TEXAS A&M IPU WORKSHOP

GRAPHCORE

GRAPHCORE'S SOLUTION

Hardware

IPU processor
designed for AI

Software

Poplar SDK and
development tools

Platform

IPU platforms
Available in the cloud

THE INTELLIGENCE PROCESSING UNIT (IPU)
WHAT MAKES IT DIFFERENT?

	CPU	GPU	IPU
Parallelism	Designed for scalar processing	SIMD/SIMT architecture. Designed for large blocks of dense contiguous data	Massively parallel MIMD architecture. High performance/efficiency for future ML trends
		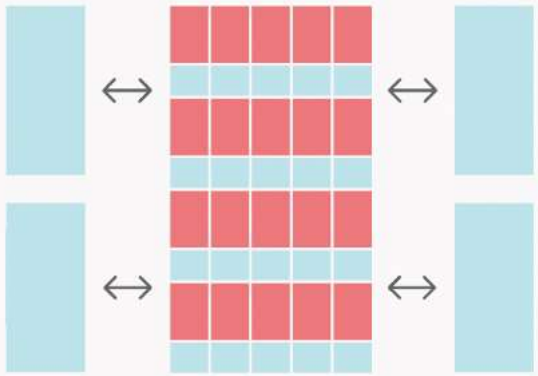	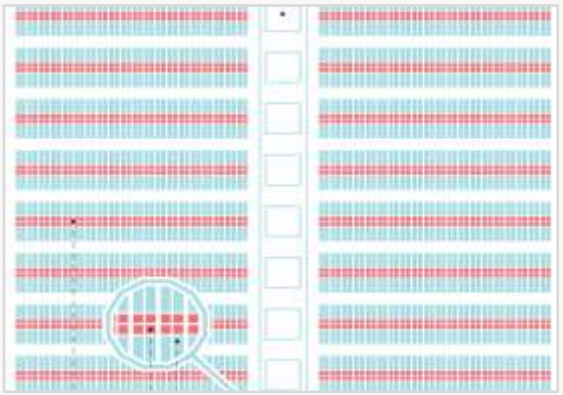
Memory Bandwidth	Off-chip memory	Model and Data spread across off-chip and small on-chip cache and shared memory (2TB/s for A100 HBM)	Main Model & Data in tightly coupled large locally distributed SRAM (~65 TB/s for Bow IPU)

