Chapter 1: Introduction

- **About me**
- **Evaluation in this course**
- **Why do we have this course, and How do I organize?**
  - Modern Society needs more computation power
  - Organize the topics with an example
- **Resources related to this course**

Chapter 1: Introduction

- **About me**
- **Evaluation in this course**
- **Why do we have this course, and How do I organize?**
  - Modern Society needs more computation power
  - Organize the topics with an example
- **Resources related to this course**

You'd better to know HPComputing

- **Are you confident with your future?**
  - Programming is just the primary
  - Big Data + Data Analytics is HOT now
    - ➢ Math is absolute the key
      - ✓ LA + Optimization → ML/DM
  - HPC – Parallel + Distributed – is the potential skill to extend your career domain

2nd year — Intern

1st year — Many courses

Scoring!

- **Scores**
  - 100 pts
    - 80 pts from projects
      - ✓ 3 times to implement the concurrent programming projects with **MPI**, **CUDA** and **MR** (**Hybrid**)
      - ✓ 1 from your selected projects
        - » Weather forecasting, DL, GC (Rendering), CBIR, E-commerce,
        - » If distributed OS, **GREAT**!
    - 20 pts from the paper report
      - ✓ Your choice for interesting and challenging topics
        - » But focusing on “Design & Implementation” better with **code analysis** and **documenting**

≥ 95	A+
≥ 90	A
≥ 80	B

≥ 70	C
≥ 60	D
< 60	E

Reports

- **Design and Implementation**
  - ~~DNS – Domain Name Server~~
  - Globus Toolkit
  - Napster
  - Big Data
    - Hadoop, Map – Reduce
    - Zookeeper, Spark, MADlib, …
    - YARN, Mesos, Kubernetes …
    - HAWQ

  - Tencent's Peacock
    - Peacock: Learning Long-Tail Topic Features for Industrial Applications
  - Baidu's Padlepadle….

You can start now

1. **Team up [组队]**
   - No more than **5** students
2. **Learn the Numeric solution/Method for Heat Equation**
   - I'll also introduce later
   - **Sequential program** for **HE**/**NM**
3. **Choose your framework**
   - **MPI**, **CUDA/GPU**, **Cluster**/**MR**
   - Install and configure the programming environment
   - Learn how to program on those frameworks
4. **Finish the programming and other reports**

No Cheating!

Even Not Passing is better than Cheating!

Chapter 1: Introduction

- **About me**
- **Evaluation in this course**
- **Why do we have this course, and How do I organize?**
  - Modern Society needs more computation power
    - More Computation Power, more choices for jobs and advanced research
  - Organize the topics with an example
- **Resources related to this course**

Modern Society needs more computation power

- **To some extent, the evolution of our human depends on the extension of our brains**
  - To imagine, conclude, compute, construct etc.

- **Among those “skills”, computation is important, and here come COMPUTERS**

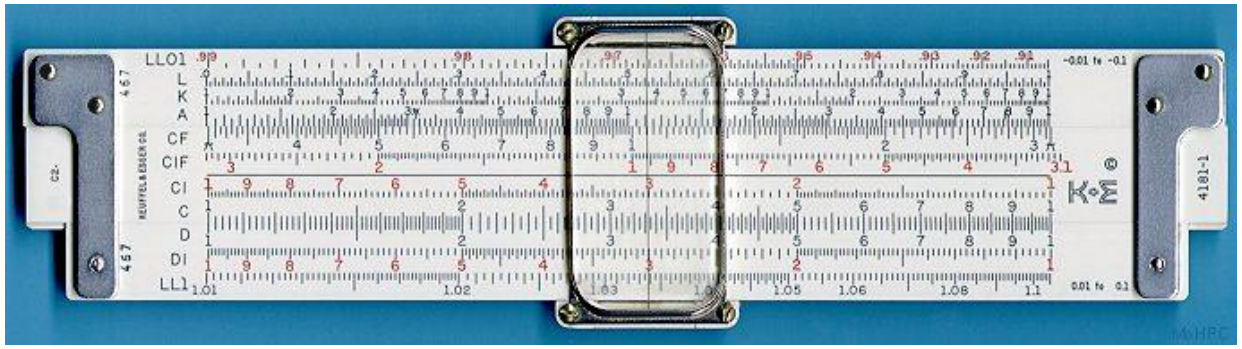

1833 Charles Babbage & Analytical Engine

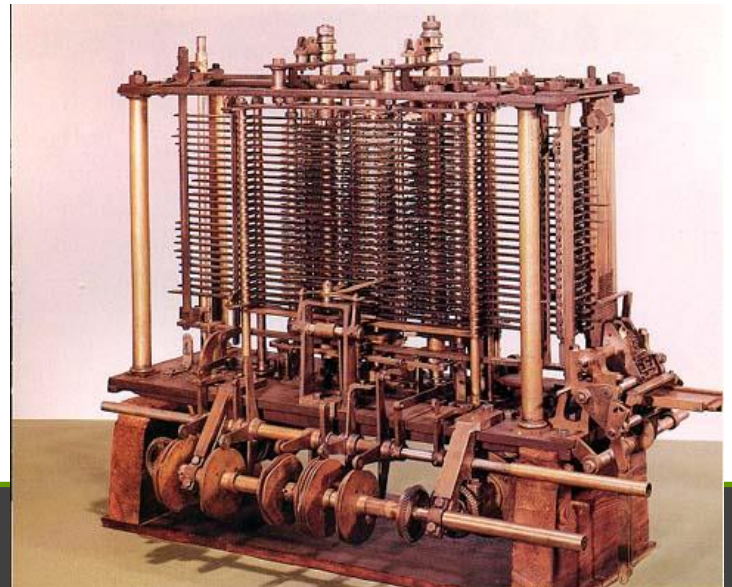

Atanasoff-Berry Computer (ABC) 1942

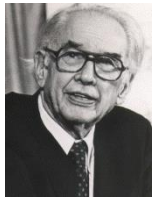

von Neumann architecture is the basis of modern computers

- With 4 generations

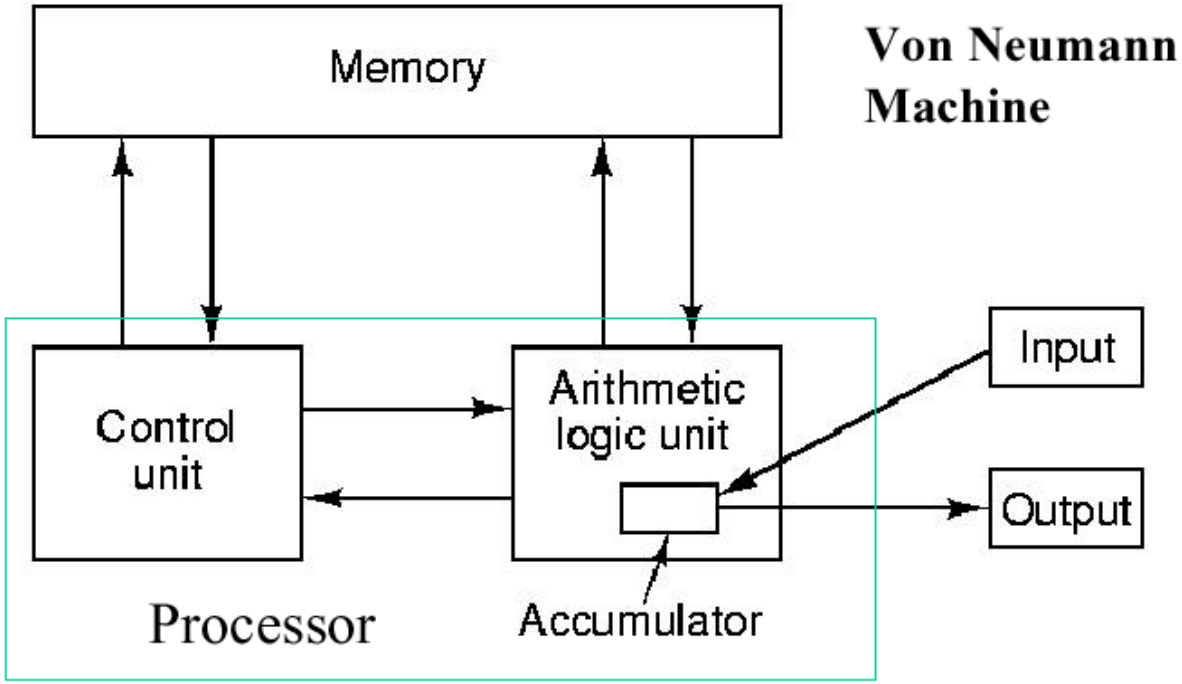

fig from mathdl.maa.org

Many applications needs more computation power

- **With HUGE data to be processed**
  - Scientific computing, business applications etc.

