

1 The Computational Cost of Asynchronous Neural 2 Communication

3 **Yael Hitron**

4 Weizmann Institute of Science, Israel
5 yael.hitron@weizmann.ac.il

6 **Merav Parter**

7 Weizmann Institute of Science, Israel
8 merav.parter@weizmann.ac.il

9 **Gur Perri**

10 Weizmann Institute of Science, Israel
11 gur.perri@weizmann.ac.il

12 — Abstract —

13 Biological neural computation is inherently asynchronous due to large variations in neuronal spike
14 timing and transmission delays. So-far, most theoretical work on neural networks assumes the
15 *synchronous* setting where neurons fire simultaneously in discrete rounds. In this work we aim at
16 understanding the barriers of asynchronous neural computation from an algorithmic perspective. We
17 consider an extension of the widely studied model of synchronized spiking neurons [Maass, Neural
18 Networks 97] to the asynchronous setting by taking into account edge and node delays.

19 ■ **Edge Delays:** We define an asynchronous model for spiking neurons in which the latency values
20 (i.e., transmission delays) of non self-loop edges *vary* adversarially over time. This extends the
21 recent work of [Hitron and Parter, ESA'19] in which the latency values are restricted to be fixed
22 over time. Our first contribution is an impossibility result that implies that the assumption that
23 self-loop edges have no delays (as assumed in Hitron and Parter) is indeed necessary. Interestingly,
24 in real biological networks self-loop edges (a.k.a. autapse) are indeed free of delays, and the
25 latter has been noted by neuroscientists to be crucial for network synchronization.

26 To capture the computational challenges in this setting, we first consider the implementation of
27 a single NOT gate. This simple function already captures the fundamental difficulties in the
28 asynchronous setting. Our key technical results are space and time upper and lower bounds
29 for the NOT function, our time bounds are *tight*. In the spirit of the distributed synchronizers
30 [Awerbuch and Peleg, FOCS'90] and following [Hitron and Parter, ESA'19], we then provide a
31 general synchronizer machinery. Our construction is very modular and it is based on efficient
32 circuit implementation of threshold gates. The complexity of our scheme is measured by the
33 overhead in the number of neurons and the computation time, both are shown to be polynomial
34 in the largest latency value, and the largest incoming degree Δ of the original network.

35 ■ **Node Delays:** We introduce the study of asynchronous communication due to variations in
36 the response rates of the neurons in the network. In real brain networks, the *round duration*
37 varies between different neurons in the network. Our key result is a simulation methodology
38 that allows one to transform the above mentioned synchronized solution under edge delays into
39 a synchronized under node delays while incurring a small overhead w.r.t space and time.

40 **2012 ACM Subject Classification** Theory of computation → Design and analysis of algorithms

41 **Keywords and phrases** asynchronous communication, asynchronous computation, spiking neurons,
42 synchronizers

43 **Digital Object Identifier** 10.4230/LIPIcs.CVIT.2016.23

44 **Acknowledgements** We are thankful to Roei Tell and Gil Cohen for helpful discussions on Boolean
45 Circuits. We also thank Yoram Moses for referring us to related work on asynchronous digital
46 circuits and discussing the connections between the two settings.



© Y.Hitron, M.paerter, G.Perri;
licensed under Creative Commons License CC-BY
42nd Conference on Very Important Topics (CVIT 2016).

Editors: John Q. Open and Joan R. Access; Article No. 23; pp. 23:1–23:47



Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

47 **1 Introduction**

48 Understanding how the brain works, as a computational device, is a central challenge of
 49 modern neuroscience, artificial intelligence, and lately also in theoretical computer science
 50 and distributed computing [18, 19, 20, 17, 16, 32, 6, 39, 37]. This line of work usually
 51 assumes a simple synchronized model [23, 24] in which neurons fire simultaneously in discrete
 52 rounds in response to their neighboring neurons that fired in the previous round. This
 53 model, while being very convenient for algorithm design, does not take into account the
 54 inherent *asynchronous* nature of neural communication. In the neuroscience literature it
 55 has been noted that the asynchronous nature of these networks mostly stems from two
 56 independent sources [29]: *edge delays* (known as response latency¹.) [35, 5] and *node delays*
 57 (known as refractory period) [34, 3]. In this paper, we aim at understanding the compu-
 58 tational cost incurred by such asynchronous communication. The overhead is measured
 59 by the overhead in the number of neurons and computation time required to compute
 60 a certain function. We believe that understanding the computational power, limitations
 61 and the connections between these models go beyond the setting of spiking neurons, and
 62 might also be relevant for the theory of digital logic design and circuit computation in general.

63
 64 **The Standard Synchronous Model [23, 24]:** Before describing our asynchronous models,
 65 we first revise the standard synchronous model formally defined by Maass. In this model, the
 66 network evolves in discrete, synchronous *rounds* as a Markov chain where each neuron u in
 67 the network is a probabilistic threshold gate with a threshold (or bias) value $b(u)$. In every
 68 round t , the firing probability of neuron u only depends on the firing status of its incoming
 69 neighbors in the preceding round $t - 1$. Formally, the potential $\text{pot}_t(u)$ of neuron u in round
 70 t is defined by the weighted sum over its incoming neighbors that fired in round $t - 1$. The
 71 neuron u fires in round t with probability that depends on the quantity $\text{pot}_t(u) - b(u)$.

72 **1.1 Asynchronous Computation with Bounded Edge Delays**

73 We define an asynchronous model with edge delays bounded² by some given integer L . The
 74 dynamic of the network \mathcal{N} is specified by a latency function $\ell : V \times V \times \mathbb{N} \rightarrow \mathbb{N}_{\leq L}$ interpreted
 75 as follows: For every neuron u firing in round τ , its spike reaches its outgoing neighbor v
 76 within $\ell(u, v, \tau)$ rounds where $\ell(u, v, \tau) \in [1, L]$ might be chosen adversarially for every $u \neq v$,
 77 and every round τ . The network solution \mathcal{N} should then output the desired solution for *any*
 78 adversarial choice of the latency function ℓ . Setting $L = 1$ yields the standard synchronous
 79 model.

80 Asynchronous computation with edge delays was recently introduced by Hitron and
 81 Parter [12]. Their model is similar to ours only that in their model, the edge latencies are
 82 required to be *fixed* over time, whereas in our model the adversary is allowed to change it
 83 from round to round. The model of Hitron and Parter includes an additional restriction
 84 on the adversary (that sets the latency values) by requiring that self-spikes, i.e., of the
 85 form $\langle u, u, \tau \rangle$ have no latency and arrive within a single round to their destination. This
 86 assumption is justified in [12] by the experimental evidence that self-spikes in brain networks
 87 have almost no delays [14]. It is commonly believed in the neuroscience community that this
 88 no-delay property of self-edges is in fact essential for network synchronization [33, 21, 40, 8].

¹ Throughout, we use the terms edge-delay and latency interchangeably.

² This bound is crucial as will be later implied by our lower bound results that depend on L . E.g., both the computation time and the size of the network in this model *must* depend on L .

89 In this work, we provide a theoretical support for this hypothesis by showing that without
 90 such an assumption, one cannot even implement a single AND gate in this model. This
 91 impossibility result holds already in a setting where $L = 2$, the edge latencies are fixed over
 92 time (as assumed in [12]), and the computation time and network size are allowed to be
 93 unbounded. For this reason, we will mostly consider in this paper *nice* latency functions in
 94 which all self-spikes have a latency value of one round.

95 1.2 Asynchronous Computation with Bounded Node Delays

96 We next turn to consider an alternative source for asynchronous communication due to
 97 variations in the response timing of the individual neurons in the network. In real brain
 98 networks, for every neuron u there is a predefined time interval between two consecutive
 99 spikes of u . The length of this interval, which we call *round*, varies considerably among
 100 different neurons in the network [34, 3]. This poses the challenge of creating a synchronized
 101 response at the network level. To account this behavior, we consider a model in which
 102 network's evolution proceeds in *seconds*. A second in this context is simply the smallest
 103 measurable time unit. For a given integer $T \geq 1$, the dynamics is specified by a node-
 104 delay function $t : V \rightarrow \mathbb{N}_{\leq T}$ interpreted as follows: the round duration of each neuron
 105 v consists of $t(v)$ seconds. Specifically, the i^{th} round of v is defined by the time interval
 106 $R_i(v) = [(i-1)t(v) + 1, i \cdot t(v)]$ for every $i \geq 1$. The neuron u fires in the second $i \cdot t(v)$ (i.e.,
 107 at the end of its i^{th} round) only if the total potential due to spikes arriving in the interval
 108 $R_i(v)$ is sufficiently large. The network solution \mathcal{N} should then output the desired solution
 109 for *any* adversarial choice of the node-delay function t . The input parameter T sets a bound
 110 on the differences between the round duration over all neurons in the network. Setting $T = 1$
 111 yields the standard synchronous model.

112 Observe that the edge and node delay models do not imply one another. In the edge
 113 latency model, even though the spikes arrive in adversarially chosen rounds, all neurons in
 114 round τ still depend only on the spikes arriving in round τ . Thus, the duration of a round for
 115 all the neurons in the network is the same: a single tick (or a second) of the global clock. In
 116 contrast, in the node delay model, the adversary selects the round duration for each neuron
 117 which has two physical interpretations. First, it determines the time duration over which
 118 the potential due to arriving spikes is accumulated. In addition, it also determines the time
 119 interval between two consecutive spikes by the given neuron. In one of our most technically
 120 involved results, we show a non-trivial reduction between the edge delay and the node delay
 121 models, provided that the input network satisfies certain properties.

122 Finally, we note that this model has several variations which are in fact supported by
 123 our simulation results (from edge delays to node delays). In particular, one can consider
 124 a more elaborated setting in which the duration of each round per node varies in an
 125 adversarial manner *over time*, i.e., the node-delay function in such a model is of the form
 126 $t : V \times \mathbb{N}_{\geq 1} \rightarrow \mathbb{N}_{\leq T}$ interpreted as follows: the i^{th} round duration on each neuron v consists
 127 of $t(v, i)$ seconds. For example, in such a model the i^{th} round of node u can consists of 2
 128 seconds while its $(i+1)^{\text{th}}$ round might consist of 100 seconds. In another variation, both edge
 129 and node delays are combined and the dynamic is specified by an L -bounded edge latency
 130 function $\ell : V \times V \times \mathbb{N}_{\geq 1} \rightarrow \mathbb{N}_{\leq L}$ and a T -bounded node-delay function $t : V \times \mathbb{N}_{\geq 1} \rightarrow \mathbb{N}_{\leq T}$.

131 1.3 Synchronizers

132 In the spirit of Awerbuch and Peleg's synchronizers for distributed networks [2] and the recent
 133 work of [12], our primary goal with respect to upper bound results is to provide a general

134 simulation methodology that takes any n -neuron network \mathcal{N} that *solves* the problem in the
 135 standard synchronized setting (i.e., in which all spikes arrive within a single round) and
 136 transforms it into an “analogous” network $\text{Sync}(\mathcal{N})$ in the edge delay setting while incurring
 137 a small overhead in the number of neurons and the computation time (w.r.t the base network
 138 \mathcal{N}).

139 For the setting in which the edge latencies are fixed over time and bounded by some integer
 140 L , Hitron and Parter [12] showed such a simulation using their efficient constructions for
 141 neural timers and counters. Their synchronized network solution $\text{Sync}(\mathcal{N})$ has $O(n + L \log L)$
 142 neurons and $O(rL^3)$ rounds where r is the computation time of the base network \mathcal{N} . While
 143 being quite efficient in terms of space and time, the synchronizers of [12] are heavily built up
 144 on the strong assumption that the latency values of the edges are fixed over time. In this
 145 paper, we aimed at understanding the edge latency setting in its most general form, and ask:

146 *Can one provide a general synchronization scheme in a setting that allows the latency*
 147 *of the network edges to vary in an adversarial manner in each round?*

148 A priori it is not so clear if one can compute even basic Boolean functions without assuming
 149 that the latency values are fixed. We answer this question in the affirmative by presenting
 150 a modular synchronization scheme using a different approach than that taken in [12]. The
 151 benefit of this approach is in its modularity. We start by understanding the implementation
 152 of a single NOT gate in this model in terms of upper and lower bounds on the space and the
 153 time of the computation. We then use this synchronized NOT solution as a building block in
 154 the final synchronized network solution. Specifically, the synchronized NOT gates are used
 155 to build synchronized circuits (and hence threshold gates) which in turn combined into a
 156 whole synchronized network solution. The space and time overheads incurred by our solution
 157 are polynomial in L (the bound on the latency) and Δ , the maximal incoming degree in the
 158 base network \mathcal{N} .

159 We next turn to consider synchronizers for the node delay model. Our approach is based
 160 on showing a *simulation* result that takes any synchronized solution $\text{sync}_E(\mathcal{N}, L)$ obtained
 161 by the synchronizer in the L -bounded *edge* latency model, and transforms it into a syn-
 162 chronized solution $\text{sync}_V(\mathcal{N}, T)$ that works in the T -bounded *node* delay model for $L = \Theta(T^2)$.
 163

164 **Remark.** It is noteworthy that in contrast to the *distributed* setting of Awerbuch and
 165 Peleg [2] where the network size does not depend on the latencies, in the neural setting it is
 166 not the case. As our lower bound constructions, both the computation time and the network
 167 size must depend (in fact, polynomially) on the largest edge latency. For this reason, for
 168 any practical purposes, the study of asynchronous communication, in general, must assume
 169 bounded delays.

170 1.4 Our Results

171 We study the cost and limitations of asynchronous neural computation in a biologically
 172 plausible yet simple model of spiking neural networks. Our main focus is in the edge latency
 173 model where the dynamic is specified by an L -bounded function $\ell : V \times V \times \mathbb{N} \rightarrow \mathbb{N}_{\leq L}$. The
 174 node delay model is concerned only towards the end of the paper (Appendix C), as it is
 175 handled via reduction to the edge delay setting. In the first part of the paper, we show
 176 several negative results for the L -bounded model. This includes an impossibility results for
 177 delay on self-loop edges, as well as size and time lower bound on an implementation of a NOT
 178 gate in this model. In the second part, we consider the construction of synchronizers in this
 179 generalized setting. We note that these constructions are self-contained and are technically

180 different from [13].