EXECUTION MODEL

COMPUTATIONAL GRAPH

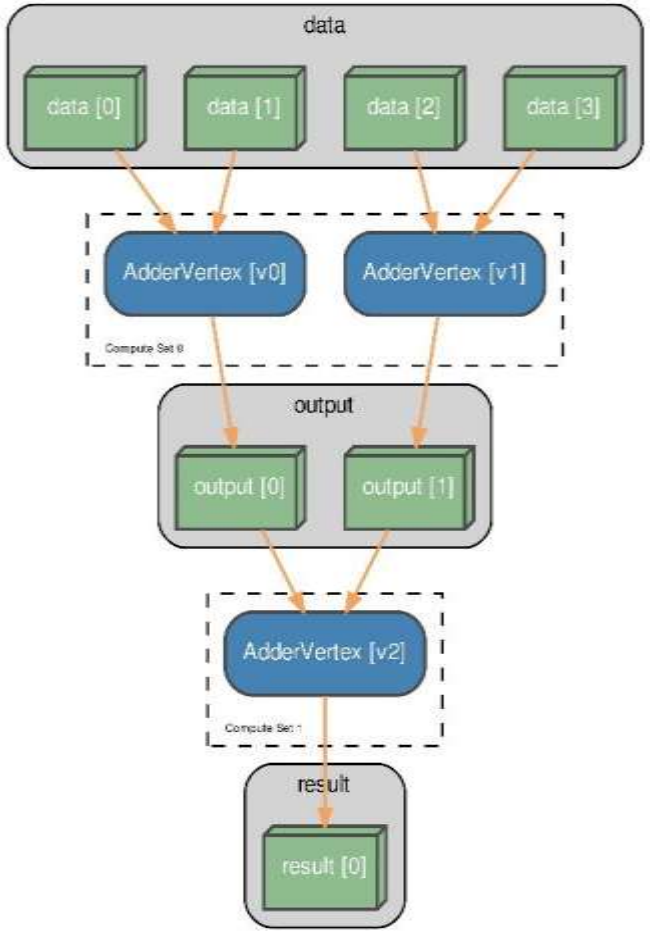

BSP SCHEDULE

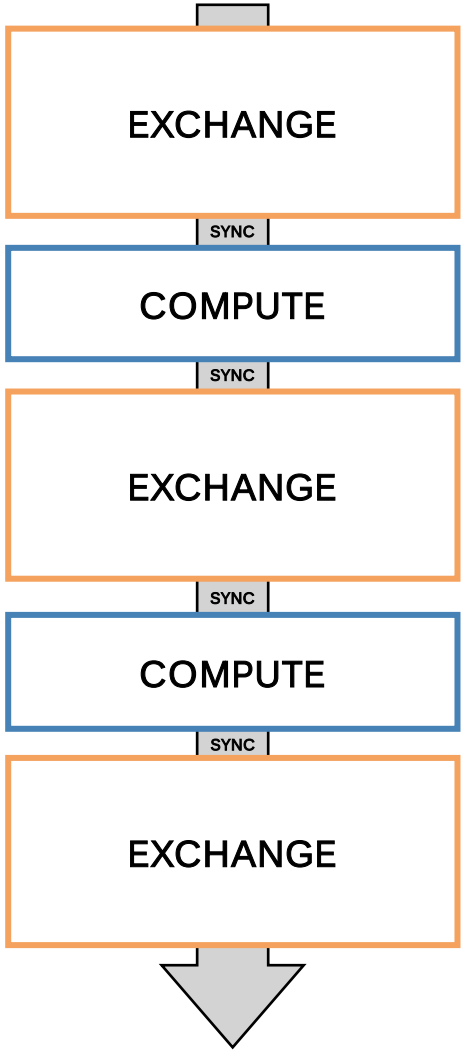

OPTIMIZED IPU EXECUTION

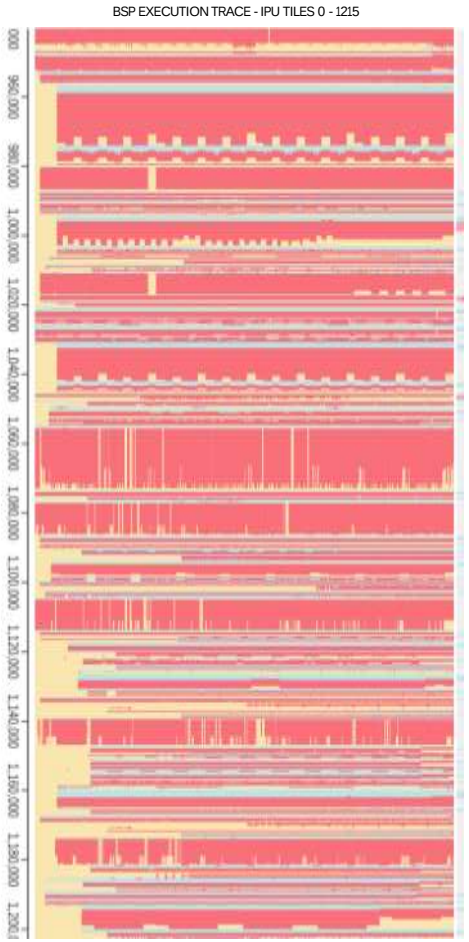

OUTPUT FROM POPVISION GRAPH ANALYSER

GRAPHCORE

SCALING ACROSS DEVICES

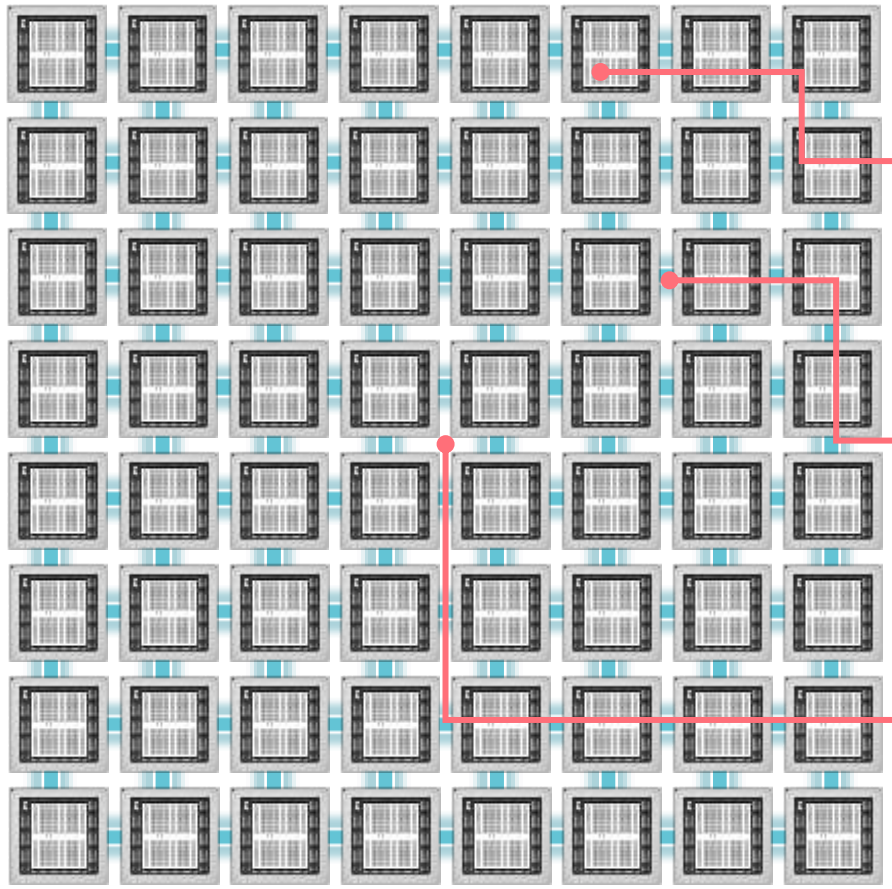

UP TO 64 IPU DEVICES USABLE AS A SINGLE LARGE IPU FROM APPLICATIONS

565248 FULLY INDEPENDENT WORKERS, 57.6GB IN-PROCESSOR MEMORY™, LEVERAGING OVER 3.8 TRILLION TRANSISTORS

EXCHANGE

- THE IPU EXCHANGE MODEL EXTENDS OFF DEVICE
  - ALLOWS TILE MESSAGING ACROSS IPU DEVICES
- SUPPORTS OPTIMIZED CROSS IPU SYNCHRONISATION
- DRIVEN BY COMPILER SUPPORT IN SOFTWARE

IPU-LINK™

- PROVIDES 320 GB/S IPU TO IPU BANDWIDTH
- SUPPORTS COMMUNICATION BETWEEN IPUS
- LAYOUT FULLY SOFTWARE CONFIGURABLE
- SUPPORTS POINT TO POINT TILE MESSAGING

GCD

- GRAPH CONTROL DOMAIN FOR APPLICATIONS
- CREATES A SINGLE LARGE IPU SOFTWARE TARGET
- FULLY CONFIGURABLE PARTITIONING OF IPUS
- BOTH DATA PARALLEL AND MODEL PARALLEL