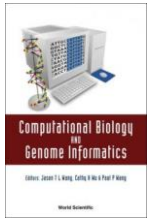

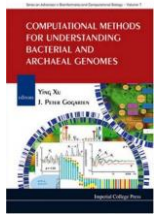

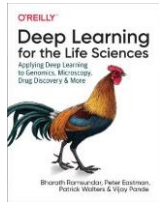

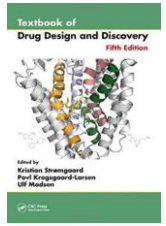

Government-Classified Work

Government - Research

OAK RIDGE NATIONAL LABORATORY

Instituto Nacional de Meteorología

(Severe) Weather Prediction & Climate Modeling

Automotive Design & Safety

Drug Discovery & Genomic Research

Aircraft/Spacecraft Design & Fuel-Efficiency

Oil Exploration & Energy Research

Basic Scientific Research

“Exascale” Applications at Berkeley Lab (LBNL)

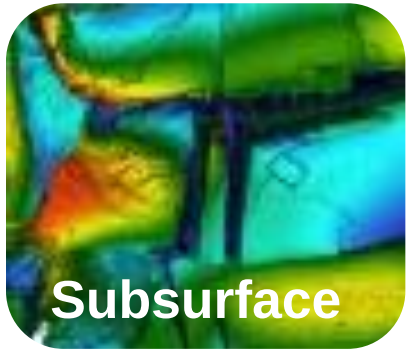

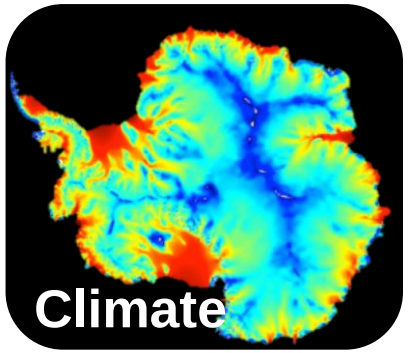

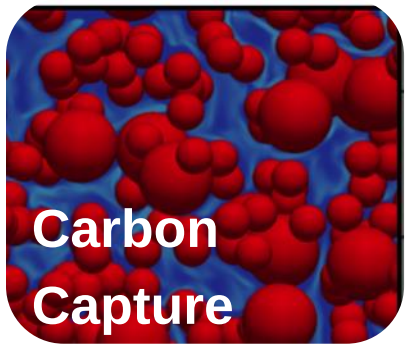

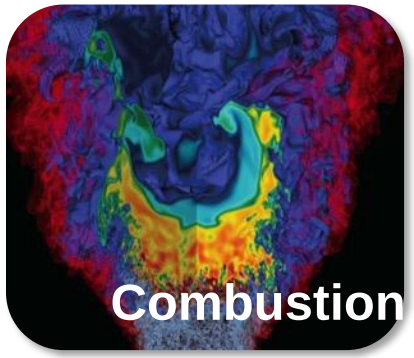

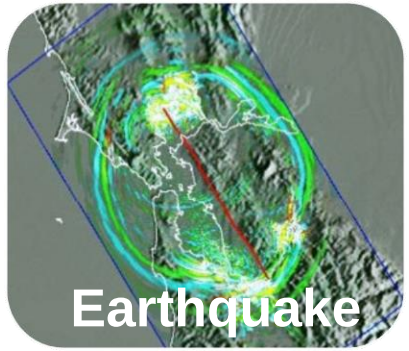

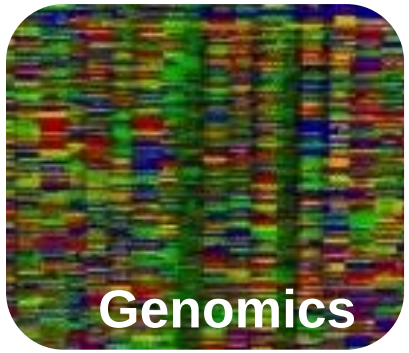

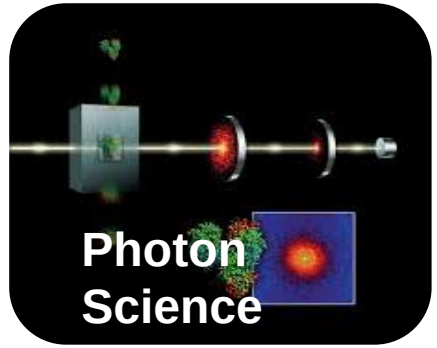

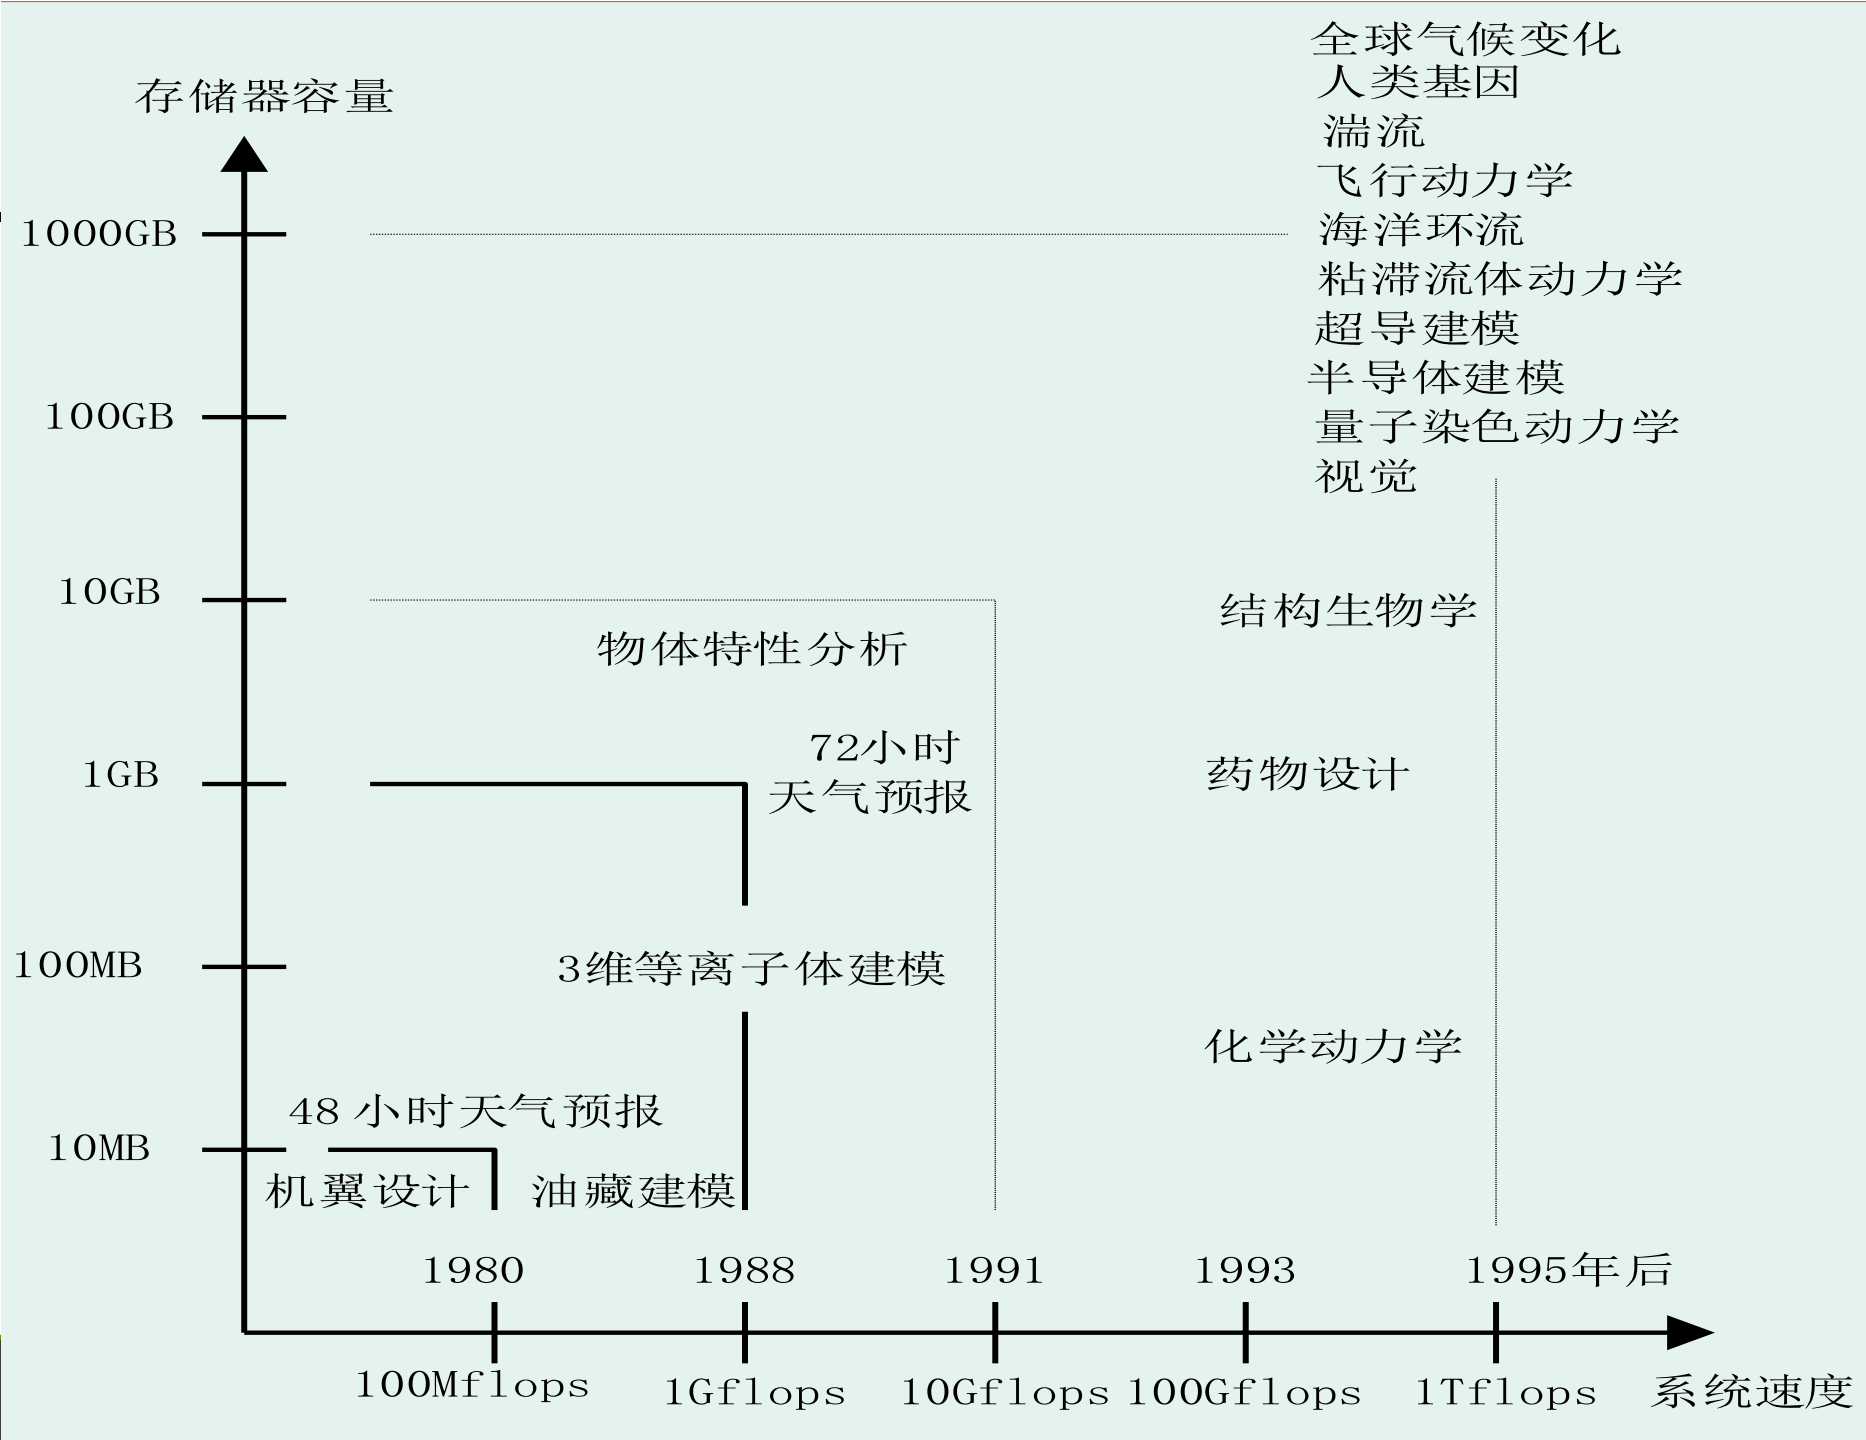
存储器容量
1000GB
100GB
10GB
1GB
100MB
10MB
全球气候变化
人类基因
湍流
飞行动力学
海洋环流
粘滞流体动力学
超导建模
半导体建模
量子染色动力学
视觉
结构生物学
物体特性分析
72小时
天气预报
药物设计
3维等离子体建模
化学动力学
48小时天气预报
机翼设计
油藏建模
1980
1988
1991
1993
1995年后
100Mflops
1Gflops
10Gflops
100Gflops
1Tflops
系统速度

