

Biology Is All You Need

David A. Markowitz, PhD | Program Manager | November 16, 2022



Intelligence Advanced Research Projects Activity

 I
 A
 R
 P
 A

 Creating Advantage through Research and Technology







Agenda







DNA Data Storage Neural Computing

UNCLASSIFIED

1



Part 1: Data Storage



Intelligence Advanced Research Projects Activity

Creating Advantage through Research and Technology



The Storage Crisis



Global Data Demand

Global Storage Supply



Source: IDC Global DataSphere Forecast, 2021-2025.

INTELLIGENCE ADVANCED RESEARCH PROJECTS ACTIVITY (IARPA)



Why Not Just Make More Storage?



Data centers are *very* resource intensive



Source: CERN

750,000 sq ft over 110 acres Power: 200 MW 5-year Cost: \$100+ million

Scaling to a *zettabyte* (1,000x) is impossible with this model



Wish List for Next-Generation Storage



Higher storage density



Higher data stability





No degradation

INTELLIGENCE ADVANCED RESEARCH PROJECTS ACTIVITY (IARPA)



Biology Has Already Solved This Problem







DNA sequences are stable for <u>700k+years</u> (under the right conditions)

Orlando, L. et al. Nature (2013) http://dx.doi.org/10.1038/nature12323

Data figure courtesy Victor Zhirnov, Semiconductor Research Corporation, and Karin Strauss, Microsoft Corporation

DNA Data Storage Workflow

LS YEARS





Largest <u>published</u> DNA data archive: ~200MB

Organick et al. Nat Biotechnol (2018) https://doi.org/10.1038/nbt.4079

• Required 9 separate synthesis runs



Molecular Information Storage (MIST) Program



Goals: By 2025, make DNA data storage

• Scalable 20 MB/system → 1 TB/system

• Economical $$100k/GB \rightarrow $1/GB$

 Practical for enterprise archival use (end-to-end workflows on a tabletop)







Present



Source: CERN

Future



Source: Wikipedia

INTELLIGENCE ADVANCED RESEARCH PROJECTS ACTIVITY (IARPA)



Science Fiction \rightarrow Science Fact







DNA Synthesis Capacity:128M oligosWrite Capacity:2 GB dataCost:\$100/run







Credit: Tara Brown Photography / University of Washington

INTELLIGENCE ADVANCED RESEARCH PROJECTS ACTIVITY (IARPA)



MIST Synthesis Platforms







MIST Sequencing Platforms



Modified Commercial Instrument





Novel Biosensor Platform









Channel Error Model



"Base" level errors

- Substitutions: ACTG → ACTA
- Insertions: ACTG → ACCTG
- Deletions: ACTG → ACG
- Block insertions, block deletions, ...
- Per-base error rate may be 5% for DNA write and read channels

"Oligo" level errors

- Erasures: ACTG, ACAC, GGTC → ACTG, GGTC
- Secondary structure, bias



MIST Codec



LOS Alamos

https://github.com/lanl/adscodex

Open source:

"ADS Codex"

- Developed by MIST partners at LANL
- Tunable robustness to *any* DNA write/read error model

Inner Codec



Outer Codec

Data Oligos

Column _o	Column ₁	Column ₂	Column ₃	Column₄	Oligo _a
Column ₁	Column ₂	Column ₃	Column₄	Column _o	Oligo _b
Column ₂	Column ₃	Column₄	Column _o	Column ₁	Oligo _c
Column ₃	Column₄	Column _o	Column ₇	Column ₂	Oligo _d
Column₄	Column _o	Column ₁	Column ₂	Column ₃	Oligo _e
Column _o	Column ₁	Column ₂	Column ₃	Column ₄	Oligo _f

Erasure Oligos

INTELLIGENCE ADVANCED RESEARCH PROJECTS ACTIVITY (IARPA)



Industry Is Building On These Advances



INTELLIGENCE ADVANCED RESEARCH PROJECTS ACTIVITY (IARPA)



Cost of Writing and Storing 1PB



Source: DNA Data Storage Alliance

No and a second

INTELLIGENCE ADVANCED RESEARCH PROJECTS ACTIVITY (IARPA)



Part 1: Conclusion



- Biology has solved fundamental problems in data storage
- The semiconductor industry offers tools to write/read DNA at scale
- <u>Coming Soon:</u> Archival DNA storage appliances
- Benefits: Keep all the world's data, cheaply and compactly, in DNA