SCALING ACROSS SYSTEMS

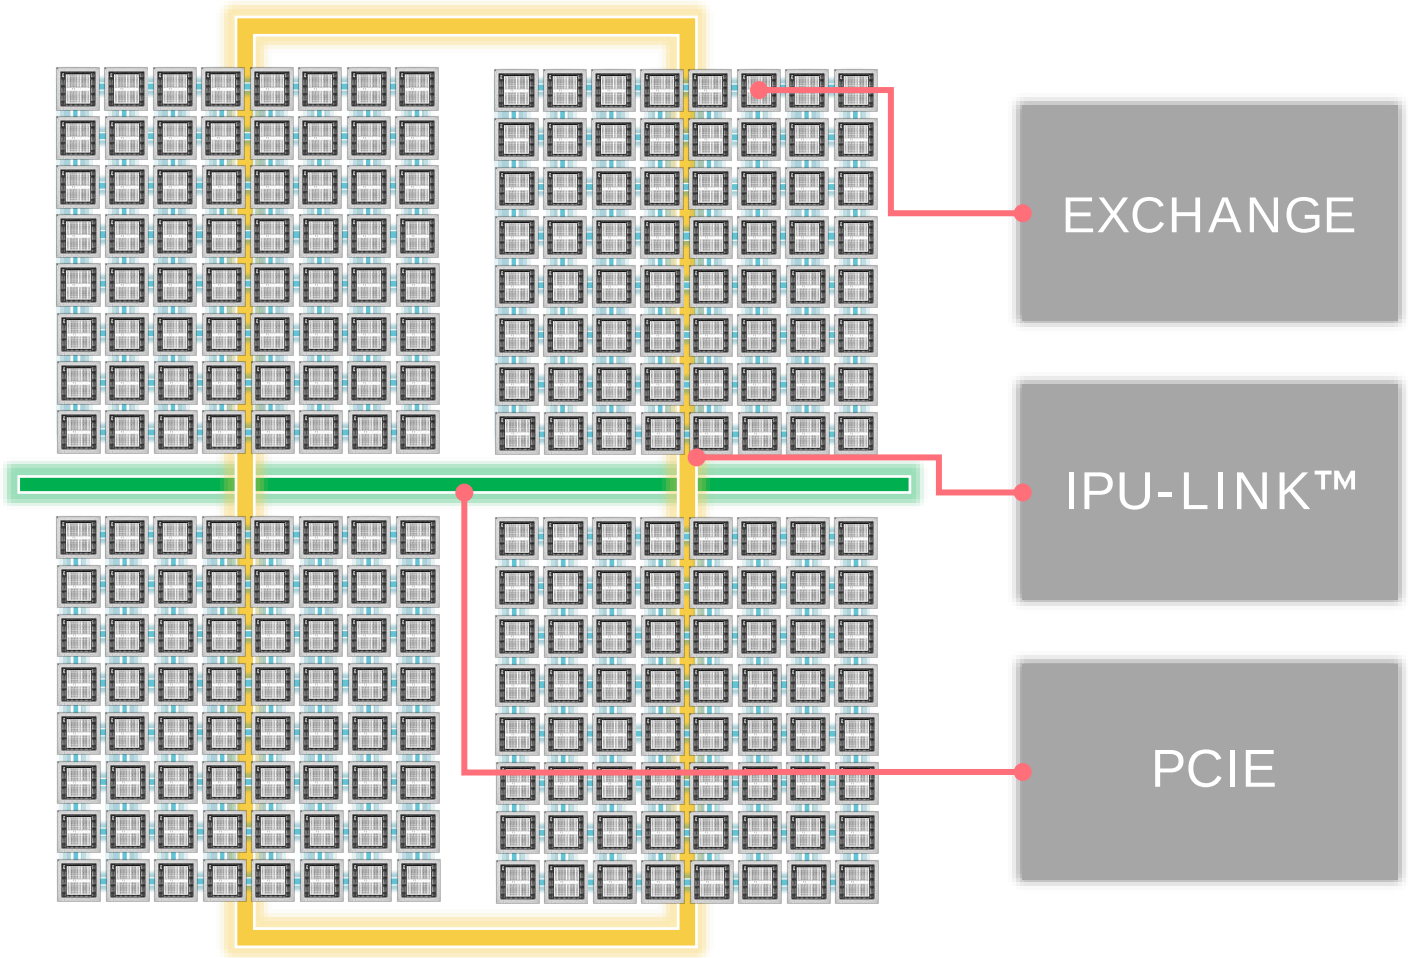

256 IPU APPLICATION TARGET BUILT FROM INTERCONNECTED 64 IPU DOMAINS

- IPU EXCHANGE SUPPORT ACROSS DOMAINS
  - DRIVEN BY COMPILER SUPPORT IN SOFTWARE
- ENABLES APPLICATION COLLECTIVES SUPPORT
- ALLOWS SCALING UP TO 64000 IPU DEVICES

- IPU LINK™ CAN BE EXTENDED ACROSS DOMAINS
- SUPPORTS OPTIMIZED IPU LINK™ COLLECTIVES
- ALLOWS REPLICATION ACROSS SYSTEMS
- SUPPORTS A STANDARD IPU SOFTWARE MODEL

- IPUS CAN ACCESS MEMORY AND DEVICES OVER PCIE
- ALLOWS INTERFACING WITH HOST BASED SOFTWARE
- APPLICATIONS CAN BUILD ON HOST NETWORKING
- ALLOWS SCALING IN STANDARD SERVER PLATFORMS