Techniques behind them

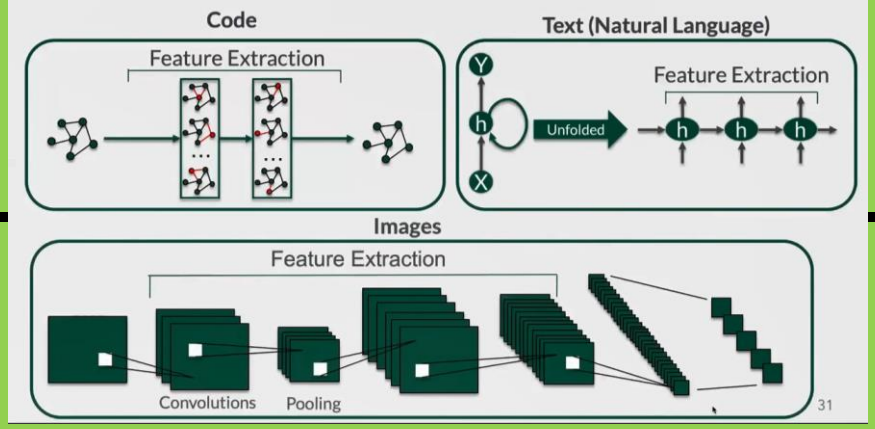

- **GPGPU, MPP, Cluster, … Data Center**
  - They are also used in HPC, DL, etc.
    - 2012年，在ImageNet ILSVRC (ImageNet Large-Scale Visual Recognition Challenge)竞赛中，深度学习的一个CNN (Convolutional Neural Network)模型 – AlexNet – 获得了冠军，识别率远远超过第二名（top 5 错误了16%，第二名为26%）。其后，深度学习轰然闯入机器学习的方方面面，沛然难挡。
      Alex Krizhevsky, Ilya Sutskever, Geoffrey E. Hinton. ImageNet classification with deep convolutional neural networks. Advances in Neural Information Processing Systems (NIPS). 2012(25): 84--90. https://dl.acm.org/doi/10.1145/3065386.
    - 2020年5月，OpenAI发布了无监督的转化语言模型GPT-3。这个模型包含1750亿个参数，训练数据量达到了45 TB（1万亿单词量）
    - 2021年6月，北京智源大会上，清华大学计算机科学与技术系长聘教授、计算机系副主任唐杰发布了智源“悟道”万亿参数智能预训练模型。唐杰介绍称，悟道2.0是目前全球最大的大规模智能模型系统，该系统参数数量已超过1.75万亿

Techniques behind them

- GPGPU, MPP, Cluster, … Data Center

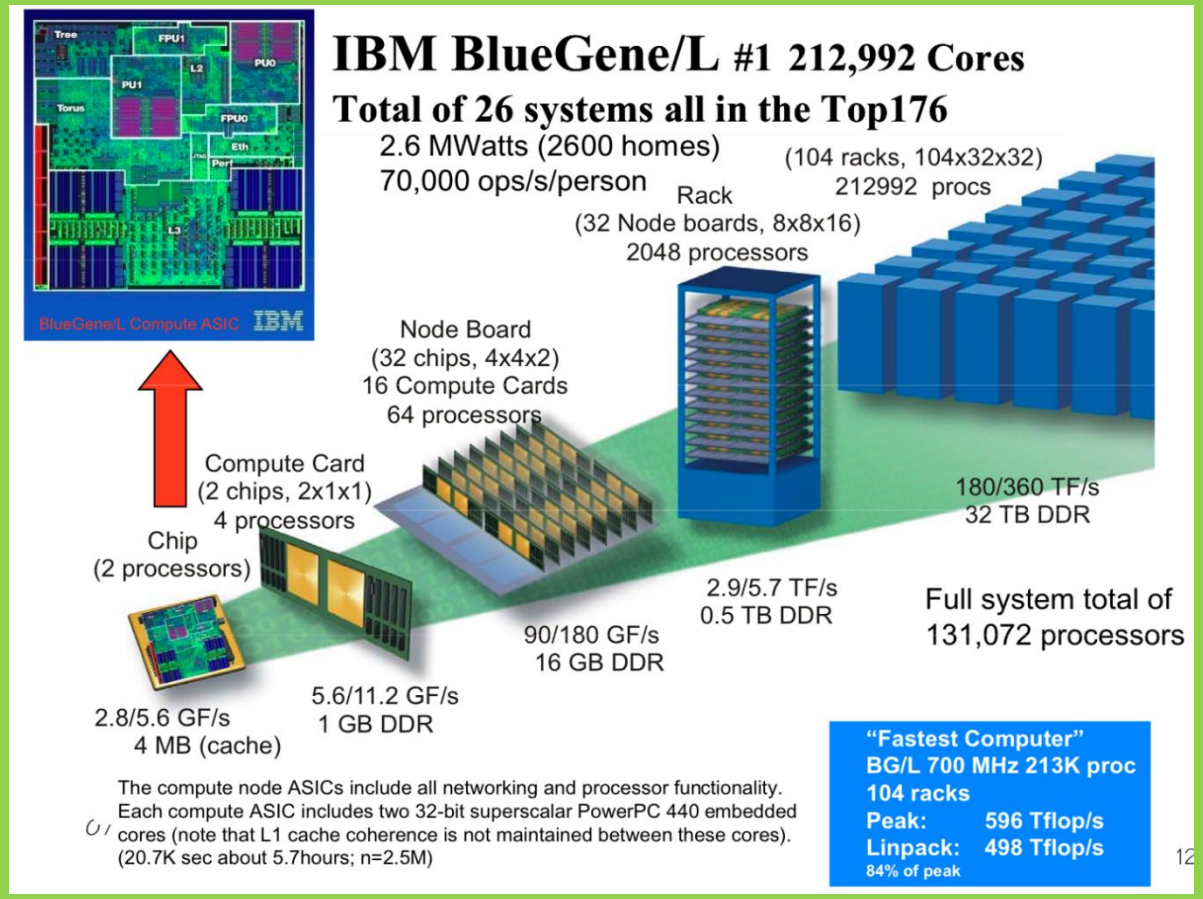

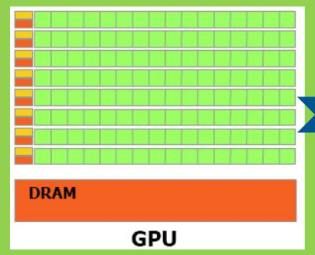

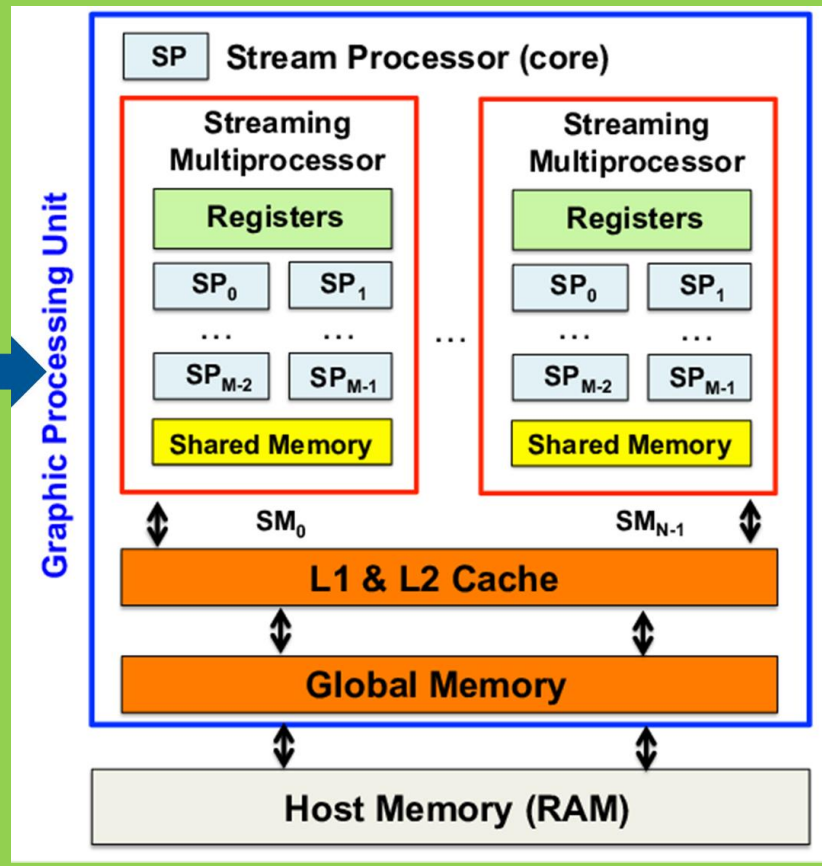

Machine learning demands more computing

300,000x increase from 2011 (AlexNet) to 2018 (AlphaGoZero)

