

DONDERS INSTITUTE
BRAIN (INSPIRED) COMPUTING

EXAMPLE PROBLEMS

- $n \times n$ chess
- $n \times n$ Go
- Box packing
- Map coloring
- Traveling salesman
- $n \times n$ Sudoku
- Graph isomorphism
- Factoring
- Discrete logarithm
- Graph connectivity
- Testing if a number is prime
- Matchmaking

Efficiently solved by quantum computer

Efficiently solved by classical computer

Harder

Where do I fit?

Towards Neuromorphic Complexity Analysis
Johan Kwisthout, Associate PI, Donders Institute

Radboud University Radboudumc

EXAMPLE PROBLEMS

- $n \times n$ chess
- $n \times n$ Go
- Box packing
- Map coloring
- Traveling salesman
- $n \times n$ Sudoku
- Graph isomorphism
- Factoring
- Discrete logarithm
- Graph connectivity
- Testing if a number is prime
- Matchmaking

Efficiently solved by quantum computer

Efficiently solved by classical computer

Harder

Where do I fit?

Intel Core i9 X-series

Brain-inspired computing

- In October 2015, the White House Office of Science and Technology Policy released the following Grand Challenge for Future Computing:
- “Create a new type of computer that can proactively interpret and learn from data, solve unfamiliar problems using what it has learned, and operate with the energy efficiency of the human brain”*

Neural-inspired systems for **non-von Neumann** computational architectures

Continuous data transfer between memory and CPU

Heat generation bottleneck!

Neuromorphic architectures

What is a neuromorphic architecture?

- Principles most people agree on
 - Neurons and Synapses as basic components
 - Co-located memory and computation
 - Potentially energy efficient (“orders of magnitude”)
- Different schools of thought
 - Analog hardware based (“memristors”)
 - Spike-based digital systems (“spiking neural networks”)

BrainScaleS SpiNNaker Loihi

Neuromorphic architectures

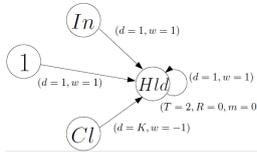
What are these neuromorphic architecture good for?

- Probably well suited for...
 - Event-driven sensors / actuators (neuromorphic robotics)
 - Energy-critical applications that allow for less precision
- Probably not the best architecture for...
 - “Deep” classification and pattern recognition (outperformed by convolutional DNNs)
 - Applications that value precision over energy usage
- Is “machine learning” the only possible application?

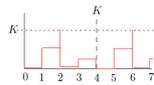
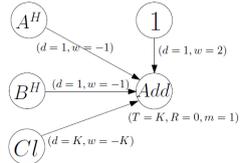
Remember the GPU...

AlphaGo

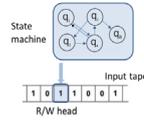
Example problem: adding two numbers mod K



- Circuit ensuring that once a neuron fires, it fires until reset by the clock
- Adder circuit spikes on carry-out, neuron potential at clock intervals encodes $A+B \text{ mod } K$

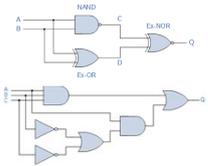


Beyond Turing



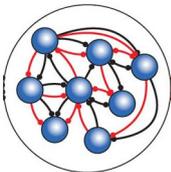
- **Turing Machine M_L**
- Input I encoded (in binary) on the tape
- State machine M_L implements algorithm
- Formally: recognizes languages $L \subset \{0,1\}^*$
- Canonical question: Does M_L accept $I \in L$ using resources (time/space) at most R ?

- **Family of Boolean Circuits $C_{L,||}$**
- Input I encoded as special input gates
- Circuit (different circuit per input size $||$) implements algorithm
- Formally: recognizes languages $L \subset \{0,1\}^*$
- Canonical question: Does, for every I , the corresponding circuit $C_{L,||}$ accept $I \in L$ using resources (time/space) at most R ?



