HTCondor in Foxconn Hyperscale AI HPC Cloud

Cheng-Han Du, Ph.D.
Foxconn Corporation

Outline

- **Foxconn AI HPC Cloud Solution**
  - Introduction
  - AI/HPC Cloud Stack – Hardware & Software
  - Applications
- **HTCondor in Foxconn Cloud**
  - Research Case – Bayesian Optimization of Deep Learning Hyperparameters

Foxconn HPC Solution

- Collaborate with Microsoft: “TECHNOLOGY OF THE YEAR” award of InfoWorld in 2018
- Solution in multiple scales
  - HGX-1, HGX-2, and workstations: standalone usage.
  - M100 Rack: All-in-one HPC box for on-premise solution
  - M500 Cloud: Ultimate scalable HPC system for public cloud
    - Serving computations of wide variety of scales
    - Massively scalable performance (HPC)
    - Huge number of smaller-scale and independent tasks for many customers (***HTC***)

HGX-1 / HGX-2 / WorkStation → M100 Rack → M500 Cluster

All-in-one HPC in a Rack

- For on-premise datacenter
- Integrated 100TB+ Storage close to high performance GPU and CPU with 100G high speed fabric
- Ready-to-use private cloud management software
- Integrated machine learning platform software

Fii Cloud & HPC

A
Smartphone
Wearable
device

B
Tablet /NB

C
Mechanical/
Component/
Automation

D
Industrial4.0 +
Smart Office+PC

E
Smart TV,
TV Wall,
Game Console
Consumer Robot

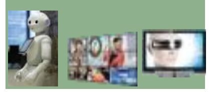

Fii

H
E-commence
/ O2O

富连网
货真 价实 良品
水活 服务

I
Investment

J
Administrative
Support &
Service

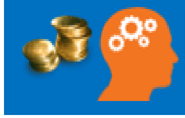

K
Small, Medium
& Large Size
Display

L
Key Component
Industry Alliance

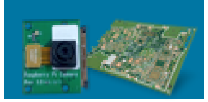

M
Health, Life &
Prevention

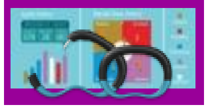

S
Semiconductor
IC design

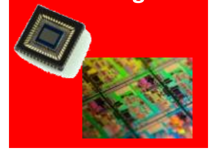

TOP 500

Big Twin

Home / News / Foxconn Builds Taiwan's Largest Supercomputer

Foxconn Builds Taiwan's Largest Supercomputer

Michael Feldman | January 4, 2018 17:02 CET

E-mail Tweet Like G+1 Share

A six-petaflop system has been installed at Taiwan's Kaohsiung Software Technology Park, making it the most powerful supercomputer in the country.

The machine was constructed by Hon Hai Precision Industry Company, aka Foxconn Technology Group, the electronics contract manufacturer that assembles the iPhone and iPad for Apple. Based in New Taipei, Foxconn is Taiwan's largest company, with a revenue of $131.8 billion (2016). The multinational enterprise operates across Asia, Europe, South America, and North American, and as of 2015, employed over 1.3 million workers. Besides Apple, Foxconn customers have included Microsoft, Intel, Hewlett-Packard, Dell, Sony, Amazon, and Google, among others.

As an original design manufacturer (ODM), Foxconn has constructed various servers before, but this appears to be the first time the company has built an entire system. No information was provided on the makeup of the new supercomputer in

TAIPEI TIMES

AI智能物聯網平台開發人才養成班 - 待業者
AI+IoT時代來臨！如何在世界競爭下殺出重圍，提升關鍵職能

Home | Front Page | Taiwan News | Business | Editorials | Sp

Home / Business

Wed, Dec 27, 2017 - Page 12 News List

Print Mail Facebook Twitter plurk funp

Hon Hai launches its first HPC with focus on health

ORDINARY PEOPLE: AI medical diagnosis technology in the US can cost up to US$8,000, a prohibitive cost that Hon Hai hopes to reduce for general public access

By Lisa Wang / Staff reporter

Objectives

Foxconn AI HPC Solution

Democratizing HPC

- Reduce usage difficulty of HPC resource
- Users focus more on application, less on IT issues

Streamlined for Deployment

- Utilize big data IT, cloud platform
- ML/DL/HPC as a service

Multidisciplinary Collaboration

- Close collaboration with customers and domain experts
- Cloud deployment consulting for customer
- Codesign of multiple applications: numerical simulation and AI with HPC

Foxconn AI HPC Stack

Applications

1. High-Resolution (8K) Image/Video Object Detection
2. Industrial Application / Surface Defect Detection
3. Medical Application / Deep learning / Industrial AI / NGS