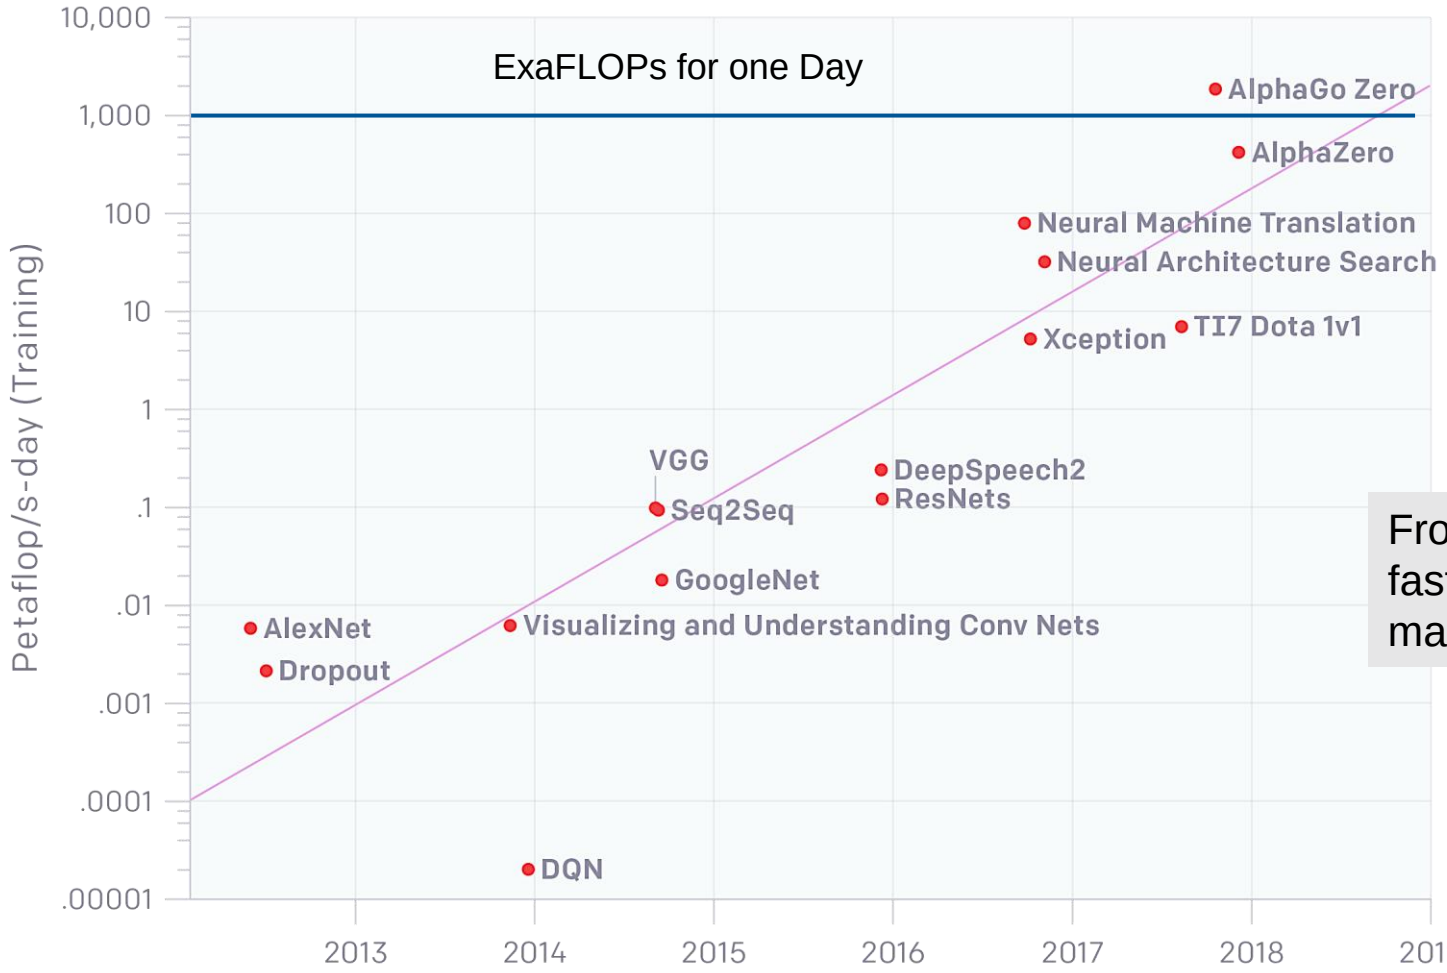

From 2011-2017 the fastest Top500 machine grew < 10x

https://blog.openai.com/ai-and-compute/

Techniques behind them – 也同样用于商务

- GPGPU, MPP, Cluster, … Data Center

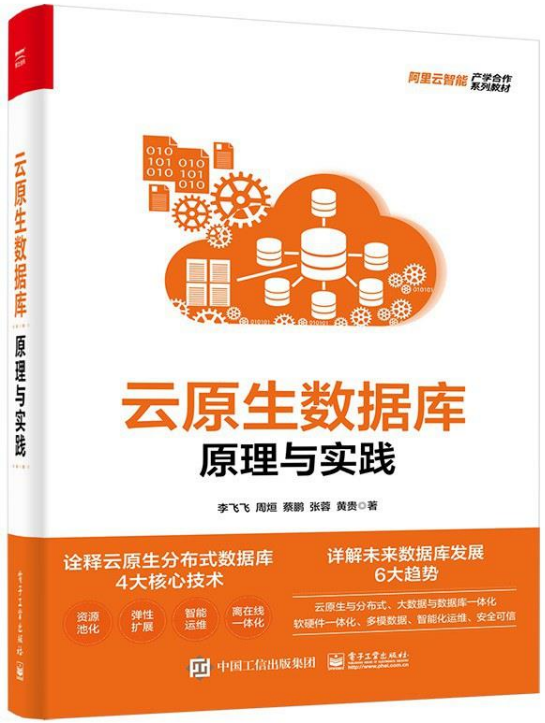
阿里云智能 产学合作系列教材
云原生数据库
原理与实践
李飞飞 周煊 蔡鹏 张蓉 黄贵 著
诠释云原生分布式数据库4大核心技术
资源池化
弹性扩展
智能运维
离在线一体化
详解未来数据库发展6大趋势
云原生与分布式、大数据与数据库一体化
软硬件一体化、多模数据、智能化运维、安全可信
中国工信出版集团


阿里云
2021阿里云开发者大会
阿里云数据库开源计划：打造云原生分布式数据库生态
MySQL生态
ApsaraDB GalaxySQL
X-Engine
Distributed File System
基于X-Paxos的PolarDB分布式
PG生态
DataNodes(PG11)
SQL
HLC
Cluster Manager
Coordinator Nodes
Tnx Manager
数据库正在加速云化。云原生以及分布式技术正在重塑数据库整个技术栈；
阿里云在自身互联网业务和云数据库服务有丰富的实践经验；
在高可用、分布式、云原生、存计分离有技术积累；
这些技术以组件和系统的方式开放出来，与开源社区一起共建云原生分布式数据库生态。
高可用MySQL：使用Paxos协议同步日志的高可用数据库
ApsaraDB GalaxySQL: 阿里巴巴MySQL分支
•超10年的技术积累，阿里核心业务场景的考验
•阿里云RDS MySQL服务的核心基础
•线程池、并发控制等优化技术增强事务处理吞吐与热点处理能力
•安全，可靠，稳定与基于运维的支撑体系
X-Engine:阿里巴巴自主研发的数据库存储引擎
•兼容MySQL生态的事务存储引擎
•自动识别冷热数据分离与分层存储
•支持共享存储，一写多读云原生架构
存储成本2~5倍
X-Paxos：确保数据强一致的高可用协议库
•分布式系统的数据一致性保障
•批量协议打包与流水线发送机制提升复制数据的吞吐
•网络条件探测自适应发包，跨数据中心长传链路优化
PolarDB for PostgreSQL: PolarDB基于PG生态研发的企业级分布式数据库
•基于混合逻辑时钟(HLC)的全局一致性
•分布式事务(分布式MVCC+2PC)
•分布式SQL计算(兼容单机SQL功能，分布式DDL和分区表支持)
•使用X-Paxos同步日志的高可用方案
•插件化分布式sharding提升扩展规模

