

COMMUNICATING AND COMPUTING WITH SPIKES IN NEUROMORPHIC
SYSTEMS

A DISSERTATION
SUBMITTED TO THE DEPARTMENT OF ELECTRICAL ENGINEERING
AND THE COMMITTEE ON GRADUATE STUDIES
OF STANFORD UNIVERSITY
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

Sam Fok
June 2018

© 2018 by Sam Fok. All Rights Reserved.

Re-distributed by Stanford University under license with the author.

This dissertation is online at: <http://purl.stanford.edu/qk018rf9117>

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

Kwabena Boahen, Primary Adviser

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

Oussama Khatib

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

Krishna Shenoy

Approved for the Stanford University Committee on Graduate Studies.

Patricia J. Gumport, Vice Provost for Graduate Education

This signature page was generated electronically upon submission of this dissertation in electronic format. An original signed hard copy of the signature page is on file in University Archives.

Abstract

We provide an overview of neuromorphic engineering and describe two contributions to Braindrop, a state-of-the-art neuromorphic system. First, we describe a method for performing summing and weighting of spike trains by accumulative thinning, a deterministic procedure for merging and dropping spikes. Previous methods relied on probabilistic thinning, which results in Poissonian statistics. As a result, when the thinned spike-train is filtered with a first-order low-pass synapse, the signal-to-noise ratio (SNR) scales as the square-root of its rate. For our accumulative thinning method, the SNR depends on the weight w ; it scales linearly in the best-case scenario ($w \rightarrow 0$) and as the square-root in the worst-case ($w \rightarrow 1$). We find that a three-quarter power scaling minimizes energy consumption.

Second, we present a serial H-tree router for two-dimensional (2D) arrays. Existing routing mechanisms for 2D arrays either use low-overhead grids with one or two shared wires per row or column (e.g., RAM) or high-overhead meshes with many wires connecting neighboring clients (e.g., supercomputers). Neither is suitable for intermediate-complexity clients (e.g., small clusters of silicon neurons). We present a router tailored to 2D arrays of such clients. It uses a tree laid out in a fractal pattern (H-tree), which requires less wiring per signal than a grid, and adopts serial-signaling, which keeps link-width constant, regardless of payload size. To route from the tree's leaves to its root (or vice versa), each node prepends (consumes) a delay-insensitive 1-of-4 code that signals the route's previous (next) branch; additional codes carry payload. We employ this serial H-tree router to service a 16×16 array of silicon-neuron clusters, each with 16 spike-generating analog somas, 4 spike-consuming analog synapses, and one 128-bit SRAM. Fabricated in a 28-nm CMOS process, the router communicates 26.8M soma-generated and 18.3M synapse-targeted spikes per second while occupying 43% of the client's $35.1 \times 36.1 \mu\text{m}^2$.

Acknowledgments

I thank Chris Eliasmith and Terry Stewart of the Center for Theoretical Neuroscience at the University of Waterloo for the development and assistance with the Neural Engineering Framework that motivated the Braindrop project. I thank Rajit Manohar of the Computer Systems Lab at Yale University for instruction on asynchronous circuit design and for the CAD tools that allow their synthesis and verification in a synchronous world. I thank my colleagues at the Brains in Silicon Lab of Stanford University for countless coffee walks and their stimulating discussions. In particular, I could not have asked for better teammates than Ben Benjamin, Alex Neckar, and Nick Oza on the Braindrop project. And finally I have been immensely fortunate to have Kwabena Boahen as advisor and thank him for his tireless drive for understanding. The works herein were made possible by grants from the Office of Naval Research and the National Science Foundation.

Contents

Abstract	iv
Acknowledgments	v
1 Artificial Intelligence and Neuromorphic Engineering	1
2 Summing and Weighting Spike Trains	5
2.1 Summing and Weighting	6
2.2 Discrete to Continuous	8
2.2.1 Poisson Process	15
2.2.2 Periodic Process	16
2.2.3 p-Thinned Periodic Process	17
2.2.4 d-Thinned Poisson Process	19
2.2.5 Approximation Quality	21
2.3 Optimizing for Power	21
2.4 Discussion	22
3 A Serial H-Tree Router for Two-Dimensional Arrays	24
3.1 Router Functionality and Overhead	24
3.2 Tree Paths versus Grid Addresses	26
3.3 Serial Communication Protocol	28
3.4 Router Logical Design	29
3.4.1 Transmitter	29
3.4.2 Receiver	31
3.5 Router Application	33
3.5.1 Transmitter Leaf	33
3.5.2 Receiver Leaf	33
3.5.3 Serial-Parallel Conversion	34
3.6 Synthesis and Validation	35

3.7 Discussion	36
4 Conclusions	43
A Spike Summing and Weighting	44
A.1 periodic SNR approximation	44
A.2 p-thinning SNR approximation	45
A.3 d-thinning SNR expansion	45
B AER Transmitter Design Space	47
B.1 AEXT Control Data decomposed (CD)	47
B.1.1 AEXT CD noTW CYC	47
B.1.2 AEXT CD noTW	56
B.1.3 AEXT CD TW	61
B.2 AEXT Control Data Combined	67
B.3 AEXT ASPR NODE	67
B.4 AEXT ASPR PFWD/MERGE (PM)	67
B.5 AEXT ASPR PM PFWD unpipelined	68
B.6 AEXT ASPR PM PFWD pipelined hq	69
B.7 AEXT ASPR PM PFWD hu	71
B.8 AEXT ASPR PM MERGE unpipelined	73
B.9 AEXT ASPR PM MERGE pipelined a.a	74
B.10 AEXT ASPR PM MERGE pipelined ah	76
B.11 AEXT ASPR PFWD PREPEND/FWD/SIMPLE_MERGE	78
B.12 AEXT ASPR PFWD PREPEND	78
B.13 AEXT PFWD FWD	78
B.14 AEXT ASPR PFWD SIMPLE_MERGE	78
B.15 AEXT ASPR MERGE	79
C AER Receiver Design Space	80
C.1 Receiver tree structure	80
C.2 AERV CD noTW cyclic control (CYC)	82
C.3 AERV CD noTW CYC NODE	83
C.4 AERV CD noTW CYC LEAF	88
C.5 AERV CD noTW CYC LEAF (no data)	92
C.6 AERV Control Data decomposed (CD) no tailword (noTW)	94
C.7 AERV CD noTW SPLIT	95
C.8 AERV CD noTW CTRL	97
C.9 AERV CD noTW LEAF	98

C.10 AERV ASPR BCAST pipelined	99
C.11 AERV ASPR BCAST unpipelined	101
C.12 AERV PSAR	102
C.13 AERV PSAR decomposed into ROUTE, READ_HEAD, FWD_BODY (RHB)	102
C.14 AERV PSAR RHB ROUTE unpipelined	103
C.15 AERV PSAR RHB READ_HEAD	104
C.16 AERV PSAR RHB FWD_BODY unpipelined	105
C.17 AERV PSAR RHB FWD_BODY pipelined	106
C.18 AERV PSAR decomposed into ROUTE PULL_CTRL PULL (RCP)	107
C.19 AERV PSAR RCP ROUTE	108
C.20 AERV PSAR PULL_CTRL	109
C.21 AERV PSAR RCP PULL unpipelined	110
D AER Interface Design Space	112
D.1 OUT e1ofN	112
D.2 OUT a1ofN	113
D.3 Deserializer	114
D.4 Ring Deserializer	114
D.5 SPLIT	116
D.6 NODE	117
D.7 C	119
D.8 RING	119
D.9 CHAIN Deserializer	124
D.10 HEAD	125
D.11 NODE	127
D.12 TAIL	130
D.13 Split Chain Deserializer	131
D.14 SPLIT	133
D.15 NODE	134
D.16 Serializer	136
D.17 Ring serializer	136
D.18 Ring serializer NODE	137
D.19 Ring serializer sequencer	139
D.20 Ring serializer 2	140
D.21 Ring serializer 2 NODE	141
D.22 Ring serializer 2 MERGE	141
D.23 Chain serializer	143
D.24 Chain serializer NODE	144

D.25 Chain serializer TAIL	146
D.26 Chain serializer C	147
D.27 Serial merge	147
D.28 Memory	148
E Serial Router Supplement	150
E.1 Half-Cycle Slack Buffer	150
E.2 Greedy but Fair N-way Arbiter using Tree/Ring Sequencing	151
E.3 NODE	152
E.3.1 PRODUCER	152
E.3.2 CONSUMER	155
Bibliography	157

List of Tables

2.1	Renewal Process ISI CV and Synapse SNR	9
3.1	Node Transistor Counts & Binary:4-ary Ratios for Large N	28
3.2	Handshaking Expansion (HSE) Syntax	38
3.3	Communicating Hardware Processes (CHP) Syntax	38

List of Figures

1.1	Artificial Neural Network	2
1.2	Digital Logic	3
1.3	Ion Channels and Subthreshold Transistors	4
2.1	Spiking Neural Network	6
2.2	Summing by merging.	7
2.3	Probabilistic (P) vs Deterministic (D) Thinning.	9
2.4	Renewal Process-Driven Synapse	12
2.5	Rate-Estimate SNR.	13
2.6	Theory vs numerical simulations.	16
2.7	SNR Approximation.	21
2.8	d-Thinned Poisson Power.	23
3.1	Grid Addresses and Tree Paths	25
3.2	Serial Link Description at Two Levels of Abstraction	27
3.3	Router Process Decomposition	39
3.4	Four-Way Arbiter	40
3.5	Serial-Parallel Conversion	40
3.6	Tree Router Layout	41
3.7	Postlayout Transmitter and Receiver SPICE Simulations	42
3.8	Fabrication and validation	42

1

Artificial Intelligence and Neuromorphic Engineering

Humans have long analogized the human brain to their computational tools and sought to imbue their tools with human-like, artificial intelligence (AI). For the industrial age, the brain was like a steam-powered system of pipes and actuators driving the body; for the digital age, the brain is like a digital computer, storing and processing bits of data. However, instead of conceptualizing the brain in terms of the current means of computation, neuromorphic engineers move in the other direction and conceive of computational methods from the structure and function of the brain. The neuromorphic approach comes with good reason: on a power budget of 20W (slightly more than one-half cup of sugar per day), the brain, with 86 billion neurons [18] and thousands of synaptic connections per neuron, still roundly beats existing methods at tasks we take for granted like walking around, having intelligible conversations, and making sense of the world.

That the brain has useful lessons for computation was not as widely accepted in the recent past as it is today in 2018. However, since AlexNet won the ImageNet challenge in 2012 and beat the nearest competitor by an unheard of margin (10.8 percent) [23], brain-inspired artificial neural networks (ANNs) and the hardware to run them have returned the mainstream conscious as a viable approach for AI. ANNs like AlexNet are loosely based on the connectivity of cortex in the brain and contain from thousands of simple units, called "neurons", arranged and connected in layers (Fig. 1.1). Each neuron computes a simple nonlinear function of the sum of its inputs, and each connection weights the source neuron's output, or activation, and delivers it as input to a destination neuron.

Prior to AlexNet's success, algorithms geared towards AI were largely not brain-inspired and hardware design efforts were concentrated on powerful, general-purpose central processing units (CPUs), which operate synchronously and digitally: a central clock governs the flow of data and




Figure 1.1: Artificial Neural Network

Inputs are passed through layers of neurons and connections. Neurons compute simple nonlinear functions of their summed input while connections weight signals passing through them.

instructions through the computer and numbers are represented in binary. The 0s and 1s of binary correspond to the power supply voltage and ground of digital logic (Fig. 1.2). Further, these CPUs instantiate part of the broader Von Neumann computer architecture, which separates state (the memory) from computation (the CPU) and excel at sequential programs, or programs that require executing steps one at a time. Although Von Neumann (of the Von Neumann architecture) himself noted the differences between a computer’s architecture and the brain’s as well as the potential of a more brain-like architecture [44], there was little economic incentive to develop alternative computing architectures. The computing industry adopted the silicon semiconductor metal-oxide-semiconductor field-effect transistor (MOSFET) as their physical substrate, which permitted exponentially increasing transistor counts in a chip over time—dubbed Moore’s Law. When computer architects could expect to double a designs’ performance every two to three years by switching to the latest silicon manufacturing process, alternative designs remained out of consideration. As a result, after decades of Moore’s law, the digital computer approach has become synonymous with computation; the information age is synonymous with the digital age.

However, physical devices have physical limits, so developing brain-like architectures makes more economic sense as the semiconductor industry takes longer and spends more to shrink transistors. Specifically, at the smallest transistor sizes, the cost per transistor no longer decreases with transistor size, so it is no longer economically viable to rely on device shrinkage for performance gains. In the vacuum after nearly 50 years of Moore’s law, AlexNet succeeded by matching a brain-inspired ANN structure, the convolutional neural network (CNN), with non-CPU hardware, the graphical processing unit (GPU), that was well-suited to the network’s architecture. An ANN operates by passing inputs through thousands of neurons and millions of parameters simultaneously, and moving all of that data around is a poor fit for CPUs which are bottlenecked by the separation between state




Figure 1.2: Digital Logic

LEFT: Digital computing relies on transistor networks (black boxes) to switch an output node between a connection to the power supply (top, blue), representing the binary value 1, and a connection to ground (bottom, red), representing the binary value 0. RIGHT: In this simple digital circuit, an inverter, the output is the inverted input.

and computation. A GPU better serves ANNs by matching the structure of the ANN dataflow—GPU inputs flow through many simple cores in parallel instead of one complicated core sequentially. AlexNet’s success on a GPU convincingly demonstrating the utility of brain-inspired computational approaches and hardware.

Though people now recognize the value of highly-parallel, brain-like computing, there remain significant differences between the brain’s architecture and today’s hardware. To realize an ANN’s thresholding, weighting, and summing operations efficiently, computer architects have developed application-specific integrated circuits (ASICs) [21], but these ASICs still rely on the predominant, synchronous-digital paradigm. However, from Carver Mead’s work in the 1980s and 90s [34], neuromorphic engineers recognize that the physics underlying transistors and the physics underlying ion channels in a biological neurons membrane are qualitatively matched. Specifically, electrons or holes passing through a subthreshold MOSFET and ions passing through a neuron’s membrane are both diffusion driven phenomena [33]. As a result, both are voltage controlled current sources with an exponential dependence of current on voltage (Fig. 1.3). That is, subthreshold MOSFETs are analogous to ion currents through a neuron’s membrane.

With subthreshold MOSFETs as their building blocks, neuromorphic engineers push the brain analogy deeper than network-level structures and build silicon neurons that follow the inherent dynamics of biological neurons—spiking dynamics. In contrast to simple and static artificial neurons, biological neurons are continuous dynamical systems governed by electrochemical gradients and voltage-sensitive ion-channel proteins embedded in their membranes. Their most prominent behavior is the generation and emission of stereotyped, traveling depolarizations (spikes) along their projections (axons) to other neurons.

Neuromorphic engineers construct large scale spiking neural networks (SNNs) to leverage the minimal power consumption of subthreshold MOSFETs for useful computation and to understand the value of dynamics and spiking for brain-like computation with a forward engineering approach. While the predominant ANN approach running on synchronous-digital hardware has led to impressive achievements in image classification, speech classification, game playing, driving autonomy,

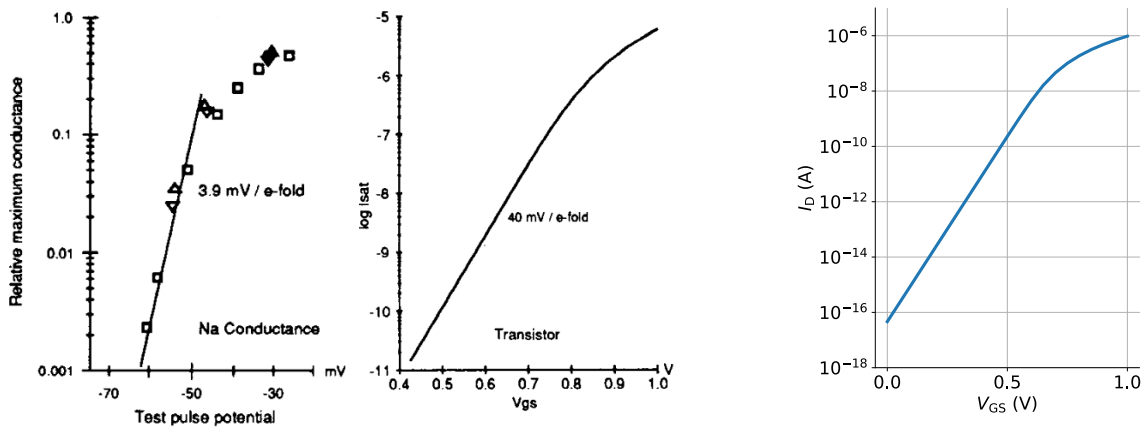


Figure 1.3: Ion Channels and Subthreshold Transistors

LEFT: Adapted from [34]. Comparison between sodium channel conductance in a neuron for a given membrane potential and a transistor’s source to drain current for a given gate voltage show similar exponential relationships between voltage and current. RIGHT: the gate voltage to drain current relationship in a more modern, 28nm process transistor.

and many other fields, practitioners are still working to push such AI systems into highly power-constrained environments (and hence limited connectivity) that require real-time operation. The femtowatt power for subthreshold MOSFETs provide an attractive means for computing on minimal power budgets (Fig. 1.3, RIGHT), and by systematically building SNNs from low-power spiking dynamical circuits, neuromorphic engineers also address the critique [20] leveled at Neuroscience that reverse engineering and modeling techniques have yet to yield a mechanistic understanding of the brain. Such a mechanistic understanding is inherent to the systematic construction of SNNs practiced by neuromorphic engineers.

Prior to my joining the field, neuromorphic engineers developed efficient and scalable neuromorphic systems as culminated in Neurogrid with a million neurons and a billion synapses [2]. However, a million neurons and billion synapses are difficult to use without systematic means of mapping higher level computations onto the spiking neurons. By combining Neurogrid with an existing means, the Neural Engineering Framework (NEF) [15], we demonstrated the NEF principles on Neurogrid [12] and demoed a robot arm controller with Neurogrid’s spiking silicon neurons [35]. From the lessons of marrying Neurogrid and NEF, we began our next project to build Braindrop, a chip architected specifically for large scale SNNs as systematically arranged with NEF. In this thesis, I present two contributions to the Braindrop chip that push forward our understanding of spiking communication and computation in neuromorphic systems. Chapter 2 describes the statistical consequences of spiking communication and computation in first-order dynamical systems, and Chapter 3 describes a router network for communicating spikes and programming packets.

2

Summing and Weighting Spike Trains

Similar to ANNs like AlexNet, SNNs rely on numerous neurons (providing thresholding) and connections (providing weighting and summing) (Fig. 2.1, ARCHITECTURE & WEIGHT MATRIX). Whereas these synchronous-digital designs operate at high-precision (scales exponentially with the number of bits) thousands of times faster than real-time (processing image or speech data in large batches), the hybrid analog-digital brain operates at vastly lower-precision (1-3 bits per spike [9]) in real-time. Similarly, ANN algorithms perform well at low precisions [13, 19] and real-time operation is appropriate for applications at the edge (processing a single user’s speech).

Neuromorphic engineers seek to realize efficient and robust hardware architectures and computational frameworks by emulating the brain’s hybrid analog-digital approach, trading precision and speed for energy efficiency [2, 3, 15]. Specifically, spiking-analog somas provide thresholding while consuming femtoamps of current, and all-or-nothing spikes—a unary representation—permit weighting with a random number generator and a comparator as well as summing by merging with a multiplexer (Fig. 2.1, OPERATIONS). Although such probabilistic (p) thinning, is cheap to implement,¹ it may be just as costly in energy as a conventional, binary multiply-accumulate unit when controlling for precision [29] due to quadratic—as opposed to logarithmic—scaling.

In this chapter, we describe an alternative, deterministic (d) thinning method and demonstrate its advantages over probabilistic thinning by analyzing each method’s statistical consequences. In Section 2.1, we describe the weighting and summing operations and analyze each operation’s output statistics by treating spikes as points in a temporal point-process. In Section 2.2, we analyze the

¹A spike train with rate λ is thinned by dropping a fraction $(1 - w)$ of the spikes to produce a spike train with rate $w\lambda$, thus weighting the spike train by w . Although this method applies to positive signals, it is readily extended to negative signals by permitting spikes to be signed. Zero is still implicitly represented by an absence of spikes (i.e., zero represents itself).

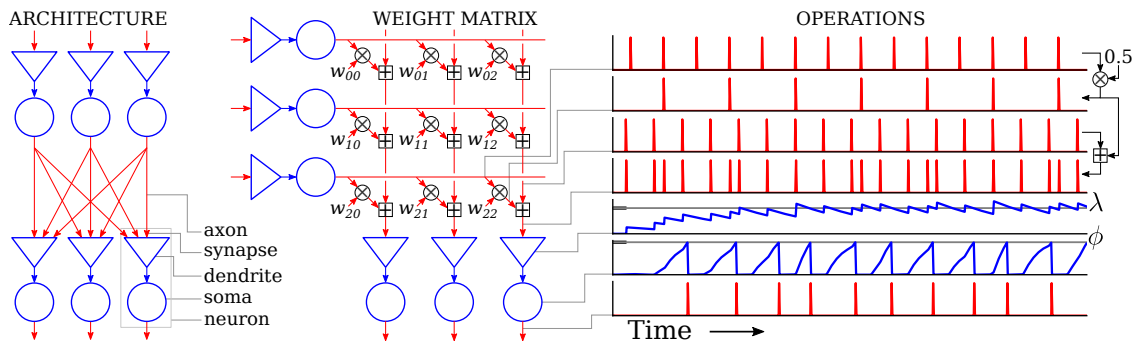


Figure 2.1: Spiking Neural Network

ARCHITECTURE: Spikes flow through axons and synapses (red) while currents flow through dendrites and somas (blue). **WEIGHT MATRIX:** Spike trains are weighted (\otimes) and summed (\boxplus) as they travel from one layer to the next. **OPERATIONS:** Weighting is by thinning and summing is by merging. Conversion from spike trains to currents is by low-pass filtering in synapses, whose currents settle around the input spike-rate λ . Conversion from currents back to spikes is by integration and thresholding in somas, whose spiking threshold is ϕ .

statistical effects of thinning spike-trains on synaptic signals and show that the two methods produce different SNR scaling. In Section 2.3, we analyze and optimize the power requirements for the two methods. In Section 2.4, we discuss our results and conclude the chapter.

2.1 Summing and Weighting

An accumulator implements summing and weighting in an ANN by combining merging with d-thinning. Merging produces a more Poissonian spike-train, while d-thinning produces a more periodic spike-train. It is crucial that merging occur before thinning. That is, instead of merging the outputs of N accumulators together, we merge N spike trains and feed the resulting spike train into a single accumulator (unlike in Figure 2.1, WEIGHT MATRIX). This ordering preserves the accumulators' more periodic-like spike-train statistics.

Summing by merging (Fig. 2.1, MERGING) produces more Poissonian spike-trains (Fig. 2.1, PERIODIC & POISSON INPUT). That is, the interspike-intervals (ISIs) become independent and identically distributed (IID) exponential random variables (i.e. from a renewal process). Consider the case when all N spike-trains have the same rate $\frac{\lambda}{N}$. Pick one as a reference and superpose the remaining $N - 1$ on it with random offsets T_j , drawn independently from a uniform distribution on $\{0, \frac{N}{\lambda}\}$ (i.e. within the reference spike-train's first ISI). The resulting, merged spike-train's first ISI is given by $\Delta T_1 = \min(T_1, T_2, \dots, T_{N-1})$. Its cumulative distribution function (CDF) is




Figure 2.2: Summing by merging.

MERGING: Spike-trains are merged, to producing an output spike-train whose rate is the sum of the input rates. PERIODIC vs POISSON INPUT: As N , the number of (equal-rate) spike trains increases, the cumulative distribution functions (CDFs) of empirical output ISIs (T_{out}), transition from degenerate (dashed line) to exponential (dotted line) versus remaining exponential. CDFs include 10,000 samples of T_{out} , normalized by their mean.

$$\begin{aligned}
 P(\Delta T_1 \leq t) &= P(\min(T_1, T_2, \dots, T_{N-1}) \leq t) \\
 &= 1 - P(T_1 > t \wedge T_2 > t \wedge \dots \vee T_{N-1} > t) \\
 &= 1 - \int_t^{\frac{N}{\lambda}} \int_t^{\frac{N}{\lambda}} \dots \int_t^{\frac{N}{\lambda}} \left(\frac{\lambda}{N}\right)^N dt_1 dt_2 \dots dt_N \\
 &= 1 - \left(\frac{\lambda}{N}\right)^N \left(\frac{N}{\lambda} - t\right)^N \\
 &= 1 - \left(1 - \frac{\lambda t}{N}\right)^N \underset{N \rightarrow \infty}{=} 1 - e^{-\lambda t}
 \end{aligned}$$

recognizing that $\lim_{N \rightarrow \infty} \left(1 + \frac{x}{N}\right)^N = e^x$. This exponential distribution describes other ISIs in the merged spike-train as well since its first ISI is not special (i.e., its ISIs are identically distributed). It also arises when the spike rates are not equal, provided that consecutive spikes are never from the same neuron, which is satisfied if no neuron fires more than N times faster than the mean.

To see that ISIs are also independent for sufficiently large N , consider the merged, equal-rate periodic spike-trains again. The number of unique ISIs is exactly N ; once all N unique ISIs have been traversed, they will repeat. However, this will take a longer and longer time as N tends to infinity. Once this time-scale exceeds several synaptic time constants, the ISIs can be effectively considered to be independent.²

p and d-thinning (Fig. 2.1, THINNING), produce more Poissonian and more periodic spike-trains for decreasing weight w , respectively (Fig. 2.1, P-THIN & D-THIN). That is, the resulting spike-train's

²It is crucial that the individual spike-train *rates* tend to 0 as the number of spike-trains tends to infinity in contrast to scaling the spike *amplitudes* (represented by delta functions). If only the spikes themselves are scaled, the superposed independent sources will not converge to a Poisson process—even in the limit [26]. As a consequence, spike times will synchronize more strongly through multiple, feed-forward layers of a neural network [14].

ISIs' (X) coefficient of variation ($CV(X)$) approaches 1 and 0, respectively.³ An ISI in a p-thinned spike-train, is given by $\Delta T_{\text{out}} = \sum_j^S \Delta T_{\text{in}_j}$, where ΔT_{in_j} are the IID pre-thinned ISIs, and S is the number of ISIs skipped before a spike makes it through. Since S is geometrically distributed with parameter w ,

$$\begin{aligned} \mathbf{E}[\Delta T_{\text{out}}] &= \mathbf{E}[\Delta T_{\text{in}}]/w \\ \text{Var}(\Delta T_{\text{out}}) &= (w\text{Var}(\Delta T_{\text{in}}) + \mathbf{E}[\Delta T_{\text{in}}]^2(1-w))/w^2 \\ \text{CV}_{\text{pthin}}(\Delta T_{\text{out}}) &= \sqrt{w \frac{\text{Var}(\Delta T_{\text{in}})}{\mathbf{E}[\Delta T_{\text{in}}]^2} + (1-w)} = \sqrt{w\text{CV}(\Delta T_{\text{in}})^2 + (1-w)} \end{aligned}$$

w interpolates between $\text{CV}_{\text{poi}}(\Delta T_{\text{out}}) = 1$ and $\text{CV}(\Delta T_{\text{in}})$. As $w \rightarrow 0$, $\text{CV}_{\text{pthin}}(\Delta T_{\text{out}}) \rightarrow 1$, while as $w \rightarrow 1$, $\text{CV}_{\text{pthin}}(\Delta T_{\text{out}}) \rightarrow \text{CV}(\Delta T_{\text{in}})$ (i.e, all input spikes are output). For example, when a periodic spike-train ($\text{CV}_{\text{per}}(\Delta T_{\text{out}}) = 1$) is p-thinned, $\text{CV}_{\text{pthin}}(\Delta T_{\text{out}}) = \sqrt{1-w}$ (Table 2.1, p-thinned periodic), becoming more Poisson as $w \rightarrow 0$.

In the case of d-thinning, the spike train becomes more periodic with $\Delta T_{\text{out}} = \sum_j^{1/w} \Delta T_{\text{in}_j}$.⁴ The sum now contains a deterministic number ($1/w$) of ΔT_{in} ,

$$\begin{aligned} \mathbf{E}[\Delta T_{\text{out}}] &= \mathbf{E}[\Delta T_{\text{in}}]/w \\ \text{Var}(\Delta T_{\text{out}}) &= \text{Var}(\Delta T_{\text{in}})/w \\ \text{CV}_{\text{dthin}}(\Delta T_{\text{out}}) &= \sqrt{w \frac{\text{Var}(\Delta T_{\text{in}})}{\mathbf{E}[\Delta T_{\text{in}}]^2}} = \sqrt{w}\text{CV}(\Delta T_{\text{in}}) \end{aligned}$$

Now w interpolates between $\text{CV}_{\text{per}}(\Delta T_{\text{out}}) = 0$ and $\text{CV}(\Delta T_{\text{in}})$. As $w \rightarrow 0$, $\text{CV}_{\text{dthin}}(\Delta T_{\text{out}}) \rightarrow 0$, while as $w \rightarrow 1$, $\text{CV}_{\text{dthin}}(\Delta T_{\text{out}}) \rightarrow \text{CV}(\Delta T_{\text{in}})$ (as with CV_{pthin}). For example, when a Poisson spike-train ($\text{CV}_{\text{poi}}(\Delta T_{\text{in}}) = 1$) is d-thinned, $\text{CV}_{\text{dthin}}(\Delta T_{\text{out}}) = \sqrt{w}$, becoming more periodic as $w \rightarrow 0$ (Table 2.1, d-thinned Poisson).

2.2 Discrete to Continuous

Spiking neural networks encode signals in their spike rates, which requires estimating the underlying spike rate from the train of spikes with methods that respect causality and track temporal changes in the rate. The first requirement precludes estimating the instantaneous spike-rate as it requires

³ $CV(X) = \sqrt{\text{Var}(X)}/\mathbf{E}[X]$, where $\mathbf{E}[X]$ is X 's expectation (i.e, mean), and $\text{Var}(X)$ is its variance. For ISIs in a periodic spike-train, $\text{CV}_{\text{per}}(X) = 0$, and for ISIs in a Poisson spike-train, $\text{CV}_{\text{poi}}(X) = 1$ (Table 2.1).

⁴When ΔT_{in} is exponentially distributed, ΔT_{out} , will be gamma—or more specifically Erlang—distributed as the sum of IID exponentials.

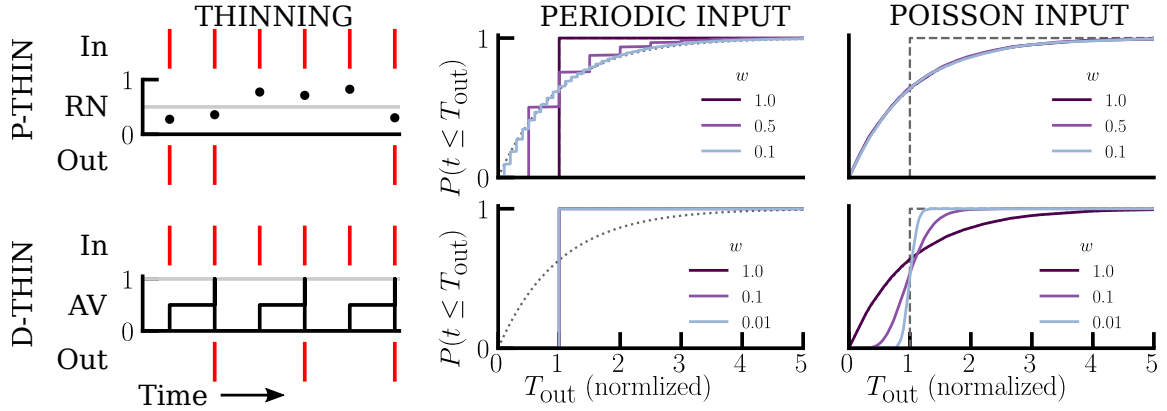


Figure 2.3: Probabilistic (P) vs Deterministic (D) Thinning.

P-THIN: Each input spike draws a random number (RN) from a uniform distribution in $\{0:1\}$. If RN is less than the weight (0.5), then the spike is forwarded out. D-THIN: Each input spike increments an accumulated value (AV) by the weight (0.5). If AV exceeds a threshold (1), then the spike is forwarded out and AV is decremented by the threshold. P-THIN PERIODIC vs D-THIN POISSON: As the weight decreases, cumulative distribution functions (CDFs) of empirical output ISIs (T_{out}), sampled and plotted as in Fig. 2.1, transition from degenerate (dashed lines) to exponential (dotted lines) versus exponential to degenerate. P-THIN POISSON and D-THIN PERIODIC: CDFs do not change with the weight.

Table 2.1: Renewal Process ISI CV and Synapse SNR

Renewal Process	ISI CV	Synapse SNR
Poisson	1	$\sqrt{2\lambda\tau}$
periodic	0	$\sqrt{\frac{2\lambda\tau}{\coth(\frac{1}{2\lambda\tau}) - 2\lambda\tau}}$
p-thinned periodic	$\sqrt{1-w}$	$\sqrt{\frac{2\lambda\tau}{1-p+p\coth(\frac{p}{2\lambda\tau}) - 2\lambda\tau}}$
d-thinned Poisson	\sqrt{w}	$\sqrt{\frac{2\lambda\tau}{\frac{(1+k\lambda\tau)^k + (k\lambda\tau)^k}{(1+k\lambda\tau)^k - (k\lambda\tau)^k} - 2\lambda\tau}}$

knowing the next future spike-time. The second requirement limits the number of ISI samples available to the time window times the spike-rate.

We consider methods that obtain an estimate, X , of the rate by computing a running average of spikes, weighting a spike that happened T_j seconds in the past by $h(T_j)$. That is,

$$X = \sum_{j=0}^{\infty} h(T_j) \quad (2.1)$$

defined at a point in time after an infinite number of spikes from a fixed-rate spike-train have arrived, which allows us to ignore the initial transient and focus on the steady-state mean. We quantify the estimate's quality using its signal-to-noise ratio:

$$\text{SNR}(X) = \mathbf{E}[X] / \sqrt{\text{Var}(X)} \quad (2.2)$$

We consider the simplest implementation of $h(T_j)$: a first-order, low-pass filter (LPF) driven by a train of impulses, $s(t) = \sum_j \delta(t - t_j)$, where t_j is the j th spike's time. That is,

$$\tau \frac{dx}{dt} = -x + s \quad (2.3)$$

where $x(t)$ is the LPF's state (and output) and τ is its time-constant. Upon receiving the j th spike, $\delta(t - t_j)$, $x(t)$ jumps by $1/\tau$ and then decays at the rate $1/\tau$ —only $x(t)$'s *mean* value settles to the spike rate—so larger τ result in cleaner but slower estimates. In this case, $h(t) = \frac{1}{\tau} e^{-t/\tau} u(t)$, which yields weights that decay by a factor of e every τ seconds; $u(t)$ is the unit-step function. In addition to depending on τ , the rate-estimate's variability also depends on the incoming ISIs' variability (Fig. 2.2).

We derive $\text{SNR}(X)$ for low-pass filtered Poisson, periodic, p-thinned periodic, and d-thinned Poisson spike-trains (Table 2.1 & Fig. 2.2). We do not consider p-thinned Poisson nor d-thinned periodic spike-trains because their statistics do not change with the weight (see Fig. 2.1). It suffices to find the first and second moments of X for each process since $\text{Var}(X) = \mathbf{E}[X^2] - \mathbf{E}[X]^2$. To express T_j in terms of IID random variables, we define $\Delta T_j = T_j - T_{j-1}$ for $j > 0$ so that $T_j = \sum_{n=0}^j \Delta T_n$. Note that $\Delta T_0 = T_0$, the elapsed time from the most recent spike, is not an ISI and so is not generally drawn from the same distribution as the ISIs. Substituting into (2.1) yields

$$\begin{aligned}
X &= \sum_{j=0}^{\infty} h \left(\sum_{l=0}^j \Delta T_l \right) \\
&= \frac{1}{\tau} e^{-\frac{\Delta T_0}{\tau}} + \frac{1}{\tau} e^{-\frac{\Delta T_0 + \Delta T_1}{\tau}} + \frac{1}{\tau} e^{-\frac{\Delta T_0 + \Delta T_1 + \Delta T_2}{\tau}} + \dots \\
&= \frac{1}{\tau} \sum_{j=0}^{\infty} \prod_{l=0}^j e^{-\frac{\Delta T_l}{\tau}}
\end{aligned} \tag{2.4}$$

Therefore,

$$\begin{aligned}
X^2 &= \left(\frac{1}{\tau} \sum_{j=0}^{\infty} \prod_{l=0}^j e^{-\frac{\Delta T_l}{\tau}} \right)^2 \\
&= \frac{1}{\tau^2} \sum_{j=0}^{\infty} \prod_{l=0}^j e^{-\frac{2\Delta T_l}{\tau}} + \frac{2}{\tau^2} \sum_{j=0}^{\infty} \sum_{l=j+1}^{\infty} \prod_{m=0}^j e^{-\frac{2\Delta T_m}{\tau}} \prod_{n=j+1}^l e^{-\frac{\Delta T_n}{\tau}} \quad \text{breaking into diagonal and cross terms,} \\
&= \frac{1}{\tau^2} \sum_{j=0}^{\infty} \prod_{l=0}^j e^{-\frac{2\Delta T_l}{\tau}} + \frac{2}{\tau^2} \sum_{j=0}^{\infty} \prod_{m=0}^j e^{-\frac{2\Delta T_m}{\tau}} \sum_{l=j+1}^{\infty} \prod_{n=j+1}^l e^{-\frac{\Delta T_n}{\tau}}
\end{aligned}$$

Assuming ΔT_0 is identically distributed to the ISIs (i.e., that the ISIs are memoryless as in the Poisson process), from (2.4),

$$\begin{aligned}
\mathbf{E}[X] &= \mathbf{E} \left[\frac{1}{\tau} \sum_{j=0}^{\infty} \prod_{l=0}^j e^{-\frac{\Delta T_l}{\tau}} \right] \\
&= \frac{1}{\tau} \sum_{j=0}^{\infty} \mathbf{E} \left[\prod_{l=0}^j e^{-\frac{\Delta T_l}{\tau}} \right] \\
&= \frac{1}{\tau} \sum_{j=0}^{\infty} \mathbf{E} \left[e^{-\frac{\Delta T}{\tau}} \right]^{j+1} \quad \Delta T_j \text{ are IID} \\
&= \frac{1}{\tau} \sum_{j=1}^{\infty} \varphi_{\Delta T} \left(\frac{j}{\tau} \right)
\end{aligned} \tag{2.5}$$

noting that $\mathbf{E} \left[e^{-\Delta T/\tau} \right] = \varphi_{\Delta T}(s)|_{s=i/\tau}$, the characteristic function of ΔT evaluated at i/τ .⁵ From (2.5),

⁵The characteristic function is related to the Fourier transform: $\varphi_X(s) = \mathcal{F}_{f_X} \left(\frac{-s}{2\pi} \right)$ since $\varphi_X(s) = \int e^{isx} f_X(x) dx$ for random variable X with probability density function $f_X(x)$ and $\mathcal{F}_g(\omega) = \int e^{-2\pi i x \omega} g(x) dx$ for function $g(x)$.




Figure 2.4: Renewal Process-Driven Synapse

TOP: Spikes generated by Poisson, p-thinned periodic ($p = 0.7$), d-thinned Poisson ($k = 20$), and periodic point-processes. The underlying spike rates are matched at $\lambda = 5$ spikes per τ . BOTTOM: Synaptic response to spikes is governed by first-order dynamics (2.3). The mean output converges to λ , and, in this example, the variability decreases as we go from Poisson to periodic (ordered as above).

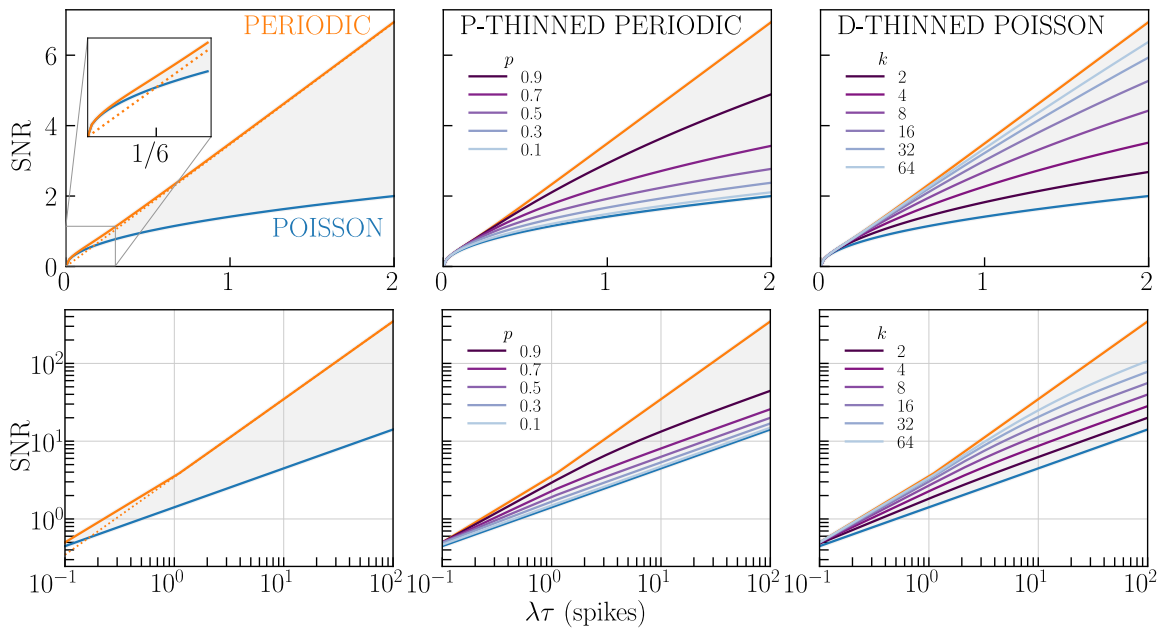


Figure 2.5: Rate-Estimate SNR.

PERIODIC & POISSON: SNR scales linearly and as the square-root, respectively, with $\lambda\tau$, the number of spikes per synaptic time-constant τ , for $\lambda\tau > 1/6$. Below this number (inset), periodic switches from linear (dotted line) to square-root. P-THINNED PERIODIC: SNR transitions from periodic to Poisson behavior (shaded region) as p decreases. D-THINNED POISSON: SNR transitions from Poisson to periodic behavior as k increases.

$$\begin{aligned}
& \mathbf{E}[X^2] \\
&= \mathbf{E} \left[\frac{1}{\tau^2} \sum_{j=0}^{\infty} \prod_{l=0}^j e^{-\frac{2\Delta T_l}{\tau}} + \frac{2}{\tau^2} \sum_{j=0}^{\infty} \prod_{m=0}^j e^{-\frac{2\Delta T_m}{\tau}} \sum_{l=j+1}^{\infty} \prod_{n=j+1}^l e^{-\frac{\Delta T_n}{\tau}} \right] \\
&= \frac{1}{\tau^2} \sum_{j=1}^{\infty} \mathbf{E} \left[e^{-\frac{2\Delta T}{\tau}} \right]^j + \frac{2}{\tau^2} \sum_{j=1}^{\infty} \mathbf{E} \left[e^{-\frac{2\Delta T}{\tau}} \right]^j \sum_{l=j+1}^{\infty} \mathbf{E} \left[e^{-\frac{\Delta T}{\tau}} \right]^{l-j} \\
&\text{substituting } m = l - j \\
&= \frac{1}{\tau^2} \sum_{j=1}^{\infty} \mathbf{E} \left[e^{-\frac{2\Delta T}{\tau}} \right]^j + \frac{2}{\tau^2} \sum_{j=1}^{\infty} \mathbf{E} \left[e^{-\frac{2\Delta T}{\tau}} \right]^j \sum_{m=1}^{\infty} \mathbf{E} \left[e^{-\frac{\Delta T}{\tau}} \right]^m \\
&= \frac{1}{\tau^2} \sum_{j=1}^{\infty} \varphi_{\Delta T} \left(\frac{2i}{\tau} \right)^j \left(1 + 2 \sum_{m=1}^{\infty} \varphi_{\Delta T} \left(\frac{i}{\tau} \right)^m \right) \tag{2.6}
\end{aligned}$$

When ΔT_0 is not identically distributed to the ISIs (i.e., the ISI distribution has memory), we split out ΔT_0 from (2.4) and (2.5),

$$\begin{aligned}
X &= \frac{1}{\tau} \sum_{j=0}^{\infty} \prod_{l=0}^j e^{-\frac{\Delta T_l}{\tau}} = \frac{1}{\tau} e^{-\frac{\Delta T_0}{\tau}} \left(1 + \sum_{j=1}^{\infty} \prod_{l=1}^j e^{-\frac{\Delta T_l}{\tau}} \right) \\
X^2 &= \frac{1}{\tau^2} e^{-\frac{2\Delta T_0}{\tau}} \left(1 + 2 \sum_{j=1}^{\infty} \prod_{l=1}^j e^{-\frac{\Delta T_l}{\tau}} + \sum_{j=1}^{\infty} \prod_{l=1}^j e^{-\frac{2\Delta T_l}{\tau}} + 2 \sum_{j=1}^{\infty} \prod_{m=1}^j e^{-\frac{2\Delta T_m}{\tau}} \sum_{l=j+1}^{\infty} \prod_{n=j+1}^l e^{-\frac{\Delta T_n}{\tau}} \right)
\end{aligned}$$

which results in,

$$\begin{aligned}
\mathbf{E}[X] &= \frac{1}{\tau} \mathbf{E} \left[e^{-\frac{\Delta T_0}{\tau}} \right] \left(1 + \sum_{j=1}^{\infty} \mathbf{E} \left[e^{-\frac{\Delta T}{\tau}} \right]^j \right) \\
&= \frac{1}{\tau} \varphi_{\Delta T_0} \left(\frac{i}{\tau} \right) \left(1 + \sum_{j=1}^{\infty} \varphi_{\Delta T} \left(\frac{i}{\tau} \right)^j \right)
\end{aligned} \tag{2.7}$$

$$\begin{aligned}
\mathbf{E}[X^2] &= \frac{1}{\tau^2} \mathbf{E} \left[e^{-\frac{2\Delta T_0}{\tau}} \right] \left(1 + 2 \sum_{j=1}^{\infty} \mathbf{E} \left[e^{-\frac{\Delta T}{\tau}} \right]^j \right. \\
&\quad \left. + \sum_{j=1}^{\infty} \mathbf{E} \left[e^{-\frac{2\Delta T}{\tau}} \right]^j + 2 \sum_{j=1}^{\infty} \mathbf{E} \left[e^{-\frac{2\Delta T}{\tau}} \right]^j \sum_{l=1}^{\infty} \mathbf{E} \left[e^{-\frac{\Delta T}{\tau}} \right]^l \right) \\
&= \frac{1}{\tau^2} \varphi_{\Delta T_0} \left(\frac{2i}{\tau} \right) \left(1 + 2 \sum_{j=1}^{\infty} \varphi_{\Delta T} \left(\frac{i}{\tau} \right)^j \right. \\
&\quad \left. + \sum_{j=1}^{\infty} \varphi_{\Delta T} \left(\frac{2i}{\tau} \right)^j + 2 \sum_{j=1}^{\infty} \varphi_{\Delta T} \left(\frac{2i}{\tau} \right)^j \sum_{l=1}^{\infty} \varphi_{\Delta T} \left(\frac{i}{\tau} \right)^l \right)
\end{aligned} \tag{2.8}$$

2.2.1 Poisson Process

With Poisson input, SNR scales with the square-root of the number of spikes per synaptic time-constant. For an exponential distribution with rate parameter λ , $\varphi_{\Delta T}(s) = \frac{\lambda}{\lambda - is}$. Using (2.2) and $s = i/\tau$,

$$\mathbf{E}[X] = \frac{1}{\tau} \sum_{j=1}^{\infty} \left(\frac{\lambda\tau}{\lambda\tau + 1} \right)^j = \frac{1}{\tau} \frac{\frac{\lambda\tau}{\lambda\tau + 1}}{1 - \frac{\lambda\tau}{\lambda\tau + 1}} = \lambda$$

That is, the expected rate-estimate matches the input rate. Using (2.6) and $s = 2i/\tau$,

$$\begin{aligned}
\mathbf{E}[X^2] &= \frac{1}{\tau^2} \sum_{j=1}^{\infty} \left(\frac{\lambda\tau}{\lambda\tau + 2} \right)^j \left(1 + 2 \sum_{m=1}^{\infty} \left(\frac{\lambda\tau}{\lambda\tau + 1} \right)^m \right) \\
&= \frac{1}{\tau^2} \frac{\lambda\tau}{2} (1 + 2\lambda\tau) = \frac{\lambda}{2\tau} + \lambda^2
\end{aligned}$$

Therefore, from (2.2),

$$\begin{aligned}
\text{Var}(X) &= \lambda/2\tau \\
\text{SNR}_{\text{poi}}(X) &= \sqrt{2\lambda\tau}
\end{aligned}$$

(Table 2.1, Poisson & Fig. 2.2, POISSON), which we confirm with numerical simulations (Fig. 2.2.1,




Figure 2.6: Theory vs numerical simulations.

