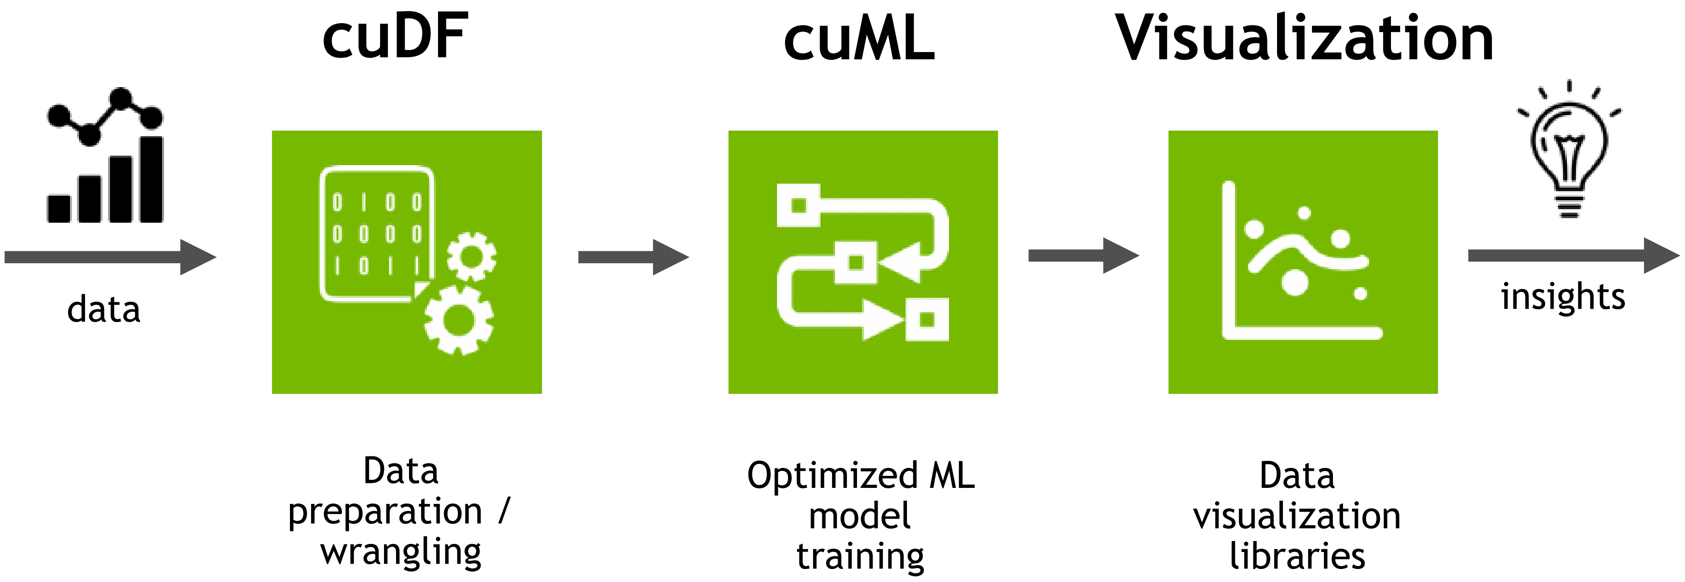

PILLARS OF RAPIDS PERFORMANCE

CUDA Architecture

Massively parallel processing

NVLink/NVSwitch

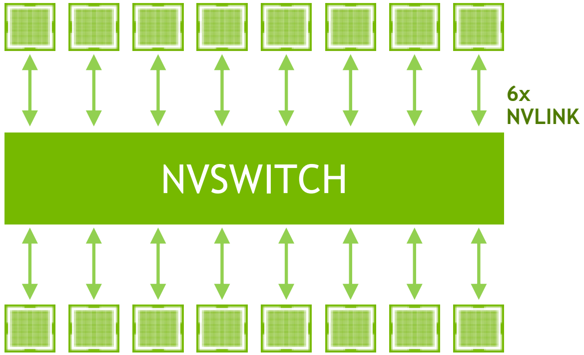

High speed connecting between GPUs for distribute algorithms

Memory Architecture

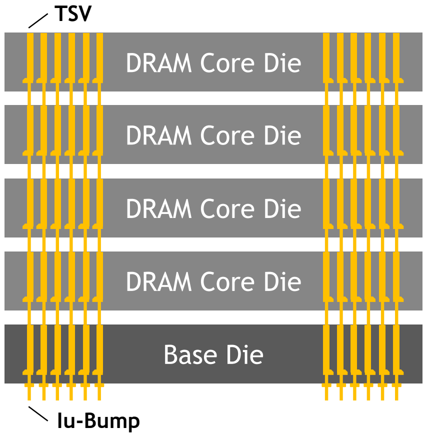

Large virtual GPU memory, high-speed memory

FASTER INSIGHTS FOR MACHINE LEARNING

HGX-2 544X Speedup Compared to CPU-Only Server Nodes

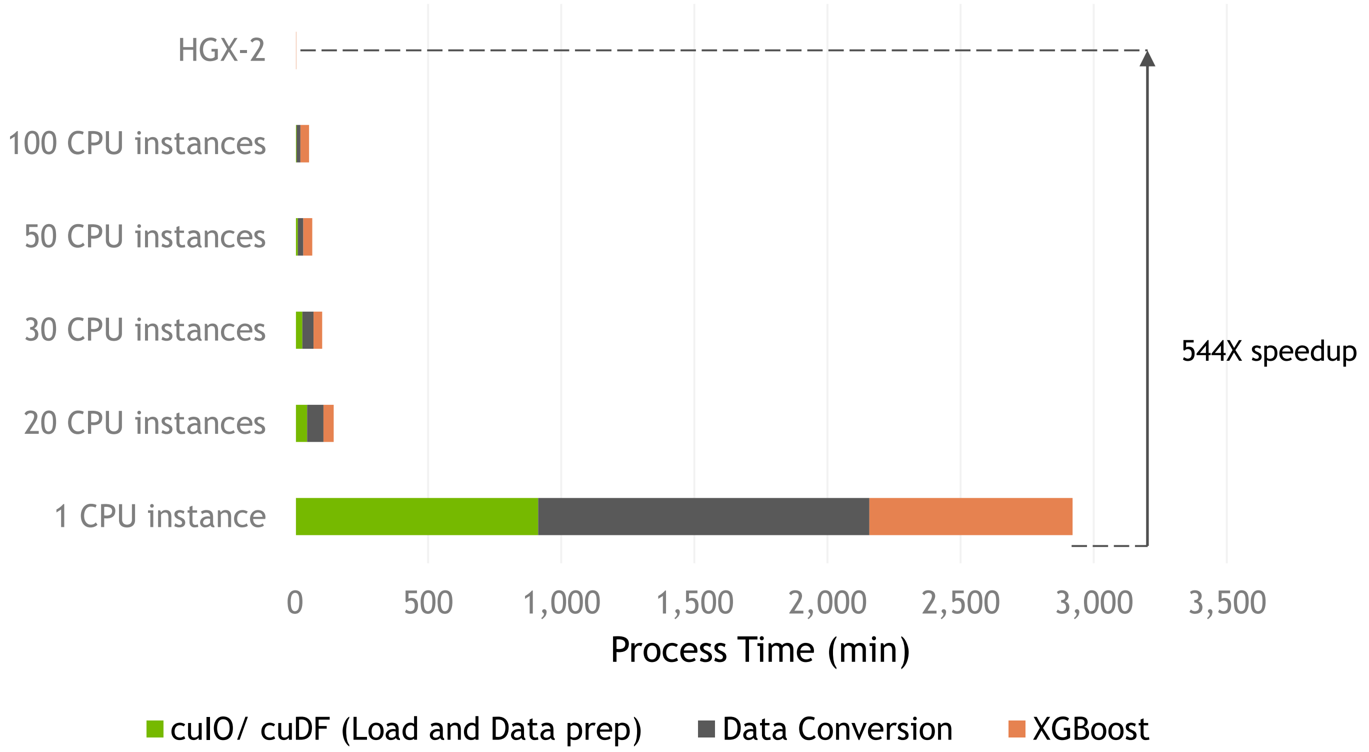

GPU Measurements Completed on DGX-2 running RAPIDS
CPU: 20 CPU cluster- comparison is prorated to 1 CPU (61 GB of memory, 8 vCPUs, 64-bit platform), Apache Spark
US Mortgage Data Fannie Mae and Freddie Mac 2006-2017 | 146M mortgages
Benchmark 200GB CSV dataset | Data preparation includes joins, variable transformations

GPU APPROACH WINS MIT GRAPH CHALLENGE

Two years in a row!

2017 GraphChallenge 2018

NVIDIA wins "Highest Performance" in the Static Graph Challenge

Triangle Counting K-Truss Counting

Static Graph Challenge on GPU

Mauro Bisson, Massimiliano Fatica
NVIDIA Corporation
Santa Clara, CA 95050, USA

***Abstract*—This paper presents the details of a CUDA implementation of the Subgraph Isomorphism Graph Challenge, a new effort aimed at driving progress in the graph analytics field. challenge consists of two graph analytics: triangle counting and k-truss. We present our CUDA implementation of the graph triangle counting operation and of the k-truss subgraph decomposition. Both implementations share the same codebase taking advantage of a set intersection