华为的


华为openGauss开源数据库实战
openGauss数据库源码解析
openGauss数据库核心技术
数据库原理与技术实践教程——基于华为GaussDB
数据库原理与技术——基于华为GaussDB


华为云智能数据湖FusionInsight围绕数据全生命周期构建
生态伙伴
东华软件
中软国际
华傲数据
汇智通信
先进数通
长亮科技
宇信科技
北明软件
太极股份
电信十所
文思海辉
易华录
亚信科技
神州数码
东方国信
中油瑞飞
数据使能
资产沉淀
数据运营（DAYU）
数据集成管理
数据治理
数据服务
数据集成开发
统一资产目录
统一安全管理
数据隐私保护
数据湖查询分析引擎（HetuEngine）
数据虚拟化
跨源跨域连接
统一SQL入口
多元计算
大数据 MRS
数据仓库 GaussDB(DWS)
+AI
ModelArts
数据存储引擎（CarbonData）
物理机
虚机
容器
多样性算力
语音视频
图形图像
文本
时序
其他
统一存储


2020天猫双11全球狂欢季
2020天猫双11全球狂欢季 (11.1-11)
订单创建峰值
58.3
万笔/秒

而 DL 也成为了当下“精准营销”/“计算广告” 的核心算法

- 40%+的Google 收益来自广告收益!
  - Google的年利润

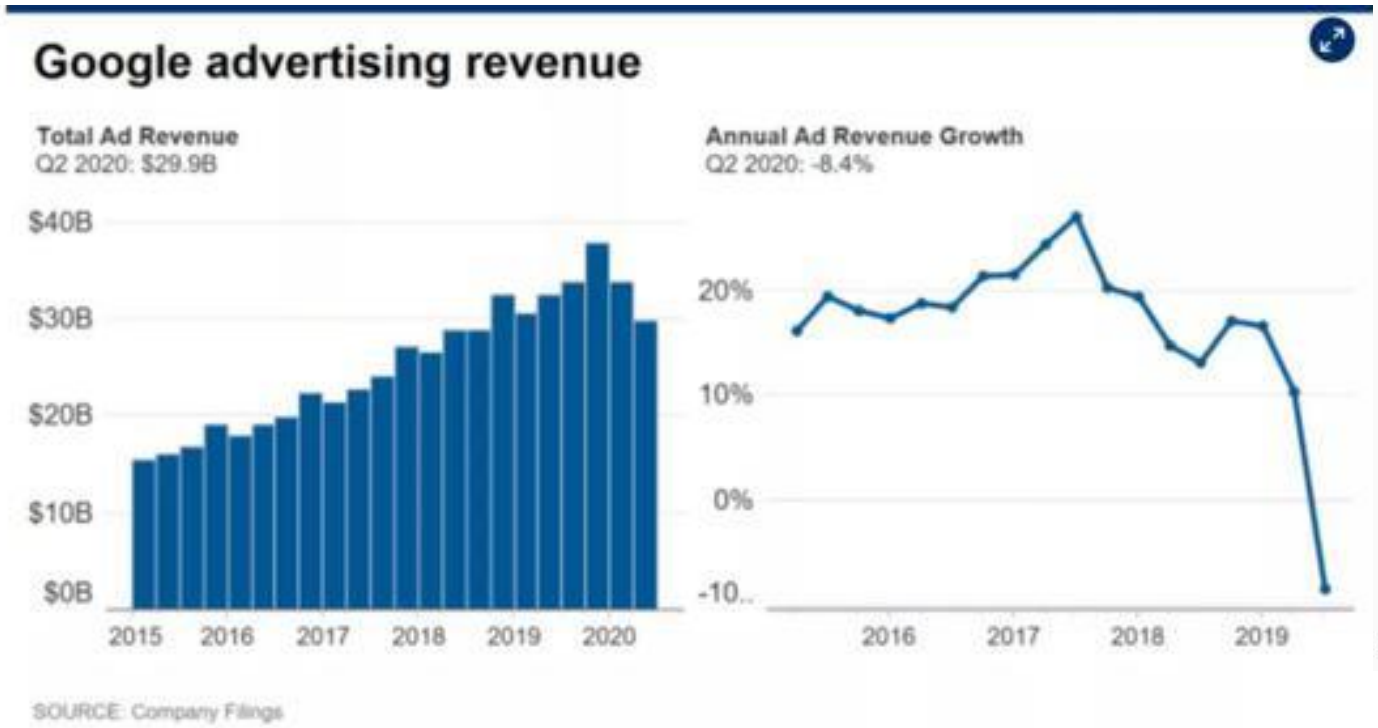

https://tech.sina.com.cn/i/2019-12-04/doc-iihnzhfz3490699.shtml

http://tech.sina.com.cn/csj/2020-07-31/doc-iivhuipn6034442.shtml

精准营销/计算广告背后的是Large Scale Computing

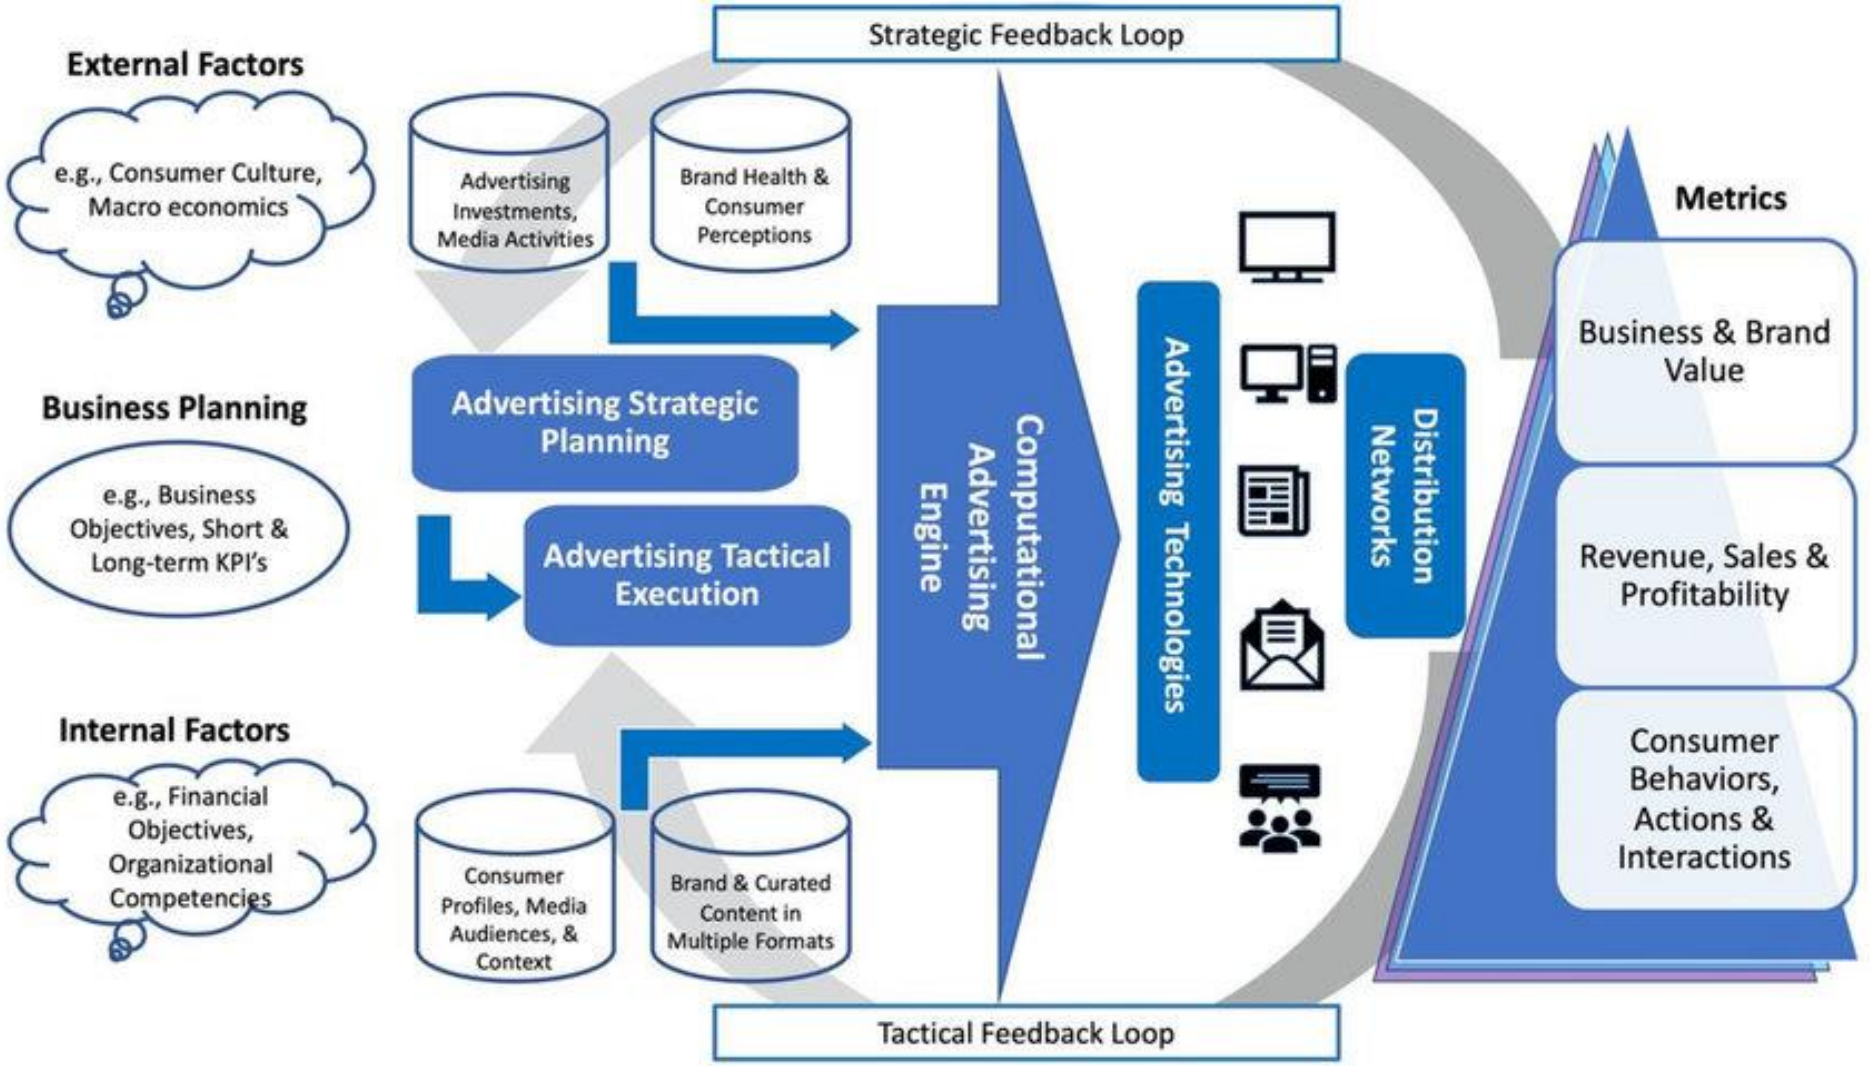

精准营销/计算广告背后的是Large Scale Computing

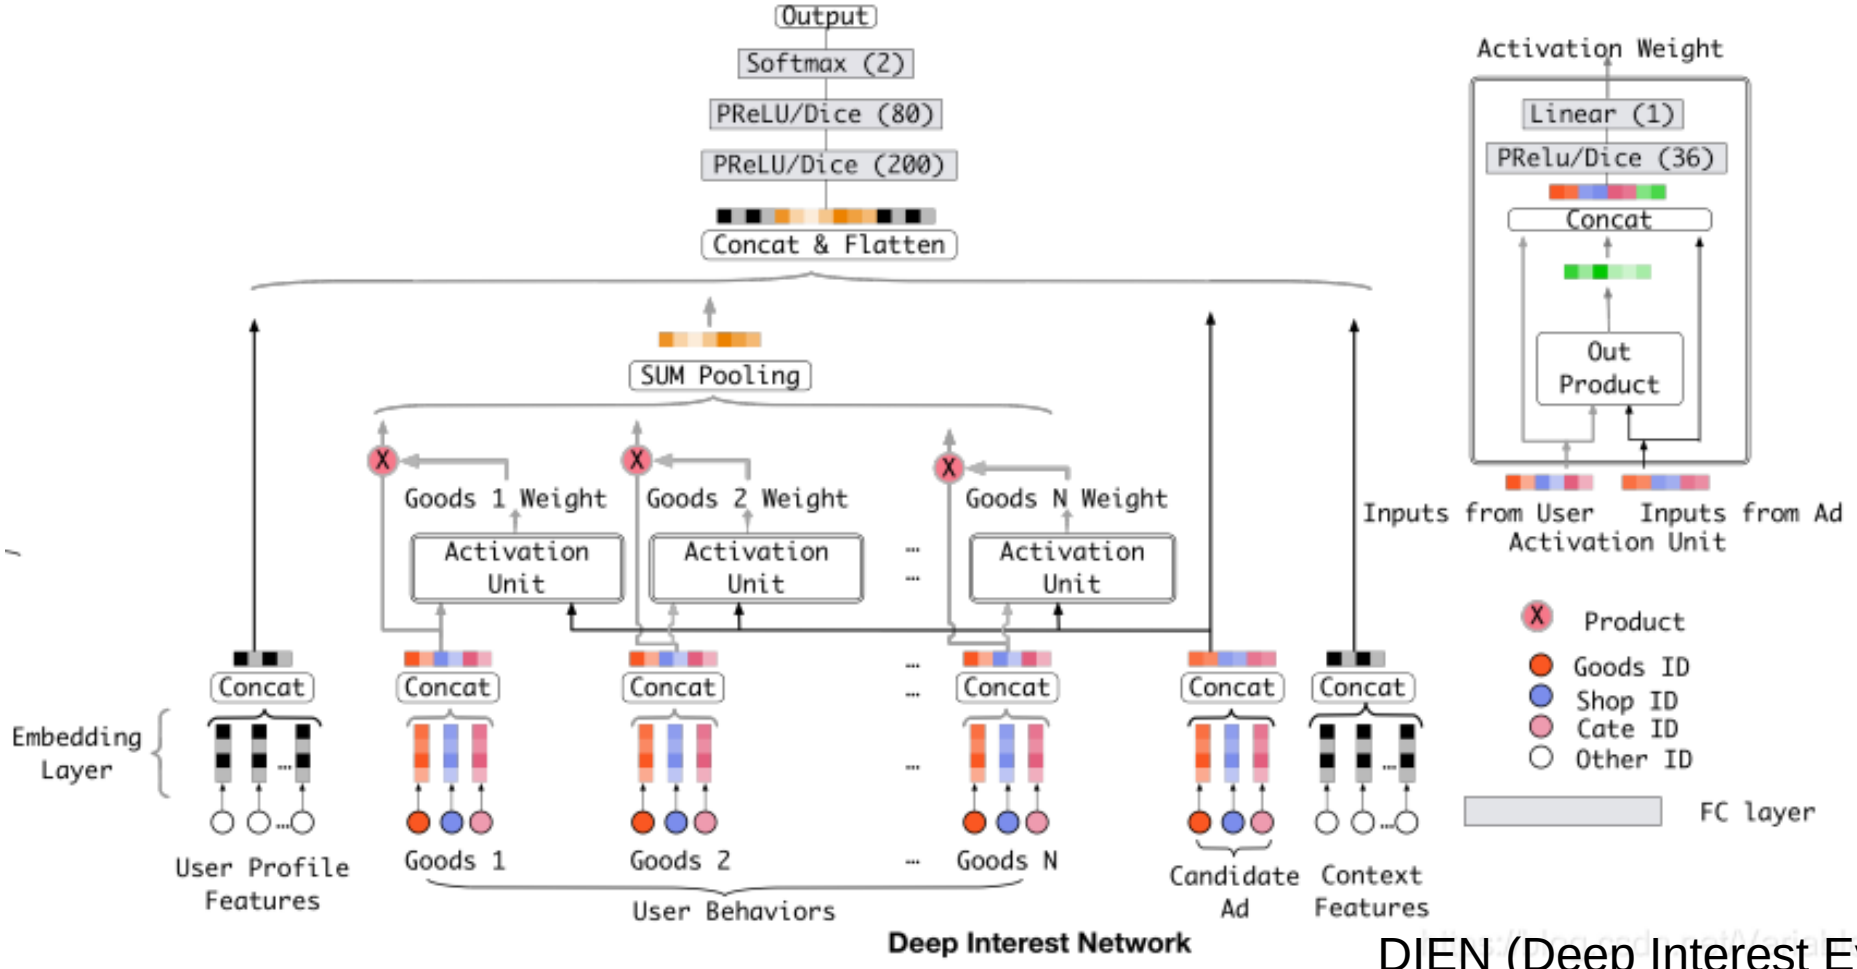

DIEN (Deep Interest Evolution Network)

https://blog.csdn.net/VariableX/article/details/108887709

大纲

- **前言**
  - 为什么需要“大规模计算” [HPC, DL, Business platform system, Cloud已经合流]