Empirical SNR measured and averaged over 10,000 spike-train realizations with $30\lambda\tau$ spikes each. The results (circles) match the corresponding theoretical SNR (solid lines).

POISSON). Note that the relevant quantity is $\lambda\tau$, the number of spikes arriving within a time constant, not the spike rate λ or the synaptic time constant τ individually.

2.2.2 Periodic Process

With periodic input, SNR scales linearly with $\lambda\tau$. While the ISIs are equal, ΔT_0 is uniformly distributed between 0 and $1/\lambda$, so that $\varphi_{\Delta T_0}(s) = \frac{e^{is/\lambda} - 1}{is/\lambda}$. From T_0 into the past, every subsequent spike is $1/\lambda$ further from the present, so from (2.7) and (2.8),

$$\begin{aligned} \mathbf{E}[X] &= \frac{1}{\tau} \frac{e^{-\frac{1}{\lambda\tau}} - 1}{-1/\lambda\tau} \left(1 + \sum_{j=1}^{\infty} e^{-\frac{j}{\lambda\tau}} \right) = \lambda \\ \mathbf{E}[X^2] &= \frac{1}{\tau^2} \frac{e^{-\frac{2}{\lambda\tau}} - 1}{-2/\lambda\tau} \left(1 + 2 \sum_{j=1}^{\infty} e^{-\frac{j}{\lambda\tau}} + \sum_{j=1}^{\infty} e^{-\frac{2j}{\lambda\tau}} + 2 \sum_{j=1}^{\infty} e^{-\frac{2j}{\lambda\tau}} \sum_{l=1}^{\infty} e^{-\frac{l}{\lambda\tau}} \right) \\ &= \frac{\lambda}{2\tau} \frac{1 + e^{-\frac{1}{\lambda\tau}}}{1 - e^{-\frac{1}{\lambda\tau}}} = \frac{\lambda}{2\tau} \coth\left(\frac{1}{2\lambda\tau}\right) \end{aligned}$$

Therefore, from (2.2),

$$\begin{aligned} \text{Var}(X) &= \frac{\lambda}{2\tau} \coth\left(\frac{1}{2\lambda\tau}\right) - \lambda^2 = \frac{\lambda^2}{2\lambda\tau} \left(\coth\left(\frac{1}{2\lambda\tau}\right) - 2\lambda\tau \right) \\ \text{SNR}_{\text{per}}(X) &= \sqrt{\frac{2\lambda\tau}{\coth\left(\frac{1}{2\lambda\tau}\right) - 2\lambda\tau}} \end{aligned}$$

(Table 2.1, periodic & Fig. 2.2, PERIODIC), which we confirm with numerical simulations (Fig. 2.2.1, PERIODIC).

When $\lambda\tau$ is small and large, $\text{SNR}_{\text{per}}(X)$ scales as the square-root and linearly, respectively. As $\lambda\tau$ approaches zero, $\text{SNR}_{\text{per}}(X)$ scales similarly to Poisson SNR ($\sqrt{\lambda\tau}$):

$$\lim_{\lambda\tau \rightarrow 0} \text{SNR}_{\text{per}}(X) = \lim_{\lambda\tau \rightarrow 0} \sqrt{\frac{2\lambda\tau}{\coth\left(\frac{1}{2\lambda\tau}\right) - 2\lambda\tau}} = \sqrt{2\lambda\tau}$$

(i.e., periodic and Poisson inputs are indistinguishable from an SNR perspective). For high $\lambda\tau$, the limit is not as straightforward— $\lim_{\lambda\tau \rightarrow \infty} \coth(1/2\lambda\tau)$ is not defined—but we find that $\text{SNR}_{\text{per}} \rightarrow 2\sqrt{3}\lambda\tau$ (see Appendix A.1). That is, SNR scales linearly with $\lambda\tau$ at high $\lambda\tau$. These two approximations intersect at $\lambda\tau = 1/6$ (Fig. 2.2, PERIODIC, inset). The low $\lambda\tau$ approximation is within 10% and 1% of the actual SNR for $\lambda\tau < 0.0950$ and $\lambda\tau < 0.00995$, respectively. The high $\lambda\tau$ approximation is within 10% and 1% of the actual SNR for $\lambda\tau > 0.253$ and $\lambda\tau > 0.902$, respectively. In summary,

$$\text{SNR}_{\text{per}}(\lambda\tau) \approx \begin{cases} \sqrt{2\lambda\tau} & \lambda\tau \leq 1/6 \\ 2\sqrt{3}\lambda\tau & \lambda\tau > 1/6 \end{cases}$$

2.2.3 p-Thinned Periodic Process

With p-thinned periodic input, SNR transitions from periodic to Poisson with decreasing p ($= w$). Instead of using (2.7) and (2.8) directly, we modify X by introducing a random (indicator) variable

$$I_j = \begin{cases} 1 & \text{with probability } p \\ 0 & \text{with probability } 1 - p \end{cases}$$

to indicate whether the j th pre-thinned spike is kept or not:

$$X = \frac{1}{\tau} e^{-\frac{\Delta T_0}{\tau}} \left(I_0 + \sum_{j=1}^{\infty} e^{-\frac{j}{\lambda_p \tau}} I_j \right)$$

where ΔT_0 is the time to the first, pre-thinned spike and λ_p is the rate of the pre-thinned spike-train. Only spikes whose corresponding indicator is 1 will affect the current state. Therefore,

$$\begin{aligned}
\mathbf{E}[X] &= \mathbf{E} \left[\frac{1}{\tau} e^{-\frac{\Delta T_0}{\tau}} \left(I_0 + \sum_{j=1}^{\infty} e^{\frac{-j}{\lambda_p \tau}} I_j \right) \right] \\
&= \frac{1}{\tau} \mathbf{E} \left[e^{-\frac{\Delta T_0}{\tau}} \right] \left(\mathbf{E}[I_0] + \sum_{j=1}^{\infty} e^{\frac{-j}{\lambda_p \tau}} \mathbf{E}[I_j] \right) \quad (T_0, I_j \text{ are independent}) \\
&= \frac{1}{\tau} \frac{1 - e^{-\frac{1}{\lambda_p \tau}}}{1/\lambda_p \tau} \left(p + \sum_{j=0}^{\infty} e^{\frac{-j}{\lambda_p \tau}} p \right) = p \lambda_p = \lambda
\end{aligned}$$

For $\mathbf{E}[X^2]$,

$$\begin{aligned}
\mathbf{E}[X^2] &= \mathbf{E} \left[\left(\frac{1}{\tau} e^{-\frac{\Delta T_0}{\tau}} \left(I_0 + \sum_{j=1}^{\infty} e^{\frac{-j}{\lambda_p \tau}} I_j \right) \right)^2 \right] \\
&= \frac{\lambda_p}{2\tau} \left(1 - e^{-2/\lambda_p \tau} \right) \mathbf{E} \left[\left(\sum_{j=0}^{\infty} e^{\frac{-j}{\lambda_p \tau}} I_j \right)^2 \right]
\end{aligned}$$

The expectation can be broken into diagonal and cross terms.

$$\begin{aligned}
\mathbf{E} \left[\left(\sum_{j=0}^{\infty} e^{\frac{-j}{\lambda_p \tau}} I_j \right)^2 \right] &= \mathbf{E} \left[\sum_{j=0}^{\infty} e^{\frac{-2j}{\lambda_p \tau}} I_j^2 + \sum_{j=0, j \neq l}^{\infty} \sum_{l=0}^{\infty} e^{\frac{-(j+l)}{\lambda_p \tau}} I_j I_l \right] \\
&= \sum_{j=0}^{\infty} e^{\frac{-2j}{\lambda_p \tau}} \mathbf{E}[I_j^2] + \sum_{j=0, j \neq l}^{\infty} \sum_{l=0}^{\infty} e^{\frac{-(j+l)}{\lambda_p \tau}} \mathbf{E}[I_j I_l] \\
&= \sum_{j=0}^{\infty} e^{\frac{-jk}{\lambda_p \tau}} p + \sum_{j=0, j \neq l}^{\infty} \sum_{l=0}^{\infty} e^{\frac{-(j+l)}{\lambda_p \tau}} p^2 \quad (I_j, I_l \text{ independent for } j \neq l) \\
&= \frac{p}{1 - e^{-\frac{2}{\lambda_p \tau}}} + \frac{p^2}{\left(1 - e^{-\frac{1}{\lambda_p \tau}} \right)^2} - \frac{p^2}{1 - e^{-\frac{2}{\lambda_p \tau}}}
\end{aligned}$$

After substitution back into $\mathbf{E}[X^2]$,

$$\mathbf{E}[X^2] = \frac{\lambda p}{2\tau} \left(1 - p + p \frac{1 + e^{-1/\lambda p \tau}}{1 - e^{-1/\lambda p \tau}} \right) = \frac{\lambda}{2\tau} \left(1 - p + p \coth \left(\frac{p}{2\lambda\tau} \right) \right)$$

Therefore,

$$\begin{aligned} \text{Var}(X) &= \frac{\lambda}{2\tau} \left(1 - p + p \coth \left(\frac{p}{2\lambda\tau} \right) - 2\lambda\tau \right) \\ \text{SNR}_{\text{pthin}}(X) &= \sqrt{\frac{2\lambda\tau}{1 - p + p \coth \left(\frac{p}{2\lambda\tau} \right) - 2\lambda\tau}} \end{aligned}$$

(Table 2.1, p-thinned periodic & Fig. 2.2, P-THIN.), which we confirm with numerical simulations (Fig. 2.2.1, P-THIN.). As $p \rightarrow 1$, $\text{SNR}_{\text{pthin}} \rightarrow \text{SNR}_{\text{per}}$ (i.e., no spikes are dropped), and as $p \rightarrow 0$, $\text{SNR}_{\text{pthin}} \rightarrow \text{SNR}_{\text{poi}}$ (see Appendix A.2). Therefore,

$$\text{SNR}_{\text{pthin}}(X) \approx \begin{cases} \text{SNR}_{\text{poi}}(X) & \text{at low } p \\ \text{SNR}_{\text{per}}(X) & \text{at high } p \end{cases}$$

2.2.4 d-Thinned Poisson Process

With d-thinned Poisson input, SNR transitions from Poisson to periodic behavior with increasing k ($= 1/w$). The d-thinned Poisson process is not memoryless, necessitating the use of (2.7) and (2.8) and considering ΔT_0 and ΔT separately. Expressing $\Delta T_0 = U\Delta T^*$, where U selects uniformly within random interval ΔT^* (i.e., $U \sim \text{Uniform}(0, 1)$), which is not identically distributed to ΔT , an example of the *inspection paradox* [40]. Although we select a point in time without bias, our point in time is biased towards falling within larger ISIs simply because they are larger. Accounting for this bias, ΔT^* 's probability density function (PDF) is $f_{\Delta T^*}(t) = \frac{t f_{\Delta T}(t)}{\mathbf{E}[\Delta T]}$. Scaling density function $f_{\Delta T}(t)$ by t captures this bias and dividing by $\mathbf{E}[\Delta T]$ renormalizes the scaled density to a valid probability density (i.e. $\int f_{\Delta T^*} = \int t f_{\Delta T}(t) dt = \mathbf{E}[\Delta T]$). As a result,

$$\begin{aligned} \mathbf{E} \left[e^{-\frac{\Delta T_0}{\tau}} \right] &= \lambda\tau \left(1 - \left(\frac{k\lambda\tau}{1 + k\lambda\tau} \right)^k \right) \\ \mathbf{E} \left[e^{-\frac{2\Delta T_0}{\tau}} \right] &= \frac{\lambda\tau}{2} \left(1 - \left(\frac{k\lambda\tau}{2 + k\lambda\tau} \right)^k \right) \end{aligned}$$

and since $\varphi_{\Delta T}(s) = \left(\frac{k\lambda}{k\lambda - is} \right)^{k\lambda}$, (2.7) and (2.8) produce

$$\begin{aligned} \mathbf{E}[X] &= \frac{1}{\tau} \lambda \tau \left(1 - \left(\frac{k\lambda\tau}{1+k\lambda\tau} \right)^k \right) \sum_{j=0}^{\infty} \left(\frac{k\lambda\tau}{k\lambda\tau+1} \right)^{k\lambda j} = \lambda \\ \mathbf{E}[X^2] &= \frac{1}{\tau^2} \frac{\lambda\tau}{2} \left(1 - \left(\frac{k\lambda\tau}{1+k\lambda\tau} \right)^k \right) \sum_{j=0}^{\infty} \left(\frac{k\lambda\tau}{2+k\lambda\tau} \right)^{kj} \left(1 + 2 \sum_{l=1}^{\infty} \left(\frac{k\lambda\tau}{1+k\lambda\tau} \right)^{kl} \right) \\ &= \frac{\lambda}{2\tau} \frac{1 + \left(\frac{k\lambda\tau}{1+k\lambda\tau} \right)^k}{1 - \left(\frac{k\lambda\tau}{1+k\lambda\tau} \right)^k} = \frac{\lambda}{2\tau} \frac{(1+k\lambda\tau)^k + (k\lambda\tau)^k}{(1+k\lambda\tau)^k - (k\lambda\tau)^k} \end{aligned}$$

therefore,

$$\begin{aligned} \text{Var}(X) &= \frac{\lambda}{2\tau} \left(\frac{(1+k\lambda\tau)^k + (k\lambda\tau)^k}{(1+k\lambda\tau)^k - (k\lambda\tau)^k} - \frac{2k\lambda\tau}{k} \right) \\ \text{SNR}_{\text{dthin}}(X) &= \sqrt{\frac{2\lambda\tau}{\frac{(1+k\lambda\tau)^k + (k\lambda\tau)^k}{(1+k\lambda\tau)^k - (k\lambda\tau)^k} - 2\lambda\tau}} \end{aligned}$$

(Table 2.1, d-thinned Poisson & Fig. 2.2, D-THIN.), which we confirm with numerical simulations (Fig. 2.2.1, D-THIN.). Alternatively (see Appendix A.3), $\text{SNR}_{\text{dthin}} = h \text{SNR}_{\text{poi}}(X)$, where

$$h = \sqrt{\frac{1 + k^2\lambda\tau + \dots + \frac{1}{2}k^{k-1}(k-1)(\lambda\tau)^{k-2} + k^k(\lambda\tau)^{k-1}}{1 + (k^2-2)\lambda\tau + \dots + \frac{1}{6}k^{k-2}(k-1)(k+4)(\lambda\tau)^{k-2} + k^{k-1}(\lambda\tau)^{k-1}}}$$

At high $\lambda\tau$, d-thinning preserves the input's information, while at low $\lambda\tau$, a d-thinned Poisson spike-train is indistinguishable from a Poisson spike-train with the same rate:

$$h \underset{\lambda\tau \rightarrow \infty}{=} \sqrt{\frac{k^k(\lambda\tau)^{k-1}}{k^{k-1}(\lambda\tau)^{k-1}}} = \sqrt{k} \quad \text{and} \quad h \underset{\lambda\tau \rightarrow 0}{=} 1$$

Therefore, $\text{SNR}_{\text{dthin}}(X) \underset{\lambda\tau \rightarrow \infty}{=} \sqrt{k} \text{SNR}_{\text{poi}}(X) = \sqrt{2k\lambda\tau}$, which is the input spike-train's SNR, and $\text{SNR}_{\text{dthin}}(X) \underset{\lambda\tau \rightarrow 0}{=} \text{SNR}_{\text{poi}}(X)$.

At $k = 1$, $\text{SNR}_{\text{dthin}}(X) = \text{SNR}_{\text{poi}}(X)$ (i.e., all Poissonian input spikes are output). As $k \rightarrow \infty$, $\text{SNR}_{\text{dthin}}(X) \rightarrow \text{SNR}_{\text{per}}(X)$. However, $k \rightarrow \infty$ means processing many input, pre-thinned spikes for each output spike. When accounting for the energy costs associated with processing spikes, $\text{SNR}_{\text{dthin}}$ scales as $\sqrt[3]{(3c/2b)(\lambda\tau)^2}$ where c is the fanout of the thinned spike-train and b is the cost of processing each pre-thinned spike. After optimizing for costs, we find that $\text{SNR}_{\text{dthin}}$ scales as $(\lambda\tau)^{3/4}$.




Figure 2.7: SNR Approximation.

Relative position of $\text{SNR}_{\text{pthin}}$ and $\text{SNR}_{\text{dthin}}$ between SNR_{poi} and SNR_{per} computed via (2.9). P-THINNED PERIODIC: For low $\lambda\tau$, $\text{SNR}_{\text{pthin}}$ is more closely approximated by SNR_{per} than SNR_{poi} (curves are near $r_{\text{proc}} = 1$), but as $\lambda\tau$ increases or p decreases, $\text{SNR}_{\text{pthin}}$ becomes more closely approximated by SNR_{poi} (curves approach $r_{\text{proc}} = 0$). D-THINNED POISSON: For low $\lambda\tau$, $\text{SNR}_{\text{dthin}}$ is more closely approximated by SNR_{per} , but as $\lambda\tau$ increases, $\text{SNR}_{\text{dthin}}$ becomes more closely approximated by SNR_{poi} . However, as k increases (i.e., w decreases), $\text{SNR}_{\text{dthin}}$ moves closer to SNR_{per} , in contrast to $\text{SNR}_{\text{pthin}}$.

2.2.5 Approximation Quality

While p-thinned periodic approaches Poisson and d-thinned Poisson approaches periodic as $w \rightarrow 0$, the convergence depends on $\lambda\tau$. To quantify this, consider

$$r_{\text{proc}}(\lambda\tau, w) = \frac{\text{SNR}_{\text{proc}}(X(\lambda\tau, w)) - \text{SNR}_{\text{poi}}(X(\lambda\tau))}{\text{SNR}_{\text{per}}(X(\lambda\tau)) - \text{SNR}_{\text{poi}}(X(\lambda\tau))} \quad (2.9)$$

which measures where the SNR of proc (either pthin or dthin) lies between the Poisson and periodic limits at a given $\lambda\tau$ (Fig. 2.2.5). At $r_{\text{proc}} = 0$, $\text{SNR}_{\text{proc}} = \text{SNR}_{\text{poi}}$, at $r_{\text{proc}} = 1$, $\text{SNR}_{\text{proc}} = \text{SNR}_{\text{per}}$, and at $r_{\text{proc}} = 0.5$, SNR_{proc} is halfway between SNR_{per} and SNR_{poi} .

2.3 Optimizing for Power

The advantages of weighted-summation by d over p-thinning ultimately translate into lower power consumption. Power is given by $P = c_{\text{in}}\lambda_{\text{in}} + c_{\text{out}}\lambda_{\text{out}}$ where $c_{\text{in(ou)}}$ and $\lambda_{\text{in(ou)}}$ are the energy cost per input (output) spike and input (output) spike rate, respectively. c_{in} for p-thinning is set as each

input spike incurs a soma communication, weight lookup, random sample, and comparison (between the weight and the sample). c_{in} for d-thinning is set as each input spike incurs a soma communication, weight lookup, accumulation lookup, addition (between the weight and accumulation state), and thresholding. c_{out} is the same for both p and d-thinning: each output spike incurs a target lookup and set of synapse communications (i.e. the spike fans out).

Power for summing-and-weighting as implemented by merging and p-thinning is minimized when $p = 1$, (i.e., not weighting at all). Due to the Poisson statistics ($\text{SNR}_{\text{poi}}(X) = \sqrt{2\lambda\tau}$),

$$P_{\text{pthin}} = c_{\text{in}} \frac{\text{SNR}_{\text{tgt}}^2}{2\tau p} + c_{\text{out}} \frac{\text{SNR}_{\text{tgt}}^2}{2\tau} = \frac{1}{2\tau} \left(\frac{c_{\text{in}}}{p} + c_{\text{out}} \right) \text{SNR}_{\text{tgt}}^2$$

for a target SNR (SNR_{tgt}) since $\lambda_{\text{out}} = \text{SNR}_{\text{tgt}}^2/2\tau$ and $\lambda_{\text{in}} = \lambda_{\text{out}}/p = \text{SNR}_{\text{tgt}}^2/2\tau p$. However, regardless of c_{in} or c_{out} , power consumption is minimized when $p = 1$ since p is restricted to $\{0, 1\}$.

In contrast, d-thinning incentives weighting for sufficient $\lambda\tau$. Since $\text{SNR}_{\text{dthin}}(X) \underset{\lambda\tau \rightarrow 0}{=} \sqrt{2\lambda\tau}$, and $\text{SNR}_{\text{dthin}}(X) \underset{\lambda\tau \rightarrow \infty}{=} \sqrt{2k\lambda\tau}$,

$$\begin{aligned} P_{\text{dthin}} \underset{\lambda\tau \rightarrow 0}{=} c_{\text{in}} \frac{\text{SNR}_{\text{tgt}}^2}{2\tau} k + c_{\text{out}} \frac{\text{SNR}_{\text{tgt}}^2}{2\tau} &= \frac{1}{2\tau} (c_{\text{in}}k + c_{\text{out}}) \text{SNR}_{\text{tgt}}^2 \\ P_{\text{dthin}} \underset{\lambda\tau \rightarrow \infty}{=} c_{\text{in}} \frac{\text{SNR}_{\text{tgt}}^2}{2\tau} + c_{\text{out}} \frac{\text{SNR}_{\text{tgt}}^2}{2k\tau} &= \frac{1}{2\tau} \left(c_{\text{in}} + \frac{c_{\text{out}}}{k} \right) \text{SNR}_{\text{tgt}}^2 \end{aligned}$$

At low $\lambda\tau$, d-thinning power is minimized by $k = 1$ (i.e., not weighting at all) as with p-thinning. However, at high $\lambda\tau$, power decreases with increasing k (i.e., weighting improves power consumption). Between the extremes, the minimum-power k switches from $k = 1$ to $k > 1$ depending on c_{in} and c_{out} (Fig. 2.3).

2.4 Discussion

Although neuromorphic engineers look to biology for inspiration and biological neurons show trial-to-trial variability often modeled by Poisson statistics, Poisson spike-trains are poor at communicating information about their underlying rate. Therefore, injecting noise (via p-thinning) simply to mimic Poisson statistics observed in biological data (e.g., trial-to-trial variability of synaptic transmission) is an ill-founded pursuit (outside the corresponding biological context [8, 16, 43]).

We might seem to be claiming to violate the data-processing inequality as the accumulator "cleans up" a Poisson spike-train, but that's only an artifact of comparing the SNRs between spike trains resulting from p-thinning and d-thinning. If we instead compared the p-thinning and d-thinning SNR to their, respective, pre-thinned SNRs, we would see that both procedures decrease SNR (i.e.,




Figure 2.8: d-Thinned Poisson Power.

Set $c_{\text{in}} = \alpha$ and $c_{\text{out}} = 1 - \alpha$ so that $0 \leq \alpha \leq 1$ sweeps linearly between relative output and input energy costs. LEFT: $\alpha = 0.5$ (i.e., input and output spikes cost the same) iso-power contours. The minimum power curve passes through each contour's peak. At low $\lambda\tau$, the minimum power curve follows $\text{SNR}_{\text{poi}}(X)$, but as $\lambda\tau$ increases, it lifts above $\text{SNR}_{\text{poi}}(X)$ as weighting lowers the power requirements for a given SNR. CENTER and RIGHT: Minimum power curves vary with α . The relative input-to-output cost determines when it becomes more efficient to start weighting to reduce the output spike-rate.

we respect the data-processing inequality). Our claim is on the relative drop in SNR—d-thinning produces a smaller relative decrease in SNR than p-thinning. Therefore, given a choice of p-thinning or d-thinning for weighting and a choice of summing before weighting or weighting before summing, summing and then d-thinning is the clear favorite. It results in $\text{SNR}_{\text{dthin}}$ scaling with the $3/4$ power of $\lambda\tau$ instead of as the square root—the case of p-thinning or (d-thinning) before summing. When Poisson statistics arise, they must arise from with contextual mechanisms embedded within larger frameworks.

3

A Serial H-Tree Router for Two-Dimensional Arrays

3.1 Router Functionality and Overhead

Advances in CMOS fabrication processes enable increases in the number and complexity of computational units in highly distributed and parallel architectures (e.g., neuromorphic processors; [2,37,39]), which calls for a corresponding increase in scalability and sophistication of routing mechanisms. Router area should be a reasonable fraction of the total system—router architecture is therefore dictated by client complexity. In this regard, high-overhead routers (e.g., meshes with parallel interfaces) capable of communicating arbitrary data-types at high bandwidths, are unsuited for intermediate-complexity clients with lower data-rate requirements. To provide multiple-data-type functionality for such clients, we adapt existing low-overhead routers.

Low-overhead routers contain a transmitter and a receiver [4]. The transmitter merges data from all of the clients into a single stream and adds source-identifying addresses to each datum to form a packet. The receiver takes a stream of packets, parses each destination-identifying address, and delivers the datum to the specified client.

For N clients, a low-overhead router’s circuitry scales as $\mathcal{O}(\sqrt{N})$ by sharing resources. Clients are tiled in a two-dimensional (2D) array and share row and column wires within the array and transceiver circuitry at the edge of the array [4,27] (Fig. 3.1, GRID ADDR). Sharing works correctly if certain timing assumptions are met [5,6], but these assumptions are difficult to satisfy for long wires, which are susceptible to phenomena such as *charge relaxation*, whereby a significant voltage difference arises between the wire’s two ends [25]. For this reason, grids do not readily scale to large arrays.

Our router presented herein switches from grid addresses to tree paths (Fig. 3.1, TREE PATH),

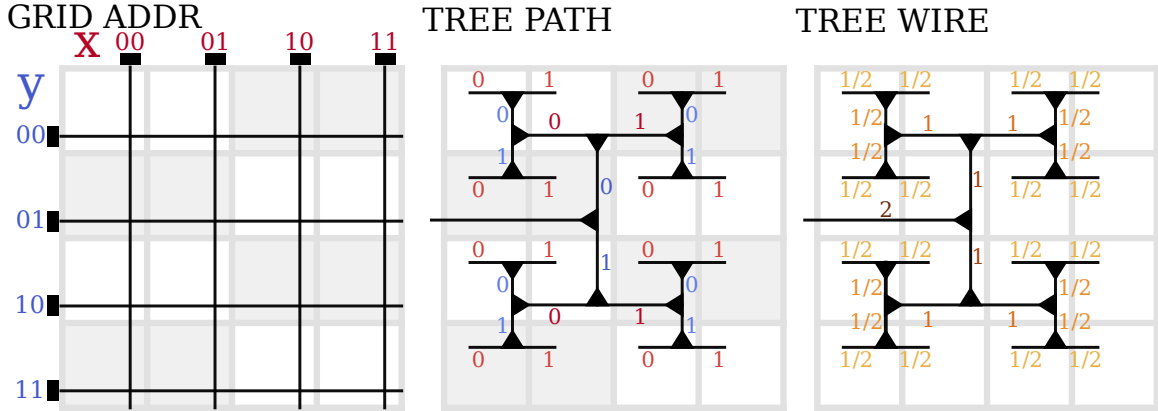


Figure 3.1: Grid Addresses and Tree Paths

Clients (white and gray squares) are tiled in a 2D array and routed to (or from) using a grid or a tree. **GRID ADDR**: A client’s address is encoded by concatenating its x and y positions (in binary). Addressing circuitry is placed at the array’s edge (black rectangles) and scales as $\mathcal{O}(\sqrt{N})$ for N clients ($N = 16$ shown). **TREE PATH**: A client’s path is encoded by traversing the tree from the root to leaf (indexed by $n = 3$ and $n = 0$, respectively, in Algorithms 1 & 2). Each up and left (down and right) branch appends a 0 (1). Shaded squares indicate differences between tree and grid binary-number assignments (e.g., the bottom left client’s grid-address is 0011, but its tree-path is 1010). Routing circuitry is embedded within the array (black triangles) and scales as $\mathcal{O}(N)$. **TREE WIRE**: Wire segments are annotated with their lengths.

trading an increase in logic circuitry for enhanced scalability and functionality. The increase in logic circuitry—from $\mathcal{O}(\sqrt{N})$ to $\mathcal{O}(N)$ for N clients—is worthwhile for emerging intermediate-complexity clients that use thick-oxide transistors for ultra-low power analog computation and much smaller thin-oxide transistors for ultra-fast digital communication [3]. The enhanced scalability arises because its asynchronous implementation’s timing assumptions are easily met. And the enhanced functionality arises because its serial protocol supports multiple datatypes, whereas the grid’s parallel protocol limits payload size. For backward compatibility, converting grid addresses into tree paths and vice versa is straightforward (Algorithms 1 & 2).

In Section 3.2, we describe how grid addresses and tree paths are encoded, show that both require $\mathcal{O}(N)$ wiring, and justify our choice of a 4-ary tree over a binary tree. In Section 3.3, we describe the serial link our router uses. In Section 3.4, we describe the logical design of the router’s nodes. In Section 3.5, we describe how the router’s leaves were customized for a neuromorphic application. In Section 3.6, we describe the router’s logical and physical synthesis and its verification and validation. In Section 3.7, we conclude the chapter with a discussion of our results.

Algorithm 1 Converts Path (p) to Address (x,y) <hr/> Require: $l = \text{length}(p)$ for $n = 0$ to $l/2 - 1$ do $x[n] \leftarrow p[2n]$ $y[n] \leftarrow p[2n + 1]$ end for	Algorithm 2 Converts Address (x,y) to Path (p) <hr/> Require: $l = \text{length}(p)$ for $n = 0$ to $l/2 - 1$ do $p[2n] \leftarrow x[n]$ $p[2n + 1] \leftarrow y[n]$ end for
--	--

3.2 Tree Paths versus Grid Addresses

A 2D array can be routed to (or from) using a grid or a tree. We consider N clients, of unit width and height, arranged on a square grid (Fig. 3.1, WIRING).

Given equal link-widths, the tree requires less wiring than the grid. To calculate the length of the tree’s wiring, W_t , we start from the N unit centers: $N \frac{1}{2}$ -unit segments project horizontally from each center. At the second level, $\frac{N}{2} \frac{1}{2}$ -unit segments project vertically. At the third level, $\frac{N}{4}$ 1-unit segments project horizontally. This geometric pattern continues up to the root; each level alternates between horizontal and vertical orientation and halves the number of segments from the previous, lower level, while doubling the segment-length every other level. Overall, we have

$$W_t = \frac{1}{2} \left(N + \frac{N}{2} \right) + 1 \left(\frac{N}{4} + \frac{N}{8} \right) + 2 \left(\frac{N}{16} + \frac{N}{32} \right) + \dots = \frac{3}{2} N$$

as N scales up (coloring matches Fig. 3.1, WIRING). For comparison, in the grid, each client adds 2 units of wire so that $W_g = 2N$. Therefore, $W_t = \frac{3}{4} W_g$: the tree uses up to 25% less wiring than the grid.¹ While the tree’s segments become longer as we move from leaves to root, they are shared among more and more leaves.

However, the primary trade-off is the tree’s larger transistor-count ($\mathcal{O}(N)$ versus $\mathcal{O}(\sqrt{N})$ for the grid), determined by the node-count times the transistors-per-node. To reduce the node-count, we opted for a 4-ary tree over a binary tree. A binary tree has $N - 1$ nodes whereas a 4-ary tree has $\frac{N-1}{3}$ nodes. In general, a k -ary tree has $\frac{N-1}{k-1}$ nodes and $\log_k(N)$ levels. Consequently, switching divides the node-count by three, halves the number of levels, halves the latency, and doubles the unpipelined throughput.

If switching from binary to 4-ary doubles the transistors-per-node,² and divides the node count by three, we would expect to decrease the overall transistor count by 33%. However, nodes are not homogeneous; leaf nodes are tailored to clients’ needs. The total transistor count in an N -client

¹In emerging 3D processes, with wire-segments traveling along three axes, segment-count still halves at each level towards the root, but segment-length only doubles every third level (c.f. every other level in 2D). As a result, $W_t = \frac{7}{24} W_g \approx 0.29 W_g$: the tree uses up to 71% less wiring.

²Doubling occurs if combinational gates (e.g., NANDS or NORs)—whose transistor count is $2 \times$ their fan-in—dominate. For sequential gates—whose state-holding elements are not replicated—the increase is sublinear. When gates are treed to build wider gates, the increase is supralinear.

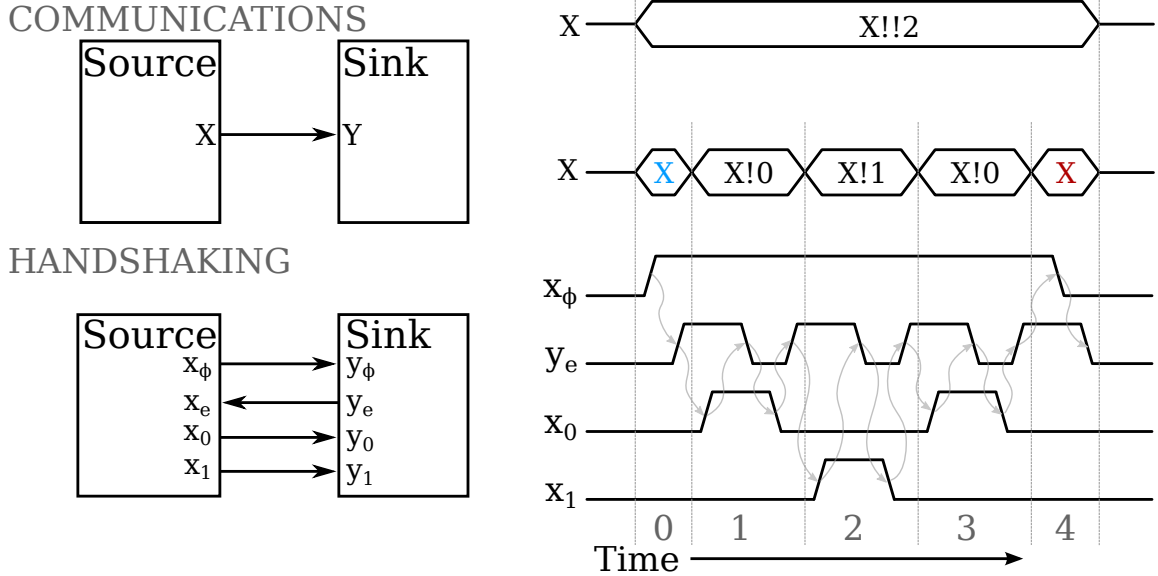


Figure 3.2: Serial Link Description at Two Levels of Abstraction

HANDSHAKING: *Source* drives control line x_ϕ and data lines x_0 and x_1 . *Sink* drives control line y_e . Arrows point from driver to listener. Two-phase handshakes (time slots 0 and 4) initiate and terminate packet communication; four-phase handshakes (1, 2, and 3) send the packet's bits as 1-of-2 codes. All transitions are acknowledged (curved gray arrows), so the protocol is delay-insensitive. y_e 's last transition is acknowledged by x_ϕ 's initial transition in the next packet. A three-bit packet (010) is communicated in this example, but the protocol supports arbitrary-sized packets and 1-of- D codes using D data lines. **COMMUNICATIONS:** A channel connects *Source*'s output port (X) to *Sink*'s input port (Y). *Source*'s dataless communications (X and X) initiate and terminate packet transmission; such communications are colored blue and red, respectively, throughout the text. Its datafull communications ($X!0$ and $X!1$) send the packet's bits. The entire communication sequence may be consolidated into the single operation ($X!!2$).

k -ary tree whose leaf and intermediate nodes have T_{Lk} and T_{Ik} transistors each, respectively, is

$$T_{\text{tot}k} = \frac{N-1}{k-1} \frac{(1-k/N)T_{Ik} + (k-1)T_{Lk}}{k-k/N} \quad (3.1)$$

For $T_{Lk} = T_{Ik} = T_k$, this expression reduces to $\frac{N-1}{k-1}T_k$: the total number of nodes times the transistors-per-node. Note that the ratio of leaf to intermediate nodes is $k-1:1-k/N$, which approaches 1:1 and 3:1 for binary and 4-ary trees, respectively, as N increases. Thus, based on T_{Lk} 's and T_{Ik} 's values for our designs (Tab. 3.1),³ which have different mixes of combinational and sequential logic and treed gates, switching from binary to 4-ary increases the average transistor count of the transmitter's and receiver's nodes by $2.6\times$ and $1.6\times$, respectively. As a result, their overall transistor count reduces by 13.3% and 45%, respectively (see Tab. 3.1).

³The leaf node's communication is dataless—it requests or acknowledges.

k	Transmitter		Receiver	
	2	4	2	4
T_{Lk}	78	208	30	54
T_{Ik}	91	255	64	148
T_{tot4}/T_{tot2}	0.867		0.550	

We built our serial tree router with quasi-delay-insensitive (QDI) circuits. The only timing assumption made is the *isochronic fork*. Signal-propagation delay along branches of such forks are assumed to be *equally* insignificant (hence the *iso* combining form; precise definitions may be found in [30,31]). This assumption is the minimal one necessary for useful computation with asynchronous circuits (i.e., Turing complete). No assumptions are made about signal-propagation delays through gates or nonisochronic wires, except that they are positive and finite.

3.3 Serial Communication Protocol

To keep link-width constant, we use serial communication. The path-length grows as we move from leaf to root in a tree. Hence, codes communicated over links closer to the root have more bits than those communicated over links closer to the leaves. A parallel protocol thus requires wider links (i.e. more wires) towards the root, whereas a serial protocol makes do with a constant width. Further, the latter allows us to communicate more than just the encoded paths; we can communicate data (e.g., configuration settings) as well.

Our serial-link follows a fully delay-insensitive version of the *bundled-data* protocol in [6] (Fig. 3.2, HANDSHAKING). For example, the following source generates a random bitstream and segments it into packets of arbitrary length (see Table 3.2 for syntax):

$$\begin{aligned}
& x_\phi \uparrow; [x_e] \uparrow; *[[\text{true} \longrightarrow x_0 \uparrow; [\neg x_e]; x_0 \downarrow; [x_e] \\
& \quad | \text{true} \longrightarrow x_1 \uparrow; [\neg x_e]; x_1 \downarrow; [x_e] \\
& \quad | \text{true} \longrightarrow x_\phi \downarrow; [\neg x_e]; x_\phi \uparrow; [x_e]]]
\end{aligned}$$

Handshakes that demarcate the beginning and end of packet transmission are colored blue and red, respectively. If the x_ϕ branch is executed immediately, or consecutively, the packet contains no data. A sink that consumes the source's data operates as follows.

$$\begin{aligned}
& [y_\phi] \downarrow; y_e \uparrow; *[[y_0 \vee y_1 \longrightarrow y_e \downarrow; [\neg y_0 \wedge \neg y_1]; y_e \uparrow \\
& \quad \square \quad \neg y_\phi \longrightarrow y_e \downarrow; [y_\phi] \downarrow; y_e \uparrow]]
\end{aligned}$$

Selection (deterministic) is used instead of arbitration (nondeterministic) because the source guarantees mutual exclusion between the branches.

At a higher level of abstraction, we describe the source's and sink's operation simply in terms of *communications* on ports connected by a channel (Fig. 3.2, COMMUNICATIONS). For the source:

$$X; *[[\text{true} \longrightarrow X!0 | \text{true} \longrightarrow X!1 | \text{true} \longrightarrow X; X]]$$

X and X correspond to two-phase handshakes (on x_ϕ and x_e) that demarcate the packet (see Table 3.3 for notation). $X!$ corresponds to repeated four-phase handshakes (on $x_{0,1}$ and x_e) that send the payload. For the sink:

$$Y; *[[\overline{Y?} \longrightarrow Y? \parallel \overline{Y} \longrightarrow Y; Y]]$$

Note the unconventional use of the probe to check whether a *datafull* communication is pending. This probe ($\overline{Y?}$) corresponds to $y_0 \vee y_1$, whereas the *dataless* communication probe (\overline{Y}) corresponds to $\neg y_\phi$.

We introduce $??$ and $!!$ operators to describe serial read and write communications concisely (Fig. 3.2, COMMUNICATIONS). The source and sink are described as

$$*[[\text{true} \longrightarrow X!!\text{null} \mid \text{true} \longrightarrow X!!\text{Rand}()]] \parallel *[[Y??]]$$

where `null` is an empty string (i.e. the packet is empty) and `Rand()` returns a random, nonnegative integer.

3.4 Router Logical Design

The router consists of a transmitter and a receiver, both composed of a tree of nodes (Fig. 3.4.1). The transmitter merges packets from the clients into a single stream for transmission to the environment. The receiver does the inverse; it splits each packet in the stream off to the targeted client. For conciseness, we describe the nodes' operation for a binary tree ($\text{TX}(2)$ and $\text{RV}(2)$). It is straightforward to extend these processes to a k -ary tree ($\text{TX}(k)$ and $\text{RV}(k)$).

For a design space exploration see appendices B, C, and D. Designs were evaluated by their transistor costs assuming an array of 4096 neurons, 1024 synapses, and 256 memory banks. With each group of 1 synapse and 4 neurons uses 28 bits of memory.

3.4.1 Transmitter

A transmitter node merges packet streams from its children into a single packet stream for its parent (another node one level closer to the root, unless the node is itself the root):

$$\begin{aligned} & \text{TX}(2) \\ & \equiv *[[\overline{C_0} \longrightarrow P!!(0 \oplus C_0??) \mid \overline{C_1} \longrightarrow P!!(1 \oplus C_1??)]] \end{aligned}$$

A packet from port $C_{0,1}$ is interpreted as a string; $a \oplus b$ prepends a to b (e.g., $1 \oplus 01 = 101$). Prepending a child's port index at that node in the tree to its data builds the overall path from leaf to root.

Expanding $!!$ and $??$ operators, and separating out arbitration, $\text{TX}(2)$ becomes

$$\begin{aligned}
& * [[\overline{C_0} \longrightarrow P; P!0; C_0; [\overline{C_0} \longrightarrow P; C_0] \\
& \quad | \overline{C_1} \longrightarrow P; P!1; C_1; [\overline{C_1} \longrightarrow P; C_1]] \\
& \parallel * [[\overline{C_0?} \longrightarrow P!C_0? \parallel \overline{C_1?} \longrightarrow P!C_1?]]
\end{aligned}$$

Communications that demarcate when packet transmission begins and ends at the child and parent ports are colored blue and red, respectively. Putting $P!\{0, 1\}$ before $C_{0,1}$ ensures that the child's index is prepended to the packet before the child's data are forwarded with the $P!C_{0,1?}$ communications.

Further expansion yields the following HSE.

$$\begin{aligned}
& * [[c_{0\phi} \longrightarrow p_\phi\uparrow; [p_e]; p_0\uparrow; [\neg p_e]; p_0\downarrow; [p_e]; c_{0e}\uparrow; \\
& \quad [\neg c_{0\phi}]; p_\phi\downarrow; [\neg p_e]; c_{0e}\downarrow \\
& \quad | c_{1\phi} \longrightarrow p_\phi\uparrow; [p_e]; p_1\uparrow; [\neg p_e]; p_1\downarrow; [p_e]; c_{1e}\uparrow; \\
& \quad [\neg c_{1\phi}]; p_\phi\downarrow; [\neg p_e]; c_{1e}\downarrow]], \\
& * [[c_{00} \longrightarrow p_0\uparrow; [\neg p_e]; c_{0e}\downarrow; [\neg c_{00}]; p_0\downarrow; [p_e]; c_{0e}\uparrow \\
& \quad \parallel c_{01} \longrightarrow p_1\uparrow; [\neg p_e]; c_{1e}\downarrow; [\neg c_{01}]; p_1\downarrow; [p_e]; c_{1e}\uparrow \\
& \quad \parallel c_{10} \longrightarrow p_0\uparrow; [\neg p_e]; c_{0e}\downarrow; [\neg c_{10}]; p_0\downarrow; [p_e]; c_{0e}\uparrow \\
& \quad \parallel c_{11} \longrightarrow p_1\uparrow; [\neg p_e]; c_{1e}\downarrow; [\neg c_{11}]; p_1\downarrow; [p_e]; c_{1e}\uparrow]]
\end{aligned}$$

Note that the initial parent communication completes ($[p_e]$) and a code is transmitted to the parent before the initial child communication is acknowledged ($c_{0e}\uparrow$ or $c_{1e}\uparrow$). After that, the selection process relays the child's data.

