JUNE 11, 2018

Embrace Disruption

Pete Beckman
Senior Scientist, Argonne National Laboratory
Co-Director, Northwestern University / Argonne Institute for Science and Engineering

U.S. DEPARTMENT OF ENERGY
Argonne National Laboratory is a U.S. Department of Energy laboratory managed by UChicago Argonne, LLC.

Argonne NATIONAL LABORATORY

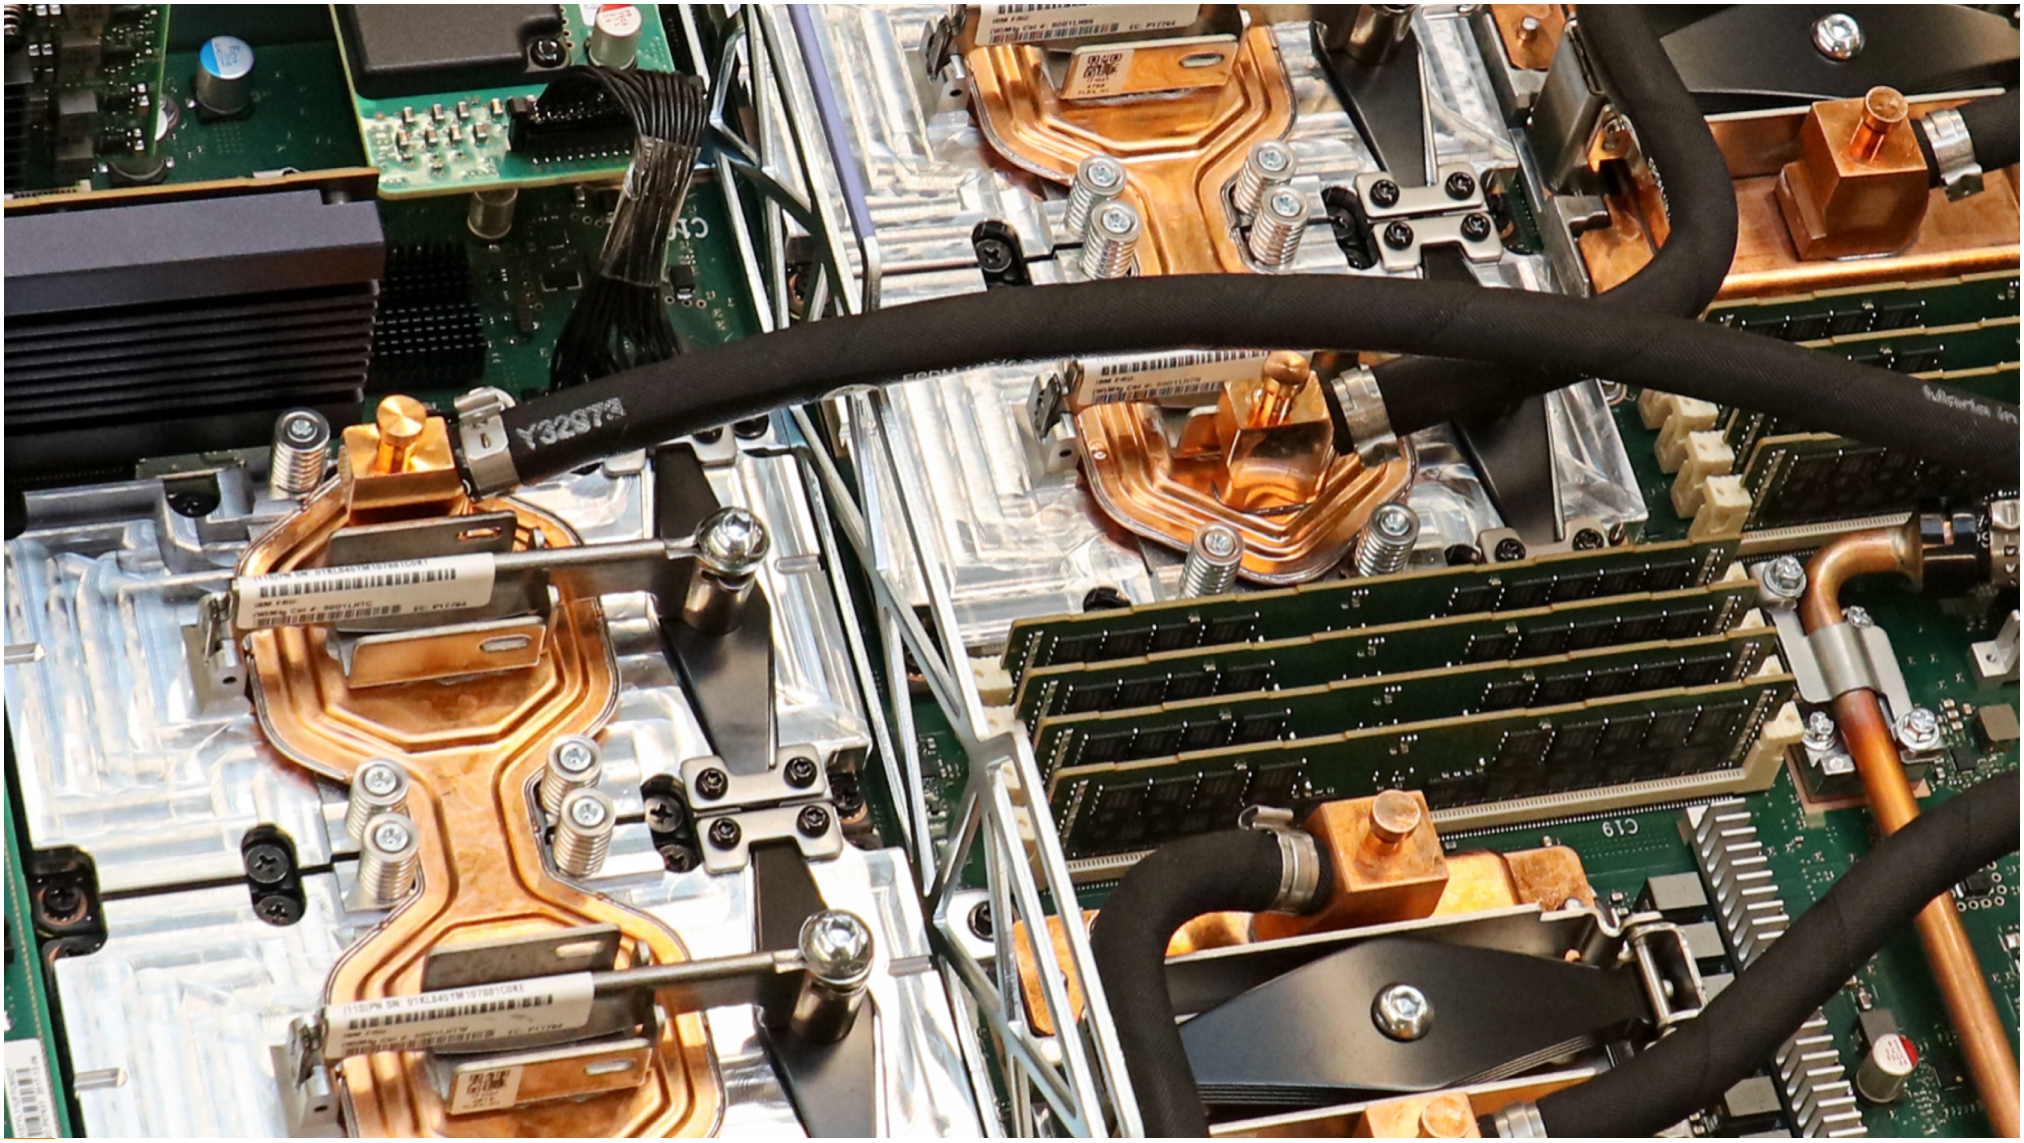

An $O(N)$ Sorting Algorithm: Machine Learning Sorting

Hanqing Zhao[1,2] and Yuehan Luo[3]