PRODUCT DETAILS

BOW IPU PROCESSOR

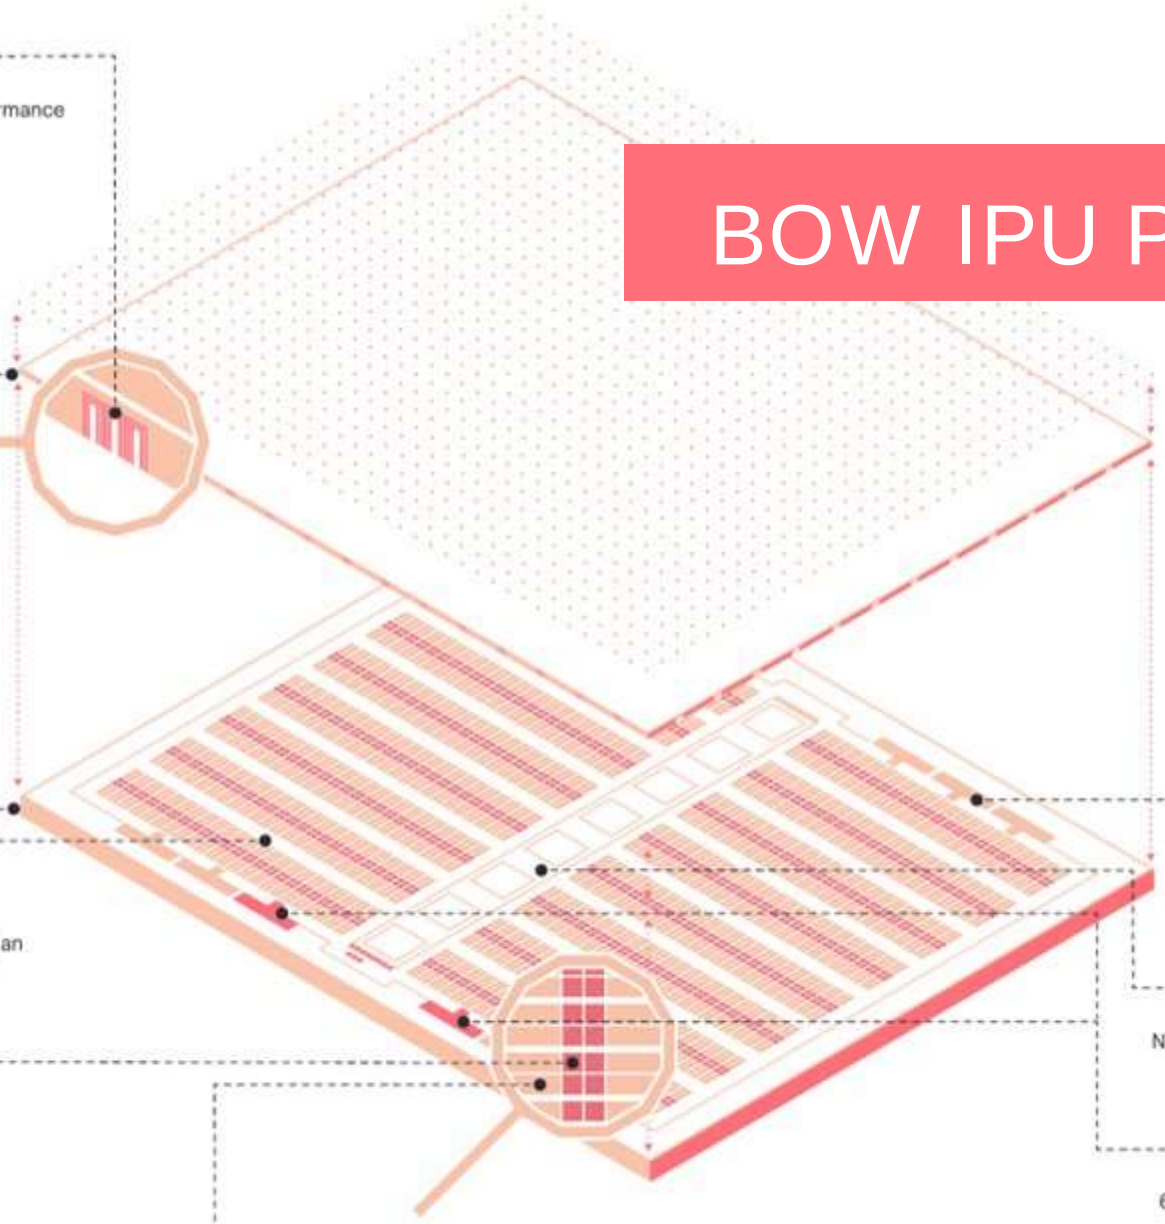

Deep Trench Capacitor

Efficient power delivery
Enables increase in operational performance

Wafer-On-Wafer

Advanced silicon 3D stacking technology

Closely coupled power delivery die

Higher operating frequency and enhanced overall performance

IPU-Tiles™

1472 independent IPU-Tiles™ each with an IPU-Core™ and In-Processor-Memory™

IPU-Core™

1472 independent IPU-Core™

8832 independent program threads executing in parallel

In-Processor-Memory™

900MB In-Processor-Memory™ per IPU

65.4TB/s memory bandwidth per IPU

IPU-Links™

10x IPU-Links,
320GB/s chip to chip bandwidth

IPU-Exchange™

11 TB/s all to all IPU-Exchange™
Non-blocking, any communication pattern

PCIe

PCI Gen4 x16
64 GB/s bidirectional bandwidth to host

BOW-2000 IPU MACHINE

1U blade form factor delivering 1.4 PetaFLOPS AI Compute

Disaggregated AI/ML accelerator platform

Excellent performance & TCO leveraging In-Processor memory & IPU-Exchange

IPU-Links scale to Bow Pod64

Expansion to Bow Pod256 and beyond with IPU-GW Links

BOW IPU-2000

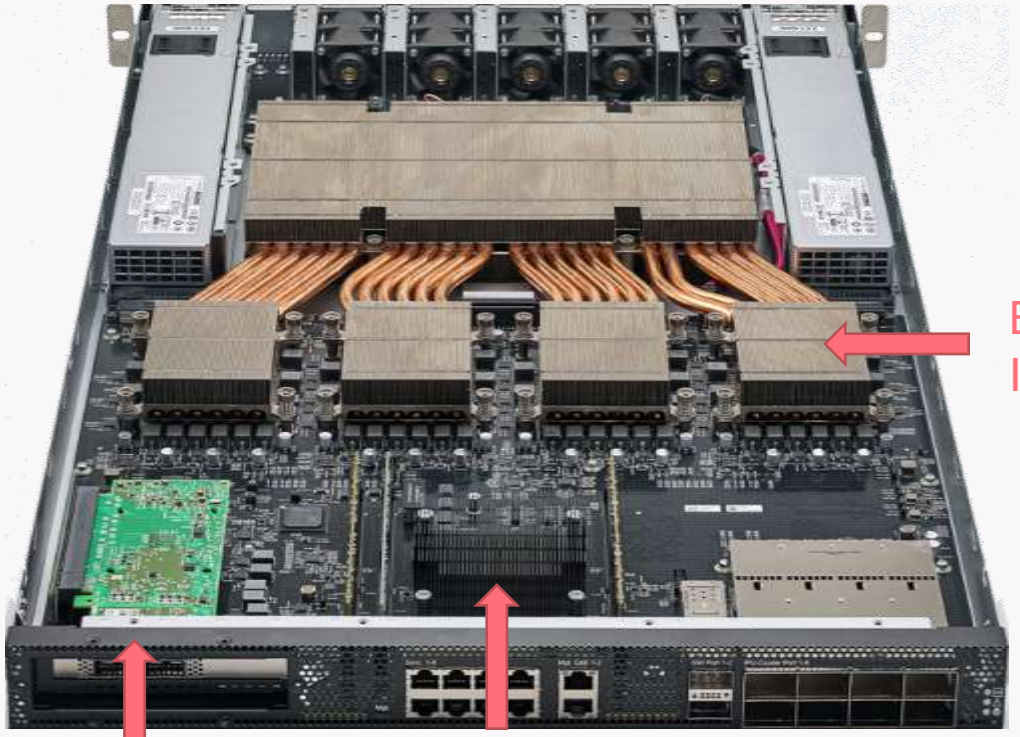

BOW: 3RD GENERATION IPU SYSTEMS

SHIPPING TO CUSTOMERS TODAY

BOW POD$_{16}$

4x Bow-2000
5.6 PetaFLOPS
1 CPU server

BOW POD$_{64}$

16x Bow-2000
22.4 PetaFLOPS
1-4 CPU server(s)

BOW POD$_{256}$

64x Bow-2000
89.6 PetaFLOPS
4-16 CPU server(s)

BOW-2000: THE BUILDING BLOCK OF LARGE PODS

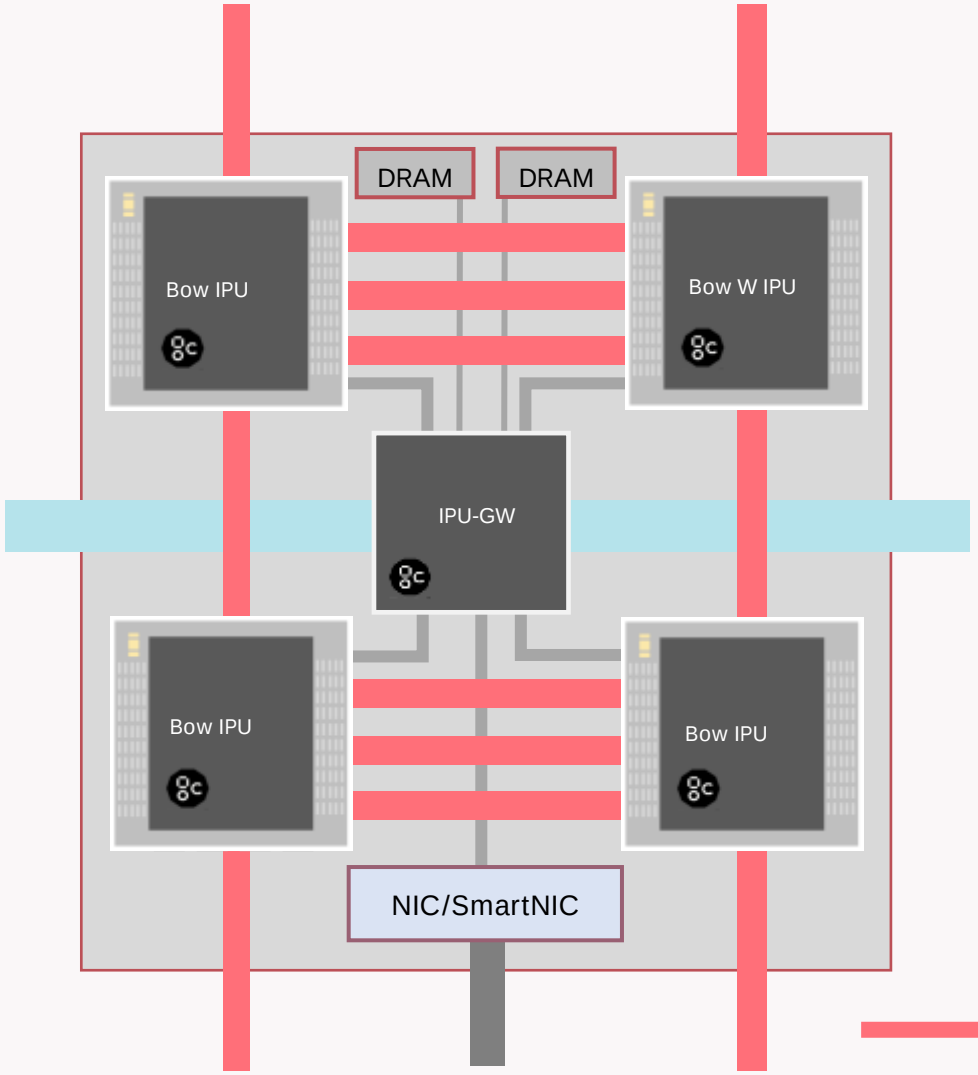

COMPUTE

4x Bow IPUs

- 1.4 $PFLOP_{16}$ compute
- 5,888 processor cores
- > 35,000 independent parallel threads

DATA

Exchange Memory

- 3.6GB In-Processor-Memory @ 260 TB/s
- 128GB Streaming Memory DRAM (up to 256GB) @ 20 GB/s

COMMUNICATIONS

IPU-Fabric managed by IPU-GW

- Host-Link – 100GE to Poplar Server for standard data center networking
- IPU-Link – 2D Torus for intra-POD64 communication
- GW-Link - 2x 100Gbps Gateway-Links for rack-to-rack – flexible topology

x16 IPU-Link [64GB/s]

Host-Link Network I/F [100Gbps]

IPU-GW Link [100Gbps]

x8 PCIe G4 [32GB/s]