181

182 **Negative Results:** We first show that without assuming a minimal latency value on
 183 the self-loop edges, one cannot compute $AND(x, y)$ given two Boolean inputs x and y , even
 184 when the edge latency values are fixed over time and the largest latency is $L = 2$.

► **Theorem 1** (Impossibility for Arbitrary Latency Functions). *There exists no network that computes $AND(x, y)$ in a setting that allows the adversary to pick latencies in $\{1, 2\}$ for all edges in the network.*

185 The proof goes by showing that for any given candidate network solution \mathcal{N} , there exists
 186 a bad latency function ℓ under which \mathcal{N} fails to compute $AND(x, y)$. This holds even when
 187 the latencies are fixed over time. From that point on, we restrict attention to *nice* latency
 188 functions, that assign latency value 1 to the self-loop edges in the network.

189 ► **Definition 2.** *A latency function is nice if $\ell(u, u, \tau) = 1$ for every $u \in V$ and round τ .*

190 Our key technical contributions are lower bounds on the network size and the computation
 191 time for computing the $NOT(x)$ function in the L -bounded setting. Informally speaking, the
 192 $NOT(x)$ function appears to be a “complete” function for the purpose of synchronization in
 193 this asynchronous model. Indeed, our $NOT(x)$ implementation captures most of the essence
 194 of the L -bounded model. To obtain the final synchronization scheme we mainly glue together
 195 synchronized NOT units. For this reason, we spend much attention into understanding the
 196 tightness of our constructions by providing nearly matching lower bound results.

► **Lemma 3** (Size and Time Lower Bounds for Async. Computation of $NOT(x)$). *Any network that computes $NOT(x)$ in the L -bounded asynchronous setting must use $\Omega(L)$ neurons and $\Omega(L^3)$ time (the time lower bound is tight).*

197 This should be compared with the size lower bound of $\Omega(\log L)$ shown by [13] for their
 198 simplified asynchronous setting.

199

200 **Positive Results:** Our end result is a synchronizer that given any network \mathcal{N} in the
 201 standard synchronous setting and an integer L , computes a network $\text{sync}_E(\mathcal{N}, L)$ that
 202 performs the “same” computation as \mathcal{N} in the L -bounded edge delay setting.

► **Theorem 4** (Synchronizers for Edge Delays). *There exists a synchronizer that given a network \mathcal{N} with n neurons, maximum in-degree Δ , and maximum edge latency L , constructs a network $\text{sync}_E(\mathcal{N}, L)$ that has an “analogous” execution in the L -bounded edge-delay setting with a total number of $\tilde{O}(L^4 \cdot \text{poly}(\Delta) \cdot n)$ neurons and a time overhead^a of $\tilde{O}(L^5 \cdot \log \Delta)$.*

^a The \tilde{O} hides a factor of $\text{poly}(\log(n \cdot r))$, where r is the number of simulation rounds.

203 Although the construction is inspired by the work of Awerbuch and Peleg [2], the imple-
 204 mentation is very different as our neurons, unlike processors in a distributed network, are
 205 memoryless. Thus, they cannot aggregate the incoming messages as in [2]. Our construction
 206 is also different than that of [13], as the latter crucially depends on having fixed latencies
 207 over time.

208 For the node delay model, in Appendix C we show that given a network \mathcal{N} in the standard
 209 synchronous setting and an integer T , one can compute an analogous network $\text{sync}_V(\mathcal{N}, T)$
 210 in the node-delay model with bounded node delay T by taking the following approach. Apply
 211 the algorithm of Theorem 4 with \mathcal{N} and $L = \Theta(T^2)$. This results with a network $\text{sync}_E(\mathcal{N}, L)$
 212 that performs the same computation as \mathcal{N} in the L -bounded edge delay model. The desired
 213 network $\text{sync}_V(\mathcal{N}, T)$ is then obtained by multiplying the edge weights of a carefully defined
 214 edge *subset* in $\text{sync}_E(\mathcal{N}, L)$ by a factor of T . The quite delicate analysis then shows that
 215 the network $\text{sync}_V(\mathcal{N}, T)$ indeed simulates the original network \mathcal{N} upon any selection of
 216 the node-delay function $t : V \rightarrow \mathbb{N}_{\leq T}$. By setting $L = O(T^2)$ in Theorem 4, we show the
 217 following for the node delay model:

► **Theorem 5 (Synchronizers for Node Delays).** *There exists a synchronizer that given a network \mathcal{N} with n neurons, maximum in-degree Δ , and maximum node delay T , constructs a network $\text{sync}_V(\mathcal{N}, T)$ that has an “analogous” execution in the T -bounded node-delay setting with a total number of $\tilde{O}(T^8 \cdot \text{poly}(\Delta) \cdot n)$ neurons and a time overhead of $\tilde{O}(T^{10} \cdot \log \Delta)$.*

218 We note that our preference to take a modular approach rather than an optimized one
 219 inevitably leads to suboptimal space and time bounds in both Theorems 4 and 5. For example,
 220 Theorem 5 is shown via a simulation result, which further deepens our understanding of the
 221 connections between these models. We believe that by employing a more direct approach
 222 for building synchronizers in the node-delay model, one should get a considerably improved
 223 dependency in the delay bound T .

224 1.5 Our Approach in a Nutshell

225 We next provide the high level ideas for the key contributions. Throughout, unless stated
 226 otherwise we consider the edge delay model where the dynamics is specified by an arbitrary
 227 latency function.

228
 229 **Size and Time Lower Bound for Computing $NOT(x)$.** A network \mathcal{N} with input
 230 neuron x and an output neuron z computes the function $NOT(x)$ in the asynchronous setting
 231 if the following holds: when $x = 0$, the output z fires in at *least* one round regardless of the
 232 latency function; and when $x = 1$ the output z never fires for any latency function. To show
 233 a size lower bound of $\Omega(L)$ we take the following approach. First, we reduce any network
 234 \mathcal{N} that computes $NOT(x)$ into a simpler and yet not larger network \mathcal{N}_{simple} . In the latter
 235 network the only inhibitor is the input x which also has a self-loop of large positive weight,
 236 and outgoing edges of very large negative weight to all the excitatory neurons in \mathcal{N} . The
 237 second part of the proof shows a lower bound for \mathcal{N}_{simple} using its specialized structure. We
 238 will assume towards a contradiction that the in-degree of each neuron in \mathcal{N}_{simple} is less than
 239 L and exhibit two conflicting latency functions ℓ_0, ℓ_1 that satisfy the following. If \mathcal{N}_{simple}
 240 computes $NOT(x)$ with ℓ_0 and with $x = 0$, then it must fail to compute $NOT(x)$ with the
 241 function ℓ_1 and $x = 1$. To compute these latency functions, we partition the simulation into
 242 blocks T_0, \dots , each containing L rounds. In each phase i , we set the latency values for all the
 243 edges and all the rounds in block T_i . Throughout, we will keep the invariant that there exists
 244 *no* neuron that fires when $x = 0$ and with ℓ_0 , but does not fire when $x = 1$ and with ℓ_1 . By
 245 the correctness of the network, the output neuron must fire at least once when $x = 0$, thus
 246 leading to the contradiction. In the very high level, the fact that the in-degree of each vertex
 247 is small is used in order to spread over the at most L incoming spikes of each neuron u in a

248 balanced manner over the L rounds of the block. This will prevent the firing of a neuron z
 249 when $x = 0$. Our lower bound is complemented by an upper bound of $O(L^2)$ as described next.

250

251 **The Generalized Synchronization Scheme.** The scheme is based on gradual steps.

252

253 **Step I: Synchronization of NOT and OR Gates.** We start by considering the asyn-
 254 chronous computation of simple Boolean functions $NOT(x)$ and $OR(x_1, \dots, x_\ell)$ with a small
 255 number of neurons. The key challenge is in implementing the NOT gate. When $x = 1$,
 256 the output gate is required not to fire (i.e., output 0) throughout the *entire* execution. In
 257 contrast, when $x = 0$ the output gate should fire at least once during the execution. The
 258 construction is combinatorial and uses a similar logic to the lower bound arguments. It
 259 contains a collection of $L + 1$ neurons with outgoing edges to the output z . The above
 260 mentioned lower bound result shows that the incoming degree of z must be at least $L - 2$.


261

262 **Step II: Synchronization of a Boolean Circuit.** Any Boolean circuit \mathcal{A} can be imple-
 263 mented by NOT and OR gates. To simulate the computation of \mathcal{A} in the asynchronous
 264 setting, we replace each gate g_i in \mathcal{A} by its synchronized implementation $\text{Sync}(g_i)$ constructed
 265 in Step (I). For a gate g_i in layer j with incoming gates $g_{i,1}, \dots, g_{i,k}$, the input to the
 266 sub-network $\text{Sync}(g_i)$ are the output neurons of the sub-networks $\text{Sync}(g_{i,1}), \dots, \text{Sync}(g_{i,k})$.
 267 The synchronization between the layers of the circuit is governed by a directed chain of
 268 $O(dL^3)$ neurons. The head of the chain fires in the first round of the simulation and activates
 269 the network. The sub-networks $\text{Sync}(g_i)$ of gates g_i in layer j are activated by the $\Theta(j \cdot L^3)$
 270 neuron in this chain. These parameters are set so that we can be sure that the modules of
 271 layer j are activated, only *after* the spikes from the output neurons of the previous layer
 272 have reached the input of this layer. Overall the synchronized transformed network $\text{Sync}(\mathcal{A})$
 273 has $O(d \cdot L^3 + m \cdot L^2)$ neurons, where d is the depth of \mathcal{A} and m is the number of gates. The
 274 overtime in the computation is $O(d \cdot L^4)$ rounds.

275

276 **Step III: Synchronization of a Single (Probabilistic) Threshold Gate.** To syn-
 277 chronize a single deterministic threshold gate, we use the fact that a threshold gate with
 278 incoming degree Δ can be implemented by a Boolean circuit with $\text{poly}(\Delta)$ neurons and depth
 279 $O(\log \Delta)$. This allows us to use the synchronized construction of the previous step. Turning
 280 to probabilistic threshold gates, here it is much less clear how to implement such a gate by a
 281 Boolean circuit. We take the following approach. First, we use the fact from [20] that a spik-
 282 ing neuron³ u with bias $b(u)$ is equivalent to a *deterministic* neuron u' whose bias is sampled
 283 from the Logistic distribution with mean $b(u)$. Therefore our key challenge is in sampling
 284 a value from a given Logistic distribution. To do that, we use a collection of k (input-less)
 285 spiking neurons each fires independently with probability half. These neurons provide us the
 286 random bits for this process of sampling. In fact, these fair coins tosses allows one to sample
 287 a value *almost* uniformly at random in the range $[0, 2^k]$. We will then use the method of
 288 inverse transform sampling to convert this almost-uniform sampled value to a value that is
 289 sampled from the Logistic distribution up to a small error in the sampling. Using Taylor
 290 expansion of the natural log function, we implement this Uniform to Logistic transformation
 291 by a collection of simple arithmetic operations applied on a collection of Boolean neurons.
 292 The total error in our sampling is set to be small enough so that the output distribution
 293 of the Boolean circuit is almost indistinguishable from that of the probabilistic threshold gate.

³ The probabilistic threshold gates of SNN.



■ **Figure 1** A road-map for synchronizing spiking neural networks.

294

295 **Grand Finale: Synchronization of a Spiking Neural Network.** Finally, given an
 296 SNN network \mathcal{N} of (probabilistic) threshold gates the synchronized network $\text{sync}(\mathcal{N})$ is ob-
 297 tained as follows. Each threshold gate g_i in \mathcal{N} is replaced by its synchronized implementation
 298 $\text{Sync}(g_i)$. The key challenge is in synchronizing these modules so that every neuron v in \mathcal{N}
 299 (i.e., not only the output neuron) has an equivalent neuron v' in $\text{sync}(\mathcal{N})$ that simulates v for
 300 any possible latency function throughout the entire simulation. See Fig. 1 for an illustration.

301

302 **From Edge Delays to Node Delays.** Given a network \mathcal{N} to be simulated and an
 303 integer T , our goal is to build a synchronizer $\text{sync}_V(\mathcal{N}, T)$ that *simulates* \mathcal{N} in the T -
 304 bounded node-delay model. To do that, we first compute a network $\text{sync}_E(\mathcal{N}, L)$ that
 305 simulates \mathcal{N} in the L -bounded edge-delay model for $L = \Theta(T^2)$. Then the output network
 306 $\text{sync}_V(\mathcal{N}, T)$ is obtained by dividing some of the edge weights in $\text{sync}_E(\mathcal{N}, L)$ by a factor
 307 of T . The correctness argument is based on showing that for every node-delay function
 308 $t : V \rightarrow \mathbb{N}_{\leq T}$, there exists an edge-delay function $\ell : V \times V \times \mathbb{N}_{\geq 0} \rightarrow \mathbb{N}_{\leq L}$ such that the
 309 execution of the network $\text{sync}_E(\mathcal{N}, L)$ with the edge-delay function ℓ (in the edge-delay
 310 model) is *similar* to the execution of the network $\text{sync}_V(\mathcal{N}, T)$ with the node-delay function
 311 t (in the node-delay model). Since the network $\text{sync}_E(\mathcal{N}, L)$ simulates the original network
 312 \mathcal{N} for any edge-delay function, it will imply that the network $\text{sync}_V(\mathcal{N}, T)$ simulates the
 313 original network \mathcal{N} for any node-delay function as desired.

314

315 **Additional Related Work.** Asynchronized communication in spiking neural networks has
 316 been studied in several settings. Maass [22, 25] considered a quite elaborated model for
 317 deterministic neural networks with *arbitrary* response *functions* for the edges, and a vector
 318 firing times for all neuron. The approach of [22] mostly concerned the computational power
 319 of this model upon choosing the best parameters for the network. I.e., showing feasibility
 320 results for various functions. In contrast, in this work our goal is to *bound* the computation
 321 time and the network size under this asynchronous setting. Khun et al. [15] studied the
 322 asynchronous dynamics under the stochastic model of DeVille and Peskin [7].

323 Turning to the setting of logical circuits, there is a long line of work on the asynchronous
 324 setting under various model assumptions [1, 11, 36, 4] that do not quite fit the memory-less
 325 setting of spiking neurons. A more related work to our setting is by Martin, Manohar and
 326 Moses [28, 26, 27] who studied the computational power of asynchronous digital circuits. In
 327 particular, they characterize the necessary and sufficient conditions for a valid operation of a
 328 given circuit in the asynchronous setting. For example, they showed that if all edges and
 329 nodes suffer from an unbounded delay then the computational power of the circuit must be
 330 very limited. The focus in our work is quite different. Instead of studying the computational
 331 power of the asynchronous setting, we bound the computational overhead for solving concrete
 332 problems.

