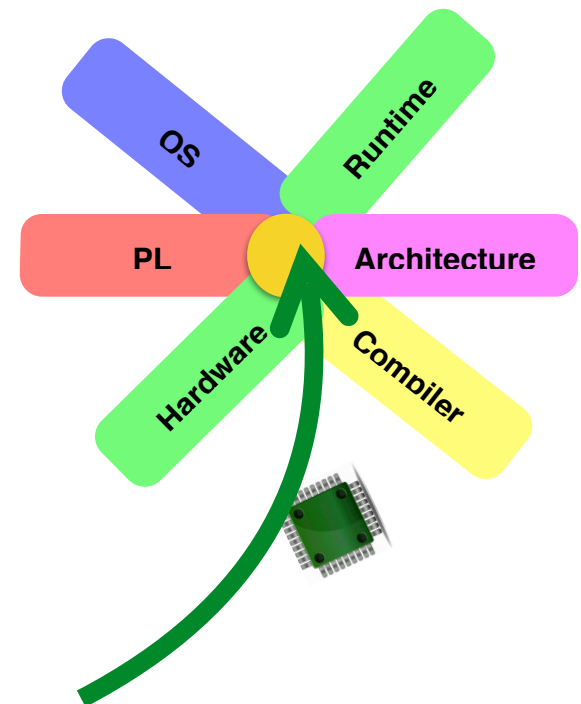
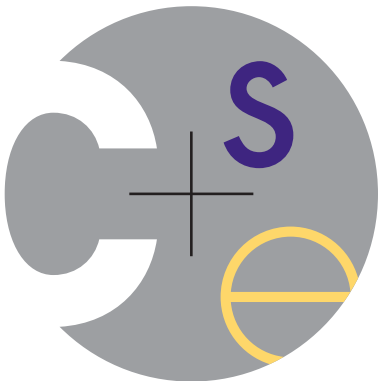


Imperfection is Beautiful and Efficient

Luis Ceze

University of Washington

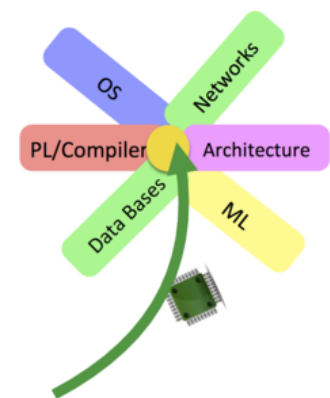


I do “computer architecture”

oh, you make computers
more beautiful?



Sort of, I like making them
more beautiful and efficient!



So, what's computer architecture research?

Finding new ways of building computers that are **more efficient**, **faster**, can meet **new constraints** etc...

core *problems*:

performance
reliability
energy



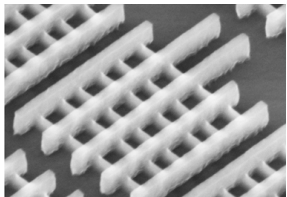
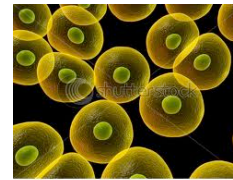
application domains:

big
small
tiny
nano



device technology:

transistors
3d chips
photonics
storage cells



(fast, non-volatile!)



Facebook Stores 240 Billion Photos And Adds 350 Million More A Day



JULIE BORT



JAN. 30, 2013, 6:48 PM

🔥 1,265

💬 1



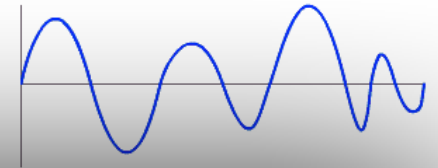
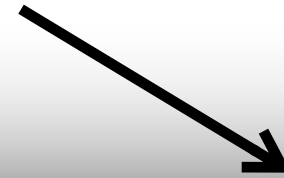
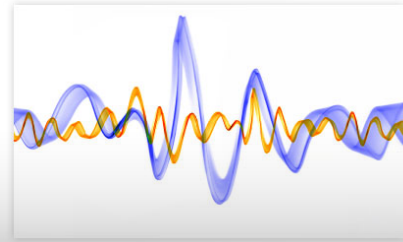
Netflix Now Accounts for 34 Percent of US Internet Traffic at Peak Times

Associated Press, May 14, 2014

YouTube: 100 hours of video uploaded per minute. 4 billion video views a day.



image, sound
and video processing



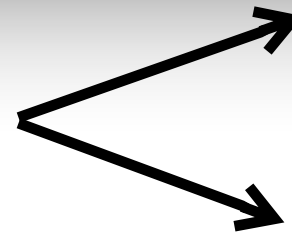
These applications consume a lot
(most?) of our **cycles/bytes/bandwidth!**

Often **input data is inexact by nature**
(from sensors)

They have **multiple acceptable outputs**

simulations, games,
search, machine learning

```
0101010100110  
100110011010100  
101001101011010  
111011110101001  
100010110010010  
001001000010001  
1010101010101  
001001000010001
```



```
0110101  
1111010  
1111010  
0110101  
1111010  
1111010
```



Yet, we design computer systems like this..

Lower layers work hard to expose a **general, reliable, precise and (mostly) deterministic** interface.

At a **big** efficiency cost!

Application

Language

Compiler

ISA/Architecture

Circuits

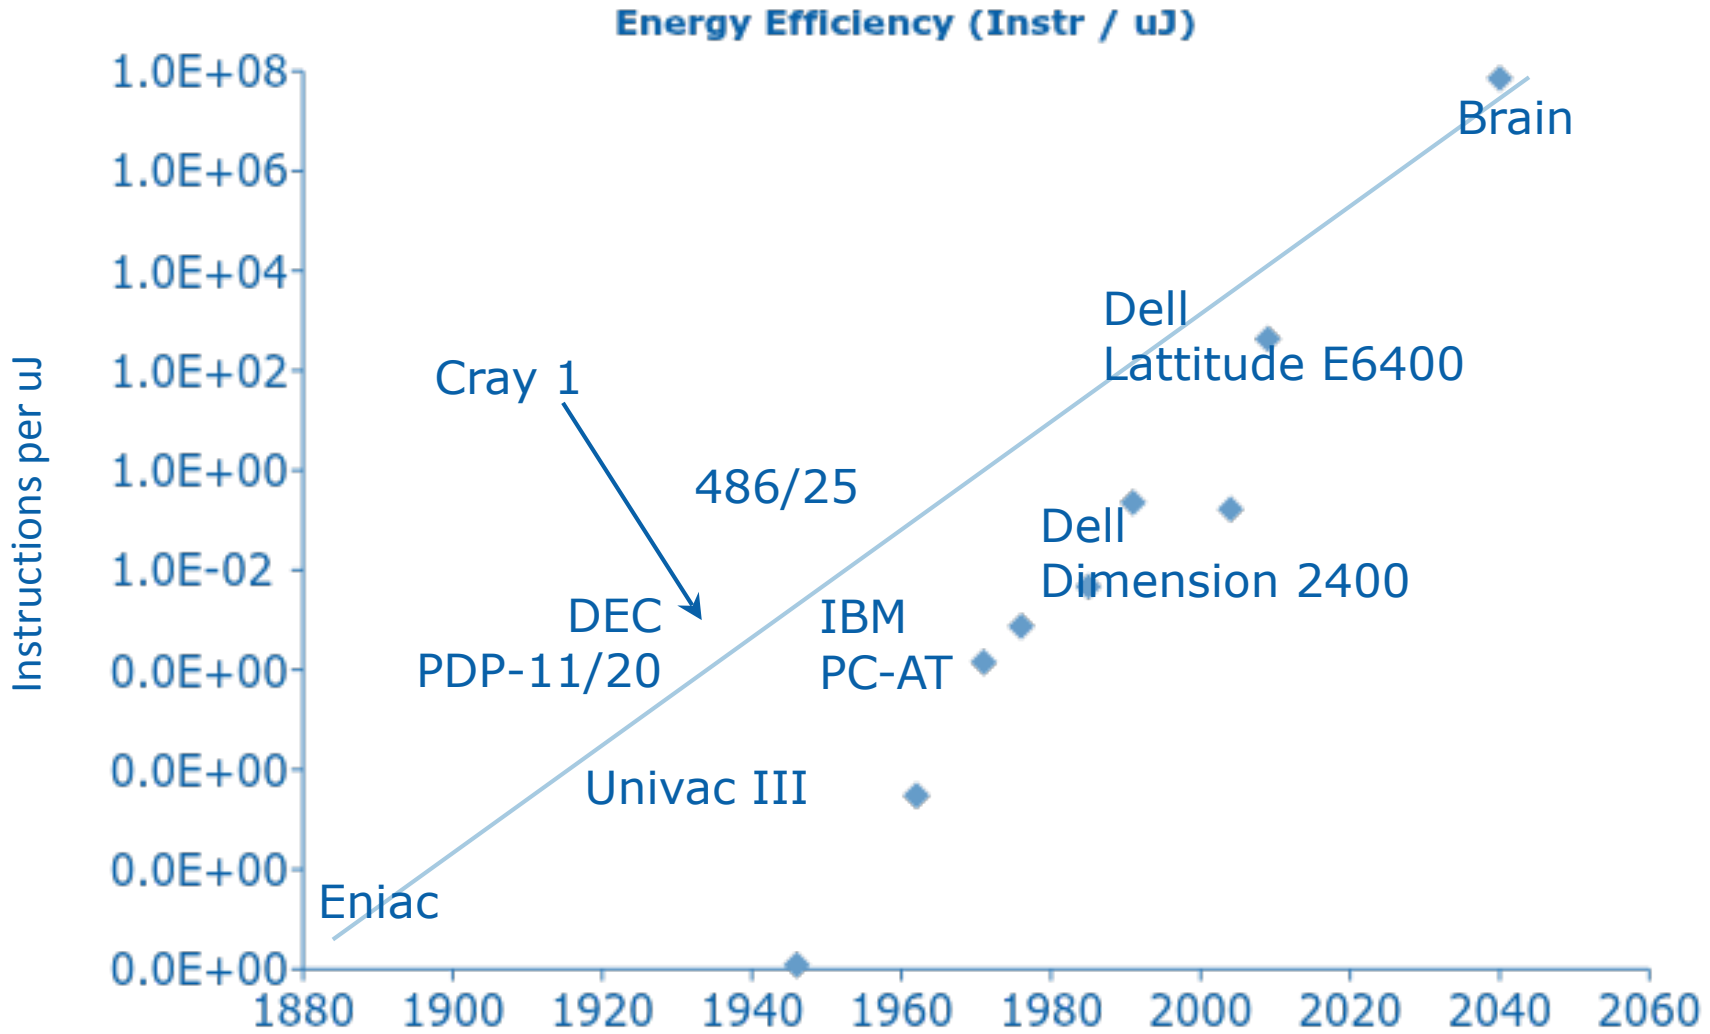
Physics



250W

Energy Efficiency Scaling

Joshua Smith, UW



What is “approximate computing”?