[1] *Department of Modern Physics, University of Science and Technology of China, Hefei 230026, China*
[2] *School of Physical Science and Technology, and Key Laboratory for Magnetism and Magnetic Materials of MOE, Lanzhou University, Lanzhou, Gansu 730000, China*[*]
[3] *Department of Applied Mathematics and Statistics, State University of New York at Stony Brook, Stony Brook, NY, 11794, USA*

We propose an $O(N)$ sorting algorithm based on Machine Learning method, which shows a huge potential for sorting big data. This sorting algorithm can be applied to parallel sorting and is suitable for GPU or TPU acceleration. Furthermore, we apply this algorithm to sparse hash table.

INTRODUCTION

Sorting, as a fundamental operation on data, has attracted intensively interest from the beginning of computing [1]. Lots of excellent algorithms have been designed, however, it's been proven that sorting algorithms based on comparison have a fundamental requirement of $\Omega(N \log N)$ comparisons, which means $O(N \log N)$ time complexity. Recent years, with the emergence of big data (even terabytes of data), efficiency becomes more important for data processing, and researchers have put many

algorithms
as human
cal image
accomplish
Neural n
the biology
layer and l
connecting
according
flect the re
is a mappi

cerebras

GRAPHCORE

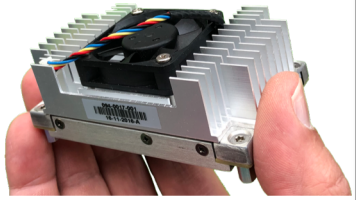

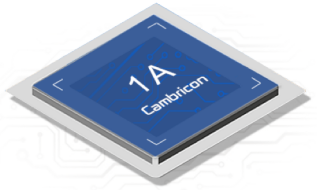

groq

MYTHIC

The Case for Learned Index Structures

Tim Kraska*
MIT
Cambridge, MA
kraska@mit.edu

Alex Beutel
Google, Inc.
Mountain View, CA
alexbeutel@google.com

Ed H. Chi
Google, Inc.
Mountain View, CA
edchi@google.com

Jeffrey Dean
Google, Inc.
Mountain View, CA
jeff@google.com

Neoklis Polyzotis
Google, Inc.
Mountain View, CA
npolyzotis@google.com

Abstract

Indexes are models: a B-Tree-Index can be seen as a model to map a key to the position of a record within a sorted array, a Hash-Index as a model to map a key to a position of a record within an unsorted array, and a BitMap-Index as a model to indicate if a data record exists or not. In this exploratory research paper, we start from this premise and posit that all existing index structures can be replaced with other

1) Why are we still tuning?