BOW POD16 DIRECT ATTACH

Server with x4 Bow-2000

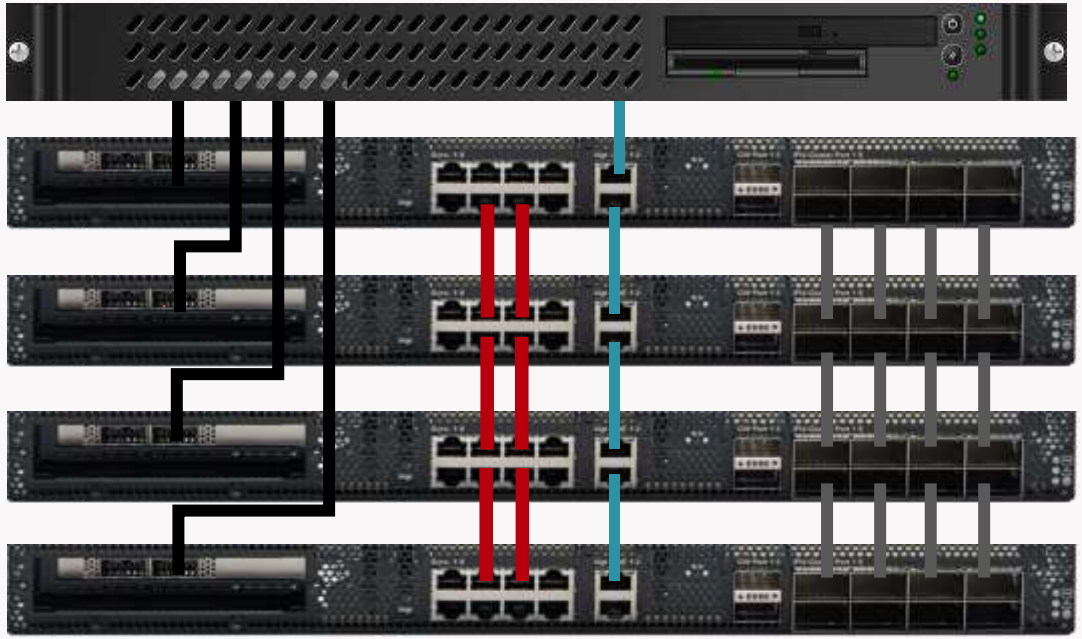

- Convenient cost effective evaluation platform
- Available through Graphcore channel for on-premise or Graphcloud
- Wide range of benchmarks and examples for Bow Pod_{16} performance evaluation
- Scale-out with Bow Pod_{64} and beyond

Host-Link 100GE network interface (QSFP, 1.0m)

1GbE Management (Cat5, 1.5m)

Sync-Link (Cat5, 0.15m)

IPU-Link (OSFP, 0.3m)

BOW POD64 REFERENCE DESIGN

Pre-Qualified 64-IPU Design with Reference Server and Switches

- Up to 16 Bow-2000 platforms
- Reference architecture supports different server requirements based on workload
- Bow Pod_{64} Configuration:
  - 64 IPUs
  - 22.4 PFLOPs @ FP16.16
  - ~58GB IPU In-Processor memory
  - ~7TB Streaming Memory
- Bow Pod Host disaggregation
  - Flexibly connect required host server compute over fabric
- 2D-Torus topology
  - Maximizes bandwidth across IPU-Links
  - All-Reduce 2x faster than mesh topology
- Scalable to 64K Bow IPUs

