Efficient Digital Neurons for Large Scale Cortical Architectures

James E. Smith

2014 ACM/IEEE 41st International Symposium on Computer Architecture (ISCA)

Presented by Rasmus Lüscher

Executive Summary

Motivation: Large scale architectures are needed to emulate the neocortex to support research studying the operation of the brain.

Problem: Existing models of “complex” two stage neurons are **more accurate** but **less efficient** than “simple” one stage neuron models.

Goal: Provide implementations of **efficient** digital neurons that could be used in development of future large-scale cortical architectures.

Key Contributions:

- Four neurons are implemented in a manner that allows side-by-side comparison.
- Proposed a new **two stage neuron** that is **biologically accurate** and is almost as **efficient** as a one stage neuron.
- Method to compare implementations

Outline

Background

Problems & Goals

Digital Neuron Models

Comparison

Efficient Two Stage Neuron

Key Takeaways

Strengths & Weaknesses

Thoughts & Ideas

Open Discussion

Background

- **Background**
  - Introduction
  - Biological Neuron Behavior
  - Hodgkin Huxley Model
- Problems & Goals
- Digital Neuron Models
- Comparison
- Efficient Two Stage Neuron
- Key Takeaways
- Strengths & Weaknesses
- Thoughts & Ideas
- Open Discussion

Introduction

- Researchers want to understand how brains work.
  - Large scale experimentation required to discover its computational paradigm.
  - Once paradigm is revealed there will be demand for practical implementations of new type of computer.

- Neuron based computing might be able to exploit some of the brain's advantages.
  - Massive parallelism
  - Very energy efficient

*The neocortex is the sensing, thinking, perceiving part of the brain.

Biological Neurons

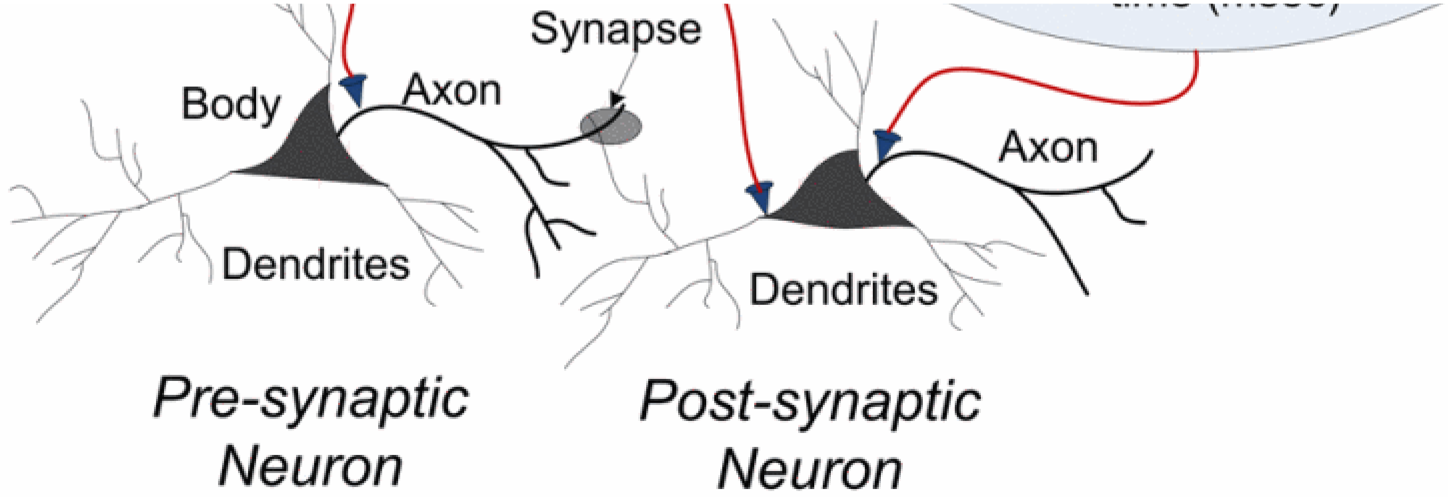

Hodgkin Huxley Model

- Mathematical Model
- Earned Nobel Prize in medicine 1963
- Characterizes Membrane Potential (voltage inside a neuron)
- Ancestor to most neuron models today (biological and digital)

Excitatory synapses produce positive PSPs -> **increasing** total Membrane Potential

Inhibitory synapses produce negative PSPs -> **decreasing** total Membrane Potential

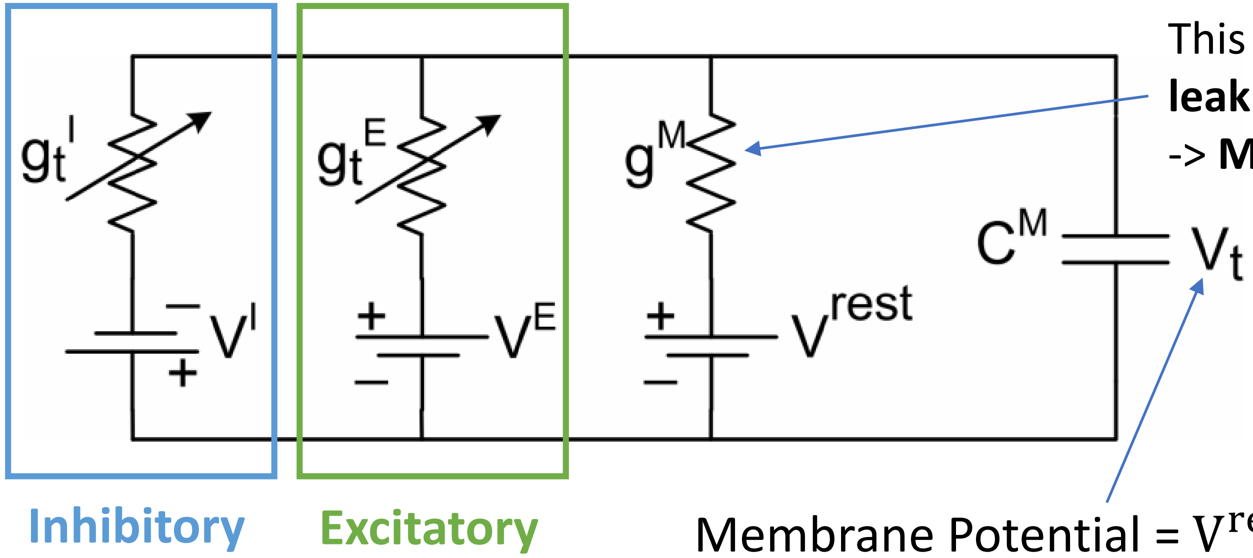

This capacitor models ions steadily **leaking** out of the neuron body -> **Membrane Potential decay**