Part 2: Computing



Intelligence Advanced Research Projects Activity

Creating Advantage through Research and Technology





The Computing Crisis



• AI models are growing exponentially

• Resource requirements are extreme

• Scaling 1,000x and beyond looks prohibitive, even with AI accelerators



Trainable Parameters

Sources: arXiv:2202.05924 arXiv:2206.05229



Power Is Already A Bottleneck



SIGN IN / UP	The A Register [®]
ON-PREM	
The US's bigges	st datacenter market is short o

electricity

Power transmission bottlenecks could delay Northern Virginia DC projects into 2026

A Brandon Vigliarolo

Fri 29 Jul 2022 16:30 UTC

21

Q

UPDATED The largest datacenter market in the US is running into trouble: There isn't enough power transmission capacity in the region to handle all the bit barn projects.

That's the word from Andy Power, president and CFO at Digital Realty, a real estate investment trust that owns and maintains more than 290 datacenters around the globe. Power shared the news, which he described as "very recently breaking" in an <u>earnings call</u> yesterday.

Dominion Energy, the primary power provider in that market, Power said, had contacted Digital Realty and its other major customers in the region to inform them of a "pinch point" in eastern Loudoun County, Virginia, that could result in delivery delays stretching all the way to 2026.

"If this is to come to fruition as we recently learned, it will obviously likely be a slowdown in delivery of new supply in what is our largest and the largest and most consistently in demand datacenter market in the world," Power said.

The apparent cause isn't due to electricity generation, Power said, but transmission – there simply aren't enough lines to move enough energy.

Loudon County <u>describes itself</u> as the largest concentration of datacenters in the world, and with more than 25 million square feet of datacenter space and four million more in development, it might be right.

Login / Create Account Q Menu =

HOME > NEWS > CRITICAL POWER

Report: Home building to halt in West London, due to data center power demands

All the electricity capacity is already used up, GLA tells developers

July 28, 2022 By: Peter Judge O Be the first to comment

The Greater London Authority has told developers that new housing projects in West London could be banned till 2035, because data centers have taken all the electricity capacity.

New projects are being rejected in three west London boroughs because the grid has run out of capacity for new homes, the GLA said according to a report in the <u>Financial Times</u>. The Authority says it will take more than a decade to bulk up grid capacity to allow new developments.



work will take years.

Three boroughs - Hillingdon, Ealing, and Hounslow will have to wait a decade for electricity supplies to reach a level where new homes can be built, the GLA says, in a note to developers seen by the FT. Applicants for electrical power have been told there is not enough capacity for a new connection until 2035.

The GLA says data centers are to blame for acute pressure on the west London grid, as many have been built along the M4 corridor, close to fiber optic cables. A single data center can use the same amount of electricity as thousands of homes, leaving a shortage for other customers.

Data centers use electricity "equivalent of towns or small cities," the GLA note says, "to power servers and ensure resilience in service". The National Grid and utility SSEN are upgrading power networks, but the

London has set targets for house building which depend on activity across the city, and the three boroughs built 11 percent of London's housing supply in 2019–2020 (some 5,000 homes). This means the ban will increase an already chronic housing shortage, the FT says.

Major recent developments in the area include <u>Virtus' Stockley Park campus</u>, which opened three years ago with four data centers, an <u>Ark Data Centres development in Park Royal</u> and numerous facilities

THE SAME TIMES THURSDAY NOVEMBER 3 2022 BUSINESS Doubt over tech giants' €2bn data investment

Microsoft and Amazon are looking elsewhere after grid connection lockout

Laura Roddy Sunday August 21 2022, 12.01am BST, The Sunday Times

he tech giants Microsoft and Amazon have put €2 billion of investment in three Irish data centres under review.

A decision to explore alternative locations was made shortly after the Commission for the Regulation of Utilities (CRU) announced in November last year that there would be a moratorium placed on data centres in the greater Dublin region due to grid constraints.

An industry source said that Microsoft was sizing up alternative locations in London, Frankfurt and Madrid, while Amazon has decided to build its data centre near London. While both companies had received planning permission for their facilities, it is understood that Amazon had not received a connection from EirGrid and had been told it would not be eligible for one.