2 The Synchronous and Asynchronous SNN Models

A deterministic neuron u is modeled by a *deterministic* threshold gate. Letting $b(u)$ to be the threshold value of u , then u outputs 1 if the weighted sum of its incoming neighbors exceeds $b(u)$. A *spiking neuron* is modeled by a probabilistic threshold gate which fires with a sigmoidal probability that depends on the difference between its weighted incoming sum and $b(u)$.

Neural Network Definition. A *Neural Network* (NN) $\mathcal{N} = \langle X, Z, Y, w, b \rangle$ consists of n input neurons $X = \{x_1, \dots, x_n\}$, m output neurons $Y = \{y_1, \dots, y_m\}$, and k auxiliary neurons $Z = \{z_1, \dots, z_k\}$. In spiking neural network (SNN), the neurons can be either deterministic threshold gates or probabilistic threshold gates. The directed weighted synaptic connections between $V = X \cup Z \cup Y$ are described by the weight function $w : V \times V \rightarrow \mathbb{R}$. A weight $w(u, v) = 0$ indicates that a connection is not present between neurons u and v . Finally, for any neuron v , the value $b(v) \in \mathbb{R}$ is the threshold value (activation bias). The in-degree of every input neuron x_i is zero, i.e., $w(u, x) = 0$ for all $u \in V$ and $x \in X$. Additionally, each neuron is either inhibitory or excitatory: if v is inhibitory, then $w(v, u) \leq 0$ and if v is excitatory, then $w(v, u) \geq 0$ for every u . This restriction arises from the biological structure of the neurons.

Network Dynamics in the Synchronous Setting. The network evolves in discrete, synchronous rounds as a Markov chain. The firing probability of every neuron in round τ depends on the firing status of its neighbors in round $\tau - 1$, via a standard sigmoid function, with details given below. For each neuron u , and each round $\tau \geq 0$, let $\sigma_\tau(u) = 1$ if u fires (i.e., generates a spike) in round τ . Let $\sigma_0(u)$ denote the initial firing state of the neuron. The firing state of each input neuron x_j in each round is the input to the network. For each non-input neuron u and every round $\tau \geq 1$, let $\text{pot}(u, \tau)$ denote the membrane potential at round τ and $p(u, \tau)$ denote the firing probability ($\Pr[\sigma_\tau(u) = 1]$), calculated as $\text{pot}(u, \tau) = \sum_{v \in V} w(v, u) \cdot \sigma_{\tau-1}(v) - b(u)$ and $p(u, \tau) = \frac{1}{1 + e^{-\text{pot}(u, \tau)/\lambda}}$ where $\lambda > 0$ is a *temperature parameter* which determines the steepness of the sigmoid. Clearly, λ does not affect the computational power of the network, thus we set $\lambda = 1$.

2.1 Network Dynamics in the Edge Delay Setting

The dynamic of the network is governed by a latency function $\ell : V \times V \times \mathbb{N} \rightarrow \mathbb{N}$ interpreted as follows. For every directed edge $e = (u, v)$ and round τ , a spike generated by u in round τ arrived at v after $\ell(u, v, \tau)$ rounds. In the synchronous setting, $\ell(u, v, \tau) = 1$ for every u, v, τ .

For every neuron v and round τ , let $A(v, \tau) = \{(v, \tau') \mid v \in V, \tau' + \ell(v, u, \tau') = \tau\}$ denote all the spike events that if occur, arrive to u at round τ . The state of u in round τ is given by:

$$\text{pot}(u, \tau) = \sum_{(v, \tau') \in A(v, \tau)} w(v, u) \cdot \sigma_{\tau'}(v) - b(u) \quad \text{and} \quad \sigma_\tau(u) = 1 \text{ iff } \text{pot}(v, \tau) \geq 0. \quad (1)$$

If u is a probabilistic threshold gate then it fires with probability $p(u, \tau) = \frac{1}{1 + e^{-\text{pot}(u, \tau)}}$. When $\ell(u, v, \tau) = \ell(u, v, \tau')$ for every u, v and $\tau' \neq \tau$, we may omit τ and write $\ell(u, v)$.

► **Definition 6 (The L -bounded Edge-Delay Setting).** *Given is a network \mathcal{N} and an integer L . It is assumed the network contains a special neuron, the starter, that fires in the first round of the simulation. The dynamic is determined by a latency function ℓ . This function ℓ can be chosen arbitrarily among all L -bounded nice functions.*

376 ► **Definition 7** (Computation of a Boolean Function in the L -bounded Edge-Delay Setting). Let
 377 $f : \{0, 1\}^n \rightarrow \{0, 1\}^k$ be a Boolean function. A network \mathcal{N} with n input neurons x_1, \dots, x_n
 378 and k output neurons z_1, \dots, z_k computes f in this setting if for every nice L -bounded function
 379 ℓ and for every fixed possible assignment to the input neurons b_1, \dots, b_n the following holds:
 380 (i) If $f_i(b_1, \dots, b_n) = 1$, then there exists a round in which z_i fires, where $f_i(\cdot)$ is the i^{th} bit in
 381 the output of f . (ii) If $f_i(b_1, \dots, b_n) = 0$ then z_i does not fire throughout the entire execution.
 382 Furthermore, the network \mathcal{N} computes the function f in r rounds if \mathcal{N} computes f , and for
 383 every nice L -bounded function ℓ , input bits b_1, \dots, b_n and index i such that $f_i(b_1, \dots, b_n) = 1$,
 384 z_i fires in some round $\tau \leq r$.

385 Note that by this definition, for a network \mathcal{N} that computes a Boolean function f within
 386 r rounds, to evaluate the output of the function it is sufficient to inspect the state of the
 387 output bits over the first r rounds of the network's simulation. Furthermore, as edge delays
 388 are allowed to be chosen in an adversarial manner, one cannot hope for having all output
 389 neurons to fire in the exact same around. One mechanism that we use to keep the out-
 390 put neurons fire simultaneously is by using self-loop edges whose latency values is fixed to be 1.

391
 392 **Synchronizers.** A synchronizer ν is an algorithm that gets as input a network \mathcal{N} and
 393 outputs a network $\mathcal{N}' = \text{sync}(\mathcal{N})$ that contains all the neurons of \mathcal{N} , plus additional auxiliary
 394 neurons. One of the auxiliary neurons in \mathcal{N}' is a *starter* neuron that fires in the *first* round of
 395 the simulation. The network \mathcal{N}' works in the asynchronous setting and should have *similar*
 396 *execution* to \mathcal{N} in the sense that for every neuron $v \in V(\mathcal{N})$, the firing pattern of v in the
 397 asynchronous network should be similar to the one in the synchronous network. The output
 398 network \mathcal{N}' simulates each round of the network \mathcal{N} by a *phase*.

399 ► **Definition 8** (Phases). We partition the execution of \mathcal{N}' into phases $1, 2, \dots$, using a
 400 function $r : V(\mathcal{N}) \times \mathbb{N} \rightarrow \mathbb{N}$ that defines the beginning of phase p , i.e. the p^{th} phase is the
 401 round interval $[r(v, p), r(v, p + 1))$.

402 ► **Definition 9** (Similar Executions (Deterministic Networks)). The synchronous execution Π
 403 of a deterministic network \mathcal{N} is specified by a list of states $\Pi = \{\sigma_1, \dots\}$ where each σ_i is
 404 a binary vector describing the firing status of the neurons in round i . The asynchronous
 405 execution of the network $\mathcal{N}' = \text{sync}_E(\mathcal{N}, L)$ with a latency function ℓ denoted by $\Pi'(\ell)$ is
 406 defined analogously only when applying the asynchronous dynamic of Eq. (1). The execution
 407 $\Pi'(\ell)$ is divided into phases according the a function $r : V(\mathcal{N}) \times \mathbb{N} \rightarrow \mathbb{N}$.

408 The network \mathcal{N} and the pair $\langle \mathcal{N}', \ell \rangle$ have a **similar execution** if $V(\mathcal{N}) \subseteq V(\mathcal{N}')$, and
 409 in addition, a neuron $v \in V(\mathcal{N})$ fires in round p in the execution Π iff v fires during phase p
 410 in $\Pi'(\ell)$. The networks \mathcal{N} and \mathcal{N}' are **similar** if \mathcal{N} and $\langle \mathcal{N}', \ell \rangle$ have a similar execution for
 411 every nice latency function ℓ .

412 Note that specifically, if a synchronous network \mathcal{N} computes a Boolean function f by round
 413 r and \mathcal{N} and \mathcal{N}' are similar, then \mathcal{N}' computes f by phase r . Therefore, if we know that
 414 each phase is of at most q rounds, we get that \mathcal{N}' computes f in $r \cdot q$ rounds.

415 Finally, we note that the extension for randomized networks with probabilistic gates is
 416 quite straightforward if one simply fixed the random coins used by the neurons over the
 417 simulation. That is, to be able to faithfully compare the simulation of two random networks,
 418 one has to fix the random coins of both of the simulations to be the same. For this reason,
 419 given an input randomized network \mathcal{N} in the synchronized model, we maintain all the random
 420 coins generated by the neurons of the network over the simulation. These random coins
 421 are then fed to the network \mathcal{N}' (i.e., obtained by applying our synchronizers). Since we

422 compare two randomized networks that use the same set of random coins, we can treat
 423 these networks as deterministic. In Appendix C we provide the analogous definitions for the
 424 T -bounded node-delay model. Throughout the main paper, we consider only the edge-delay
 425 model and to avoid cumbersome notation that synchronized network solutions for this model
 426 are denoted by $\text{sync}(\mathcal{N})$ (rather than $\text{sync}_E(\mathcal{N}, L)$).

427 3 Negative Results

428 **Impossibility Result for Arbitrary Latency Functions.** We start by considering
 429 Theorem 1, and show that if the latency values are allowed to be set in an adversarial manner
 430 in $\{1, 2\}$, then there exists no network that computes the AND of two Boolean inputs. In
 431 Appendix A, we show:

432 ► **Lemma 10.** *Given input neurons x, y and an output neuron z , there is no network*
 433 *computing $\text{AND}(x, y)$ under every latency function $\ell : V \times V \rightarrow \{1, 2\}$.*

434 In the high level, we show that one can set the latency values such that all the spikes that
 435 *depend* on the value of x (resp., y) arrive at odd (resp., even) rounds. Therefore, at any
 436 round, there is no neuron that fires as a function of *both* x and y .

437 Size and Time Lower Bound

438 In this section we show the proof for Lemma 3. Here we focus on the size lower bound
 439 although the high level proof strategy for the time lower bound is quite similar. The time
 440 lower bound is presented in Appendix A.1. Our proof strategy is as follows. First we reduce
 441 any network \mathcal{N} that computes $\text{NOT}(x)$ in the asynchronous setting, to a network $\mathcal{N}_{\text{simple}}$
 442 with a simpler structure that makes it easier to make arguments on it. The second part of
 443 the argument shows the lower bound for simple networks. All missing proofs of this section
 444 are in Appendix A.

445 ► **Definition 11 (Strong Neurons and Simple Networks).** *A neuron u is strong in a given*
 446 *network if $w(u, u) \geq b(u)$, and otherwise it is weak. Note that specifically, an excitatory*
 447 *neuron u with $b(u) \leq 0$ is strong. Given a single input neuron x , we say that a network*
 448 *\mathcal{N} is simple if the following hold: (i) x is a strong neuron and has an outgoing edge of*
 449 *infinite negative weights to all other neurons in the network; and (ii) all other neurons are*
 450 *excitatory.*

451 We note that the simple network is not a *legally defined* neural network: the input neuron
 452 has an incoming edge (self-loop), and it is an inhibitor with a positive self-loop. However,
 453 this network definition is only for the sake of the analysis and as such, it is not restricted to
 454 follow any rule.

455
 456 **Reduction to Simple Networks.** Given a network \mathcal{N} with an input neuron x , define
 457 $\mathcal{N}_{\text{simple}}$ as follows. Exclude all the inhibitory neurons from $\mathcal{N}_{\text{simple}}$ and take all edges
 458 between excitatory neurons to be as in \mathcal{N} . Then, add a self-loop of infinite weight to the
 459 input neuron x , and connect it to every neuron with infinite⁴ negative weight.

⁴ By infinite we mean large enough so that when the spike by x arrives at some neuron v , v would not fire.

460 ► **Lemma 12.** *If \mathcal{N} computes $NOT(x)$ within r rounds starting with the initial state $\bar{\sigma}$,*
 461 *then also \mathcal{N}_{simple} computes it within r rounds, when starting with the initial states as in $\bar{\sigma}$*
 462 *restricted to the vertices of \mathcal{N}_{simple} .*

463 The proof goes by claiming that for any latency function ℓ_{simple} for \mathcal{N}_{simple} , we can show
 464 the existence of a latency function ℓ for \mathcal{N} whose performance is only *worse* than that of
 465 \mathcal{N}_{simple} with ℓ_{simple} . That is, when $x = 0$ (resp., $x = 1$) then the potential of all neurons in
 466 $\mathcal{N}_{simple}, \ell_{simple}$ is not decreased (resp. increased) when compared to \mathcal{N}, ℓ . Since the network
 467 \mathcal{N} computes $NOT(x)$ with the latency function ℓ within r rounds, we conclude that also
 468 \mathcal{N}_{simple} computes it with the latency function ℓ_{simple} within at most r rounds.

469 Fix a $NOT(x)$ network \mathcal{N} . For an integer r , a latency function ℓ for \mathcal{N} is *r-good* with the
 470 initial configuration $\bar{\sigma}$, if the network computes $NOT(x)$ within r rounds. I.e., when $x = 0$,
 471 the output of \mathcal{N} fires in some round $\tau \leq r$, and when $x = 1$ it never fires when all latencies
 472 are given based on ℓ . If ℓ is *r-good* for some integer r we say it is *good*, and otherwise the
 473 latency function is *bad* (the network fails to compute $NOT(x)$). Note that in order for a
 474 network to compute $NOT(x)$ within r rounds, it is required that any latency function is
 475 *r-good* for a fixed initial configuration.