Building **acceptable** systems out of unreliable/
inaccurate hardware and software components

Efficiency and
performance

Output accuracy



Application

Language

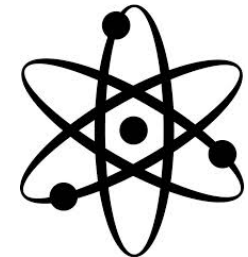
Compiler

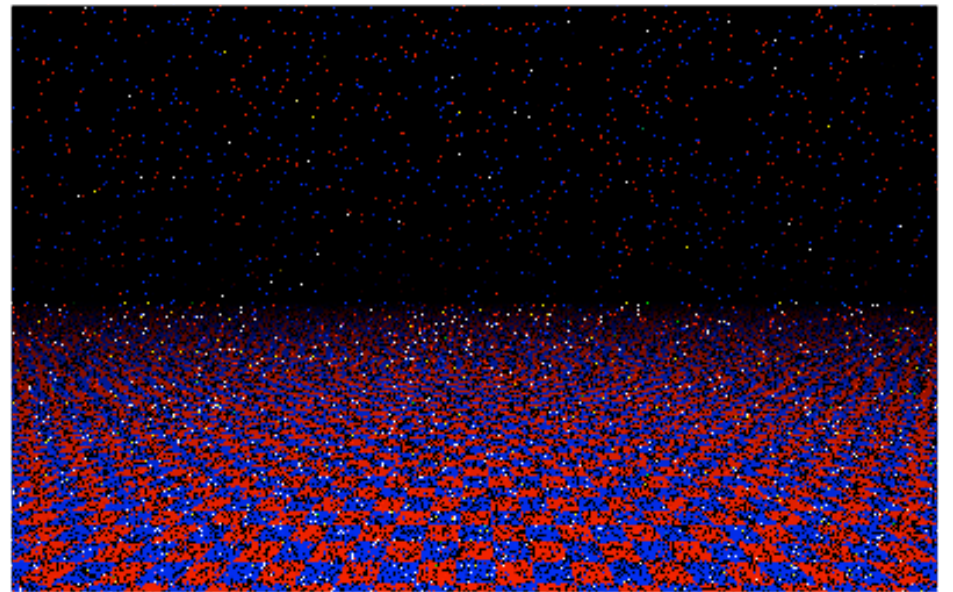
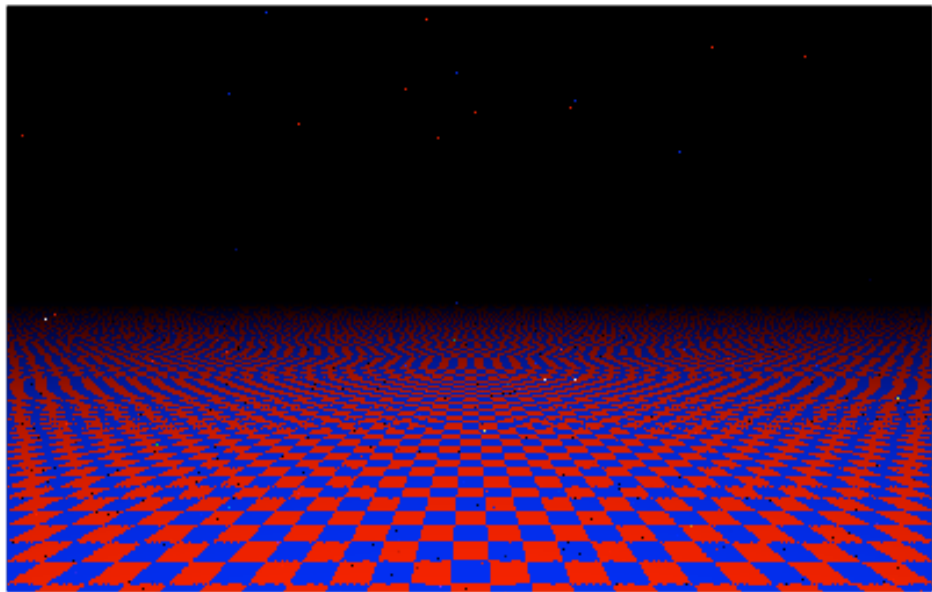
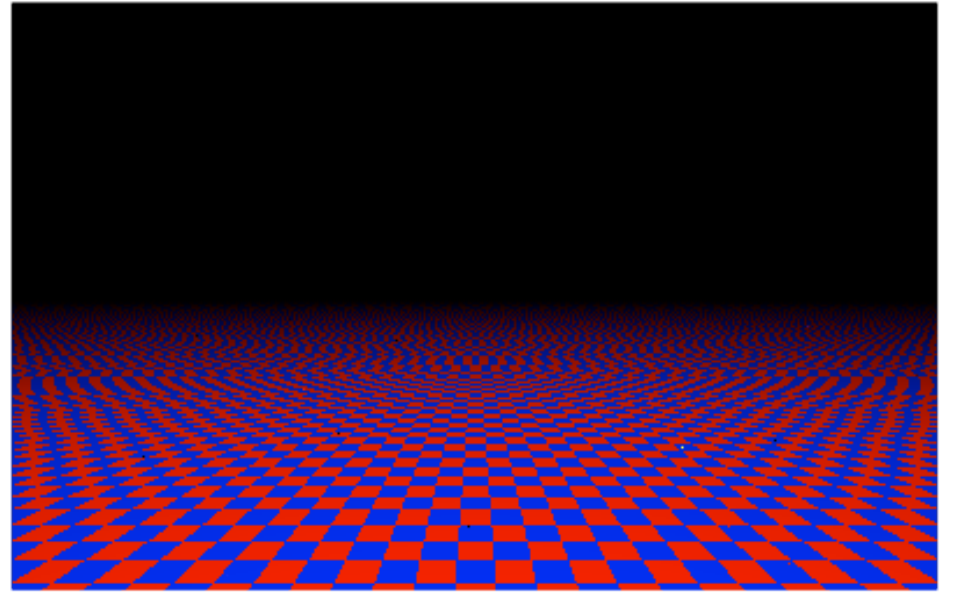
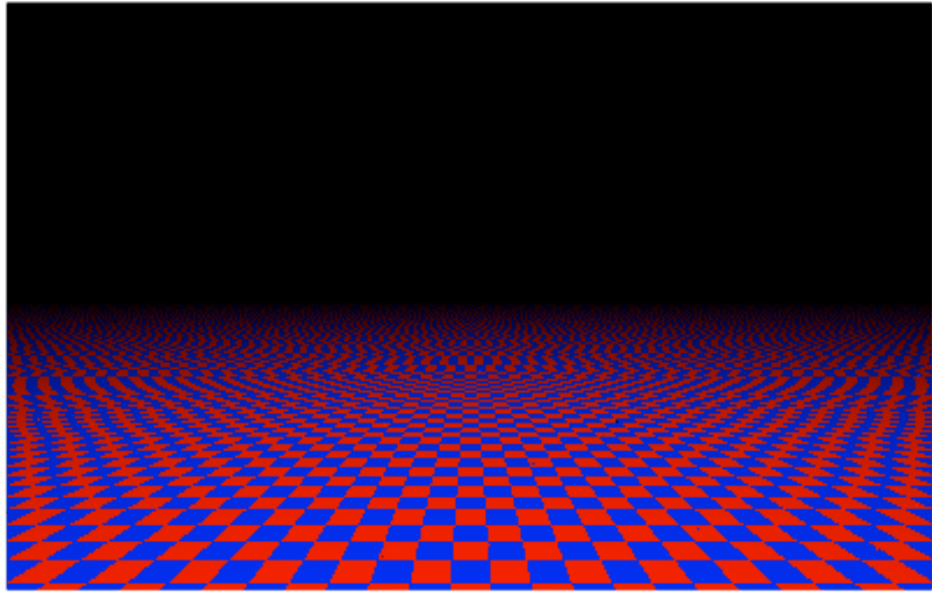
ISA/Architecture