We proceed by factorizing the arbitration process into the arbiter itself and the remaining child-parent communication:

$$\begin{aligned}
& * [[c_{0\phi} \longrightarrow s_0\uparrow; [\neg c_{0\phi}]; s_0\downarrow | c_{1\phi} \longrightarrow s_1\uparrow; [\neg c_{1\phi}]; s_1\downarrow]], \\
& * [[s_0 \wedge \neg u \longrightarrow p_\phi\uparrow; [p_e]; w_0\uparrow; p_0\uparrow; [\neg p_e]; u\uparrow; w_0\downarrow; \\
& \quad p_0\downarrow; [p_e]; c_{0e}\uparrow; [\neg s_0]; p_\phi\downarrow; [\neg p_e]; c_{0e}\downarrow; u\downarrow \\
& \quad \parallel s_1 \wedge \neg u \longrightarrow p_\phi\uparrow; [p_e]; w_1\uparrow; p_1\uparrow; [\neg p_e]; u\uparrow; w_1\downarrow; \\
& \quad p_1\downarrow; [p_e]; c_{1e}\uparrow; [\neg s_1]; p_\phi\downarrow; [\neg p_e]; c_{1e}\downarrow; u\downarrow]]
\end{aligned}$$

$s_{0,1}$ are introduced to store the selection result; $w_{0,1}$ are introduced to distinguish the state immediately after $[p_e]$ (prepending the index) from that immediately after $[p_e]$ (acknowledging the child); and u is introduced to preserve mutual exclusion in the selection process when its branches are implemented as concurrent processes. It prevents the s_1 branch (i.e. $p_\phi\uparrow; [p_e]; \dots$) from beginning before the s_0 branch completes (i.e. $p_\phi\downarrow; [\neg p_e]; \dots$) when $c_{1\phi}$ is high and the arbiter executes $s_1\uparrow$ immediately after $s_0\downarrow$.

Our 4-ary transmitter tree's node uses a four-way arbiter (Fig. 3.4.1). Three mutual-exclusion elements are interconnected in a binary decision-tree by handshaking circuitry ($c_{0:3\phi}$ and $s_{0:3}$ connect to the two ARB2's $_c_{0,1i}$ inputs and $c_{0,1o}$ outputs, respectively) [32]. CHP and HSE are omitted

for brevity. For comparison, a binary tree's node requires just one mutual-exclusion element—with no additional overhead. Although the arbiter design used here contains no pipelining, a greedy, but fair arbiter is described in Appendix E.2.

The transmitter node's HSE (sans ARB(4)) is implemented by the following production rule set (PRS).

$$\begin{array}{ll}
\neg u \wedge (s_0 \vee s_1) & \rightarrow p_\phi \uparrow \\
(c_{0e} \wedge \neg s_0) \vee (c_{1e} \wedge \neg s_1) & \rightarrow p_\phi \downarrow \\
\\
s_0 \wedge p_e \wedge \neg u & \rightarrow w_0 \uparrow & u & \rightarrow w_0 \downarrow \\
s_1 \wedge p_e \wedge \neg u & \rightarrow w_1 \uparrow & u & \rightarrow w_1 \downarrow \\
\\
c_{00} \vee c_{10} \vee w_0 & \rightarrow p_0 \uparrow & \neg(c_{00} \vee c_{10} \vee w_0) & \rightarrow p_0 \downarrow \\
c_{01} \vee c_{11} \vee w_1 & \rightarrow p_1 \uparrow & \neg(c_{01} \vee c_{11} \vee w_1) & \rightarrow p_1 \downarrow \\
\\
(w_0 \vee w_1) \wedge \neg p_e & \rightarrow u \uparrow & \neg(c_{0e} \vee c_{1e} \vee p_\phi) & \rightarrow u \downarrow \\
\\
s_0 \wedge u \wedge p_e \wedge \neg c_{1e} & \rightarrow c_{0e} \uparrow & \neg p_e & \rightarrow c_{0e} \downarrow \\
s_1 \wedge u \wedge p_e \wedge \neg c_{0e} & \rightarrow c_{1e} \uparrow & \neg p_e & \rightarrow c_{1e} \downarrow
\end{array}$$

3.4.2 Receiver

A receiver node splits a packet stream from its parent into packet streams for its children (another node one level closer to the leaves, unless the node is a leaf itself):

$$\begin{array}{l}
\text{RV}(2) \\
\equiv *[[P??(s, d) \bullet [s = 0 \rightarrow C_0!!d \mid s = 1 \rightarrow C_1!!d]]]
\end{array}$$

It uses the packet's first word (written into s) to decide which child to send the remainder of the packet (written into d); s has 1 bit for a binary tree or 2 bits for a 4-ary tree (RV(4)).

We expand RV(2)'s ?? and !! communications as follows.

$$\begin{array}{l}
P;P?s \bullet [s = 0 \rightarrow C_0 \mid s = 1 \rightarrow C_1]; \\
*[[\overline{P?} \wedge s = 0 \rightarrow C_0!P? \\
\quad \mid \overline{P?} \wedge s = 1 \rightarrow C_1!P? \\
\quad \mid \overline{P} \rightarrow P \bullet [s = 0 \rightarrow C_0 \mid s = 1 \rightarrow C_1]; \\
\quad P;P?s \bullet [s = 0 \rightarrow C_0 \mid s = 1 \rightarrow C_1]]]
\end{array}$$

This process can be expanded further as

$$\begin{aligned}
& * [[p_0 \wedge s_0 \longrightarrow c_{00}\uparrow; [\neg c_{0e}]; p_e\downarrow; [\neg p_0]; c_{00}\downarrow; [c_{0e}]; p_e\uparrow \\
& \quad \square p_1 \wedge s_0 \longrightarrow c_{01}\uparrow; [\neg c_{0e}]; p_e\downarrow; [\neg p_1]; c_{01}\downarrow; [c_{0e}]; p_e\uparrow \\
& \quad \square p_0 \wedge s_1 \longrightarrow c_{10}\uparrow; [\neg c_{1e}]; p_e\downarrow; [\neg p_0]; c_{10}\downarrow; [c_{1e}]; p_e\uparrow \\
& \quad \square p_1 \wedge s_1 \longrightarrow c_{11}\uparrow; [\neg c_{1e}]; p_e\downarrow; [\neg p_1]; c_{11}\downarrow; [c_{1e}]; p_e\uparrow \\
& \quad \square \neg p_\phi \longrightarrow s_0\downarrow, s_1\downarrow; \\
& \quad \quad c_{0\phi}\downarrow, c_{1\phi}\downarrow; [\neg c_{0e} \wedge \neg c_{1e}]; p_e\downarrow; [p_\phi]; p_e\uparrow; \\
& \quad \quad [p_0 \longrightarrow s_0\uparrow; p_e\downarrow; [\neg p_0]; c_{0\phi}\uparrow; [c_{0e}]; p_e\uparrow \\
& \quad \quad \square p_1 \longrightarrow s_1\uparrow; p_e\downarrow; [\neg p_1]; c_{1\phi}\uparrow; [c_{1e}]; p_e\uparrow]]]
\end{aligned}$$

After reset, the process resumes at \square .

To realize these five branches as five concurrent processes, we must preclude the the first four from starting immediately after the fifth process executes $s_{0,1}\uparrow$. We accomplish this by replacing $s_{0,1}$ in their guards with $c_{0,1\phi}$, which also indicate the selected child.

$$\begin{aligned}
& * [[p_0 \wedge c_{0\phi} \longrightarrow c_{00}\uparrow; [\neg c_{0e}]; p_e\downarrow; [\neg p_0]; c_{00}\downarrow; [c_{0e}]; p_e\uparrow \\
& \quad \square p_1 \wedge c_{0\phi} \longrightarrow c_{01}\uparrow; [\neg c_{0e}]; p_e\downarrow; [\neg p_1]; c_{01}\downarrow; [c_{0e}]; p_e\uparrow \\
& \quad \square p_0 \wedge c_{1\phi} \longrightarrow c_{10}\uparrow; [\neg c_{1e}]; p_e\downarrow; [\neg p_0]; c_{10}\downarrow; [c_{1e}]; p_e\uparrow \\
& \quad \square p_1 \wedge c_{1\phi} \longrightarrow c_{11}\uparrow; [\neg c_{1e}]; p_e\downarrow; [\neg p_1]; c_{11}\downarrow; [c_{1e}]; p_e\uparrow \\
& \quad \square \neg p_\phi \longrightarrow s_0\downarrow, s_1\downarrow; ss\downarrow; v\downarrow; \\
& \quad \quad c_{0\phi}\downarrow, c_{1\phi}\downarrow; [\neg c_{0e} \wedge \neg c_{1e}]; p_e\downarrow; [p_\phi]; p_e\uparrow; \\
& \quad \quad [p_0 \longrightarrow s_0\uparrow; ss\uparrow; p_e\downarrow; [\neg p_0]; v\uparrow; c_{0\phi}\uparrow; [c_{0e}]; p_e\uparrow \\
& \quad \quad \square p_1 \longrightarrow s_1\uparrow; ss\uparrow; p_e\downarrow; [\neg p_1]; v\uparrow; c_{1\phi}\uparrow; [c_{1e}]; p_e\uparrow]]]
\end{aligned}$$

v is introduced to allow $c_{0,1\phi}$ to be combinational and ss is introduced to reduce the length of p_e 's pull-up and pull-down chains (see PRS below).

The following PRS implements the receiver node's HSE.

$$\begin{aligned}
p_\phi \wedge \neg ss \vee c_{0e} \vee c_{1e} & \longrightarrow p_e\uparrow \\
(\neg p_\phi \vee ss) \wedge \neg c_{0e} \wedge \neg c_{1e} & \longrightarrow p_e\downarrow \\
s_0 \vee s_1 & \longrightarrow ss\uparrow & \neg s_0 \wedge \neg s_1 & \longrightarrow ss\downarrow \\
p_0 \wedge \neg v & \longrightarrow s_0\uparrow & \neg p_\phi & \longrightarrow s_0\downarrow \\
p_1 \wedge \neg v & \longrightarrow s_1\uparrow & \neg p_\phi & \longrightarrow s_1\downarrow \\
ss \wedge \neg p_0 \wedge \neg p_1 & \longrightarrow v\uparrow & \neg ss & \longrightarrow v\downarrow \\
v \wedge s_0 & \longrightarrow c_{0\phi}\uparrow & \neg v \vee \neg s_0 & \longrightarrow c_{0\phi}\downarrow \\
v \wedge s_1 & \longrightarrow c_{1\phi}\uparrow & \neg v \vee \neg s_1 & \longrightarrow c_{1\phi}\downarrow \\
p_0 \wedge c_{0\phi} & \longrightarrow c_{00}\uparrow & \neg p_0 \vee \neg c_{0\phi} & \longrightarrow c_{00}\downarrow \\
p_1 \wedge c_{0\phi} & \longrightarrow c_{01}\uparrow & \neg p_1 \vee \neg c_{0\phi} & \longrightarrow c_{01}\downarrow \\
p_0 \wedge c_{1\phi} & \longrightarrow c_{10}\uparrow & \neg p_0 \vee \neg c_{1\phi} & \longrightarrow c_{10}\downarrow \\
p_1 \wedge c_{1\phi} & \longrightarrow c_{11}\uparrow & \neg p_1 \vee \neg c_{1\phi} & \longrightarrow c_{11}\downarrow
\end{aligned}$$

3.5 Router Application

We connected a 2D array of spiking-neuron clusters to a datapath using our asynchronous serial tree-router, a natural choice for spike communication. A product of continuous, noisy analog dynamics at biological timescales, spikes are relatively infrequent (sub-kHz) and asynchronous (there’s no clock). Each cluster contains 16 spike-generating soma circuits, 4 spike-consuming synapse circuits, and a configuration memory.

Clusters are tiled in a 16×16 array. To service the 4,096 somas, 1,024 synapses, and 256 memories, the transmitter’s and receiver’s trees are six ($4^6 = 4096$) and five ($4^5 = 1024$) levels deep, respectively. Half of the receiver’s 1,024 output ports suffice to service all 1,024 synapses because each port supplies 2 bits (a 1-of-4 code) whereas each synapse needs only 1 bit (indicates whether a spike is excitatory or inhibitory). Thus, 512 ports are left over to service the 256 memories. We customized the transmitter’s and receiver’s leaf nodes to suit this application as follows.

3.5.1 Transmitter Leaf

The transmitter leaf transmits a soma’s spike up the tree by creating a packet containing the soma’s index. With no other data to convey, the transmitter node is simplified to

$$\text{TXL}(2) \equiv *[[\overline{C_0} \rightarrow C_0 \bullet (P; P!0)]; [\overline{C_0} \rightarrow C_0 \bullet P] \\ | \overline{C_1} \rightarrow C_1 \bullet (P; P!1)]; [\overline{C_1} \rightarrow C_1 \bullet P]]$$

Note that our design actually instantiates TXL(4). We omit its HSE and PRS for brevity.

Somas lock up the transmitter during spike emission. When emitting a spike, a soma initiates packet transmission with a two-phase handshake ($C_{0,1}$). Afterwards, the soma enters a refractory period for up to a few milliseconds before executing another two-phase handshake ($\overline{C}_{0,1}$) to terminate transmission. If communications within the transmitter are slackless, the soma will lock up the transmitter during its refractory period. To prevent this, we insert a buffer (i.e., latch) between the soma and the transmitter’s leaf.

3.5.2 Receiver Leaf

The receiver leaf services four synapses as well as a configuration memory (via a deserializer). We repurpose two of the receiver node’s 2-bit ports to service the four synapses and use a third port to communicate with the memory:

$$\begin{aligned}
\text{RVL}(4) &\equiv P;P?s \bullet [s = 2 \longrightarrow C_2 \llbracket s \neq 2 \longrightarrow \text{skip} \rrbracket]; \\
&\quad *[\overline{P?} \wedge s = 0 \longrightarrow C_0!P? \\
&\quad \quad \llbracket \overline{P?} \wedge s = 1 \longrightarrow C_1!P? \\
&\quad \quad \llbracket \overline{P?} \wedge s = 2 \longrightarrow C_2!P? \\
&\quad \quad \overline{P} \longrightarrow P \bullet [s = 2 \longrightarrow C_2 \llbracket s \neq 2 \longrightarrow \text{skip} \rrbracket]; \\
&\quad \quad P;P?s \bullet [s = 2 \longrightarrow C_2 \llbracket s \neq 2 \longrightarrow \text{skip} \rrbracket]]
\end{aligned}$$

Only the third port (C_2) continues with the serial protocol. HSE and PRS are omitted for brevity.

Synapses, like somas, require buffering. Depending on its analog biasing, a synapse may take up to a few milliseconds to acknowledge an input spike. We add a full cycle of slack to the otherwise slackless communication from root to synapse. One half-cycle is built into the leaf; a standard weak-precharge half-buffer provides the other.

The configuration memory accepts 6 bits of address and 2 bits of data (its 128 bits are organized into eight rows and eight 2-bit-wide columns). These 8 bits are encoded in four 1-of-4 codes that the deserializer receives in series from the receiver leaf and presents in parallel to the memory.

3.5.3 Serial-Parallel Conversion

The deserializer converts M sequentially delivered 1-of-4 codes into a $M \times 1$ -of-4 parallel code using a chain of M DES (Fig. 3.4.2, DESERIAL). For 1-of-2 codes, DE's HSE is:

$$\begin{aligned}
&*[[s_i]; [x_0 \longrightarrow y_0 \uparrow; x_a \uparrow; [\neg x_0]; s_o \uparrow; x_a \downarrow; [\neg s_i]; y_0 \downarrow; s_o \downarrow \\
&\quad \llbracket x_1 \longrightarrow y_1 \uparrow; x_a \uparrow; [\neg x_1]; s_o \uparrow; x_a \downarrow; [\neg s_i]; y_1 \downarrow; s_o \downarrow \rrbracket]
\end{aligned}$$

For each conversion, $s_{i,o}$ propagate an event along the chain twice. The first time, serial input codes are latched to build the parallel output. The second time, the parallel output is cleared, which happens once a c-element that closes the chain receives the environment's acknowledge. PRS for a 1-of-2 version of DE is as follows.

$$\begin{array}{llll}
\neg s_o \wedge s_i \wedge x_0 & \rightarrow y_0 \uparrow & \neg s_i & \rightarrow y_0 \downarrow \\
\neg s_o \wedge s_i \wedge x_1 & \rightarrow y_1 \uparrow & s_o \vee \neg v_y & \rightarrow x_a \downarrow \\
\neg s_o \wedge v_y & \rightarrow x_a \uparrow & \neg s_i & \rightarrow y_1 \downarrow \\
y_0 \vee y_1 & \rightarrow v_y \uparrow & \neg y_0 \wedge \neg y_1 & \rightarrow v_y \downarrow \\
v_y \wedge \neg x_0 \wedge \neg x_1 & \rightarrow s_o \uparrow & \neg v_y & \rightarrow s_o \downarrow
\end{array}$$

v_y is introduced to shorten transistor chains.

The serializer does the converse of the deserializer: It uses a chain of M SES to slice a $M \times 1$ -of-4 parallel code into M 1-of-4 codes and forwards them sequentially to the environment using SEQ (Fig. 3.4.2, SERIAL). For 1-of-2 codes, SE's HSE is:

```

*[[si];
[x0 → y0↑; [ya]; u↑; y0↓; [¬ya]; so↑; [¬si]; u↓; [¬x0]; so↓
[x1 → y1↑; [ya]; u↑; y1↓; [¬ya]; so↑; [¬si]; u↓; [¬x1]; so↓
]]

```

u is added to distinguish states before and after the $y_{0,1}$ - y_a handshake. As with the deserializer, $s_{i,o}$ propagate an event along the chain twice for each conversion. The first time, each SE relays a code to SEQ. The second time, each SE checks that its parallel-input slice is cleared. SE's PRS is:

$$\begin{array}{lll}
x_0 \wedge \neg u \wedge s_i \rightarrow y_0 \uparrow & u \wedge \neg s_o & \rightarrow y_0 \downarrow \\
x_1 \wedge \neg u \wedge s_i \rightarrow y_1 \uparrow & u \wedge \neg s_o & \rightarrow y_1 \downarrow \\
s_i \wedge y_a \rightarrow u \uparrow & \neg s_i & \rightarrow u \downarrow \\
u \wedge \neg y_a \rightarrow s_o \uparrow & \neg u \wedge \neg x_0 \wedge \neg x_1 & \rightarrow s_o \downarrow
\end{array}$$

SEQ's HSE is:

```

*[[si]; so↑; [x0 ∨ x1]; yφ↑; [¬si ∧ ye]; yφ↓; [¬ye]; so↓],
*[[x0 ∧ ye → y0↑; [¬ye]; xa↑; [¬x0]; y0↓; [ye]; xa↓
[x1 ∧ ye → y1↑; [¬ye]; xa↑; [¬x1]; y1↓; [ye]; xa↓]]

```

By closing the chain, it initiates $(y_\phi \uparrow; [y_e])$ and terminates $(y_\phi \downarrow; [\neg y_e])$ packet transmission on the event's first and second pass, respectively. In between, it forwards codes that SES provide. SEQ's PRS is:

$$\begin{array}{lll}
x_0 \vee x_1 \rightarrow y_\phi \uparrow & \neg s_i \wedge y_e \rightarrow y_\phi \downarrow \\
(y_0 \vee y_1) \wedge \neg y_e \rightarrow x_a \uparrow & y_e \rightarrow x_a \downarrow \\
y_e \wedge x_0 \rightarrow y_0 \uparrow & \neg x_0 \rightarrow y_0 \downarrow \\
y_e \wedge x_1 \rightarrow y_1 \uparrow & \neg x_1 \rightarrow y_1 \downarrow \\
\neg s_i \wedge \neg y_e \wedge \neg y_\phi \rightarrow s_o \downarrow & s_i \rightarrow s_o \uparrow
\end{array}$$

3.6 Synthesis and Validation

For logical synthesis, we described logical hierarchy and PRS in the Asynchronous Compiler Tools (ACT) language.⁴ We verified logical correctness with PRSIM, a discrete-event simulator that executes PRS with randomized delays [1], and then checked for logical-physical consistency in the presence of transistor parasitic capacitances with CoSIM, a PRSIM-SPICE co-simulator [1].

For physical synthesis, we decomposed our ACT into standard cells and generated their layouts with cellTK [22]. Encounter (*Cadence*) place-and-routed lower-level router circuitry—4 transmitter leaves (to service 16 somas), their parent node, a receiver leaf (to service 4 synapses and an SRAM), and a deserializer (to interface with the SRAM)—in the lower 43% ($547 \mu\text{m}^2$) of the neuron-cluster

⁴Available at https://github.com/samfok/AER_serial_tree_router

tile ($1,261\mu\text{m}^2$). Of this router area, 14% ($76\mu\text{m}^2$) was reserved (cutout) for higher-level circuitry (Fig. 3.4.2, TILE).

We placed tiles in a 16×16 -array and placed the router’s higher-level nodes in their cutouts, along with repeaters to drive long wires (Fig. 3.4.2, TILE 16×16 , and H-TREES). We extracted parasitic resistances and capacitances from this layout and performed simulations to check for spurious transitions and to predict the router’s maximum throughput.

Our postlayout simulations predicted that the transmitter and receiver could communicate up to 42.5 and 50.8 Mspike/s, respectively (Fig. 3.6).⁵ Codes from (or to) nodes lower in the tree take longer (e.g., 4.31 ns from the transmitter tree’s leaves versus 1.28 ns from its root) because the number of communications involved increases (from 6 at the leaf to 1 at the root). On average, 4.5 four-phase communications are performed, including one for the 2 two-phase communications that demarcate the packet. At 4 phases per communication and 4 transitions per phase, transversing six nodes involves 432 transitions.⁶ Thus, the 42.5 Mspike/s cycle-rate corresponds to 56 ps per transition, in line with expectations for a 28-nm process.

Post-fabrication in a 28-nm, fully depleted, silicon-on-insulator process, we brought the chip up and validated the router’s functionality. (Fig. 3.6). From two chips, we measured maximum throughputs of 27.4 and 26.1 Mspikes/s for the transmitter and 18.1 and 18.5 Mspikes/s for the receiver.⁷ Differences between simulations the chip measurements are explained by additional delays introduced by unpipelined datapath communications.

3.7 Discussion

Pioneering researchers developed transmitters and receivers to write and read spikes to and from 1D or 2D arrays of silicon neurons using the *address-event representation* (AER; [4, 24, 27, 41]). A neuron’s address is transmitted every time it spikes, hence the name *address-event*. In 1D, the spike is identified by a unique address assigned to each neuron. In 2D, the spike is identified by the neuron’s row and column addresses and, in first-generation designs, these addresses are transmitted in parallel.

Second-generation designs communicated row and column addresses in series. In addition to saving wires by multiplexing, this so-called *word-serial* protocol supports packets with an arbitrary number of words. Thus, additional column addresses could be appended to communicate multiple spikes read from or written to the same row in parallel [5, 6, 25]. This so-called *burst-mode* offered higher throughput, servicing arrays containing as many as 64k somas and 256k synapses [2] at rates up to 43.4M spike/s (ignoring off-chip delays) [7]. Array or chip addresses could be prepended

⁵In operation, somas generate up to 500 spike/s each, and synapses consume up to 1000 spikes/s each, so we expect the transmitter and receiver to communicate 2 and 1 Mspikes/s, respectively

⁶In our PRSIM simulations, we counted 422 transitions for the transmitter and serializer and 481 transitions for the deserializer and receiver.

⁷Equivalently, the fabricated router can service arrays of up to 53.5k somas and 18.3k synapses, respectively

to further expand the address-space. Thus, an address-event-based router could service multiple arrays distributed across multiple chips [11, 36, 38]. Further, data as well as addresses could be communicated over the link (or bus) connecting the transmitter to the receiver [17]. In this fashion, multiple spikes read in parallel from small groups of neurons have been transmitted using a single dataword, boosting throughput, which had plateaued at 50M spike/s [10], to 300M spike/s [42].

To communicate configuration datawords to or from individual neurons—or clusters thereof—we could widen the neuronal interface, but the additional bandwidth would be largely wasted. Datawords use all of the wires but occur rarely, whereas spikes occur frequently but only use one (e.g., a soma’s output) or two (e.g., a synapse’s excitatory or inhibitory input) wires. We thus keep the neuronal interface narrow and transmit data serially, saving wires by taking more time. In addition to supporting multiple data-types efficiently, a serial protocol places no limit on the number of bits a dataword can have, unlike a parallel protocol.

We did away with timing-assumptions by switching from row-column addresses to tree paths. Striking a balance between node-count and node-complexity, we chose a 4-ary over a binary tree, which reduced transistor-count by 19.1% overall. Returns diminish for higher degrees because realizing wider gates requires treeing narrower gates (with no more than four transistors in series).⁸ Although its thin-oxide transistors outnumber the neuron-cluster’s thick-oxide transistors 1.9:1, the router takes up only 43% of the total area because thick-oxide transistors are much larger than thin oxide transistors.

Throughput may be enhanced substantially by pipelining the otherwise slackless communication from leaf to root (transmitter) or root to leaf (receiver) and between router and datapath. Pipelining can be added to the current design at no additional area cost by replacing repeaters with latches or placing latches in unused tile cutouts (see Appendix E.1. Subsequent codes would take no more time than the first one, which takes 1.28 (transmitter) or 1.8 ns (receiver) (Fig. 3.6). Therefore, with seven communications per spike, pipelining would increase throughput from 42.5M to 111.6M spike/s (transmitter) or from 50.7M to 79.4M spike/s (receiver).

⁸Gates with longer chains operate much slower, increasing the duration that downstream gates pass short-circuit currents.

	Assignment
$x\uparrow / x\downarrow$	Set boolean variable x to true / false
	Program Composition
$s_1; s_2$	Execute segment s_1 and then s_2
s_1, s_2	Execute s_1 concurrently with s_2
$*[s]$	Execute s repeatedly
	Boolean Operations
$x / \neg x$	Return the value of x / negated value of x
$e_1 \wedge e_2$	Return the logical-and of e_1 and e_2
$e_1 \vee e_2$	Return the logical-or of e_1 and e_2
	Branching
$[e]$	Wait until boolean expression e is true
$[e_1 \rightarrow s_1]$	When e_1 becomes true, execute s_1
$[e_1 \rightarrow s_1 \mid e_2 \rightarrow s_2]$	If boolean e_1 (e_2) is true, execute s_1 (s_2) If both are true, execute either s_1 or s_2 If both are false, wait
$[e_1 \rightarrow s_1 \square e_2 \rightarrow s_2]$	If boolean e_1 (e_2) is true, execute s_1 (s_2) Assume e_1 and e_2 cannot both be true If both are false, wait

Table 3.2: Handshaking Expansion (HSE) Syntax

	Assignment
$x := d$	Set variable x to d 's value
	Communication
X	Communicate on port X (dataless)
$X!x$	Write value of x to X
$X?x$	Read value from X to x
$Y!X?$	Read value from X and write it to port Y
\bar{X}	<i>True</i> if a communication is pending and <i>false</i> if not
	Program Composition
$S_1; S_2$	Execute segment S_1 and then S_2
$S_1 \parallel S_2$	Execute S_1 in parallel with S_2
$S_1 \bullet S_2$	Overlap the execution of S_1 and S_2 (called <i>bullet</i>)
$*[S]$	Execute S repeatedly
	Boolean Operations
\neg, \wedge, \vee	Same as in Table 3.2
$x = d$	Return <i>true</i> if x 's value equals d 's and <i>false</i> if not
	Branching
$\rightarrow, \mid, \square$	Same as in Table 3.2

Table 3.3: Communicating Hardware Processes (CHP) Syntax

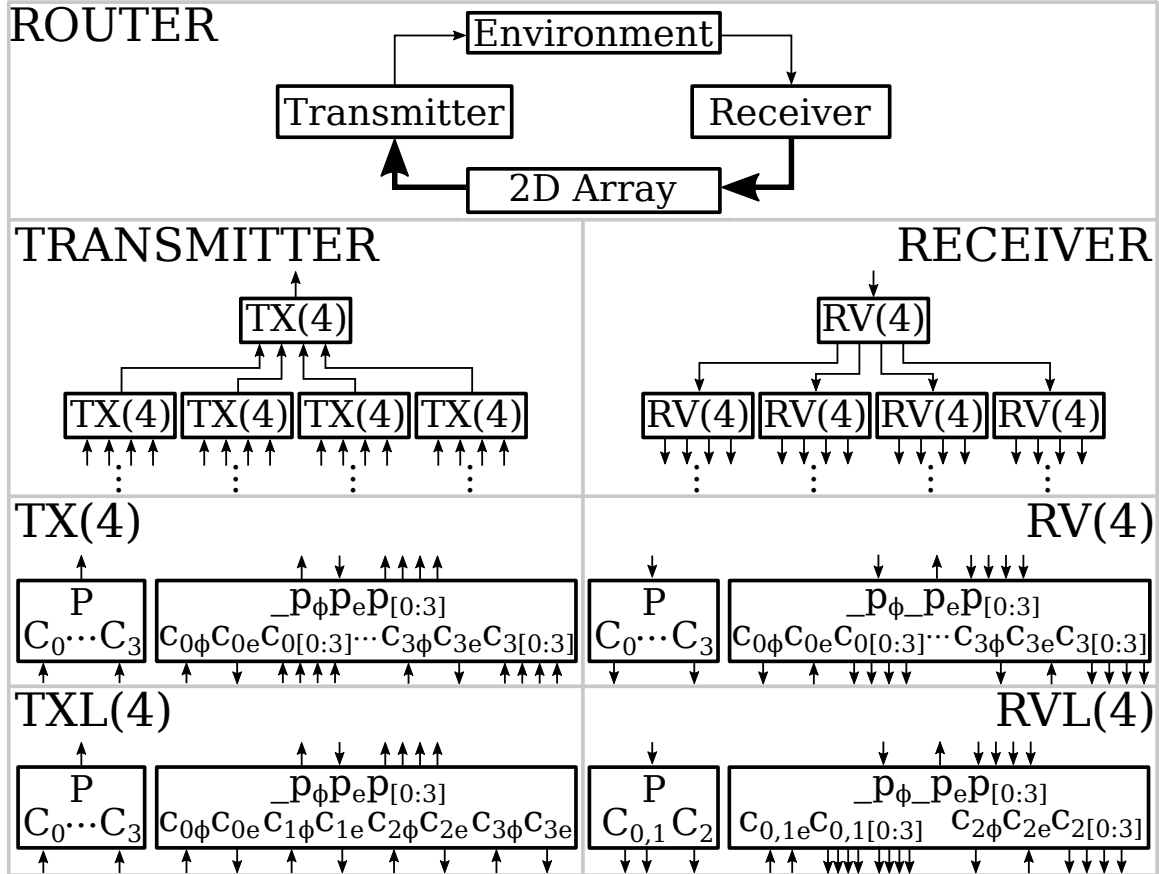


Figure 3.3: Router Process Decomposition

ROUTER: Facilitates communication between clients tiled in a 2D array and an external environment using a transmitter and a receiver. TRANSMITTER and RECEIVER: A pair of 4-ary trees provide an input and output port at their leaves for each client. TX(4) and RV(4): CHP ports (left) and HSE signals (right) that interface processes running in TRANSMITTER's and RECEIVER's nodes with their environment. TXL(4) and RVL(4): Same as previous but for processes in the leaves.

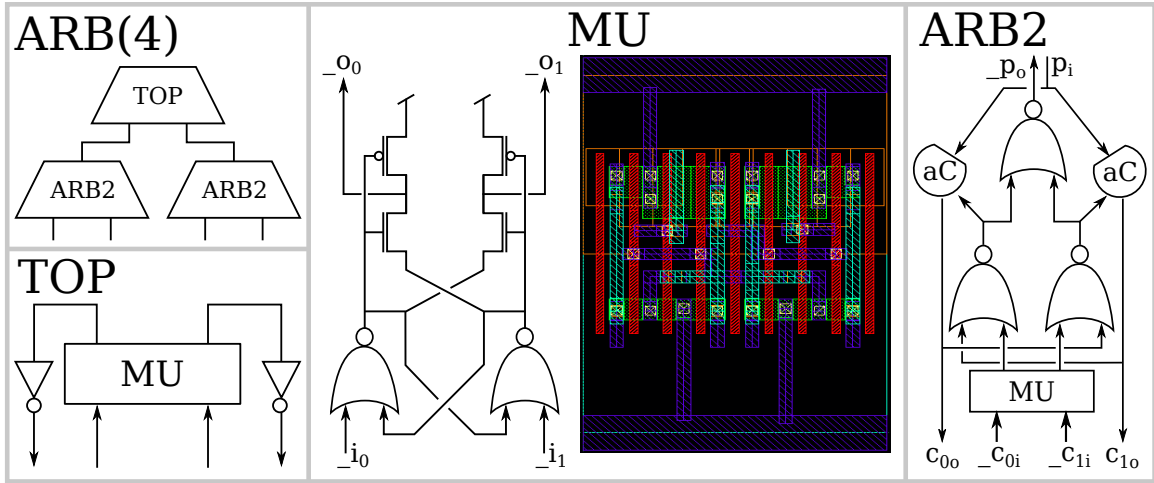


Figure 3.4: Four-Way Arbiter

ARB(4): Selects one of four clients with one TOP and two ARB2s; k clients require $k - 2$ ARB2s connected in a binary tree. TOP: Performs two-way selection with a MU. MU: Selects one of two active-low (indicated by underscore prefix) inputs ($_i_0$ and $_i_1$) using cross-coupled NOR gates. Four additional transistors filter out metastable signals before toggling the outputs ($_o_0$ and $_o_1$). MU's custom standard-cell layout is shown. ARB2: Relays its childrens' requests to its parent and relays its parent's grant to a requesting child, selected beforehand by MU. The two, lower NOR gates ensure that handshakes on $c_{0i,o}$ and $c_{1i,o}$ do not overlap; aC are asymmetric C-elements.

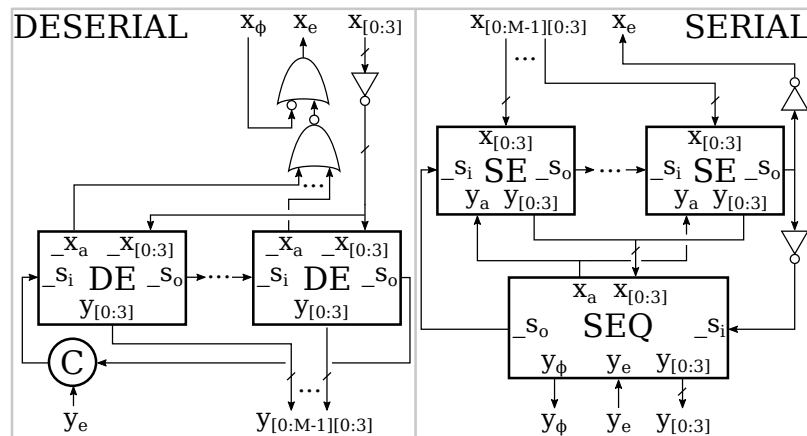


Figure 3.5: Serial-Parallel Conversion

DESERIAL: Serial input fans out to a chain of M DES. An event moves from one to the next with each serial input; it loops back around through the C-element when parallel output occurs. SERIAL: Parallel input is divided among a chain of M SES. As an event moves from one to the next, it outputs its data to SEQ. The event loops back around through SEQ when serial transmission is complete.

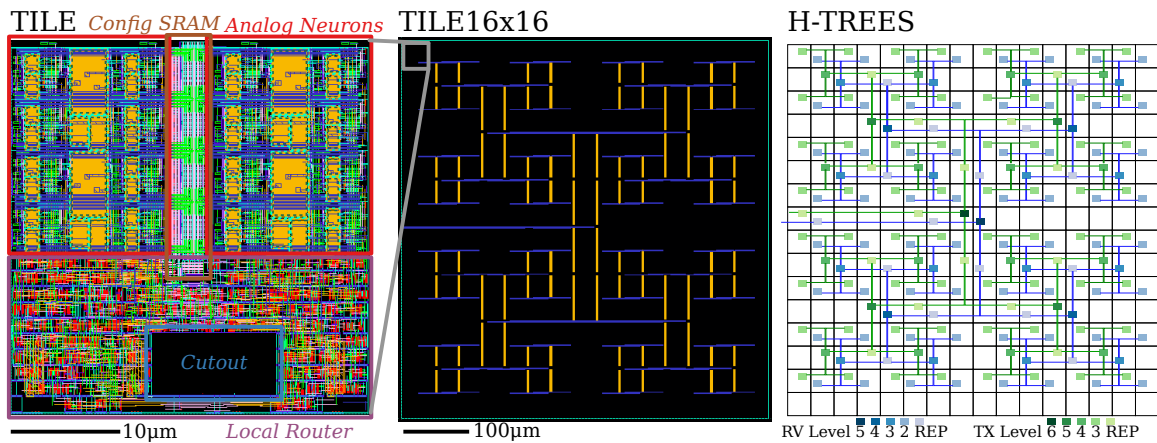


Figure 3.6: Tree Router Layout

TILE: Neuron cluster and low-level, local router circuits. *Analog Neurons*: Circuitry for 16 somas and 4 synapses. *Config SRAM*: Sixty-four 2-bit words—tiled in 8 rows and 8 columns—for analog circuitry configuration. *Local Router*: Four transmitter-tree leaves, their parent, a receiver-tree leaf, and a deserializer (SRAM interface). *Cutout*: Populated as needed with the transmitter or receiver trees' higher-level nodes or repeaters. TILE16×16: Full 2D array. Digital signals enter and exit on its left side, where the datapath is attached. To minimize crosstalk with analog circuitry, H-tree wires runs over the TILES' *Local Routers* and *Config SRAMs*. A one-tile horizontal displacement between the two H-trees makes wiring possible with just metal layers 5 (yellow) and 6 (purple). H-TREES: Placement of two H-trees' higher-level nodes and repeaters (green for transmitter and blue for receiver) in tiles (white squares) with routing overlaid. TX: Transmitter; RV: Receiver; REP: Repeater.

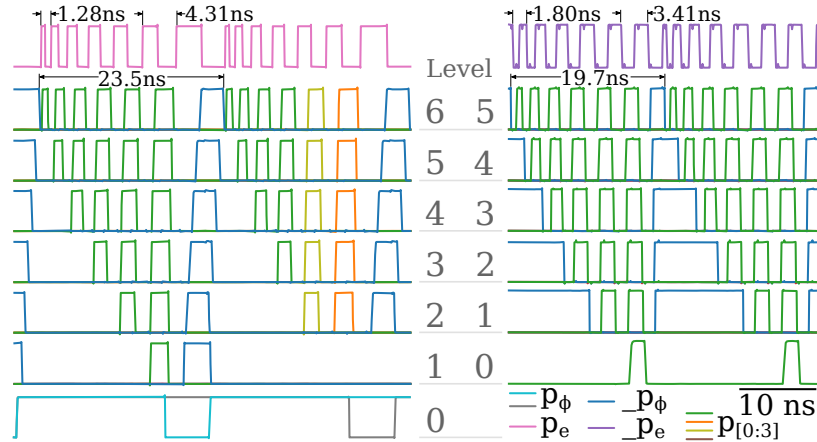


Figure 3.7: Postlayout Transmitter and Receiver SPICE Simulations

Left: SOMA 0 and 9 (000000 and 000021 in 4-ary) spike simultaneously (Level 0). Their p_ϕ signals propagate up to Level 1 (only SOMA 0's parent is shown) and Level 2, where SOMA 0's subtree is selected. Thus, its p_ϕ signal propagates to the root (Level 6). Enabled by the environment's p_e signal (top), which propagates down the tree (not shown), each node forwards its requesting child's 1-of-4 coded index ($p_{0:3}$) and then forwards indices forwarded by the child (they are all 0 for SOMA 0). Each node clears its p_ϕ once its child clears its p_ϕ , signaling that there are no more indices to be forwarded. A node is then free to select another requesting subtree, as happens at Level 2 for SOMA 9's subtree. *Right:* The environment sends two inhibitory spikes to SYNAPSE 0 by injecting two packets containing its path appended with 0 (i.e., 000000) at the root. Each node selects the child indexed by the first 1-of-4 code ($p_{0:3}$) and forwards the remaining codes to that child after lowering p_ϕ . Note that the leaf (Level 1) does not propagate p_ϕ to the synapse (Level 0).




Figure 3.8: Fabrication and validation

Left: Test chip containing analog neurons, the router presented herein, and a digital datapath. *Center:* Test board piggybacked on an FPGA development board (Opal Kelly) that provides a USB link to a host computer. *Right:* Visualization of a 32×32 -soma patch of the chip's spiking activity. Each small square represents a soma; its brightness reflects the soma's spike rate. The four, bright soma clusters are receiving excitatory input from spikes delivered to nearby synapses.

4

Conclusions

In this thesis, I have described the theoretical underpinnings for Braindrop's accumulative hardware for spike-train weighting and summing and have detailed the physical router hardware for communicating spikes as well as programming packets to and from Braindrop's neuron array. With its completion, Braindrop affords neuromorphic engineers both a sufficient number of neurons for nontrivial tasks and, different from Neurogrid, a well-mapped, systematic means of configuring those neurons. It is my hope that Braindrop enables neuromorphic engineers to convincingly demonstrate to the wider world that the brain still has something to teach us about computing.

Appendix A

Spike Summing and Weighting

A.1 periodic SNR approximation

For $\lambda\tau \rightarrow \infty$, $\text{SNR}_{\text{periodic}} \rightarrow 2\sqrt{3}\lambda\tau$. Seeing this is not so straightforward since $\lim_{\lambda\tau \rightarrow \infty} \coth\left(\frac{1}{2\lambda\tau}\right)$ is not defined, so we move forward using Taylor series approximations. Recalling that $\coth\left(\frac{1}{2\lambda\tau}\right) = (1 + e^{-1/\lambda\tau})(1 - e^{-1/\lambda\tau})^{-1}$,

$$\begin{aligned} \text{SNR}(X) &= \sqrt{\frac{2\lambda\tau}{\frac{1+e^{-1/\lambda\tau}}{1-e^{-1/\lambda\tau}} - 2\lambda\tau}} = \sqrt{\frac{2\lambda\tau}{\frac{1+1-\frac{1}{\lambda\tau}+\frac{1}{2}\frac{1}{(\lambda\tau)^2}-\frac{1}{6}\frac{1}{(\lambda\tau)^3}+\dots}{1-1+\frac{1}{\lambda\tau}-\frac{1}{2}\frac{1}{(\lambda\tau)^2}+\frac{1}{6}\frac{1}{(\lambda\tau)^3}+\dots} - 2\lambda\tau}} \\ &= \sqrt{\frac{1}{\frac{\frac{1}{2\lambda\tau} - \frac{1}{\lambda\tau} + \frac{1}{2}\frac{1}{(\lambda\tau)^2} - \frac{1}{6}\frac{1}{(\lambda\tau)^3} + \dots}{\frac{1}{\lambda\tau} - \frac{1}{2}\frac{1}{(\lambda\tau)^2} + \frac{1}{6}\frac{1}{(\lambda\tau)^3} + \dots} - 1}} \\ &= \sqrt{\frac{1}{\frac{2 - \frac{1}{\lambda\tau} + \frac{1}{2}\frac{1}{(\lambda\tau)^2} - \frac{1}{6}\frac{1}{(\lambda\tau)^3} + \dots}{2 - \frac{1}{\lambda\tau} + \frac{1}{3}\frac{1}{(\lambda\tau)^2} - \frac{1}{12}\frac{1}{(\lambda\tau)^3} + \dots} - 1}} \end{aligned}$$

The Taylor series approximations begin to differ with the $\frac{1}{(\lambda\tau)^2}$ coefficients; we drop higher order terms that converge to 0 much faster.

$$\begin{aligned}
\text{SNR}(X) & \stackrel{\lambda\tau \rightarrow \infty}{=} \sqrt{\frac{1}{\frac{2 - \frac{1}{\lambda\tau} + \frac{1}{2} \frac{1}{(\lambda\tau)^2}}{2 - \frac{1}{\lambda\tau} + \frac{1}{3} \frac{1}{(\lambda\tau)^2}} - 1}} = \sqrt{\frac{1}{\frac{2 - \frac{1}{\lambda\tau} + \frac{1}{3} \frac{1}{(\lambda\tau)^2} + \frac{1}{6} \frac{1}{(\lambda\tau)^2}}{2 - \frac{1}{\lambda\tau} + \frac{1}{3} \frac{1}{(\lambda\tau)^2}} - 1}} \\
& = \sqrt{\frac{1}{1 + \frac{\frac{1}{6} \frac{1}{(\lambda\tau)^2}}{2 - \frac{1}{\lambda\tau} + \frac{1}{3} \frac{1}{(\lambda\tau)^2}} - 1}} \stackrel{\lambda\tau \rightarrow \infty}{=} \sqrt{\frac{1}{\frac{\frac{1}{6} \frac{1}{(\lambda\tau)^2}}{2}}} = 2\sqrt{3}\lambda\tau
\end{aligned}$$

A.2 p-thinning SNR approximation

As p tends to 0, the p-thinned periodic SNR tends towards the Poisson SNR. As before, we Taylor series expand coth and find where terms begin to differ to find $\lim_{p \rightarrow 0} \text{SNR}(X)$.

$$\begin{aligned}
\lim_{p \rightarrow 0} \text{SNR}(X) & = \lim_{p \rightarrow 0} \sqrt{\frac{2\lambda\tau}{1 + p \frac{1 + (1 - \frac{p}{\lambda\tau} + \frac{1}{2} (\frac{p}{\lambda\tau})^2 - \frac{1}{6} (\frac{p}{\lambda\tau})^3 + \dots)}{1 - (1 - \frac{p}{\lambda\tau} + \frac{1}{2} (\frac{p}{\lambda\tau})^2 - \frac{1}{6} (\frac{p}{\lambda\tau})^3 + \dots)}} - 2\lambda\tau}} \\
& = \lim_{p \rightarrow 0} \sqrt{\frac{2\lambda\tau}{1 + 2\lambda\tau \frac{2 - \frac{p}{\lambda\tau} + \frac{1}{2} (\frac{p}{\lambda\tau})^2 - \frac{1}{6} (\frac{p}{\lambda\tau})^3 + \dots}{2 - \frac{p}{\lambda\tau} + \frac{1}{3} (\frac{p}{\lambda\tau})^2 - \frac{1}{12} (\frac{p}{\lambda\tau})^3 - \dots}} - 2\lambda\tau}} \\
& = \lim_{p \rightarrow 0} \sqrt{\frac{2\lambda\tau}{1 + 2\lambda\tau \frac{2 - \frac{p}{\lambda\tau} + \frac{1}{3} (\frac{p}{\lambda\tau})^2 + \frac{1}{6} (\frac{p}{\lambda\tau})^2}{2 - \frac{p}{\lambda\tau} + \frac{1}{3} (\frac{p}{\lambda\tau})^2}} - 2\lambda\tau}} \\
& = \lim_{p \rightarrow 0} \sqrt{\frac{2\lambda\tau}{1 + 2\lambda\tau \left(1 + \frac{\frac{1}{6} (\frac{p}{\lambda\tau})^2}{2 - \frac{p}{\lambda\tau} + \frac{1}{3} (\frac{p}{\lambda\tau})^2}\right)} - 2\lambda\tau}} = \sqrt{2\lambda\tau}
\end{aligned}$$

A.3 d-thinning SNR expansion

We express $\text{SNR}_{\text{dthin}}(X)$ in terms of $\text{SNR}_{\text{poi}}(X)$, by considering $\text{SNR}_{\text{dthin}}(X) = \text{SNR}_{\text{poi}}(X)/g$ and expanding binomials:

$$\begin{aligned}
g^2 & = \frac{(1 + k\lambda\tau)^k + (k\lambda\tau)^k}{(1 + k\lambda\tau)^k - (k\lambda\tau)^k} - 2\lambda\tau = \frac{\sum_{j=0}^k \binom{k}{j} (k\lambda\tau)^j + (k\lambda\tau)^k}{\sum_{j=0}^k \binom{k}{j} (k\lambda\tau)^j - (k\lambda\tau)^k} - 2\lambda\tau \\
& = \frac{1 + (k^2 - 2)\lambda\tau + \dots + \frac{1}{6} k^{k-2} (k-1)(k+4)(\lambda\tau)^{k-2} + k^{k-1} (\lambda\tau)^{k-1}}{1 + k^2\lambda\tau + \dots + \frac{1}{2} k^{k-1} (k-1)(\lambda\tau)^{k-2} + k^k (\lambda\tau)^{k-1}}
\end{aligned}$$

By taking a fourth-order Taylor series approximation in the denominator of $\text{SNR}_{\text{dthin}}$,

$$\begin{aligned} \text{SNR}(X) &\approx \sqrt{\frac{2\lambda\tau}{1 + \frac{k^4 - 10k^2 + 9}{15(2\lambda\tau)^4} + \frac{-k^4 + 20k^2 - 19}{45(2\lambda\tau)^3} + \frac{k^2 - 1}{3(2\lambda\tau)^2} + \frac{k^2 - 1}{3(2\lambda\tau)}}} \\ &\approx \sqrt{2\lambda\tau / (1 + k^2/3(2\lambda\tau))} \quad \text{assuming } k \gg 1 \text{ and } \lambda\tau \gg 1 \end{aligned}$$

Appendix B

AER Transmitter Design Space

This appendix explores the router’s AER transmitter (AEXT) design space. From the implemented transmitter design described in Section 3.4.1, the designs described here are earlier iterations and listed in approximately reverse chronological order.