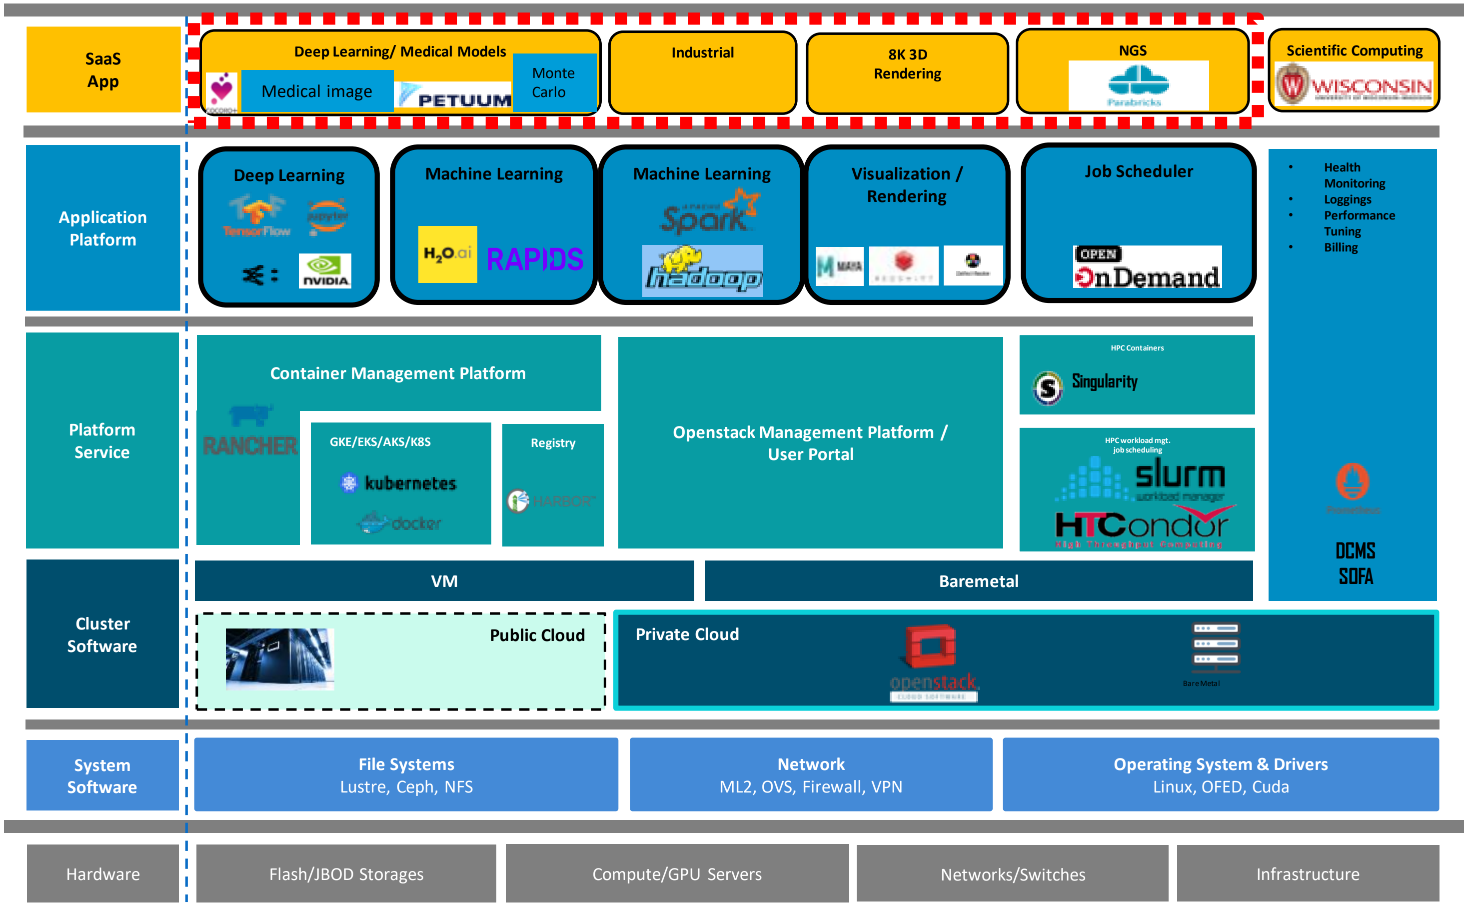

Foxconn AI HPC Stack

Machine Learning / Deep Learning Service

1. Multi-node multi-GPU cluster environment
2. Training and inference workflow
3. Deep learning inference application platform
4. Performance autotuning (FAUST, SOFA, POTATO)