476
 477 **Lower Bound for Simple Networks.** Assume towards contradiction that there ex-
 478 ists a *simple* network $\mathcal{N} = \mathcal{N}_{simple}$ with maximum in-degree $\Delta_{in} < L - 2$ that computes
 479 $NOT(x)$. I.e., there exists an initial configuration $\bar{\sigma}$ for all neurons but x such that every
 480 latency function ℓ is good for $\langle \mathcal{N}, \bar{\sigma} \rangle$. In what follows, we define two *conflicting* latency
 481 functions ℓ_0 and ℓ_1 , such that if ℓ_0 is good when the initial state of x is 0, then it implies
 482 that ℓ_1 is bad when the initial state of x is 1.

483
 484 **Defining the Latency Functions ℓ_0 and ℓ_1 .** Recall that for every $b \in \{0, 1\}$, $\bar{\sigma}_b = [b, \bar{\sigma}]$ is
 485 the initial state vector where x has the initial state b and the initial states of all other neurons
 486 is specified by the vector $\bar{\sigma}$. The construction of ℓ_0, ℓ_1 is inductive. To avoid cumbersome
 487 notation, we start the simulation in round -1 rather than in round 0. For this first round
 488 -1 , let $\ell_b(v, u, -1) = 1$ for $v \neq x$, and $\ell_b(x, u, -1) = L$ for every u and $b \in \{0, 1\}$. Thus, the
 489 positive spikes (by any $v \neq x$) fired in round -1 arrive in round 0, and the negative spikes of
 490 x arrive in round $L - 1$.

491 To define the latency of the edges in the remaining rounds $\tau \geq 0$, we partition them into
 492 blocks, each of size L rounds where the i^{th} block is $T_i = [iL, iL + (L - 1)]$ for every $i \geq 0$. We
 493 continue in steps $i = 1, \dots$ where in step i , the latency values of $\ell_0(e, \tau), \ell_1(e, \tau)$ are defined
 494 for every edge e in the network \mathcal{N} , and for every round $\tau \in T_i$. For every $b \in \{0, 1\}$ and a
 495 block T_i , define $A_{i,b}$ as the set of neurons that fire (hence *active*) in the first round of T_i
 496 when executing \mathcal{N} with the initial configuration $\bar{\sigma}_b$, and the latency function ℓ_b . Throughout
 497 the process of defining the latency functions, we maintain these invariants at the beginning
 498 of step i :

- 499 ■ (I1) All the positive spikes generated at any round before the interval T_i arrive to their
 500 destination by the first round of T_i . Furthermore, all negative spikes generated at any
 501 round before the interval T_i arrive to their destination either by the first round of T_i or
 502 on the last round of T_i , namely, round $iL + (L - 1)$.
- 503 ■ (I2) $A_{i,0} \setminus \bigcup_{i' < i} A_{i',1} = \emptyset$.

504 We define the latency for the rounds in T_i , and then show that the invariants are maintained.

505 **Defining the latency function ℓ_1 for T_i .** For every self-loop edge e and every $\tau \in T_i$,
 506 let $\ell_1(e, \tau) = 1$. For every edge $e = (x, u)$ where $u \neq x$, and $\tau \in T_i \setminus \{iL + (L - 1)\}$, let
 507 $\ell_1(e, \tau) = (iL + (L - 1)) - \tau$, i.e., the spike of x arrives u in the last round of the interval

508 T_i . For $e = (x, u)$ and $\tau = iL + (L - 1)$, let $\ell_1(e, \tau) = L$ so that the spike arrives in the last
 509 round of the next block, i.e., round $(i + 1)L + (L - 1)$. For every other edge $e = (v, u)$ with
 510 $v \neq x$, let $\ell_1(e, \tau) = (i + 1)L - \tau$, i.e., the spike arrives at the first *first* round of the next
 511 block T_{i+1} .

512 **Defining the latency function ℓ_0 for T_i .** As for ℓ_1 , for a self-loop edge e , we set
 513 $\ell_0(e, \tau) = 1$. For an edge $e = (x, u)$ we set $\ell_0(e, \tau)$ arbitrarily (since $x = 0$, those values
 514 are meaningless). We now fix a neuron u , and set the latency values of all its incoming
 515 edges (v, u) . Since we have already defined the latency values of all edges up to block T_i ,
 516 at the beginning of step i , the sets $A_{i,0}, A_{i,1}$ can be computed. Let g_1, \dots, g_ω be the weak
 517 incoming neighbors of u in $A_{i,0}$, and h_1, \dots, h_s be the strong incoming neighbors of u in
 518 $A_{i,0}$. We Consider two cases. The neuron u is said to have a *dominant* neighbor if it has
 519 a neighbor with a *sufficiently* large incoming weight, where the precise weight threshold
 520 depend on whether the incoming neighbor is weak or strong. Specifically, it has a dominant
 521 neighbor if it has either a weak neighbor g_j with $w(g_j, u) \geq b(u)$, or a strong neighbor h_j
 522 with $w(h_j, u) \geq b(u)/(L - 1)$.

523 ■ **Case 1: u has a dominant neighbor.** Let $\ell_0(e, \tau) = (i + 1)L - \tau$ for every incoming
 524 edge $e = (v, u)$. That is, we schedule all the incoming spikes of u in this block to arrive
 525 at u in the *first* round of the next block T_{i+1} .

526 ■ **Case 2: u has no dominant neighbor.** Since $\deg(u) < L - 2$, we have that $\omega + s < L - 2$,
 527 and in particular $\omega \leq L - 2$. For each weak neuron g_j , set $\ell_0(g_j, u, iL) = j + 1$. That
 528 is, the spike from g_j in round iL is scheduled to arrive at u in round $iL + (j + 1)$. For
 529 each strong neighbor h_j , we split all the spikes generated by h_j during the $\omega + 1$ rounds
 530 $iL, \dots, iL + \omega$ in a balanced manner over $L - (\omega + 1)$ rounds. Specifically, we set the
 531 latency values of the at most $\omega + 1$ spikes by h_j during the rounds $iL, \dots, iL + \omega$ such
 532 that in each round $\tau \in [iL + (\omega + 2), (i + 1)L]$, u receives at most $(\omega + 1)/(L - (\omega + 1))$
 533 spikes from h_j ⁵. For every $\tau \in [iL + (\omega + 1), iL + (L - 1)]$, let $\ell_0(h_j, u, \tau) = 1$, i.e., the
 534 spike arrives one round later. The latency of all the remaining edges e and rounds τ in
 535 T_i is set to $\ell_0(e, \tau) = (i + 1)L - \tau$, so that it arrives in round $(i + 1)L$.

536 In AppendixA.1.2, we prove that the invariants hold by induction on the number of rounds.
 537 Since the output z is required to fire when $x = 0$ but must not fire when $x = 1$, we get the
 538 desired contradiction. In Appendix A.1, we show the time lower bound of $\Omega(L^3)$ rounds.
 539 This bound is tight, and the construction while having a similar high-level ideas is slightly
 540 more involved than the size lower bound.

541 4 Upper Bounds

542 4.1 Synchronization of Logic Gates and Boolean Circuits

543 First observe that the simple implementation of an OR-gate works also in the asynchronous
 544 setting.

545 ▷ **Observation 13 (OR gate).** Given input neurons x_1, \dots, x_n and output neuron z , there
 546 exists a deterministic network OR_{sync} with *no* auxiliary neurons, that computes the OR gate
 547 of x_1, \dots, x_n using L rounds. I.e, it holds that: (i) If $\sigma_0(x_1) \vee \dots \vee \sigma_0(x_n) = 0$, then $\sigma_t(z) = 0$
 548 for every t , and (ii) If $\sigma_0(x_1) \vee \dots \vee \sigma_0(x_n) = 1$, then there exists a round $t \in [1, L]$ such that
 549 $\sigma_t(z) = 1$. Moreover, if an input neuron fires in round τ , the output neuron z fires in some
 550 round $t \in [\tau + 1, \tau + L]$.

⁵ For simplicity, we assume that $(\omega + 1)$ divides $(L - (\omega + 1))$

551 We next consider the more technically involved setting of synchronizing a NOT gate.

552 ► **Lemma 14** (NOT gate). *There is a network NOT_{sync} of size $O(L^2)$ with input neuron x*
 553 *and output z , that computes $\text{NOT}(x)$ within $O(L^3)$ rounds. I.e., it holds that: (i) If $\sigma_0(x) = 1$,*
 554 *then $\sigma_t(z) = 0$ for every t , and (ii) If $\sigma_0(x) = 0$, then there exists a round $t \in [L, \Theta(L^3)]$*
 555 *such that $\sigma_t(z) = 1$.*

556 The following synchronous implementation assumes that the network contains a special
 557 starter neuron v^* that fires at the beginning of the simulation, regardless of the input value
 558 of x . Later on in Section 4.3, when presenting the complete synchronization scheme, this
 559 neuron v^* will receive the starting firing signal from the *global* pulse generator.

560

561 **Network Description.** The network consists of the following components, see Figure 2.

- 562 1. A *chain* $C = [c_0 = v^*, \dots, c_{5L^2}]$ containing $5L^2 + 1$ neurons. The head of the chain is
 563 the starter neuron that fires in the first round. For every $i \geq 0$, the neuron c_i has bias
 564 $b(c_i) = 1$. Moreover, for every $i \geq 1$ the neuron c_i has an incoming edge from c_{i-1} with
 565 weight 1.
- 566 2. A *memory neuron* m that remembers the initial state of x . The memory neuron has a
 567 positive incoming edge from x , as well as a self-loop both with weight 1 and bias $b(m) = 1$.
- 568 3. A *reset* inhibitory neuron r with an edge from m of weight $w(m, r) = 1$, and bias $b(r) = 1$.
- 569 4. A collection of $L + 1$ intermediate neurons v_0, \dots, v_L that are connected to the output
 570 neuron z , where each v_i has an incoming edge from the neuron $c_{5 \cdot iL} \in C$ with weight
 571 $w(c_{5 \cdot iL}, v_i) = 1$, a self-loop of weight 1 and bias $b(v_i) = 1$. In addition, each v_i has a
 572 negative incoming edge from the reset neuron r with weight $w(r, v_i) = -\infty$. Finally, each
 573 v_i has an edge to z with weight $w(v_i, z) = 1$ and bias $b(z) = L + 1$.

574 The correctness of the construction and the proof of Lemma 14 are deferred to Appendix B.1.

575

576 **Synchronization of a Boolean Circuit.** Given the synchronized sub-networks of Obser-
 577 vation 13 and Lemma 14, we now show how to synchronize a Boolean circuit that contains
 578 OR and NOT gates.

579 ► **Lemma 15.** *Given a Boolean circuit \mathcal{A} of OR and NOT gates with n inputs, k outputs,*
 580 *m gates and depth d , there exists a deterministic network $\mathcal{N} = \text{sync}(\mathcal{A})$ with input neurons*
 581 *x_1, \dots, x_n , output neurons z_1, \dots, z_k and $O(dL^3 + mL^2)$ auxiliary neurons, that computes \mathcal{A}*
 582 *in $O(dL^4)$ rounds. I.e., it holds that (i) If $[\mathcal{A}(\sigma_0(x_1), \dots, \sigma_0(x_n))]_i = 0$ then $\sigma_t(z_i, \mathcal{N}) = 0$*
 583 *for every t ; and (ii) If $[\mathcal{A}(\sigma_0(x_1), \dots, \sigma_0(x_n))]_i = 1$, then⁶ there exists $t \in [1, O(dL^4)]$ such*
 584 *that $\sigma_t(z_i, \mathcal{N}) = 1$.*

585 In the high-level, the network $\mathcal{N} = \text{sync}(\mathcal{A})$ is obtained by replacing each gate g_i with its
 586 synchronized module $\text{Sync}(g_i)$. The input neurons to the gate modules in layer j of \mathcal{A} are the
 587 output neurons of the gate modules of layer $j - 1$ in \mathcal{A} . The network then contains a chain
 588 of length $O(d \cdot L^3)$ to control the synchronization between layers: the modules of layer j are
 589 activated only after the modules of the previous layer have completed their computation.
 590 See Appendix B.2.

591 4.2 Synchronization of a Single Threshold Gate

592 **Deterministic Threshold Gate.** Given a deterministic threshold gate g with Δ inputs,
 593 one can implement g using a Boolean Circuit with $\text{poly}(\Delta)$ gates and depth $O(\log \Delta)$ (see

⁶ For a vector of n bits $x \in \{0, 1\}^n$, let $[x]_i$ denote the i^{th} bit of x . I.e., if $x = (x_1, \dots, x_n)$, then $[x]_i = x_i$.

594 Appendix B.3). Combining with the construction described in Lemma 15 we show the
 595 following:

596 ► **Lemma 16.** *Given a weighted threshold gate $g = f(x_1, \dots, x_\Delta)$, there exists a network $\mathcal{N} =$
 597 $\text{Sync}(g)$ with Δ input neurons x_1, \dots, x_Δ , an output neuron z , and $O(\log \Delta \cdot L^3 + \text{poly}(\Delta) \cdot L^2)$
 598 auxiliary neurons that computes f within $O(\log \Delta \cdot L^4)$ rounds. I.e. the output z fires in
 599 round $\tau \in [2, O(\log \Delta \cdot L^4)]$ if and only if $f(\sigma_0(x_1), \dots, \sigma_0(x_\Delta)) = 1$.*

600 **Probabilistic Threshold Gate.** We next turn to consider the more challenging setting of
 601 probabilistic threshold gates. To synchronize such gates, we first describe how to implement
 602 them by using a Boolean *circuit* \mathcal{A} that contains two types of gates: deterministic threshold
 603 gates, and input-less gates which outputs 1 with probability $1/2$. We hereafter denote
 604 the latter gates by uniformly random gates⁷. The output distributions of the probabilistic
 605 threshold gate and the output gate of \mathcal{A} will be very close up to a small additive error of $\epsilon \in$
 606 $(0, 1)$. The synchronized probabilistic gate will be obtained by applying the synchronization
 607 scheme of Lemma 15 on the circuit \mathcal{A} .