B.1 AEXT Control Data decomposed (CD)

In this design, the control and data are separated; there is a control tree and a data tree (cf. AEXT ASPR and AEXT PSAR with combined control and data).

B.1.1 AEXT CD noTW CYC

The transmitter (AEXT) control-data decomposed (CD) without tailword (noTW) cyclic signaling (CYC) design make more efficient use of the control signaling (relative to AEXT CD noTW) by doing away with the data enable/acknowledge entirely.

Radix 2 accounting (2047 intermediate nodes, 2048 leaf nodes):

intermediate nodes			
component	transistors/component	components/node	transistors/node
NODE	90	1	
total transistors/intermediate node			90
leaf nodes			
component	transistors/component	components/node	transistors/node
LEAF	74	1	74
total transistors/leaf node			74

(90 transistors/intermediate node * 2047 intermediate nodes + 74 transistors/leaf node * 2048 leaf nodes) / 4096 neurons = **82.0 transistors/neuron**

Radix 4 accounting (341 intermediate nodes, 1024 leaf nodes):

intermediate nodes			
component	transistors/component	components/node	transistors/node
NODE	274	1	274
total transistors/intermediate node			274
leaf nodes			
component	transistors/component	components/node	transistors/node
LEAF	218	1	218
total transistors/leaf node			218

(274 transistors/intermediate node * 341 intermediate nodes + 218 transistors/leaf node * 1024 leaf nodes) / 4096 neurons = **77.3 transistors/neuron**

AEXT CD noTW CYC NODE

Intermediate node of AEXT tree.

```

*[[c0 → po↑; [pi];
   w0↑; [¬pi]; u↑; w0↓; [pi];
   c0o↑; [¬c0]; po↓; [¬pi]; c0o↓; u↓
  [c1 → po↑; [pi];
   w1↑; [¬pi]; u↑; w1↓; [pi];
   c1o↑; [¬c1]; po↓; [¬pi]; c1o↓; u↓
]]

*[[c00 ∨ c10 ∨ w0 → p0↑; [¬pi]; c0o↓; [¬c00 ∧ ¬c10 ∧ ¬w0]; p0↓; [pi ∧ c0]; c0o↑
  [c01 ∨ c11 ∨ w1 → p1↑; [¬pi]; c1o↓; [¬c01 ∧ ¬c10 ∧ ¬w1]; p1↓; [pi ∧ c1]; c1o↑
]]

```

It's helpful to consider the projection of the HSE on to the parent control and data lines.

```

*[po↑; [pi];
  [¬pi]; [pi];
  (P↑; [¬pi]; P↓; [pi]) × (m - 1)
  po↓; [¬pi]
]

```

The first line propagates the child request up the tree and waits for the parents to acknowledge.

The second line is the node outputting a new head word.

The third line repeats $(m - 1)$ times where m is this node's level in the tree.

The fourth line propagates the child reset up the tree.

$$\begin{aligned}
\neg u \wedge (c0 \vee c1) &\quad \rightarrow po\uparrow \\
(c0o \wedge \neg c0) \vee (c1o \wedge \neg c1) &\rightarrow po\downarrow \\
\\
c0 \wedge pi \wedge \neg u &\rightarrow w0\uparrow & c1 \wedge pi \wedge \neg u &\rightarrow w1\uparrow \\
u &\rightarrow w0\downarrow & u &\rightarrow w1\downarrow \\
\\
(w0 \vee w1) \wedge \neg pi &\rightarrow u\uparrow \\
\neg c0o \wedge \neg c1o \wedge \neg po &\rightarrow u\downarrow \\
\\
c0 \wedge u \wedge pi \wedge \neg c1o &\rightarrow c0o\uparrow & c1 \wedge u \wedge pi \wedge \neg c0o &\rightarrow c1o\uparrow \\
\neg pi &\rightarrow c0o\downarrow & \neg pi &\rightarrow c1o\downarrow \\
\\
c00 \vee c10 \vee w0 &\rightarrow p0\uparrow & c01 \vee c11 \vee w1 &\rightarrow p1\uparrow \\
\neg c00 \wedge \neg c10 \wedge \neg w0 &\rightarrow p0\downarrow & \neg c01 \wedge \neg c11 \wedge \neg w1 &\rightarrow p1\downarrow
\end{aligned}$$

Radix 2 transistor approximate accounting:

rule	transistor count	comments
$c[0, 1]$	12	2-way arbiter
p_o	11	
$w[0, 1]$	16	
u	10	
$c[0, 1]_o$	18	
$p[0, 1]$	12	
total	79	

Radix 4 transistor approximate accounting:

rule	transistor count	comments
$c[0, 1, 2, 3]$	92	4-way unpipelined arbiter
p_o	19	
$w[0, 1, 2, 3]$	32	
u	10	
$c[0, 1, 2, 3]_o$	44	
$p[0, 1, 2, 3]$	40	
total	237	

CMOS-implementable PRS

$$\begin{array}{lll}
_u \wedge (c0 \vee c1) & \rightarrow _po\downarrow & \neg_po \rightarrow po\uparrow \\
(\neg_c0o \wedge \neg c0) \vee (\neg_c1o \wedge \neg c1) & \rightarrow _po\uparrow & _po \rightarrow po\downarrow \\
\\
c0 \wedge pi \wedge _u & \rightarrow _w0\downarrow & c1 \wedge pi \wedge _u \rightarrow _w1\downarrow \\
\neg_u & \rightarrow _w0\uparrow & \neg_u \rightarrow _w1\uparrow \\
\\
(\neg_w0 \vee \neg_w1) \wedge \neg pi & \rightarrow u\uparrow & \neg u \rightarrow _u\uparrow \\
_c0o \wedge _c1o \wedge _po & \rightarrow u\downarrow & u \rightarrow _u\downarrow \\
\\
c0 \wedge u \wedge pi \wedge _c1o & \rightarrow _c0o\downarrow & c1 \wedge u \wedge pi \wedge _c0o \rightarrow _c1o\downarrow \\
\neg pi & \rightarrow _c0o\uparrow & \neg pi \rightarrow _c1o\uparrow \\
\\
\neg_c0o \rightarrow c0o\uparrow & \neg_c1o \rightarrow c1o\uparrow & \\
_c0o \rightarrow c0o\downarrow & _c1o \rightarrow c1o\downarrow & \\
\\
\neg_c00 \vee \neg_c10 \vee \neg_w0 & \rightarrow p0\uparrow & \neg_c01 \vee \neg_c11 \vee \neg_w1 \rightarrow p1\uparrow \\
_c00 \wedge _c10 \wedge _w0 & \rightarrow p0\downarrow & _c01 \wedge _c11 \wedge _w1 \rightarrow p1\downarrow \\
\\
\neg p0 \rightarrow _p0\uparrow & \neg p1 \rightarrow _p1\uparrow & \\
p0 \rightarrow _p0\downarrow & p1 \rightarrow _p1\downarrow &
\end{array}$$

Note that the root NODE does not create $_p[0, 1]$. We simply present a normal-sense p_i , p_o , and $p[0, 1]$ interface to the environment.

We could make another another version of CMOS-implementable PRS to alternate with this version and eliminate the inverters creating p_o , $c[0, 1]_o$, and $p[0, 1]$. However, that version would have a PMOS pull up chain for $c[0, 1]_o$ that doesn't scale. With a radix 2 tree the chain is already 4 transistors long. With a radix 4 tree the chain is 6 transistors long. We want to keep PMOS chains 3 long or shorter.

Radix 2 transistor accounting:

rule	transistor count	comments
$c[0, 1]$	12	2-way arbiter
$-p_o$	11	
p_o	2	
$-w[0, 1]$	16	
u	10	
$-u$	2	
$-c[0, 1]_o$	18	
$c[0, 1]_o$	4	
$p[0, 1]$	12	
$-p[0, 1]$	4	
total	91	

Radix 4 transistor accounting:

rule	transistor count	comments
$c[0, 1, 2, 3]$	92	4-way unpipelined arbiter
$-p_o$	19	
p_o	2	
$-w[0, 1, 2, 3]$	32	
u	10	
$-u$	2	
$-c[0, 1, 2, 3]_o$	44	
$c[0, 1, 2, 3]_o$	8	
$p[0, 1, 2, 3]$	40	
$-p[0, 1, 2, 3]$	8	
total	257	

AEXT CD noTW CYC NODE (reference implementation)

Intermediate node of AEXT tree

$$\begin{aligned}
& * [[c0 \longrightarrow q0\uparrow; po\uparrow; [pi]; \\
& \quad w0\uparrow; [\neg pi]; u\uparrow; w0\downarrow; [pi]; \\
& \quad c0o\uparrow; [\neg c0]; q0\downarrow; po\downarrow; u\downarrow; [\neg pi]; c0o\downarrow \\
& \quad \llbracket c1 \longrightarrow q1\uparrow; po\uparrow; [pi]; \\
& \quad \quad w1\uparrow; [\neg pi]; u\uparrow; w1\downarrow; [pi]; \\
& \quad \quad c1o\uparrow; [\neg c1]; q1\downarrow; po\downarrow; u\downarrow; [\neg pi]; c1o\downarrow \\
& \quad \rrbracket \\
&]] \\
& * [[c00 \vee c10 \vee w0 \longrightarrow p0\uparrow; [\neg pi]; c0o\downarrow; [\neg c00 \wedge \neg c10 \wedge \neg w0]; p0\downarrow; [pi \wedge q0]; c0o\uparrow \\
& \quad \llbracket c01 \vee c11 \vee w1 \longrightarrow p1\uparrow; [\neg pi]; c1o\downarrow; [\neg c01 \wedge \neg c11 \wedge \neg w1]; p1\downarrow; [pi \wedge q1]; c1o\uparrow \\
& \quad \rrbracket \\
&]] \\
& c0 \wedge \neg c1o \rightarrow q0\uparrow \quad c1 \wedge \neg c0o \rightarrow q1\uparrow \\
& \neg c0 \wedge c0o \rightarrow q0\downarrow \quad \neg c1 \wedge c1o \rightarrow q1\downarrow \\
& q0 \vee q1 \rightarrow po\uparrow \\
& \neg q0 \wedge \neg q1 \rightarrow po\downarrow \\
& q0 \wedge pi \wedge \neg u \rightarrow w0\uparrow \quad q1 \wedge pi \wedge \neg u \rightarrow w1\uparrow \\
& u \rightarrow w0\downarrow \quad u \rightarrow w1\downarrow \\
& (w0 \vee w1) \wedge \neg pi \rightarrow u\uparrow \\
& \neg po \rightarrow u\downarrow \\
& q0 \wedge u \wedge pi \rightarrow c0o\uparrow \quad \neg q1 \wedge u \wedge pi \rightarrow c1o\uparrow \\
& (\neg u \vee p0 \vee p1) \wedge \neg pi \rightarrow c0o\downarrow \quad (\neg u \vee p0 \vee p1) \wedge \neg pi \rightarrow c1o\downarrow \\
& c00 \vee c10 \vee w0 \rightarrow p0\uparrow \quad c01 \vee c11 \vee w1 \rightarrow p1\uparrow \\
& \neg c00 \wedge \neg c10 \wedge \neg w0 \rightarrow p0\downarrow \quad \neg c01 \wedge \neg c11 \wedge \neg w1 \rightarrow p1\downarrow
\end{aligned}$$

Radix 2 transistor accounting:

rule	transistor count	comments
$c[0, 1]$	12	2-way arbiter
$q[0, 1]$	16	
po	4	
$w[0, 1]$	16	
u	8	
$c[0, 1]_o$	22	
$p[0, 1]$	12	
total	90	

Radix 4 transistor accounting:

rule	transistor count	comments
$c[0, 1, 2, 3]$	92	4-way unpipelined arbiter
$q[0, 1, 2, 3]$	40	
p_o	8	
$w[0, 1, 2, 3]$	32	
u	10	
$c[0, 1, 2, 3]_o$	52	
$p[0, 1, 2, 3]$	40	
total	274	

Radix 2 transistor accounting:

AEXT CD noTW CYC LEAF

Leaf node of AEXT tree

```

*[[c0 → po↑; [pi];
   p0↑; [¬pi]; u↑; p0↓; [pi];
   c0o↑; [¬c0]; po↓; [¬pi]; c0o↓; u↓
  [c1 → po↑; [pi];
   p1↑; [¬pi]; u↑; p1↓; [pi];
   c1o↑; [¬c1]; po↓; [¬pi]; c1o↓; u↓
  ]]

```

PRS

$$\begin{aligned} \neg u \wedge c0 \vee c1 &\quad \rightarrow \text{po}\uparrow \\ (c0o \wedge \neg c0) \vee (c1o \wedge \neg c1) &\quad \rightarrow \text{po}\downarrow \\ \\ c0 \wedge pi \wedge \neg u &\quad \rightarrow \text{p0}\uparrow & c1 \wedge pi \wedge \neg u &\quad \rightarrow \text{p1}\uparrow \\ u &\quad \rightarrow \text{p0}\downarrow & u &\quad \rightarrow \text{p1}\downarrow \\ \\ (p0 \vee p1) \wedge \neg pi &\quad \rightarrow u\uparrow \\ \neg c0o \wedge \neg c1o \wedge \neg po &\quad \rightarrow u\downarrow \\ \\ c0 \wedge u \wedge pi \wedge \neg c1o &\quad \rightarrow c0o\uparrow & c1 \wedge u \wedge pi \wedge \neg c0o &\quad \rightarrow c1o\uparrow \\ \neg u \wedge \neg pi &\quad \rightarrow c0o\downarrow & \neg u \wedge \neg pi &\quad \rightarrow c1o\downarrow \end{aligned}$$

Radix 2 transistor approximate accounting:

rule	transistor count	comments
$c[0, 1]$	12	2-way arbiter
p_o	11	
$p[0, 1]$	16	
u	10	
$c[0, 1]_o$	20	
total	69	

Radix 4 transistor approximate accounting:

rule	transistor count	comments
$c[0, 1, 2, 3]$	92	4-way unpipelined arbiter
p_o	17	
$p[0, 1, 2, 3]$	32	
u	14	
$c[0, 1, 2, 3]_o$	48	
total	203	

CMOS-implementable PRS

$$\begin{aligned} _u \wedge c0 \vee c1 & \rightarrow _po \downarrow & \neg_po & \rightarrow po \uparrow \\ (\neg_c0o \wedge \neg c0) \vee (\neg_c1o \wedge \neg c1) & \rightarrow _po \uparrow & _po & \rightarrow po \downarrow \end{aligned}$$

$$\begin{aligned} c0 \wedge pi \wedge _u & \rightarrow _p0 \downarrow & c1 \wedge pi \wedge _u & \rightarrow _p1 \downarrow \\ \neg_u & \rightarrow _p0 \uparrow & \neg_u & \rightarrow _p1 \downarrow \end{aligned}$$

$$\begin{aligned} (\neg_p0 \vee \neg_p1) \wedge \neg pi & \rightarrow u \uparrow & \neg u & \rightarrow _u \uparrow \\ _c0o \wedge _c1o \wedge _po & \rightarrow u \downarrow & u & \rightarrow _u \downarrow \end{aligned}$$

$$\begin{aligned} c0 \wedge u \wedge pi \wedge _c1o & \rightarrow _c0o \downarrow & \neg_c0o & \rightarrow c0o \uparrow \\ \neg u \wedge \neg pi & \rightarrow _c0o \uparrow & _c0o & \rightarrow c0o \downarrow \end{aligned}$$

$$\begin{aligned} c1 \wedge u \wedge pi \wedge _c0o & \rightarrow _c1o \downarrow & \neg_c1o & \rightarrow c1o \uparrow \\ \neg u \wedge \neg pi & \rightarrow _c1o \uparrow & _c1o & \rightarrow c1o \downarrow \end{aligned}$$

Radix 2 transistor accounting:

rule	transistor count	comments
$c[0, 1]$	12	2-way arbiter
$-p_o$	11	
p_o	2	
$-p[0, 1]$	16	
u	10	
$-u$	2	
$-c[0, 1]_o$	20	
$c[0, 1]_o$	4	
total	77	

Radix 4 transistor accounting:

rule	transistor count	comments
$c[0, 1, 2, 3]$	92	4-way unpipelined arbiter
$-p_o$	17	
p_o	2	
$-p[0, 1, 2, 3]$	32	
u	14	
$-u$	2	
$-c[0, 1, 2, 3]_o$	48	
$c[0, 1, 2, 3]_o$	8	
total	215	

AEXT CD noTW CYC LEAF (reference implementation)

Leaf node of AEXT tree

```

*[[c0 → q0↑; po↑; [pi];
   p0↑; [¬pi]; u↑; p0↓; [pi];
   c0o↑; [¬c0]; q0↓; po↓; u↓; [¬pi]; c0o↓
  [c1 → q1↑; po↑; [pi];
   p1↑; [¬pi]; u↑; p1↓; [pi];
   c1o↑; [¬c1]; q1↓; po↓; u↓; [¬pi]; c1o↓
  ]]

```

$$c0 \wedge \neg c1o \rightarrow q0\uparrow \qquad c1 \wedge \neg c0o \rightarrow q1\uparrow$$

$$\neg c0 \wedge c0o \rightarrow q0\downarrow \qquad \neg c1 \wedge c1o \rightarrow q1\downarrow$$

$$q0 \vee q1 \rightarrow po\uparrow$$

$$\neg q0 \wedge \neg q1 \rightarrow po\downarrow$$

$$\begin{array}{ll} q0 \wedge pi \wedge \neg u \rightarrow p0\uparrow & q1 \wedge pi \wedge \neg u \rightarrow p1\uparrow \\ u \rightarrow p0\downarrow & u \rightarrow p1\downarrow \end{array}$$

$$\begin{array}{ll} (p0 \vee p1) \wedge \neg pi \rightarrow u\uparrow \\ \neg po \rightarrow u\downarrow \end{array}$$

$$\begin{array}{ll} q0 \wedge u \wedge pi \rightarrow c0o\uparrow & q1 \wedge u \wedge pi \rightarrow c1o\uparrow \\ \neg u \wedge \neg pi \rightarrow c0o\downarrow & \neg u \wedge \neg pi \rightarrow c1o\downarrow \end{array}$$

Radix 2 transistor accounting:

rule	transistor count	comments
$c[0, 1]$	12	2-way arbiter
$q[0, 1]$	16	
p_o	4	
$p[0, 1]$	16	
u	8	
$c[0, 1]_o$	18	
total	74	

Radix 4 transistor accounting:

rule	transistor count	comments
$c[0, 1, 2, 3]$	92	4-way unpipelined arbiter
$q[0, 1, 2, 3]$	40	
p_o	8	
$p[0, 1, 2, 3]$	32	
u	10	
$c[0, 1, 2, 3]_o$	36	
total	218	

B.1.2 AEXT CD noTW

This design has no tail word.

Radix 2 accounting (2047 intermediate nodes, 2048 leaf nodes):

intermediate nodes			
component	transistors/component	components/node	transistors/node
CTRL	86	1	86
MERGE	28	1	28
total transistors/intermediate node			114
leaf nodes			
component	transistors/component	components/node	transistors/node
LEAF	64	1	64
total transistors/leaf node			64

(114 transistors/intermediate node * 2047 intermediate nodes + 64 transistors/leaf node * 2048 leaf nodes) / 4096 neurons = **89.0 transistors/neuron**

Radix 4 accounting (341 intermediate nodes, 1024 leaf nodes):

intermediate nodes			
component	transistors/component	components/node	transistors/node
CTRL	246	1	246
MERGE	60	1	60
total transistors/intermediate node			306
leaf nodes			
component	transistors/component	components/node	transistors/node
LEAF	204	1	204
total transistors/leaf node			204

(306 transistors/intermediate node * 341 intermediate nodes + 204 transistors/leaf node * 1024 leaf nodes) / 4096 neurons = **76.5 transistors/neuron**

AEXT CD noTW CTRL

```

*[[C0 • P; M; C0 • P
  | C1 • P; M; C1 • P
]]

*[[c0 → po↑; [pi]; c0o↑;
  mw↑; [we]; w0↑; [¬we]; w0↓; mw↓;
  m0↑; [¬c0]; po↓; [¬pi]; c0o↓; m0↓
 [c1 → po↑; [pi]; c1o↑;
  mw↑; [we]; w1↑; [¬we]; w1↓; mw↓;
  m1↑; [¬c1]; po↓; [¬pi]; c1o↓; m1↓
]]

```

$$\begin{aligned}
& * [[c0 \longrightarrow q0\uparrow; po\uparrow; [pi]; \\
& \quad mw\uparrow; c0o\uparrow; [we]; w0\uparrow; [\neg we]; m0\uparrow; mw\downarrow; w0\downarrow \\
& \quad [\neg c0]; q0\downarrow; po\downarrow; [\neg pi]; m0\downarrow; c0o\downarrow \\
& \quad] [c1 \longrightarrow q1\uparrow; po\uparrow; [pi]; \\
& \quad \quad mw\uparrow; c1o\uparrow; [we]; w1\uparrow; [\neg we]; m1\uparrow; mw\downarrow; w1\downarrow \\
& \quad \quad [\neg c1]; q1\downarrow; po\downarrow; [\neg pi]; m1\downarrow; c1o\downarrow \\
& \quad]]
\end{aligned}$$

sequence

$$mx\downarrow; cxo\downarrow$$

required to lower control before releasing child to lower word line.

$$\begin{aligned}
c0 \wedge \neg m0 \wedge \neg c1o &\rightarrow q0\uparrow & c1 \wedge \neg m1 \wedge \neg c0o &\rightarrow q1\uparrow \\
\neg c0 \wedge m0 \wedge \neg w0 &\rightarrow q0\downarrow & \neg c1 \wedge m1 \wedge \neg w1 &\rightarrow q1\downarrow
\end{aligned}$$

$$\begin{aligned}
q0 \vee q1 &\rightarrow po\uparrow \\
\neg q0 \wedge \neg q1 &\rightarrow po\downarrow
\end{aligned}$$

$$\begin{aligned}
q0 \wedge mw &\rightarrow c0o\uparrow & q1 \wedge mw &\rightarrow c1o\uparrow \\
\neg m0 \wedge \neg pi &\rightarrow c0o\downarrow & \neg m1 \wedge \neg pi &\rightarrow c1o\downarrow
\end{aligned}$$

$$\begin{aligned}
c0o \wedge we &\rightarrow w0\uparrow & c1o \wedge we &\rightarrow w1\uparrow \\
\neg mw &\rightarrow w0\downarrow & \neg mw &\rightarrow w1\downarrow
\end{aligned}$$

$$\begin{aligned}
pi \wedge \neg m0 \wedge \neg m1 &\rightarrow mw\uparrow \\
\neg pi \vee m0 \vee m1 &\rightarrow mw\downarrow
\end{aligned}$$

$$\begin{aligned}
w0 \wedge \neg we &\rightarrow m0\uparrow & w1 \wedge \neg we &\rightarrow m1\uparrow \\
\neg pi &\rightarrow m0\downarrow & \neg pi &\rightarrow m1\downarrow
\end{aligned}$$

Radix 2 transistor accounting:

rule	transistor count	comments
$c[0, 1]$	12	2-input arbiter
$q[0, 1]$	20	
po	4	
$c[0, 1]_o$	16	
$w[0, 1]$	14	
mw	6	
$m[0, 1]$	14	
total	86	

Radix 4 transistor accounting:

rule	transistor count	comments
$c[0, 1, 2, 3]$	92	4-input arbiter
$q[0, 1, 2, 3]$	48	
p_o	8	
$c[0, 1, 2, 3]_o$	32	
$w[0, 1, 2, 3]$	28	
mw	10	
$m[0, 1, 2, 3]$	28	
total	246	

AEXT CD noTW MERGE

```

*[[pe];
  [m0 → c0e↑;
    [c00 → p0↑; [¬pe]; c0e↓; [¬c00]; p0↓
    []c01 → p1↑; [¬pe]; c0e↓; [¬c01]; p1↓]
  []m1 → c1e↑;
    [c10 → p0↑; [¬pe]; c1e↓; [¬c10]; p0↓
    []c11 → p1↑; [¬pe]; c1e↓; [¬c11]; p1↓]
  []mw → cwe↑
    [cw0 → p0↑; [¬pe]; cwe↓; [¬cw0]; p0↓
    []cw1 → p1↑; [¬pe]; cwe↓; [¬cw1]; p1↓]
  ]]

pe ∧ m0    → c0e↑      pe ∧ mw    → cwe↑
¬pe ∨ ¬m0  → c0e↓      ¬pe ∨ ¬mw  → cwe↓

pe ∧ m1    → c1e↑
¬pe ∨ ¬m1  → c1e↓

c00 ∨ c10 ∨ cw0    → p0↑      c01 ∨ c11 ∨ cw1    → p1↑
¬c00 ∧ ¬c10 ∧ ¬cw0 → p0↓      ¬c01 ∧ ¬c11 ∧ ¬cw1 → p1↓

```

Radix 2 transistor accounting:

rule	transistor count	comments
$c[0, 1, w]e$	12	
$p[0, 1]$	12	
total	24	

Radix 4 transistor accounting:

rule	transistor count	comments
$c[0, 1, 2, 3, w]e$	20	
$p[0, 1, 2, 3]$	40	
total	60	

AEXT CD noTW LEAF

```

*[[C0 • P; W!0; C0 • P
  | C1 • P; W!1; C1 • P
]]

*[[c0 → po↑; [pi]; c0o↑;
   [we]; w0↑; [¬we]; w0↓
   [¬c0]; po↓; [¬pi]; c0o↓
  ||c1 → po↑; [pi]; c1o↑;
   [we]; w1↑; [¬we]; w1↓
   [¬c1]; po↓; [¬pi]; c1o↓
]]

*[[c0 → q0↑; po↑; [pi]; c0o↑;
   [¬c0 ∧ we]; w0↑; [¬we]; q0↓; po↓; [¬pi]; w0↓;
   c0o↓
  ||c1 → q1↑; po↑; [pi]; c1o↑;
   [¬c1 ∧ we]; w1↑; [¬we]; q1↓; po↓; [¬pi]; w1↓;
   c1o↓
]]

```

The sequence of

$$[\neg we]; po\downarrow; [\neg pi]; wx\downarrow;$$

is very important. Parent nodes need to reset control of merge before lowering word line.

$$\begin{array}{ll}
c0 \wedge \neg c1o \rightarrow q0\uparrow & c0 \wedge \neg c1o \rightarrow q1\uparrow \\
w0 \wedge \neg we \rightarrow q0\downarrow & w1 \wedge \neg we \rightarrow q1\downarrow \\
q0 \vee q1 \rightarrow po\uparrow & \\
\neg q0 \wedge \neg q1 \rightarrow po\downarrow & \\
q0 \wedge \neg c0 \wedge we \rightarrow w0\uparrow & q1 \wedge \neg c1 \wedge we \rightarrow w1\uparrow \\
\neg pi \rightarrow w0\downarrow & \neg pi \rightarrow w1\downarrow
\end{array}$$

$$\begin{array}{ll}
q0 \wedge pi & \rightarrow c0o\uparrow \\
\neg w0 \wedge \neg pi & \rightarrow c0o\downarrow
\end{array}
\qquad
\begin{array}{ll}
q1 \wedge pi & \rightarrow c1o\uparrow \\
\neg w1 \wedge \neg pi & \rightarrow c1o\downarrow
\end{array}$$

Radix 2 transistor accounting:

rule	transistor count	comments
$c[0, 1]$	12	2-input arbiter
$q[0, 1]$	16	
p_o	4	
$w[0, 1]$	16	
$c[0, 1]_o$	16	
total	64	

Radix 4 transistor accounting:

rule	transistor count	comments
$c[0, 1, 2, 3]$	92	4-input arbiter
$q[0, 1, 2, 3]$	40	
p_o	8	
$w[0, 1, 2, 3]$	32	
$c[0, 1, 2, 3]_o$	32	
total	204	

B.1.3 AEXT CD TW

This design has a tail word.

Radix 2 accounting (4095 nodes / 4096 neurons):

component	transistors/component	components/node	transistors/node
CTRL	84	1	84
MERGE	28	1	28
FWDT	15	2	30
total transistors/node			142

$$142 \text{ transistors/node} * 4095 \text{ nodes} / 4096 \text{ neurons} = 142.0 \text{ transistors/neuron}$$

We also need

$$8 \text{ transistors} / \text{INT} * 1 \text{ INT} / \text{neuron} = 8 \text{ transistors/neuron}$$

This gives us

$$142 + 8 = \mathbf{150 \text{ transistors/neuron}}$$

Radix 4 transistor accounting (1365 nodes / 4096 neurons):

component	transistors/component	components/node	transistors/node
CTRL	238	1	238
MERGE	68	1	68
FWDT	15	4	60
total transistors/node			366

366 transistors/node * 1365 nodes / 4096 neurons = 122.0 transistors/neuron

We also need

8 transistors / INT * 1 INT / neuron = 8 transistors/neuron

This gives us

122 + 8 = **130 transistors/neuron**

AEXT CD TW CTRL

Control.

```

*[[ $\overline{C0} \rightarrow P; mh \bullet H!0; m0 \bullet F0; C0$ 
  |  $\overline{C1} \rightarrow P; mh \bullet H!1; m1 \bullet F1; C1$ 
]]

*[[ $c0 \rightarrow po\uparrow; [pi];$ 
   $mw\uparrow; [we]; w0\uparrow; [\neg we]; mw\downarrow; w0\downarrow;$ 
   $m0\uparrow; f0o\uparrow; [f0i]; m0\downarrow; f0o\downarrow; [\neg f0i]$ 
   $c0o\uparrow; [\neg c0]; po\downarrow; [\neg pi]; c0o\downarrow$ 
   $\llbracket c1 \rightarrow po\uparrow; [pi];$ 
   $mw\uparrow; [we]; w1\uparrow; [\neg we]; mw\downarrow; w1\downarrow;$ 
   $m1\uparrow; f1o\uparrow; [f1i]; m1\downarrow; f1o\downarrow; [\neg f1i]$ 
   $c1o\uparrow; [\neg c1]; po\downarrow; [\neg pi]; c1o\downarrow$ 
]]

*[[ $c0 \rightarrow po\uparrow; [pi];$ 
   $(c0o\uparrow; [\neg c0]),$ 
   $(mw\uparrow; [we]; w0\uparrow; [\neg we]; m0\uparrow; mw\downarrow; w0\downarrow;$ 
   $f0o\uparrow; [f0i]; po\downarrow; [\neg pi]);$ 
   $m0\downarrow; f0o\downarrow; c0o\downarrow; [\neg f0i]$ 
   $\llbracket c1 \rightarrow po\uparrow; [pi];$ 
   $(c1o\uparrow; [\neg c1]),$ 
   $(mw\uparrow; [we]; w1\uparrow; [\neg we]; m1\uparrow; mw\downarrow; w1\downarrow;$ 
   $f1o\uparrow; [f1i]; po\downarrow; [\neg pi]);$ 
   $m1\downarrow; f1o\downarrow; c1o\downarrow; [\neg f1i]$ 
]]

```

$$(c0 \vee c1) \wedge \neg f1i \wedge \neg f0i \rightarrow po\uparrow$$

$$f1i \vee f0i \rightarrow po\downarrow$$

$$pi \wedge \neg m0 \wedge \neg m1 \rightarrow mw\uparrow \quad w1 \wedge \neg we \rightarrow m1\uparrow$$

$$\neg pi \vee m0 \vee m1 \rightarrow mw\downarrow \quad \neg pi \wedge \neg c1 \rightarrow m1\downarrow$$

$$w0 \wedge \neg we \rightarrow m0\uparrow$$

$$\neg pi \wedge \neg c0 \rightarrow m0\downarrow$$

$$c0o \wedge we \rightarrow w0\uparrow \quad c1o \wedge we \rightarrow w1\uparrow$$

$$\neg mw \rightarrow w0\downarrow \quad \neg mw \rightarrow w1\downarrow$$

$$m0 \wedge \neg w0 \rightarrow f0o\uparrow \quad m1 \wedge \neg w1 \rightarrow f1o\uparrow$$

$$\neg m0 \vee w0 \rightarrow f0o\downarrow \quad \neg m1 \vee w1 \rightarrow f1o\downarrow$$

$$c0 \wedge pi \wedge \neg c1o \rightarrow c0o\uparrow \quad c1 \wedge pi \wedge \neg c0o \rightarrow c1o\uparrow$$

$$\neg f0o \wedge \neg pi \rightarrow c0o\downarrow \quad \neg f1o \wedge \neg pi \rightarrow c1o\downarrow$$

Radix 2 transistor accounting:

rule	transistor count	comments
$c[0, 1]$	12	2-input arbiter
po	10	
mw	6	
$m[0, 1]$	16	
$w[0, 1]$	14	
$f[0, 1]_o$	8	
$c[0, 1]_o$	18	
total	84	

Radix 4 transistor accounting:

rule	transistor count	comments
$c[0, 1, 2, 3]$	92	4-input arbiter
po	16	
mw	10	
$m[0, 1, 2, 3]$	32	
$w[0, 1, 2, 3]$	28	
$f[0, 1, 2, 3]_o$	16	
$c[0, 1, 2, 3]_o$	44	
total	238	

AEXT CD TW MERGE

Controlled merge.

$$\begin{aligned}
& *[[\overline{M0} \longrightarrow * [P!(C0?)]] \\
& \quad \boxed{\overline{M1} \longrightarrow * [P!(C1?)]} \\
& \quad \boxed{\overline{Mh} \longrightarrow P!(H?)} \\
& \quad]] \\
& *[[pe]; \\
& \quad [m0 \longrightarrow c0e\uparrow; \\
& \quad \quad [c00 \longrightarrow p0\uparrow; [\neg pe]; c0e\downarrow; [\neg c00]; p0\downarrow \\
& \quad \quad \boxed{c01 \longrightarrow p1\uparrow; [\neg pe]; c0e\downarrow; [\neg c01]; p1\downarrow} \\
& \quad \quad \boxed{c0t \longrightarrow pt\uparrow; [\neg pe]; c0e\downarrow; [\neg c0t]; pt\downarrow}] \\
& \quad [m1 \longrightarrow c1e\uparrow; \\
& \quad \quad [c10 \longrightarrow p0\uparrow; [\neg pe]; c1e\downarrow; [\neg c10]; p0\downarrow \\
& \quad \quad \boxed{c11 \longrightarrow p1\uparrow; [\neg pe]; c1e\downarrow; [\neg c11]; p1\downarrow} \\
& \quad \quad \boxed{c1t \longrightarrow pt\uparrow; [\neg pe]; c1e\downarrow; [\neg c1t]; pt\downarrow}] \\
& \quad [mw \longrightarrow cwe\uparrow \\
& \quad \quad [cw0 \longrightarrow p0\uparrow; [\neg pe]; cwe\downarrow; [\neg cw0]; p0\downarrow \\
& \quad \quad \boxed{cw1 \longrightarrow p1\uparrow; [\neg pe]; cwe\downarrow; [\neg cw1]; p1\downarrow}] \\
& \quad]]
\end{aligned}$$

$$\begin{aligned}
pe \wedge m0 & \longrightarrow c0e\uparrow & pe \wedge mw & \longrightarrow cwe\uparrow \\
\neg pe \vee \neg m0 & \longrightarrow c0e\downarrow & \neg pe \vee \neg mw & \longrightarrow cwe\downarrow
\end{aligned}$$

$$\begin{aligned}
pe \wedge m1 & \longrightarrow c1e\uparrow \\
\neg pe \vee \neg m1 & \longrightarrow c1e\downarrow
\end{aligned}$$

$$\begin{aligned}
c00 \vee c10 \vee cw0 & \longrightarrow p0\uparrow & c0t \vee c1t & \longrightarrow pt\uparrow \\
\neg c00 \wedge \neg c10 \wedge \neg cw0 & \longrightarrow p0\downarrow & \neg c0t \wedge \neg c1t & \longrightarrow pt\downarrow
\end{aligned}$$

$$\begin{aligned}
c01 \vee c11 \vee cw1 & \longrightarrow p1\uparrow \\
\neg c01 \wedge \neg c11 \wedge \neg cw1 & \longrightarrow p1\downarrow
\end{aligned}$$

Radix 2 transistor accounting:

rule	transistor count	comments
$c[0, 1, w]e$	12	
$p[0, 1]$	12	
pt	4	
total	28	

Radix 4 transistor accounting:

rule	transistor count	comments
$c[0, 1, 2, 3, w]e$	20	
$p[0, 1, 2, 3]$	40	
pt	8	
total	68	

AEXT CD TW WORD

Output a word.

```

* $[Y!(X?)]$ 

* $[[ye]; xe\uparrow;$ 
   $[x0 \rightarrow y0\uparrow; \neg ye]; xe\downarrow; [\neg x0]; y0\downarrow$ 
   $][x1 \rightarrow y1\uparrow; \neg ye]; xe\downarrow; [\neg x1]; y1\downarrow$ 
 $]]$ 

 $ye \rightarrow xe\uparrow$ 
 $\neg ye \rightarrow xe\downarrow$ 

 $x0 \rightarrow y0\uparrow$        $x1 \rightarrow y1\uparrow$ 
 $\neg x0 \rightarrow y0\downarrow$    $\neg x1 \rightarrow y1\downarrow$ 

```

Radix 2 transistor accounting:

rule	transistor count	comments
total	0	all wires

Radix 4 transistor accounting:

rule	transistor count	comments
total	0	all wires

AEXT CD TW FWDT

Forward and detect tail.

```

* $[C\uparrow;$ 
   $[X = 0 \rightarrow Y!X$ 
   $][X = 1 \rightarrow Y!X$ 
   $][X = t \rightarrow Y!X; C\downarrow$ 
 $]]$ 

```

```

*[[ci ∧ ye]; xe↑;
  [x0 → y0↑; [¬ye]; xe↓; [¬x0]; y0↓
  [x1 → y1↑; [¬ye]; xe↓; [¬x1]; y1↓
  [xt → yt↑; [¬ye]; xe↓; co↑; [¬xt ∧ ¬ci]; yt↓; co↓
]]

```

how do I express this in CHP?

$$\begin{array}{ll} ci \wedge ye & \rightarrow xe \uparrow & yt \wedge \neg xe & \rightarrow co \uparrow \\ \neg ci \vee \neg ye & \rightarrow xe \downarrow & \neg yt \vee xe & \rightarrow co \downarrow \end{array}$$

$$\begin{array}{ll} x0 & \rightarrow y0 \uparrow & xt & \rightarrow yt \uparrow \\ \neg x0 & \rightarrow y0 \downarrow & \neg xt \wedge \neg ci & \rightarrow yt \downarrow \end{array}$$

$$\begin{array}{l} x1 \rightarrow y1 \uparrow \\ \neg x1 \rightarrow y1 \downarrow \end{array}$$

Radix 2 transistor accounting:

rule	transistor count	comments
xe	4	
co	4	
$y[0, 1]$	0	wires
yt	7	
total	15	

Radix 4 transistor accounting:

rule	transistor count	comments
xe	4	
co	4	
$y[0, 1, 2, 3]$	0	wires
yt	7	
total	15	

AEXT CD TW INT

Neuron interface.

```
*[N; C; D]
```

```
*[[ni]; co↑; [ci ∧ de]; dt↑; no↑;
  [¬ni]; co↓; [¬ci ∧ ¬de]; dt↓; no↓]
```

$$\begin{array}{ll}
ni \rightarrow co\uparrow & dt \rightarrow no\uparrow \\
\neg ni \rightarrow co\downarrow & \neg dt \rightarrow no\downarrow \\
\\
ci \wedge de \rightarrow dt\uparrow & \\
\neg ci \wedge \neg de \rightarrow dt\downarrow &
\end{array}$$

Transistor accounting:

rule	transistor count	comments
c_o	0	wires
dt	8	
n_o	0	wires
total	8	

B.2 AEXT Control Data Combined

In these designs, the control and data are not decomposed.

B.3 AEXT ASPR NODE

The active-sender, passive-receiver (ASPR) design assumes that the children actively send data to their passively receiving parents. The following describes a monolithic node process encapsulating that idea.

```

NODE ≡
*[[h →
  [C0 → s := 0, P!(0);
  | C1 → s := 1, P!(1)];
  h := false
[] ¬h →
  [s = 0 → C0?x; P!x
  [] s = 1 → C1?x; P!x
  ]; h := x.tail
]]

```

B.4 AEXT ASPR PFWD/MERGE (PM)

The monolithic NODE can be decomposed into PFWD, which prepends a word to the packet indicating which branch the packet is coming from and MERGE processes, and MERGE, which arbitrates between incoming packet streams and outputs them one at a time.

Radix 2 accounting (4095 nodes / 4096 neurons):

component	transistors/component	components/node	transistors/node
PFWD	53	2	106
MERGE	90	1	90
total transistors/node			196

196 transistors/node * 4095 nodes / 4096 neurons = **196.0 transistors/neuron**

The leaf node PFWDs only need to communicate the tail bit and their prepend bit and can leave off the other bit. This saves 2 production rules, or 14 transistors per node. With 2048 leaf nodes, this saves us 28672 transistors.

(196 transistors/node * 4095 nodes - 28672) / 4096 neurons = **189.0 transistors/neuron**

Radix 4 transistor accounting (1365 nodes / 4096 neurons):

component	transistors/component	components/node	transistors/node
PFWD	73	4	292
MERGE	288	1	288
total transistors/node			580

580 transistors/node * 1365 nodes / 4096 neurons = **193.3 transistors/neuron**

The leaf node PFWDs only need to communicate the tail bit and their prepend bit and can leave off the other bit. This saves 4 production rules, or 32 transistors per node. With 1024 leaf nodes, this saves us 32768 transistors.

(698 transistors/node * 4095 nodes - 32768) / 4096 neurons = **185.3 transistors/neuron**

B.5 AEXT ASPR PM PFWD unpipelined

HSE

$$\begin{aligned}
 & * [h \wedge (x0 \vee x1 \vee xt) \longrightarrow q\uparrow; yp\uparrow; [yi]; h\downarrow; yp\downarrow; [\neg yi]; q\downarrow \\
 & \quad \llbracket \neg h \wedge \neg q \wedge x0 \longrightarrow y0\uparrow; [yi]; xo\uparrow; [\neg x0]; y0\downarrow; [\neg yi]; xo\downarrow \\
 & \quad \llbracket \neg h \wedge \neg q \wedge x1 \longrightarrow y1\uparrow; [yi]; xo\uparrow; [\neg x1]; y1\downarrow; [\neg yi]; xo\downarrow \\
 & \quad \llbracket \neg h \wedge \neg q \wedge xt \longrightarrow yt\uparrow; [yi]; xo\uparrow; [\neg xt]; h\uparrow; yt\downarrow; [\neg yi]; xo\downarrow
 \end{aligned}$$

PRS

$$\begin{aligned}
 \neg xt \wedge yt & \rightarrow h\uparrow & h \wedge (x0 \vee x1 \vee xt) & \rightarrow q\uparrow \\
 q \wedge yi & \rightarrow h\downarrow & \neg h \wedge \neg yi & \rightarrow q\downarrow \\
 \\
 \neg q \wedge yi & \rightarrow xo\uparrow \\
 q \vee \neg yi & \rightarrow xo\downarrow
 \end{aligned}$$

$$\begin{aligned}
h \wedge q &\rightarrow yp\uparrow \\
\neg h \wedge q \wedge yi &\rightarrow yp\downarrow \\
\neg h \wedge \neg q \wedge x0 &\rightarrow y0\uparrow & \neg h \wedge \neg q \wedge xt &\rightarrow yt\uparrow \\
\neg q \wedge \neg x0 &\rightarrow y0\downarrow & h &\rightarrow yt\downarrow \\
\neg h \wedge \neg q \wedge x1 &\rightarrow y1\uparrow \\
\neg q \wedge \neg x1 &\rightarrow y1\downarrow
\end{aligned}$$

Radix 2 transistor accounting:

rule	transistor count	comments
h	8	
q	10	
x_o	4	
yp	5	OR'ed with a y rule below which also provides the staticizer
$y[0,1]$	18	
yt	8	
total	53	

Radix 4 transistor accounting:

rule	transistor count	comments
h	8	
q	12	
x_o	4	
yp	5	OR'ed with a y rule below which also provides the staticizer
$y[0,1,2,3]$	36	
yt	8	
total	73	

B.6 AEXT ASPR PM PFWD pipelined hq

This version has fewer state variables than AEXT PFWD hu, but the pull-up and pull-down chains are too long.