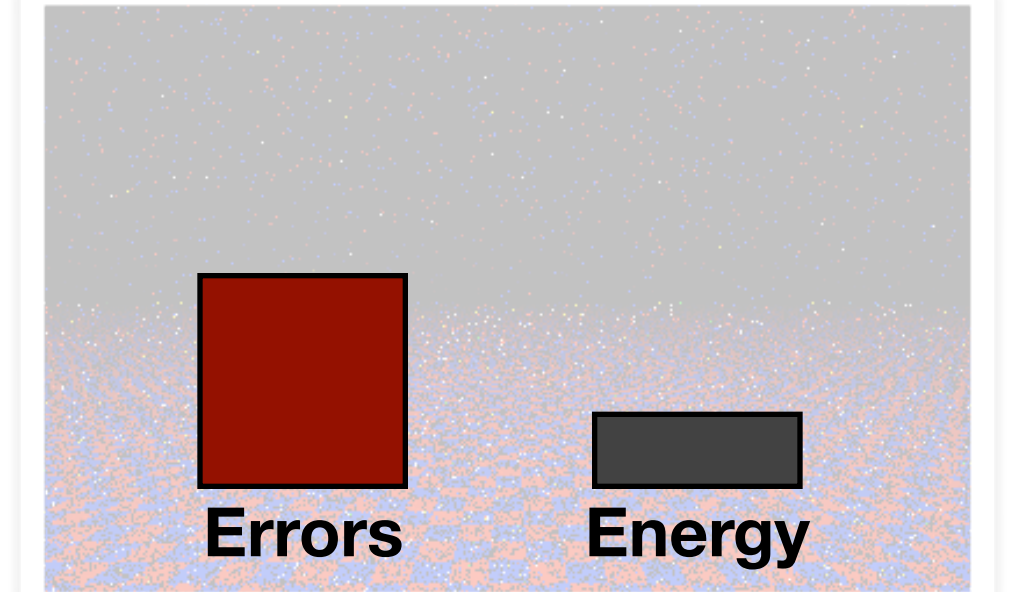
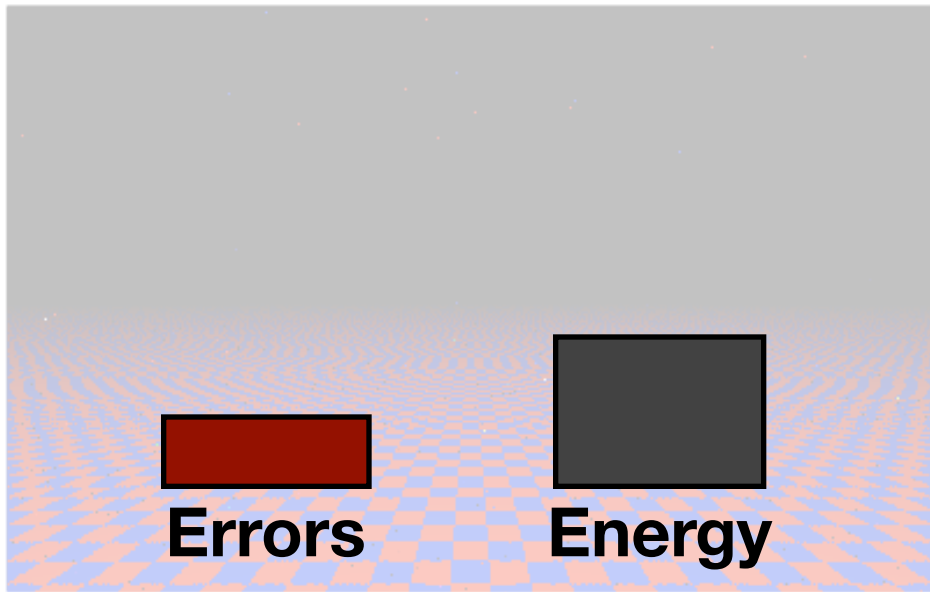
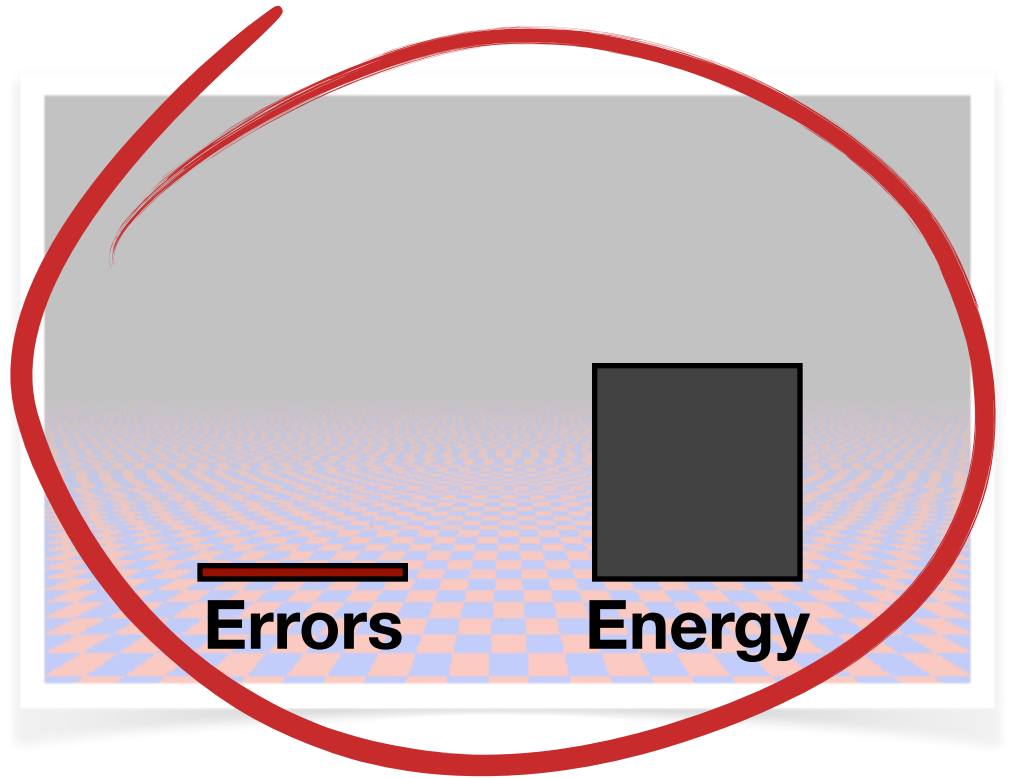
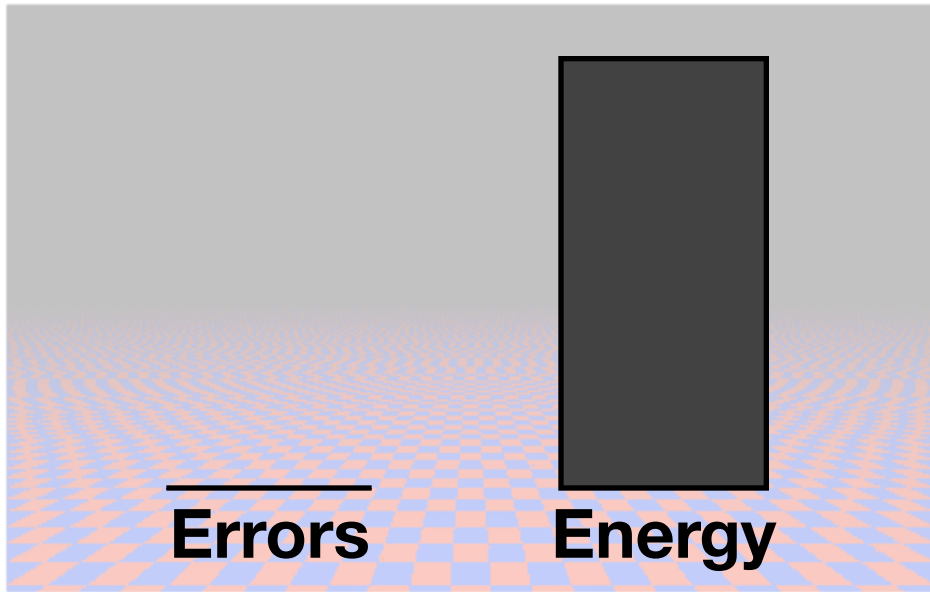
Circuits

Physics

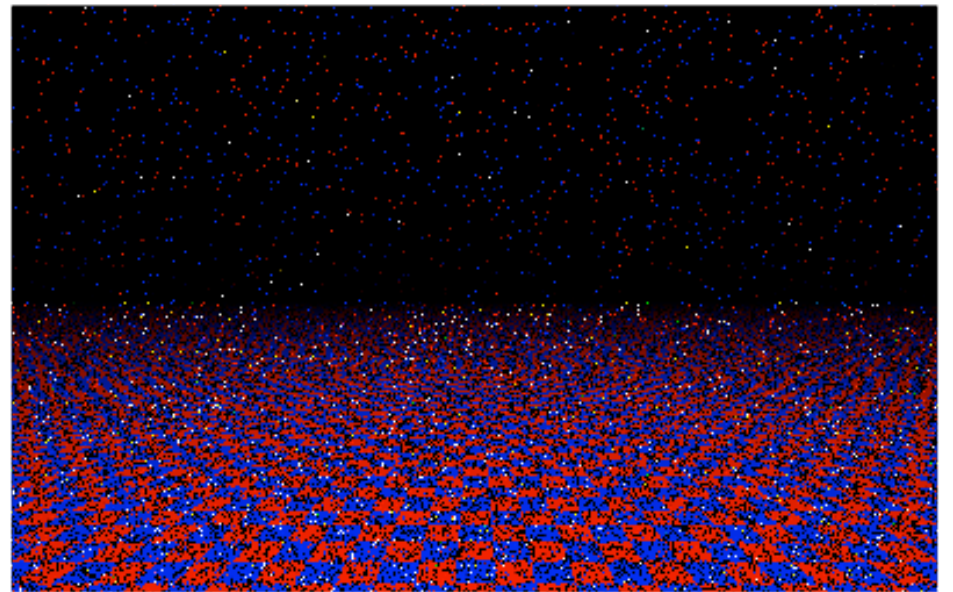
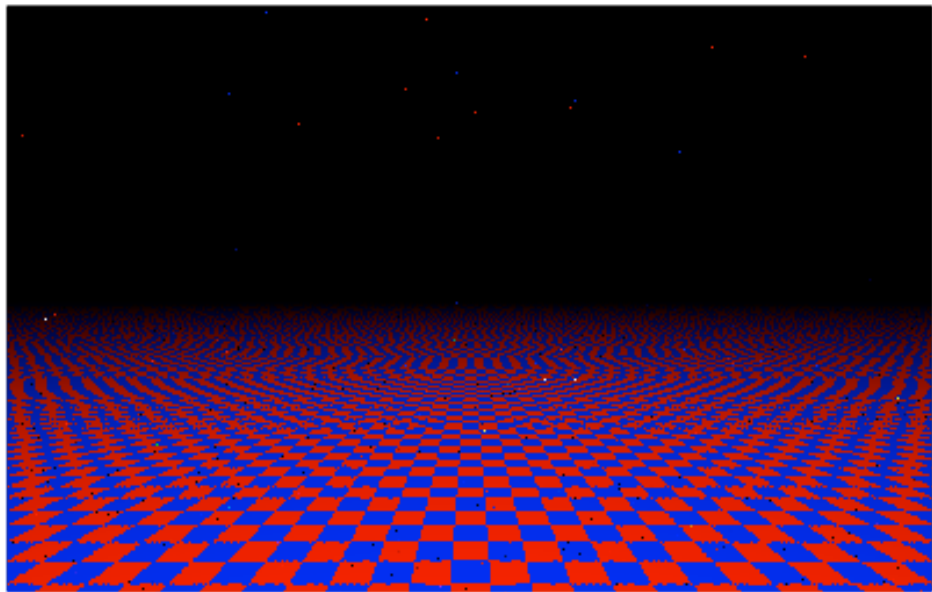
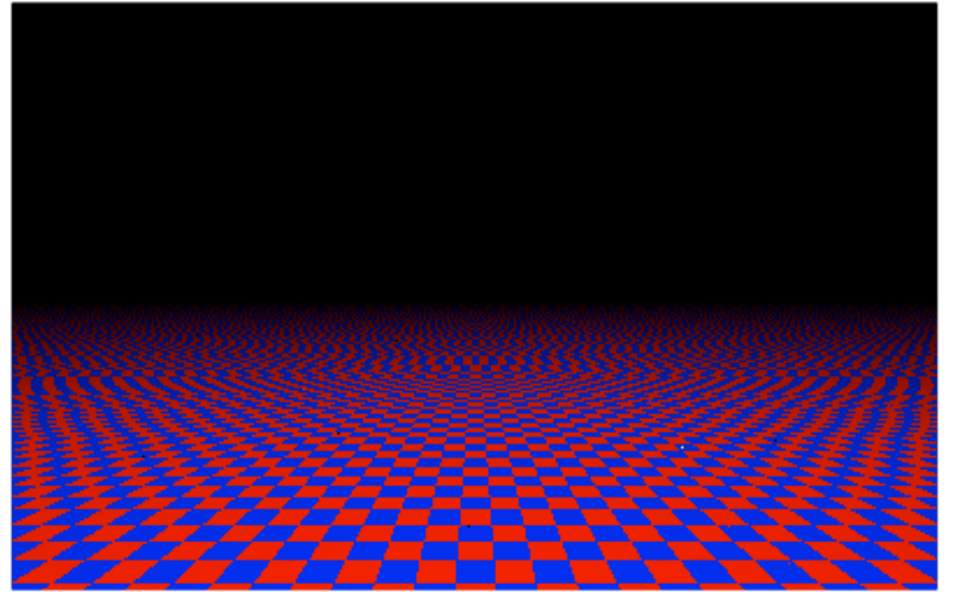
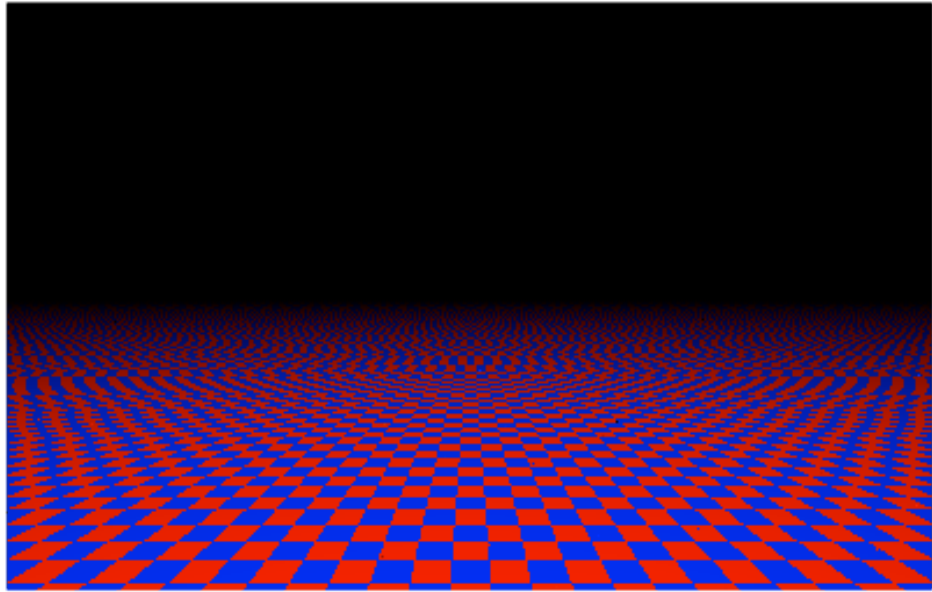
In essence, it is about **specializing computation, storage and communication** to properties of the data and the algorithm. Enables better use of underlying substrate.