Membrane Potential = V^{rest} + $\sum$ Excitatory - $\sum$ Inhibitory

Problems & Goals

Background

Problems & Goals

Digital Neuron Models

Comparison

Efficient Two Stage Neuron

Key Takeaways

Strengths & Weaknesses

Thoughts & Ideas

Open Discussion

- Problems
- Goals

Problems

- Energy efficiency is very important for large computational systems
  - Efficient neurons -> energy can be used somewhere else in the system
- Different neuron models to choose from
  - Depending on assumptions complex models might be needed
  - Existing **complex** models are more accurate but about 10 times less efficient than simple models
  - One might need to make a **compromise**

Goals

- Help future research by providing efficient implementations for various Digital Neuron Models
- Provide method to compare and evaluate implementations
- Propose a more efficient complex digital neuron model
  - Help avoid needing to make compromise

Digital Neuron Models

- Background
- Problems & Goals
- **Digital Neuron Models**
  - Models in this paper
  - Two Stage vs. One Stage
  - Base Model (LIF)
  - Implementations
- Comparison
- Efficient Two Stage Neuron
- Key Takeaways
- Strengths & Weaknesses
- Thoughts & Ideas
- Open Discussion

Models presented in this paper

- Four different existing models covered in this paper
  - Two **complex** models
    - incorporate both membrane potential decay and synapse conductance decay
  - Two **simple** models
    - only have membrane potential decay

- The models are described in order of complexity (complex -> simple)

Two Stage vs. One Stage Models

Two Stage Models

- Complex
- Two stages of exponential decay
  - Synapse Conductance
  - Membrane Potential
- PSP is like in biological neuron

One Stage Models

- Simple
- One stage of exponential decay
  - Only Membrane Potential
- PSP gets more abstract
  - Infinitely steep slope of leading edge
  - Results in different timing behavior

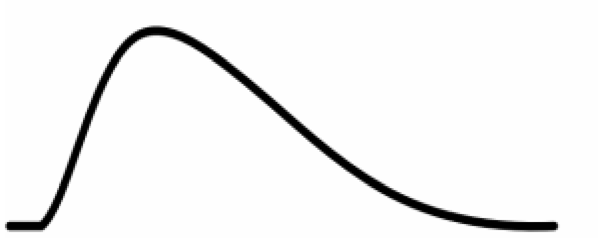

Two stage models (DLIF and DSRM0)

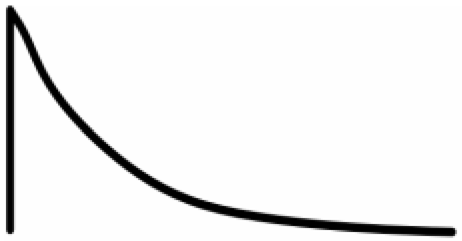

One stage model (SLIF)

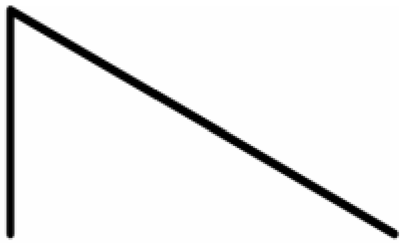

One stage model (LLIF)

PSP = **p**ost **s**ynaptic **p**otential resulting from a single input spike

Leaky Integrate and Fire (LIF) Model

- Base model for the models in this paper

- Incorporates a membrane potential that decays exponentially with some time constant

DLIF

- **Two stage** model
- A LIF implementation that also implements **decaying** synapse potential (synapse adds potential over some time after activation)
- Most complex model in the context of this paper