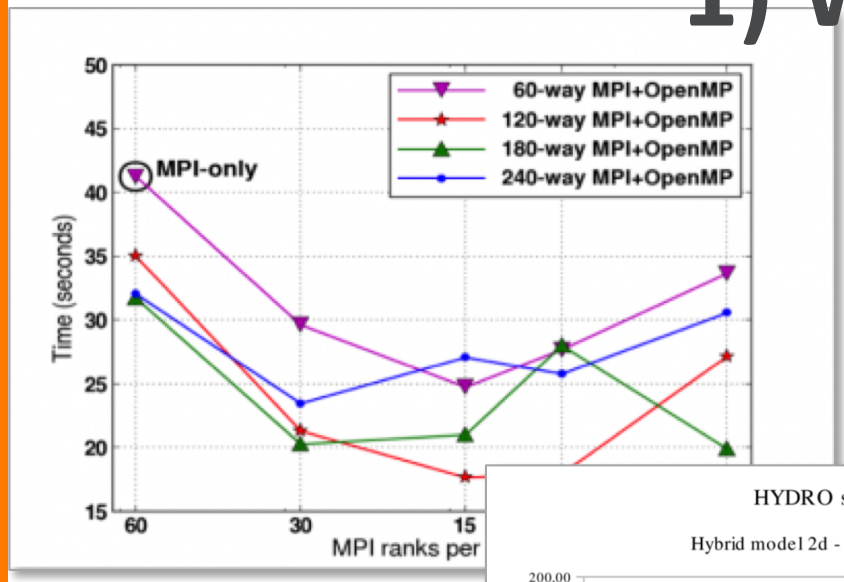

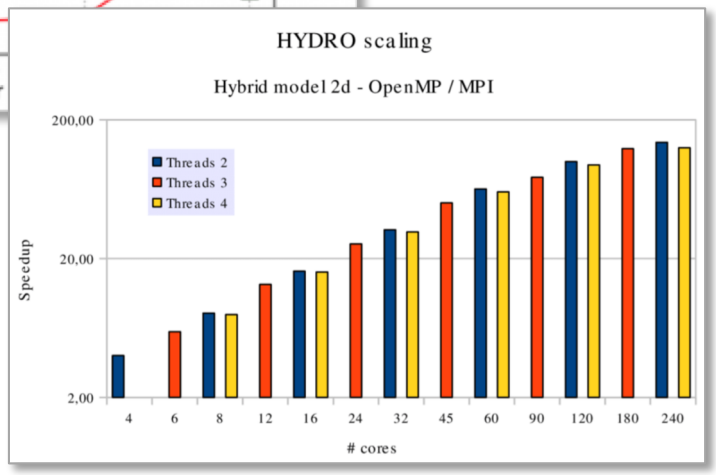

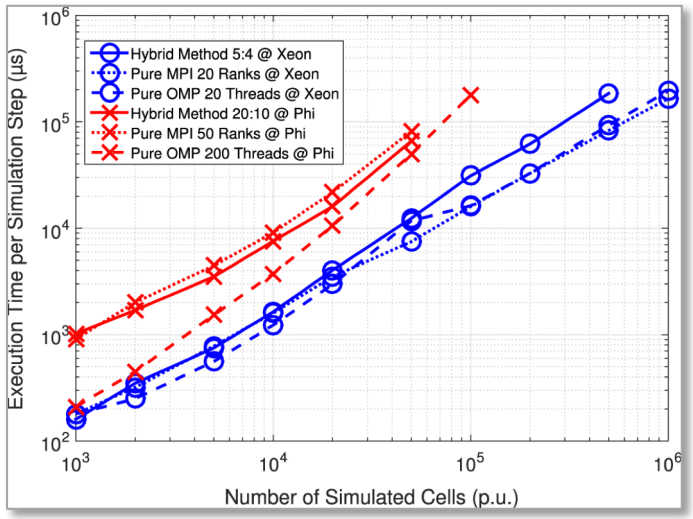

2) Why are we still driving?