608 Our key result might be of independent interest in the context of Boolean circuits:

609 ► **Lemma 17.** *Given a probabilistic threshold gate g with Δ inputs, and an error para-
 610 meter $\epsilon \in (0, 1)$, there exists a Boolean circuit with depth $\text{poly}(\log \Delta, \log(1/\epsilon))$ and a total
 611 $\text{poly}(\Delta, \log(1/\epsilon))$ deterministic gates. In addition, there is a collection of $O(\log(1/\epsilon))$ uni-
 612 formly random gates (each outputs 1 independently w.p. $1/2$), and an output gate g' that
 613 approximates g in the following sense. Letting $p(\bar{x}), p'(\bar{x})$ be the probability that g, g' output 1
 614 given input \bar{x} , it holds that $|p(\bar{x}) - p'(\bar{x})| \leq \theta(\epsilon)$ for any fixed assignment of input \bar{x} .*

615 Our starting point is the following useful fact from [20]:


616 ▷ **Observation 18.** Let g_1 be a *probabilistic* gate with an incoming weighted sum W and
 617 bias b_1 . Let g_2 be a *deterministic* threshold gate with incoming weighted sum W and bias
 618 b_2 , where b_2 is sampled from the Logistic distribution with mean b_1 and scale 1. Then
 619 $\Pr[g_1 = 1] = \Pr[g_2 = 1] = 1/(1 + e^{-(W-b_1)})$.

620 The observation holds as the cumulative density function of the Logistic distribution is a
 621 Sigmoidal function. Since we already know how to implement a deterministic threshold
 622 gate using a Boolean circuit, the key challenge is in sampling a value from the Logistic
 623 distribution using a small number of uniformly random gates (i.e., fair coins). This is done
 624 in two key steps. First, using $O(\log 1/\epsilon)$ uniformly random gates, we sample a value from an
 625 ϵ^4 -discretization of the uniform distribution⁸. Then, we use the method of inverse transform
 626 sampling to sample from a distribution that is $\Theta(\epsilon)$ -close (in L_1 norm) to the Sigmoidal
 627 distribution. For a value r sampled u.a.r in $[0, 1]$, a sample from the Logistic distribution with
 628 mean b and scale 1 is given by $b + \ln(r/(1-r))$. To compute the expression $b + \ln(r/(1-r))$
 629 using a Boolean circuit, we approximate the $\ln(x)$ function using the first $O(\log 1/\epsilon)$ terms
 630 of the Taylor expansion. The almost-Logistic sample will serve as the bias of a deterministic
 631 threshold gate and will be fed to the Boolean circuit of Lemma 16. The full description is
 632 given in Appendix B.4.

633 We can then synchronize the Boolean Circuit as described in Lemma 17.

⁷ A uniformly random gate is a fair coin, in contrast to probabilistic threshold gate that outputs 1 based on a Sigmoidal distribution.

⁸ Our sample is equivalent to sampling a value from the uniform distribution and then rounding it to the closest value of the form $i \cdot \epsilon^4$ for some integer i .



■ **Figure 2** Left: synchronized network of a single NOT gate. Middle: A synchronized network for a Boolean circuit. Right: The transformation of a single neuron v_i in the synchronized network for the given SNN.

634 ► **Corollary 19.** *Given a probabilistic threshold gate g with Δ inputs, and an error parameter*
 635 *$\epsilon \in (0, 1)$, there exists a network $\mathcal{N} = \text{Sync}(g)$ with Δ input neurons x_1, \dots, x_Δ , an output*
 636 *neuron z , and $\text{poly}(\Delta, \log 1/\epsilon) \cdot L^3$ auxiliary neurons such that z approximates the gate*
 637 *g within $\text{poly}(\log \Delta, \log 1/\epsilon) \cdot L^4$ rounds in the following sense. For any fixed input \bar{x} ,*
 638 *with probability at least $1 - \Theta(\epsilon)$, it holds that g outputs 1 iff z fires in some round in*
 639 *$[1, \text{poly}(\log \Delta, \log 1/\epsilon) \cdot L^4]$.*

640 4.3 The Complete Synchronization Scheme

641 The complete synchronization scheme and the proof of Theorem 4 are given in Appendix B.5.
 642 In the high level, the construction has two parts: a global pulse generator, and a specific
 643 adaptation of the given network \mathcal{N} into a network $\text{sync}(\mathcal{N})$, see Figure 2.

644 The *pulse generator* is implemented by a directed cycle of length $k = \tilde{O}(L^4 \log \Delta)$. The
 645 input layer and output layer in $\text{sync}(\mathcal{N})$ are *exactly* as in \mathcal{N} . Let V be the neurons of \mathcal{N} . For
 646 each auxiliary neuron $v_i \in V$, we add its synchronized sub-network $\text{Sync}(v_i)$ from Lemma
 647 16 and Cor. 19. Recall that each neuron in \mathcal{N} implements either a threshold gate or a
 648 probabilistic threshold gate. For each such $v_i \in V$, we also add an AND module AND_i , which
 649 receives input from the sub-network $\text{Sync}(v_i)$ and the pulse generator. The neuron v_i is set
 650 to be the output neuron of this AND_i module.

651 — References —

- 652 1 Douglas B Armstrong, Arthur D Friedman, and Premachandran R Menon. Design of asynchron-
 653 ous circuits assuming unbounded gate delays. *IEEE Transactions on Computers*, 100(12):1110–
 654 1120, 1969.
- 655 2 Baruch Awerbuch and David Peleg. Network synchronization with polylogarithmic overhead.
 656 In *31st Annual Symposium on Foundations of Computer Science, St. Louis, Missouri, USA,*
 657 *October 22-24, 1990, Volume II*, pages 514–522, 1990.
- 658 3 Michael J Berry II and Markus Meister. Refractoriness and neural precision. In *Advances in*
 659 *Neural Information Processing Systems*, pages 110–116, 1998.
- 660 4 Tobias Bjerregaard and Shankar Mahadevan. A survey of research and practices of network-
 661 on-chip. *ACM Computing Surveys (CSUR)*, 38(1):1, 2006.
- 662 5 Sami Boudkkazi, Edmond Carlier, Norbert Ancri, Olivier Caillard, Pierre Giraud, Laure
 663 Fronzaroli-Molinieres, and Dominique Debanne. Release-dependent variations in synaptic
 664 latency: a putative code for short-and long-term synaptic dynamics. *Neuron*, 56(6):1048–1060,
 665 2007.

- 666 6 Chi-Ning Chou, Kai-Min Chung, and Chi-Jen Lu. On the algorithmic power of spiking neural
667 networks. In *10th Innovations in Theoretical Computer Science Conference, ITCS 2019,*
668 *January 10-12, 2019, San Diego, California, USA*, pages 26:1–26:20, 2019.
- 669 7 RE Lee DeVille and Charles S Peskin. Synchrony and asynchrony in a fully stochastic neural
670 network. *Bulletin of mathematical biology*, 70(6):1608–1633, 2008.
- 671 8 Huawei Fan, Yafeng Wang, Hengtong Wang, Ying-Cheng Lai, and Xingang Wang. Autapses
672 promote synchronization in neuronal networks. *Scientific reports*, 8(1):580, 2018.
- 673 9 Martin Fürer. Faster integer multiplication. *SIAM Journal on Computing*, 39(3):979–1005,
674 2009.
- 675 10 Johan Håstad. On the size of weights for threshold gates. *SIAM J. Discrete Math.*,
676 7(3):484–492, 1994. URL: <https://doi.org/10.1137/S0895480192235878>, doi:10.1137/
677 S0895480192235878.
- 678 11 Scott Hauck. Asynchronous design methodologies: An overview. *Proceedings of the IEEE*,
679 83(1):69–93, 1995.
- 680 12 Yael Hitron and Merav Parter. Counting to ten with two fingers: Compressed counting with
681 spiking neurons. *ESA*, 2019.
- 682 13 Yael Hitron and Merav Parter. Counting to ten with two fingers: Compressed counting with
683 spiking neurons. *CoRR*, abs/1902.10369, 2019. URL: <http://arxiv.org/abs/1902.10369>,
684 arXiv:1902.10369.
- 685 14 Kaori Ikeda and John M Bekkers. Autapses. *Current Biology*, 16(9):R308, 2006.
- 686 15 Fabian Kuhn, Joel Spencer, Konstantinos Panagiotou, and Angelika Steger. Synchrony
687 and asynchrony in neural networks. In *Proceedings of the twenty-first annual ACM-SIAM*
688 *symposium on Discrete algorithms*, pages 949–964. SIAM, 2010.
- 689 16 Robert A. Legenstein, Wolfgang Maass, Christos H. Papadimitriou, and Santosh Srinivas
690 Vempala. Long term memory and the densest k-subgraph problem. In *9th Innovations in*
691 *Theoretical Computer Science Conference, ITCS 2018, January 11-14, 2018, Cambridge, MA,*
692 *USA*, pages 57:1–57:15, 2018.
- 693 17 Nancy Lynch and Cameron Musco. A basic compositional model for spiking neural networks.
694 *arXiv preprint arXiv:1808.03884*, 2018.
- 695 18 Nancy Lynch, Cameron Musco, and Merav Parter. Computational tradeoffs in biological
696 neural networks: Self-stabilizing winner-take-all networks. In *Proceedings of the 8th Conference*
697 *on Innovations in Theoretical Computer Science (ITCS)*, 2017.
- 698 19 Nancy Lynch, Cameron Musco, and Merav Parter. Spiking neural networks: An algorithmic
699 perspective. In *5th Workshop on Biological Distributed Algorithms (BDA 2017)*, July 2017.
- 700 20 Nancy A. Lynch, Cameron Musco, and Merav Parter. Neuro-ram unit with applications to
701 similarity testing and compression in spiking neural networks. In *31st International Symposium*
702 *on Distributed Computing, DISC 2017, October 16-20, 2017, Vienna, Austria*, pages 33:1–33:16,
703 2017.
- 704 21 Jun Ma, Xinlin Song, Wuyin Jin, and Chuni Wang. Autapse-induced synchronization in a
705 coupled neuronal network. *Chaos, Solitons & Fractals*, 80:31–38, 2015.
- 706 22 Wolfgang Maass. Lower bounds for the computational power of networks of spiking neurons.
707 *Electronic Colloquium on Computational Complexity (ECCC)*, 1(19), 1994. URL: <http://eccc.hpi-web.de/eccc-reports/1994/TR94-019/index.html>.
- 708 23 Wolfgang Maass. On the computational power of noisy spiking neurons. In *Advances in Neural*
709 *Information Processing Systems 8 (NIPS)*, 1996.
- 710 24 Wolfgang Maass. Networks of spiking neurons: the third generation of neural network models.
711 *Neural Networks*, 10(9):1659–1671, 1997.
- 712 25 Wolfgang Maass. Paradigms for computing with spiking neurons. In *Models of Neural Networks*
713 *IV*, pages 373–402. Springer, 2002.
- 714 26 Rajit Manohar and Yoram Moses. Analyzing isochronic forks with potential causality. In
715 *21st IEEE International Symposium on Asynchronous Circuits and Systems, ASYNC 2015*,
716

- 717 *Mountain View, CA, USA, May 4-6, 2015*, pages 69–76, 2015. URL: <https://doi.org/10.1109/ASYNC.2015.19>, doi:10.1109/ASYNC.2015.19.
- 718
- 719 **27** Rajit Manohar and Yoram Moses. The eventual c-element theorem for delay-insensitive
720 asynchronous circuits. In *2017 23rd IEEE International Symposium on Asynchronous Circuits
721 and Systems (ASYNC)*, pages 102–109. IEEE, 2017.
- 722 **28** Alain J Martin. The limitations to delay-insensitivity in asynchronous circuits. In *Beauty is
723 our business*, pages 302–311. Springer, 1990.
- 724 **29** Robert Miller. Time and the brain. *CRC Press*, 2000.
- 725 **30** Saburo Muroga. *Threshold logic and its applications*. Wiley, 1971.
- 726 **31** Yu P Ofman. On the algorithmic complexity of discrete functions. In *Doklady Akademii Nauk*,
727 volume 145, pages 48–51. Russian Academy of Sciences, 1962.
- 728 **32** Christos H. Papadimitriou and Santosh S. Vempala. Random projection in the brain and
729 computation with assemblies of neurons. In *10th Innovations in Theoretical Computer Science
730 Conference, ITCS 2019, January 10-12, 2019, San Diego, California, USA*, pages 57:1–57:19,
731 2019. URL: <https://doi.org/10.4230/LIPIcs.ITCS.2019.57>, doi:10.4230/LIPIcs.ITCS.
732 2019.57.
- 733 **33** Huixin Qin, Jun Ma, Chunni Wang, and Ying Wu. Autapse-induced spiral wave in network of
734 neurons under noise. *PloS one*, 9(6):e100849, 2014.
- 735 **34** Alexa Riehle, Sonja Grün, Markus Diesmann, and Ad Aertsen. Spike synchronization and rate
736 modulation differentially involved in motor cortical function. *Science*, 278(5345):1950–1953,
737 1997.
- 738 **35** BL Sabatini and WG Regehr. Timing of synaptic transmission. *Annual review of physiology*,
739 61(1):521–542, 1999.
- 740 **36** Jens Sparsø. Asynchronous circuit design—a tutorial. In *Chapters 1-8 in "Principles of
741 asynchronous circuit design-A systems Perspective"*. Kluwer Academic Publishers, 2001.
- 742 **37** Lili Su, Chia-Jung Chang, and Nancy Lynch. Spike-based winner-take-all computation:
743 Fundamental limits and order-optimal circuits. *Neural Computation*, 2019.
- 744 **38** Christopher S. Wallace. A suggestion for a fast multiplier. *IEEE Trans. Electronic Computers*,
745 13(1):14–17, 1964. URL: <https://doi.org/10.1109/PGEC.1964.263830>, doi:10.1109/PGEC.
746 1964.263830.
- 747 **39** Barbeeba Wang and Nancy Lynch. Integrating temporal information to spatial information in
748 a neural circuit. *arXiv preprint arXiv:1903.01217*, 2019.
- 749 **40** Ergin Yilmaz, Mahmut Ozer, Veli Baysal, and Matjaž Perc. Autapse-induced multiple
750 coherence resonance in single neurons and neuronal networks. *Scientific Reports*, 6:30914,
751 2016.

752 **A** Missing Proofs for the Negative Results

753 **Impossibility Result, Proof of Lemma 10:** Assume towards contradiction there exists
754 a network $\mathcal{N} = (V, \{x, y\}, \{z\}, w, b)$ that computes the AND gate of the initial states x_0, y_0
755 of the inputs x and y . It is then required that if $x_0 \wedge y_0 = 1$, then there exists a round
756 in which z fires, and if $x_0 \wedge y_0 = 0$ then z is idle throughout the execution. Our goal is
757 to show the existence of a bad assignment of latency values to the edges of \mathcal{N} . Such bad
758 assignment exists even if we fix the latency of each edge to the same value throughout the
759 entire execution. We begin with some quick observations.