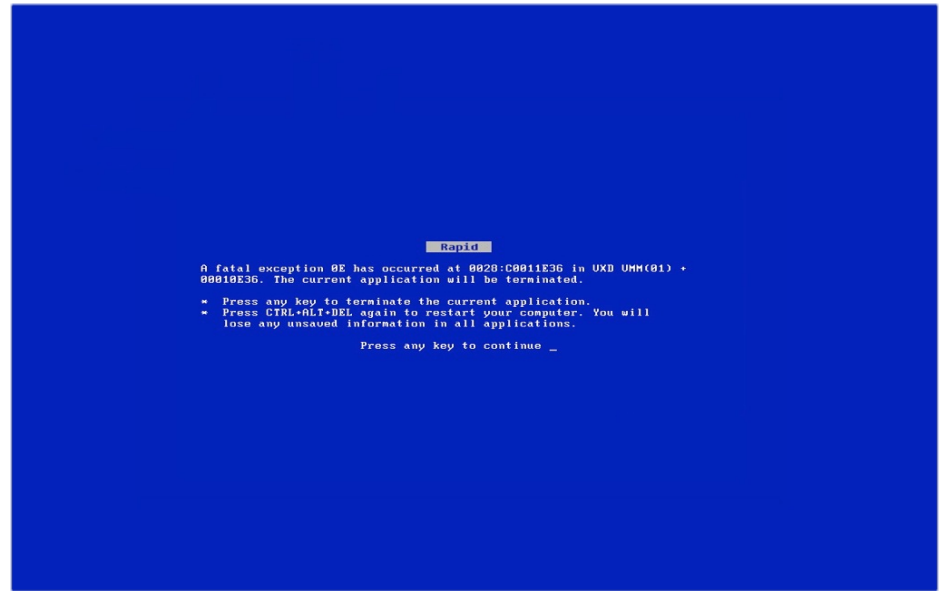
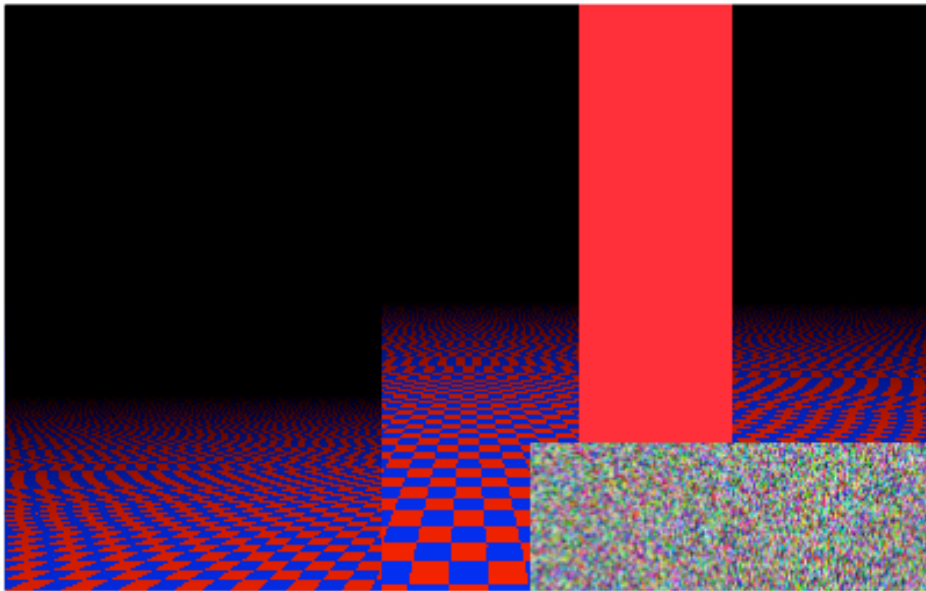
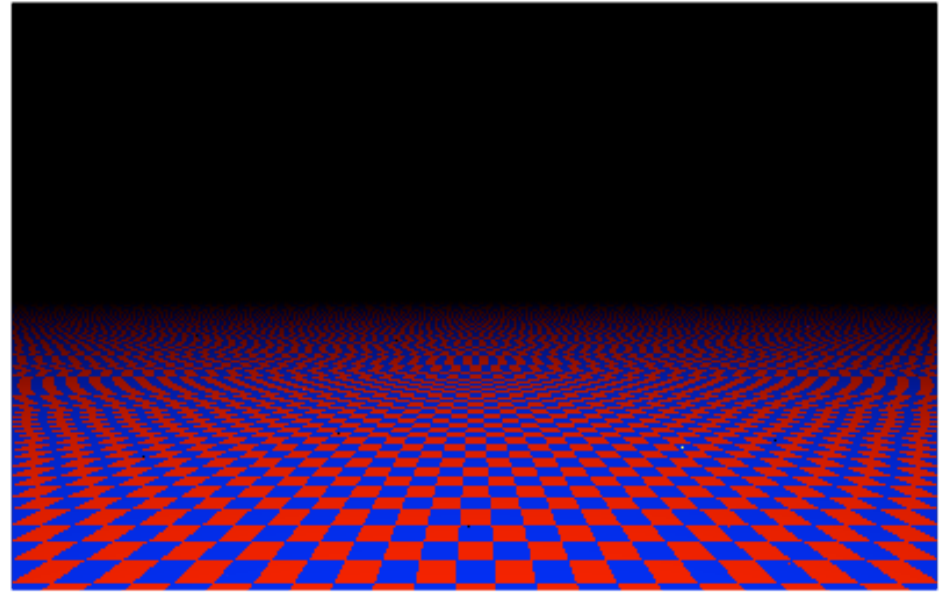