    - 导入 – 科学计算(天气预报)，DL, 互联网平台(Google, Amazon, Alibaba, MeiTuan, …)
- **基础篇**
  - 并发程序的样子 – Divide & Conquer, Model & Challenges, PCAM, Data/Task, …
    - 天气预报的计算
  - 运行环境
    - 硬件 – 自己梳理的3个方案 – Shared/Unshared Memory, Hybrid
    - 系统软件 – 协议栈, Modern OS, Distributed Job Scheduler, GTM等
- **算法级篇**
  - OpenMP, MPI, CUDA (DL的实现), Big Data 中的MR/Spark等 (只涉及在Big Data SDK之上的编程；大数据本身的介绍放到后一部分)
- **系统级篇 – 互联网平台的实现**
  - “秒杀”的技术架构
  - 计算广告
  - 系统架构 (HTAP等)
    - Flink, ClickHouse, MaxCompute, ELK …

Some Particularly Challenging Computations

- **Science**
  - Global climate modeling
  - Biology: genomics; protein folding; drug design
  - Astrophysical modeling
  - Computational Chemistry
  - Computational Material Sciences and Nanosciences
- **Engineering**
  - Semiconductor design
  - Earthquake and structural modeling
  - Computation fluid dynamics (airplane design)
  - Combustion (engine design)
  - Crash simulation
- **Business**
  - Financial and economic modeling
  - Transaction processing, web services and search engines