- 760 ■ The state of a neuron v in round τ , namely, $\sigma_\tau(v)$, is fully determined by the network, the
761 latency function, the input and the initial state σ_0 , that is $\sigma_\tau(v) = H(\mathcal{N}, \ell, \sigma_0, x_0, y_0, v, \tau)$
762 for some function H .
- Given the network \mathcal{N} and a latency function ℓ , the state of a neuron v in round τ is a
function of the *previous* states of its incoming neighbors denoted as $u_1 \dots u_k$:

$$\sigma_\tau(v) = F_v(\sigma_{\tau-\ell(u_1,v)}(u_1), \dots, \sigma_{\tau-\ell(u_k,v)}(u_k)).$$

► **Definition 20.** For a neuron v and round τ , we say that the state $\sigma_\tau(v)$ is x -independent (equivalently y -independent) if its value does not depend on the initial state of x , i.e. if

$$H(\mathcal{N}, \ell, \sigma_0, x_0 = 0, y_0, v, \tau) = H(\mathcal{N}, \ell, \sigma_0, x_0 = 1, y_0, v, \tau) .$$

763 ▷ **Observation 21.** A concatenation of x -independent functions is also x -independent.
 764 Specifically, for round τ and neuron v with incoming neighbors u_1, \dots, u_k , it holds that if
 765 $\sigma_{\tau-\ell(u_1,v)}(u_1), \dots, \sigma_{\tau-\ell(u_k,v)}(u_k)$ are x -independent then $\sigma_\tau(v)$ is also x -independent.

766 Given the network \mathcal{N} we set the edge latencies as follows:

$$767 \quad \ell(u, v) = \begin{cases} 1 & \text{if either } u = x \text{ or } v = x, \text{ but not both.} \\ 2 & \text{otherwise.} \end{cases}$$

768 We next show that for every neuron $v \in V$, its firing state in each round is either x -independent
 769 or y -independent. Specifically, the firing state of z in each round does not depend on both
 770 x_0 and y_0 . This will contradict the assumption that z computes an AND gate of x_0 and y_0 .

771 ▷ **Claim 22.** For every round $\tau \geq 1$ it holds that: (1) For every $v \in V \setminus \{x\}$, the firing state
 772 $\sigma_\tau(v)$ is x -independent if τ is even and y -independent if τ is odd. (2) $\sigma_\tau(x)$ is x -independent
 773 if τ is odd and y -independent if τ is even.

Proof. By induction on the round τ . For $\tau = 1$, since all outgoing edges from y have latency 2 (except for the edge (y, x) , if exists), in round 1 no neuron $v \in V \setminus \{x\}$ received a spike from y and therefore $\sigma_1(v)$ is y -independent. Because the edge from x to itself has latency 2, and the edge from y to x has latency 1, in round 1 the neuron x can receive a signal from y but not from x and therefore also $\sigma_1(x)$ is x -independent. For $\tau = 2$ and $v \neq x$, since the edges from x have latencies 1, and all other edges have latency 2, it holds that

$$\sigma_2(v) = F_v(\sigma_1(x), \sigma_0(y), \sigma_0(u_1), \dots, \sigma_0(u_k)),$$

where u_1, \dots, u_k are the neighbors of v in $V \setminus \{x\}$. Because $\sigma_1(x)$ is x -independent, and $\sigma_0(u_i)$ are the initial states, by Observation 21 we can conclude that $\sigma_2(v)$ is x -independent. Next, for the input neuron x , since all its incoming edges (except for the self-loop) have latency 1 it holds that

$$\sigma_2(x) = F_x(\sigma_0(x), \sigma_1(u_1), \dots, \sigma_1(u_k)).$$

774 Because $\sigma_1(v)$ is y -independent for all $v \neq x$, we conclude that $\sigma_2(x)$ is y -independent as
 775 well.

Assume the claim holds for every round $\tau' < \tau$ and we will show the claim holds for round τ as well. For $v \neq x$ with incoming neighbors u_1, \dots, u_k in $V \setminus \{x\}$, by the definition of the latencies it holds that

$$\sigma_\tau(v) = F_v(\sigma_{\tau-1}(x), \sigma_{\tau-2}(u_1), \dots, \sigma_{\tau-2}(u_k)).$$

If τ is even, so is $\tau - 2$ and by the induction assumption $\sigma_{\tau-2}(u_i)$ are x -independent. Because $\tau - 1$ is odd, $\sigma_{\tau-1}(x)$ is x -independent. Hence, by Observation 21 we conclude that $\sigma_\tau(v)$ is also x -independent. Similarly, if τ is odd, then by the induction assumption $\sigma_{\tau-2}(u_i)$ and $\sigma_{\tau-1}(x)$ are y -independent, and therefore $\sigma_\tau(v)$ is also y -independent. Then, for the neuron x , it holds that

$$\sigma_\tau(x) = F_x(\sigma_{\tau-2}(x), \sigma_{\tau-1}(u_1), \dots, \sigma_{\tau-1}(u_k)).$$

776 If τ is odd, by the induction assumption $\sigma_{\tau-2}(x)$ as well as $\sigma_{\tau-1}(u_1), \dots, \sigma_{\tau-1}(u_k)$ are
 777 x -independent and therefore $\sigma_\tau(x)$ is x -independent. On the other hand, if x is even, then
 778 by the induction assumption $\sigma_{\tau-2}(x)$ and $\sigma_{\tau-1}(u_1), \dots, \sigma_{\tau-1}(u_k)$ are y -independent and
 779 therefore $\sigma_\tau(x)$ is also y -independent. \blacktriangleleft

780 Since in each round the output neuron z is either x -independent or y -independent, this
 781 contradicts the assumption and Lemma 10 follows.

782 A.1 Size Lower Bound for Computing $NOT(x)$

783 A.1.1 Reduction to Simple Networks, Proof of Lemma 12:

784 The reduction is based on the following notion of domination between two configurations.

785 **Domination.** Given a network \mathcal{N} , a latency function ℓ , and a vector of starting states $\bar{\sigma}$
 786 for all neurons but the input x , for $b \in \{0, 1\}$ define $\text{pot}_b(u, \tau, \mathcal{N}, \ell, \bar{\sigma})$ as the potential of
 787 neuron u in round τ in the simulation of \mathcal{N} with the initial vector state $[b, \bar{\sigma}]$, i.e., with the
 788 initial state of x is being b and all other initial states are as in $\bar{\sigma}$. When $\bar{\sigma}$ is clear from
 789 the context, we may omit it, and simply write $\text{pot}_b(u, \tau, \mathcal{N}, \ell)$. Given networks $\mathcal{N}_1, \mathcal{N}_2$ with
 790 vertices V_1, V_2 and latency functions ℓ_1, ℓ_2 respectively, we say that $\langle \mathcal{N}_1, \bar{\sigma}_1 \rangle$ and $\langle \mathcal{N}_2, \bar{\sigma}_2 \rangle$
 791 are *compatible* if $V_1 \subseteq V_2$ and σ_1 and σ_2 agree on the mutual vertices of V_1 , i.e., $\sigma_1(u) = \sigma_2(u)$
 792 for every $u \in V_1$.
 793

794 In our arguments, we consider a pair of compatible configurations $\langle \mathcal{N}_1, \bar{\sigma}_1 \rangle$ and $\langle \mathcal{N}_2, \bar{\sigma}_2 \rangle$
 795 along with latency functions ℓ_1, ℓ_2 for these configurations. We say that $\langle \mathcal{N}_1, \bar{\sigma}_1, \ell_1 \rangle$ *dominates*
 796 $\langle \mathcal{N}_2, \bar{\sigma}_2, \ell_2 \rangle$ if $\langle \mathcal{N}_1, \bar{\sigma}_1 \rangle$ and $\langle \mathcal{N}_2, \bar{\sigma}_2 \rangle$ are compatible and in addition:

- 797 ■ $\text{pot}_1(u, \tau, \mathcal{N}_1, \ell_1, \bar{\sigma}_1) \leq \text{pot}_1(u, \tau, \mathcal{N}_2, \ell_2, \bar{\sigma}_2)$ for every $u \in V_1 \setminus \{x\}$ and $\tau \geq 0$.
- 798 ■ $\text{pot}_0(u, \tau, \mathcal{N}_1, \ell_1, \bar{\sigma}_1) \geq \text{pot}_0(u, \tau, \mathcal{N}_2, \ell_2, \bar{\sigma}_2)$ for every $u \in V_1 \setminus \{x\}$ and $\tau \geq 0$.

Let V, V_{simple} be the vertex sets of \mathcal{N} and $\mathcal{N}_{\text{simple}}$ respectively. Let $\bar{\sigma}_{\text{simple}}$ be the initial
 state vector that agrees with $\bar{\sigma}$ on all vertices in V_{simple} . Thus, $\langle \mathcal{N}_{\text{simple}}, \bar{\sigma}_{\text{simple}} \rangle$ and $\langle \mathcal{N}, \bar{\sigma} \rangle$
 are compatible. For a number of rounds r , a latency function ℓ is *r -good* for $\langle \mathcal{N}, \bar{\sigma} \rangle$ if \mathcal{N}
 computes $NOT(x)$ within r rounds under ℓ when starting from the initial state vector $\bar{\sigma}$.
 Our proof strategy is as follows. We will show that every latency function ℓ_{simple} is r -good
 for $\langle \mathcal{N}_{\text{simple}}, \bar{\sigma}_{\text{simple}} \rangle$, by showing that there exists a function ℓ such that

$$\langle \mathcal{N}_{\text{simple}}, \bar{\sigma}_{\text{simple}}, \ell_{\text{simple}} \rangle \text{ dominates } \langle \mathcal{N}, \bar{\sigma}, \ell \rangle.$$

799 Given a latency function ℓ_{simple} for the network $\mathcal{N}_{\text{simple}}$, let ℓ be a latency function for \mathcal{N}
 800 which is similar on excitatory neurons and gives inhibitory neurons the latency value of the
 801 neuron x . I.e., $\ell(v, u, \tau) = \ell_{\text{simple}}(v, u, \tau)$ for every pair of excitatory neurons. In addition,
 802 $\ell(v', u, \tau) = \ell_{\text{simple}}(x, u, \tau)$ for every inhibitory neuron v' , and a neuron $u \in V_{\text{simple}}$. All
 803 remaining latency values (i.e., the incoming edges to the inhibitors of \mathcal{N}) can be chosen
 804 arbitrarily.

805 We show by induction on the round τ , that (i) $\text{pot}_1(u, \tau, \mathcal{N}_{\text{simple}}, \ell_{\text{simple}}) \leq \text{pot}_1(u, \tau, \mathcal{N}, \ell)$
 806 for every $u \in V_{\text{simple}} \setminus \{x\}$, $\tau \geq 0$, and (ii) $\text{pot}_0(u, \tau, \mathcal{N}_{\text{simple}}, \ell_{\text{simple}}) \geq \text{pot}_0(u, \tau, \mathcal{N}, \ell)$ for
 807 every $u \in V_{\text{simple}} \setminus \{x\}$ and $\tau \geq 0$. For $\tau = 0$, this is true as the potential values in $\tau = 0$ are
 808 simply the initial states, and the vector of initial states of $\bar{\sigma}$ and $\bar{\sigma}_{\text{simple}}$ are compatible. For
 809 the induction step, let $\tau \geq 1$, and assume correctness for all $\tau' \leq \tau - 1$. We will prove the
 810 claims for round τ . Let u be a neuron in $V_{\text{simple}} \setminus \{x\}$, and let v_1, \dots, v_k be its incoming
 811 excitatory neighbors.

813 **The initial state of x is 0:** By the induction assumption, it holds that
 814 $\text{pot}_0(u, \tau', \mathcal{N}_{simple}, \ell_{simple}) \geq \text{pot}_0(u, \tau', \mathcal{N}, \ell)$ for every round $\tau' \leq \tau - 1$ and every neuron
 815 $v \in \{v_1, \dots, v_k\}$. Thus every excitatory neuron v_i that fires in round τ' in the simulation of
 816 \mathcal{N} , also fires in round τ' in the simulation of \mathcal{N}_{simple} for every $\tau' \leq \tau - 1$. Combining with the
 817 definition of the latency function ℓ , we get that each spike from v_i that arrives to u at round
 818 τ of the simulation of \mathcal{N} also arrives u in round τ in the simulation of \mathcal{N}_{simple} . Let ω be an
 819 inhibitory incoming neighbor of u in \mathcal{N} , then ω does not exist in \mathcal{N}_{simple} . Also note that since
 820 $\sigma_0(x) = 0$, a negative spike from x never arrives at u in the simulation of network \mathcal{N}_{simple} .
 821 Therefore, no negative spikes arrive at u in \mathcal{N}_{simple} . Summing over the positive and negative
 822 spike weights, we get that $\text{pot}_0(u, \tau, \mathcal{N}_{simple}, \ell_{simple}) \geq \text{pot}_0(u, \tau, \mathcal{N}, \ell)$ for every $u \in V_1 \setminus \{x\}$.

823
 824 **The initial state of x is 1:** By the induction assumption it holds that
 825 $\text{pot}_1(u, \tau', \mathcal{N}_{simple}, \ell_{simple}) \leq \text{pot}_1(u, \tau', \mathcal{N}, \ell)$ for every round $\tau' \leq \tau - 1$ and every $u \in$
 826 $\{v_1, \dots, v_k\}$. Thus every v_i that fires in round τ' in the simulation of \mathcal{N}_{simple} , also fires in
 827 round τ' in the simulation of \mathcal{N} . Combining with the definition of the latency function ℓ ,
 828 we get that each spike from v_i that arrives at u in round τ of the simulation of \mathcal{N}_{simple} also
 829 arrives at u in round τ in the simulation of \mathcal{N} .