CHP

$PFWD \equiv$
 $h := true;$
 $*[[h \wedge \bar{X} \longrightarrow Y!(header); h\downarrow$
 $\quad \square \neg h \wedge \bar{X} \longrightarrow Y!(X?) \bullet [X = t \longrightarrow h\uparrow];$
 $\quad]$
 $]$

HSE

$*[h \wedge (x0 \vee x1 \vee xt) \longrightarrow yp\uparrow; [yi]; q\uparrow; yp\downarrow; [\neg yi]; h\downarrow; q\downarrow$
 $\quad \square \neg h \wedge x0 \longrightarrow y0\uparrow; xo\uparrow; [yi]; y0\downarrow; [\neg x0]; xo\downarrow; [\neg yi]$
 $\quad \square \neg h \wedge x1 \longrightarrow y1\uparrow; xo\uparrow; [yi]; y1\downarrow; [\neg x1]; xo\downarrow; [\neg yi]$
 $\quad \square \neg h \wedge xt \longrightarrow yt\uparrow; xo\uparrow; [yi]; h\uparrow; yt\downarrow; [\neg xt]; xo\downarrow; [\neg yi]]$

PRS

$h \wedge yi \wedge yp \rightarrow q\uparrow \quad yt \wedge xo \wedge yi \rightarrow h\uparrow$
 $\neg h \quad \quad \rightarrow q\downarrow \quad q \wedge \neg yi \quad \rightarrow h\downarrow$

 $\neg h \wedge (y0 \vee y1 \vee yt) \quad \rightarrow xo\uparrow$
 $\neg x0 \wedge \neg x1 \wedge \neg xt \wedge \neg y0 \wedge \neg y1 \wedge \neg yt \rightarrow xo\downarrow$

 $h \wedge \neg q \wedge \neg yi \wedge \neg xo \wedge (x0 \vee x1 \vee xt) \rightarrow yp\uparrow$
 $q \quad \quad \quad \quad \quad \quad \quad \rightarrow yp\downarrow$

 $\neg h \wedge \neg q \wedge \neg yi \wedge x0 \wedge \neg xo \rightarrow y0\uparrow \quad \neg h \wedge \neg q \wedge \neg yi \wedge x1 \wedge \neg xo \rightarrow y1\uparrow$
 $yi \wedge xo \quad \quad \quad \quad \quad \quad \quad \rightarrow y0\downarrow \quad yi \wedge xo \quad \quad \quad \quad \quad \quad \quad \rightarrow y1\downarrow$

 $\neg h \wedge \neg q \wedge \neg yi \wedge xt \wedge \neg xo \rightarrow yt\uparrow$
 $h \wedge xo \quad \quad \quad \quad \quad \quad \quad \rightarrow yt\downarrow$

Radix 2 transistor accounting:

rule	transistor count	comments
q	8	
h	9	
xo	14	
yp	8	OR'ed with a y rule below which also provides the staticizer
$y[0, 1]$	22	
yt	11	
total	72	

Radix 4 transistor accounting:

rule	transistor count	comments
q	8	
h	9	
xo	20	
yp	10	OR'ed with a y rule below which also provides the staticizer
$y[0, 1, 2, 3]$	44	
yt	11	
total	102	

B.7 AEXT ASPR PM PFWD hu

This version has more state variables than AEXT PFWD hq, but has reasonable pull-up and pull-down chains.

CHP

$$\begin{aligned}
 PFWD \equiv & \\
 & * [[h \wedge \bar{X} \longrightarrow Y!(\text{header}); h\downarrow \\
 & \quad \square \neg h \wedge \bar{X} \longrightarrow X?u \bullet Y!u, [u = t \longrightarrow h\uparrow] \\
 & \quad] \\
 &]
 \end{aligned}$$

HSE

$$\begin{aligned}
 & * [[h \wedge (x0 \vee x1 \vee xt) \longrightarrow yp\uparrow; [yi]; h\downarrow; yp\downarrow; [\neg yi] \\
 & \quad \square \neg h \wedge x0 \longrightarrow u0\uparrow; (xo\uparrow; [\neg x0]), (y0\uparrow; [yi]); u0\downarrow; (y0\downarrow; [\neg yi]), xo\downarrow \\
 & \quad \square \neg h \wedge x1 \longrightarrow u1\uparrow; (xo\uparrow; [\neg x1]), (y1\uparrow; [yi]); u1\downarrow; (y1\downarrow; [\neg yi]), xo\downarrow \\
 & \quad \square \neg h \wedge xt \longrightarrow ut\uparrow; (xo\uparrow; [\neg xt]), (yt\uparrow; h\uparrow; [yi]); ut\downarrow; (yt\downarrow; [\neg yi]), xo\downarrow \\
 & \quad] \\
 &]
 \end{aligned}$$

$$\begin{aligned}
& * [[h \wedge (x0 \vee x1 \vee xt) \longrightarrow yp\uparrow; [yi]; h\downarrow; yp\downarrow; [\neg yi] \\
& \quad \square \neg h \wedge x0 \longrightarrow u0\uparrow; [\neg x0 \wedge yi]; u0\downarrow; [\neg yi] \\
& \quad \square \neg h \wedge x1 \longrightarrow u1\uparrow; [\neg x1 \wedge yi]; u1\downarrow; [\neg yi] \\
& \quad \square \neg h \wedge xt \longrightarrow ut\uparrow; [\neg xt \wedge h \wedge yi]; ut\downarrow; [\neg yi] \\
& \quad] \\
&]
\end{aligned}$$

$$\begin{aligned}
& * [u0 \longrightarrow xo\uparrow, y0\uparrow; [\neg u0]; y0\downarrow, xo\downarrow \\
& \quad \square u1 \longrightarrow xo\uparrow, y1\uparrow; [\neg u1]; y1\downarrow, xo\downarrow \\
& \quad \square ut \longrightarrow xo\uparrow, (yt\uparrow; h\uparrow); [ut\downarrow]; yt\downarrow, xo\downarrow \\
& \quad]
\end{aligned}$$

PRS

$$\begin{aligned}
yt & \rightarrow h\uparrow \\
yp \wedge yi & \rightarrow h\downarrow
\end{aligned}$$

$$\begin{aligned}
h \wedge (x0 \vee x1 \vee xt) \wedge \neg yi \wedge \neg yt & \rightarrow yp\uparrow \\
\neg h \wedge yi \wedge \neg un & \rightarrow yp\downarrow
\end{aligned}$$

$$\begin{aligned}
u0 \vee u1 \vee ut & \rightarrow xo\uparrow \\
\neg u0 \wedge \neg u1 \wedge \neg ut & \rightarrow xo\downarrow
\end{aligned}$$

$$\begin{aligned}
\neg h \wedge x0 \wedge \neg yi & \rightarrow u0\uparrow & \neg h \wedge x1 \wedge \neg yi & \rightarrow u1\uparrow \\
\neg x0 \wedge yi & \rightarrow u0\downarrow & \neg x1 \wedge yi & \rightarrow u1\downarrow
\end{aligned}$$

$$\begin{aligned}
\neg h \wedge xt \wedge \neg yi & \rightarrow ut\uparrow \\
h \wedge \neg xt \wedge yi & \rightarrow ut\downarrow
\end{aligned}$$

$$\begin{aligned}
u0 & \rightarrow y0\uparrow & u1 & \rightarrow y1\uparrow \\
\neg h \wedge \neg u0 & \rightarrow y0\downarrow & \neg h \wedge \neg u1 & \rightarrow y1\downarrow
\end{aligned}$$

$$\begin{aligned}
ut & \rightarrow yt\uparrow \\
\neg ut & \rightarrow yt\downarrow
\end{aligned}$$

Radix 2 transistor accounting:

rule	transistor count	comments
h	7	
yp	9	OR'ed with a y rule below which also provides the staticizer
xo	6	
$u[0,1]$	18	
ut	10	
$y[0,1]$	14	
yt	6	
total	70	

Radix 4 transistor accounting:

rule	transistor count	comments
h	7	
yp	11	OR'ed with a y rule below which also provides the staticizer
xo	10	
$u[0,1,2,3]$	36	
ut	10	
$y[0,1,2,3]$	28	
yt	6	
total	108	

B.8 AEXT ASPR PM MERGE unpipelined

HSE

$$*[[\neg a0 \wedge \neg a1 \wedge (c00 \vee c01 \vee c0t) \longrightarrow a0\uparrow \\ \mid \neg a0 \wedge \neg a1 \wedge (c10 \vee c11 \vee c1t) \longrightarrow a1\uparrow]]$$

$$*[[a0 \wedge c00 \longrightarrow p0\uparrow; [pi]; c0o\uparrow; [\neg c00]; p0\downarrow; [\neg pi]; c0o\downarrow \\ \llbracket a0 \wedge c01 \longrightarrow p1\uparrow; [pi]; c0o\uparrow; [\neg c01]; p1\downarrow; [\neg pi]; c0o\downarrow \\ \llbracket a0 \wedge c0t \longrightarrow pt\uparrow; [pi]; c0o\uparrow; [\neg c0t]; a0\downarrow; pt\downarrow; [\neg pi]; c0o\downarrow \\ \llbracket a1 \wedge c10 \longrightarrow p0\uparrow; [pi]; c1o\uparrow; [\neg c10]; p0\downarrow; [\neg pi]; c1o\downarrow \\ \llbracket a1 \wedge c11 \longrightarrow p1\uparrow; [pi]; c1o\uparrow; [\neg c11]; p1\downarrow; [\neg pi]; c1o\downarrow \\ \llbracket a1 \wedge c1t \longrightarrow pt\uparrow; [pi]; c1o\uparrow; [\neg c1t]; a1\downarrow; pt\downarrow; [\neg pi]; c1o\downarrow \\ \rrbracket]]$$

PRS

$$\begin{array}{ll}
(c00 \vee c01 \vee c0t) \wedge \neg a0 \rightarrow a0i\uparrow & a0o \wedge \neg a1 \wedge \neg pi \rightarrow a0\uparrow \\
a0 & \rightarrow a0i\downarrow & \neg a0o \wedge pt \wedge \neg c0t \rightarrow a0\downarrow \\
\\
(c10 \vee c11 \vee c1t) \wedge \neg a1 \rightarrow a1i\uparrow & a1o \wedge \neg a0 \wedge \neg pi \rightarrow a1\uparrow \\
a1 & \rightarrow a1i\downarrow & \neg a1o \wedge pt \wedge \neg c1t \rightarrow a1\downarrow \\
\\
(a0 \wedge c00 \vee a1 \wedge c10) \rightarrow p0\uparrow & (a0 \wedge c0t \vee a1 \wedge c1t) \rightarrow pt\uparrow \\
(a0 \wedge \neg c00 \vee a1 \wedge \neg c10) \rightarrow p0\downarrow & \neg a0 \wedge \neg a1 \rightarrow pt\downarrow \\
\\
(a0 \wedge c00 \vee a1 \wedge c10) \rightarrow p0\uparrow & \\
(a0 \wedge \neg c00 \vee a1 \wedge \neg c10) \rightarrow p0\downarrow & \\
\\
a0 \wedge pi \rightarrow c0o\uparrow & a1 \wedge pi \rightarrow c1o\uparrow \\
\neg a0 \vee \neg pi \rightarrow c0o\downarrow & \neg a1 \vee \neg pi \rightarrow c1o\downarrow
\end{array}$$

Radix 2 transistor accounting:

rule	transistor count	comments
$a[0, 1]i$	16	can be combinational with radix 2
$a[0, 1]o$	12	2-way arbiter
$a[0, 1]$	20	
$p[0, 1]$	24	
pt	10	
$c[0, 1]o$	8	
total	90	

Radix 4 transistor accounting:

rule	transistor count	comments
$a[0, 1, 2, 3]i$	44	
$a[0, 1, 2, 3]o$	92	4-way unpipelined arbiter
$a[0, 1, 2, 3]$	40	
$p[0, 1, 2, 3]$	80	
pt	16	
$c[0, 1, 2, 3]o$	16	
total	288	

B.9 AEXT ASPR PM MERGE pipelined a_a

I don't like this version because the pullup chains for the state variables are too long and won't scale to higher radix encoding.

CHP

MERGE \equiv
 $*[[h \rightarrow [\overline{C0} \rightarrow a := 0 \mid \overline{C1} \rightarrow a := 1]; h\downarrow$
 $\quad \llbracket \neg h \wedge a = 0 \rightarrow P!(C0?)$
 $\quad \llbracket \neg h \wedge a = 1 \rightarrow P!(C1?)$
 $\quad \rrbracket$
 $]]$

HSE

$*[[\neg a0 \wedge \neg a1 \wedge (c00 \vee c01 \vee c0t) \rightarrow a0\uparrow$
 $\quad \mid \neg a0 \wedge \neg a1 \wedge (c10 \vee c11 \vee c1t) \rightarrow a1\uparrow]]$

$*[[a0 \wedge c00 \rightarrow p0\uparrow; c0o\uparrow; [pi \wedge \neg c00]; p0\downarrow; c0o\downarrow; [\neg pi]$
 $\quad \llbracket a0 \wedge c01 \rightarrow p1\uparrow; c0o\uparrow; [pi \wedge \neg c01]; p1\downarrow; c0o\downarrow; [\neg pi]$
 $\quad \llbracket a0 \wedge c0t \rightarrow pt\uparrow; c0o\uparrow; [pi \wedge \neg c0t]; a0\downarrow; pt\downarrow; c0o\downarrow; [\neg pi]$
 $\quad \llbracket a1 \wedge c10 \rightarrow p0\uparrow; c1o\uparrow; [pi \wedge \neg c10]; p0\downarrow; c1o\downarrow; [\neg pi]$
 $\quad \llbracket a1 \wedge c11 \rightarrow p1\uparrow; c1o\uparrow; [pi \wedge \neg c11]; p1\downarrow; c1o\downarrow; [\neg pi]$
 $\quad \llbracket a1 \wedge c1t \rightarrow pt\uparrow; c1o\uparrow; [pi \wedge \neg c1t]; a1\downarrow; pt\downarrow; c1o\downarrow; [\neg pi]$
 $\quad \rrbracket$

PRS

$\neg a0 \wedge (c00 \vee c01 \vee c0t) \wedge \neg c0o \rightarrow a0i\uparrow$	$a0o \wedge \neg a1 \wedge \neg c1o \wedge \neg pi \rightarrow a0\uparrow$
$a0 \rightarrow a0i\downarrow$	$\neg a0o \wedge pi \wedge pt \wedge \neg c0t \rightarrow a0\downarrow$
$\neg a1 \wedge (c10 \vee c11 \vee c1t) \wedge \neg c1o \rightarrow a1i\uparrow$	$a1o \wedge \neg a0 \wedge \neg c0o \wedge \neg pi \rightarrow a1\uparrow$
$a1 \rightarrow a1i\downarrow$	$\neg a1o \wedge pi \wedge pt \wedge \neg c1t \rightarrow a1\downarrow$
$\neg pi \wedge (a0 \wedge c00 \vee a1 \wedge c10) \rightarrow p0\uparrow$	$\neg pi \wedge (a0 \wedge c0t \vee a1 \wedge c1t) \rightarrow pt\uparrow$
$pi \wedge (\neg a0 \wedge \neg c00 \vee \neg a1 \wedge \neg c10) \rightarrow p0\downarrow$	$\neg a0 \wedge \neg a1 \rightarrow pt\downarrow$
$\neg pi \wedge (a0 \wedge c01 \vee a1 \wedge c11) \rightarrow p1\uparrow$	
$pi \wedge (\neg a0 \wedge \neg c01 \vee \neg a1 \wedge \neg c11) \rightarrow p1\downarrow$	
$a0 \wedge (p0 \vee p1 \vee pt) \rightarrow c0o\uparrow$	$a1 \wedge (p0 \vee p1 \vee pt) \rightarrow c1o\uparrow$
$\neg p0 \wedge \neg p1 \wedge \neg pt \rightarrow c0o\downarrow$	$\neg p0 \wedge \neg p1 \wedge \neg pt \rightarrow c1o\downarrow$

Radix 2 transistor accounting:

rule	transistor count	comments
$a[0, 1]i$	20	
$a[0, 1]o$	12	2-way arbiter
$a[0, 1]$	24	
$p[0, 1]$	28	
pt	11	
$c[0, 1]o$	22	
total	117	

Radix 4 transistor accounting:

rule	transistor count	comments
$a[0, 1, 2, 3]i$	48	
$a[0, 1, 2, 3]o$	92	4-way unpipelined arbiter
$a[0, 1, 2, 3]$	48	
$p[0, 1, 2, 3]$	88	
pt	17	
$c[0, 1, 2, 3]o$	60	
total	353	

B.10 AEXT ASPR PM MERGE pipelined ah

This one has acceptable pullup/pulldown chains, but I'm worried about making it CMOS implementable

HSE

$$\begin{aligned} & * [[\neg a0 \wedge (c00 \vee c01 \vee c0t) \longrightarrow a0\uparrow; h\downarrow; [\neg a0]; h\uparrow \\ & \quad | \neg a1 \wedge (c10 \vee c11 \vee c1t) \longrightarrow a1\uparrow; h\downarrow; [\neg a1]; h\uparrow]] \end{aligned}$$

$$\begin{aligned} & * [[a0 \wedge c00 \longrightarrow p0\uparrow; c0o\uparrow; [pi \wedge \neg c00]; p0\downarrow; c0o\downarrow; [\neg pi] \\ & \quad [a0 \wedge c01 \longrightarrow p1\uparrow; c0o\uparrow; [pi \wedge \neg c01]; p1\downarrow; c0o\downarrow; [\neg pi] \\ & \quad [a0 \wedge c0t \longrightarrow pt\uparrow; c0o\uparrow; [pi \wedge \neg c0t]; a0\downarrow; pt\downarrow; c0o\downarrow; [\neg pi] \\ & \quad [a1 \wedge c10 \longrightarrow p0\uparrow; c1o\uparrow; [pi \wedge \neg c10]; p0\downarrow; c1o\downarrow; [\neg pi] \\ & \quad [a1 \wedge c11 \longrightarrow p1\uparrow; c1o\uparrow; [pi \wedge \neg c11]; p1\downarrow; c1o\downarrow; [\neg pi] \\ & \quad [a1 \wedge c1t \longrightarrow pt\uparrow; c1o\uparrow; [pi \wedge \neg c1t]; a1\downarrow; pt\downarrow; c1o\downarrow; [\neg pi] \\ & \quad]] \end{aligned}$$

PRS

$$\begin{array}{ll}
(c00 \vee c01 \vee c0t) \wedge \neg a0 \rightarrow a0i\uparrow & h \wedge a0o \wedge \neg pi \rightarrow a0\uparrow \\
\neg h \wedge a0 \rightarrow a0i\downarrow & pt \wedge pi \wedge c0o \wedge \neg c0t \wedge \neg a0o \rightarrow a0\downarrow \\
\\
(c10 \vee c11 \vee c1t) \wedge \neg a1 \rightarrow a1i\uparrow & h \wedge a1o \wedge \neg pi \rightarrow a1\uparrow \\
\neg h \wedge a1 \rightarrow a1i\downarrow & pt \wedge pi \wedge c1o \wedge \neg c1t \wedge \neg a1o \rightarrow a1\downarrow \\
\\
\neg a0 \wedge \neg a1 \rightarrow h\uparrow & \\
a0 \vee a1 \wedge \neg c0o \wedge \neg c1o \rightarrow h\downarrow & \\
\\
\neg pi \wedge (a0 \wedge c00 \vee a1 \wedge c10) \wedge \neg h \rightarrow p0\uparrow & \neg pi \wedge (a0 \wedge c0t \vee a1 \wedge c1t) \wedge \neg h \rightarrow pt\uparrow \\
pi \wedge (a0 \wedge \neg c00 \vee a1 \wedge \neg c10) \rightarrow p0\downarrow & \neg a0 \wedge \neg a1 \rightarrow pt\downarrow \\
\\
\neg pi \wedge (a0 \wedge c01 \vee a1 \wedge c11) \wedge \neg h \rightarrow p1\uparrow & \\
pi \wedge (a0 \wedge \neg c01 \vee a1 \wedge \neg c11) \rightarrow p1\downarrow & \\
\\
a0 \wedge (p0 \vee p1 \vee pt) \rightarrow c0o\uparrow & a1 \wedge (p0 \vee p1 \vee pt) \rightarrow c1o\uparrow \\
\neg p0 \wedge \neg p1 \wedge \neg pt \rightarrow c0o\downarrow & \neg p0 \wedge \neg p1 \wedge \neg pt \rightarrow c1o\downarrow
\end{array}$$

Radix 2 transistor accounting:

rule	transistor count	comments
$a[0, 1]i$	20	
$a[0, 1]o$	12	2-way arbiter
$a[0, 1]$	24	
h	10	
$p[0, 1]$	30	
pt	12	
$c[0, 1]o$	22	
total	130	

Radix 4 transistor accounting:

rule	transistor count	comments
$a[0, 1, 2, 3]i$	48	
$a[0, 1, 2, 3]o$	92	4-way unpipelined arbiter
$a[0, 1, 2, 3]$	48	
h	14	
$p[0, 1, 2, 3]$	52	
pt	16	
$c[0, 1, 2, 3]o$	100	
total	370	

B.11 AEXT ASPR PFWD PREPEND/FWD/SIMPLE_MERGE

PREPEND and FWD are further decomposed into PREPEND, FWD, and SIMPLE_MERGE.

B.12 AEXT ASPR PFWD PREPEND

HSE

$$*[[si]; yp\uparrow; [yi]; yp\downarrow; [\neg yi]; so\uparrow; [\neg si]; so\downarrow]$$

B.13 AEXT PFWD FWD

HSE

$$\begin{aligned} &*[[x0 \vee x1 \vee xt]; so\uparrow; [si]; so\downarrow; [\neg si]; \\ &\quad [x0 \longrightarrow y0\uparrow; [yi]; y0\downarrow; [\neg yi] \\ &\quad [x1 \longrightarrow y1\uparrow; [yi]; y1\downarrow; [\neg yi] \\ &\quad [xt \longrightarrow yt\uparrow; [yi]; yt\downarrow; [\neg yi] \\ &\quad]] \end{aligned}$$

B.14 AEXT ASPR PFWD SIMPLE_MERGE

Assumes that its inputs are mutually exclusive.

HSE

$$\begin{aligned} &*[[c00 \longrightarrow p0\uparrow; [pi]; c0o\uparrow; [\neg c00]; p0\downarrow; [\neg pi]; c0o\downarrow \\ &\quad [c01 \longrightarrow p1\uparrow; [pi]; c0o\uparrow; [\neg c01]; p1\downarrow; [\neg pi]; c0o\downarrow \\ &\quad [c0t \longrightarrow pt\uparrow; [pi]; c0o\uparrow; [\neg c0t]; pt\downarrow; [\neg pi]; c0o\downarrow \\ &\quad [c10 \longrightarrow p0\uparrow; [pi]; c1o\uparrow; [\neg c00]; p0\downarrow; [\neg pi]; c1o\downarrow \\ &\quad [c11 \longrightarrow p1\uparrow; [pi]; c1o\uparrow; [\neg c01]; p1\downarrow; [\neg pi]; c1o\downarrow \\ &\quad [c1t \longrightarrow pt\uparrow; [pi]; c1o\uparrow; [\neg c0t]; pt\downarrow; [\neg pi]; c1o\downarrow \\ &\quad]] \end{aligned}$$

PRS

$$\begin{aligned} c00 \vee c10 &\longrightarrow p0\uparrow & c0t \vee c1t &\longrightarrow pt\uparrow \\ \neg c00 \wedge \neg c10 &\longrightarrow p0\downarrow & \neg c0t \wedge \neg c1t &\longrightarrow pt\downarrow \end{aligned}$$

$$\begin{aligned} c01 \vee c11 &\longrightarrow p1\uparrow \\ \neg c01 \wedge \neg c11 &\longrightarrow p1\downarrow \end{aligned}$$

$$\begin{array}{ll}
pi \wedge (c00 \vee c01 \vee c0t) \rightarrow c0o\uparrow & pi \wedge (c10 \vee c11 \vee c1t) \rightarrow c1o\uparrow \\
\neg pi & \rightarrow c0o\downarrow \quad \neg pi & \rightarrow c1o\downarrow
\end{array}$$

B.15 AEXT ASPR MERGE

MERGE sequences between outputting two serialized packet streams.

$$\begin{array}{l}
MERGE \equiv \\
*[[\neg l \wedge \neg r \rightarrow \\
\quad [\bar{L} \rightarrow l := true \\
\quad | \bar{R} \rightarrow r := true \\
\quad] \\
\quad [] l \vee r \rightarrow \\
\quad \quad [l \rightarrow O!(L?) \\
\quad \quad [] r \rightarrow O!(R?) \\
\quad]] \\
]
\end{array}$$

Appendix C

AER Receiver Design Space

This appendix explores the AER receiver (AERV) design space in reverse chronological order.

C.1 Receiver tree structure

The receiver tree structure is dictated by its interface with the synapse and neuron/synapse configuration memory. We could place a memory for each group of 1 synapse and 4 neurons,

Tree structure:


1 NODE
4 NODE
16 NODE
64 NODE
256 LEAF
1024 SYN 1024 DESERIAL
 1024 MEM4

Accounting:

component	transistors/component	components	transistors	comments
AERV NODE	152	85	12920	
AERV LEAF	152	256	38912	
DESERIAL	216	1024	221184	6 1-of-2 words
OR	4	1024	4096	
total			277112	
total/neuron			67.65	

This design is expensive. We are required to use the 1-of-2 instead of 1-of-4 deserializer. The deserializer costs are derived from the number of words, which are derived from the shape of the memory as detailed in Section D.28.

To reduce overhead and use a 1-of-4 deserializer, we bundle 2 synapses (8 neurons) into each port and consolidate the memories for the 16 neurons and their synapses into a single memory:



By consolidating the memories, we further save a port on each leaf node.

Tree structure:

- 1 NODE
- 4 NODE
- 16 NODE
- 64 NODE
- 256 LEAF(3)
- 512 SYN2 256 DESERIAL
- 256 MEM16

Accounting:

component	transistors/component	components	transistors	comments
AERV NODE	152	85	12920	
AERV LEAF(3)	114	256	29184	
DESERIAL	126	256	32256	4 1-of-4 words
OR	4	512	2048	
total			86408	
total/neuron			21.10	

We could continue this line of thought by consolidating synapses and memories until all configuration data is stored in a single, monolithic memory.

Tree structure:

- 1 NODE(3)
- 2 NODE DESERIAL
- 8 NODE MEM4096
- 32 NODE
- 128 LEAF
- 128 SYN8

Accounting:

component	transistors/component	components	transistors	comments
AERV NODE(3)	114	1	114	
AERV NODE	152	42	6384	
AERV LEAF	152	128	19456	
DESERIAL	352	1	352	
OR	4	512	2048	
total			28354	
total/neuron			6.92	

This design is the cheapest, but would be difficult to layout and wire.

C.2 AERV CD noTW cyclic control (CYC)

Builds off AERV CD noTW (Section C.6) by reusing the control lines to request and acknowledge the data as is done in AEXT CD noTW cyc (Section B.1.1).

Radix 2 accounting:

intermediate nodes			
component	transistors/component	components/node	transistors/node
NODE	60	1	60
total transistors/intermediate node			60
leaf nodes			
component	transistors/component	components/node	transistors/node
LEAF	60	1	60
total transistors/leaf node			60

$(60 \text{ transistors/intermediate node} * 511 \text{ intermediate nodes} + 60 \text{ transistors/leaf node} * 512 \text{ leaf nodes}) / 4096 \text{ neurons} = \mathbf{15.0 \text{ transistors/neuron}}$

Radix 4 accounting:

intermediate nodes			
component	transistors/component	components/node	transistors/node
NODE	152	1	152
total transistors/intermediate node			152
leaf nodes			
component	transistors/component	components/node	transistors/node
LEAF	152	1	152
total transistors/leaf node			152

(152 transistors/intermediate node * 256 intermediate nodes + 152 transistors/leaf node * 85 leaf nodes) / 4096 neurons = **12.7 transistors/neuron**

C.3 AERV CD noTW CYC NODE

HSE

```

*[[pφ → p0↑;
   [p0 → u0↑; uu↑; p0↓; [¬p0]; v↑; c0φ↑; [c0i]; p0↑
   [p1 → u1↑; uu↑; p0↓; [¬p1]; v↑; c1φ↑; [c1i]; p0↑
   ]
  [p0 ∧ c0φ → c00↑; [¬c0i]; p0↓; [¬p0]; c00↓; [c0i];; p0↑
  [p1 ∧ c0φ → c01↑; [¬c0i]; p0↓; [¬p1]; c01↓; [c0i];; p0↑
  [p0 ∧ c1φ → c10↑; [¬c1i]; p0↓; [¬p0]; c10↓; [c1i];; p0↑
  [p1 ∧ c1φ → c11↑; [¬c1i]; p0↓; [¬p1]; c11↓; [c1i];; p0↑
  [¬pφ → u0↓, u1↓; uu↓, (c0φ↓, c1φ↓; [¬c0i ∧ ¬c1i]; p0↓), v↓
  ]]

```

PRS

```

u0 ∨ u1 → uu↑
¬u0 ∧ ¬u1 → uu↓

pφ ∧ ¬uu ∨ c0i ∨ c1i → p0↑
(¬pφ ∨ uu) ∧ ¬c0i ∧ ¬c1i → p0↓

p0 ∧ ¬v → u0↑
¬pφ → u0↓

p1 ∧ ¬v → u1↑
¬pφ → u1↓

uu ∧ ¬p0 ∧ ¬p1 → v↑
¬uu → v↓

v ∧ u0 → c0φ↑
¬v ∨ ¬u0 → c0φ↓

v ∧ u1 → c1φ↑
¬v ∨ ¬u1 → c1φ↓

```

$$p0 \wedge c0\phi \rightarrow c00\uparrow$$

$$\neg p0 \vee \neg c0\phi \rightarrow c00\downarrow$$

$$p1 \wedge c0\phi \rightarrow c01\uparrow$$

$$\neg p1 \vee \neg c0\phi \rightarrow c01\downarrow$$

$$p0 \wedge c1\phi \rightarrow c10\uparrow$$

$$\neg p0 \vee \neg c1\phi \rightarrow c10\downarrow$$

$$p1 \wedge c1\phi \rightarrow c11\uparrow$$

$$\neg p1 \vee \neg c1\phi \rightarrow c11\downarrow$$

Radix 4 transistor approximate accounting:

rule	transistor count	comments
uu	8	
p_o	12	
$u[0, 1, 2, 3]$	28	
v	10	
$c[0, 1, 2, 3]\phi$	16	
$c[0, 1, 2, 3][0, 1, 2, 3]$	64	
total	138	

Alternative 0: HSE

```

*[[ $p_i$ ];  $po\uparrow$ ;
  [ $p0 \rightarrow u0\uparrow; uu\uparrow; po\downarrow; [\neg p0]; c0o\uparrow; cco\uparrow; u0\downarrow; uu\downarrow; [c0i]; cci\uparrow;$ 
    $po\uparrow; [\neg p_i]; c0o\downarrow; cco\downarrow, ([\neg c0i]; cci\downarrow); po\downarrow$ 
  [ $p1 \rightarrow u1\uparrow; uu\uparrow; po\downarrow; [\neg p1]; c1o\uparrow; cco\uparrow; u1\downarrow; uu\downarrow; [c1i]; cci\uparrow;$ 
    $po\uparrow; [\neg p_i]; c1o\downarrow; cco\downarrow, ([\neg c1i]; cci\downarrow); po\downarrow$ 
  ]];
]

*[[ $c0o \wedge p0 \rightarrow c00\uparrow; [\neg c0i]; cci\downarrow; po\downarrow; [\neg p0]; c00\downarrow; [c0i]; cci\uparrow; po\uparrow$ 
  [ $c0o \wedge p1 \rightarrow c01\uparrow; [\neg c0i]; cci\downarrow; po\downarrow; [\neg p1]; c01\downarrow; [c0i]; cci\uparrow; po\uparrow$ 
  [ $c1o \wedge p0 \rightarrow c10\uparrow; [\neg c1i]; cci\downarrow; po\downarrow; [\neg p0]; c10\downarrow; [c1i]; cci\uparrow; po\uparrow$ 
  [ $c1o \wedge p1 \rightarrow c11\uparrow; [\neg c1i]; cci\downarrow; po\downarrow; [\neg p1]; c11\downarrow; [c1i]; cci\uparrow; po\uparrow$ 
  ]]]

```

PRS

$$(pi \wedge \neg cco \vee cci) \wedge \neg uu \rightarrow po\uparrow$$

$$(\neg pi \vee cco) \wedge \neg cci \vee uu \rightarrow po\downarrow$$

$$p0 \wedge \neg cco \rightarrow u0\uparrow$$

$$cco \rightarrow u0\downarrow$$

$$p1 \wedge \neg cco \rightarrow u1\uparrow$$

$$cco \rightarrow u1\downarrow$$

$$u0 \vee u1 \rightarrow uu\uparrow$$

$$\neg u0 \wedge \neg u1 \rightarrow uu\downarrow$$

$$u0 \wedge \neg p0 \rightarrow c0o\uparrow$$

$$\neg pi \rightarrow c0o\downarrow$$

$$u1 \wedge \neg p1 \rightarrow c1o\uparrow$$

$$\neg pi \rightarrow c1o\downarrow$$

$$c0o \vee c1o \rightarrow cco\uparrow$$

$$\neg c0o \wedge \neg c1o \rightarrow cco\downarrow$$

$$c0i \vee c1i \rightarrow cci\uparrow$$

$$\neg c0i \wedge \neg c1i \rightarrow cci\downarrow$$

$$c0o \wedge p0 \rightarrow c00\uparrow$$

$$\neg c0o \vee \neg p0 \rightarrow c00\downarrow$$

$$c1o \wedge p0 \rightarrow c10\uparrow$$

$$\neg c1o \vee \neg p0 \rightarrow c10\downarrow$$

$$c0o \wedge p1 \rightarrow c01\uparrow$$

$$\neg c0o \vee \neg p1 \rightarrow c01\downarrow$$

$$c1o \wedge p1 \rightarrow c11\uparrow$$

$$\neg c1o \vee \neg p1 \rightarrow c11\downarrow$$

Radix 2 transistor approximate accounting:

rule	transistor count	comments
p_o	8	
$u[0, 1]$	14	
uu	4	could flatten in p_o
$c[0, 1]_o$	14	
cc_o	4	
cc_i	4	could flatten in p_o
$c[0, 1][0, 1]$	16	
total	64	60 if cc_i and uu flattened in p_o

Radix 4 transistor approximate accounting:

rule	transistor count	comments
p_o	8	
$u[0, 1, 2, 3]$	28	
uu	8	
$c[0, 1, 2, 3]_o$	28	
cc_o	8	
cc_i	8	
$c[0, 1, 2, 3][0, 1, 2, 3]$	64	
total	152	flattening cc_i and uu would make p_o pullup chain too long

Radix 3, 1-of-4 out transistor approximate accounting: This is used in some of the receiver designs.

rule	transistor count	comments
p_o	8	
$u[0, 1, 2]$	21	
uu	6	
$c[0, 1, 2]_o$	21	
cc_o	6	
cc_i	6	
$c[0, 1, 2][0, 1, 2, 3]$	48	
total	116	114 if cc_i or uu flattened in p_o (can't flatten both; pullup chain too long)

CMOS-implementable PRS

$$\neg_cci \rightarrow _cci\uparrow$$

$$_cci \rightarrow _cci\downarrow$$

$$\neg_cco \rightarrow _cco\uparrow$$

$$_cco \rightarrow _cco\downarrow$$

$$(pi \wedge _cco \vee _cci) \wedge _uu \rightarrow _po\downarrow$$

$$(\neg pi \vee \neg_cco) \wedge \neg_cci \vee \neg_uu \rightarrow _po\uparrow$$

$$\neg p0 \rightarrow _p0\uparrow \quad \neg p1 \rightarrow _p1\uparrow$$

$$p0 \rightarrow _p0\downarrow \quad p1 \rightarrow _p1\downarrow$$

$$\neg_cco \rightarrow _cco\uparrow$$

$$_cco \rightarrow _cco\downarrow$$

$$\neg_p0 \wedge \neg_cco \rightarrow u0\uparrow \quad \neg_p1 \wedge \neg_cco \rightarrow u1\uparrow$$

$$_cco \rightarrow u0\downarrow \quad _cco \rightarrow u1\downarrow$$

$$u0 \vee u1 \rightarrow _uu\downarrow$$

$$\neg u0 \wedge \neg u1 \rightarrow _uu\uparrow$$

$$u0 \wedge _p0 \rightarrow _c0o\downarrow$$

$$\neg pi \rightarrow _c0o\uparrow$$

$$u1 \wedge _p1 \rightarrow _c1o\downarrow$$

$$\neg pi \rightarrow _c1o\uparrow$$

$$\neg _c0o \vee \neg _c1o \rightarrow cco\uparrow$$

$$_c0o \wedge _c1o \rightarrow cco\downarrow$$

$$c0i \vee c1i \rightarrow _cci\downarrow$$

$$\neg c0i \wedge \neg c1i \rightarrow _cci\uparrow$$

$$\neg _c0o \wedge \neg _p0 \rightarrow c00\uparrow$$

$$_c0o \vee _p0 \rightarrow c00\downarrow$$

$$\neg _c1o \wedge \neg _p0 \rightarrow c10\uparrow$$

$$_c1o \vee _p0 \rightarrow c10\downarrow$$

$$\neg _c0o \wedge \neg _p1 \rightarrow c01\uparrow$$

$$_c0o \vee _p1 \rightarrow c01\downarrow$$

$$\neg _c1o \wedge \neg _p1 \rightarrow c11\uparrow$$

$$_c1o \vee _p1 \rightarrow c11\downarrow$$

$$\neg _c0o \rightarrow c0o\uparrow$$

$$_c0o \rightarrow c0o\downarrow$$

$$\neg _c1o \rightarrow c1o\uparrow$$

$$_c1o \rightarrow c1o\downarrow$$

$$\neg _po \rightarrow po\uparrow$$

$$_po \rightarrow po\downarrow$$

Radix 2 transistor accounting:

rule	transistor count	comments
$_cci$	2	
$_cco$	2	
po	8	
$_p[0, 1]$	4	
$_cco$	2	
$u[0, 1]$	14	
$_uu$	4	
$_c[0, 1]_o$	14	
cco	4	
$_cci$	4	
$c[0, 1][0, 1]$	16	
$c[0, 1]_o$	4	
po	2	
total	80	

Radix 4 transistor accounting:

rule	transistor count	comments
$_cc_i$	2	
$_cc_o$	2	
p_o	8	
$_p[0, 1, 2, 3]$	8	
$_cc_o$	2	
$u[0, 1, 2, 3]$	28	
$_uu$	8	
$_c[0, 1, 2, 3]_o$	28	
cc_o	8	
$_cc_i$	8	
$c[0, 1, 2, 3][0, 1, 2, 3]$	64	
$c[0, 1, 2, 3]_o$	8	
p_o	2	
total	176	

C.4 AERV CD noTW CYC LEAF

This leaf design does not transmit the c_o signals to the neuron. Rather data just shows up on the cx lines and the neurons acknowledge on the c_i lines.

```

*[[ $p_i$ ];  $po\uparrow$ ;
  [ $p0 \rightarrow u0\uparrow; uu\uparrow; po\downarrow; [\neg p0]; c0\uparrow; cc\uparrow; u0\downarrow; uu\downarrow;$ 
     $po\uparrow; [\neg p_i]; c0\downarrow; cc\downarrow; po\downarrow$ 
  [ $p1 \rightarrow u1\uparrow; uu\uparrow; po\downarrow; [\neg p1]; c1\uparrow; cc\uparrow; u1\downarrow; uu\downarrow;$ 
     $po\uparrow; [\neg p_i]; c1\downarrow; cc\downarrow; po\downarrow$ 
  ]];
]

*[[ $c0 \wedge p0 \rightarrow c00\uparrow; [c0_i]; cci\uparrow; po\downarrow; [\neg p0]; c00\downarrow; [\neg c0_i]; cci\downarrow; po\uparrow$ 
  [ $c0 \wedge p1 \rightarrow c01\uparrow; [c0_i]; cci\uparrow; po\downarrow; [\neg p1]; c01\downarrow; [\neg c0_i]; cci\downarrow; po\uparrow$ 
  [ $c1 \wedge p0 \rightarrow c10\uparrow; [c1_i]; cci\uparrow; po\downarrow; [\neg p0]; c10\downarrow; [\neg c1_i]; cci\downarrow; po\uparrow$ 
  [ $c1 \wedge p1 \rightarrow c11\uparrow; [c1_i]; cci\uparrow; po\downarrow; [\neg p1]; c11\downarrow; [\neg c1_i]; cci\downarrow; po\uparrow$ 
  ]]
```

PRS

$$(p_i \vee cco) \wedge \neg uu \wedge \neg cci \rightarrow po\uparrow$$

$$\neg p_i \wedge \neg cco \vee uu \vee cci \rightarrow po\downarrow$$

$$\begin{array}{ll}
\neg cc \wedge p0 \rightarrow u0\uparrow & \neg cc \wedge p1 \rightarrow u1\uparrow \\
cc \rightarrow u0\downarrow & cc \rightarrow u1\downarrow \\
\\
u0 \vee u1 \rightarrow uu\uparrow & \\
\neg u0 \wedge \neg u1 \rightarrow uu\downarrow & \\
\\
u0 \wedge \neg p0 \rightarrow c0\uparrow & u1 \wedge \neg p1 \rightarrow c1\uparrow \\
\neg pi \rightarrow c0\downarrow & \neg pi \rightarrow c1\downarrow \\
\\
c0 \vee c1 \rightarrow cc\uparrow & c0i \vee c1i \rightarrow cci\uparrow \\
\neg c0 \wedge \neg c1 \rightarrow cc\downarrow & \neg c0i \wedge \neg c1i \rightarrow cci\downarrow \\
\\
c0 \wedge p0 \rightarrow c00\uparrow & c1 \wedge p0 \rightarrow c10\uparrow \\
\neg c0 \vee \neg p0 \rightarrow c00\downarrow & \neg c1 \vee \neg p0 \rightarrow c10\downarrow \\
\\
c0 \wedge p1 \rightarrow c01\uparrow & c1 \wedge p1 \rightarrow c11\uparrow \\
\neg c0 \vee \neg p1 \rightarrow c01\downarrow & \neg c1 \vee \neg p1 \rightarrow c11\downarrow
\end{array}$$

Radix 2 transistor approximate accounting:

rule	transistor count	comments
p_o	8	
$u[0, 1]$	14	
uu	4	could flatten in p_o
$c[0, 1]_o$	14	
cc_o	4	
cc_i	4	could flatten in p_o
$c[0, 1][0, 1]$	16	
total	64	60 if cc_i and uu flattened in p_o

Radix 4 transistor approximate accounting:

rule	transistor count	comments
p_o	8	
$u[0, 1, 2, 3]$	28	
uu	8	
$c[0, 1, 2, 3]_o$	28	
cc_o	8	
cc_i	8	
$c[0, 1, 2, 3][0, 1, 2, 3]$	64	
total	152	flattening cc_i and uu would make p_o pullup chain too long

Radix 3, 1-of-4 out transistor approximate accounting: This is used in some of the receiver designs.

rule	transistor count	comments
p_o	8	
$u[0, 1, 2]$	21	
uu	6	
$c[0, 1, 2]_o$	21	
cc_o	6	
cc_i	6	
$c[0, 1, 2][0, 1, 2, 3]$	48	
total	116	114 if cc_i or uu flattened in p_o (can't flatten both; pullup chain too long)

CMOS-implementable PRS

$$\neg pi \rightarrow _pi\uparrow$$

$$pi \rightarrow _pi\downarrow$$

$$\neg_cc i \rightarrow _cc i\uparrow$$

$$_cc i \rightarrow _cc i\downarrow$$

$$(\neg_pi \vee \neg_cc) \wedge \neg uu \wedge \neg_cc i \rightarrow po\uparrow$$

$$_pi \wedge _cc \vee uu \vee _cc i \rightarrow po\downarrow$$

$$_cc \wedge p0 \rightarrow _u0\downarrow \quad _cc \wedge p1 \rightarrow _u1\downarrow$$

$$\neg_cc \rightarrow _u0\uparrow \quad \neg_cc \rightarrow _u1\uparrow$$

$$\neg_u0 \vee \neg_u1 \rightarrow uu\uparrow$$

$$_u0 \wedge _u1 \rightarrow uu\downarrow$$

$$\neg_u0 \wedge \neg p0 \rightarrow c0\uparrow \quad \neg_u1 \wedge \neg p1 \rightarrow c1\uparrow$$

$$_pi \rightarrow c0\downarrow \quad _pi \rightarrow c1\downarrow$$

$$c0 \vee c1 \rightarrow _cc\downarrow \quad c0i \vee c1i \rightarrow _cc i\downarrow$$

$$\neg c0 \wedge \neg c1 \rightarrow _cc\uparrow \quad \neg c0i \wedge \neg c1i \rightarrow _cc i\uparrow$$

$$c0 \wedge p0 \rightarrow _c00\downarrow \quad c1 \wedge p0 \rightarrow _c10\downarrow$$

$$\neg c0 \vee \neg p0 \rightarrow _c00\uparrow \quad \neg c1 \vee \neg p0 \rightarrow _c10\uparrow$$

$$c0 \wedge p1 \rightarrow _c01\downarrow \quad c1 \wedge p1 \rightarrow _c11\downarrow$$

$$\neg c0 \vee \neg p1 \rightarrow _c01\uparrow \quad \neg c1 \vee \neg p1 \rightarrow _c11\uparrow$$

$$\begin{array}{ll} \neg_c00 \rightarrow c00\uparrow & \neg_c10 \rightarrow c10\uparrow \\ _c00 \rightarrow c00\downarrow & _c10 \rightarrow c10\downarrow \end{array}$$

$$\begin{array}{ll} \neg_c01 \rightarrow c01\uparrow & \neg_c11 \rightarrow c11\uparrow \\ _c01 \rightarrow c01\downarrow & _c11 \rightarrow c11\downarrow \end{array}$$

Alternative 1: HSE

$$\begin{array}{l} *[[p\phi \longrightarrow po\uparrow; \\ \quad [p0 \longrightarrow u0\uparrow; uu\uparrow; po\downarrow; [\neg p0]; v\uparrow; c0s\uparrow; , po\uparrow \\ \quad \llbracket p1 \longrightarrow u1\uparrow; uu\uparrow; po\downarrow; [\neg p1]; v\uparrow; c1s\uparrow; , po\uparrow \\ \quad \rrbracket \\ \quad \llbracket p0 \wedge c0s \longrightarrow c00\uparrow; [c0i]; cci\uparrow; po\downarrow; [\neg p0]; c00\downarrow; [\neg c0i]; cci\downarrow; po\uparrow \\ \quad \llbracket p1 \wedge c0s \longrightarrow c01\uparrow; [c0i]; cci\uparrow; po\downarrow; [\neg p1]; c01\downarrow; [\neg c0i]; cci\downarrow; po\uparrow \\ \quad \llbracket p0 \wedge c1s \longrightarrow c10\uparrow; [c1i]; cci\uparrow; po\downarrow; [\neg p0]; c10\downarrow; [\neg c1i]; cci\downarrow; po\uparrow \\ \quad \llbracket p1 \wedge c1s \longrightarrow c11\uparrow; [c1i]; cci\uparrow; po\downarrow; [\neg p1]; c11\downarrow; [\neg c1i]; cci\downarrow; po\uparrow \\ \quad \llbracket \neg p\phi \longrightarrow u0\downarrow, u1\downarrow; uu\downarrow, c0s\downarrow, c1s\downarrow; v\downarrow, po\downarrow \\ \quad \rrbracket \\]] \end{array}$$

$$\begin{array}{l} u0 \vee u1 \rightarrow uu\uparrow \\ \neg u0 \wedge \neg u1 \rightarrow uu\downarrow \end{array}$$

$$\begin{array}{l} c0i \vee c1i \rightarrow cci\uparrow \\ \neg c0i \wedge \neg c1i \rightarrow cci\downarrow \end{array}$$

$$\begin{array}{l} p\phi \wedge \neg uu \vee v \wedge \neg cci \rightarrow po\uparrow \\ (\neg p\phi \vee uu) \wedge (\neg v \vee cci) \rightarrow po\downarrow \end{array}$$

$$\begin{array}{l} p0 \wedge \neg v \rightarrow u0\uparrow \\ \neg p\phi \rightarrow u0\downarrow \end{array}$$

$$\begin{array}{l} p1 \wedge \neg v \rightarrow u1\uparrow \\ \neg p\phi \rightarrow u1\downarrow \end{array}$$

$$\begin{array}{l} uu \wedge \neg p0 \wedge \neg p1 \rightarrow v\uparrow \\ \neg uu \wedge \neg c0s \wedge \neg c1s \rightarrow v\downarrow \end{array}$$

$$\begin{array}{l} v \wedge u0 \rightarrow c0s\uparrow \\ \neg v \vee \neg u0 \rightarrow c0s\downarrow \end{array}$$

$$\begin{array}{l} v \wedge u1 \rightarrow c1s\uparrow \\ \neg v \vee \neg u1 \rightarrow c1s\downarrow \end{array}$$

$$p0 \wedge c0s \rightarrow c00\uparrow$$

$$\neg p0 \vee \neg c0s \rightarrow c00\downarrow$$

$$p1 \wedge c0s \rightarrow c01\uparrow$$

$$\neg p1 \vee \neg c0s \rightarrow c01\downarrow$$

$$p0 \wedge c1s \rightarrow c10\uparrow$$

$$\neg p0 \vee \neg c1s \rightarrow c10\downarrow$$

$$p1 \wedge c1s \rightarrow c11\uparrow$$

$$\neg p1 \vee \neg c1s \rightarrow c11\downarrow$$

Radix 2 transistor approximate accounting:

rule	transistor count	comments
uu	4	
cc_i	4	could flatten in p_o
p_o	8	
$u[0, 1]$	14	
v	10	
$c[0, 1]s$	8	
$c[0, 1][0, 1]$	16	
total	69	67 if cc_i flattened in p_o

Radix 4 transistor approximate accounting:

rule	transistor count	comments
uu	8	
cc_i	8	could flatten in p_o
p_o	8	
$u[0, 1, 2, 3]$	28	
v	14	
$c[0, 1, 2, 3]s$	16	
$c[0, 1, 2, 3][0, 1, 2, 3]$	64	
total	146	144 if cc_i flattened in p_o

C.5 AERV CD noTW CYC LEAF (no data)

If we don't need to send any data to the neurons we can make very cheap leaf nodes.