- **Defense**
  - Nuclear weapons -- test by simulations
  - Cryptography

SDSC: San Diego Supercomputer Center
SAC: Strategic Applications Collaborations

Modeling Galaxy Dynamics and Evolution

Project Leaders

- Lars Hernquist, Harvard and John Dubinski, U Toronto
- Stuart Johnson and Bob Leary, **SDSC SAC** Program

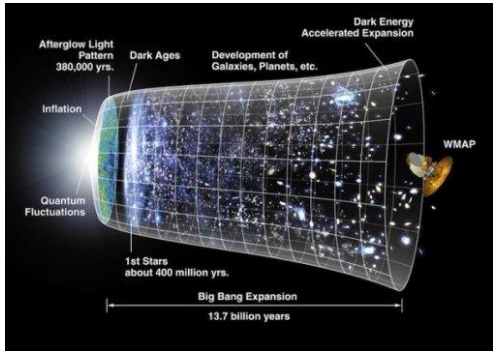

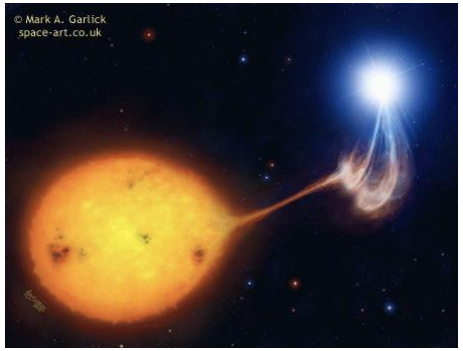

First images from Blue Horizon Simulation

- 24M particles = 10M stars + 2M dark matter halo in each galaxy
  - Working on 120M particle run
- Run on all 1,152 processors during acceptance

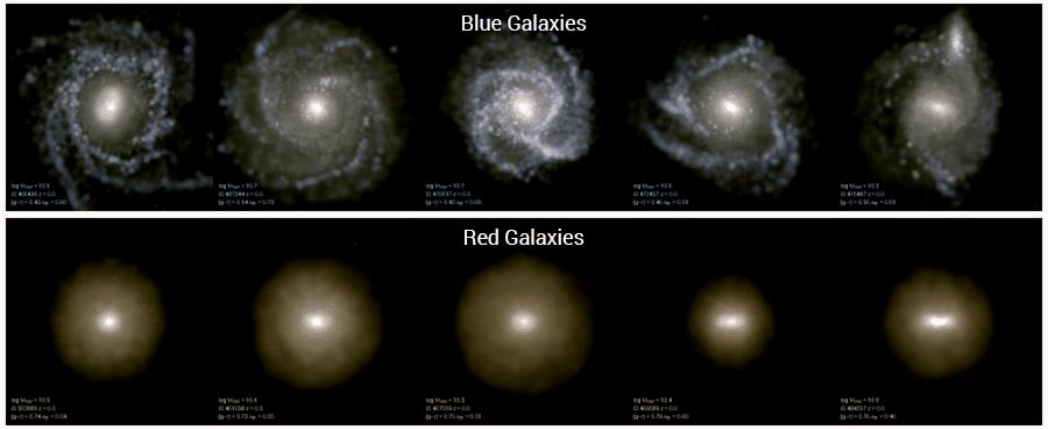

"Due to SAC efforts, our simulations run two to three times faster, we can ask more precise questions and get better answers" ...Hernquist

Chapter 1: Introduction

- **About me**
- **Evaluation in this course**
- **Why do we have this course, and How do I organize?**
  - Modern Society needs more computation power
  - Organize the topics with an example
- **Resources related to this course**