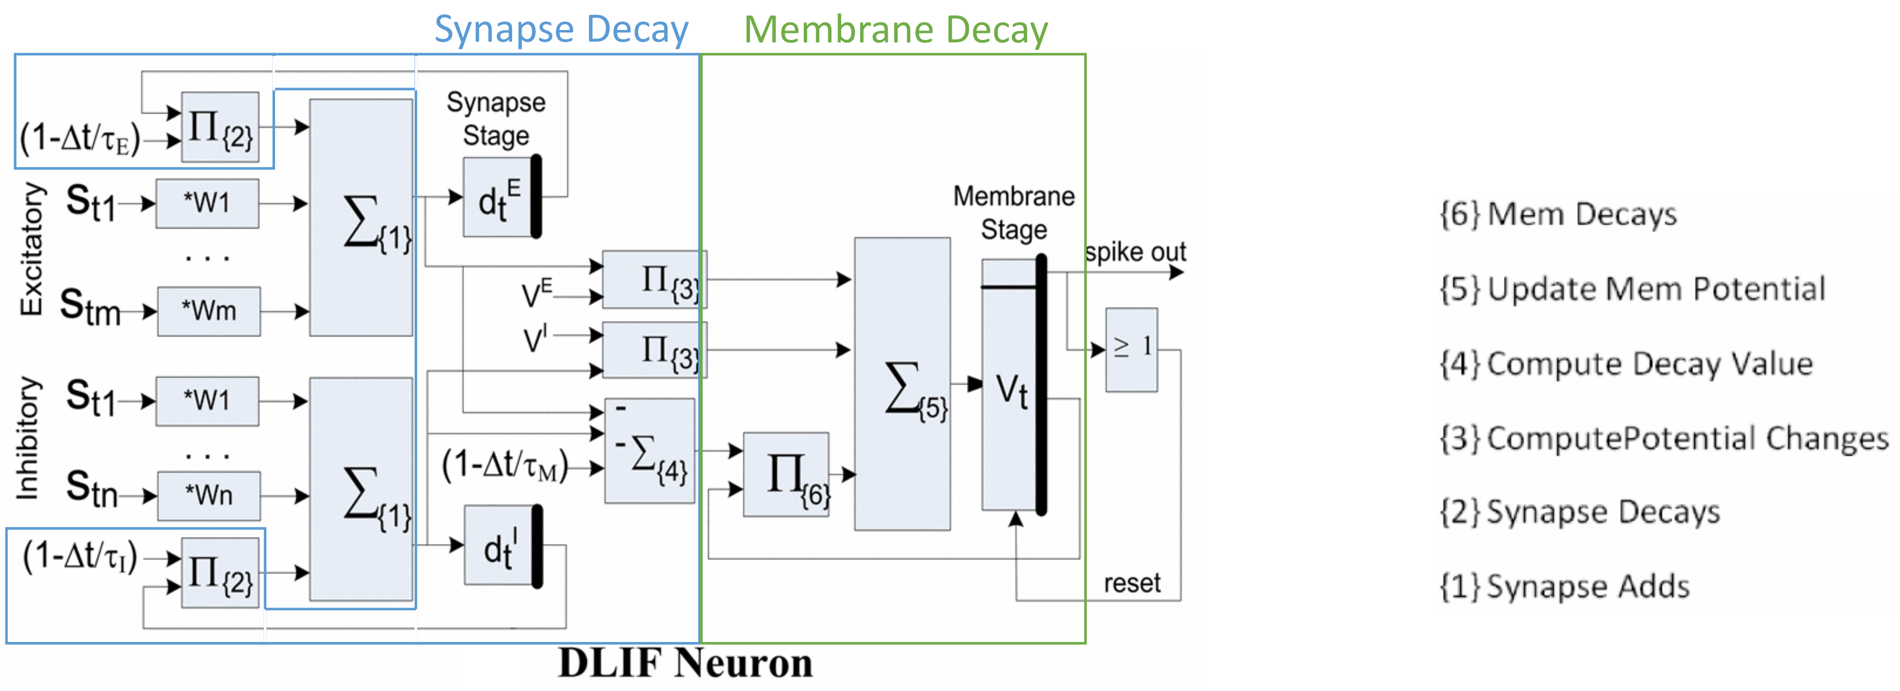

DSRM0

- **Two stage** model
- Like DLIF but spike responses are assumed to be independent (SRM0)
  - Less “unnecessary” operations -> **more efficient**
- **Same behavior** as DLIF

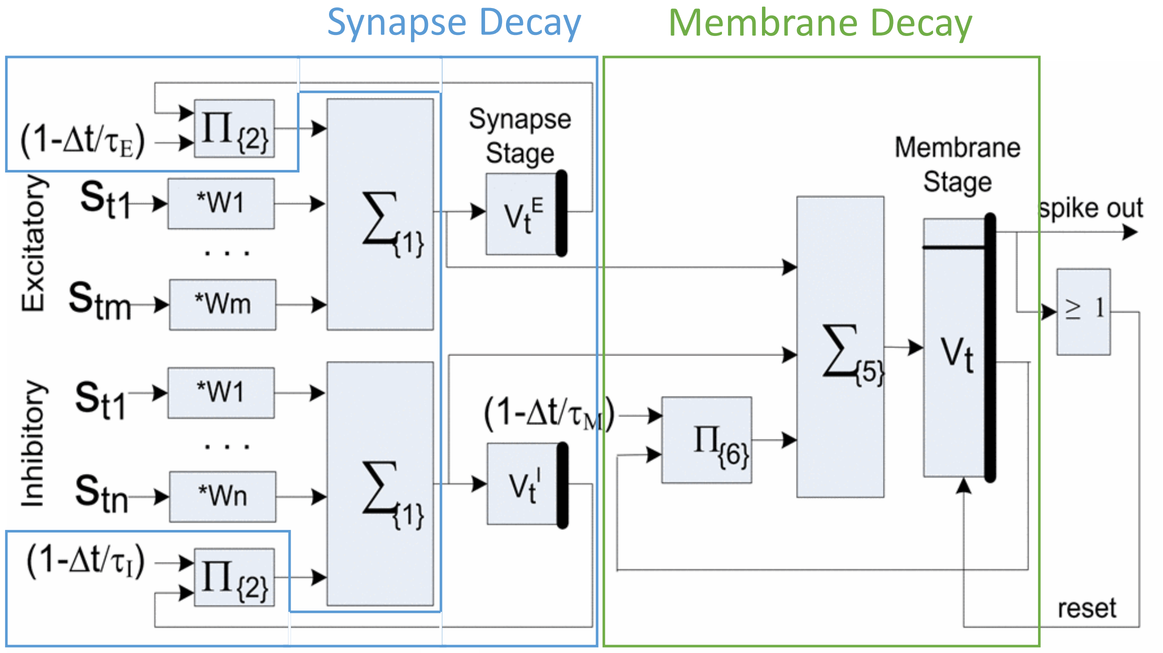

DSRM0 Neuron

{6} Mem Decays

{5} Update Mem Potential

{4} Compute Decay Value

{3} ComputePotential Changes

{2} Synapse Decays

{1} Synapse Adds

SRM0 = “zeroth order” of the more general Spike Response Model

LIF with Step Inputs (SLIF) Model

- **One stage** model
- Uses step inputs (synapse adds potential instantly)

Membrane Decay

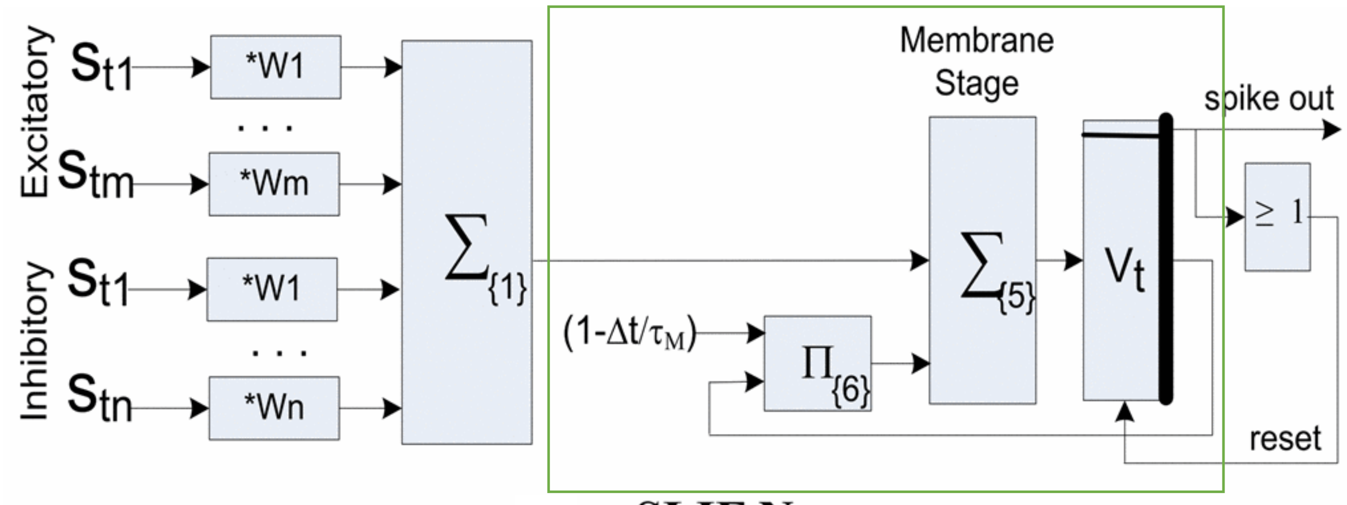

SLIF Neuron

{6} Mem Decays
{5} Update Mem Potential
{4} Compute Decay Value
{3} ComputePotential Changes
{2} Synapse Decays
{1} Synapse Adds