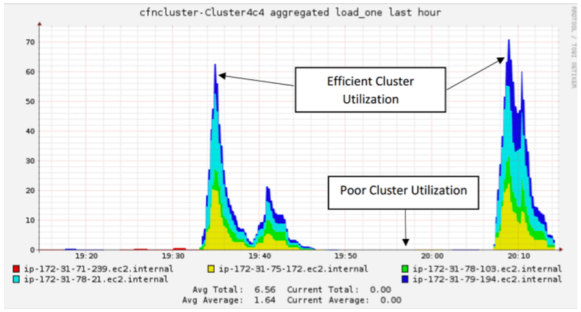

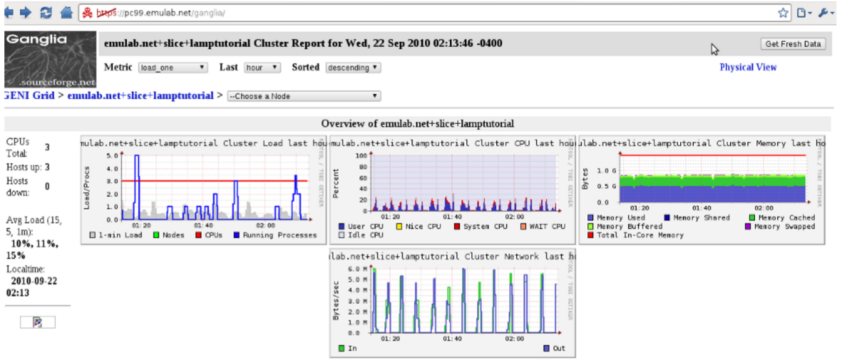

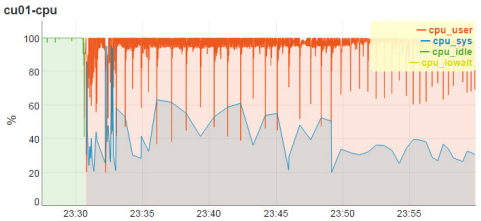

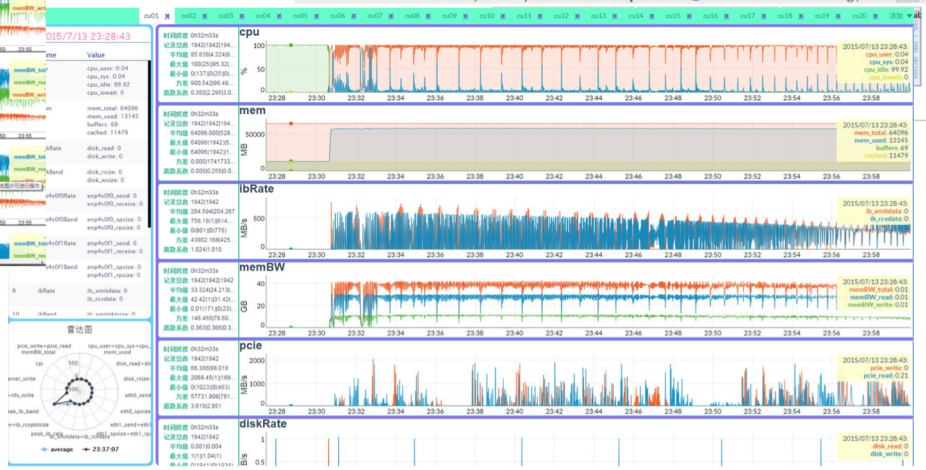

Argonne NATIONAL LABORATORY

3) Why are we still designing?