**But approximation needs to be done
carefully... or...**





Three important questions

1 What and how to approximate?

Language

2 How good is my output?

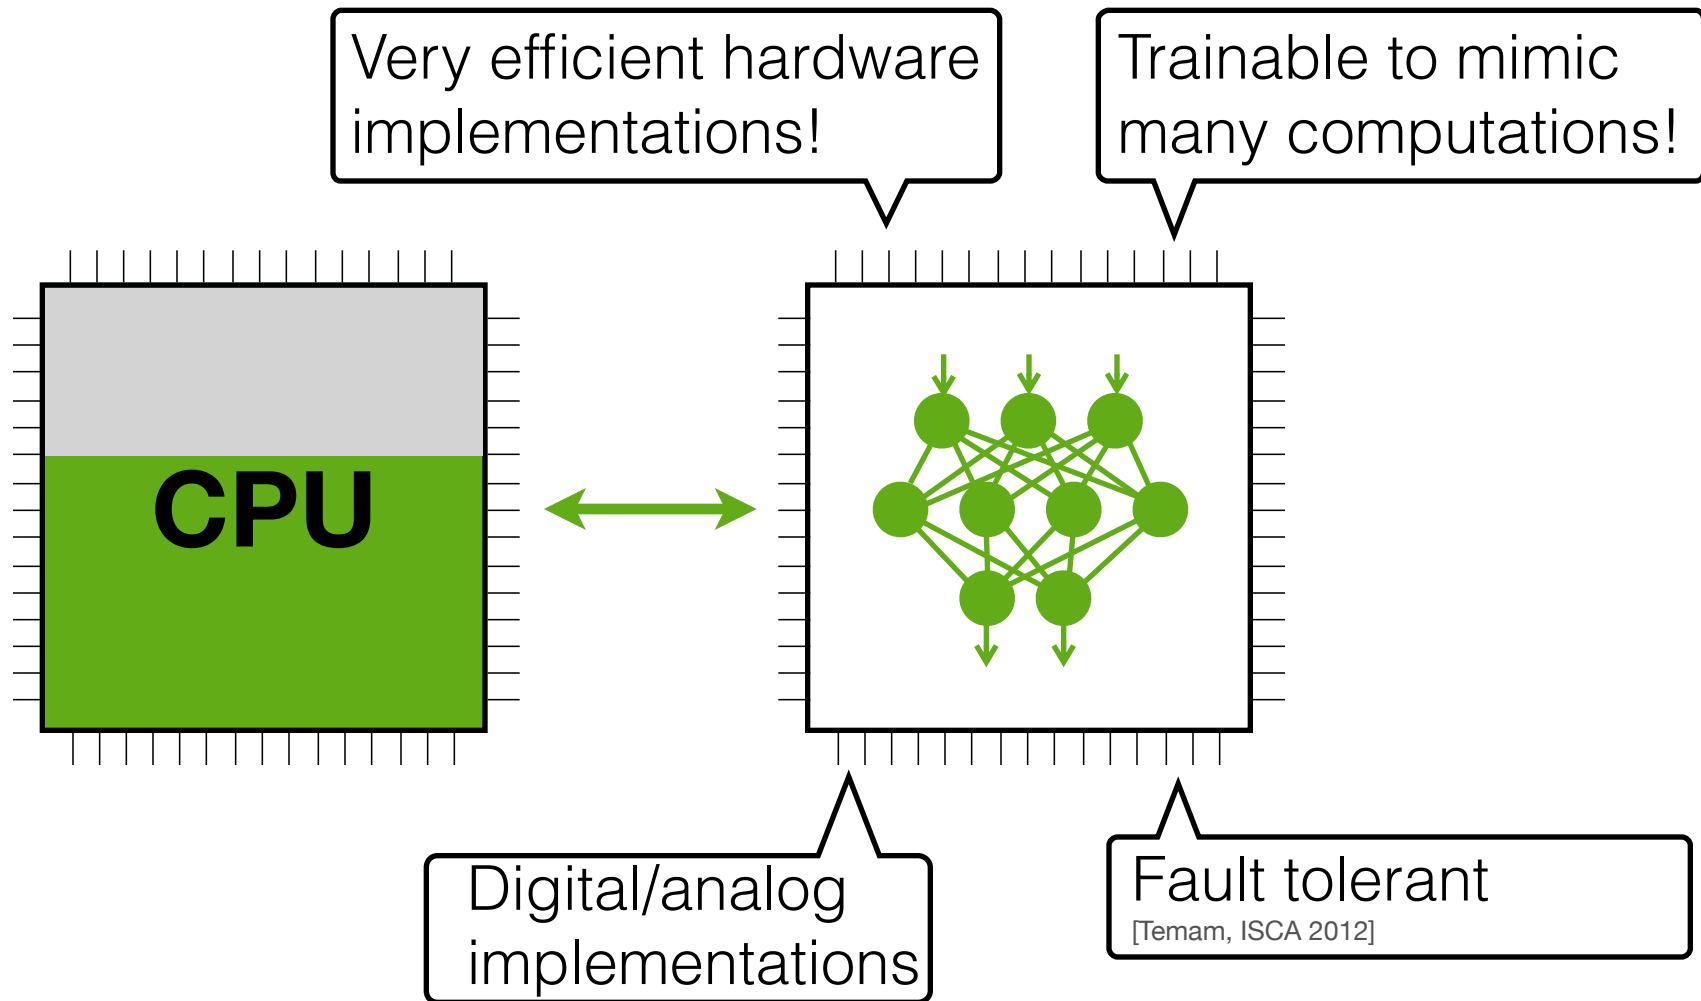
Compiler

3 How to take advantage of it?

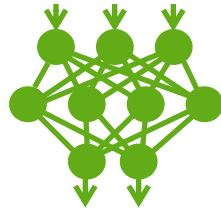
Runtime

Hardware

Example: Neural Networks as General Approximate Accelerators



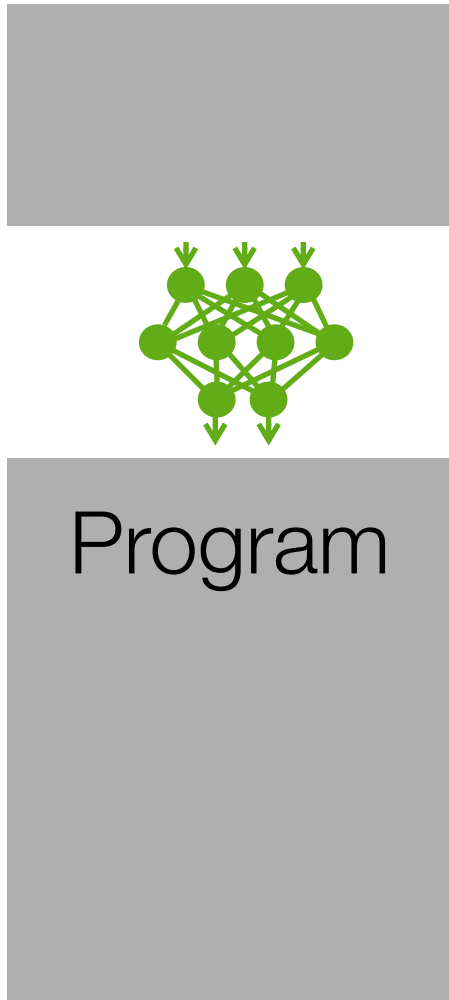
Neural acceleration



Find an approximate program component

Compile the program and train a neural network

Neural acceleration



Find an approximate program component

Compile the program and train a neural network

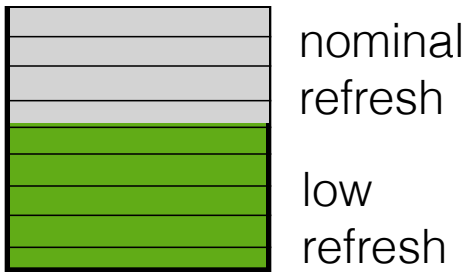
Execute on a fast Neural Processing Unit (NPU)

3-20x performance/efficiency improvement (FPGA SoC prototype).

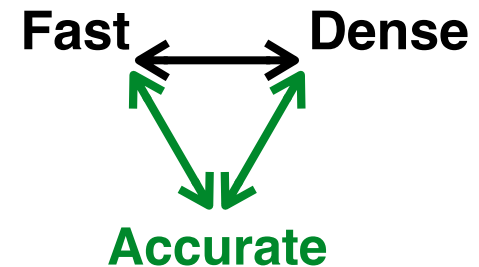
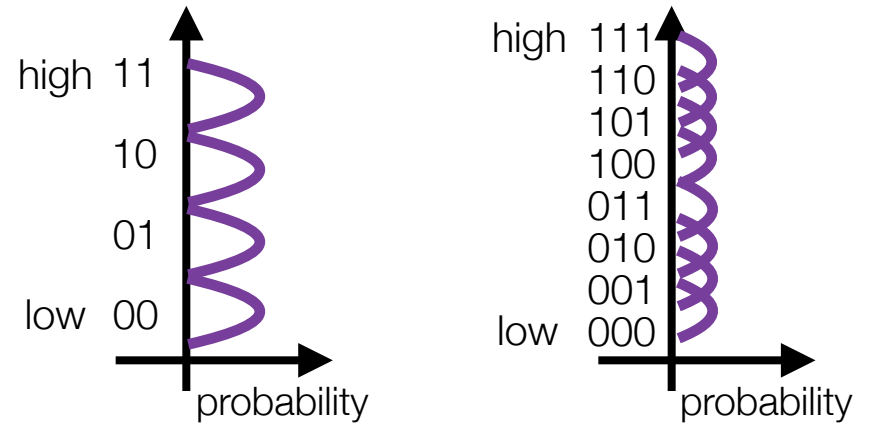
Approximation beyond the CPU

Multi-level solid state cells

DRAM



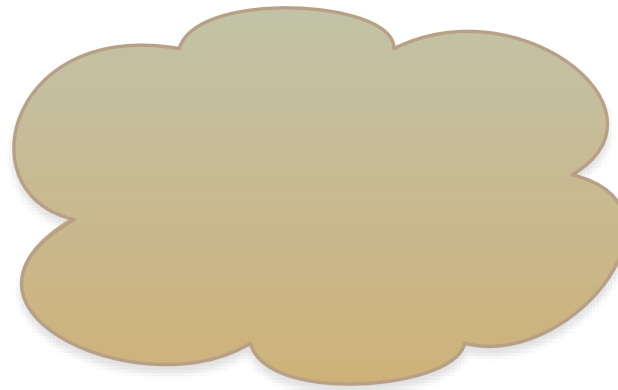
Wireless networking



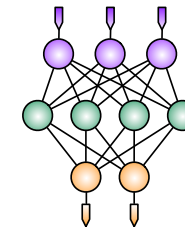
And end-to-end story



Approximate
near-sensor
processing



Approximate
communication



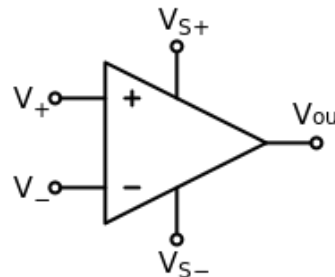
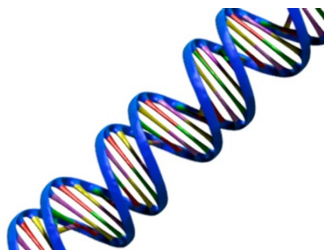
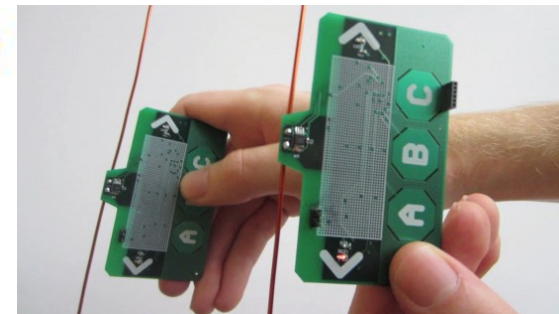
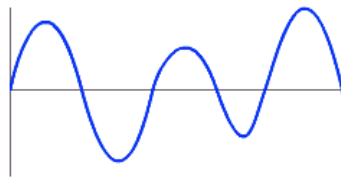
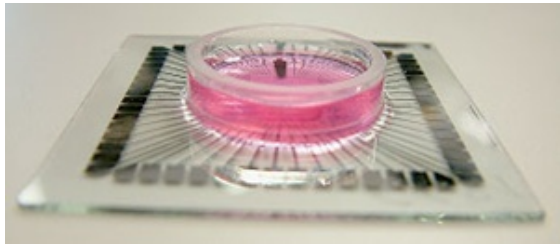
$$o_j = \text{sigmoid}\left(\sum_i w_{ji} x_{ji}\right)$$