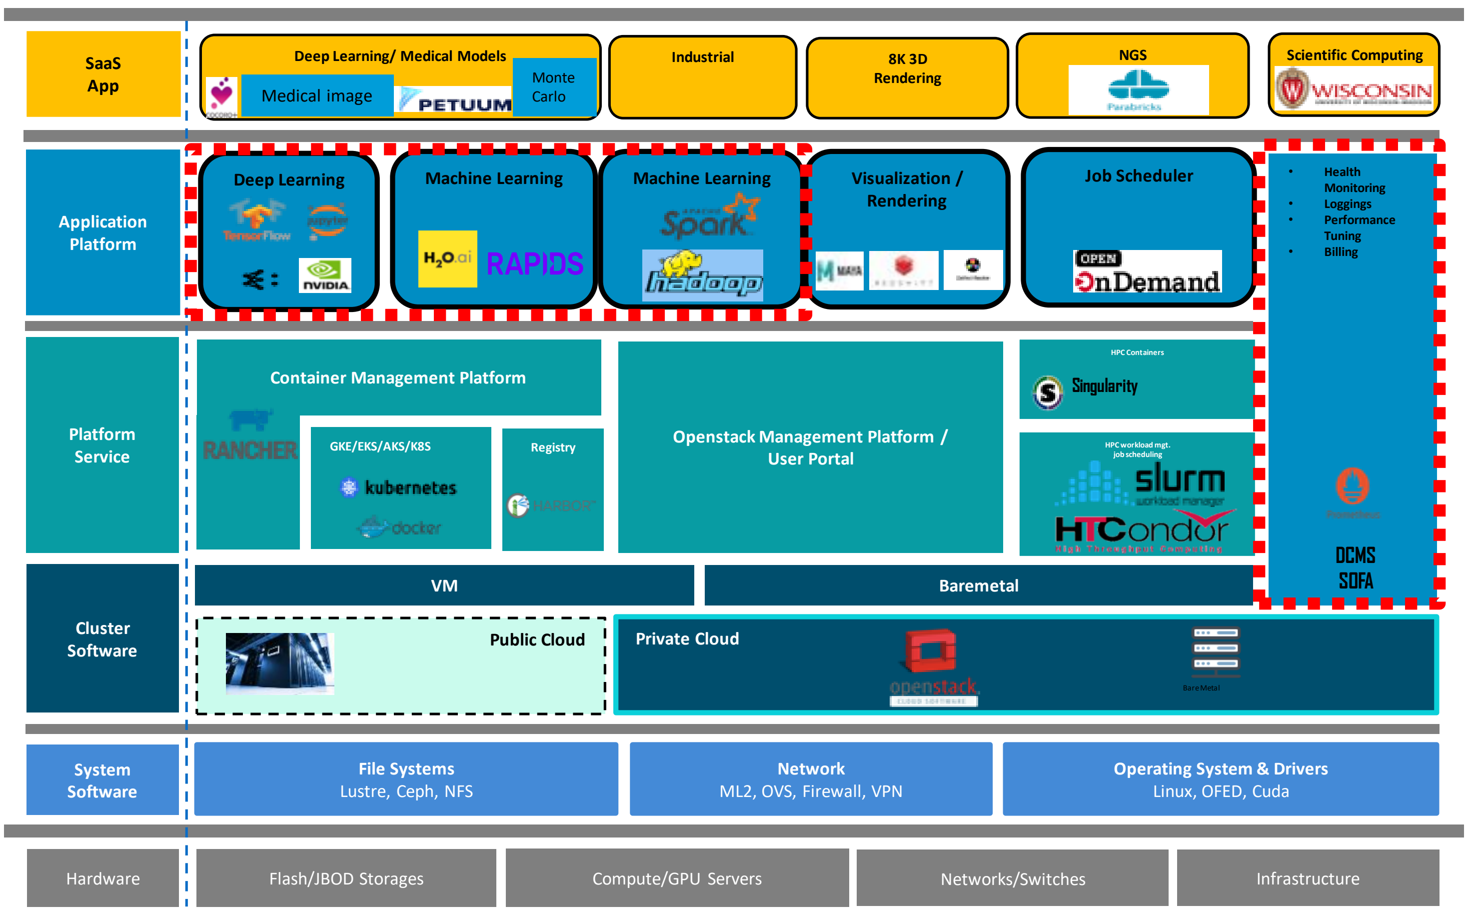

Foxconn AI HPC Stack

HPC Service

1. Optimized job scheduler and virtualization environment according to specified performance-energy criteria
2. Job dispatch in hybrid HPC environment, including on-premise HPC clusters and public cloud data centers

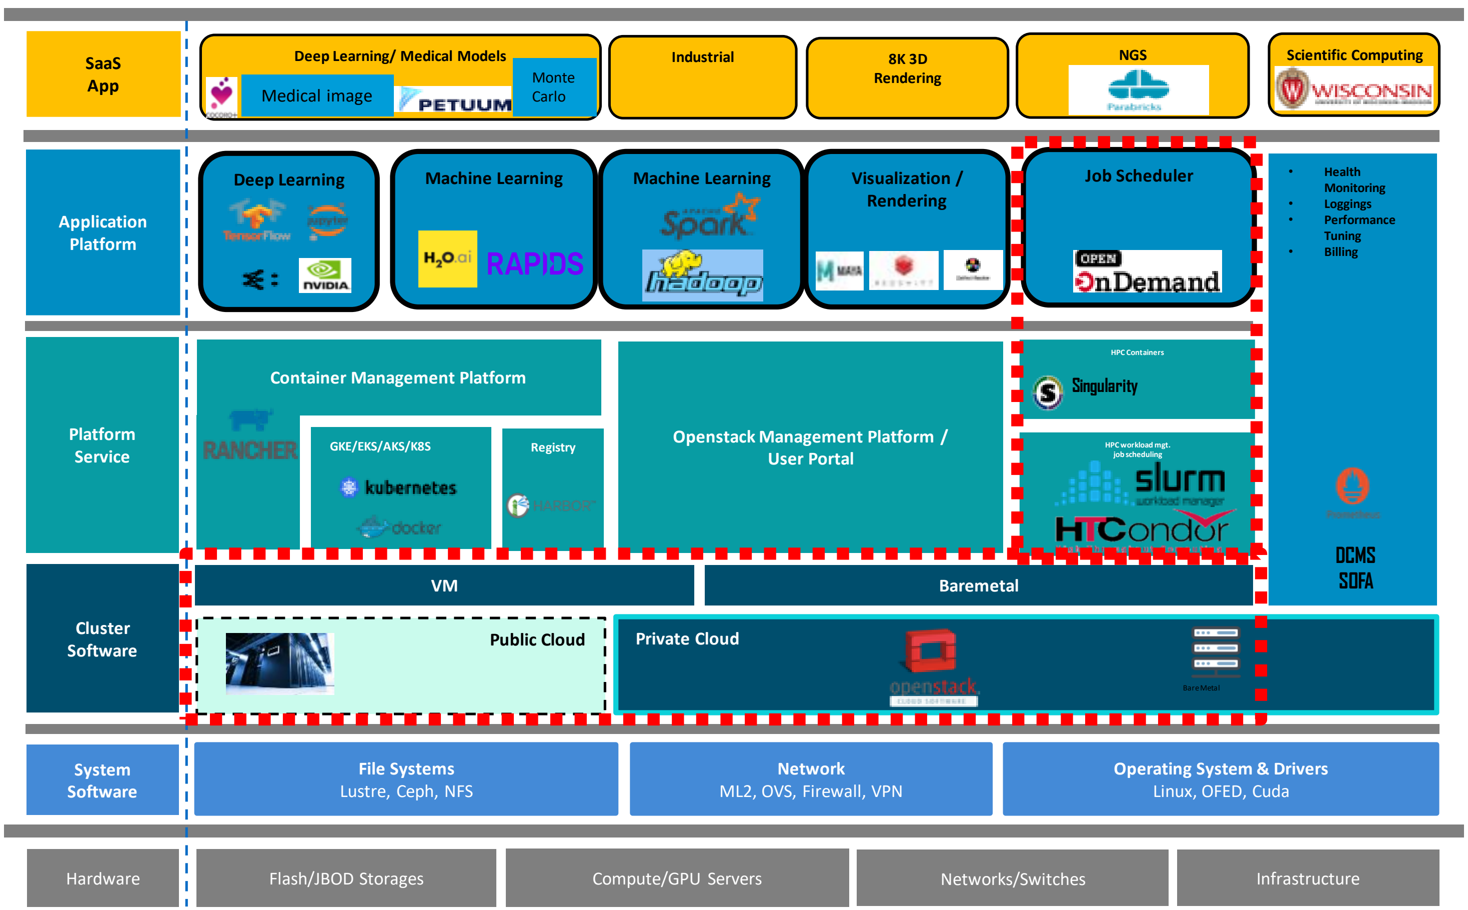

Hybrid Cluster Management (In Progress)