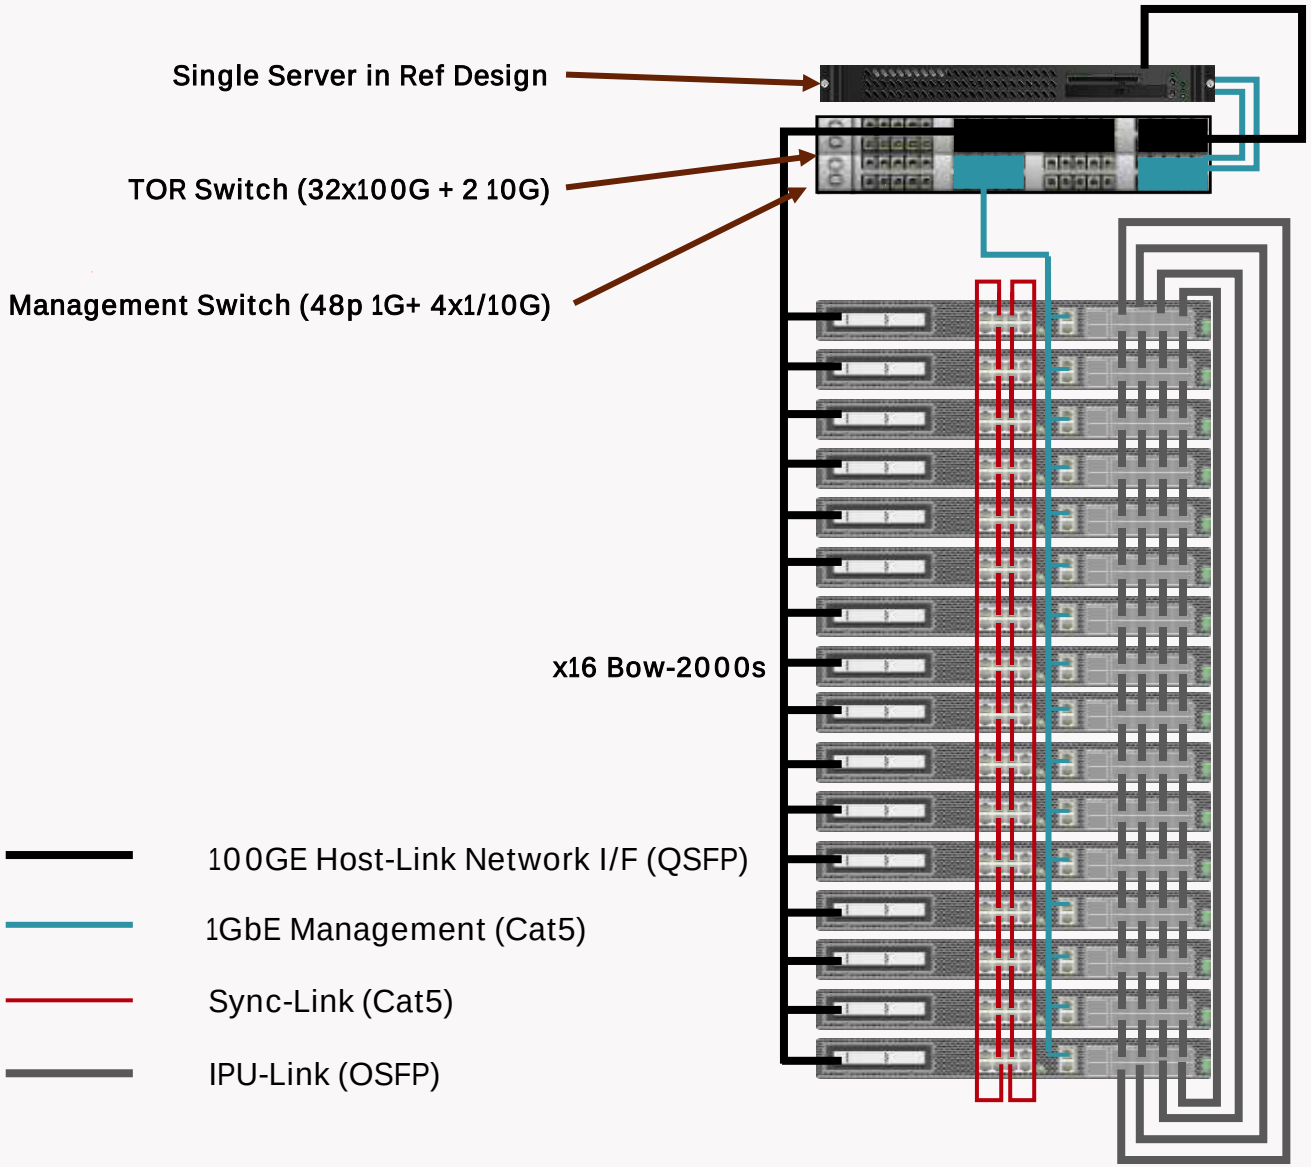

MODELS AND SOFTWARE

GRAPHCORE SOFTWARE ECOSYSTEM

WORLD CLASS DEVELOPER RESOURCES FOR IPU USERS

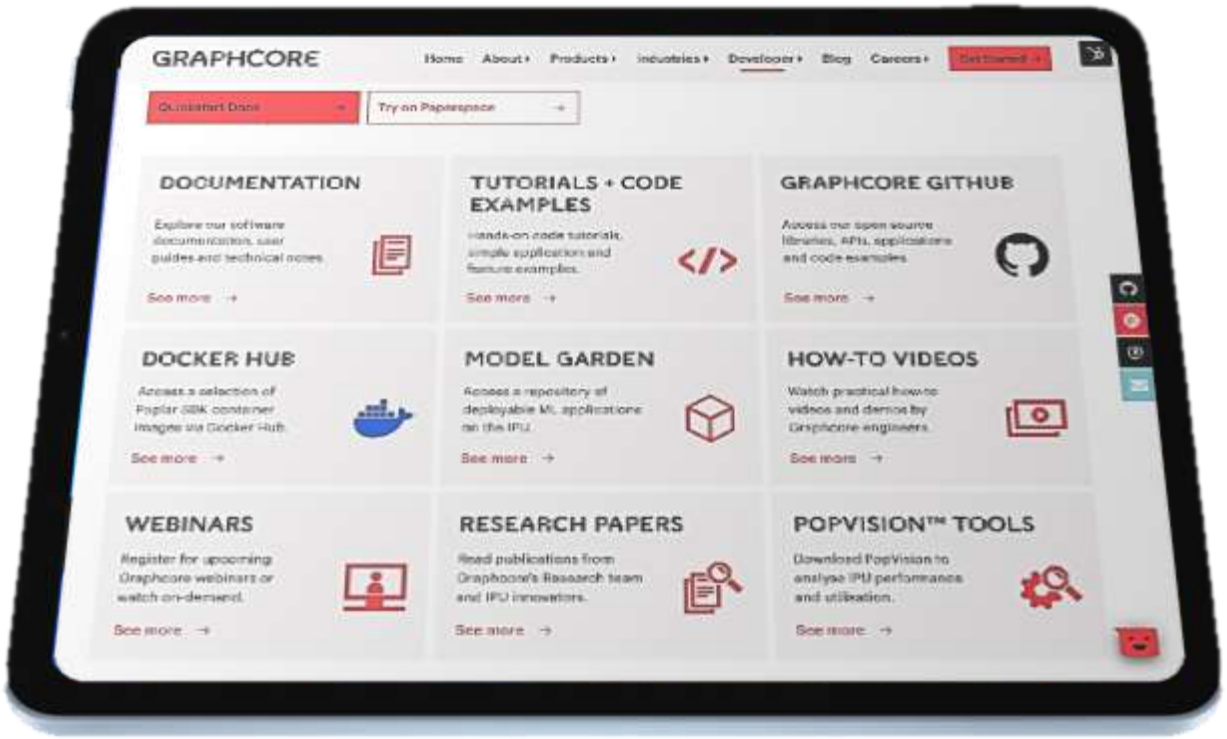

WWW.GRAPHCORE.AI/DEVELOPER

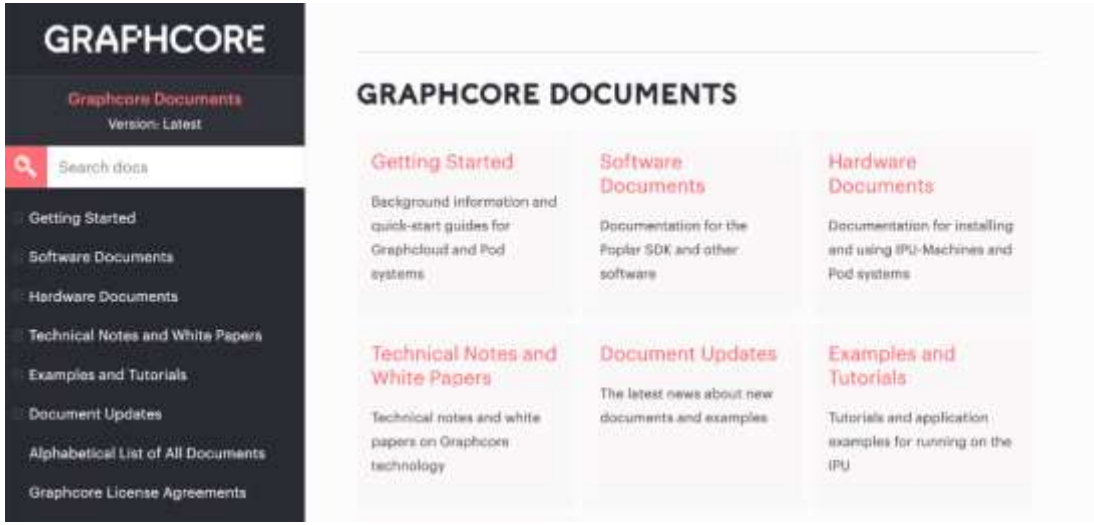

GRAPHCORE SOFTWARE MATURITY

NLP/TRANSFORMERS
IMAGE CLASSIFICATION/CNNS
OBJECT DETECTION
LARGE MODELS
MLPERF
CONDITIONAL SPARSITY
GNNS

ML APPLICATIONS

TUTORIALS
CODE EXAMPLES
DOCUMENTATION
VIDEOS
NATIVE IPU CODERS PROGRAM
APPS PORTFOLIO

DEVELOPER ECOSYSTEM

FRONTENDS

JUPYTER NOTEBOOKS

INFERENCE DEPLOYMENT TOOLKIT

FRAMEWORKS

ONNX
HALO
Keras
PaddlePaddle

FW BACKENDS

XLA
POPART+
POPDIST

POPLAR®

PARTITIONER
POPIR
POPIT
POPLIBS
GCL
POPLAR
GRAPH ENGINE
GRAPH COMPILER

DRIVERS

GC DEVICE ACCESS LAYER
IPUOF DRIVER
PCIe DRIVER

POPLAR® SDK

GRAPH ANALYZER
SYSTEM ANALYZER
DEBUGGER
DEVELOPMENT ENVIRONMENT

POPVISION TOOLS

V-IPU

SYSTEM MONITORING

PROMETHEUS
GRAFANA

JOB DEPLOYMENT

K8S
SLURM

SYSTEM SOFTWARE

GRAPHCORE

POPLIBS GRAPH LIBRARIES

Complete set of open-source libraries of Machine Learning primitives and building blocks

- Over 50 optimized functions for common machine learning models
- More than 750 high performance compute elements
- Simple C++ graph building API
- Implement any application
- Full control flow support