Deep neural nets

Approximate
storage

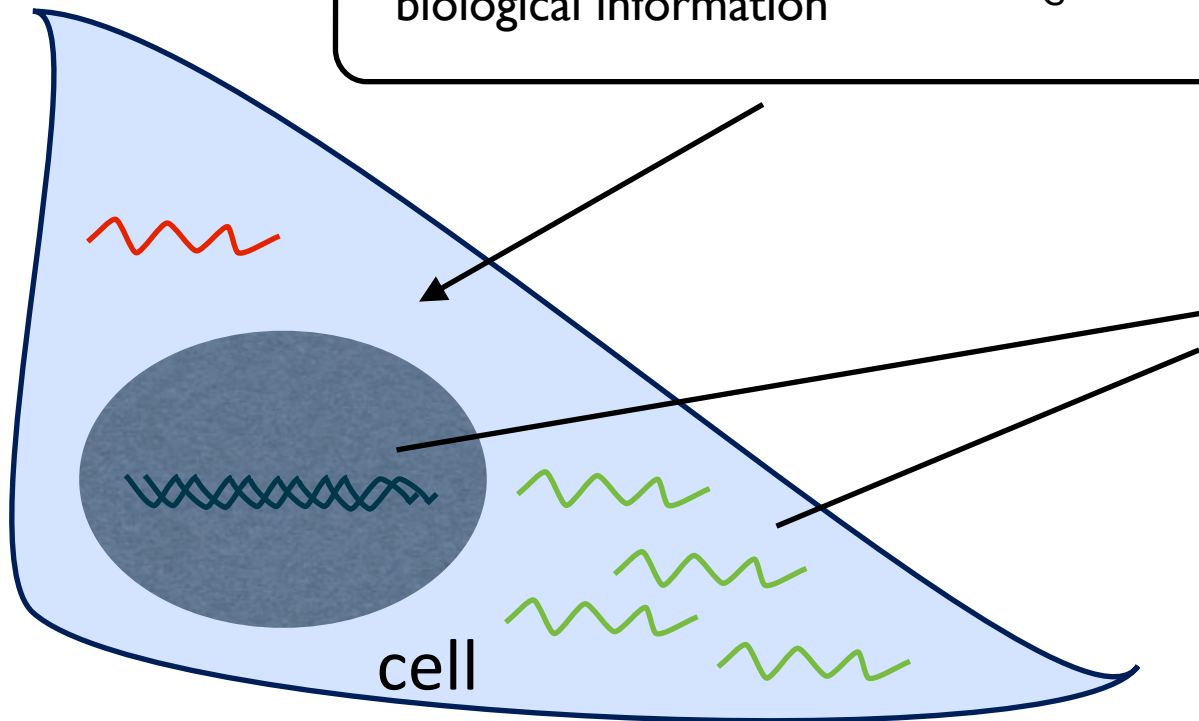
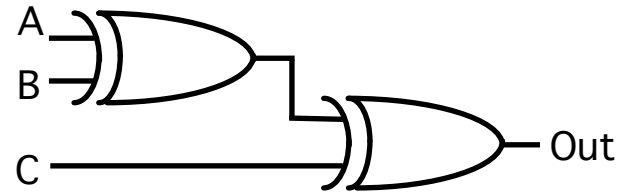
Beyond CMOS, beyond digital, beyond efficiency

Science of approximate computing



Embedded controllers for biochemical systems: “smart” therapeutics and diagnostics

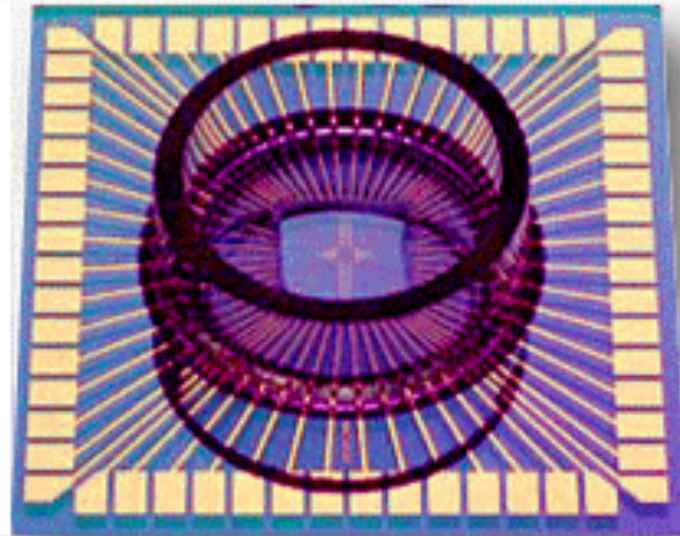
A cell-based “computer” needs to be biocompatible, and sense, analyze and act on biological information



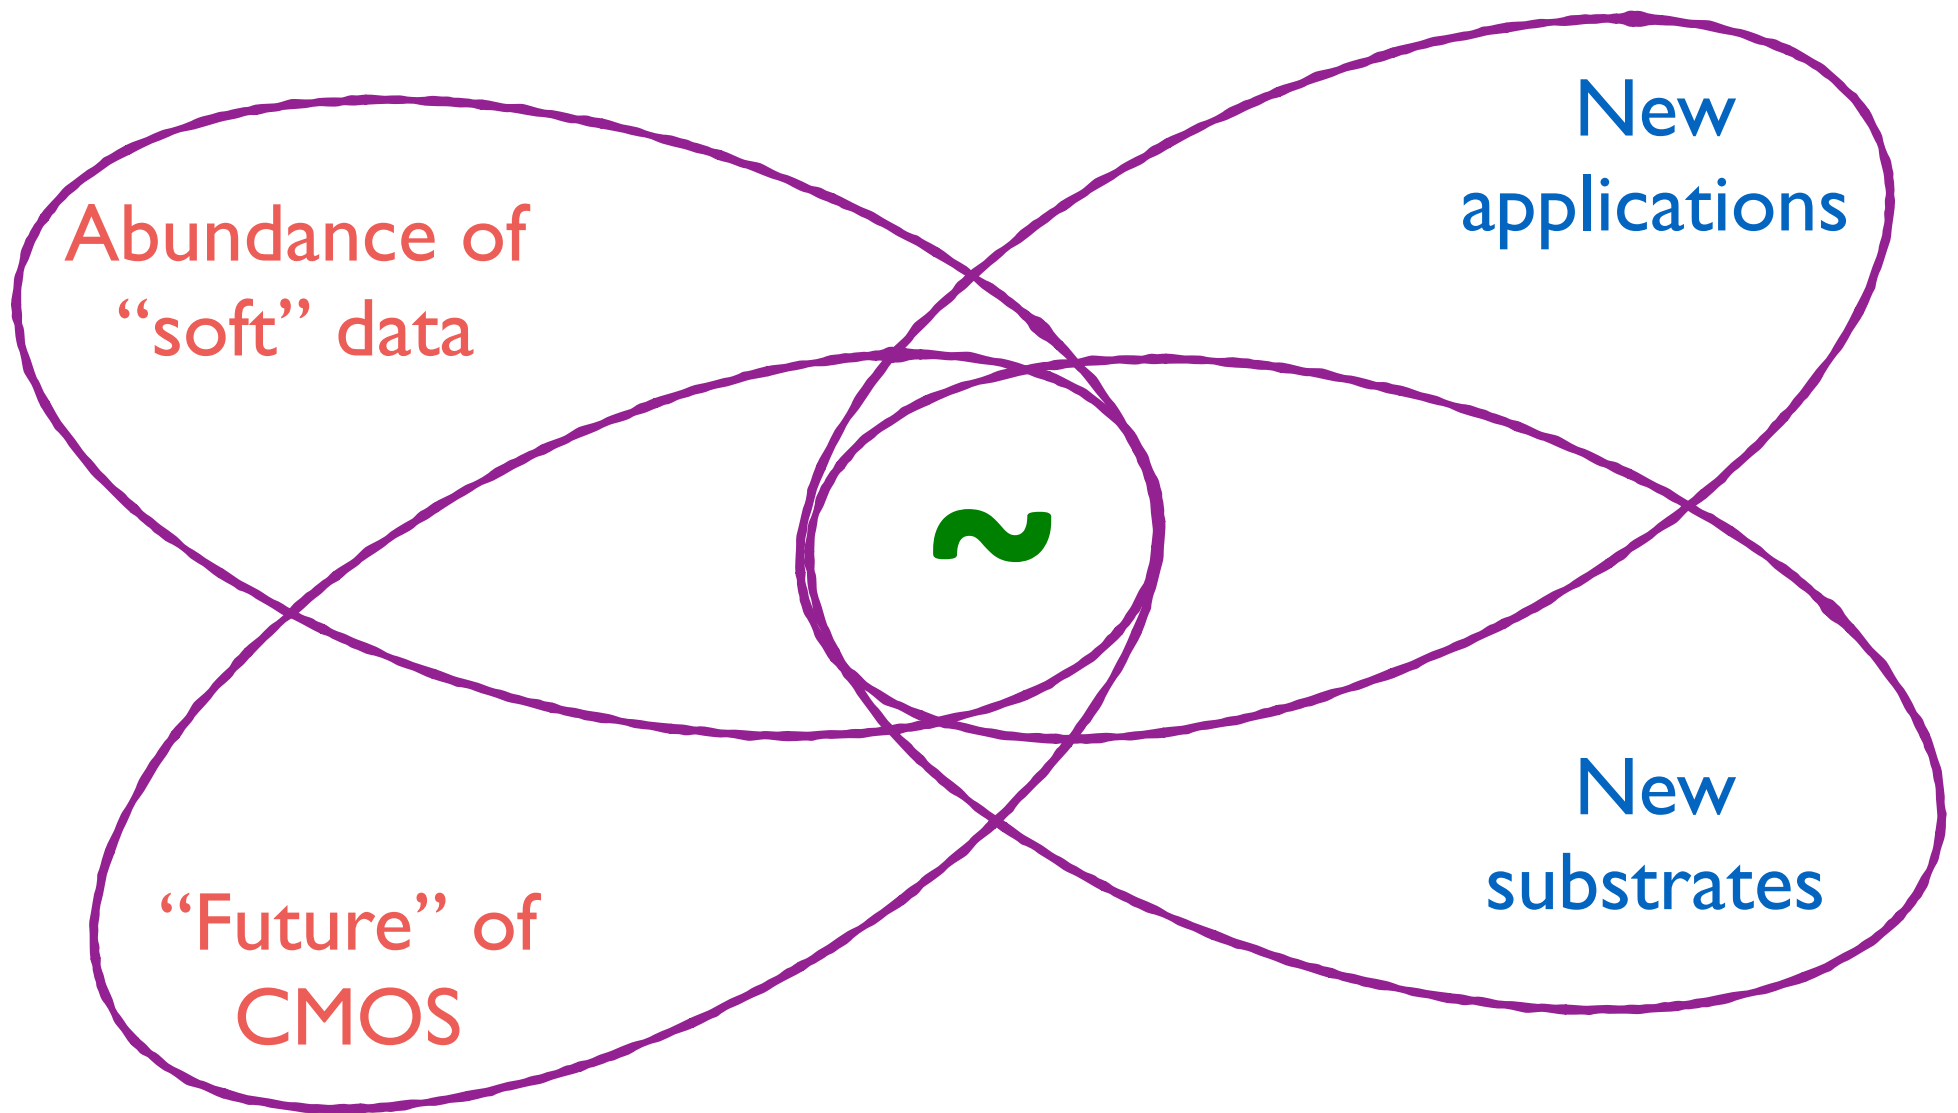
Biological Information is encoded in the sequences and amounts of biomolecules (DNA, RNA, proteins, etc.)