- Coexist HTCondor, Slurm, and K8S cluster under an OpenStack tenant
- Allow Slurm, K8S, and HTCondor batch jobs to share the same cluster hardware resources

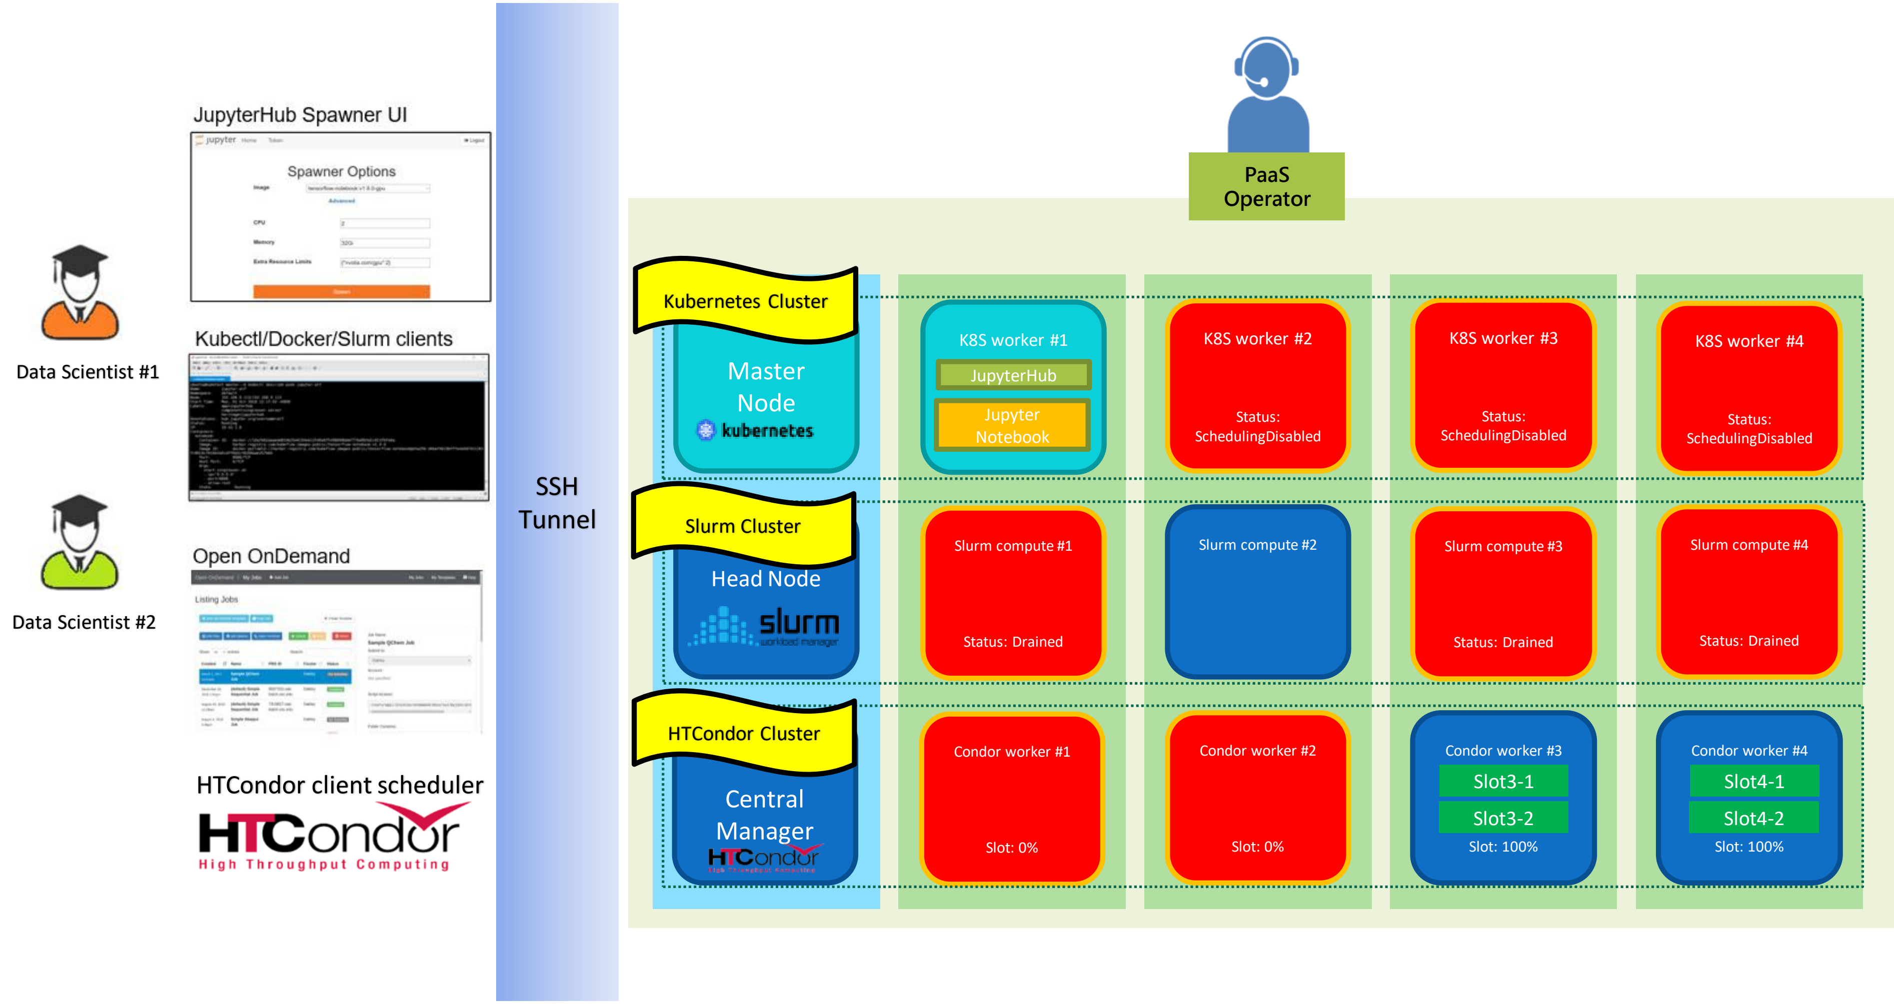

HGX-1 Architecture

Modular GPU computing solution

- Reconfigurable **scale-up** system
- High-density GPU clusters for space saving
- CPU node
  - 32 cores in total