Linear Leak Integrate and Fire (LLIF)

- **One stage** model
- Like SLIF but uses linear decay for membrane potential
  - **more efficient** than SLIF but still **similar behavior**
- Simplest model in the context of this paper

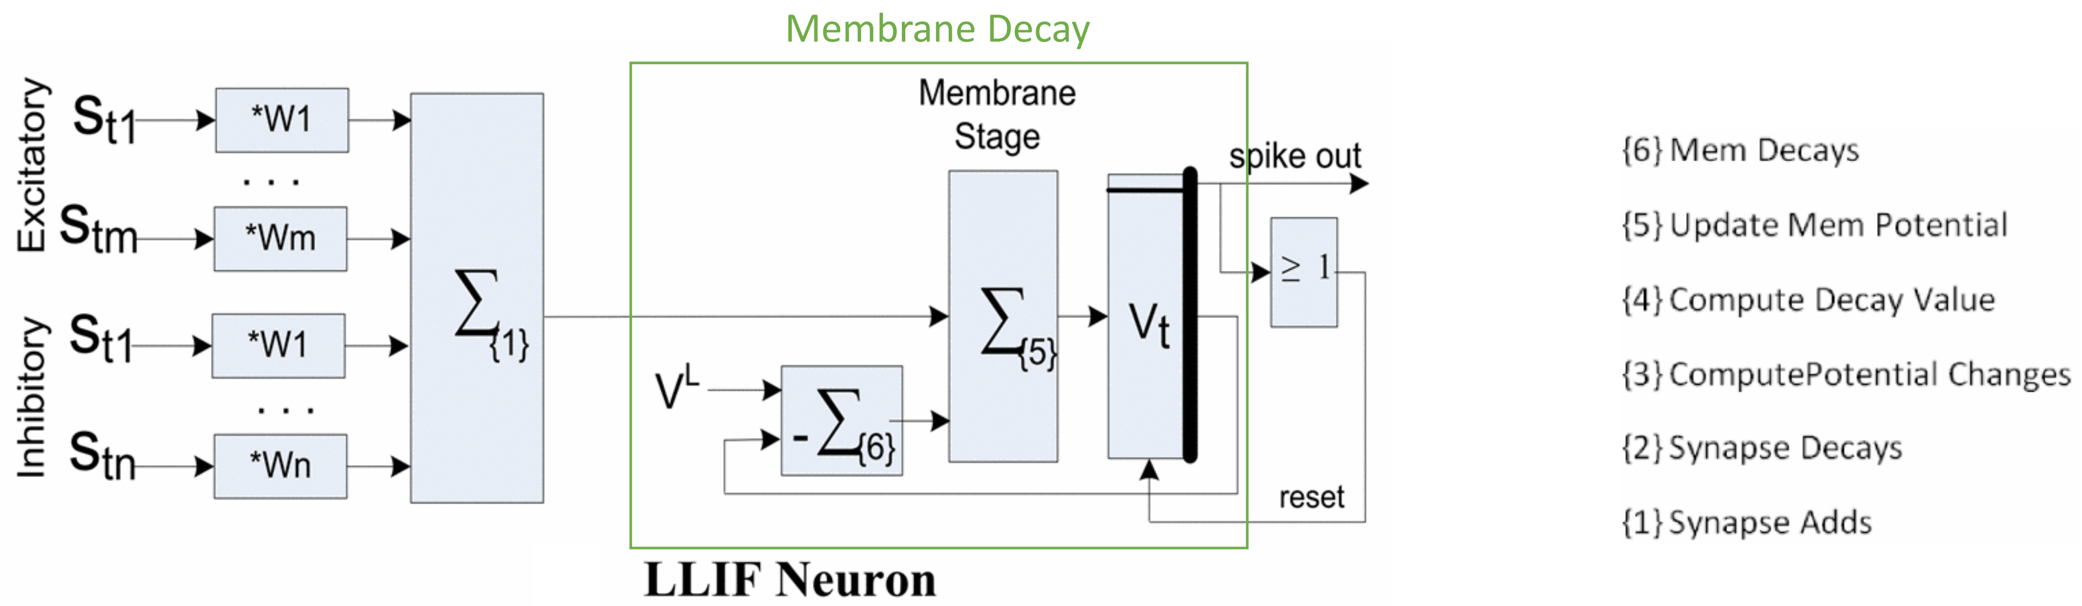

Comparison

Background

Problems & Goals

Digital Neuron Models

Comparison

- Methodology
- Behavior
- Efficiency

Efficient Two Stage Neuron

Key Takeaways

Strengths & Weaknesses

Thoughts & Ideas

Open Discussion

Methodology

- Simulator: GNU Octave (Matlab clone)
- Models have 100 synapses each
  - 80% excitatory
  - 20% inhibitory
- Input: correlated, randomly distributed input spike trains
- A coincidence measure is used to compare output spike trains
- Models are calibrated to produce similar input output behavior
  - Most complex neuron (DLIF) used as reference

Input

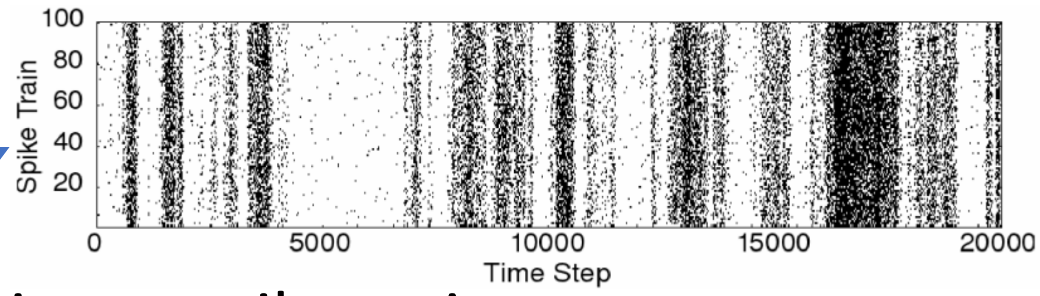

Output(DLIF)