Biology Has Already Solved This Problem



>100T synapses (plus many other parameters)



• Human brain consumes ~20W of power, of which nearly half (9W) is lost to heat

Levy & Calvert, PNAS 2021 https://doi.org/10.1073/pnas.2008173118

 Power cost of "training" a brain for 18 years is equivalent to powering <u>0.3 home for 1 year</u>

 Proof by existence that general intelligence can be implemented – and efficiently



Bridging The Performance Gap



Neuromorphic Hardware



Inspired by the brain: Event-based communication Parallel sparse compute Processing-in-memory

- As DNN inference accelerators, can be >100x more efficient than GPUs [1]
- But... we need **1,000,000x** better power efficiency to rival biology
- <u>Gap:</u> Neuromorphic computing models lack explicit guidance from the brain's
 - Data representations
 - Data transformations
 - Learning rules

[1] April 2022: https://www.eetasia.com/scaling-and-low-power-requirements-for-neuromorphic-computing/



Machine Intelligence from Cortical Networks (MICrONS) Program



Goal: Revolutionize machine learning by identifying and then exploiting the mathematical functions underlying the algorithms of the brain.





MICrONS Functional Imaging





Baylor College of Medicine

INTELLIGENCE ADVANCED RESEARCH PROJECTS ACTIVITY (IARPA)

Blank 0.3 – 0.5 sec

Image 0.5 sec 5 min



MICrONS Structural Imaging





Automatic collection of 40 nm tissue sections onto GridTape

Low-res scan of a section in one GridTape slot

Detail view of one tile within the 1x1 mm ROI



ULEBRATTING 15 YEARS

MICrONS Image Alignment











Automated Segmentation & Synapse Identification





Image Segmentation Agglomeration Across Sections Morphological Recovery Synapse Identification

Visualizations courtesy Jeff Lichtman

INTELLIGENCE ADVANCED RESEARCH PROJECTS ACTIVITY (IARPA)

UNCLASSIFIED

27



MICrONS Data Volume





INTELLIGENCE ADVANCED RESEARCH PROJECTS ACTIVITY (IARPA)



Automated Segmentation & Synapse Identification





~**3000** NVIDIA **T4 GPU** (~2 day completion time) Produced ~3 PB of additional data (affinity map, PSD map)



"Chunkflow: Distributed Hybrid Cloud Processing of Large 3D Images by Convolutional Nets" https://arxiv.org/abs/1904.10489

INTELLIGENCE ADVANCED RESEARCH PROJECTS ACTIVITY (IARPA)

UNCLASSIFIED

29



MICrONS Automated Image Segmentation







Layer 4



Pia, Layer 2/3



Layer 6 / white matter









MICrONS-Explorer.org



 $\leftarrow \rightarrow C$ 目 ☆ ⊘ 👱 🖆 ○ A https://www.microns-explorer.org/ MICrONS Explorer Home Data Tools Gallery About A functional connectome containing 200,000 cells, 75,000 neurons with physiology, and 523 million synapses. This resource provides interactive visualizations of anatomical and functional data that span all 6 layers of mouse primary visual cortex and 3 higher visual areas (LM, AL, RL) within a cubic millimeter volume.

Free access courtesy of:

Soogle Cloud Public Datasets



EXPLORE



What Learning Algorithm Does The Brain Use?



- **Prediction:** Most learning models predict graded analog synapse sizes
- MICRONS Data: Excitatory synapses in cortex do not match this prediction



Best fitting model:

Sum of a binary variable and an analog variable drawn from a log-normal distribution.

"Binary and analog variation of synapses between cortical pyramidal neurons" https://doi.org/10.1101/2019.12.29.890319

INTELLIGENCE ADVANCED RESEARCH PROJECTS ACTIVITY (IARPA)

How Does Computation Map to Cellular Compartments?

 MICRONS Data: Inhibitory neurons target specific cellular compartments to shape computation by excitatory neurons.

with data from Nuno daCosta (AIBS)

How Does Long-Range Connectivity Relate to Function?

Long-Range Feedback

INTELLIGENCE ADVANCED RESEARCH PROJECTS ACTIVITY (IARPA)

Part 2: Conclusion

- Biology has solved fundamental problems in efficient computing
- Neuroscience offers tools to map the brain's solutions
- Available Today: Free, public maps of cortical activity and connectivity
- **Benefits:** Clear guidance for ML/AI algorithm development

Conclusion

Intelligence Advanced Research Projects Activity

Biology Is All You Need

- Biology has solved fundamental problems in data storage & computing
- These solutions can have profound real-world benefits
- Bridging biology, engineering, and computer science is the key