*[[pi]; $po\uparrow$;
 $[p0 \rightarrow c0o\uparrow; [c0i]; po\downarrow; [\neg p0]; u\uparrow; po\uparrow; [\neg pi]; c0o\downarrow[\neg c0i]; u\downarrow; po\downarrow$
 $[p1 \rightarrow c1o\uparrow; [c1i]; po\downarrow; [\neg p1]; u\uparrow; po\uparrow; [\neg pi]; c1o\downarrow[\neg c1i]; u\downarrow; po\downarrow$
 $]$;
 $]$

PRS

$(pi \wedge \neg c0i \wedge \neg c1i) \vee u \rightarrow po\uparrow$
 $(\neg pi \vee c0i \vee c1i) \wedge \neg u \rightarrow po\downarrow$

$p0 \rightarrow c0o\uparrow$ $p1 \rightarrow c1o\uparrow$
 $\neg pi \rightarrow c0o\downarrow$ $\neg pi \rightarrow c1o\downarrow$

$c0i \wedge \neg p0 \vee c1i \wedge \neg p1 \rightarrow u\uparrow$
 $(\neg c0i \vee p0) \wedge (\neg c1i \vee p1) \rightarrow u\downarrow$

Radix 2 transistor approximate accounting:

rule	transistor count	comments
po	8	
$c[0, 1]_o$	12	
$u[0, 1]$	8	
total	28	

Radix 4 transistor approximate accounting:

rule	transistor count	comments
po	12	
$c[0, 1, 2, 3]_o$	24	
$u[0, 1, 2, 3]$	16	
total	52	

CMOS-implementable PRS

$\neg c0i \rightarrow \neg c0i\uparrow$ $\neg c1i \rightarrow \neg c1i\uparrow$
 $c0i \rightarrow \neg c0i\downarrow$ $c1i \rightarrow \neg c1i\downarrow$

$(pi \wedge \neg c0i \wedge \neg c1i) \vee u \rightarrow \neg po\downarrow$
 $(\neg pi \vee \neg c0i \vee \neg c1i) \wedge \neg u \rightarrow po\uparrow$

$p0 \rightarrow \neg c0o\downarrow$ $p1 \rightarrow \neg c1o\downarrow$
 $\neg pi \rightarrow \neg c0o\uparrow$ $\neg pi \rightarrow \neg c1o\uparrow$

$$\neg_c0i \wedge \neg p0 \vee \neg_c1i \wedge \neg p1 \rightarrow u\uparrow$$

$$(_c0i \vee p0) \wedge (_c1i \vee p1) \rightarrow u\downarrow$$

$$\neg_c0o \rightarrow c0o\uparrow \quad \neg_c1o \rightarrow c1o\uparrow$$

$$_c0o \rightarrow c0o\downarrow \quad _c1o \rightarrow c1o\downarrow$$

$$\neg_po \rightarrow po\uparrow$$

$$_po \rightarrow po\downarrow$$

C.6 AERV Control Data decomposed (CD) no tailword (noTW)

Separating control from data and removing the tail word reduced the number of transistors in the transmitter. We'll try to apply these same techniques to the receiver. Specifically, we'll want something that can interface with the control/data decomposed, no tail word, cyclic control, transmitter developed in Section B.1.1.

The accounting depends on whether we need the receiver to deliver payload or not. Without payload, we can simplify the leaf node circuitry. With payload, we'll need a more complicated interface with the neuron to be developed.

First we'll consider the case without payload:

Radix 2 accounting (2047 intermediate nodes, 2048 leaf nodes):

intermediate nodes			
component	transistors/component	components/node	transistors/node
SPLIT	30	1	30
CTRL	36	1	36
total transistors/intermediate node			66
leaf nodes			
component	transistors/component	components/node	transistors/node
LEAF	30	1	30
total transistors/leaf node			30

(66 transistors/intermediate node * 2047 intermediate nodes + 30 transistors/leaf node * 2048 leaf nodes) / 4096 neurons = **48.0 transistors/neuron**

Radix 4 accounting (341 intermediate nodes, 1024 leaf nodes):

intermediate nodes			
component	transistors/component	components/node	transistors/node
SPLIT	90	1	90
CTRL	68	1	68
total transistors/intermediate node			158
leaf nodes			
component	transistors/component	components/node	transistors/node
LEAF	58	1	58
total transistors/leaf node			58

(158 transistors/intermediate node * 341 intermediate nodes + 58 transistors/leaf node * 1024 leaf nodes) / 4096 neurons = **27.7 transistors/neuron**

Now we'll consider the case where we have to deliver payload. In this case, we cannot use the simplified leaf nodes because there is data to be sent to the neuron. In addition, we'll need to develop more circuitry per neuron to set bits. We haven't specified what or how data will be set, so this will be developed in the future.

Radix 2 accounting (4095 nodes):

$$(66 \text{ transistors/node} * 4095 \text{ nodes}) / 4096 \text{ neurons} = \mathbf{66.0 \text{ transistors/neuron}}$$

Radix 4 accounting (1365 nodes):

$$(158 \text{ transistors/node} * 1365 \text{ nodes}) / 4096 \text{ neurons} = \mathbf{52.6 \text{ transistors/neuron}}$$

C.7 AERV CD noTW SPLIT

$$*[\overline{C0} \wedge S = 0 \longrightarrow C0!(P?)$$

$$\quad \overline{C1} \wedge S = 1 \longrightarrow C1!(P?)$$

$$\quad \overline{C2} \wedge S = 2 \longrightarrow C2!(P?)$$

]]

$$\begin{aligned}
& * [[c0e \longrightarrow pe\uparrow; \\
& \quad [s0 \wedge p0 \longrightarrow c00\uparrow; [\neg c0e]; pe\downarrow; [\neg p0]; c00\downarrow \\
& \quad [s0 \wedge p1 \longrightarrow c01\uparrow; [\neg c0e]; pe\downarrow; [\neg p1]; c01\downarrow \\
& \quad] \\
& \quad [c1e \longrightarrow pe\uparrow; \\
& \quad \quad [s1 \wedge p0 \longrightarrow c10\uparrow; [\neg c1e]; pe\downarrow; [\neg p0]; c10\downarrow \\
& \quad \quad [s1 \wedge p1 \longrightarrow c11\uparrow; [\neg c1e]; pe\downarrow; [\neg p1]; c11\downarrow \\
& \quad \quad] \\
& \quad [c2e \longrightarrow pe\uparrow; \\
& \quad \quad [s2 \wedge p0 \longrightarrow c20\uparrow; [\neg c2e]; pe\downarrow; [\neg p0]; c20\downarrow \\
& \quad \quad [s2 \wedge p1 \longrightarrow c21\uparrow; [\neg c2e]; pe\downarrow; [\neg p1]; c21\downarrow \\
& \quad \quad] \\
&]]
\end{aligned}$$

$$\begin{aligned}
c0e \vee c1e \vee c2e & \longrightarrow pe\uparrow \\
\neg c0e \wedge \neg c1e \wedge \neg c2e & \longrightarrow pe\downarrow
\end{aligned}$$

$$\begin{aligned}
s0 \wedge p0 & \longrightarrow c00\uparrow & s1 \wedge p1 & \longrightarrow c11\uparrow \\
\neg s0 \vee \neg p0 & \longrightarrow c00\downarrow & \neg s1 \vee \neg p1 & \longrightarrow c11\downarrow
\end{aligned}$$

$$\begin{aligned}
s0 \wedge p1 & \longrightarrow c01\uparrow & s2 \wedge p0 & \longrightarrow c20\uparrow \\
\neg s0 \vee \neg p1 & \longrightarrow c01\downarrow & \neg s2 \vee \neg p0 & \longrightarrow c21\downarrow
\end{aligned}$$

$$\begin{aligned}
s1 \wedge p0 & \longrightarrow c10\uparrow & s2 \wedge p1 & \longrightarrow c21\uparrow \\
\neg s1 \vee \neg p0 & \longrightarrow c10\downarrow & \neg s2 \vee \neg p1 & \longrightarrow c21\downarrow
\end{aligned}$$

Radix 2 transistor accounting:

rule	transistor count	comments
pe	6	
$c[0, 1, 2][0, 1]$	24	
total	30	

Radix 4 transistor accounting:

rule	transistor count	comments
pe	10	
$c[0, 1, 2, 3, 4][0, 1, 2, 3]$	80	
total	90	

C.8 AERV CD noTW CTRL

$$\begin{aligned}
& *[[\bar{P} \longrightarrow X?u \bullet S := 2 \\
& \quad [u = 0 \longrightarrow S := 0; C0; P \\
& \quad \llbracket u = 1 \longrightarrow S := 1; C1; P \\
& \quad] \\
&]] \\
& *[[pi]; xe\uparrow; s2\uparrow; \\
& \quad [x0 \longrightarrow u0\uparrow \llbracket x1 \longrightarrow u1\uparrow]; xe\downarrow; [\neg x0 \wedge \neg x1]; s2\downarrow; \\
& \quad [u0 \longrightarrow s0\uparrow; c0o\uparrow; [c0i]; po\uparrow; [\neg pi]; u0\downarrow; s0\downarrow; c0o\downarrow; [\neg c0i] \\
& \quad \llbracket u1 \longrightarrow s1\uparrow; c1o\uparrow; [c1i]; po\uparrow; [\neg pi]; u1\downarrow; s1\downarrow; c1o\downarrow; [\neg c1i] \\
& \quad]; po\downarrow \\
&] \\
& pi \wedge \neg u0 \wedge \neg u1 \rightarrow xe\uparrow \\
& \neg pi \vee u0 \vee u1 \rightarrow xe\downarrow \\
& xe \vee x0 \vee x1 \rightarrow s2\uparrow \\
& \neg xe \wedge \neg x0 \wedge \neg x1 \rightarrow s2\downarrow \\
& x0 \rightarrow u0\uparrow \quad x1 \rightarrow u1\uparrow \\
& \neg pi \rightarrow u0\downarrow \quad \neg pi \rightarrow u1\downarrow \\
& u0 \wedge \neg s2 \rightarrow s0\uparrow \quad u1 \wedge \neg s2 \rightarrow s1\uparrow \\
& \neg u0 \vee s2 \rightarrow s0\downarrow \quad \neg u1 \vee s2 \rightarrow s1\downarrow \\
& s0 \rightarrow c0o\uparrow \quad s1 \rightarrow c1o\uparrow \\
& \neg s0 \rightarrow c0o\downarrow \quad \neg s1 \rightarrow c1o\downarrow \\
& c0i \vee c1i \rightarrow po\uparrow \\
& \neg c0i \wedge \neg c1i \rightarrow po\downarrow
\end{aligned}$$

Radix 2 transistor accounting:

rule	transistor count	comments
xe	6	
$s2$	6	
$u[0, 1]$	12	
$s[0, 1]$	8	
$c[0, 1]_o$	0	$s[0, 1] = c[0, 1]_o$
po	4	
total	36	

Radix 4 transistor accounting:

rule	transistor count	comments
xe	10	
$s4$	10	
$u[0, 1, 2, 3]$	24	
$s[0, 1, 2, 3]$	16	
$c[0, 1, 2, 3]_o$	0	$s[0, 1, 2, 3] = c[0, 1, 2, 3]_o$
po	8	
total	68	

C.9 AERV CD noTW LEAF

In the case that we don't need to deliver payload to the neuron, we can use this LEAF process to interface with the neuron.

$$\begin{aligned}
 & * [[\bar{P} \rightarrow X?u \\
 & \quad [u = 0 \rightarrow C0; P \\
 & \quad [u = 1 \rightarrow C1; P \\
 & \quad] \\
 &]]
 \end{aligned}$$

$$\begin{aligned}
 & * [[pi]; xe\uparrow; \\
 & \quad [x0 \rightarrow u0\uparrow [x1 \rightarrow u1\uparrow]; xe\downarrow; [\neg x0 \wedge \neg x1]; \\
 & \quad [u0 \rightarrow c0o\uparrow; [c0i]; po\uparrow; [\neg pi]; u0\downarrow; c0o\downarrow; [\neg c0i] \\
 & \quad [u1 \rightarrow c1o\uparrow; [c1i]; po\uparrow; [\neg pi]; u1\downarrow; c1o\downarrow; [\neg c1i] \\
 & \quad]; po\downarrow \\
 &]
 \end{aligned}$$

$$\begin{aligned}
 pi \wedge \neg u0 \wedge \neg u1 & \rightarrow xe\uparrow \\
 \neg pi \vee u0 \vee u1 & \rightarrow xe\downarrow
 \end{aligned}$$

$$\begin{aligned}
 x0 & \rightarrow u0\uparrow & x1 & \rightarrow u1\uparrow \\
 \neg pi & \rightarrow u0\downarrow & \neg pi & \rightarrow u1\downarrow
 \end{aligned}$$

$$\begin{aligned}
 u0 \wedge \neg x0 & \rightarrow c0o\uparrow & u1 \wedge \neg x1 & \rightarrow c1o\uparrow \\
 \neg u0 \vee x0 & \rightarrow c0o\downarrow & \neg u1 \vee x1 & \rightarrow c1o\downarrow
 \end{aligned}$$

$$\begin{aligned}
 c0i \vee c1i & \rightarrow po\uparrow \\
 \neg c0i \wedge \neg c1i & \rightarrow po\downarrow
 \end{aligned}$$

Radix 2 transistor accounting:

rule	transistor count	comments
xe	6	
$u[0, 1]$	12	
$c[0, 1]_o$	8	
p_o	4	
total	30	

Radix 4 transistor accounting:

rule	transistor count	comments
xe	10	
$u[0, 1, 2, 3]$	24	
$c[0, 1, 2, 3]_o$	16	
p_o	8	
total	58	

C.10 AERV ASPR BCAST pipelined

HSE

strict cpcp

$BCAST \equiv$

*[[$p0 \rightarrow c00\uparrow, c10\uparrow; po\uparrow; [c0i \wedge c1i \wedge \neg p0]; c00\downarrow, c10\downarrow; po\downarrow; [\neg c0i \wedge \neg c1i];$
 $\llbracket p1 \rightarrow c01\uparrow, c11\uparrow; po\uparrow; [c0i \wedge c1i \wedge \neg p1]; c01\downarrow, c11\downarrow; po\downarrow; [\neg c0i \wedge \neg c1i];$
 $\llbracket pt \rightarrow c0t\uparrow, c1t\uparrow; po\uparrow; [c0i \wedge c1i \wedge \neg pt]; c0t\downarrow, c1t\downarrow; po\downarrow; [\neg c0i \wedge \neg c1i];$
 \rrbracket]]

$BCAST \equiv$

*[[$p0 \rightarrow q0\uparrow; c00\uparrow, c10\uparrow; po\uparrow; [c0i \wedge c1i \wedge \neg p0]; q0\downarrow; c00\downarrow, c10\downarrow; po\downarrow; [\neg c0i \wedge \neg c1i]$
 $\llbracket p1 \rightarrow q1\uparrow; c01\uparrow, c11\uparrow; po\uparrow; [c0i \wedge c1i \wedge \neg p1]; q1\downarrow; c01\downarrow, c11\downarrow; po\downarrow; [\neg c0i \wedge \neg c1i]$
 $\llbracket pt \rightarrow qt\uparrow; c0t\uparrow, c1t\uparrow; po\uparrow; [c0i \wedge c1i \wedge \neg pt]; qt\downarrow; c0t\downarrow, c1t\downarrow; po\downarrow; [\neg c0i \wedge \neg c1i]$
 \rrbracket]]

PRS

$$\begin{array}{ll}
q0 \rightarrow c00\uparrow & q0 \rightarrow c10\uparrow \\
\neg q0 \rightarrow c00\downarrow & \neg q0 \rightarrow c10\downarrow \\
\\
q1 \rightarrow c01\uparrow & q1 \rightarrow c11\uparrow \\
\neg q1 \rightarrow c01\downarrow & \neg q1 \rightarrow c11\downarrow \\
\\
qt \rightarrow c0t\uparrow & qt \rightarrow c1t\uparrow \\
\neg qt \rightarrow c0t\downarrow & \neg qt \rightarrow c1t\downarrow \\
\\
\neg c0i \wedge \neg c1i \wedge p0 \rightarrow q0\uparrow & \neg c0i \wedge \neg c1i \wedge pt \rightarrow qt\uparrow \\
c0i \wedge c1i \wedge \neg p0 \rightarrow q0\downarrow & c0i \wedge c1i \wedge \neg pt \rightarrow qt\downarrow \\
\\
\neg c0i \wedge \neg c1i \wedge p1 \rightarrow q1\uparrow & \\
c0i \wedge c1i \wedge \neg p1 \rightarrow q1\downarrow &
\end{array}$$

HSE

output ordering cpcp parallelized

$$\begin{array}{l}
BCAST \equiv \\
*[[p0 \longrightarrow ([\neg c0i]; c00\uparrow), ([\neg c1i]; c10\uparrow); po\uparrow; [\neg p0]; ([c0i]; c00\downarrow), ([c1i]; c10\downarrow); po\downarrow; \\
\quad [p1 \longrightarrow ([\neg c0i]; c01\uparrow), ([\neg c1i]; c11\uparrow); po\uparrow; [\neg p1]; ([c0i]; c01\downarrow), ([c1i]; c11\downarrow); po\downarrow; \\
\quad [pt \longrightarrow ([\neg c0i]; c0t\uparrow), ([\neg c1i]; c1t\uparrow); po\uparrow; [\neg pt]; ([c0i]; c0t\downarrow), ([c1i]; c1t\downarrow); po\downarrow; \\
\quad]]
\end{array}$$

PRS

$$\begin{array}{ll}
\neg c0i \wedge p0 \rightarrow c00\uparrow & \neg c1i \wedge p0 \rightarrow c10\uparrow \\
c0i \wedge \neg p0 \rightarrow c00\downarrow & c1i \wedge \neg p0 \rightarrow c10\downarrow \\
\\
\neg c0i \wedge p1 \rightarrow c01\uparrow & \neg c1i \wedge p1 \rightarrow c11\uparrow \\
c0i \wedge \neg p1 \rightarrow c01\downarrow & c1i \wedge \neg p1 \rightarrow c11\downarrow \\
\\
\neg c0i \wedge pt \rightarrow c0t\uparrow & \neg c1i \wedge pt \rightarrow c1t\uparrow \\
c0i \wedge \neg pt \rightarrow c0t\downarrow & c1i \wedge \neg pt \rightarrow c1t\downarrow \\
\\
VN(C) \rightarrow po\uparrow & \\
\neg VN(C) \rightarrow po\downarrow &
\end{array}$$

po is the output of a VN detector

HSE

swap ordering of p and c in reset output ordering cppc parallelized

BCAST \equiv
 $*[[p0 \longrightarrow ([\neg c0i]; c00\uparrow), ([\neg c1i]; c10\uparrow); (po\uparrow; [\neg p0]; po\downarrow); ([c0i]; c00\downarrow), ([c1i]; c10\downarrow)$
 $\quad \llbracket p1 \longrightarrow ([\neg c0i]; c01\uparrow), ([\neg c1i]; c11\uparrow); (po\uparrow; [\neg p1]; po\downarrow); ([c0i]; c01\downarrow), ([c1i]; c11\downarrow)$
 $\quad \llbracket pt \longrightarrow ([\neg c0i]; c0t\uparrow), ([\neg c1i]; c1t\uparrow); (po\uparrow; [\neg pt]; po\downarrow); ([c0i]; c0t\downarrow), ([c1i]; c1t\downarrow)$
 $\quad \rrbracket]]$

PRS

$\neg c0i \wedge p0 \quad \rightarrow \quad c00\uparrow$	$\neg c1i \wedge p0 \quad \rightarrow \quad c10\uparrow$
$c0i \wedge \neg p0 \wedge \neg po \rightarrow c00\downarrow$	$c1i \wedge \neg p0 \wedge \neg po \rightarrow c10\downarrow$
$\neg c0i \wedge p1 \quad \rightarrow \quad c01\uparrow$	$\neg c1i \wedge p1 \quad \rightarrow \quad c11\uparrow$
$c0i \wedge \neg p1 \wedge \neg po \rightarrow c01\downarrow$	$c1i \wedge \neg p1 \wedge \neg po \rightarrow c11\downarrow$
$\neg c0i \wedge pt \quad \rightarrow \quad c0t\uparrow$	$\neg c1i \wedge pt \quad \rightarrow \quad c1t\uparrow$
$c0i \wedge \neg pt \wedge \neg po \rightarrow c0t\downarrow$	$c1i \wedge \neg pt \wedge \neg po \rightarrow c1t\downarrow$
$(c00 \wedge c10 \vee c01 \wedge c11 \vee c0t \wedge c1t) \wedge (p0 \vee p1 \vee pt) \rightarrow po\uparrow$	
$\neg p0 \wedge \neg p1 \wedge \neg pt \quad \rightarrow \quad po\downarrow$	

instability on down phases of c because p input can rise at anytime Could probably fix with state variables.

C.11 AERV ASPR BCAST unpipelined**HSE**

BCAST \equiv
 $*[[p0 \longrightarrow c00\uparrow, c10\uparrow; [c0i \wedge c1i]; po\uparrow; [\neg p0]; c00\downarrow, c10\downarrow; [\neg c0i \wedge \neg c1i]; po\downarrow$
 $\quad \llbracket p1 \longrightarrow c01\uparrow, c11\uparrow; [c0i \wedge c1i]; po\uparrow; [\neg p1]; c01\downarrow, c11\downarrow; [\neg c0i \wedge \neg c1i]; po\downarrow$
 $\quad \llbracket pt \longrightarrow c0t\uparrow, c1t\uparrow; [c0i \wedge c1i]; po\uparrow; [\neg pt]; c0t\downarrow, c1t\downarrow; [\neg c0i \wedge \neg c1i]; po\downarrow$
 $\quad \rrbracket]]$

PRS

$$\begin{array}{ll} p0 \rightarrow c00\uparrow & p0 \rightarrow c10\uparrow \\ \neg p0 \rightarrow c00\downarrow & \neg p0 \rightarrow c10\downarrow \end{array}$$

$$\begin{array}{ll} p1 \rightarrow c01\uparrow & p1 \rightarrow c11\uparrow \\ \neg p1 \rightarrow c01\downarrow & \neg p1 \rightarrow c11\downarrow \end{array}$$

$$\begin{array}{ll} pt \rightarrow c0t\uparrow & pt \rightarrow c1t\uparrow \\ \neg pt \rightarrow c0t\downarrow & \neg pt \rightarrow c1t\downarrow \end{array}$$

$$\begin{array}{l} c0i \wedge c1i \rightarrow po\uparrow \\ \neg c0i \wedge \neg c1i \rightarrow po\downarrow \end{array}$$

C.12 AERV PSAR

This makes the circuitry much simpler

C.13 AERV PSAR decomposed into ROUTE, READ_HEAD, FWD_BODY (RHB)

ROUTE sends a parent's signal to one of its children depending on which child requests. Assumes requests are mutually exclusive.

READ_HEAD reads the head word and signals FWD_BODY which way to forward the body packet

FWD_BODY forwards words to the children based on command from DEC

Radix 2 accounting (4095 nodes / 4096 neurons):

component	transistors/component	components/node	transistors/node
ROUTE	39	1	39
READ_HEAD	29	1	29
FWD_BODY	51	1	51
total transistors/node			119

119 transistors/node * 4095 nodes / 4096 neurons = **119.0 transistors/neuron**

Radix 4 transistor accounting (1365 nodes / 4096 neurons):

component	transistors/component	components/node	transistors/node
ROUTE	71	1	71
READ_HEAD	53	1	53
FWD_BODY	125	1	125
total transistors/node			249

249 transistors/node * 1365 nodes / 4096 neurons = **83.0 transistors/neuron**

However, we can still send a payload to the neurons with 1-of-2 data instead of 1-of-4 data at the leaf nodes. There are 1024 leaf nodes. This will simplify the leaf node ROUTE and FWD_BODY components because their children only need to see the 1-bit payload and tail. Each leaf node ROUTE can lose 1 bit (i.e. 2 data lines or 2 asymmetric c-elements or 14 transistors). Each leaf node FWD_BODY can lose 1 bit (i.e. 2 data lines or 2 AND-gates or 8 transistors) for each of 4 children. Therefore we can subtract

$$1024 * (14 + 8 * 4) = 47104 \text{ transistors.}$$

Leaving out the high bit from the leaf nodes yields

(249 transistors/node * 1365 nodes - 47104 transistors) / 4096 neurons = **71.5 transistors/neuron**

C.14 AERV PSAR RHB ROUTE unpipelined

Note that when communicating with READ_HEAD, ROUTE does not need to send the tail bit; READ_HEAD should never see a tail bit.

CHP

```
*[[ $\overline{C0!}; C0!(P?)$ 
  [ $\overline{H!}; H!(P?)$ ]
]
```

HSE

```
*[[ $c0e \vee he$ ];  $pe \uparrow$ 
  [ $p0 \wedge c0e \rightarrow c00 \uparrow; [\neg c0e]; pe \downarrow; [\neg p0]; c00 \downarrow$ 
  [ $p1 \wedge c0e \rightarrow c01 \uparrow; [\neg c0e]; pe \downarrow; [\neg p1]; c01 \downarrow$ 
  [ $pt \wedge c0e \rightarrow c0t \uparrow; [\neg c0e]; pe \downarrow; [\neg pt]; c0t \downarrow$ 
  [ $p0 \wedge he \rightarrow h0 \uparrow; [\neg he]; pe \downarrow; [\neg p0]; h0 \downarrow$ 
  [ $p1 \wedge he \rightarrow h1 \uparrow; [\neg he]; pe \downarrow; [\neg p1]; h1 \downarrow$ 
  ]
]
```

PRS

$$c0e \vee he \rightarrow pe\uparrow$$

$$\neg c0e \wedge \neg he \rightarrow pe\downarrow$$

$$p0 \wedge c0e \rightarrow c00\uparrow \quad p0 \wedge he \rightarrow h0\uparrow$$

$$\neg p0 \rightarrow c00\downarrow \quad \neg p0 \rightarrow h0\downarrow$$

$$p1 \wedge c0e \rightarrow c01\uparrow \quad p1 \wedge he \rightarrow h1\uparrow$$

$$\neg p1 \rightarrow c01\downarrow \quad \neg p1 \rightarrow h1\downarrow$$

$$pt \wedge c0e \rightarrow c0t\uparrow$$

$$\neg pt \rightarrow c0t\downarrow$$

Radix 2 accounting:

rule	transistor count	comments
pe	4	
$c0[0, 1, t]$	21	
$h[0, 1]$	14	
total	39	

Radix 4 transistor accounting:

rule	transistor count	comments
pe	8	
$c0[0, 1, 2, 3, t]$	35	
$h[0, 1, 2, 3]$	28	
total	71	

C.15 AERV PSAR RHB READ HEAD

HSE

$$*[[si]; xe\uparrow;$$

$$\quad [x0 \rightarrow u0\uparrow; xe\downarrow; [\neg x0]; s0\uparrow; [\neg si]; u0\downarrow; s0\downarrow$$

$$\quad [x1 \rightarrow u1\uparrow; xe\downarrow; [\neg x1]; s1\uparrow; [\neg si]; u1\downarrow; s1\downarrow$$

$$\quad]$$

$$]$$

PRS

$$si \wedge \neg u0 \wedge \neg u1 \rightarrow xe\uparrow$$

$$u0 \vee u1 \rightarrow xe\downarrow$$

$$\begin{array}{ll} x0 \rightarrow u0\uparrow & u0 \wedge \neg x0 \rightarrow s0\uparrow \\ \neg si \rightarrow u0\downarrow & \neg u0 \vee x0 \rightarrow s0\downarrow \end{array}$$

$$\begin{array}{ll} x1 \rightarrow u1\uparrow & u1 \wedge \neg x1 \rightarrow s1\uparrow \\ \neg si \rightarrow u1\downarrow & \neg u1 \vee x1 \rightarrow s1\downarrow \end{array}$$

Radix 2 accounting:

rule	transistor count	comments
xe	9	
$u[0, 1]$	12	
$s[0, 1]$	8	
total	29	

Radix 4 transistor accounting:

rule	transistor count	comments
xe	13	
$u[0, 1, 2, 3]$	24	
$s[0, 1, 2, 3]$	16	
total	53	

C.16 AERV PSAR RHB FWD_BODY unpipelined

HSE

```

*[[¬s0 ∧ ¬s1 → so↑;
  []s0 → [c0e]; pe↑;
    [ p0 → c00↑; [¬c0e]; pe↓; [¬p0]; c00↓
    []p1 → c01↑; [¬c0e]; pe↓; [¬p1]; c01↓
    []pt → c0t↑; [¬c0e]; pe↓; [¬pt]; so↓; [¬s0]; c0t↓
    ]
  []s1 → [c1e]; pe↑;
    [ p0 → c10↑; [¬c1e]; pe↓; [¬p0]; c10↓
    []p1 → c11↑; [¬c1e]; pe↓; [¬p1]; c11↓
    []pt → c1t↑; [¬c1e]; pe↓; [¬pt]; so↓; [¬s1]; c1t↓
    ]
  ]
]
```

PRS

$$\begin{array}{lll}
s0 \wedge c0e \vee s1 \wedge c1e & \rightarrow pe\uparrow & \neg c0t \wedge \neg c1t \quad \rightarrow so\uparrow \\
s0 \wedge \neg c0e \vee s1 \wedge \neg c1e & \rightarrow pe\downarrow & \neg pt \wedge (c0t \vee c1t) \rightarrow so\downarrow
\end{array}$$

$$\begin{array}{ll}
p0 \wedge s0 & \rightarrow c00\uparrow \\
\neg p0 \vee \neg s0 & \rightarrow c00\downarrow
\end{array}
\quad
\begin{array}{ll}
p0 \wedge s1 & \rightarrow c10\uparrow \\
\neg p0 \vee \neg s1 & \rightarrow c10\downarrow
\end{array}$$

$$\begin{array}{ll}
p1 \wedge s0 & \rightarrow c01\uparrow \\
\neg p1 \vee \neg s0 & \rightarrow c01\downarrow
\end{array}
\quad
\begin{array}{ll}
p1 \wedge s1 & \rightarrow c11\uparrow \\
\neg p1 \vee \neg s1 & \rightarrow c11\downarrow
\end{array}$$

$$\begin{array}{ll}
pt \wedge s0 & \rightarrow c0t\uparrow \\
\neg s0 & \rightarrow c0t\downarrow
\end{array}
\quad
\begin{array}{ll}
pt \wedge s1 & \rightarrow c1t\uparrow \\
\neg s1 & \rightarrow c1t\downarrow
\end{array}$$

Radix 2 accounting:

rule	transistor count	comments
pe	12	
s_o	9	
$c[0, 1][0, 1]$	16	
$c[0, 1]t$	14	
total	51	

Radix 4 transistor accounting:

rule	transistor count	comments
pe	20	
s_o	13	
$c[0, 1, 2, 3][0, 1, 2, 3]$	64	
$c[0, 1, 2, 3]t$	28	
total	125	

C.17 AERV PSAR RHB FWD_BODY pipelined

HSE

```

*[so↑; [s0 ∨ s1]; pe↑;
  [ p0 ∧ s0 ∧ c0e → c00↑; pe↓; [¬p0 ∧ ¬c0e]; c00↓
  [ p1 ∧ s0 ∧ c0e → c01↑; pe↓; [¬p1 ∧ ¬c0e]; c01↓
  [ pt ∧ s0 ∧ c0e → c0t↑; pe↓; [¬pt ∧ ¬c0e]; so↓; [¬s0]; c0t↓
  [ p0 ∧ s1 ∧ c1e → c10↑; pe↓; [¬p0 ∧ ¬c1e]; c10↓
  [ p1 ∧ s1 ∧ c1e → c11↑; pe↓; [¬p1 ∧ ¬c1e]; c11↓
  [ pt ∧ s1 ∧ c1e → c1t↑; pe↓; [¬pt ∧ ¬c1e]; so↓; [¬s1]; c1t↓
  ]
]

```

PRS

$$\begin{array}{ll}
s0 \vee s1 \wedge \neg q \rightarrow pe\uparrow & c00 \vee c01 \vee c0t \vee c10 \vee c11 \vee c1t \rightarrow q\uparrow \\
q \rightarrow pe\downarrow & \neg c00 \wedge \neg c01 \wedge \neg c0t \wedge \neg c10 \wedge \neg c11 \wedge \neg c1t \rightarrow q\downarrow \\
\neg q \rightarrow so\uparrow & \\
\neg pt \wedge (\neg c0e \wedge c0t \vee \neg c1e \wedge c1t) \rightarrow so\downarrow & \\
p0 \wedge s0 \wedge c0e \rightarrow c00\uparrow & p0 \wedge s1 \wedge c1e \rightarrow c10\uparrow \\
\neg p0 \wedge \neg c0e \rightarrow c00\downarrow & \neg p0 \wedge \neg c0e \rightarrow c10\downarrow \\
p1 \wedge s0 \wedge c0e \rightarrow c01\uparrow & p1 \wedge s1 \wedge c1e \rightarrow c11\uparrow \\
\neg p1 \wedge \neg c0e \rightarrow c01\downarrow & \neg p1 \wedge \neg c0e \rightarrow c11\downarrow \\
pt \wedge s0 \wedge c0e \rightarrow c0t\uparrow & pt \wedge s1 \wedge c1e \rightarrow c1t\uparrow \\
\neg s0 \rightarrow c0t\downarrow & \neg s1 \rightarrow c1t\downarrow
\end{array}$$

Radix 2 accounting:

rule	transistor count	comments
total	90	

Radix 4 transistor accounting:

rule	transistor count	comments
total	262	

C.18 AERV PSAR decomposed into ROUTE PULL_CTRL PULL (RCP)

ROUTE sends a parent's signal to one of its children depending on which child requests. Assumes requests are mutually exclusive. This is same ROUTE as above.

PULL_CTRL reads the head word and indicates which PULL should request data from ROUTE.
 PULL requests data from ROUTE and passes the data to the child.

Radix 2 accounting (4095 nodes / 4096 neurons):

component	transistors/component	components/node	transistors/node
ROUTE	62	1	62
PULL_CTRL	28	1	28
PULL	14	2	28
total transistors/node			118

118 transistors/node * 4095 nodes / 4096 neurons = **118.0 transistors/neuron**

Radix 4 transistor accounting (1365 nodes / 4096 neurons):

component	transistors/component	components/node	transistors/node
ROUTE	178	1	178
PULL_CTRL	58	1	58
PULL	14	4	56
total transistors/node			292

292 transistors/node * 1365 nodes / 4096 neurons = **97.3 transistors/neuron**

However, we can still send a payload to the neurons with 1-of-2 data instead of 1-of-4 data at the leaf nodes. There are 1024 leaf nodes. This will simplify the leaf node ROUTE component because its children only need to see the 1-bit payload and tail. Each leaf node ROUTE can lose 1 bit (i.e. 2 data lines or 2 asymmetric c-elements or 14 transistors) for each of 4 children. Therefore we can subtract $1024 * 14 * 4 = 57344$ transistors

Leaving out the high bit from the leaf nodes yields

$(292 \text{ transistors/node} * 1365 \text{ nodes} - 57344 \text{ transistors}) / 4096 \text{ neurons} = \mathbf{83.3 \text{ transistors/neuron}}$

C.19 AERV PSAR RCP ROUTE

This decomposition largely reuses the unpipelined ROUTE in Section C.14 above. For this decomposition, ROUTE connects to [radix] instances of PULL and 1 instance of PULL_CTRL.

Radix 2 accounting:

rule	transistor count	comments
pe	6	
$c[0, 1][0, 1, t]$	42	
$h[0, 1]$	14	
total	62	

Radix 4 transistor accounting:

rule	transistor count	comments
pe	10	
$c[0, 1, 2, 3][0, 1, 2, 3, t]$	140	
$h[0, 1, 2, 3]$	28	
total	178	

C.20 AERV PSAR PULL_CTRL

CHP

HSE

```

*[[s0i ∧ s1i]; xe↑;
  [ x0 → u0↑; xe↓; [¬x0]; s0o↑; [¬s0i]; u0↓; s0o↓
  [x1 → u1↑; xe↓; [¬x1]; s1o↑; [¬s0i]; u1↓; s0o↓
  ]]

```

PRS

$$\begin{array}{ll}
s0i \wedge s1i \wedge q & \rightarrow xe\uparrow & \neg u0 \wedge \neg u1 & \rightarrow q\uparrow \\
\neg s0i \vee \neg s1i \vee \neg q & \rightarrow xe\downarrow & u0 \vee u1 & \rightarrow q\downarrow
\end{array}$$

$$\begin{array}{ll}
x0 & \rightarrow u0\uparrow & x1 & \rightarrow u1\uparrow \\
\neg s0i & \rightarrow u0\downarrow & \neg s1i & \rightarrow u1\downarrow
\end{array}$$

$$\begin{array}{ll}
u0 \wedge \neg x0 & \rightarrow s0o\uparrow & u1 \wedge \neg x1 & \rightarrow s1o\uparrow \\
\neg u0 \vee x0 & \rightarrow s0o\downarrow & \neg u1 \vee x1 & \rightarrow s1o\downarrow
\end{array}$$

Radix 2 accounting:

rule	transistor count	comments
xe	8	no q
$u[0, 1]$	12	
$s[0, 1]_o$	8	
total	28	

Radix 4 transistor accounting:

rule	transistor count	comments
xe	10	
q	8	
$u[0, 1, 2, 3]$	24	
$s[0, 1, 2, 3]_o$	16	
total	58	

C.21 AERV PSAR RCP PULL unpipelined

HSE

```

* $\lceil \neg si \rightarrow so \uparrow$ ;
   $\llbracket si \rightarrow ye \uparrow$ ;  $xe \uparrow$ 
     $\llbracket x0 \rightarrow y0 \uparrow$ ;  $\lceil \neg ye \rceil$ ;  $xe \downarrow$ ;  $\lceil \neg x0 \rceil$ ;  $y0 \downarrow$ 
     $\llbracket x1 \rightarrow y1 \uparrow$ ;  $\lceil \neg ye \rceil$ ;  $xe \downarrow$ ;  $\lceil \neg x1 \rceil$ ;  $y1 \downarrow$ 
     $\llbracket xt \rightarrow yt \uparrow$ ;  $\lceil \neg ye \rceil$ ;  $xe \downarrow$ ;  $\lceil \neg xt \rceil$ ;  $so \downarrow$ ;  $\lceil \neg si \rceil$ ;  $yt \downarrow$ 
   $\rrbracket$ 
 $\rrbracket$ 

```

PRS

```

 $\neg yt \wedge \neg si \rightarrow so \uparrow$        $si \wedge ye \rightarrow xe \uparrow$ 
 $yt \wedge \neg xt \rightarrow so \downarrow$      $\neg si \vee \neg ye \rightarrow xe \downarrow$ 

 $x0 \rightarrow y0 \uparrow$        $xt \rightarrow yt \uparrow$ 
 $\neg x0 \rightarrow y0 \downarrow$    $\neg si \rightarrow yt \downarrow$ 

 $x1 \rightarrow y1 \uparrow$ 
 $\neg x1 \rightarrow y1 \downarrow$ 

```

Radix 2 transistor accounting:

rule	transistor count	comments
s_o	4	
xe	4	
$y[0, 1]$	0	wires
yt	6	
total	14	

Radix 4 transistor accounting:

rule	transistor count	comments
s_o	4	
xe	4	
$y[0, 1, 2, 3]$	0	wires
yt	6	
total	14	

Appendix D

AER Interface Design Space

The router interfaces include conversions between the serial protocol as well as deserializers and serializers.

D.1 OUT e1ofN

Interfaces AEXT/AERV serial format to e1ofN channel.

$$*[[xi \wedge ye]; xo\uparrow; [\neg xi \wedge \neg ye]; xo\downarrow]$$

$$*[[x0 \longrightarrow y0\uparrow; [\neg x0]; y0\downarrow \\ \quad \llbracket x1 \longrightarrow y1\uparrow; [\neg x1]; y1\downarrow \\ \quad \rrbracket]]$$

PRS

$$xi \wedge ye \quad \rightarrow \quad xo\uparrow \\ \neg xi \vee \neg ye \rightarrow \quad xo\downarrow$$

$$x0 \quad \rightarrow \quad y0\uparrow \quad \quad x1 \quad \rightarrow \quad y1\uparrow \\ \neg x0 \rightarrow \quad y0\downarrow \quad \quad \neg x1 \rightarrow \quad y1\downarrow$$

CMOS-implementable PRS version 0

$$xi \wedge ye \quad \rightarrow \quad _xo\downarrow \quad \quad \neg _xo \rightarrow \quad xo\downarrow \\ \neg xi \vee \neg ye \rightarrow \quad _xo\uparrow \quad \quad _xo \quad \rightarrow \quad xo\uparrow$$

$$x0 \quad \rightarrow \quad y0\uparrow \quad \quad x1 \quad \rightarrow \quad y1\uparrow \\ \neg x0 \rightarrow \quad y0\downarrow \quad \quad \neg x1 \rightarrow \quad y1\downarrow$$

CMOS-implementable PRS version 1

$$\begin{array}{ll}
xi \wedge ye \rightarrow _xo\downarrow & \neg_xo \rightarrow xo\downarrow \\
\neg xi \vee \neg ye \rightarrow _xo\uparrow & _xo \rightarrow xo\uparrow \\
\neg_x0 \rightarrow y0\uparrow & \neg_x1 \rightarrow y1\uparrow \\
_x0 \rightarrow y0\downarrow & _x1 \rightarrow y1\downarrow
\end{array}$$

Radix 2 transistor accounting:

rule	transistor count	comments
xo	4	
$y[0,1]$	0	wires
total	4	

Radix 4 transistor accounting:

rule	transistor count	comments
xo	4	
$y[0,1]$	0	wires
total	4	

D.2 OUT a1ofN

Interfaces AEXT/AERV serial format to a1ofN channel.

$$\begin{array}{l}
*[[xi]; xo\uparrow; [\neg xi]; xo\downarrow] \\
\\
*[[x0 \longrightarrow y0\uparrow; [ya]; xo\downarrow; [\neg x0]; y0\downarrow; [\neg ya]; xo\uparrow \\
\quad [x1 \longrightarrow y1\uparrow; [ya]; xo\downarrow; [\neg x1]; y1\downarrow; [\neg ya]; xo\uparrow \\
\quad]]
\end{array}$$

PRS

$$\begin{array}{ll}
xi \wedge \neg ya \rightarrow xo\uparrow & \\
\neg xi \vee ya \rightarrow xo\downarrow & \\
\\
x0 \rightarrow y0\uparrow & x1 \rightarrow y1\uparrow \\
\neg x0 \rightarrow y0\downarrow & \neg x1 \rightarrow y1\downarrow
\end{array}$$

CMOS-implementable PRS version 0

$$\begin{array}{ll}
xi \wedge _ya \rightarrow _xo\downarrow & \neg_xo \rightarrow xo\uparrow \\
\neg xi \vee \neg_ya \rightarrow _xo\uparrow & _xo \rightarrow xo\downarrow
\end{array}$$

$$\begin{array}{ll} _x0 \rightarrow _y0\downarrow & _x1 \rightarrow _y1\downarrow \\ \neg_x0 \rightarrow _y0\uparrow & \neg_x1 \rightarrow _y1\uparrow \end{array}$$

Radix 2 transistor accounting:

rule	transistor count	comments
$_x_o$	4	
x_o	2	
$_y[0, 1]$	0	wires
total	6	

Radix 4 transistor accounting:

rule	transistor count	comments
$_x_o$	4	
x_o	2	
$_y[0, 1]$	0	wires
total	6	

D.3 Deserializer

The deserializer converts 1-of-N serial data into M-1-of-N parallel data.

```

*[X?ym;
  [m < M - 1 → m := m + 1
  [] m = M - 1 → Y!y; m := 0
]
]

```

We place a deserializer at output of the transmitter to interface with the datapath circuitry. It is the first in a series of processes that communicate with the outside environment:

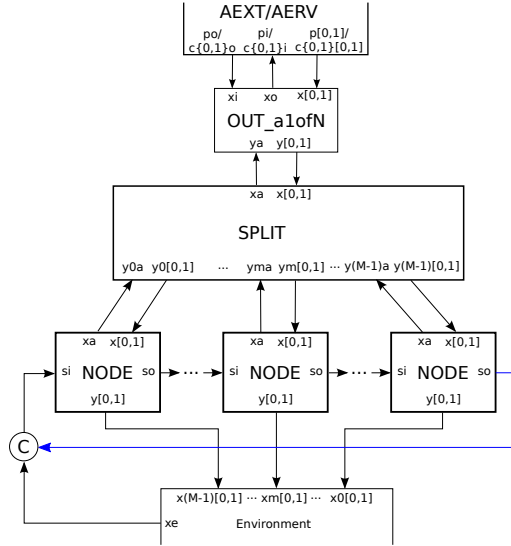
transmitter → deserializer → 1-of-4-to-1-of-2 converter (if needed) → [Datapath] → serializer → receiver

We also place a deserializer at the output of the receiver to interface with the neuron configuration memory.