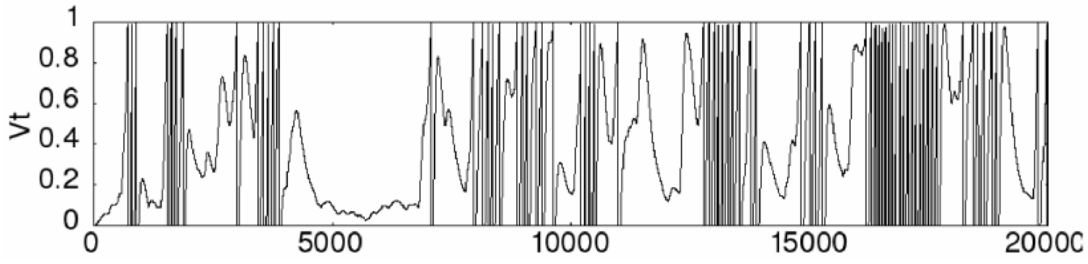

Behavior Results

- DSRM0 yields **maximum similarity** to DLIF
  - Both **two stage** neurons essentially **behave the same**
- **One stage** neurons **less similarity** to reference but still close
- Two stage neurons behave more like biological neurons than one stage neurons (clear difference in ellipse)

  -> **two stage** neurons are **more accurate**

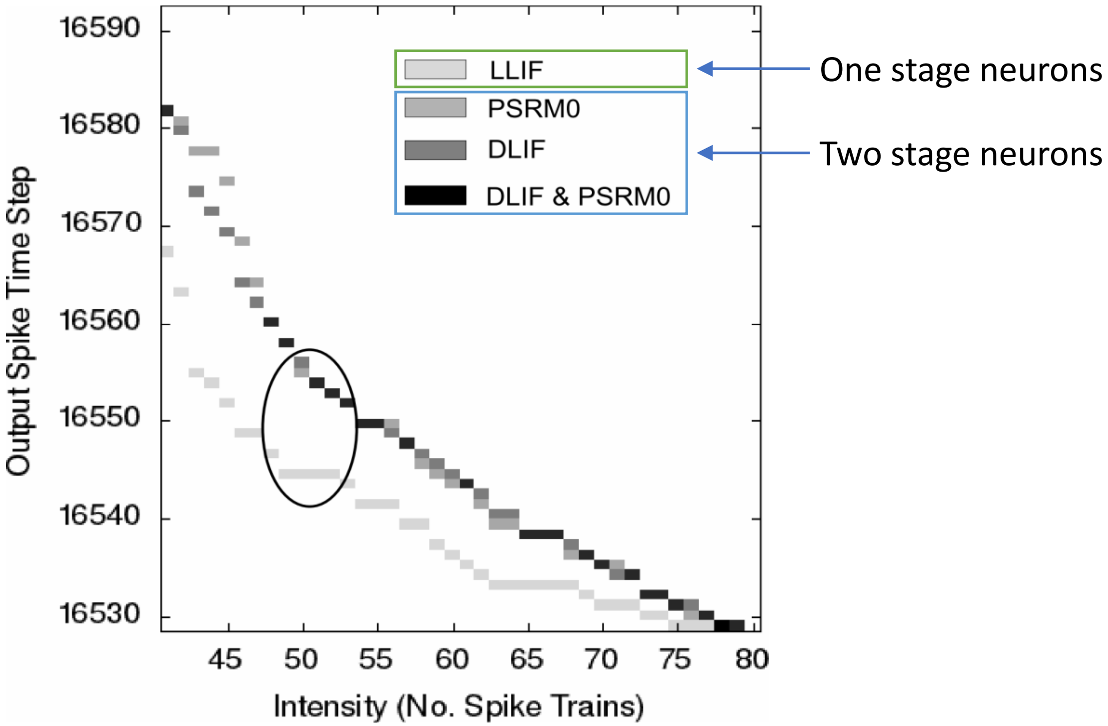

Efficiency Results

- LLIF over 10 times more efficient than DLIF
- DSRM0 requires fewer operations than DLIF but still almost 8 times more than LLIF
- SLIF requires about twice as many operations as LLIF