- Single-node with 8 or 16 GPUs

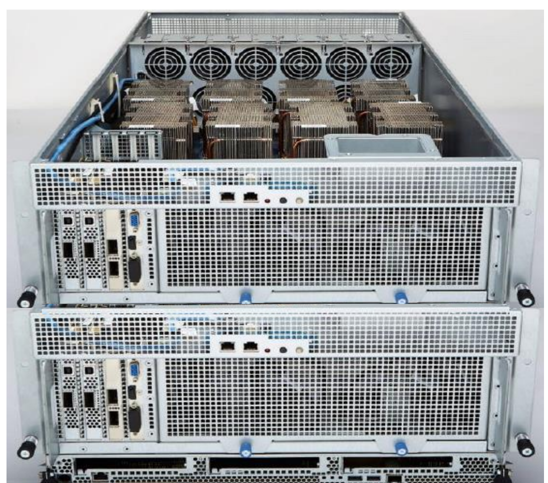

HGX-1 Stingray-HPC64, Local M.2 SSD

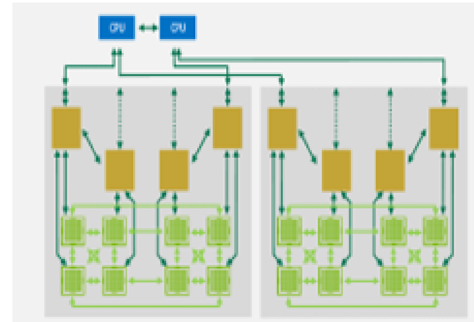

CPU: Intel(R) Xeon(R) Gold 6154 CPU @ 3.00GHz
Memory: DDR4 2666MHz 32G x 24
GPU: Tesla V100-SXM2-32GB
Driver: 410.72
OS: NGC nvcr.io/nvidia/tensorflow:18.09-py3
Cuda: cuda 10
Cudnn: 7.3.0.29-1
Tensorflow: 1.10
Benchmark: 1.9

8 GPU

ResNet50 Real, Hvd	BS=64	BS=128	BS=256	BS=512
fp32	2395	2760	2987	X
fp16	3959	5129	6029	6464

16 GPU

ResNet50 Real, Hvd	BS=64	BS=128	BS=256	BS=512
fp32	4448	5306	5847	X
fp16	7140	9459	11244	12362

Next Generation Sequencing

Parabricks NGS

- NGS pipeline speedup: 30 ~ 50X
- Sequencer-generated FASTQ files to variant call files (VCFs) with GPUs

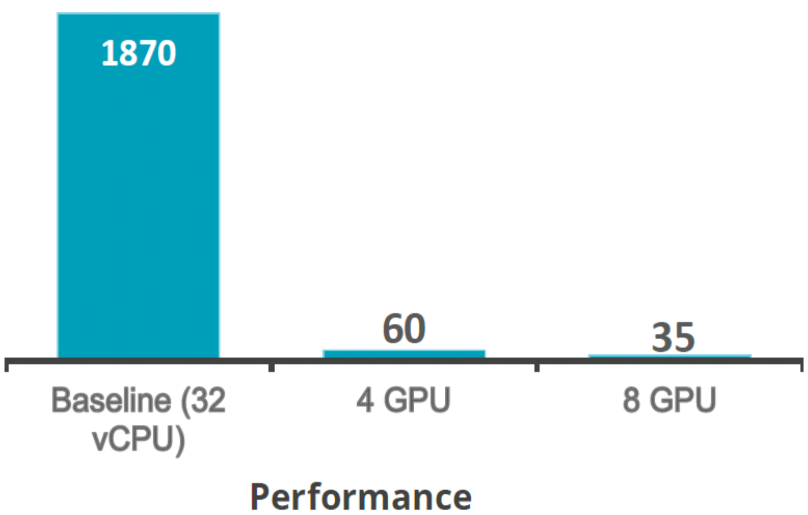

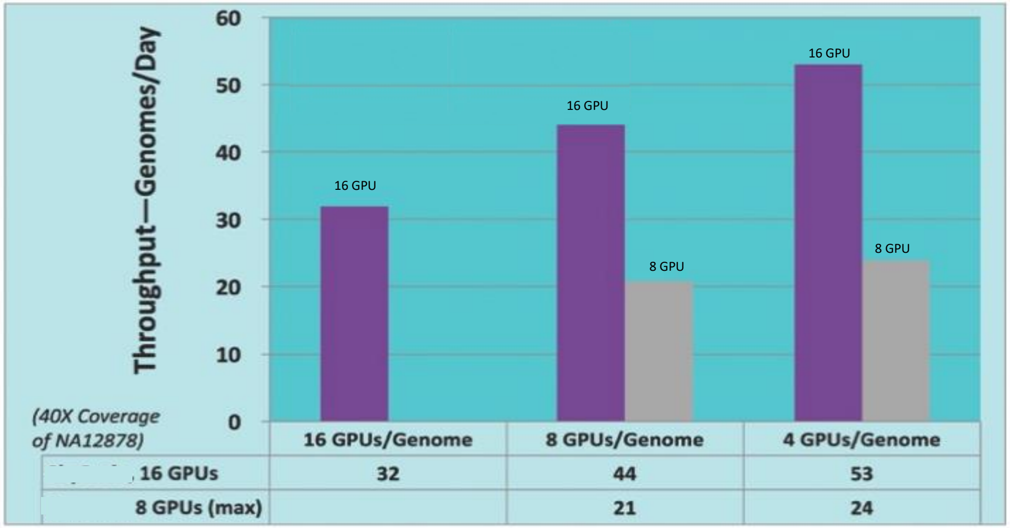

Medical Applications

Mammogram abnormality detection

Real-time instant identification

- Single image upload identification
  - Detection of abnormal position and score