D.4 Ring Deserializer

This design uses a ring of nodes receiving data from a central splitter to sequence words into their respective place in the parallel output. We decompose this process into SPLIT and NODE. This

design has a slightly cheaper 1-of-4 implementation than the chain deserializer of Section D.9. However, the data signals in SPLIT and the environment enable signal have fanouts that grow with the number of words. The below figure shows the decomposition for packets containing M 1-of-2 words.



An OUT a1ofN process (if necessary and described above) first converts the AEXT/AERV serial communication protocol to the standard a1ofN protocol.

1-of-2 approximate scaling:

component	transistors/component	components/deserializer	transistors/deserializer
OUT a1ofN	4	1	4
SPLIT	$3M - 2$	1	$3M - 2$
NODE	32	M	$32M$
C	8	1	8
approx. transistors/deserializer			$35M + 10$

1-of-4 approximate scaling:

component	transistors/component	components/deserializer	transistors/deserializer
OUT a1ofN	4	1	4
SPLIT	$3M - 2$	1	$3M - 2$
NODE	54	M	$54M$
C	8	1	8
approx. transistors/deserializer			$57M + 10$

For the transmitter to handle 4096 neurons encoded as 1-of-2 or 1-of-4 words, we would need 12 and 6 NODEs, respectively.

1-of-2 accounting:

component	transistors/component	components/deserializer	transistors/deserializer
OUT a1ofN	4	1	4
SPLIT	30	1	30
NODE	32	12	384
C	8	1	8
total transistors/deserializer			426

1-of-4 accounting:

component	transistors/component	components/deserializer	transistors/deserializer
OUT a1ofN	4	1	4
SPLIT	16	1	16
NODE	54	6	324
C	8	1	8
total transistors/deserializer			352

D.5 SPLIT

SPLIT takes incoming words and routes them to their respective locations in the parallel output.

For M words per packet,

```

*[[x0 → y00↑, ..., y(M-1)0↑; [y0a ∨ .. ∨ y(M-1)a]; xa↑;
  [¬x0]; y00↓, ..., y(M-1)1↓; [¬y0a ∧ .. ∧ ¬y(M-1)a]; xa↓
  [x1 → y01↑, ..., y(M-1)1↑; [y0a ∨ .. ∨ y(M-1)a]; xa↑;
  [¬x0]; y01↓, ..., y(M-1)1↓; [¬y0a ∧ .. ∧ ¬y(M-1)a]; xa↓
]]

```

For a 2-word packet,

```

*[[x0 → y00↑, y10↑; [y0a ∨ y1a]; xa↑;
  [¬x0]; y00↓, y01↓; [¬y0a ∧ ¬y1a]; xa↓
  [x1 → y01↑, y11↑; [y0a ∨ y1a]; xa↑;
  [¬x0]; y01↓, y11↓; [¬y0a ∧ ¬y1a]; xa↓
]]

```

PRS

```

x0 → y00↑      x1 → y01↑
¬x0 → y00↓     ¬x1 → y01↓

x0 → y10↑      x1 → y11↑
¬x0 → y10↓     ¬x1 → y11↓

```

$$y0a \vee y1a \rightarrow xa\uparrow$$

$$\neg y0a \wedge \neg y1a \rightarrow xa\downarrow$$

1-of-2 transistor approximate accounting:

rule	transistor count	comments
$y[0..M-1][0, 1]$	0	wires
xa	$8(M-1)/3$	4-ary OR-tree approx.
approx. total	$3M-2$	

1-of-4 transistor approximate accounting:

rule	transistor count	comments
$y[0..M-1][0, 1, 2, 3]$	0	wires
xa	$8(M-1)/3$	4-ary OR-tree approx.
approx. total	$3M-2$	

CMOS-implementable PRS

$$_x0 \rightarrow _y00\uparrow \quad _x1 \rightarrow _y01\uparrow$$

$$\neg_x0 \rightarrow _y00\downarrow \quad \neg_x1 \rightarrow _y01\downarrow$$

$$_x0 \rightarrow _y10\uparrow \quad _x1 \rightarrow _y11\uparrow$$

$$\neg_x0 \rightarrow _y10\downarrow \quad \neg_x1 \rightarrow _y11\downarrow$$

$$y0a \vee y1a \rightarrow _xa\downarrow$$

$$\neg y0a \wedge \neg y1a \rightarrow _xa\uparrow$$

$$\neg_y00 \rightarrow y00\uparrow \quad \neg_y01 \rightarrow y01\uparrow$$

$$_y00 \rightarrow y00\downarrow \quad _y01 \rightarrow y01\downarrow$$

$$\neg_y10 \rightarrow y10\uparrow \quad \neg_y11 \rightarrow y11\uparrow$$

$$_y10 \rightarrow y10\downarrow \quad _y11 \rightarrow y11\downarrow$$

D.6 NODE

NODE latches data from SPLIT.

```
*[[si];
  [x0 → y0↑; xa↑; [¬x0]; s↑; so↑; xa↓; [¬si]; y0↓; s↓; so↓
  [x1 → y1↑; xa↑; [¬x1]; s↑; so↑; xa↓; [¬si]; y1↓; s↓; so↓
  ]
]
```


The s state variable is necessary for bubble reshuffling. It breaks a cycle of isochronic branches with an odd number of bubbles (See Section D.8), which would be impossible to make CMOS-implementable.

PRS

$$\begin{array}{ll} \neg s \wedge si \wedge x0 \rightarrow y0\uparrow & \neg s \wedge si \wedge x1 \rightarrow y1\uparrow \\ \neg si & \rightarrow y0\downarrow \quad \neg si & \rightarrow y1\downarrow \end{array}$$

$$\begin{array}{l} \neg so \wedge vy \rightarrow xa\uparrow \\ so \vee \neg vy \rightarrow xa\downarrow \end{array}$$

$$\begin{array}{l} vy \wedge \neg x0 \wedge \neg x1 \rightarrow s\uparrow \\ \neg vy & \rightarrow s\downarrow \end{array}$$

$$\begin{array}{l} s \rightarrow so\uparrow \\ \neg s \rightarrow so\downarrow \end{array}$$

$$\begin{array}{l} y0 \vee y1 \rightarrow vy\uparrow \\ \neg y0 \wedge \neg y1 \rightarrow vy\downarrow \end{array}$$

1-of-2 transistor approximate accounting:

rule	transistor count	comments
$y[0, 1][0, 1]$	16	
xa	4	
so	8	
vy	4	
total	32	

1-of-4 transistor approximate accounting:

rule	transistor count	comments
$y[0, 1, 2, 3][0, 1, 2, 3]$	32	
xa	4	
so	10	
vy	8	
total	54	

CMOS-implementable PRS

$$\begin{array}{l} \neg \neg s \rightarrow \neg \neg s\uparrow \\ \neg s \rightarrow \neg \neg s\downarrow \end{array}$$

$$\begin{array}{ll}
\neg _s \wedge \neg _si \wedge \neg _x0 \rightarrow y0\uparrow & \neg _s \wedge \neg _si \wedge \neg _x1 \rightarrow y1\uparrow \\
_si \rightarrow y0\downarrow & _si \rightarrow y1\downarrow \\
\\
_so \wedge vy \rightarrow _xa\downarrow & \\
\neg _so \vee \neg vy \rightarrow _xa\uparrow & \\
\\
vy \wedge _x0 \wedge _x1 \rightarrow _s\downarrow & \\
\neg vy \rightarrow _s\uparrow & \\
\\
_s \rightarrow _so\downarrow & \\
\neg _s \rightarrow _so\uparrow & \\
\\
\neg y0 \rightarrow _y0\uparrow & \neg y1 \rightarrow _y1\uparrow \\
y0 \rightarrow _y0\downarrow & y1 \rightarrow _y1\downarrow \\
\\
\neg _y0 \vee \neg _y1 \rightarrow vy\uparrow & \\
_y0 \wedge _y1 \rightarrow vy\downarrow &
\end{array}$$

D.7 C

C in the ring deserializer of Section D.4 is a C-element taking in the environment enable signal and the last node's s_o signal to produce first node's s_i signal. s_i indicates whether we are in the up or down phase of the serial-to-parallel conversion.

PRS

$$\begin{array}{l}
\neg so \wedge xe \rightarrow si\uparrow \\
so \wedge \neg xe \rightarrow si\downarrow
\end{array}$$

A C-element costs 8 transistors.

CMOS-implementable PRS

$$\begin{array}{l}
_so \wedge xe \rightarrow _si\downarrow \\
\neg _so \wedge \neg xe \rightarrow _si\uparrow
\end{array}$$

D.8 RING

In the interest of bubble reshuffling, the deserial ring NODEs and C-element should be described in a single process.

```

*[[ye]; s0↑;
  [x0 → y00↑; x0a↑; [¬x0]
  []x1 → y01↑; x0a↑; [¬x1]
  ]; s01↑; s1↑; x0a↓;
...
  [x0 → ym0↑; xma↑; [¬x0]
  []x1 → ym1↑; xma↑; [¬x1]
  ]; sm(m + 1)↑; s(m + 1)↑; xma↓;
...
  [x0 → y(M - 1)0↑; x(M - 1)a↑; [¬x0]
  []x1 → y(M - 1)1↑; x(M - 1)a↑; [¬x1]
  ]; s(M - 1)M↑; sM↑; x(M - 1)a↓;
  [¬ye]; s0↓;
  y00↓, y01↓; s01↓; s1↓;
...
  ym0↓, ym1↓; sm(m + 1)↓; s(m + 1)↓
...
  y(M - 1)0↓, y(M - 1)1↓; s(M - 1)M↓; sM↓
]

```

Recall in Section D.6 that we had an apparently extraneous state variable s . If we removed the state variable, the NODE processes would share isochronic branches, which would create a cycle of isochronic branches among NODES. Further, this cycle would have an odd number of bubbles and be impossible to make CMOS-implementable.

PRS

Y data:

$$\begin{array}{ll}
\neg s01 \wedge s0 \wedge x0 \rightarrow y00\uparrow & \neg s01 \wedge s0 \wedge x1 \rightarrow y01\uparrow \\
\neg s0 & \rightarrow y00\downarrow \quad \neg s0 \quad \rightarrow y01\downarrow
\end{array}$$

...

$$\begin{array}{ll}
\neg sm(m + 1) \wedge sm \wedge x0 \rightarrow ym0\uparrow & \neg sm(m + 1) \wedge sm \wedge x1 \rightarrow ym1\uparrow \\
\neg sm & \rightarrow ym0\downarrow \quad \neg sm \quad \rightarrow ym1\downarrow
\end{array}$$

...

$$\begin{array}{ll}
\neg s(M - 1)M \wedge s(M - 1) \wedge x0 \rightarrow y(M - 1)0\uparrow & \neg s(M - 1)M \wedge s(M - 1) \wedge x1 \rightarrow y(M - 1)1\uparrow \\
\neg s(M - 1) & \rightarrow y(M - 1)0\downarrow \quad \neg s(M - 1) \quad \rightarrow y(M - 1)1\downarrow
\end{array}$$

X acknowledge:

$$\neg s01 \wedge v0y \rightarrow x0a\uparrow$$

$$s01 \vee \neg v0y \rightarrow x0a\downarrow$$

...

$$\neg sm(m+1) \wedge vmy \rightarrow xma\uparrow$$

$$sm(m+1) \vee \neg vmy \rightarrow xma\downarrow$$

...

$$\neg s(M-1)M \wedge v(M-1)y \rightarrow x(M-1)a\uparrow$$

$$s(M-1)M \vee \neg v(M-1)y \rightarrow x(M-1)a\downarrow$$

S buffer states:

$$v0y \wedge \neg x0 \wedge \neg x1 \rightarrow s01\uparrow$$

$$\neg v0y \rightarrow s01\downarrow$$

...

$$vmy \wedge \neg x0 \wedge \neg x1 \rightarrow sm(m+1)\uparrow$$

$$\neg vmy \rightarrow sm(m+1)\downarrow$$

...

$$v(M-1)y \wedge \neg x0 \wedge \neg x1 \rightarrow s(M-1)M\uparrow$$

$$\neg v(M-1)y \rightarrow s(M-1)M\downarrow$$

S input states:

$$\neg sM \wedge ye \rightarrow s0\uparrow$$

$$sM \wedge \neg ye \rightarrow s0\downarrow$$

$$s01 \rightarrow s1\uparrow$$

$$\neg s01 \rightarrow s1\downarrow$$

...

$$sm(m+1) \rightarrow s(m+1)\uparrow$$

$$\neg sm(m+1) \rightarrow s(m+1)\downarrow$$

...

$$s(M-1)M \rightarrow sM\uparrow$$

$$\neg s(M-1)M \rightarrow sM\downarrow$$

Y valid detectors:

$$y00 \vee y01 \rightarrow v0y\uparrow$$

$$\neg y00 \wedge \neg y01 \rightarrow v0y\downarrow$$

...

$$\begin{aligned} ym0 \vee ym1 &\rightarrow vmy\uparrow \\ \neg ym0 \wedge \neg ym1 &\rightarrow vmy\downarrow \end{aligned}$$

...

$$\begin{aligned} y(M-1)0 \vee y(M-1)1 &\rightarrow v(M-1)y\uparrow \\ \neg y(M-1)0 \wedge \neg y(M-1)1 &\rightarrow v(M-1)y\downarrow \end{aligned}$$

CMOS-implementable PRS

Y data:

$$\begin{aligned} \neg_s01 &\rightarrow _s01\uparrow \\ _s01 &\rightarrow _s01\downarrow \end{aligned}$$

$$\begin{array}{ll} \neg_s01 \wedge \neg_s0 \wedge \neg_x0 &\rightarrow y00\uparrow & \neg_s01 \wedge \neg_s0 \wedge \neg_x1 &\rightarrow y01\uparrow \\ _s0 &\rightarrow y00\downarrow & _s0 &\rightarrow y01\downarrow \end{array}$$

...

$$\begin{aligned} \neg_sm(m+1) &\rightarrow _sm(m+1)\uparrow \\ _sm(m+1) &\rightarrow _sm(m+1)\downarrow \end{aligned}$$

$$\begin{array}{ll} \neg_sm(m+1) \wedge \neg_sm \wedge \neg_x0 &\rightarrow ym0\uparrow & \neg_sm(m+1) \wedge \neg_sm \wedge \neg_x1 &\rightarrow ym1\uparrow \\ _sm &\rightarrow ym0\downarrow & _sm &\rightarrow ym1\downarrow \end{array}$$

...

$$\begin{aligned} \neg_s(M-1)M &\rightarrow _s(M-1)M\downarrow \\ _s(M-1)M &\rightarrow _s(M-1)M\downarrow \end{aligned}$$

$$\begin{array}{ll} \neg_s(M-1)M \wedge \neg_s(M-1) \wedge \neg_x0 &\rightarrow y(M-1)0\uparrow & \neg_s(M-1)M \wedge \neg_s(M-1) \wedge \neg_x1 &\rightarrow y(M-1)0\downarrow \\ _s(M-1) &\rightarrow y(M-1)0\downarrow & _s(M-1) &\rightarrow y(M-1)0\downarrow \end{array}$$

X acknowledge:

$$\begin{aligned} \neg_s01 \wedge \neg_v0y &\rightarrow x0a\uparrow \\ _s01 \vee _v0y &\rightarrow x0a\downarrow \end{aligned}$$

...

$$\begin{aligned} \neg_sm(m+1) \wedge \neg_vmy &\rightarrow xma\uparrow \\ _sm(m+1) \vee _vmy &\rightarrow xma\downarrow \end{aligned}$$

...

$$\begin{aligned}\neg _s(M-1)M \wedge \neg _v(M-1)y &\rightarrow x(M-1)a\uparrow \\ _s(M-1)M \vee _v(M-1)y &\rightarrow x(M-1)a\downarrow\end{aligned}$$

S buffer states:

$$\begin{aligned}\neg _v0y &\rightarrow _v0y & _v0y \wedge _x0 \wedge _x1 &\rightarrow _s01\downarrow \\ _v0y &\rightarrow _v0y & \neg _v0y &\rightarrow _s01\uparrow\end{aligned}$$

...

$$\begin{aligned}\neg _vmy &\rightarrow _vmy & _vmy \wedge _x0 \wedge _x1 &\rightarrow _sm(m+1)\downarrow \\ _vmy &\rightarrow _vmy & \neg _vmy &\rightarrow _sm(m+1)\uparrow\end{aligned}$$

...

$$\begin{aligned}\neg _v(M-1)y &\rightarrow _v(M-1)y & _v(M-1)y \wedge _x0 \wedge _x1 &\rightarrow _s(M-1)M\downarrow \\ _v(M-1)y &\rightarrow _v(M-1)y & \neg _v(M-1)y &\rightarrow _s(M-1)M\uparrow\end{aligned}$$

S input/output states:

$$\begin{aligned}_sM \wedge ye &\rightarrow _s0\downarrow \\ \neg _sM \wedge \neg ye &\rightarrow _s0\uparrow\end{aligned}$$

$$\begin{aligned}\neg _s01 &\rightarrow _s1\uparrow \\ _s01 &\rightarrow _s1\downarrow\end{aligned}$$

...

$$\begin{aligned}\neg _sm(m+1) &\rightarrow _s(m+1)\uparrow \\ _sm(m+1) &\rightarrow _s(m+1)\downarrow\end{aligned}$$

...

$$\begin{aligned}\neg _s(M-1)M &\rightarrow _sM\uparrow \\ _s(M-1)M &\rightarrow _sM\downarrow\end{aligned}$$

Y valid detectors:

$$\begin{aligned}y00 \vee y01 &\rightarrow _v0y\downarrow \\ \neg y00 \wedge \neg y01 &\rightarrow _v0y\uparrow\end{aligned}$$

...

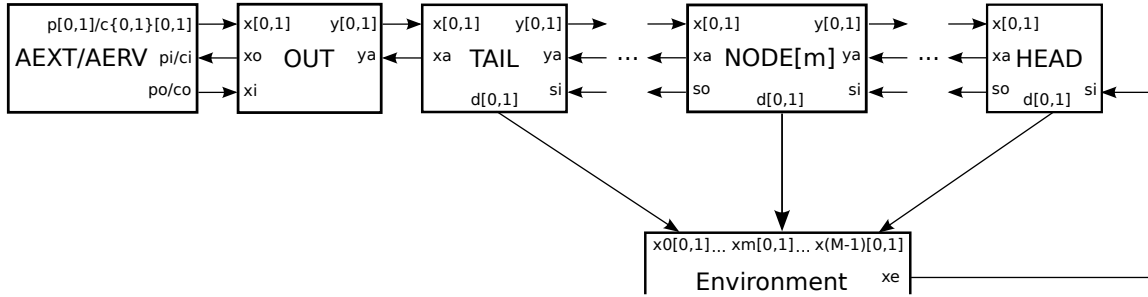
$$\begin{aligned}ym0 \vee ym1 &\rightarrow _vmy\downarrow \\ \neg ym0 \wedge \neg ym1 &\rightarrow _vmy\uparrow\end{aligned}$$

...

$$\begin{aligned}y(M-1)0 \vee y(M-1)1 &\rightarrow _v(M-1)y\downarrow \\ \neg y(M-1)0 \wedge \neg y(M-1)1 &\rightarrow _v(M-1)y\uparrow\end{aligned}$$

D.9 CHAIN Deserializer

This design uses a chain of nodes to sequence words into their respective place in the parallel output. That is, it takes as input a1ofN data and outputs eMx1ofN data. We decompose this process into HEAD, NODE, and TAIL processes. Although it has a bit more expensive 1-of-4 implementation than the split ring of Section D.4, its signal fanouts remain constant as the number of words in the packet grows. The below figure shows the decomposition for packets consisting of M 1-of-2 groups.



An OUT a1ofN process (described above) first converts the AEXT/AERV serial communication protocol to the standard a1ofN protocol. Each link in the chain outputs one of the words in the parallel output.

1-of-2 approximate scaling:

component	transistors/component	components/deserializer	transistors/deserializer
OUT a1ofN	4	1	4
HEAD	28	1	28
NODE	38	$M - 2$	$38(M - 2)$
TAIL	34	1	34
transistors/deserializer			$38M - 10$

1-of-4 approximate scaling:

component	transistors/component	components/deserializer	transistors/deserializer
OUT a1ofN	4	1	4
HEAD	50	1	50
NODE	68	$M - 2$	$68(M - 2)$
TAIL	66	1	66
transistors/deserializer			$68M - 16$

For the transmitter encoding 4096 neurons as 1-of-2 or 1-of-4 words, we would need 12 and 6 links in the chain, respectively.

1-of-2 approximate accounting:

component	transistors/component	components/deserializer	transistors/deserializer
OUT a1ofN	4	1	4
HEAD	28	1	28
NODE	38	10	380
TAIL	34	1	34
total transistors/deserializer			446

1-of-4 approximate accounting:

component	transistors/component	components/deserializer	transistors/deserializer
OUT a1ofN	4	1	4
HEAD	50	1	50
NODE	68	4	272
TAIL	66	1	66
total transistors/deserializer			392

D.10 HEAD

* $[[si \wedge X]; D\uparrow; xa\uparrow; [\neg X]; so\uparrow; xa\downarrow; [\neg si]; D\downarrow; so\downarrow]$

PRS

$si \wedge x0 \rightarrow d0\uparrow$ $si \wedge x1 \rightarrow d1\uparrow$

$\neg si \rightarrow d0\downarrow$ $\neg si \rightarrow d1\downarrow$

$\neg x0 \wedge \neg x1 \wedge vd \rightarrow so\uparrow$

$x0 \vee x1 \vee \neg vd \rightarrow so\downarrow$

$\neg so \wedge vd \rightarrow xa\uparrow$

$so \vee \neg vd \rightarrow xa\downarrow$

$d0 \vee d1 \rightarrow vd\uparrow$

$\neg d0 \wedge \neg d1 \rightarrow vd\downarrow$

1-of-2 transistor approximate accounting:

rule	transistor count	comments
$d[0, 1]$	14	
s_o	6	
xa	4	
vd	4	
total	28	

1-of-4 transistor approximate accounting:

rule	transistor count	comments
$d[0, 1, 2, 3]$	28	
s_o	10	
xa	4	
vd	8	
total	50	

CMOS-implementable PRS

$$si \wedge x0 \rightarrow _d0\downarrow \quad si \wedge x1 \rightarrow _d1\downarrow$$

$$\neg si \rightarrow _d0\uparrow \quad \neg si \rightarrow _d1\uparrow$$

$$\neg_d0 \rightarrow d0\uparrow \quad \neg_d1 \rightarrow d1\uparrow$$

$$_d0 \rightarrow d0\downarrow \quad _d1 \rightarrow d1\downarrow$$

$$\neg x0 \wedge \neg x1 \wedge \neg vd \rightarrow so\uparrow$$

$$x0 \vee x1 \vee vd \rightarrow so\downarrow$$

$$\neg so \rightarrow _so\uparrow$$

$$so \rightarrow so\downarrow$$

$$_so \wedge vd \rightarrow _xa\downarrow$$

$$\neg_so \vee \neg vd \rightarrow _xa\uparrow$$

$$\neg_d0 \vee \neg_d1 \rightarrow vd\uparrow$$

$$_d0 \wedge _d1 \rightarrow vd\downarrow$$

$$\neg vd \rightarrow _vd\uparrow$$

$$vd \rightarrow _vd\downarrow$$

$$\neg_xa \rightarrow xa\uparrow$$

$$_xa \rightarrow xa\downarrow$$

1-of-2 transistor accounting:

rule	transistor count	comments
$_d[0, 1]$	14	
$d[0, 1]$	4	
$_s_o$	6	
$_s_o$	2	
$_x_a$	4	
vd	4	
$_vd$	2	
x_a	2	
total	38	

1-of-4 transistor accounting:

rule	transistor count	comments
$_d[0, 1, 2, 3]$	28	
$d[0, 1, 2, 3]$	8	
$_s_o$	10	
$_s_o$	2	
$_x_a$	4	
vd	8	
$_vd$	2	
x_a	2	
total	64	

D.11 NODE

$$*[[\neg si \wedge X \longrightarrow Y\uparrow; [ya]; xa\uparrow; [\neg X]; Y\downarrow; [si \longrightarrow [\neg ya]; xa\downarrow] \neg ya \longrightarrow xa\downarrow]$$

$$[[si \wedge X \longrightarrow D\uparrow; xa\uparrow; [\neg X]; so\uparrow; xa\downarrow; [\neg si]; D\downarrow; so\downarrow$$

$$]]$$

PRS

$$\neg si \wedge x0 \rightarrow y0\uparrow \quad \neg si \wedge x1 \rightarrow y1\uparrow$$

$$si \vee \neg x0 \rightarrow y0\downarrow \quad si \vee \neg x1 \rightarrow y1\downarrow$$

$$si \wedge x0 \rightarrow d0\uparrow \quad si \wedge x1 \rightarrow d1\uparrow$$

$$\neg si \rightarrow d0\downarrow \quad \neg si \rightarrow d1\downarrow$$

$$\neg x0 \wedge \neg x1 \wedge vd \rightarrow so\uparrow$$

$$x0 \vee x1 \vee \neg vd \rightarrow so\downarrow$$

$$ya \vee \neg so \wedge vd \rightarrow xa \uparrow$$

$$\neg ya \wedge (so \vee \neg vd) \rightarrow xa \downarrow$$

$$d0 \vee d1 \rightarrow vd \uparrow$$

$$\neg d0 \wedge \neg d1 \rightarrow vd \downarrow$$

1-of-2 transistor approximate accounting:

rule	transistor count	comments
$y[0, 1]$	8	
$d[0, 1]$	14	
s_o	6	
xa	6	
vd	4	
total	38	

1-of-4 transistor approximate accounting:

rule	transistor count	comments
$y[0, 1, 2, 3]$	16	
$d[0, 1, 2, 3]$	28	
s_o	10	
xa	6	
vd	8	
total	68	

CMOS-implementable PRS

$$\neg x0 \rightarrow _x0 \downarrow \quad \neg x1 \rightarrow _x1 \downarrow$$

$$x0 \rightarrow _x0 \uparrow \quad x1 \rightarrow _x1 \uparrow$$

$$\neg si \wedge \neg _x0 \rightarrow y0 \uparrow \quad \neg si \wedge \neg _x1 \rightarrow y1 \uparrow$$

$$si \vee _x0 \rightarrow y0 \downarrow \quad si \vee _x1 \rightarrow y1 \downarrow$$

$$\neg si \rightarrow _si \uparrow$$

$$si \rightarrow _si \downarrow$$

$$\neg _si \wedge \neg _x0 \rightarrow d0 \uparrow \quad \neg _si \wedge \neg _x1 \rightarrow d1 \downarrow$$

$$_si \rightarrow d0 \downarrow \quad _si \rightarrow _d1 \uparrow$$

$$\neg _vd \rightarrow vd \uparrow$$

$$_vd \rightarrow vd \downarrow$$

$$_x0 \wedge _x1 \wedge vd \quad \rightarrow _so\downarrow$$

$$\neg_x0 \vee \neg_x1 \vee \neg vd \rightarrow _so\uparrow$$

$$\neg_so \rightarrow so\uparrow$$

$$_so \rightarrow so\downarrow$$

$$\neg so \rightarrow ___so\uparrow$$

$$so \rightarrow ___so\downarrow$$

$$ya \vee ___so \wedge vd \quad \rightarrow _xa\downarrow$$

$$\neg ya \wedge (\neg ___so \vee \neg vd) \rightarrow _xa\uparrow$$

$$d0 \vee d1 \quad \rightarrow _vd\downarrow$$

$$\neg d0 \wedge \neg d1 \rightarrow \neg vd\uparrow$$

$$\neg_xa \rightarrow xa\uparrow$$

$$_xa \rightarrow xa\downarrow$$

1-of-2 transistor accounting:

rule	transistor count	comments
$_x[0, 1]$	4	
$y[0, 1]$	8	
$_s_i$	2	
$d[0, 1]$	14	
vd	2	
$_s_o$	6	
s_o	2	
$___s_o$	2	
$_xa$	6	
$_vd$	4	
xa	2	
total	38	

1-of-4 transistor accounting:

rule	transistor count	comments
$_x[0, 1, 2, 3]$	8	
$y[0, 1, 2, 3]$	16	
$_s_i$	2	
$d[0, 1, 2, 3]$	28	
vd	2	
$_s_o$	10	
s_o	2	
$_xa$	6	
$_vd$	8	
xa	2	
total	84	

D.12 TAIL

$*[[\neg si \wedge X \longrightarrow Y\uparrow; [ya]; xa\uparrow; [\neg X]; Y\downarrow; [si \longrightarrow [\neg ya]; xa\downarrow] \neg ya \longrightarrow xa\downarrow]$
 $\llbracket si \wedge X \longrightarrow D\uparrow; xa\uparrow; [\neg si \wedge \neg X]; D\downarrow; xa\downarrow$
 $\rrbracket]$

PRS

$\neg si \wedge x0 \wedge \neg d0 \rightarrow y0\uparrow$ $\neg si \wedge x1 \wedge \neg d1 \rightarrow y1\uparrow$
 $si \vee \neg x0 \vee d0 \rightarrow y0\downarrow$ $si \vee \neg x1 \vee d1 \rightarrow y1\downarrow$

 $si \wedge x0 \rightarrow d0\uparrow$ $si \wedge x1 \rightarrow d1\uparrow$
 $\neg si \wedge \neg x0 \rightarrow d0\downarrow$ $\neg si \wedge \neg x1 \rightarrow d1\downarrow$

 $ya \vee d0 \vee d1 \rightarrow xa\uparrow$
 $\neg ya \wedge \neg d0 \wedge \neg d1 \rightarrow xa\downarrow$

1-of-2 transistor approximate accounting:

rule	transistor count	comments
$y[0, 1]$	12	
$d[0, 1]$	16	
xa	6	
total	34	

1-of-4 transistor approximate accounting:

rule	transistor count	comments
$y[0, 1, 2, 3]$	24	
$d[0, 1, 2, 3]$	32	
xa	10	
total	66	

CMOS-implementable PRS

$$\begin{array}{ll}
\neg si \wedge \neg x0 \wedge \neg d0 \rightarrow y0\uparrow & \neg si \wedge \neg x1 \wedge \neg d1 \rightarrow y1\uparrow \\
si \vee x0 \vee d0 \rightarrow y0\downarrow & si \vee x1 \vee d1 \rightarrow y1\downarrow \\
\neg si \rightarrow \neg si\uparrow & \\
si \rightarrow \neg si\downarrow & \\
\neg \neg si \wedge \neg x0 \rightarrow d0\uparrow & \neg \neg si \wedge \neg x1 \rightarrow d1\uparrow \\
\neg si \wedge x0 \rightarrow d0\downarrow & \neg si \wedge \neg x1 \rightarrow d1\downarrow \\
ya \vee d0 \vee d1 \rightarrow xa\downarrow & \\
\neg ya \wedge \neg d0 \wedge \neg d1 \rightarrow xa\uparrow &
\end{array}$$

1-of-2 transistor accounting:

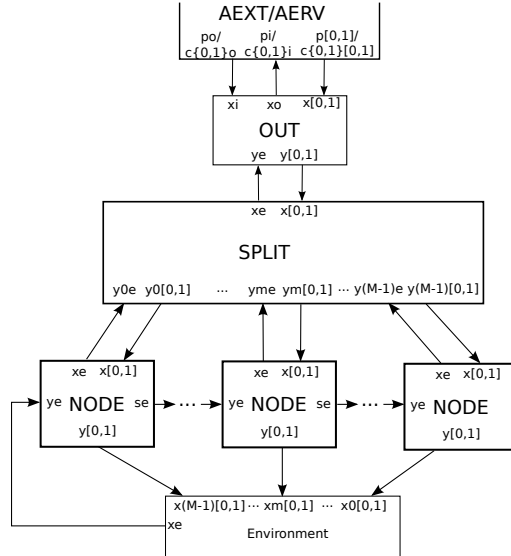
rule	transistor count	comments
$y[0, 1]$	12	
$\neg si$	2	
$d[0, 1]$	16	
$\neg xa$	6	
total	36	

1-of-4 transistor accounting:

rule	transistor count	comments
$y[0, 1, 2, 3]$	24	
$\neg si$	2	
$d[0, 1, 2, 3]$	32	
xa	10	
total	68	

D.13 Split Chain Deserializer

This design uses a chain of nodes pulling data from a central splitter to sequence words into their respective place in the parallel output. We decompose this process into SPLIT and NODE. The below figure shows the decomposition for packets containing M 1-of-2 words.



An OUT e1ofN process (described above) first converts the AEXT/AERV serial communication protocol to the standard e1ofN protocol.

1-of-2 approximate scaling:

component	transistors/component	components/deserializer	transistors/deserializer
OUT e1ofN	4	1	4
SPLIT	$17M - 2$	1	$17M - 2$
NODE (int)	28	$M - 1$	$28(M - 1)$
NODE (end)	20	1	20
approx. transistors/deserializer			$45M - 6$

1-of-4 approximate scaling:

component	transistors/component	components/deserializer	transistors/deserializer
OUT e1ofN	4	1	4
SPLIT	$31M - 2$	1	$31M - 2$
NODE (int)	50	$M - 1$	$50(M - 1)$
NODE (end)	38	1	38
approx. transistors/deserializer			$81M - 10$

For the transmitter to handle 4096 neurons encoded as 1-of-2 or 1-of-4 words, we would need 12 and 6 NODEs, respectively.

1-of-2 accounting:

component	transistors/component	components/deserializer	transistors/deserializer
OUT e1ofN	4	1	4
SPLIT	198	1	198
NODE (int)	28	11	308
NODE (end)	20	1	20
total transistors/deserializer			530

1-of-4 accounting:

component	transistors/component	components/deserializer	transistors/deserializer
OUT e1ofN	4	1	4
SPLIT	184	1	184
NODE (int)	50	5	250
NODE (end)	38	1	38
total transistors/deserializer			476

D.14 SPLIT

SPLIT takes incoming words and routes them to their respective locations in the parallel output.

For M words per packet,

```

*[[y0e → xe↑;
  [x0 → y00↑; [¬y0e]; xe↓; [¬x0]; y00↓
  [x1 → y01↑; [¬y0e]; xe↓; [¬x1]; y01↓
  ]
  [] ...
  [y(M-1)e → xe↑;
  [x1 → y(M-1)0↑; [¬y(M-1)e]; xe↓; [¬x0]; y(M-1)0↓
  [x1 → y(M-1)1↑; [¬y(M-1)e]; xe↓; [¬x1]; y(M-1)1↓
  ]
  ]
  ]

```

For a 2-word packet,


```

*[[y0e → xe↑;
  [x0 → y00↑; [¬y0e]; xe↓; [¬x0]; y00↓
  []x1 → y01↑; [¬y0e]; xe↓; [¬x1]; y01↓
  ]
 []y1e → xe↑;
  [x1 → y10↑; [¬y1e]; xe↓; [¬x0]; y10↓
  []x1 → y11↑; [¬y1e]; xe↓; [¬x1]; y11↓
  ]
]]

```

```

y0e ∨ y1e → xe↑
¬y0e ∧ ¬y1e → xe↓

```

```

y0e ∧ x0 → y00↑      y0e ∧ x1 → y01↑
¬x0      → y00↓      ¬x1      → y01↓

```

```

y1e ∧ x0 → y10↑      y1e ∧ x1 → y11↑
¬x0      → y10↓      ¬x1      → y11↓

```

1-of-2 transistor accounting:

rule	transistor count	comments
xe	$8(M - 1)/3$	4-ary OR-tree approx.
$y[0..M - 1][0, 1]$	$14M$	
approx. total	$17M - 2$	

1-of-4 transistor accounting:

rule	transistor count	comments
xe	$8(M - 1)/3$	4-ary OR-tree approx.
$y[0..M - 1][0, 1, 2, 3]$	$28M$	
approx. total	$31M - 2$	

D.15 NODE

NODE latches data from SPLIT. For beginning and intermediate NODEs,

```

*[[ye]; xe↑; [X]; Y↑; xe↓; [¬X]; se↑; [¬ye]; Y↓; se↓]

```

```

ye ∧ ¬vy → xe↑
¬ye ∨ vy → xe↓

```

$$\begin{aligned}
x0 &\rightarrow y0\uparrow & x1 &\rightarrow y1\uparrow \\
\neg ye &\rightarrow y0\downarrow & \neg ye &\rightarrow y1\downarrow \\
\neg x0 \wedge \neg x1 \wedge vy &\rightarrow se\uparrow \\
\neg vy && &\rightarrow se\downarrow \\
y0 \vee y1 &\rightarrow vy\uparrow \\
\neg y0 \wedge \neg y1 &\rightarrow vy\downarrow
\end{aligned}$$

1-of-2 transistor accounting:

rule	transistor count	comments
xe	4	
$y[0, 1]$	12	
se	8	
vy	4	
total	28	

1-of-4 transistor accounting:

rule	transistor count	comments
xe	4	
$y[0, 1, 2, 3]$	28	
se	10	
vy	8	
total	50	

The end NODE of a chain does not need to forward an se signal,

$$*[[ye]; xe\uparrow; [X]; Y\uparrow; xe\downarrow; [\neg X \wedge \neg ye]; Y\downarrow]$$

$$\begin{aligned}
ye \wedge \neg y0 \wedge \neg y1 &\rightarrow xe\uparrow \\
\neg ye \vee y0 \vee y1 &\rightarrow xe\downarrow
\end{aligned}$$

$$\begin{aligned}
x0 &\rightarrow y0\uparrow & x1 &\rightarrow y1\uparrow \\
\neg x0 \wedge \neg ye &\rightarrow y0\downarrow & \neg x1 \wedge \neg ye &\rightarrow y1\downarrow
\end{aligned}$$

1-of-2 transistor accounting:

rule	transistor count	comments
xe	6	
$y[0, 1]$	14	
total	20	

1-of-4 transistor accounting:

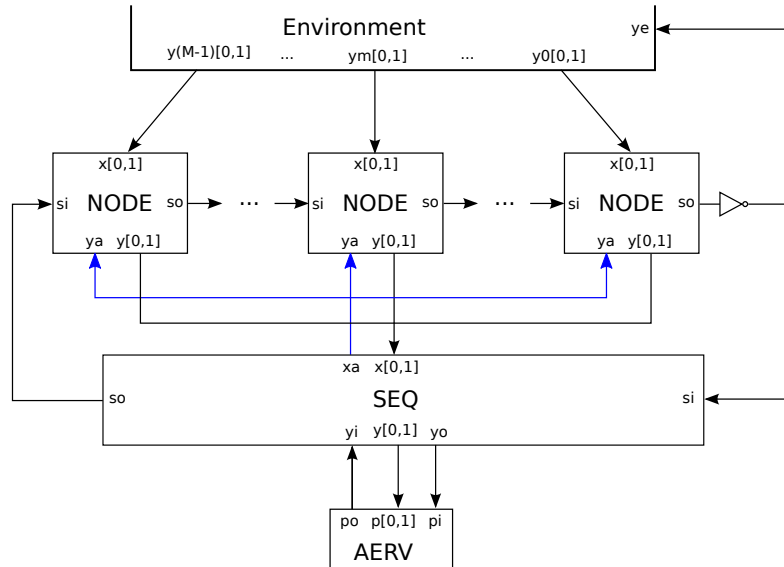
rule	transistor count	comments
xe	10	
$y[0, 1, 2, 3]$	28	
total	38	

D.16 Serializer

The serializer converts M-1-of-N parallel data into in 1-of-N serial data. We cannot create standard a1ofN or e1ofN input interfaces to the serial protocol (like we could with output interfaces in Sections D.1 and D.2) because the serial protocol requires a signal indicating the transitions between packets, which the standard a1ofN and e1ofN channels cannot provide.

D.17 Ring serializer

The ring serializer uses a ring of NODES to sequence the words of an eMx1ofN channel. A C-element indicates the transition between the up and down phases of the sequencing.



Note the shared ye and $y[0, 1]$ between the NODE processes.

1-of-2 approximate scaling:

component	transistors/component	components/serializer	transistors/serializer
NODE	34	M	$34M$
SEQ	38	1	38
approx. transistors/serializer			$34M + 38$

1-of-4 approximate scaling:

component	transistors/component	components/serializer	transistors/serializer
NODE	54	M	$54M$
SEQ	56	1	56
approx. transistors/serializer			$54M + 56$

For the receiver to handle 4096 neurons (and no data) encoded as 1-of-2 or 1-of-4 words, we would need 12 and 6 NODEs, respectively.

1-of-2 accounting:

component	transistors/component	components/serializer	transistors/serializer
NODE	34	12	408
SEQ	38	1	38
total transistors/serializer			446

1-of-4 accounting:

component	transistors/component	components/serializer	transistors/serializer
NODE	54	6	324
SEQ	56	1	56
total transistors/serializer			380

D.18 Ring serializer NODE

```

*[[si];
  [ x0 → y0↑; [ya]; u↑; y0↓; [¬ya]; so↑; [¬si]; u↓; [¬x0]; so↓
  [x1 → y1↑; [ya]; u↑; y1↓; [¬ya]; so↑; [¬si]; u↓; [¬x1]; so↓
]]

```

PRS

$$\begin{aligned}
 x0 \wedge \neg u \wedge si &\rightarrow y0\uparrow & x1 \wedge \neg u \wedge si &\rightarrow y1\uparrow \\
 u \wedge \neg so &\rightarrow y0\downarrow & u \wedge \neg so &\rightarrow y1\downarrow \\
 si \wedge ya &\rightarrow u\uparrow \\
 \neg si &\rightarrow u\downarrow
 \end{aligned}$$

$$\begin{aligned}
u \wedge \neg ya &\rightarrow so\uparrow \\
\neg u \wedge \neg x0 \wedge \neg x1 &\rightarrow so\downarrow
\end{aligned}$$

1-of-2 transistor accounting:

rule	transistor count	comments
$y[0, 1]$	18	
u	7	
so	9	
total	34	

1-of-4 transistor accounting:

rule	transistor count	comments
$y[0, 1, 2, 3]$	36	
u	7	
so	11	
total	54	

CMOS-implementable PRS

$$\neg_u \rightarrow _u\uparrow$$

$$_u \rightarrow _u\downarrow$$

$$\neg_so \rightarrow _so\uparrow$$

$$_so \rightarrow _so\downarrow$$

$$\neg_x0 \wedge \neg_u \wedge \neg_si \rightarrow y0\uparrow$$

$$\neg_x1 \wedge \neg_u \wedge \neg_si \rightarrow y1\uparrow$$

$$_u \wedge _so \rightarrow y0\downarrow$$

$$_u \wedge _so \rightarrow y1\downarrow$$

$$\neg_si \rightarrow _si\uparrow$$

$$_si \rightarrow _si\downarrow$$

$$_si \wedge ya \rightarrow _u\downarrow$$

$$\neg_si \rightarrow _u\uparrow$$

$$\neg_u \rightarrow _u\uparrow$$

$$_u \rightarrow _u\downarrow$$

$$\neg_u \wedge \neg ya \rightarrow so\uparrow$$

$$_u \wedge x0 \wedge x1 \rightarrow so\downarrow$$

$$\neg y0 \rightarrow _y0\uparrow$$

$$\neg y1 \rightarrow _y1\uparrow$$

$$y0 \rightarrow _y0\downarrow$$

$$y1 \rightarrow _y1\downarrow$$

D.19 Ring serializer sequencer

*[[si]; so ↑; [$x0 \vee x1$]; yo ↑;
 $\neg si \wedge yi$; yo ↓; [$\neg yi$]; so ↓]

*[[$x0 \longrightarrow y0$ ↑; [$\neg yi$]; xa ↑; [$\neg x0$]; $y0$ ↓; [yi]; xa ↓
 $\llbracket x1 \longrightarrow y1$ ↑; [$\neg yi$]; xa ↑; [$\neg x1$]; $y1$ ↓; [yi]; xa ↓
 \rrbracket]]

PRS

$x0 \vee x1 \rightarrow yo$ ↑
 $\neg si \wedge yi \rightarrow yo$ ↓

$si \rightarrow so$ ↑
 $\neg si \wedge \neg yi \wedge \neg yo \rightarrow so$ ↓

$yi \wedge x0 \rightarrow y0$ ↑ $yi \wedge x1 \rightarrow y1$ ↑
 $\neg x0 \rightarrow y0$ ↓ $\neg x1 \rightarrow y1$ ↓

$(y0 \vee y1) \wedge \neg yi \rightarrow xa$ ↑
 $yi \rightarrow xa$ ↓

1-of-2 transistor accounting:

rule	transistor count	comments
yo	8	
so	8	
$y[0, 1]$	14	
xa	8	
total	38	

1-of-4 transistor accounting:

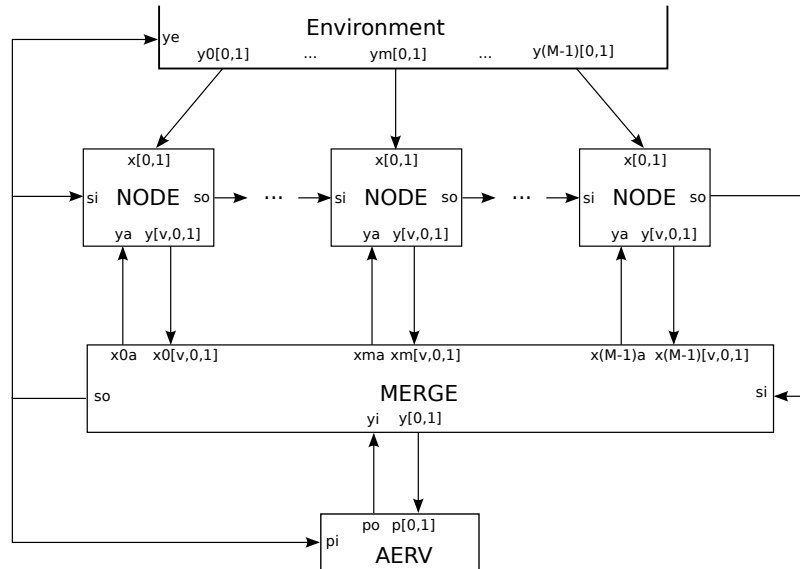
rule	transistor count	comments
yo	10	
so	8	
$y[0, 1, 2, 3]$	28	
xa	10	
total	56	

CMOS-implementable PRS

$$\begin{aligned}
 &\neg x0 \vee \neg x1 \rightarrow yo\uparrow \\
 &si \wedge yi \rightarrow yo\downarrow \\
 \\
 &\neg si \rightarrow \neg si\uparrow \\
 &si \rightarrow \neg si\downarrow \\
 \\
 &\neg si \rightarrow so\downarrow \\
 &\neg \neg si \wedge \neg yi \wedge \neg yo \rightarrow so\uparrow \\
 \\
 &\neg x0 \rightarrow \neg x0\uparrow \quad \neg x1 \rightarrow \neg x1\uparrow \\
 &x0 \rightarrow \neg x0\downarrow \quad x1 \rightarrow \neg x1\downarrow \\
 \\
 &yi \wedge \neg x0 \rightarrow y0\downarrow \quad yi \wedge \neg x1 \rightarrow y1\downarrow \\
 &\neg \neg x0 \rightarrow y0\uparrow \quad \neg \neg x1 \rightarrow y1\uparrow \\
 \\
 &(\neg y0 \vee \neg y1) \wedge \neg yi \rightarrow xa\uparrow \\
 &yi \rightarrow xa\downarrow \\
 \\
 &\neg y0 \rightarrow y0\uparrow \quad \neg y1 \rightarrow y1\uparrow \\
 &y0 \rightarrow y0\downarrow \quad y1 \rightarrow y1\downarrow
 \end{aligned}$$

D.20 Ring serializer 2

This design uses more conventional channels and more sequencing than the previous serial ring. Since there's only one serializer per neuron array, the cost per neuron is still trivial.