Graph: Number of operations needed to complete simulation

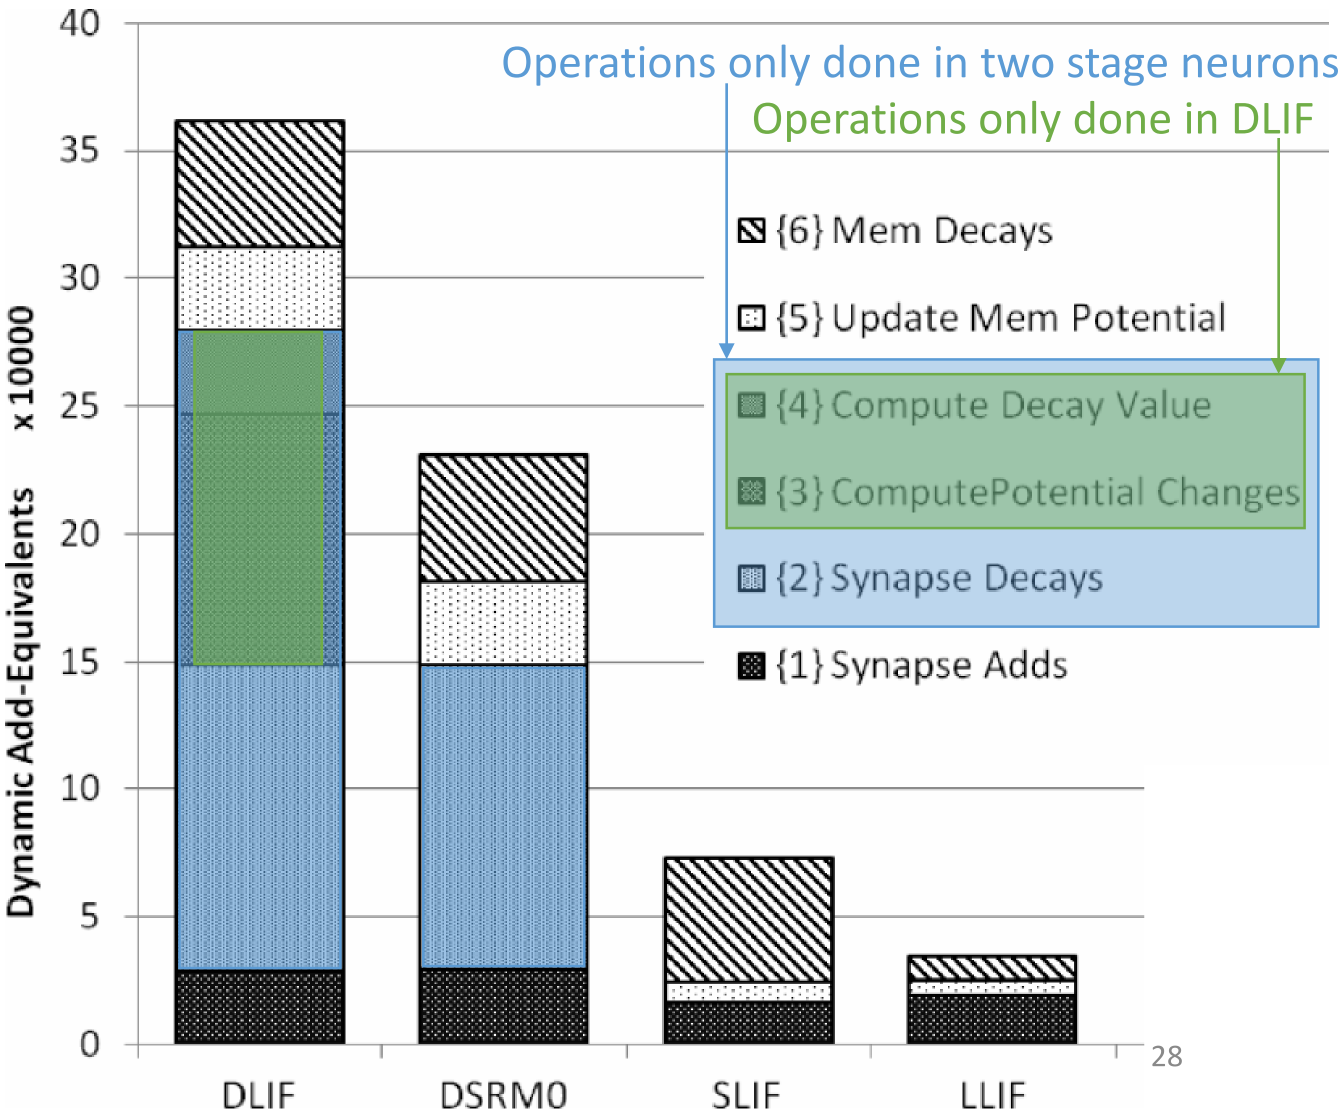

Efficient Two Stage Neuron

Proposed Improvements

- Paper proposes two ways to improve efficiency
  - Combine Synaptic Decay
  - Piecewise Linear Approximation
- Applying these changes results in a **two stage** neuron with **similar behavior** as DLIF and DSRM0 but a lot **more efficient**

Combined Synaptic Decay

- In DSRM0 excitatory and inhibitory synapse conductances decay with **different time constants** and need therefore be **calculated separately**

  Solution: **scale** inhibitory **synapse weights** so that same time constant can be applied -> synapse conductances can be combined and calculated together

- No big change in input/output behavior
- Eliminates roughly **half** of the **synaptic decay operations**

Piecewise Linear Approximation

- Idea is similar as with the LLIF neuron (simplest neuron)
  - Use single subtraction to calculate decay (rather than multiplication)
- A series of leak values is used that form a piecewise linear **approximation to exponential decay**
  - Update decrement value in between time steps (in relation to decaying value)
    - **Constant** decay operand D is **calculated once**
    - Decrement value = **binary shift of constant** D
- Can be used for synapse decay and membrane decay