830 Let ω be an inhibitory incoming neighbor of u in \mathcal{N} . By the definition of the latency
 831 function ℓ , and the fact that x fires in every round, for each spike that ω fires and arrives at
 832 u in round τ in \mathcal{N} , there is a spike from x that arrives at u in round τ in \mathcal{N}_{simple} . Since the
 833 edges from x have weight $-\infty$, we get that the negative spikes weight arriving at u in round
 834 τ in \mathcal{N}_{simple} are larger (in absolute value) than the negative spikes in \mathcal{N} . Thus, summing
 835 up the both positive and negative spike weights, we get that $\text{pot}_1(u, \tau, \mathcal{N}_{simple}, \ell_{simple}) \leq$
 836 $\text{pot}_1(u, \tau, \mathcal{N}, \ell)$ for every $u \in V_{simple} \setminus \{x\}$. This proves the induction step for round τ . We
 837 get that $\langle \mathcal{N}_{simple}, \bar{\sigma}_{simple}, \ell_{simple} \rangle$ dominates $\langle \mathcal{N}, \bar{\sigma}, \ell \rangle$.

838 Finally, we show that if $\langle \mathcal{N}_1, \bar{\sigma}_1, \ell_1 \rangle$ dominates $\langle \mathcal{N}_2, \bar{\sigma}_2, \ell_2 \rangle$ and ℓ_2 is r -good for $\langle \mathcal{N}_2, \bar{\sigma}_2 \rangle$,
 839 then also ℓ_1 is r -good for $\langle \mathcal{N}_1, \bar{\sigma}_1 \rangle$.

840 Consider the simulation of $\langle \mathcal{N}_2, \bar{\sigma}_2, \ell_2 \rangle$, and assume that the initial state of x is 0. Since ℓ_2
 841 is r -good for \mathcal{N}_2 , there is a round $\tau \leq r$ in which z fires in \mathcal{N}_2 . Since $\langle \mathcal{N}_1, \bar{\sigma}_1, \ell_1 \rangle$ dominates
 842 $\langle \mathcal{N}_2, \bar{\sigma}_2, \ell_2 \rangle$, we can apply the condition of domination for z, τ and get that z also fires in
 843 round τ in \mathcal{N}_1 . Now, assume that the initial state of x is 1. Since ℓ_2 is r -good for \mathcal{N}_2 , there
 844 is *no* round τ in which z fires in \mathcal{N}_2 . Since $\langle \mathcal{N}_1, \bar{\sigma}_1, \ell_1 \rangle$ dominates $\langle \mathcal{N}_2, \bar{\sigma}_2, \ell_2 \rangle$, we can apply
 845 the condition of domination for z and every τ , and get that z also never fires in \mathcal{N}_1 . Hence,
 846 ℓ_1 is r -good for \mathcal{N}_1 . This completes the proof of the lemma.

847 A.1.2 Size Lower Bound for Simple Networks

848 We show that the latency values defined in the step i satisfy the invariant in the beginning
 849 of step $i + 1$.

850
 851 **Proving that the Invariants Hold:** For round $i = 0$, the correctness of invariant (I1)
 852 hold since in the first round $\tau = -1$ all positive spikes are set to arrive in round 0 in both
 853 ℓ_0, ℓ_1 , while a spike from x arrives at round $\tau = L - 1$. As for the correctness of invariant
 854 (I2), note that both simulations are similar for all neurons $V \setminus \{x\}$, and, again, a spike from
 855 x arrives only in round $L - 1$. Therefore the same neurons are active in round $\tau = 0$. Hence
 856 $A_{0,0} = A_{0,1}$ and we get correctness for (I2). We now show that the invariants are preserved
 857 after each step. Assume that the correctness holds at the beginning of each step $j \leq i$, and
 858 consider now the beginning of step $i + 1$.

23:22 The Computational Cost of Asynchronous Neural Communication

- 859 ■ (I1) By the construction, in all of the cases, the values we set for ℓ_0, ℓ_1 in step i are such
 860 that all spikes except the spike from x generated at round $iL + (L - 1)$ which are generated
 861 in T_i arrive at some round $\tau \leq (i + 1)L$, i.e. by the first round of T_{i+1} . Furthermore,
 862 spikes from x generated at round $iL + (L - 1)$ is set to arrive in round $(i + 1)L + (L - 1)$.
 863 The invariant holds by combining with the correctness for all steps $i' \leq i$.
 864 ■ (I2) We start by proving the following auxiliary claim.

865 ▷ **Claim 23.** Consider the simulation of \mathcal{N} with initial state $\bar{\sigma}_b$ and latency function ℓ_b ,
 866 and let τ be round in T_i . Then:

- 867 1. For a strong neuron $u \in A_{i,1}$, u fires iff $\tau \in [iL, iL + (L - 2)] \subseteq T_i$. For a strong neuron
 868 $u \in A_{i,0}$, u fires for every $\tau \in T_i$.
 869 2. For a weak neuron $u \in A_{i,b}$, u fires iff $\tau = iL$.
 870 3. For $u \notin A_{i,b}$, u is not active in round τ .

871 **Proof. Case $b = 1$:** We start by showing that all three claims hold for $b = 1$. By
 872 the definition of ℓ_1 , the only positive spikes received by any neuron in some round
 873 $\tau \in [iL + 1, iL + (L - 1)]$ are self-loop spikes. Since a strong active neuron $u \in A_{i,1}$
 874 receives an inhibiting spike from x in round $iL + (L - 1)$, it is not active in this last round.
 875 For a weak neuron u' , its spike from the self-loop is not strong enough to make it active.
 876 Lastly, for $u \notin A_{i,1}$ since no negative spikes arrive at u in round iL we have that $b(u) > 0$.
 877 Due to the fact that no spikes arrive at u in round τ , we get that u stays inactive.

878 **Case $b = 0$:** claim (1) holds since a strong $u \in A_{i,0}$ never gets inhibited as x never
 879 fires. We will now consider claim (2) and (3) for a neuron u that is either a *weak*
 880 neuron, or a strong neuron that is *not* in $A_{i,0}$. By Invariant (I1), all the positive spikes
 881 from the previous blocks arrived by the first round of T_i . Thus if u fires in any round
 882 $\tau \in [iL + 1, iL + (L - 1)]$ (i.e., any round which is not the first one in T_i), this must be
 883 due to the incoming spikes generated in the first round of T_i . We will now prove by
 884 induction on the round τ that u does not fire in any round in $[iL + 1, iL + (L - 1)]$.

885 **Induction Base, Round $\tau = iL + 1$:** By the definition of the latency function ℓ_0 , no
 886 spike arrives at u from an incoming neighbor in that round. Therefore, if u is weak, then
 887 a spike from itself will not make it active in round τ . In addition, if $u \notin A_{i,0}$ then u did
 888 not fire in round iL , and thus receives no self-spike in round τ . Since no negative spikes
 889 arrive at u in round iL , for every $u \notin A_{i,0}$, it must hold that $b(u) > 0$.

890 **Induction Step $\tau \geq iL + 2$.** Assume that the claims (2,3) hold up to round $\tau - 1$ and
 891 consider round $\tau \geq iL + 2$. Consider first the case that u has either a weak neighbor g_j
 892 with $w(g_j, u) \geq b(u)$, or alternatively a strong neighbor h_j with $w(h_j, u) \geq b(u)/(L - 1)$.
 893 Then by the definition of ℓ_0 (Case I in our definition), all the spikes fired by the incoming
 894 neighbors of u are scheduled to arrive in the first round of T_{i+1} . Therefore, u does not
 895 receive any spike in round τ , and remains inactive.

896 Next, consider the complimentary case where all the weak neighbors g_j satisfy that
 897 $w(g_j, u) < b(u)$, and all the strong neighbors h_j satisfy $w(h_j, u) < b(u)/(L - 1)$.

898 **Case (1):** $\tau \in [iL + 2, iL + \omega + 1]$. By the induction assumption on $\tau - 1$, u did not fire
 899 in round $\tau - 1$. By the definition of ℓ_0 , u receives a spike from at most one weak neighbor
 900 g_j in round τ , and since $w(g_j, u) < b(u)$, it does not fire in this round.

901 **Case (2):** $\tau \in [iL + (\omega + 2), iL + (L - 1)]$. Let h_j be a strong active neighbor of u . By the
 902 definition of ℓ_0 , in round τ , u receives at most $(\omega + 1)/(L - (\omega + 1))$ of the spikes that fired
 903 by h_j during the interval $[iL, iL + \omega]$. Furthermore, there is one additional spike that h_j
 904 fired at round $\tau - 1$, that arrives at u in round τ . Note that u does not receive spikes from
 905 weak neighbors in round τ since all spikes from weak neighbors arrive in an earlier round

906 $\tau'' \in [iL + 2, iL + (\omega + 1)]$. In addition, since u did not fire in round $\tau - 1$, it also does
 907 not get any self spikes in round τ . Overall, u receives at most $((\omega + 1)/(L - (\omega + 1))) + 1$
 908 spikes by strong neighbors in round τ , and no other spikes (by a weak neighbor or by
 909 u). Since the spikes from the strong neighbors have weight of at most $b(u)/(L - 1)$, and
 910 there are s strong active neighbors, the overall weighted sum of the received spikes at
 911 round τ is at most

$$912 \quad s \cdot ((\omega + 1)/(L - (\omega + 1)) + 1) \cdot \frac{b(u)}{L - 1} <$$

$$913 \quad s \cdot ((\omega + 1)/s + 1) \frac{b(u)}{L - 1} = (s + \omega + 1) \frac{b(u)}{L - 1} < b(u), \quad (2)$$

914 where both inequalities follow as $s + \omega < L - 2$. Therefore, u does not get activated at
 915 round τ , claims (2)+(3) follow. ◀

916 We are now ready to prove the induction step for (I2). Assume towards contradiction
 917 that there exists a neuron $u \in A_{i+1,0} \setminus \bigcup_{i' \leq i+1} A_{i',1}$. Since u is active in the first round
 918 of T_{i+1} , using Claim 23(3), it must have an incoming neighbor that fires in the *first round*
 919 of the previous block. Let $A_{i,0}(u) = \Gamma_{in}(u) \cap A_{i,0}$ be those neighbors.

920 \triangleright **Claim 24.** For every strong neuron $v \in A_{i,0}(u)$, it holds that $w(v, u) < b(u)/(L - 1)$.

921 **Proof.** Consider such strong $v \in A_{i,0}(u)$. By Invariant (I2) for the beginning of step
 922 i , there exists a round $j \leq i$ such that $v \in A_{j,1}$. When running \mathcal{N} with initial state
 923 $\bar{\sigma}_1$ and the latency function ℓ_1 , by Claim 23(1), v fires in all of the $L - 1$ rounds of
 924 $[jL, jL + (L - 2)]$. By the construction of ℓ_1 , all these $L - 1$ spikes arrive in round $(j + 1)L$.
 925 Therefore, if $(L - 1) \cdot w(v, u) \geq b(u)$, then u is activated in round $(j + 1)L$, i.e., $u \in A_{j+1,1}$,
 926 in contradiction to the definition of u . Therefore $(L - 1) \cdot w(v, u) < b(u)$ for every strong
 927 neuron v . ◀

928 \triangleright **Claim 25.** For every weak neuron $v \in A_{i,0}(u)$, it holds that $w(v, u) < b(u)$.

929 **Proof.** Consider such weak neuron $v \in A_{i,0}(u)$. When running \mathcal{N} with initial state $\bar{\sigma}_1$
 930 and latency function ℓ_1 , by Claim 23(2), v fires once in the interval T_j , i.e., in the first
 931 round jL . By the construction of ℓ_1 this spike arrives in round $(j + 1)L$. Therefore,
 932 if $w(v, u) \geq b(u)$, then u is activated in round $(j + 1)L$, implying that $u \in A_{j+1,1}$,
 933 contradiction to the definition of u . We therefore conclude that $w(v, u) < b(u)$ for every
 934 weak neuron $v \in A_{i,0}(u)$. ◀

935 Let ω, s be the number of weak (resp., strong) neurons in $A_{i,0}(u)$. By Claims 25 and
 936 23(2), all active weak neurons in T_i fire only in the first round of that block. By the
 937 definition of the latency function, all these spikes are scheduled to the first ω rounds in
 938 T_i , and therefore none of them is scheduled to arrive on the first round of T_{i+1} . This
 939 implies that u fires in that round due to the spikes generated by its strong neighbors.
 940 By Claim 24, $w(v, u) < b(u)/(L - 1)$ for each such strong neighbor v of u . This in
 941 particular implies that u is a weak neuron, and by Claim. 23 it did not fire in the last
 942 round of T_i . By the definition of the latency function, the spikes generated by such strong
 943 neighbors are divided almost evenly among $L - \omega$ rounds, up to the first round of T_{i+1} .
 944 Each round gets at most $s \cdot (\omega/(L - \omega) + 1)$, which is strictly less than $b(u)$ by Eq. (2).
 945 Leading to contradiction for the assumption that $u \in A_{i+1,0}$.

946 Since ℓ_1 is a good latency function when starting with $x = 1$, we have that z never fires
 947 and thus $z \notin \bigcup A_{i,1}$. By applying invariant (I2) on the output neuron z for every round
 948 $\tau \geq 0$, we get that $z \notin \bigcup A_{i,0}$. By using Claim 23, we get that z never fires with ℓ_0 and
 949 $x = 0$. Contradiction to the fact that \mathcal{N} solves $NOT(x)$.

950 A.2 Time Lower Bound for Computing $NOT(x)$

951 In this section we show the following.

952 ► **Lemma 26.** *Every network \mathcal{N} that computes $NOT(x)$ in the L -bounded asynchronous*
 953 *setting requires $\Omega(L^3)$ rounds.*

954 By Lemma 12, we restrict attention to a simple network $\mathcal{N} = \mathcal{N}_{simple}$ with one input neuron
 955 x that computes $NOT(x)$. Similarly to the size lower bound, we define two *conflicting*
 956 latency functions ℓ_0 and ℓ_1 , such that if ℓ_1 is good when $x_0 = 1$, then the output neuron z of
 957 \mathcal{N} fires after $\Omega(L^3)$ rounds in the simulation with the latency function ℓ_0 and $x_0 = 0$.