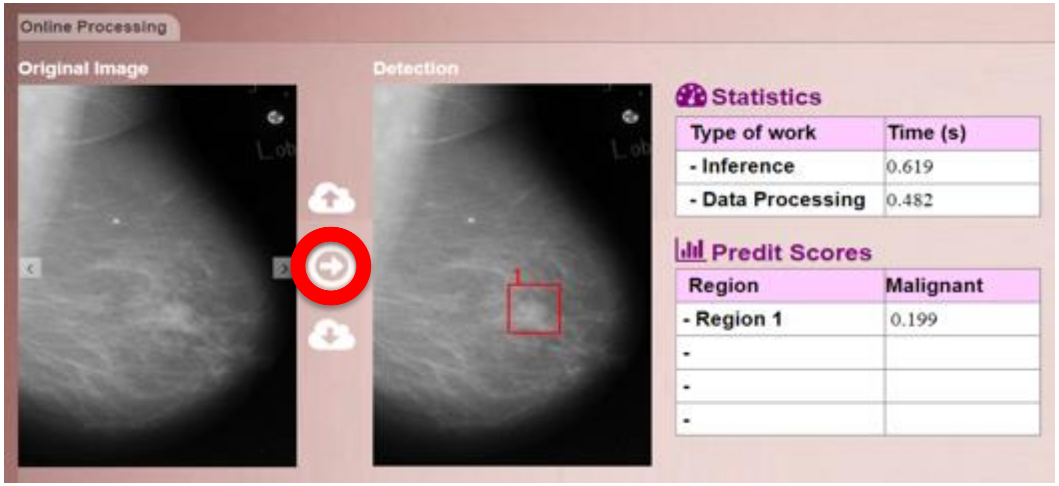

Batch identification processing

- Multiple database image recognition
  - Batch detection and anomaly location
- Batch process with 8 GPUs
  - 600,000+ images processed per day

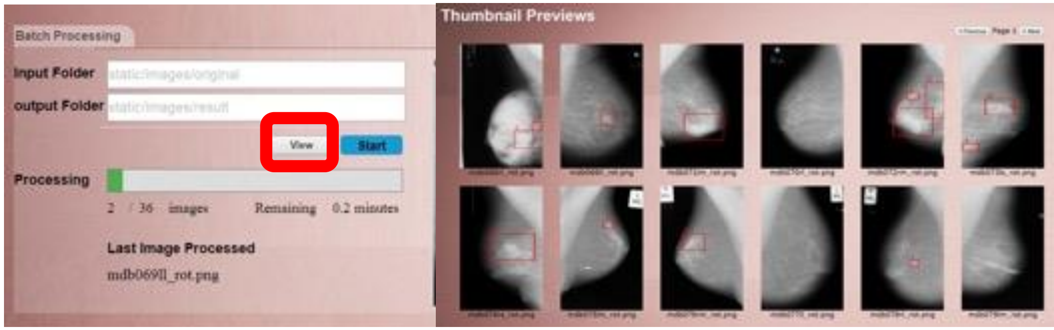

Reference:

1. Ribli *et al*., “Detecting and classifying lesions in mammograms with deep learning,” *Scientific Reports* vol. 8, 4165 (2018)
2. Digital Mammography DREAM Challenge

HGX-2 Architecture

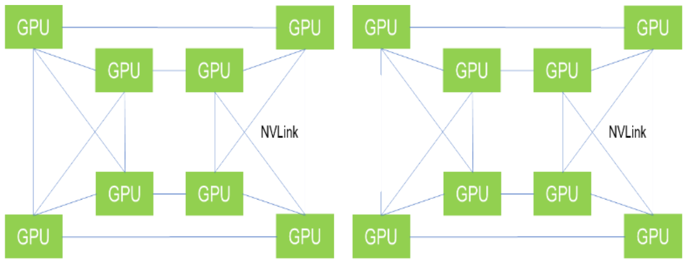

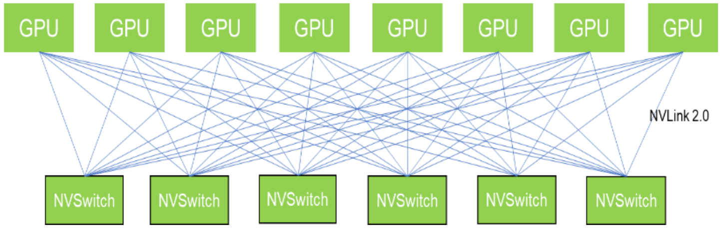

	HGX-1	*HGX-2*
GPU # in single node	8 (single box) or 16 (cascaded boxes)	16
Inter-GPU communication	NVLink ‘Cube Mesh’ + PCIe	NVLink Switch
Inter-GPU performance	1 or 2 hops in 8 GPUs of an expansion box, plus 2 hops across PCIe of different boxes	2 hops between 16 GPUs via non-blocking NVSwitch
Ideal applications	Applications require moderate GPU data exchange, e.g. most DL training models, ResNet, AlexNet, etc.	Applications require heavy GPU data exchange, e.g. Scientific simulation, etc.
TCO per GPU	Low	Higher cost due to extra NVSwitch cost and power consumption
Availability	Now	Q3, 2019

Ref: https://www.nextplatform.com/2018/04/13/building-bigger-faster-gpu-clusters-using-nvswitches/

GPU Infrastructure Consideration

On-premise GPU clusters

- Improved data security and application latency
- Limited computing resource
- High maintenance effort

Off-premise GPU data center

- Virtually unlimited computing resource for bursts of critical demands
- Low maintenance cost

HTCondor

- High throughput for massive number of tasks
- Configurable hardware demand for each application
- **Docker universe**
  - Swift deployment of complex application container
- On-demand GPU computing
  - HTCondor Annex
  - VM instance management in cloud
- Federation of clusters
- DAGMan
  - Easy orchestration of tasks