POPNN	FUNCTIONS USED IN NEURAL NETWORKS (NON-LINEARITIES, POOLING, LOSS FUNCTIONS)
POPLIN	OPTIMIZED LINEAR ALGEBRA FUNCTIONS (MATRIX MULTIPLICATION, CONVOLUTIONS)
POPOPS	FUNCTIONS FOR PERFORMING ELEMENTWISE OPERATIONS ON TENSOR DATA
POPRAND	HIGH PERFORMANCE FUNCTIONS FOR POPULATING TENSORS WITH RANDOM NUMBERS
POPUTIL	GENERAL UTILITY FUNCTIONS FOR BUILDING GRAPHS FOR IPU DEVICES
GCL	OPTIMIZED COLLECTIVES LIBRARY SUPPORTING MODEL AND DATA PARALLEL

GRAPH COMPILER

State of the art compiler for simplified IPU programming

- Handling the scheduling and work partitioning of large parallel programs including memory control:
- Optimized execution of the entire application model to run efficiently on IPU platforms
- Alleviates the burden on developers to manage data or model parallelism
- Code generation using standard LLVM

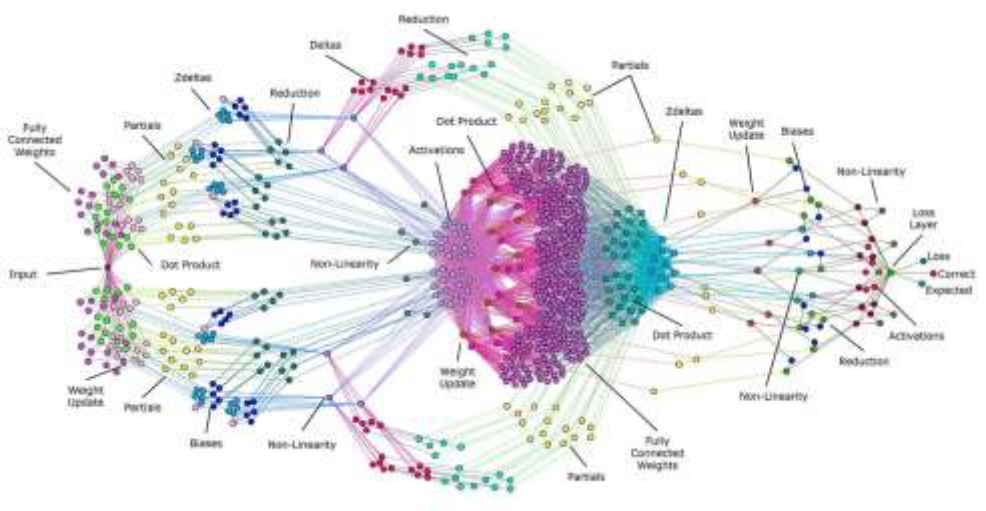

POPVISION™ TOOLS

GRAPH ANALYSER

Useful for analysing and optimising the memory use and execution performance of ML models on the IPU

SYSTEM ANALYSER

Graphical view of the timeline of host-side application execution steps

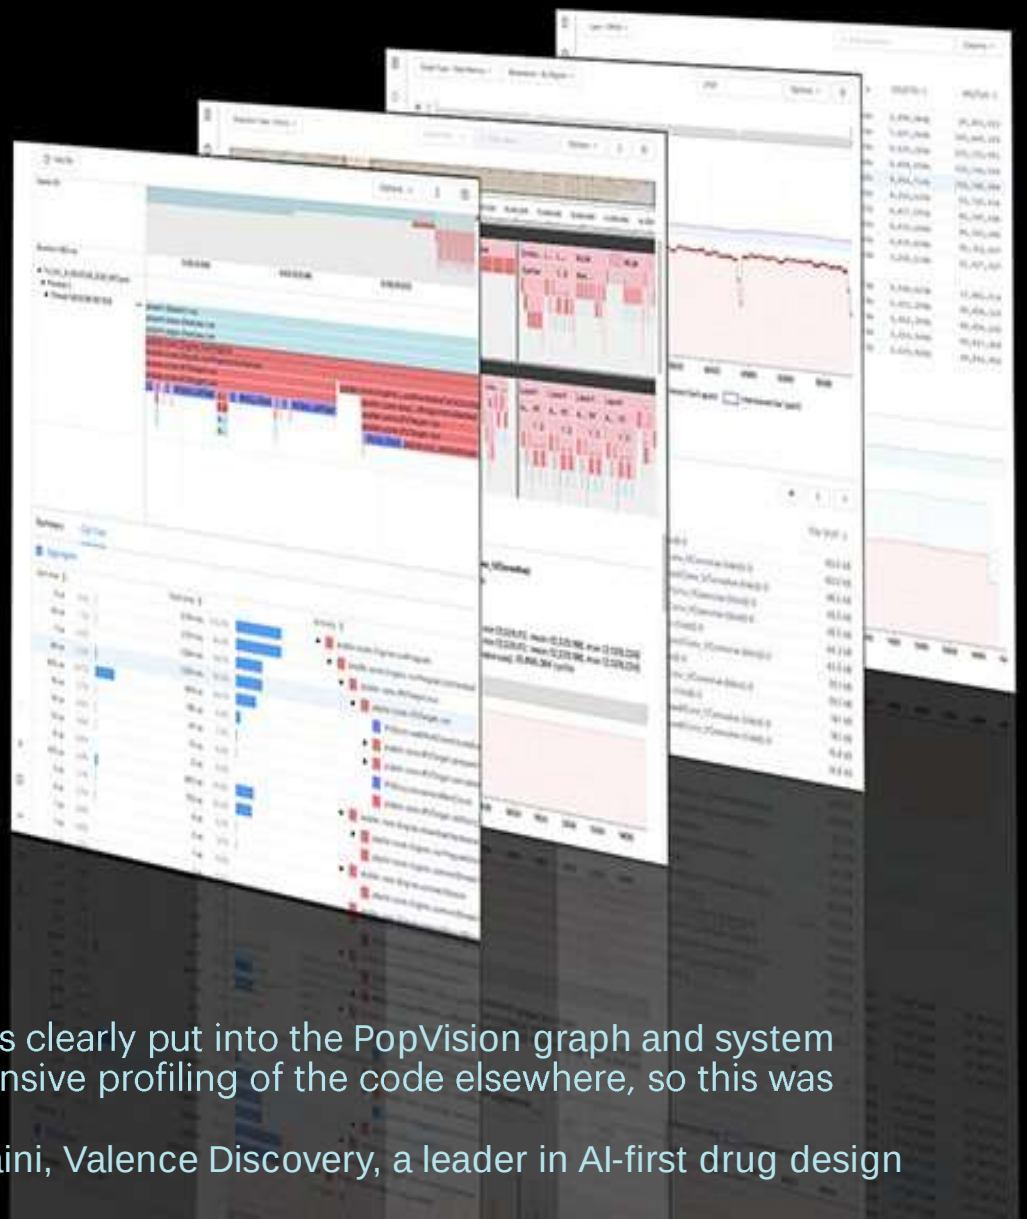

"Our team was very impressed by the care and effort Graphcore has clearly put into the PopVision graph and system analysers. It's hard to imagine getting such a helpful and comprehensive profiling of the code elsewhere, so this was really a standout feature in our IPU experience."