D.21 Ring serializer 2 NODE

$$*[[\text{si} \wedge x0 \longrightarrow y0\uparrow, yv\uparrow; [ya]; u\downarrow; y0\downarrow, yv\downarrow; [\neg ya]; so\uparrow; [\neg \text{si} \wedge \neg x0]; u\uparrow; so\downarrow \\ \llbracket \text{si} \wedge x1 \longrightarrow y1\uparrow, yv\uparrow; [ya]; u\downarrow; y1\downarrow, yv\downarrow; [\neg ya]; so\uparrow; [\neg \text{si} \wedge \neg x1]; u\uparrow; so\downarrow \\ \rrbracket]]$$

PRS

$$\begin{array}{ll} \text{si} \wedge x0 \wedge u & \rightarrow y0\uparrow & \text{si} \wedge x1 \wedge u & \rightarrow y1\uparrow \\ \neg \text{si} \vee \neg x0 \vee \neg u & \rightarrow y0\downarrow & \neg \text{si} \vee \neg x1 \vee \neg u & \rightarrow y1\downarrow \end{array}$$

$$\begin{array}{l} \text{si} \wedge (x0 \vee x1) \wedge u \quad \rightarrow yv\uparrow \\ \neg \text{si} \vee \neg x0 \wedge \neg x1 \vee \neg u \rightarrow yv\downarrow \end{array}$$

$$\begin{array}{ll} \neg \text{si} \wedge \neg x0 \wedge \neg x1 & \rightarrow u\uparrow \\ ya & \rightarrow u\downarrow \end{array}$$

$$\begin{array}{ll} \neg u \wedge \neg ya & \rightarrow so\uparrow \\ u \vee ya & \rightarrow so\downarrow \end{array}$$

CMOS-implementable PRS

$$\begin{array}{ll} \text{si} \wedge x0 \wedge u & \rightarrow _y0\downarrow & \text{si} \wedge x1 \wedge u & \rightarrow _y1\downarrow \\ \neg \text{si} \vee \neg x0 \vee \neg u & \rightarrow _y0\uparrow & \neg \text{si} \vee \neg x1 \vee \neg u & \rightarrow _y1\uparrow \end{array}$$

$$\begin{array}{l} \text{si} \wedge (x0 \vee x1) \wedge u \quad \rightarrow _yv\downarrow \\ \neg \text{si} \vee \neg x0 \wedge \neg x1 \vee \neg u \rightarrow _yv\uparrow \end{array}$$

$$\begin{array}{ll} \neg \text{si} \wedge \neg x0 \wedge \neg x1 & \rightarrow u\uparrow \\ ya & \rightarrow u\downarrow \end{array}$$

$$\begin{array}{ll} \neg u \wedge \neg ya & \rightarrow so\uparrow \\ u \vee ya & \rightarrow so\downarrow \end{array}$$

$$\begin{array}{ll} \neg _y0 \rightarrow y0\uparrow & \neg _y1 \rightarrow y1\uparrow \\ _y0 \rightarrow y0\downarrow & _y1 \rightarrow y1\downarrow \end{array}$$

D.22 Ring serializer 2 MERGE

for M words,

$$*[\neg si \wedge \neg yi]; so\uparrow; [si \wedge yi]; so\downarrow]$$

$$*[[x00 \longrightarrow u0\uparrow, [yi]; y0\uparrow; vy\uparrow; [x0v \wedge \neg yi]; x0a\uparrow; (\neg x00]; u0\downarrow; y0\downarrow; vy\downarrow), [\neg x0v]; x0a\downarrow \\ \llbracket x01 \longrightarrow u1\uparrow, [yi]; y1\uparrow; vy\uparrow; [x0v \wedge \neg yi]; x0a\uparrow; (\neg x01]; u1\downarrow; y1\downarrow; vy\downarrow), [\neg x0v]; x0a\downarrow$$

...

$$\llbracket x m 0 \longrightarrow u 0 \uparrow, [y i]; y 0 \uparrow; v y \uparrow; [x m v \wedge \neg y i]; x m a \uparrow; (\neg x m 0]; u 0 \downarrow; y 0 \downarrow; v y \downarrow), [\neg x m v]; x m a \downarrow \\ \llbracket x m 1 \longrightarrow u 1 \uparrow, [y i]; y 1 \uparrow; v y \uparrow; [x m v \wedge \neg y i]; x m a \uparrow; (\neg x m 1]; u 1 \downarrow; y 1 \downarrow; v y \downarrow), [\neg x m v]; x m a \downarrow$$

...

$$\llbracket x(M\downarrow 1)0 \longrightarrow u0\uparrow, [yi]; y0\uparrow; vy\uparrow; [x(M\downarrow 1)v \wedge \neg yi]; x(M\downarrow 1)a\uparrow; (\neg x(M\downarrow 1)0]; u0\downarrow; y0\downarrow; vy\downarrow), [\neg x(M\downarrow 1)v] \\ \llbracket x(M\downarrow 1)1 \longrightarrow u1\uparrow, [yi]; y1\uparrow; vy\uparrow; [x(M\downarrow 1)v \wedge \neg yi]; x(M\downarrow 1)a\uparrow; (\neg x(M\downarrow 1)1]; u1\downarrow; y1\downarrow; vy\downarrow), [\neg x(M\downarrow 1)v] \\ \rrbracket]$$

for $M = 2$,

$$*[\neg si \wedge \neg yi]; so\uparrow; [si \wedge yi]; so\downarrow]$$

$$*[[x00 \longrightarrow u0\uparrow, [yi]; y0\uparrow; vy\uparrow; [x0v \wedge \neg yi]; x0a\uparrow; (\neg x00]; u0\downarrow; y0\downarrow; vy\downarrow), [\neg x0v]; x0a\downarrow \\ \llbracket x01 \longrightarrow u1\uparrow, [yi]; y1\uparrow; vy\uparrow; [x0v \wedge \neg yi]; x0a\uparrow; (\neg x01]; u1\downarrow; y1\downarrow; vy\downarrow), [\neg x0v]; x0a\downarrow$$

$$\llbracket x10 \longrightarrow u0\uparrow, [yi]; y0\uparrow; vy\uparrow; [x1v \wedge \neg yi]; x1a\uparrow; (\neg x10]; u0\downarrow; y0\downarrow; vy\downarrow), [\neg x1v]; x1a\downarrow$$

$$\llbracket x11 \longrightarrow u1\uparrow, [yi]; y1\uparrow; vy\uparrow; [x1v \wedge \neg yi]; x1a\uparrow; (\neg x11]; u1\downarrow; y1\downarrow; vy\downarrow), [\neg x1v]; x1a\downarrow \\ \rrbracket]$$

$$\neg si \wedge \neg yi \rightarrow so\uparrow$$

$$si \wedge yi \rightarrow so\downarrow$$

$$x00 \vee x10 \rightarrow u0\uparrow$$

$$x01 \vee x11 \rightarrow u1\uparrow$$

$$\neg x00 \wedge \neg x10 \rightarrow u0\downarrow$$

$$\neg x01 \wedge \neg x11 \rightarrow u1\downarrow$$

$$u0 \wedge yi \rightarrow y0\uparrow$$

$$u1 \wedge yi \rightarrow y1\uparrow$$

$$\neg u0 \rightarrow y0\downarrow$$

$$\neg u1 \rightarrow y1\downarrow$$

$$y0 \vee y1 \rightarrow vy\uparrow$$

$$\neg y0 \wedge \neg y1 \rightarrow vy\downarrow$$

$$x0v \wedge \neg yi \wedge vy \rightarrow x0a\uparrow$$

$$x1v \wedge \neg yi \wedge vy \rightarrow x1a\uparrow$$

$$\neg x0v \wedge \neg vy \rightarrow x0a\downarrow$$

$$\neg x1v \wedge \neg vy \rightarrow x1a\downarrow$$

CMOS-implementable PRS

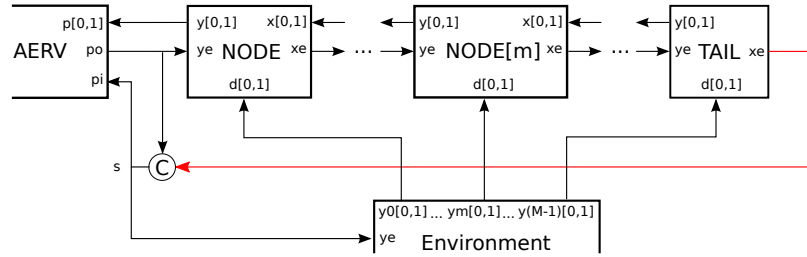
$$\neg si \wedge \neg yi \rightarrow so\uparrow$$

$$si \wedge yi \rightarrow so\downarrow$$

$$\begin{aligned}
 x00 \vee x10 &\rightarrow u0\uparrow & x01 \vee x11 &\rightarrow u1\uparrow \\
 \neg x00 \wedge \neg x10 &\rightarrow u0\downarrow & \neg x01 \wedge \neg x11 &\rightarrow u1\downarrow \\
 u0 \wedge yi &\rightarrow _y0\downarrow & u1 \wedge yi &\rightarrow _y1\downarrow \\
 \neg u0 &\rightarrow _y0\uparrow & \neg u1 &\rightarrow _y1\uparrow \\
 \neg _y0 &\rightarrow y0\uparrow & \neg _y1 &\rightarrow y1\uparrow \\
 _y0 &\rightarrow y0\downarrow & _y1 &\rightarrow y1\downarrow \\
 y0 \vee y1 &\rightarrow _vy\downarrow \\
 \neg y0 \wedge \neg y1 &\rightarrow _vy\uparrow \\
 \neg _x0v \wedge \neg yi \wedge \neg _vy &\rightarrow x0a\uparrow & \neg _x1v \wedge \neg yi \wedge \neg _vy &\rightarrow x1a\uparrow \\
 _x0v \wedge _vy &\rightarrow x0a\downarrow & _x1v \wedge _vy &\rightarrow x1a\downarrow
 \end{aligned}$$

D.23 Chain serializer

The chain serializer uses a chain of NODES to sequence the words of an eMx1ofN channel. A C-element indicates the transition between the up and down phases of the sequencing.



1-of-2 approximate scaling:

component	transistors/component	components/serializer	transistors/serializer
NODE	39	$M - 1$	$39(M - 1)$
TAIL	33	1	33
C	8	1	8
approx. transistors/serializer			$39M + 2$

1-of-4 approximate scaling:

component	transistors/component	components/serializer	transistors/serializer
NODE	65	$M - 1$	$65(M - 1)$
TAIL	53	1	53
C	8	1	8
approx. transistors/serializer			$65M - 4$

For the receiver to handle 4096 neurons (and no data) encoded as 1-of-2 or 1-of-4 words, we would need 11 and 5 NODEs, respectively.

1-of-2 accounting:

component	transistors/component	components/serializer	transistors/serializer
NODE	39	11	429
TAIL	33	1	33
C	8	1	8
total transistors/serializer			470

1-of-4 accounting:

component	transistors/component	components/serializer	transistors/serializer
NODE	65	5	325
TAIL	53	5	53
C	8	1	8
total transistors/serializer			386

D.24 Chain serializer NODE

```

*[[ye];
  [¬u →
    [d0 → y0↑; [¬ye]; u↑; y0↓
    [d1 → y1↑; [¬ye]; u↑; y1↓
    ]
  [u → xe↑;
    [x0 → y0↑; [¬ye]; xe↓; [¬x0]; y0↓
    [x1 → y1↑; [¬ye]; xe↓; [¬x1]; y1↓
    [¬d0 ∧ ¬d1 → u↓; [¬ye]; xe↓
    ]
  ]]

```

PRS

$$\begin{aligned}
ye \wedge \neg u \wedge d0 \vee x0 &\rightarrow y0\uparrow & ye \wedge \neg u \wedge d1 \vee x1 &\rightarrow y1\uparrow \\
\neg ye \wedge u \wedge \neg x0 &\rightarrow y0\downarrow & \neg ye \wedge u \wedge \neg x1 &\rightarrow y1\downarrow \\
y0 \vee y1 &\rightarrow vy\uparrow & d0 \vee d1 &\rightarrow vd\uparrow \\
\neg y0 \wedge \neg y1 &\rightarrow vy\downarrow & \neg d0 \wedge \neg d1 &\rightarrow vd\downarrow \\
vy \wedge vd &\rightarrow u\uparrow & & \\
\neg vy \wedge \neg vd &\rightarrow u\downarrow & &
\end{aligned}$$

$$\begin{aligned}
ye \wedge u \wedge \neg vy &\rightarrow xe\uparrow \\
\neg ye \wedge (\neg u \vee vy) &\rightarrow xe\downarrow
\end{aligned}$$

1-of-2 transistor approximate accounting:

rule	transistor count	comments
$y[0, 1]$	20	
u	9	
xe	10	
total	39	

1-of-4 transistor approximate accounting:

rule	transistor count	comments
$y[0, 1]$	40	
u	13	
xe	12	
total	65	

CMOS-implementable PRS

$$\begin{aligned}
\neg u &\rightarrow _u\uparrow \\
u &\rightarrow _u\downarrow
\end{aligned}$$

$$\begin{aligned}
ye \wedge _u \wedge d0 \vee x0 &\rightarrow _y0\downarrow & ye \wedge _u \wedge d1 \vee x1 &\rightarrow _y1\downarrow \\
\neg ye \wedge \neg _u \wedge \neg x0 &\rightarrow _y0\uparrow & \neg ye \wedge \neg _u \wedge \neg x1 &\rightarrow _y1\uparrow
\end{aligned}$$

$$\begin{aligned}
\neg _y0 &\rightarrow _ _y0\uparrow & \neg _y1 &\rightarrow _ _y1\uparrow \\
_y0 &\rightarrow _ _y0\downarrow & _y1 &\rightarrow _ _y1\downarrow
\end{aligned}$$

$$\begin{aligned}
_ _y0 \vee _ _y1 &\rightarrow _vy\downarrow & d0 \vee d1 &\rightarrow _vd\downarrow \\
\neg _ _y0 \wedge \neg _ _y1 &\rightarrow _vy\uparrow & \neg d0 \wedge \neg d1 &\rightarrow _vd\uparrow
\end{aligned}$$

$$\begin{aligned}
\neg _vy \wedge \neg _vd &\rightarrow u\uparrow \\
_vy \wedge _vd &\rightarrow u\downarrow
\end{aligned}$$

$$\begin{aligned}
\neg _u &\rightarrow _ _u\uparrow \\
_u &\rightarrow _ _u\downarrow
\end{aligned}$$

$$\begin{aligned}
ye \wedge _ _u \wedge _vy &\rightarrow _xe\downarrow \\
\neg ye \wedge (\neg _ _u \vee \neg _vy) &\rightarrow _xe\uparrow
\end{aligned}$$

D.25 Chain serializer TAIL

```

*[[ye];
  [¬u →
    [d0 → y0↑; u↑; [¬ye]; y0↓
    [d1 → y1↑; u↑; [¬ye]; y1↓
    ]
  ]
  [u → xe↑; [¬d0 ∧ ¬d1]; u↓; [¬ye]; xe↓
]]

```

PRS

$$\begin{array}{ll}
 ye \wedge \neg u \wedge d0 \rightarrow y0\uparrow & ye \wedge \neg u \wedge d1 \rightarrow y1\uparrow \\
 \neg ye \wedge u \rightarrow y0\downarrow & \neg ye \wedge u \rightarrow y1\downarrow \\
 \\
 y0 \vee y1 \rightarrow vy\uparrow & d0 \vee d1 \rightarrow vd\uparrow \\
 \neg y0 \wedge \neg y1 \rightarrow vy\downarrow & \neg d0 \wedge \neg d1 \rightarrow vd\downarrow \\
 \\
 vy \wedge vd \rightarrow u\uparrow & \\
 \neg vy \wedge \neg vd \rightarrow u\downarrow & \\
 \\
 ye \wedge u \wedge \neg vy \rightarrow xe\uparrow & \\
 \neg ye \wedge (\neg u \vee vy) \rightarrow xe\downarrow &
 \end{array}$$

1-of-2 transistor approximate accounting:

rule	transistor count	comments
$y[0, 1]$	16	
u	9	
xe	8	
total	33	

1-of-4 transistor approximate accounting:

rule	transistor count	comments
$y[0, 1]$	32	
u	13	
xe	8	
total	53	

CMOS-implementable PRS

$$\begin{array}{l}
 \neg u \rightarrow \neg u\uparrow \\
 u \rightarrow \neg u\downarrow
 \end{array}$$

$$\begin{aligned}
ye \wedge _u \wedge d0 &\rightarrow _y0\downarrow & ye \wedge _u \wedge d1 &\rightarrow _y1\downarrow \\
\neg ye \wedge \neg _u &\rightarrow _y0\uparrow & \neg ye \wedge \neg _u &\rightarrow _y1\uparrow \\
\neg _y0 \vee \neg _y1 &\rightarrow vy\uparrow & d0 \vee d1 &\rightarrow _vd\downarrow \\
_y0 \wedge _y1 &\rightarrow vy\downarrow & \neg d0 \wedge \neg d1 &\rightarrow _vd\uparrow \\
\neg _vd &\rightarrow _ _vd\uparrow \\
_vd &\rightarrow _ _vd\downarrow \\
\neg _vy \wedge \neg _vd &\rightarrow u\uparrow \\
_vy \wedge _vd &\rightarrow u\downarrow \\
\neg _u &\rightarrow _ _u\uparrow \\
_u &\rightarrow _ _u\downarrow \\
ye \wedge _u \wedge _vy &\rightarrow xe\uparrow \\
\neg ye \wedge (\neg _u \vee \neg _vy) &\rightarrow xe\downarrow
\end{aligned}$$

D.26 Chain serializer C

The C process in the chain serializer is a C-element and costs 8 transistors.

$$\begin{aligned}
xe \wedge po &\rightarrow s\downarrow \\
\neg xe \wedge \neg po &\rightarrow s\uparrow
\end{aligned}$$

D.27 Serial merge

The receiver is responsible for sending spikes to neurons and data to the neuron and synapse configuration memory. This design uses an arbiter to handle concurrent input so that we can send spikes and write to the configuration memory on the fly.

For M inputs and 1-of-D encoding,

$$\begin{aligned}
&*[[\\
&\quad \langle |m : M : xmi \longrightarrow yo\uparrow; [yi]; xmo\uparrow; [\neg xmi]; yo\downarrow; [\neg yi]; xmo\downarrow \rangle \\
&]] \\
&*[[\\
&\quad \langle \llbracket m : M : \langle \llbracket d : D : xmd \longrightarrow yd\uparrow; xmo\downarrow; [\neg xmd]; yd\downarrow; xmo\downarrow \rangle \rangle \\
&]]
\end{aligned}$$

For M=2 and D=2,

```

*[[x0i → yo↑; [yi]; x0o↑; [¬x0i]; yo↓; [¬yi]; x0o↓
  |x1i → yo↑; [yi]; x1o↑; [¬x1i]; yo↓; [¬yi]; x1o↓
]]

*[[x00 → y0↑; [¬yi]; x0o↓; [¬x00]; y0↓; [yi]; x0o↓
  |x01 → y1↑; [¬yi]; x0o↓; [¬x01]; y1↓; [yi]; x0o↓
  |x10 → y0↑; [¬yi]; x1o↓; [¬x10]; y0↓; [yi]; x1o↓
  |x11 → y1↑; [¬yi]; x1o↓; [¬x11]; y1↓; [yi]; x1o↓
]]

```

PRS

The parent requests and grants are handled by the standard n-way arbiter with its parent ports exposed. Otherwise,

$$\begin{array}{ll}
 x00 \vee x10 & \rightarrow y0\uparrow & x01 \vee x11 & \rightarrow y1\uparrow \\
 \neg x00 \wedge \neg x10 & \rightarrow y0\downarrow & \neg x01 \wedge \neg x11 & \rightarrow y1\downarrow
 \end{array}$$

CMOS-implementable PRS:

$$\begin{array}{ll}
 x00 \vee x10 & \rightarrow _y0\uparrow & x01 \vee x11 & \rightarrow _y1\uparrow \\
 \neg x00 \wedge \neg x10 & \rightarrow _y0\downarrow & \neg x01 \wedge \neg x11 & \rightarrow _y1\downarrow \\
 \\
 \neg_y0 & \rightarrow __y0\downarrow & \neg_y1 & \rightarrow __y1\downarrow \\
 _y0 & \rightarrow __y0\uparrow & _y1 & \rightarrow __y1\uparrow
 \end{array}$$

D.28 Memory

Each group of 4 neurons and 1 synapse needs at least 28 bits of memory. However, given the shape and size constraints of the memory, we may end up using larger memories. This memory will only be written, not read.

A memory consists of a two dimensional array of bitcells. Rows are addressed by a read/write lines, and columns are addressed by the data itself. That is, we write an entire row at a time.

The shape of the memory dictates the size of the input we deliver to it. For each write operation, we indicate the address (which write line) and data to write. The address is encoded in 1-of-2 or 1-of-4 words. The data is communicated as it will be written. It is cheaper to have write lines than data lines because the size of the encoded address scales with the logarithm of the number of write lines. So we usually maximize the number of write lines within the aspect ratio constraint of the neuron layout. Further, using 1-of-4 encoding requires that we use at least 2 data bits.

Here are some memory sizes and their required deserializer bits:

neurons/synapses	memory bits	write lines	write bits	data bits	total bits	word size	words
4/1	32	32	5	1	6	2	6
16/4	128	64	6	2	8	4	4
16/4	128	32	5	4	9	4	4.5
4096/1024	32768	16384	14	2	16	4	8

We can reduce the number of words if the memory supports banking, a level of address heirarchy on top of the address itself.

Appendix E

Serial Router Supplement

This appendix describes some components that were not used in the serial H-Tree router described in Section 3.4

E.1 Half-Cycle Slack Buffer

The serial router as laid out did not contain any pipelining. As a result, its throughput was throttled by additional unpipelined communications that extended well into the datapath. For reference, here is a buffer process that provides a half-cycle of slack. Buffers could be placed in the cutouts of the router and between the router and datapath to improve the router's throughput.

HSE

```
*[[ $x_\phi \rightarrow y_\phi \uparrow; x_e \uparrow, [y_e]$ 
   $[x_0 \rightarrow y_0 \uparrow; x_e \downarrow; [\neg y_e \wedge \neg x_0]; y_0 \downarrow; x_e \uparrow; [y_e]$ 
   $[x_1 \rightarrow y_1 \uparrow; x_e \downarrow; [\neg y_e \wedge \neg x_1]; y_1 \downarrow; x_e \uparrow; [y_e]$ 
   $[\neg x_\phi \rightarrow y_\phi \downarrow; x_e \downarrow; [\neg y_e]$ 
  ]]
```

PRS

```
 $x_\phi \wedge \neg y_e \rightarrow y_\phi \uparrow$ 
 $\neg x_\phi \wedge y_e \rightarrow y_\phi \downarrow$ 

 $y_\phi \wedge \neg y_0 \wedge \neg y_1 \rightarrow x_e \uparrow$ 
 $\neg y_\phi \vee y_0 \vee y_1 \rightarrow x_e \downarrow$ 

 $x_0 \wedge y_e \rightarrow y_0 \uparrow$ 
 $\neg x_0 \wedge \neg y_e \rightarrow y_0 \downarrow$ 

 $x_1 \wedge y_e \rightarrow y_1 \uparrow$ 
 $\neg x_1 \wedge \neg y_e \rightarrow y_1 \downarrow$ 
```

CMOS-Implementable PRS

$$\begin{array}{ll}
\neg x_\phi \rightarrow \neg x_\phi \uparrow & \neg \neg x_\phi \wedge \neg y_e \rightarrow y_\phi \uparrow \\
x_\phi \rightarrow \neg x_\phi \downarrow & \neg x_\phi \wedge y_e \rightarrow y_\phi \downarrow \\
y_\phi \wedge \neg y_0 \wedge \neg y_1 \rightarrow \neg x_e \downarrow & \neg \neg x_e \rightarrow \neg \neg x_e \uparrow \\
\neg y_\phi \vee \neg \neg y_0 \vee \neg \neg y_1 \rightarrow \neg x_e \uparrow & \neg x_e \rightarrow \neg \neg x_e \downarrow \\
x_0 \wedge y_e \rightarrow \neg y_0 \downarrow & \neg \neg y_0 \rightarrow \neg \neg y_0 \uparrow \\
\neg x_0 \wedge \neg y_e \rightarrow \neg y_0 \uparrow & \neg y_0 \rightarrow \neg \neg y_0 \downarrow \\
x_1 \wedge y_e \rightarrow \neg y_1 \downarrow & \neg \neg y_1 \rightarrow \neg \neg y_1 \uparrow \\
\neg x_1 \wedge \neg y_e \rightarrow \neg y_1 \uparrow & \neg y_1 \rightarrow \neg \neg y_1 \downarrow
\end{array}$$

4-ary Accounting

rule	transistor count	comments
$\neg x_\phi$	2	
y_ϕ	8	
$\neg x_e$	10	
$\neg \neg x_e$	2	
$\neg y[0, 1, 2, 3]$	16	staticized by $\neg \neg y[0, 1, 2, 3]$
$\neg \neg y[0, 1, 2, 3]$	16	staticizes $\neg y[0, 1, 2, 3]$
total	54	

E.2 Greedy but Fair N-way Arbiter using Tree/Ring Sequencing

Although the serial router used unpipelined 4-way arbiters for minimal transistor counts, a designer may want an N-way arbiter with pipelining. This appendix presents a pipelined, N-way arbiter which unifies the ring and tree approaches by parameterizing a tree via K , its radix. When $K = 1$, the arbiter is a flat ring, and when $K = 2$, the arbiter is a binary tree. The tree is built from NODES. The leaves of the tree service requests from N clients.

Accounting

NODES cost $30K+18$ transistors, where K is the radix of the tree.

Configuration	transistors
K -ary tree	$(30K + 18) \frac{N-1}{K-1}$
Flat ring	$30N + 18$
Binary tree	$78(N - 1)$
4-ary tree	$46(N - 1)$

E.3 NODE

For a binary tree,

$$*[[\overline{C_0} \vee \overline{C_1}]; P \bullet (\\ \overline{C_0} \rightarrow C_0; C_0 | \neg \overline{C_0} \rightarrow skip]; \\ \overline{C_1} \rightarrow C_1; C_1 | \neg \overline{C_1} \rightarrow skip) \\]]$$

For a K -ary tree,

$$*[[\langle \vee k : K : \overline{C_k} \rangle]; \\ P \bullet \langle k : K : [\overline{C_0} \rightarrow C_0; C_0 | \neg \overline{C_0} \rightarrow skip] \rangle \\]]$$

We decompose NODE into PRODUCER and CONSUMER processes.

Accounting

component	transistors/component	components/NODE	transistors/NODE
PRODUCER	$18+K$	1	$18+K$
CONSUMER	29	K	$29K$
total transistors/NODE			$30K+18$

E.3.1 PRODUCER

PRODUCER produces parent requests upon sensing child requests and sequences between CONSUMERS. For a binary tree,

$$*[[\overline{C_0} \vee \overline{C_1}]; P \bullet \langle \bullet K : k : S_k \rangle]]$$

For a K -ary tree,

$$*[[\langle \vee K : k : \overline{C_k} \rangle]; P \bullet \langle \bullet K : k : S_k \rangle]]$$

We decompose PRODUCER into CTRL and SEQ.

Accounting

component	transistors/component	components/PRODUCER	transistors/PRODUCER
CTRL	$18+K$	1	$18+K$
SEQ	0	1	0
total transistors/PRODUCER			$18+K$

CTRL

CTRL relays the child requests to the parent and initiates the sequencing between CONSUMERS.

CHP

For a binary tree,

$$*[[\overline{C_0} \vee \overline{C_1}]; P \bullet S]$$

For a K -ary tree,

$$*[[\langle \vee K : k : \overline{C_k} \rangle]; P \bullet S]$$

HSE

For a binary tree,

$$\begin{aligned} &*[[pe]; [c0r \vee c1r]; x\downarrow; \\ &\quad pr\uparrow; [\neg pe]; \\ &\quad sr\uparrow; [sa]; x\uparrow; sr\downarrow; [\neg sa]; \\ &\quad pr\downarrow] \end{aligned}$$

PRS

$$\begin{aligned} pe \wedge (c0r \vee c1r) &\rightarrow x\downarrow \\ sa &\rightarrow x\uparrow \end{aligned}$$

$$\begin{aligned} \neg x \vee sa &\rightarrow pr\uparrow \\ x \wedge \neg sa &\rightarrow pr\downarrow \end{aligned}$$

$$\begin{aligned} \neg x \wedge \neg pe &\rightarrow sr\uparrow \\ x \vee pe &\rightarrow sr\downarrow \end{aligned}$$

CMOS-implementable PRS

$$\begin{aligned} _pe \wedge (c0r \vee c1r) &\rightarrow x\downarrow \\ \neg_sa &\rightarrow x\uparrow \end{aligned}$$

$$\begin{aligned} \neg x \vee \neg_sa &\rightarrow pr\uparrow \\ x \wedge _sa &\rightarrow pr\downarrow \end{aligned}$$

$$\begin{aligned} \neg x \wedge \neg_pe &\rightarrow sr\uparrow \\ x \vee _pe &\rightarrow sr\downarrow \end{aligned}$$

$$\begin{aligned} pe &\rightarrow _pe\downarrow & sa &\rightarrow _sa\downarrow \\ \neg pe &\rightarrow _pe\uparrow & \neg sa &\rightarrow _sa\uparrow \end{aligned}$$

Accounting

rule	transistor count	comments
x	$2+K+4$	
pr	4	
sr	4	
pe	2	
sa	2	
total	$18+K$	

SEQ

SEQ sequences through the consumers.

CHP

In general,

$$SEQ(K) \equiv * [SP \bullet \langle \bullet k : K : SC_k \rangle]$$

For a binary tree,

$$* [SP \bullet SC_0 \bullet SC_1]$$

We build a K -way sequencer as K' -way sequencer tree defined recursively as

$$SEQ(K) \equiv SEQ(K') \parallel \langle \parallel k' : K' - 1 : SEQ(K/K') \rangle \parallel SEQ(K - (K' - 1)K/K')$$

where the recursion ends when $K \leq K'$. For example, a binary tree would be defined recursively as

$$SEQ(K) = SEQ(2) \parallel SEQ(K/2) \parallel SEQ(K - K/2)$$

The SC channels of the parent sequencer are the SP channels of the the child sequencers. The SC channels of the leaf sequencers connect to the consumers.

HSE

For a binary tree,

$$\begin{aligned}
& * [[spr] ; \\
& \quad sc_0 r \uparrow ; [sc_0 a] ; sc_1 r \uparrow ; [sc_1 a] ; \\
& \quad spa \uparrow ; [\neg spr] ; \\
& \quad sc_0 r \downarrow ; [\neg sc_0 a] ; sc_1 r \downarrow ; [\neg sc_1 a] ; \\
& \quad spa \downarrow \\
&]
\end{aligned}$$

In general,

$$\begin{aligned}
 & * [[spr]; \\
 & \quad \langle k' : K' : sc'_k r \uparrow; [sc'_k a] \rangle \\
 & \quad spa \uparrow; [\neg spr]; \\
 & \quad \langle k' : K' : sc'_k r \downarrow; [\neg sc'_k a] \rangle \\
 & \quad spa \downarrow \\
 &]
 \end{aligned}$$

PRS

$$\begin{aligned}
 spr & \rightarrow sc0r \uparrow & sc(k' - 1)r & \rightarrow spa \uparrow \\
 \neg spr & \rightarrow sc0r \downarrow & \neg sc(k' - 1)r & \rightarrow spa \downarrow \\
 \\
 sck'a & \rightarrow sc(k' + 1)r \uparrow \\
 \neg sck'a & \rightarrow sc(k' + 1)r \downarrow
 \end{aligned}$$

These are just wires, so all K' -way tree sequencers that implement a K -way sequencer are equivalent.

E.3.2 CONSUMER

CONSUMER checks and services requests from children.

CHP

$$\begin{aligned}
 & * [[\bar{C} \rightarrow S; C \\
 & \quad | \bar{S} \rightarrow S]]
 \end{aligned}$$

which is a variation on the precise exceptions circuit of [28].

HSE

$$\begin{aligned}
 & * [[cr \rightarrow [sr]; sa \uparrow; [\neg sr]; ce \downarrow; [\neg cr]; sa \downarrow; ce \uparrow \\
 & \quad | sr \rightarrow sa \uparrow; [\neg sr]; sa \downarrow \\
 &]]
 \end{aligned}$$

We break out nondeterministic selection.

$$\begin{aligned}
 & * [[cr \rightarrow c \uparrow; [\neg cr]; c \downarrow \\
 & \quad | sa \rightarrow s \uparrow; [\neg sa]; s \downarrow \\
 &]] \\
 & * [[c \rightarrow [sr]; sta \uparrow; [\neg sr]; ce \downarrow; [\neg c]; sta \downarrow; ce \uparrow \\
 & \quad | s \rightarrow sfa \uparrow; [\neg s]; sfa \downarrow \\
 &]] \\
 & * [[sta \vee sfa]; sa \uparrow; [\neg sta \wedge \neg sfa]; sa \downarrow]
 \end{aligned}$$

PRS

$$\begin{array}{ll} c \wedge sr \rightarrow sta\uparrow & s \rightarrow sfa\uparrow \\ \neg c \rightarrow sta\downarrow & \neg s \rightarrow sfa\downarrow \end{array}$$

$$\begin{array}{l} \neg sta \vee sr \rightarrow ce\uparrow \\ sta \wedge \neg sr \rightarrow ce\downarrow \end{array}$$

$$\begin{array}{l} sta \vee sfa \rightarrow sa\uparrow \\ \neg sta \wedge \neg sfa \rightarrow sa\downarrow \end{array}$$

CMOS-implementable PRS

$$\begin{array}{ll} c \wedge sr \rightarrow _sta\downarrow & s \rightarrow _sfa\downarrow \\ \neg c \rightarrow _sta\uparrow & \neg s \rightarrow _sfa\uparrow \end{array}$$

$$\begin{array}{l} _sta \vee sr \rightarrow _ce\downarrow \\ \neg _sta \wedge \neg sr \rightarrow _ce\uparrow \end{array}$$

$$\begin{array}{l} \neg _sta \vee \neg _sfa \rightarrow sa\uparrow \\ _sta \wedge _sfa \rightarrow sa\downarrow \end{array}$$

Accounting

rule	transistor count	comments
$[c, s]$	12	2-way active-low arbiter
$_sta$	7	
$_sfa$	2	
$_ce$	4	
sa	4	
total	29	

Bibliography

- [1] Filipp Akopyan, Carlos Tadeo, Ortega Otero, and Rajit Manohar. Hybrid synchronous-asynchronous tool flow for emerging VLSI design. In *IEEE International Workshop on Logic Synthesis*, 2016.
- [2] Ben Varkey Benjamin, Peiran Gao, Emmett McQuinn, Swadesh Choudhary, Anand R. Chandrasekaran, Jean Marie Bussat, Rodrigo Alvarez-Icaza, John V. Arthur, Paul A. Merolla, and Kwabena Boahen. Neurogrid: A mixed-analog-digital multichip system for large-scale neural simulations. *Proceedings of the IEEE*, 102(5):699–716, 2014.
- [3] Kwabena Boahen. A neuromorph’s prospectus. *Computing in Science and Engineering*, 19(2):14–28, 2017.
- [4] Kwabena A Boahen. Point-to-point connectivity between neuromorphic chips using address-events. *IEEE Transactions on Circuits & Systems II: Analog and Digital Signal Processing*, 47(5):416–434, 1999.
- [5] Kwabena A. Boahen. A burst-mode word-serial address-event link-I: transmitter design. *IEEE Transactions on Circuits and Systems I: Regular Papers*, 51(7):1269–1280, 2004.
- [6] Kwabena A Boahen. A burst-mode word-serial address-event link-II: receiver design. *IEEE Transactions on Circuits and Systems I: Regular Papers*, 51(7), 2004.
- [7] Kwabena A. Boahen. A burst-mode word-serial address-event link-III: testing and results. *IEEE Transactions on Circuits and Systems I: Regular Papers*, 51(7):1292–1300, 2004.
- [8] Martin Boerlin, Christian K. Machens, and Sophie Denève. Predictive coding of dynamical variables in balanced spiking networks. *PLoS Computational Biology*, 9:1–16, 2013.
- [9] Alexander Borst and Frdric Theunissen. Information theory and neural coding. *Nature Neuroscience*, 2(11):947–957, 1999.
- [10] Christian Brandli, Raphael Berner, Minhao Yang, Shih Chii Liu, and Tobi Delbruck. A 240 180 130 dB 3 μ s latency global shutter spatiotemporal vision sensor. *IEEE Journal of Solid-State Circuits*, 49(10):2333–2341, 2014.

- [11] T. Y W Choi, Paul A. Merolla, John V. Arthur, Kwabena A. Boahen, and Bertram E. Shi. Neuromorphic implementation of orientation hypercolumns. *IEEE Transactions on Circuits and Systems I: Regular Papers*, 52-I:1049–1060, 2005.
- [12] Swadesh Choudhary, Steven Sloan, Sam Fok, Alexander Neckar, Eric Trautmann, Peiran Gao, Terry Stewart, Chris Eliasmith, and Kwabena Boahen. Silicon neurons that compute. In *International Conference on Artificial Neural Networks*, pages 121–128, 2012.
- [13] Matthieu Courbariaux, Yoshua Bengio, and Jean-Pierre David. Training deep neural networks with low precision multiplications. *ArXiv e-prints*, 2014.
- [14] Hideyuki Cteau and Alex D. Reyes. Relation between single neuron and population spiking statistics and effects on network activity. *Physical Review Letters*, 2006.
- [15] Chris Eliasmith and Charles H Anderson. *Neural engineering: computation, representation, and dynamics in neurobiological systems*. MIT press, 2004.
- [16] William Ft Softky and Christof Koch. The highly irregular firing of cortical cells is inconsistent with temporal integration of random EPSPs. *The Journal of Neuroscience*, 13(1):334–350, 1993.
- [17] Julius Georgiou, Andreas G Andreou, and Philippe 0 Pouliquen. A mixed analog/digital asynchronous processor for cortical computations in 3D SQL-CMOS. In *2006 IEEE International Symposium on Circuits and Systems*, pages 4 pp.–, 2006.
- [18] Suzana Herculano-Houzel. The human brain in numbers: a linearly scaled-up primate brain. *Frontiers in Human Neuroscience*, 2009.
- [19] Itay Hubara, Matthieu Courbariaux, Daniel Soudry, Ran El-Yaniv, and Yoshua Bengio. Quantized neural networks: training neural networks with low precision weights and activations. *CoRR*, 2016.
- [20] Eric Jonas and Konrad Paul Kording. Could a neuroscientist understand a microprocessor? *PLoS Computational Biology*, 13(1), 2017.
- [21] Norman P. Jouppi, Al Borchers, Rick Boyle, Pierre-luc Cantin, Clifford Chao, Chris Clark, Jeremy Coriell, Mike Daley, Matt Dau, Jeffrey Dean, Ben Gelb, Cliff Young, Tara Vazir Ghaemmaghami, Rajendra Gottipati, William Gulland, Robert Hagmann, C. Richard Ho, Doug Hogberg, John Hu, Robert Hundt, Dan Hurt, Julian Ibarz, Nishant Patil, Aaron Jaffey, Alek Jaworski, Alexander Kaplan, Harshit Khaitan, Daniel Killebrew, Andy Koch, Naveen Kumar, Steve Lacy, James Laudon, James Law, David Patterson, Diemthu Le, Chris Leary, Zhuyuan Liu, Kyle Lucke, Alan Lundin, Gordon MacKean, Adriana Maggiore, Maire Mahony, Kieran Miller, Rahul Nagarajan, Gaurav Agrawal, Ravi Narayanaswami, Ray Ni, Kathy Nix, Thomas Norrie, Mark Omernick, Narayana Penukonda, Andy Phelps, Jonathan Ross, Matt Ross, Amir

- Salek, Raminder Bajwa, Emad Samadiani, Chris Severn, Gregory Sizikov, Matthew Snelham, Jed Souter, Dan Steinberg, Andy Swing, Mercedes Tan, Gregory Thorson, Bo Tian, Sarah Bates, Horia Toma, Erick Tuttle, Vijay Vasudevan, Richard Walter, Walter Wang, Eric Wilcox, Doe Hyun Yoon, Suresh Bhatia, and Nan Boden. In-datacenter performance analysis of a Tensor Processing Unit. *ACM SIGARCH Computer Architecture News*, 45(2):1–12, 2017.
- [22] Robert Karmazin, Carlos Tadeo Ortega Otero, and Rajit Manohar. cellTK: automated layout for asynchronous circuits with nonstandard cells. In *Proceedings - International Symposium on Asynchronous Circuits and Systems*, 2013.
- [23] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. ImageNet Classification with Deep Convolutional Neural Networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012.
- [24] John Lazzaro, John Wawrzynek, M. Mahowald, Massimo Sivilotti, and Dave Gillespie. Silicon auditory processors as computer peripherals. *IEEE Transactions on Neural Networks*, 4(3):523–528, 1993.
- [25] Joseph Lin and Kwabena Boahen. A delay-insensitive address-event link. In *IEEE Symposium on Asynchronous Circuits and Systems*, 2009.
- [26] Benjamin Lindner. Superposition of many independent spike trains is generally not a Poisson process. *Physical Review E - Statistical, Nonlinear, and Soft Matter Physics*, 2006.
- [27] M Mahowald. *An analog VLSI stereoscopic vision system*, volume 1. Springer, Boston, 1994.
- [28] R. Manohar, M. Nyström, and A. J. Martin. Precise exceptions in asynchronous processors. In *Proceedings - 2001 Conference on Advanced Research in VLSI, ARVLSI 2001*, 2001.
- [29] Rajit Manohar. Comparing stochastic and deterministic computing. *IEEE Computer Architecture Letters*, 14(2):119–122, 2015.
- [30] Rajit Manohar and Yoram Moses. The eventual C-element theorem for delay-insensitive asynchronous circuits. In *2017 23rd IEEE International Symposium on Asynchronous Circuits and Systems (ASYNC)*, pages 102–109, 2017.
- [31] Alain Martin. The limitations to delay-insensitivity in asynchronous circuits. In *Proceedings of the 6th MIT Conference on Advanced Research in VLSI*. MIT Press, 1990.
- [32] Alain J. Martin and Mika Nyström. Asynchronous techniques for system-on-chip design. *Proceedings of the IEEE*, 94(6):1089–1120, 2006.
- [33] Carver Mead. *Analog VLSI and neural systems*. Addison-Wesley, 1989.

- [34] Carver Mead. Neuromorphic electronic systems. *Proceedings of the IEEE*, 78(10):1629–1636, 1990.
- [35] Samir Menon, Sam Fok, Alex Neckar, Oussama Khatib, and Kwabena Boahen. Controlling articulated robots in task-space with spiking silicon neurons. In *5th IEEE RAS/EMBS International Conference on Biomedical Robotics and Biomechatronics*, pages 181–186, 2014.
- [36] Paul Merolla, John Arthur, Rodrigo Alvarez, Jean Marie Bussat, and Kwabena Boahen. A multicast tree router for multichip neuromorphic systems. *IEEE Transactions on Circuits and Systems I: Regular Papers*, 61(3):820–833, 2014.
- [37] Paul A Merolla, John V Arthur, Rodrigo Alvarez-Icaza, Andrew S Cassidy, Jun Sawada, Filipp Akopyan, Bryan L Jackson, Nabil Imam, Chen Guo, Yutaka Nakamura, Bernard Brezzo, Ivan Vo, Steven K Esser, Rathinakumar Appuswamy, Brian Taba, Arnon Amir, Myron D Flickner, William P Risk, Rajit Manohar, and Dharmendra S Modha. A million spiking-neuron integrated circuit with a scalable communication network and interface. *Science*, 345(6197):668–673, 2014.
- [38] Paul A. Merolla, John V. Arthur, Bertram E. Shi, and Kwabena A. Boahen. Expandable networks for neuromorphic chips. *IEEE Transactions on Circuits and Systems I: Regular Papers*, 54(2):301–311, 2007.
- [39] Jongkil Park, Theodore Yu, Siddharth Joshi, Christoph Maier, and Gert Cauwenberghs. Hierarchical address event routing for reconfigurable large-scale neuromorphic systems. *IEEE Transactions on Neural Networks and Learning Systems*, 28(10):2408–2422, 2017.
- [40] Sheldon Ross. *Stochastic Processes*. John Wiley & Sons, 2nd edition, 1996.
- [41] Massimo Antonio Sivilotti. *Wiring considerations in analog VLSI systems, with application to field-programmable networks*. PhD thesis, California Institute of Technology, Pasadena, CA, USA, 1991.
- [42] Bongki Son, Yunjae Suh, Sungho Kim, Heejae Jung, Jun Seok Kim, Changwoo Shin, Keunju Park, Kyoobin Lee, Jinman Park, Jooyeon Woo, Yohan Roh, Hyunku Lee, Yibing Wang, Ilya Ovsianikov, and Hyunsurk Ryu. A 640480 dynamic vision sensor with a 9m pixel and 300Meps address-event representation. In *Digest of Technical Papers - IEEE International Solid-State Circuits Conference*, pages 66–67, 2017.
- [43] Dominik Thalmeier, Marvin Uhlmann, Hilbert J. Kappen, and Raoul Martin Memmesheimer. Learning universal computations with spikes. *PLoS Computational Biology*, 12:1–29, 2016.
- [44] John Von Neumann. *The Computer and the Brain*. Yale University Press, New Haven, 3rd edition, 2012.