ugagguaguagguuguauaguu
ugagguaguagguugugugguu
ugagguaguagguuguaugguu
agagguaguagguugcauaguu
ugagguaggagguuguauaguu
ugagguaguagauuguauaguu
ugagguaguaguuuguacaguu
ugagguaguaguuuggcuguu

SIMBIOTECH: Silicon Meets Biotech



Take home: confluence of trends points to pervasive approximate computing systems

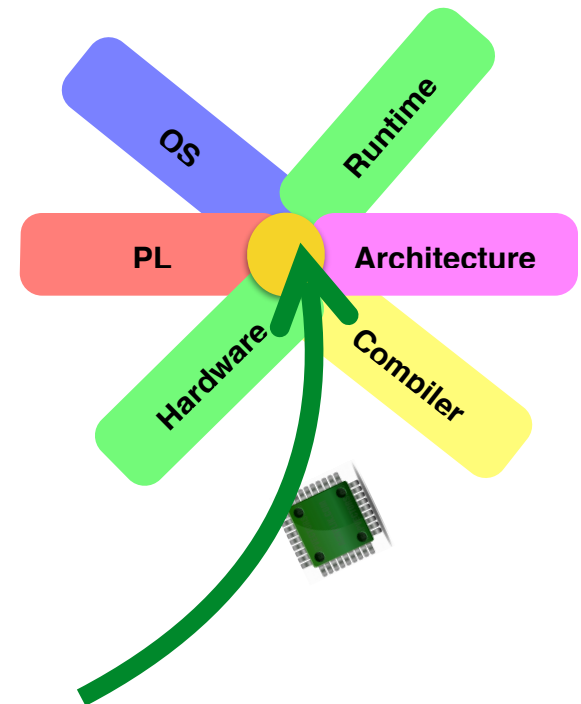
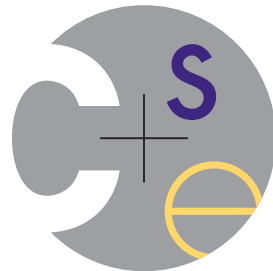


Thanks!

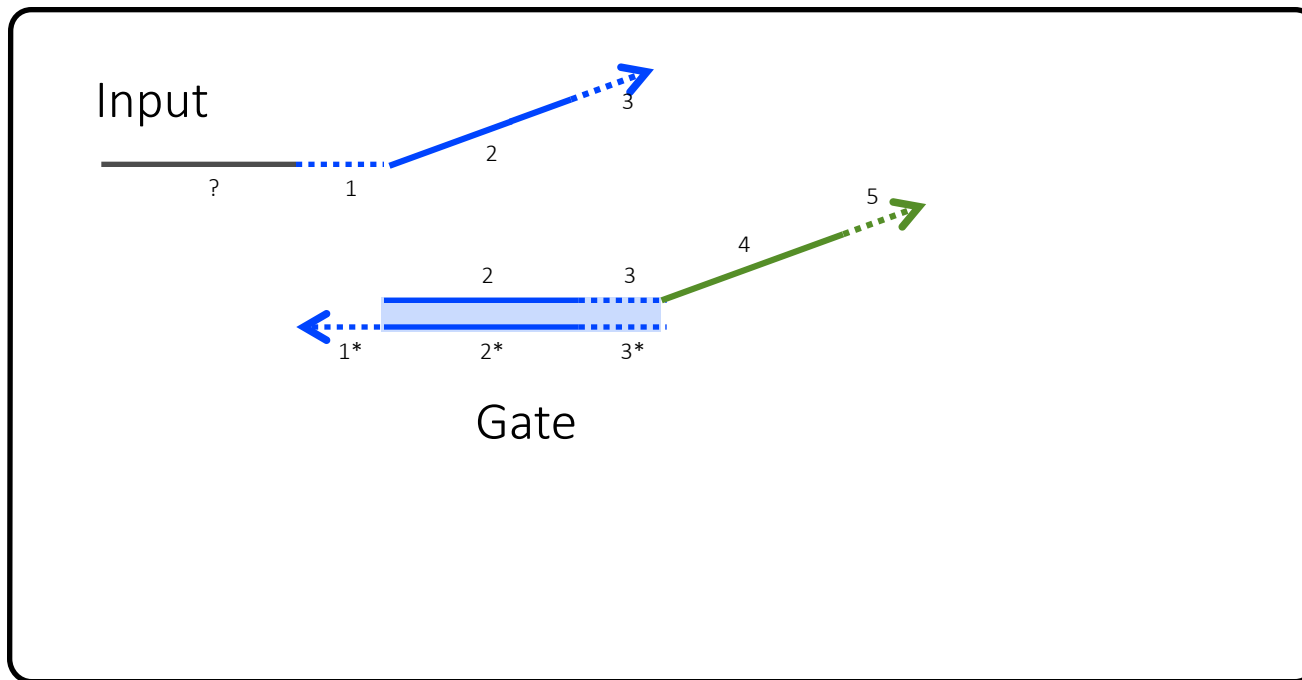
Luis Ceze

University of Washington

University of Washington

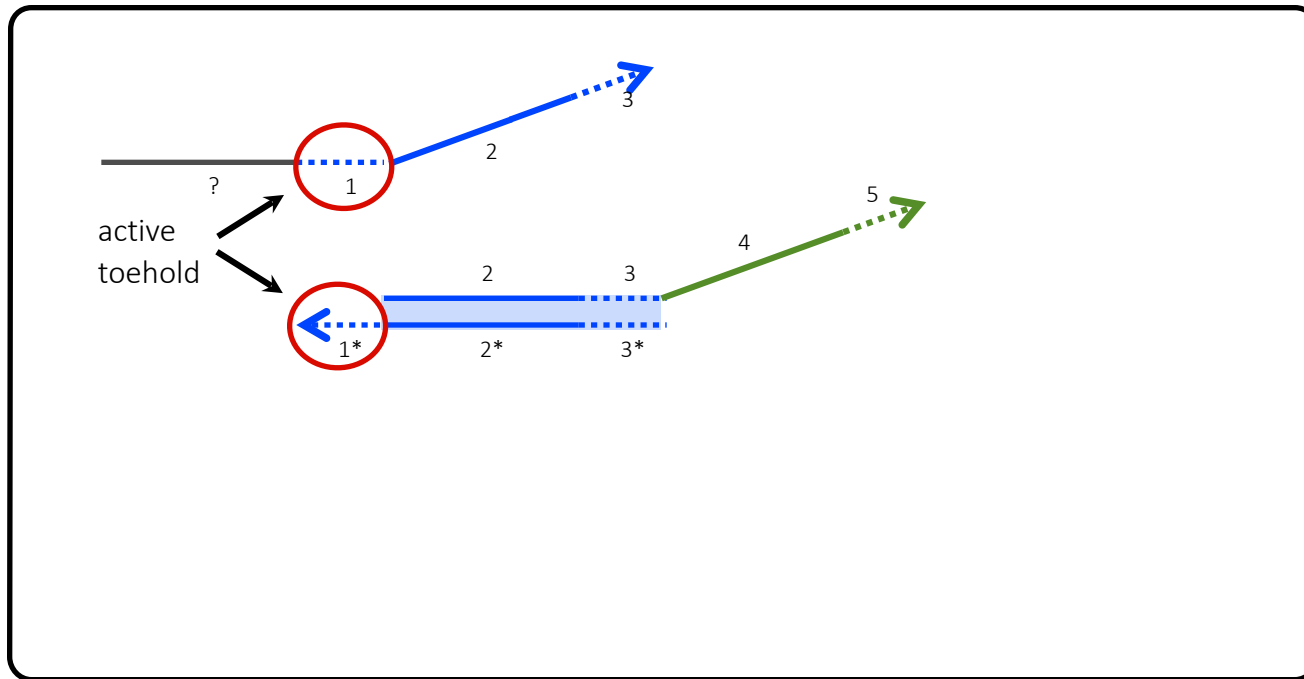


DNA strand displacement basics



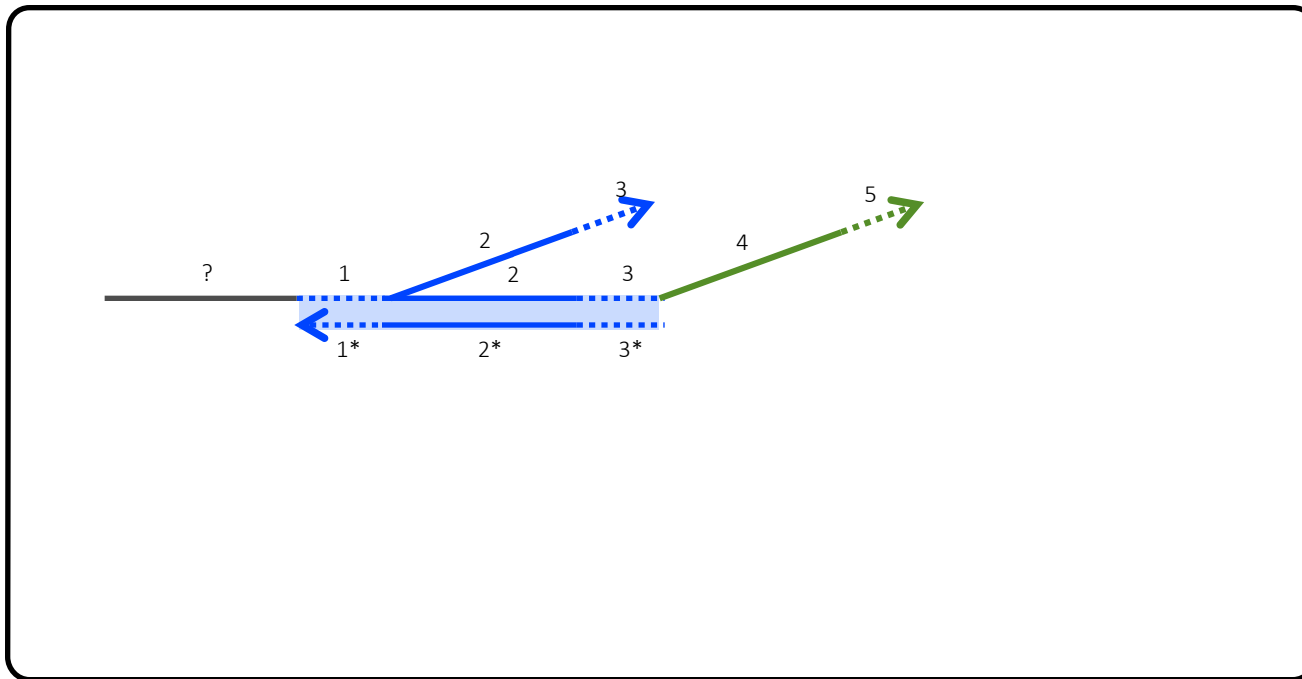
For a review see D. Y. Zhang and G. Seelig, Nature Chemistry (2011)