Public Cloud

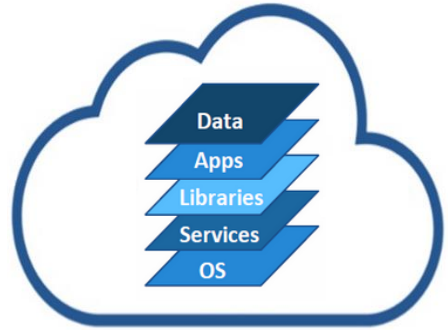

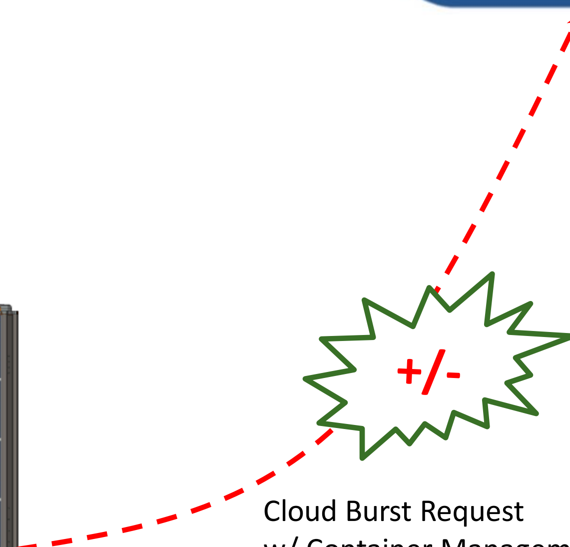

Cloud Burst Request
w/ Container Management Platform

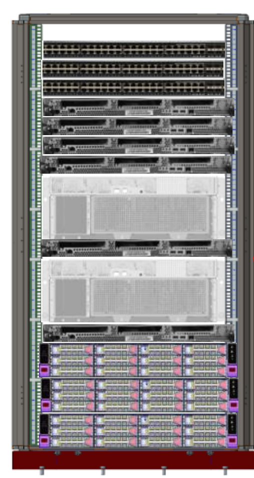

On-premise HPC Rack

Outline

- **Foxconn AI HPC Cloud Solution**
  - Company Introduction
  - AI/HPC Cloud Stack – Hardware & Software
  - Applications
- **HTCondor in Foxconn Cloud**
  - Research Case – Bayesian Optimization of Deep Learning Hyperparameters

Research Case

Task dispatch for hyperparameter optimization

Application

- Optimization of AI model or many scientific problems
- Optimization scheme
  - Grid search (parallel)
  - Random search (parallel)
  - Bayesian optimization (serial in general)

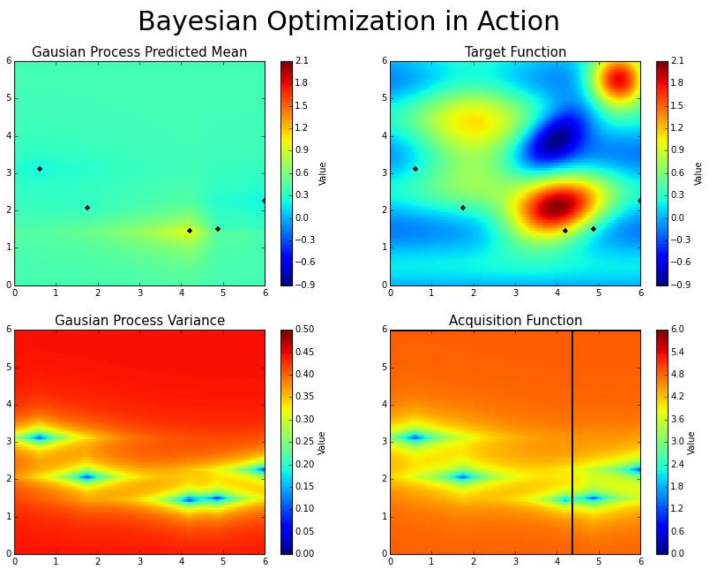

HTCondor Features

- GPU computing
- Docker universe
  - Default runtime set to “nvidia”
- Partitionable slots
- HTCondor-Python

Major Tools

- Scikit-optimize (Bayesian optimization)
- Image: nvcr.io/nvidia/tensorflow:18.12-py3 (Keras)

Reference: https://github.com/fmfn/BayesianOptimization/blob/master/examples/bayesian_optimization.gif

Research Case

Task dispatch for hyperparameter optimization

Environments and Setting

- Single-node with 8 Tesla P100 in Foxconn/DCT data center in Kaohsiung
- Mini-HTCondor
  - Single worker with partitionable resources
- Submitter
  - Python environment with HTCondor API and scikit-optimize
  - Eight observation tasks submitted in each round of optimization
  - One GPU and 4 CPU slots for each task