Dominique Beaini, Valence Discovery, a leader in AI-first drug design

STANDARD ML FRAMEWORK SUPPORT

Develop models using standard high-level frameworks or port existing models

PyTorch

TensorFlow

PyTorch Lightning

HUGGING FACE

PyG

Keras

PaddlePaddle

HALO

Existing models on alternative platforms

Easy port of high-level framework models

POPLAR®

IPU-Processor Platforms

PyG

PYTORCH GEOMETRIC FOR IPU

PyG is the ultimate library for Graph Neural Networks

Build graph learning pipelines with ease

kumo GRAPHCORE HARVARD MEDICAL SCHOOL

pyg.org

"The suitability of IPUs for running GNNs and the kind of performance advantage that Graphcore and its customers have demonstrated is really helping to accelerate the uptake of this exciting model class"

Matthias Fey – PyG creator & founder of Kumo.ai

ANNOUNCEMENT | TECHNICAL BLOG | GETTING STARTED

Apr 05, 2023 \ Popler, PyTorch, GNN

GRAPHCORE USERS CAN NOW BUILD AND RUN GNNS WITH PYTORCH GEOMETRIC

Written By:
Blazej Banaszewski

Apr 05, 2023 \ Developer, PyTorch, GNN

ACCELERATING PYG ON IPUS: UNLEASH THE POWER OF GRAPH NEURAL NETWORKS

Written By:
Blazej Banaszewski, Adam S
Akash Swamy & Mihai Polce

Apr 05, 2023 \ Developer, GNN, Paperspace

GETTING STARTED WITH PYTORCH GEOMETRIC (PYG) ON GRAPHCORE IPUS

Written By:
Adam Sanders and Arianna Saracino

RUN GNN MODELS IN PYG ON PAPERSPACE JUPYTER NOTEBOOKS

Training Dynamic Graphs
TGN
Training
Run on Gradient

Training Large Graphs
Cluster-GCN
Training
Run on Gradient

Predicting molecular properties
SchNet
Run on Gradient

Predicting molecular properties
GIN
Run on Gradient

Link Prediction training
NBFNet
Training
Run on Gradient

ENHANCED MODEL GARDEN

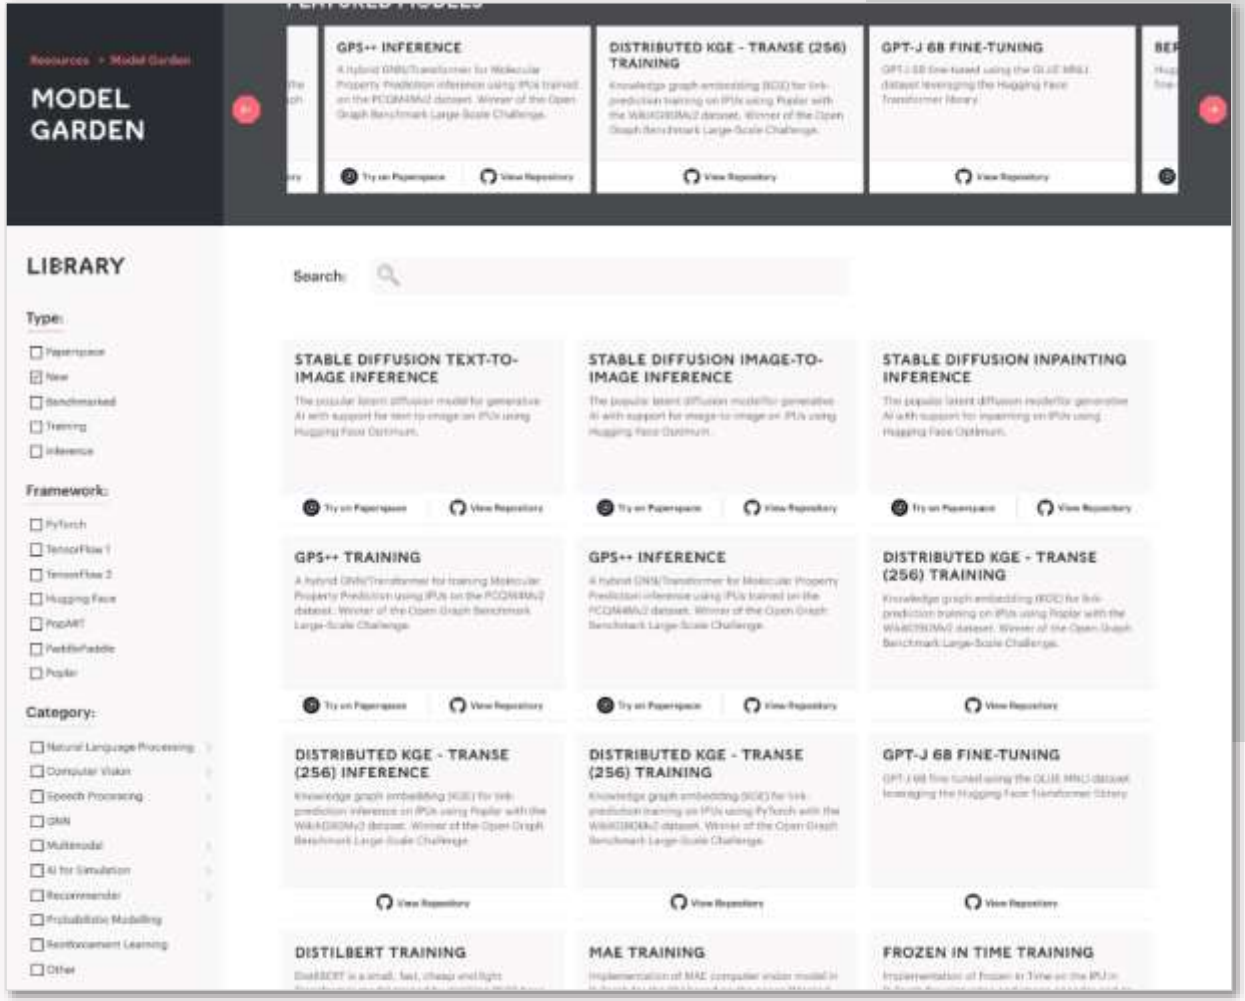

PUBLIC ACCESS TO WIDE VARIETY OF MODELS, READY TO RUN ON IPU

NEW FILTER/SEARCH CAPABILITY

DIRECT ACCESS TO GITHUB

PAPERSPACE NOTEBOOK LINKS

https://www.graphcore.ai/resources/model-garden

MODEL GARDEN COVERAGE

COMPUTER VISION

IMAGE CLASSIFICATION
- ResNet50 v1.5
- EfficientNet-B0
- EfficientNet-B4
- ResNeXt-101
- MobileNet v2
- MobileNet v3
- ViT
- DINO
- SWIN NEW
- MAE

OBJECT DETECTION
- YOLO v3
- YOLO v4
- Faster RCNN
- EfficientDet

OBJECT SEGMENTATION
- Unet (Industrial)
- Unet (Medical)

GNN
- TGN NEW
- MPNN-GIN NEW
- GPS++ NEW
- Distr. KGE NEW
- Cluster-GCN NEW
- SchNet NEW
- NBFNet NEW

AI FOR SIMULATION
- DeepMD
- DeepDriveMD
- ET0
- CosmoFlow
- ABC Covid-19

REINFORCEMENT
- RL
  Reinforcement Learning

PROBABILISTIC
- MCMC
- VAE

NLP
- BERT-Base
- BERT-Large
- Group | Packing
- GPT2
- GPT-J NEW
- RoBERTa
- Deberta
- BART
- T5
- Hubert
- DistilBERT NEW

SPEECH

STT (ASR)
- RNN-T
- Conformer

TTS
- DeepVoice3
- FastSpeech2
- FastPitch

RECOMMENDER
- Autoencoder
- DIN
- DIEN

MULTIMODAL
- LXMERT
- CLIP
- Stable Diffusion NEW
- Mini DALL-E
- Frozen In Time NEW

OTHER
- Sales Forecast
- Neural Image Fields

PyG

PyTorch

TensorFlow

Hugging Face

Keras

PaddlePaddle

POPART