DNA strand displacement basics



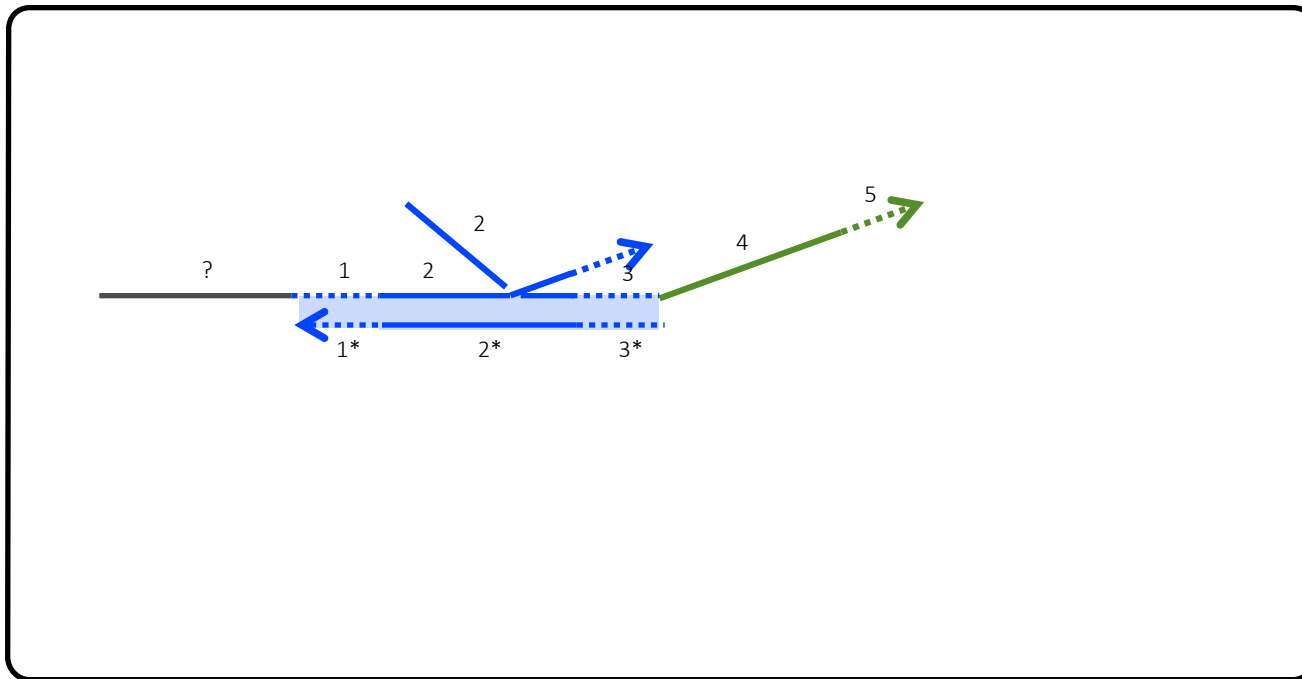
For a review see D. Y. Zhang and G. Seelig, Nature Chemistry (2011)

DNA strand displacement basics



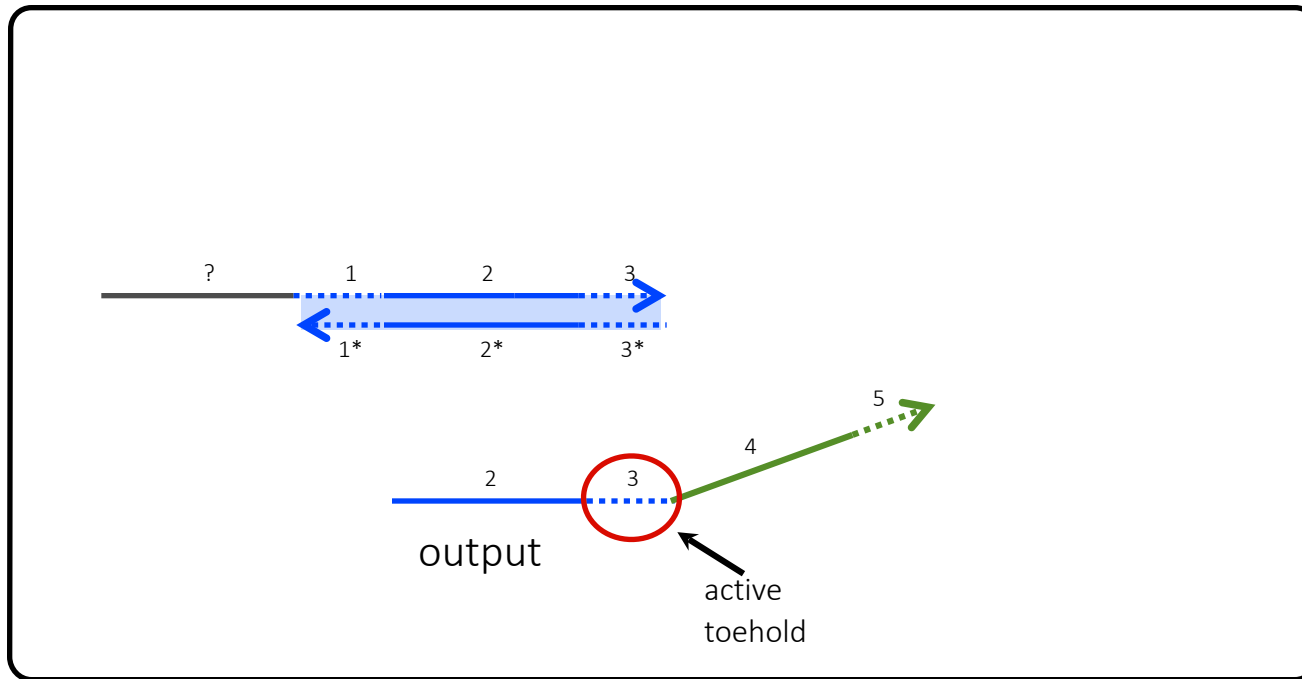
Strand displacement is initiated at the single-stranded toeholds. Toehold binding is a reversible process.

DNA strand displacement basics



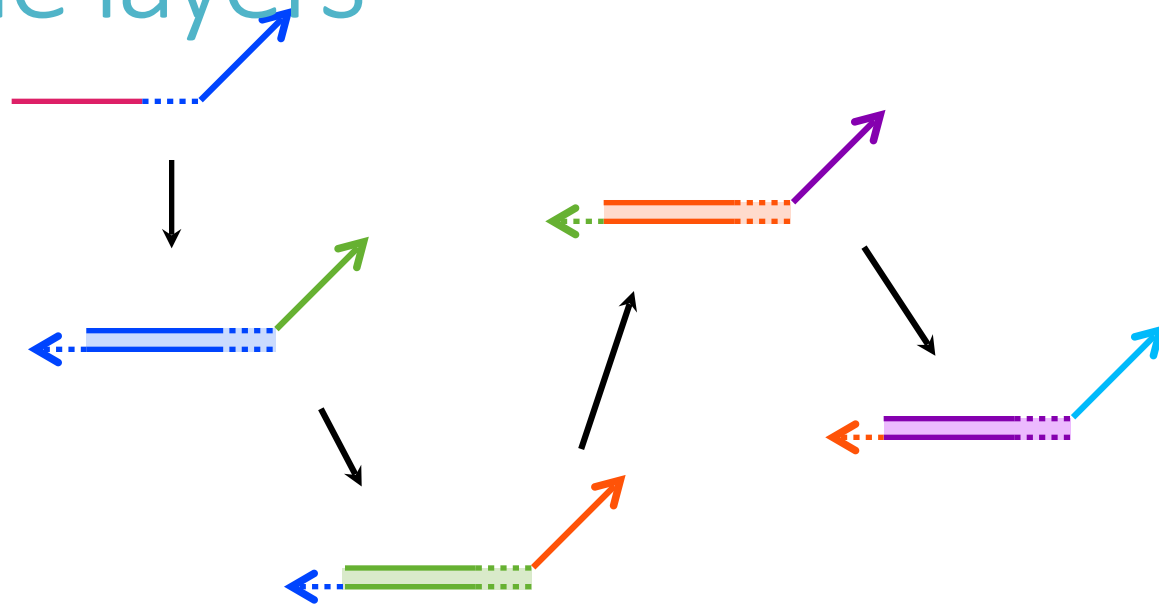
Strand displacement proceeds through a branch migration. Branch migration is a random walk.

DNA strand displacement basics



Release of the output strand is (almost) irreversible in the absence of a toehold for the reverse reaction.

Signals can propagate through multiple layers

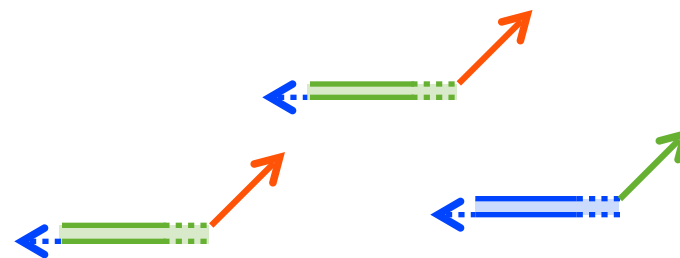


Sequence determines interactions: The chemical specificity determines the “wiring diagram.” There are $4^{20}=10^{12}$ different 20-mers but most sequences can't be used.

Take away message

We can build simple logic gates and circuits using DNA.

DNA strand displacement circuits are the largest engineered molecular circuits built so far. But they are still really small.



Classifier Generation Pipeline