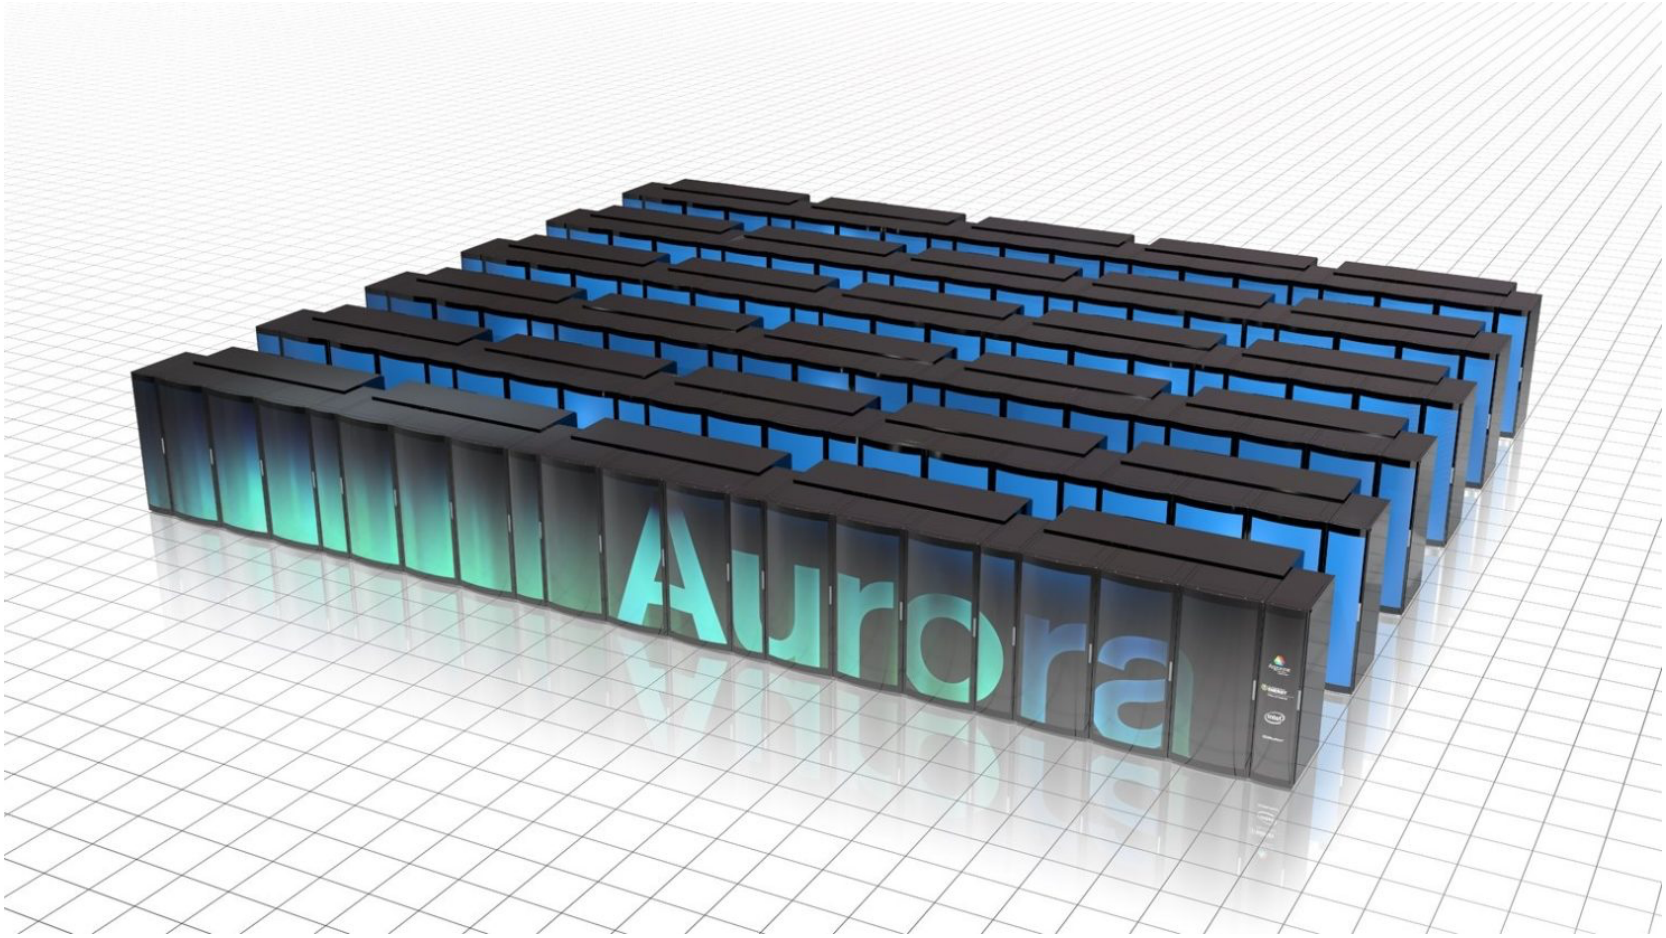

Argonne NATIONAL LABORATORY