- In short: Fewer operations are needed to calculate decay

RTL Implementation PSRM0

Combined Synaptic Decay

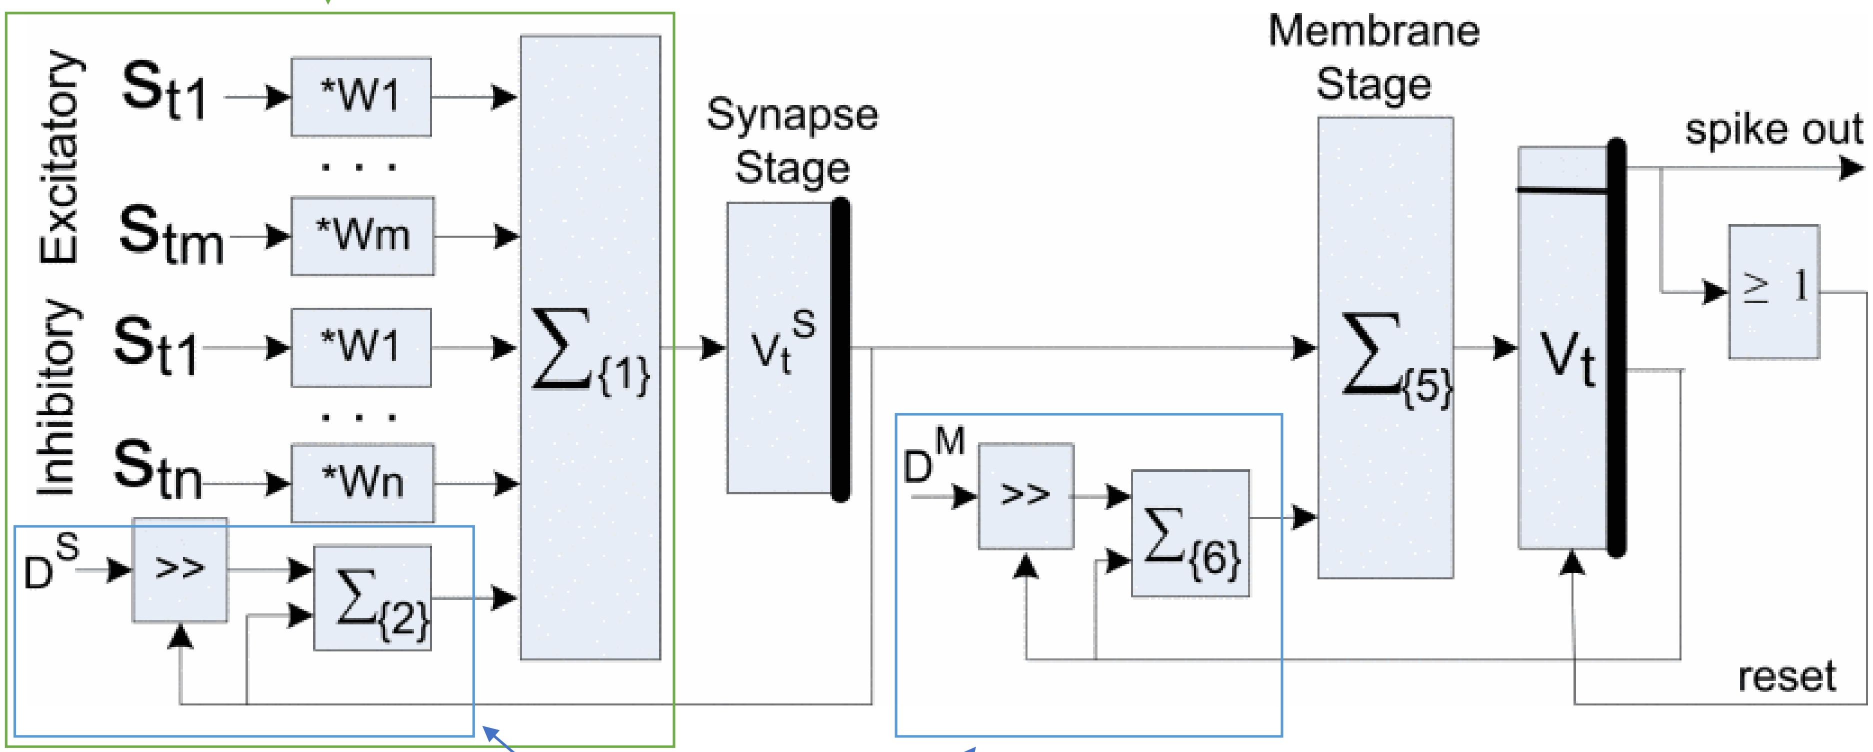

Piecewise Linear Approximation

{6} Mem Decays

{5} Update Mem Potential

{4} Compute Decay Value

{3} ComputePotential Changes

{2} Synapse Decays

{1} Synapse Adds

Results

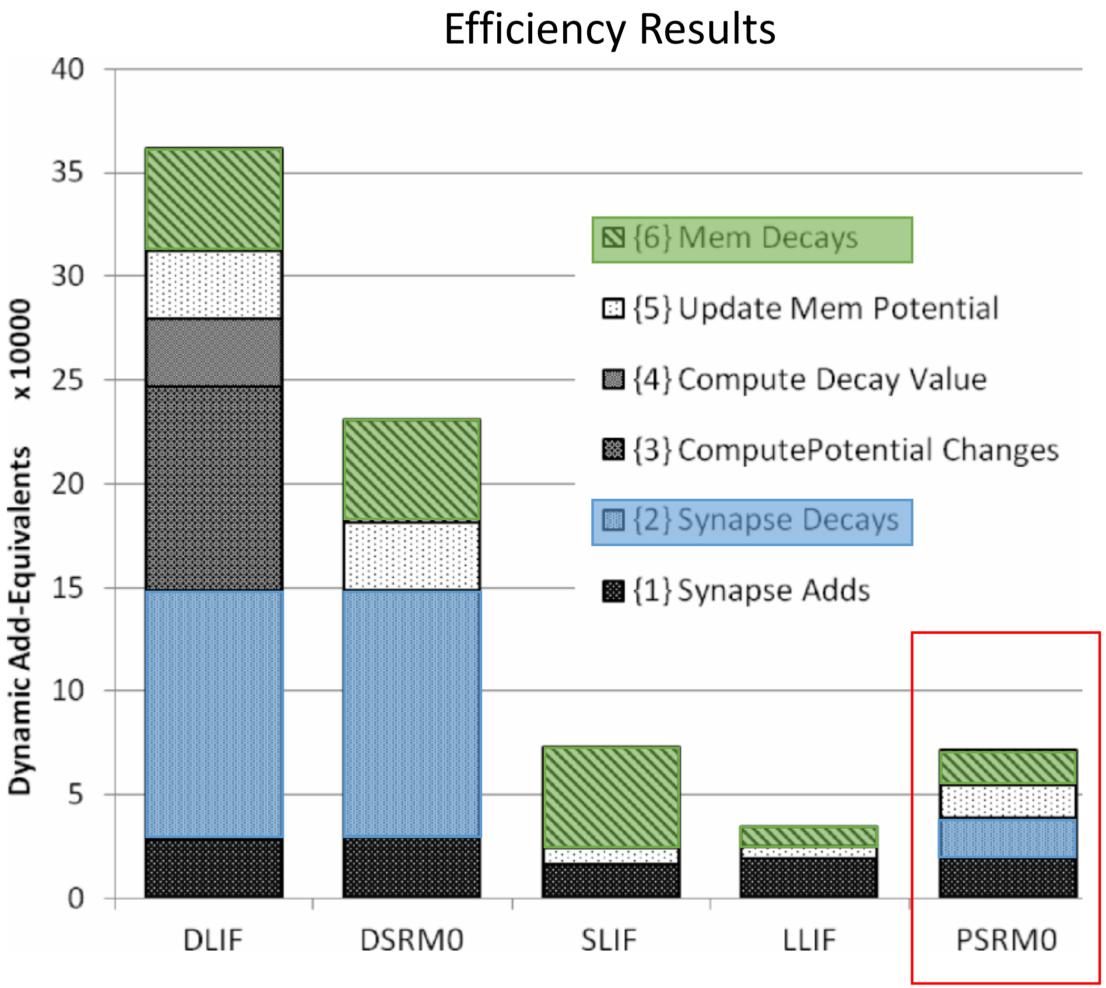
Efficiency Results
Dynamic Add-Equivalents x 10000
40
35
30
25
20
15
10
5
0
{6} Mem Decays
{5} Update Mem Potential
{4} Compute Decay Value
{3} ComputePotential Changes
{2} Synapse Decays
{1} Synapse Adds
DLIF
DSRM0
SLIF
LLIF
PSRM0