operation implemented via bitmaps. The analytics are implemented in four kernels optimized for different types of graphs. At runtime, lightweight heuristics are used to select the kernel to run based on the specific graph taken as input.**

I. Introduction

to store the graph in the Compressed Sparse Row format as it is a widely used data structure for graph processing and because its properties can be effectively exploited for the triangle counting computation. Both codes treat the input graph as undirected.

For what concerns the output, the triangle counting code produces a single number that represent the exact number of triangles found in the input graph. The k-truss code instead outputs the whole subgraph whose edges contribute to at least $k-2$ triangles.

III. Triangle counting

IMPROVING DEMAND FORECASTS

Accurate demand forecasting is a critical but challenging science for retailers requiring massive amounts of data and compute cycles. Walmart is optimizing machine learning with NVIDIA RAPIDS open-source software on GPUs.

GPUs deliver 50x faster processing speed allowing Walmart to benefit from more sophisticated algorithms, reduce forecasting errors, and increase the efficiency of its supply chain.

NVIDIA

Walmart

Thank you...

NVIDIA

- > Founded in 1993
- > Jensen Huang, Founder & CEO
- > 12,000 employees
- > $9.7B revenue in FY18

"World's Most Admired Companies"
— Fortune

"50 Smartest Companies: #1"
— MIT Tech Review

"#3 Top CEO in the World"
— Harvard Business Review

"Most Innovative Companies"
— Fast Company

NVIDIA