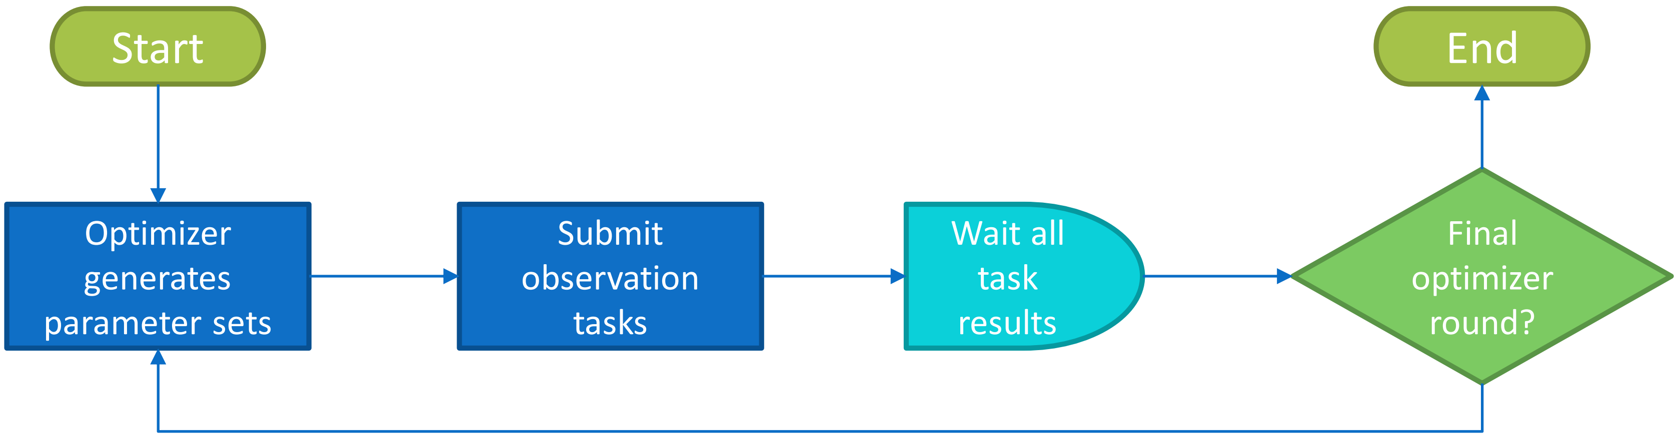

Research Case

Task dispatch for hyperparameter optimization

Test script

- Keras example: cifar10_cnn.py
  - 50 epochs
- 8 optimization rounds
- Generally underfitting
  - Baseline setting
    - Learning rate: 1.0e-4
    - Layer size: 512
    - Validation accuracy: 77.74%
  - Tuning range
    - Learning rate: [1.0e-6, 1.0]
    - Layer size: [128, 1024]

```
Score 0.779900 for parameter set  [1024, 8.049824704019605e-05]
Score 0.773200 for parameter set  [577, 8.505281624321753e-05]
Score 0.770400 for parameter set  [1018, 9.877286181054077e-05]
Score 0.767200 for parameter set  [1024, 6.368165079811596e-05]
Score 0.764600 for parameter set  [1022, 9.790266862705711e-05]
Score 0.763900 for parameter set  [1024, 6.749912931702168e-05]
Score 0.761900 for parameter set  [1024, 6.594575565833666e-05]
Score 0.757700 for parameter set  [636, 0.00012910071728580158]
Score 0.754500 for parameter set  [1024, 5.395416679452777e-05]
Score 0.753900 for parameter set  [1019, 9.833030469970878e-05]
Score 0.753800 for parameter set  [975, 0.00012207698981520957]
Score 0.751900 for parameter set  [1024, 6.789389450271147e-05]
Score 0.748900 for parameter set  [617, 6.775072733733407e-05]
Score 0.748400 for parameter set  [1024, 5.228268927494645e-05]
Score 0.747400 for parameter set  [465, 5.028506626893611e-05]
Score 0.746500 for parameter set  [968, 6.174024931710192e-05]
Score 0.746200 for parameter set  [511, 6.273130395796479e-05]
Score 0.746200 for parameter set  [1024, 5.30120813707728e-05]
Score 0.744500 for parameter set  [1024, 6.026359849627899e-05]
Score 0.744000 for parameter set  [1024, 5.151728265597643e-05]
Score 0.739500 for parameter set  [1018, 4.9762599359445486e-05]
Score 0.738800 for parameter set  [1024, 7.081208890468328e-05]
Score 0.734600 for parameter set  [405, 7.719624788188919e-05]
Score 0.734300 for parameter set  [128, 6.685671725193425e-05]
Score 0.733200 for parameter set  [213, 5.144708557740793e-05]
Score 0.732500 for parameter set  [1024, 5.264280995142357e-05]
```

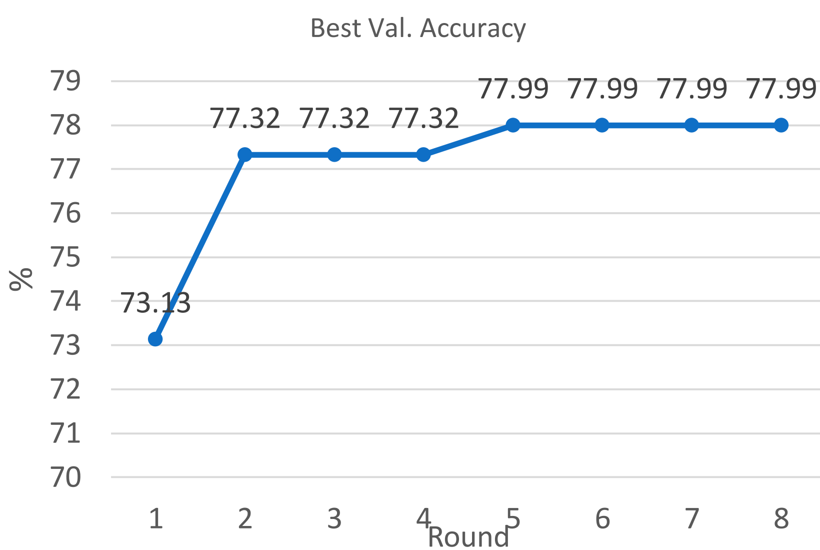

Summary

- Best accuracy in the 5th round
- Suggested learning rate: [0.00005, 0.00015]
- Suggested layer size: [512, 1024]

Extension

- Scale-out for more observations in one round
- Good results could be obtained with less rounds

Research Case

Task dispatch for hyperparameter optimization

Discussion & Further Investigation

- Network bandwidth
  - Training data upload for multiple tasks
  - ImageNet: over 100 GB of data
  - Potential solutions
    - Lustre share storage
    - Proxy
    - Docker volume mount
      - Prepare data before optimization tasks start
      - Mount volume of prepared data
      - Pro: data sharing of multiple tasks in the same node
      - Caveat: data nodes v.s. worker nodes specified by central manager

- Data Security
  - Primary concern in commercial applications
  - Potential solution
    - Encrypted filesystem

Thank You!