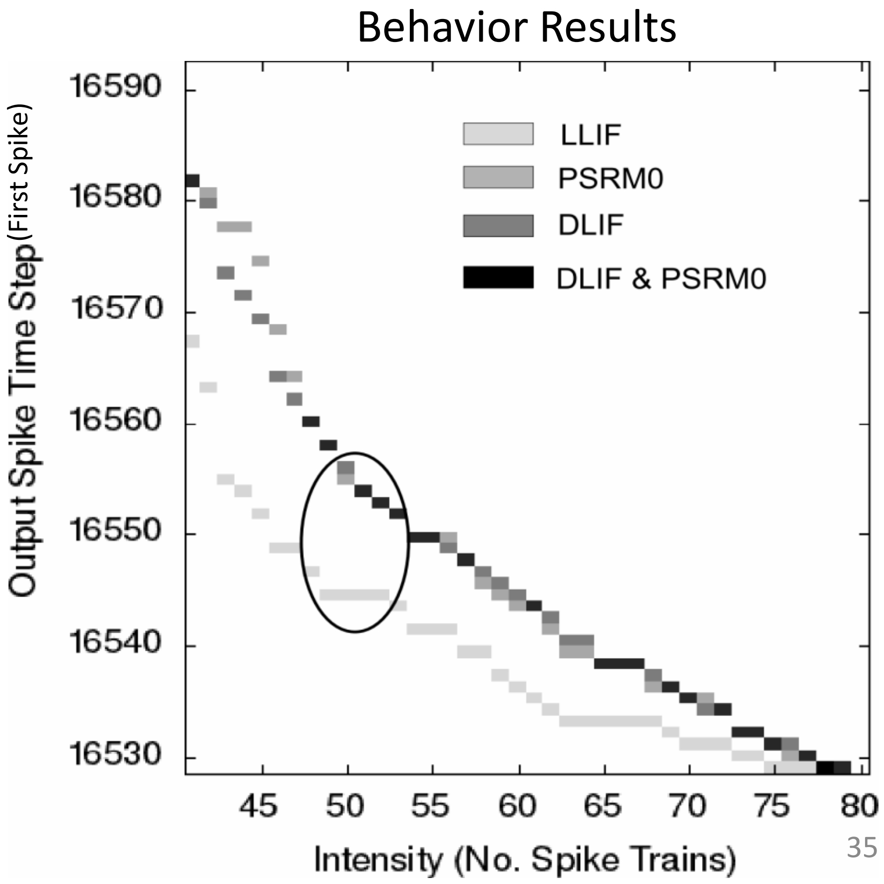
Behavior Results
Output Spike Time Step(First Spike)
16590
16580
16570
16560
16550
16540
16530
LLIF
PSRM0
DLIF
DLIF & PSRM0
45
50
55
60
65
70
75
80
Intensity (No. Spike Trains)

Key Takeaways

- Background
- Problems & Goals
- Digital Neuron Models
- Comparison
- Efficient Two Stage Neuron
- **Key Takeaways**
  - One Stage vs Two Stage
  - Best Two Stage Neuron
  - Best One Stage Neuron
  - Underlying Assumption
- Strengths & Weaknesses
- Thoughts & Ideas
- Open Discussion

Deciding between one or two stage neurons

- Behavioral difference exists between one and two stage neurons

  - Is it important for the system that the neurons are **biologically accurate**?
    - If yes, use two stage neuron
    - If no, use one stage neuron (more efficient)

  - One stage neurons may have advantages in event-driven systems

Best choice for **two stage** neurons

- DSRM0 is about 30% more efficient than DLIF

  It can be argued that the extra operations in DLIF are a biological artifact

  -> unnecessary operations

- PSRM0 is based on DSRM0 but is several factors more efficient
  - Improvements:
    1. Combining the synaptic weights
    2. Implementing piecewise linear decay
  - It still has a similar input/output behavior -> behaves like a biological neuron

-> The PSRM0 neuron is an excellent choice for large scale systems where two stage neuron behavior is desired

Best choice for **one stage** neurons

- LLIF neuron is twice as efficient as SLIF neuron
- LLIF does not appear to have significant computational disadvantages

-> LLIF is the better choice

Underlying Assumption

Paper works on the assumption that **individual neurons** will form the basic **building blocks** of future large scale systems.

Researchers might find that a **higher level of abstraction** will be a **better basis** to emulate the brains functionality and therefore individual neurons won't have to be implemented.

-> Neurons presented in this paper might become obsolete.

Executive Summary

Motivation: Large scale architectures are needed to emulate the neocortex to support research studying the operation of the brain.

Problem: Existing models of “complex” two stage neurons are **more accurate** but **less efficient** than “simple” one stage neuron models.

Goal: Provide implementations of **efficient** digital neurons that could be used in development of future large-scale cortical architectures.

Key Contributions:

- Four neurons are implemented in a manner that allows side-by-side comparison.
- Proposed a new **two stage neuron** that is **biologically accurate** and is almost as **efficient** as a one stage neuron.
- Method to compare implementations

Strengths & Weaknesses

Background

Problems & Goals

Digital Neuron Models

Comparison

Efficient Two Stage Neuron

Key Takeaways

Strengths & Weaknesses

- Strengths
- Weaknesses

Thoughts & Ideas

Open Discussion

Strengths

- Provides **efficient** implementations of **existing** neuron models
- Proposes an **efficient** implementation for a **new** two stage neuron
- Uses a nice method to compare neurons
- Mostly well-structured Paper
- Most parts are well explained

Weaknesses

- Some parts are complicated to understand
  - Could use more explanation
- Implements too many neurons
  - Could have **left out** the overly complex DLIF neuron and **go straight to** DSRM0 since they express essentially the **same behavior**

Thoughts & Ideas

Thoughts & Ideas

- Use **linear decay** on **two stage** neurons
  - Would input/output behavior still be similar?
  - Would the tradeoff be worth it?

- Use **piecewise linear decay** on **one stage** neuron
  - Would still approximate exponential decay of membrane potential
  - Slightly less efficient than LLIF but probably more accurate

Open Discussion

- Background
- Problems & Goals
- Digital Neuron Models
- Comparison
- Efficient Two Stage Neuron
- Key Takeaways
- Strengths & Weaknesses
- Thoughts & Ideas
- **Open Discussion**
  - Questions?
  - Discussion Starters

Questions?

Discussion Starters

- Can you think of other ways to improve existing digital neuron models or ideas for a new one?
- Are we ready to build a large scale system able to simulate the brain?
- Will we ever discover the computational paradigm of the brain?
  - Will this paper still be relevant when we do?
- Other Philosophical/Ethical topics? AI, Conscience, Moral, Rights …
  - Ex. Can a simulated brain think? If it thinks, does it exist?

End