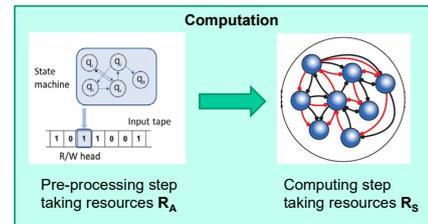
Beyond Turing

- In SNNs, input I and algorithm A are **co-located!**
- We take the circuit idea *to the extreme...*



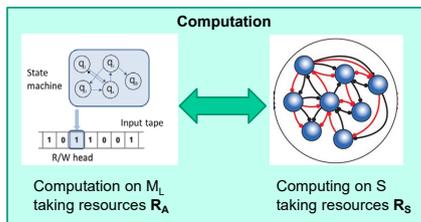
- **Collection of SNNs $S_{L,I}$**
- One network for every input I (or set of inputs $\{I\}$)
- Input and algorithm operating on it are encoded in the network structure
- Formally: recognizes languages $L \subset \{0,1\}^*$
- Accept / reject by special neurons firing
- Canonical question: Is there a resource-bounded Turing machine M_L that, given I , generates $S_{L,I}$ which decides I using resources at most R_S ?

Beyond Turing: preprocessing + computation



- **Hierarchy** of complexity classes defined by choices for R_A (time, space) and R_S (time, space, energy)
- E.g., R_A (poly time, log space), R_A (poly time, space, energy)

Computing with neuromorphic oracle



- More powerful alternative: use S as co-processor
- Formally: S is an *oracle* for Turing machine M_L

Some first theoretical results

- We have defined canonical complete problems and resource-preserving reductions between problems

SNN-HALTING

Instance: Integers t, e in unary notation; encoding $S = (N, S)$ of a spiking neural network constructable by a Turing machine $\mathcal{M}(R_A)$.

Question: Does N_{acc} fire before time step t using energy at most e ?

ORACLE SNN-HALTING

Instance: Integers t, s in unary notation; encoding of $\mathcal{M}(R_A)^{S(R_S)}$; input I .

Question: Does $\mathcal{M}(R_A)^{S(R_S)}$ accept I before time step t , accessing at most s cells on its tape?

- These problems is complete for resource bounded SNNs, resp. neuromorphic oracle machines, and can be used as starting point for reductions to actual (real-world) problems

Some first theoretical results

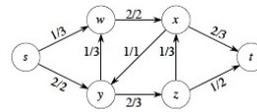
- **Neuromorphic oracles cannot do magic...**

$$P^{\text{SNN}}(\mathcal{O}(n^c), \mathcal{O}(n^c), \mathcal{O}(n^c)) = P$$

- This implies that there is no super-polynomial speedup in using a neuromorphic co-processor (which is similar as for using GPUs)
- But then again, speed is not the major consideration when we are considering neuromorphic hardware
- But **energy** is!

Some first theoretical results

Max network flow problem



This problem is P-complete, meaning that it cannot be efficiently parallelized (use only logarithmic space) on traditional machines!

MAX NETWORK FLOW
is in $L^{\text{SNN}}(\mathcal{O}(n), \mathcal{O}(n), \mathcal{O}(n))$

It can be solved in Logspace when allowed to use a neuromorphic co-processor!



MSc project
Abdullahi Ali

MAX Network Flow

- Max Network Flow is P-complete, implying it is an inherently serial problem resisting effective parallelization – takes more than $\log(n)$ space
- We can *not* solve it efficiently on a ‘stand-alone’ neuromorphic computer (working on formal proof)
- Idea: use the power of the neuromorphic oracle
 - Offload those computations that can be done in parallel
 - Specifically: finding the shortest augmenting path between s and t using “first-past-the-post” spike timing
 - We can show that the CPU then needs only $\log(n)$ space and the oracle (for dense networks) $\log(n)$ spikes

Neuromorphic complexity in theory and practice

- Theoretical research financially supported by Intel’s Neuromorphic Research Community
- Goal of INRC is to build a research community around their new Loihi chip
- We try to implement our algorithm not only in (abstract) simulator but in actual hardware
- Show the relation (or mismatch...) between formal theory and practice!



Neuromorphic complexity in theory and practice

- Theoretical model: **oracle state / oracle tape** for communication between TM and SNN
- Machine model allows for definition of **readout neurons** whose state can be written to the tape
- In reality, this doesn’t work this way: you just “upload” the network and “download” the entire SNN state → computationally costly!
- We have not yet taken **constrained** communication resources into consideration

Conclusion

- New theoretical framework
 - 1) Formal notion of “computation” in neuromorphic architectures
 - 2) Complexity classes based on resource constraints
 - 3) Hardness criteria and a means to *translate* problems into each other while keeping resources invariant
 - 4) Algorithms to show that a problem is in a specific class
- Iteration between formal model and actual architecture to ensure vital aspects are captured
- My research group is currently working on GUI to develop and ‘drag-and-drop’ design patterns (circuits) to construct algorithms