- 958 ■ The simulation with the latency function ℓ_0 is partitioned into consecutive blocks of L
 959 rounds, $T_i = [iL, iL + (L - 1)]$ for every $i \in \mathbb{N}$.
- 960 ■ The simulation with the latency function ℓ_1 is based on the notion of *important* and
 961 *unimportant* rounds. Consider the L -round interval $T_k = [k \cdot L, k \cdot L + (L - 1)]$ for $k \in \mathbb{N}$.
 962 Among the first $L/2$ rounds, there is an important round once every 16 rounds, and the rest
 963 are unimportant. Furthermore, each of the last $L/2$ rounds of the interval are unimportant.
 964 I.e., the important rounds in the interval are $\{kL + 16j \mid 16j < L/2, j \in \mathbb{N}\}$. Denote by
 965 τ_i the i^{th} important round in the simulation. Note that by definition $\tau_{i+1} - \tau_i \leq L/2$.
 966 In our arguments, the configuration of the network in the i^{th} important round τ_i of the
 967 simulation with ℓ_1 and $x_0 = 1$ will be compared against the configuration in round iL
 968 (i.e., the first round of the block T_i) in the simulation with ℓ_0 and $x_0 = 0$.
- 969 ■ Active subsets of neurons: For every $i \in \mathbb{N}$, let $A_{0,i}$ be the firing neurons (hence *active*)
 970 of round $i \cdot L$ (the first round of the block T_i) in the simulation of $\langle \mathcal{N}, \sigma_0, \ell_0 \rangle$. Similarly,
 971 let $A_{1,i}$ be the firing neurons in round τ_i of the simulation of $\langle \mathcal{N}, \sigma_1, \ell_1 \rangle$. Also define
 972 $A'_{b,i} = A_{b,i} \setminus \bigcup_{j \leq i-1} A_{b,j}$, the neurons that fire for the first time in "round" i .
- 973 ■ For every neuron u , $b \in \{0, 1\}$ and $i \in \mathbb{N}$, let $A_{b,i}(u) = A_{b,i} \cap N_{in}(u)$, $A'_{b,i}(u) = A'_{b,i} \cap N_{in}(u)$
 974 where $N_{in}(u)$ is the set of incoming neighbors of u .
- 975 ■ For a subset of neurons $V' \subseteq V$ and a neuron u , let $w(V', u) = \sum_{v \in V'} w(v, u)$. Moreover,
 976 let $\mathbf{S}(V')$ and $\mathbf{W}(V')$ be the strong⁹ and weak (respectively) neurons subsets of V' .

977 A.2.1 Defining the latency functions ℓ_0 and ℓ_1

978 Throughout, a spike event is represented by a triplet $\langle v, u, \tau \rangle$ where $v \in N_{in}(u)$ fires in round
 979 τ . Since the functions are nice, the latency values for the self spikes $\langle u, u, \tau \rangle$ for every u
 980 and τ are set to 1. For technical reasons, it is more convenient to start the simulations in
 981 round -1 , rather than in round 0. For this first round -1 , let $\ell_b(\langle v, u \rangle, -1) = 1$ for every u
 982 and every $v \neq x$, and $\ell_b(\langle x, u \rangle, -1) = L$ for every u and $b \in \{0, 1\}$. As a result, the positive
 983 spikes (by any $v \neq x$) fired in round -1 arrive to their destination in round 0, and the negat-
 984 ive spikes of x arrive in round $L-1$. We now define the latency values for the remaining spikes.

985

⁹ Recall that a neuron u is strong if $w(u, u) \geq b(u)$ and it is weak otherwise.

986 **Defining the function ℓ_0 .** Note that when $x_0 = 0$, x never fires and thus there is
 987 no need to define ℓ_0 values for the spikes of x . We define ℓ_0 iteratively in a block by block
 988 manner. Here we do not accumulate spikes and spikes generated in the i^{th} block T_i will
 989 arrive by the first round of the $(i + 1)^{\text{th}}$ block T_{i+1} . Fix a block $T_i = [iL, iL + (L - 1)]$ for
 990 $i \geq 0$ and assume that the latency values ℓ_0 for all prior spikes in rounds $\tau < iL$ have already
 991 been fixed. Thus the active set $A_{0,i}$ can be determined. First, the algorithm checks if there is
 992 a way to spread all the spikes generated in rounds of T_i among the interval $[iL + 2, (i + 1)L]$,
 993 in a way that guarantees that u will not fire in *any* of the rounds this interval. In particular,
 994 *no* spike is scheduled to arrive in round $iL + 1$. Otherwise, all spikes generated in this block
 995 are scheduled to arrive in round $(i + 1)L$ (the first round of the $(i + 1)^{\text{th}}$ block).

996
 997 **Defining the function ℓ_1 .** The definition of the function ℓ_1 is more involved. Unlike
 998 the function ℓ_0 in which all spikes generated in block T_i are scheduled by round $(i + 1)L$,
 999 here the setting is slightly more sensitive. Specifically, the scheduling algorithm of ℓ_1 will
 1000 make sure that non-self spikes arrive to their destination only in important rounds.

1001 **Spikes by the input (inhibitory neuron) x :** All spikes from x are scheduled to arrive
 1002 in the last round of the blocks T_i , namely, in rounds of the form $i \cdot L + (L - 1)$. Formally,
 1003 for every spike $\langle x, u, \tau \rangle$ where $\tau = i \cdot L + (L - 1)$ for some $i \in \mathbb{N}$, let $\ell_1(\langle x, u, \tau \rangle) = L$
 1004 thus arriving in round $\tau + L = (i + 1)L + L - 1$. For every $\tau \in [iL, iL + (L - 2)]$, let
 1005 $\ell_1(\langle x, u, \tau \rangle) = (iL + (L - 1)) - \tau$, thus arriving in round $iL + (L - 1)$ as desired.

1006 **Spikes by $v \neq x$:** The latency values are defined in a round by round fashion, such that
 1007 for every important round τ_i , every neuron u gets activated if possible. Otherwise the arrival
 1008 of the spikes towards u are postponed (when possible) to the next important round τ_{i+1} .
 1009 The spikes generated at non-important rounds will be always delayed to the next important
 1010 round. This is always possible as the distance to the next important round is at most $L/2$.
 1011 For a subset of spikes S , let $w(S) = \sum_{\langle v, u, \tau \rangle \in S} w(v, u)$ be the total weight of the spikes in S .

1012 For every important round τ_i , we will maintain a list of pending spikes $\mathcal{R}_{\tau_i}(u)$ towards u
 1013 that were not yet scheduled. In every step $\tau \geq 0$, the algorithm will schedule the spikes
 1014 generated in this round. If the round τ is important, then the algorithm will also make
 1015 decisions regarding the set of pending spikes $\mathcal{R}_{\tau_i}(u)$.

1016 We will keep the invariant that at the beginning of step τ , the latency value of all spikes
 1017 scheduled to arrive by round τ has already been determined. As we will see, the non-self
 1018 spikes will always be scheduled to arrive in important rounds. As a result, a neuron u fires in
 1019 an unimportant round τ iff u is strong and it fired in round $\tau - 1$. Initially, for every neuron
 1020 u , the algorithm adds every non-self spike $\langle v, u, -1 \rangle$ to $\mathcal{R}_{\tau_i}(u)$. For every $\tau \geq 0$, we consider
 1021 the following algorithm.

- 1022 ■ All self-spikes $\langle u, u, \tau' \rangle$ are given a latency value of $\ell(u, u, \tau') = 1$.
- 1023 ■ **Handling important rounds τ_i .** Consider a neuron u . If u fired in round $\tau_i - 1$, add
 1024 the self-spike $\langle u, u, \tau_i - 1 \rangle$ to the pending spike set $\mathcal{R}_{\tau_i}(u)$. If the total weight of its
 1025 pending spikes (towards u) is sufficiently large to make u fire, all the non-self spikes are
 1026 scheduled to arrive in τ_i . Formally, if $w(\mathcal{R}_{\tau_i}(u)) \geq b(u)$, schedule all these spikes to round
 1027 τ_i by setting $\ell(v, u, \tau') = \tau_i - \tau'$ for every spike $\langle v, u, \tau' \rangle \in \mathcal{R}_{\tau_i}(u)$.
 1028 Otherwise, if the total weight of pending spikes is small, i.e., $w(\mathcal{R}_{\tau_i}(u)) < b(u)$, the
 1029 non-self spikes are deferred to the next important round τ_{i+1} if possible (i.e., if the
 1030 latency does not exceed its upper bound L). Formally, for every non-self pending spike
 1031 $\langle v, u, \tau' \rangle \in \mathcal{R}_{\tau_i}(u)$, if $\tau_{i+1} - \tau' > L$ then let $\ell(\langle v, u, \tau' \rangle) = L$ (i.e., $\langle v, u, \tau' \rangle$ cannot be
 1032 further deferred). Otherwise, add $\langle v, u, \tau' \rangle$ to the pending spike set $\mathcal{R}_{\tau_{i+1}}(u)$ of the next
 1033 important round τ_{i+1} .

1034 Finally, all spikes generated in round τ_i are also (safely) added to the pending list
1035 $\mathcal{R}_{\tau_{i+1}}(u)$.

1036 ■ **Handling unimportant rounds.** The non-self spikes towards u generated in round
1037 τ are added to the pending spike set $\mathcal{R}_{\tau_{i+1}}(u)$ of the next important round τ_{i+1} (after
1038 round τ).

1039 First observe that the function ℓ_1 is valid: All self-spikes have a latency value of 1. Moreover,
1040 the non-self spikes have a latency value in $[1, L]$. To see this observe that for unimportant
1041 round τ , a non-self spike $\langle v, u, \tau \rangle$ is added to the pending list $\mathcal{R}_{\tau_{i+1}}(u)$ where τ_{i+1} is the next
1042 important round after τ . Due to the fact that $\tau_{i+1} - \tau_i \leq L/2$, this assignment is valid. In
1043 addition, the pending spikes $\langle v, u, \tau \rangle \in \mathcal{R}_{\tau_i}(u)$ are deferred to τ_{i+1} only if $\tau_{i+1} - \tau \leq L$.

1044 A.2.2 Proof of Lemma 26

1045 The key lemma that establishes Lemma 26 is the following:

1046 ► **Lemma 27.** *For every neuron $u \neq x$ with $u \in A'_{0,i}$ for $i < L^2/1024$, there exists some*
1047 *$i' \leq i$ such that $u \in A'_{1,i'}$.*

1048 By the correctness of the simple network \mathcal{N} , the output neuron z should not fire when $x_0 = 1$
1049 and with the latency function ℓ_1 . In other words, $z \notin A_{1,i'}$ for any i' . By Lemma 27, we get
1050 that z can only be in $A'_{0,j}$ for some $j \geq L^2/1024$, hence firing when $x_0 = 0$ only after $\Omega(L^3)$
1051 rounds. We start with the following simple observation.

1052 ▷ **Observation 28.** In the simulation of $\langle \mathcal{N}, \bar{\sigma}_0 \rangle$ with ℓ_0 , it holds for every i that: (i) each
1053 strong neuron $s \in A_{0,i}$ fires in every round of block T_i ; (ii) each weak neuron $\omega \in A_{0,i}$ which
1054 is not x fires only in the first round of block T_i ; and (iii) every neuron $v \notin A_{0,i}$ does not fire
1055 in any round of block T_i .

1056 **Proof.** (i). In the simulation with ℓ_0 with $x_0 = 0$ there are no inhibiting spikes, and if a
1057 strong neuron s fires in some round, it will keep on firing for the rest of the simulation.

1058 (ii). By the definition of the latency function ℓ_0 , no spikes from incoming neighbors of
1059 the weak neuron ω arrive in round $iL + 1$, the second round of block T_i . We will prove by
1060 induction on $\tau \in [iL + 1, iL + (L - 1)]$ that ω does not fire in round τ . For the base of the
1061 induction, since $\omega \neq x$ is excitatory and weak, it holds that $0 \leq w(\omega, \omega) < b(\omega)$, thus ω does
1062 not fire in round $iL + 1$. Assume that the claim holds up to round $\tau \geq iL + 1$ and consider
1063 round $\tau + 1$. Since ω did not fire in round τ by the induction assumption, it does not receive
1064 a self spike in round $\tau + 1$. By the definition of the function ℓ_0 , the non-self spikes that arrive
1065 in round $\tau + 1 < (i + 1)L$ cannot make ω fire. Thus ω does not fire in $\tau + 1$ and (ii) holds.

1066 (iii). Let $v \notin A_{0,i}$, i.e., v did not fire in round iL . Since v does not receive negative spikes
1067 in round iL (as the spikes of x are always scheduled to the last round of the blocks). We can
1068 then conclude that $b(v) > 0$. Since in round $iL + 1$, it receives no self-spike and no other
1069 spike, it also did not fire in round $iL + 1$. The argument then follows inductively in the same
1070 manner as in (ii). ◀

1071 We next state the following claim which is crucial to complete the key lemma.

1072 ▷ **Claim 29.** Fix a neuron $u \in A'_{0,i}$ such that for every $v \in A_{0,i-1}$ it holds that $w(v, u) < b(u)$.
1073 Then the total weight of spikes fired towards u in block T_{i-1} is at least $L \cdot b(u)/8$.

1074 We first complete the proof of Lemma 27 and only then prove Claim 29.

1075 **Proof of Lemma 27.** The proof is shown by induction on the block i . For the base case of
 1076 $i = 0$, note that the initial states and the latency functions for the neurons $V \setminus \{x\}$ in both
 1077 simulations are the same, and that spikes from x (that exist only in the simulation with ℓ_1)
 1078 arrive only in round $L - 1$. This implies that in both simulations the same neurons (except
 1079 for x) are active in round $\tau = 0$, hence $A_{0,0} = A_{1,0} \setminus \{x\}$. Now consider the block T_i for
 1080 $1 \leq i < L^2/1024$. Let $u \in A'_{0,i}$, i.e. u fires for the first time in round iL in the simulation
 1081 with ℓ_0 and $x_0 = 0$.

1082 **Case 1: There exists a previously firing dominant incoming neighbor:** First
 1083 assume that u has some incoming neighbor $v \in A_{0,i-1}$ with $w(v, u) \geq b(u)$. By definition
 1084 $v \in A'_{0,j}(u)$ for some $j \leq i - 1$, and then by the induction assumption $v \in A'_{1,i'}$ for some
 1085 $i' \leq j \leq i - 1$. By definition of the latency function ℓ_1 , since $w(v, u) \geq b(u)$, the total weight
 1086 of spikes from incoming neighbors will be sufficient to activate u in the next important
 1087 round, $\tau_{i'+1}$. Therefore, $u \in A_{1,i'+1}$, which implies $u \in A'_{1,i''+1}$ for some $i'' \leq i' + 1$. Since
 1088 $i'' \leq i' + 1 \leq i$ the condition holds.

Case 2: All previously firing incoming neighbors are not dominant: By applying
 Claim 29 on u and block T_i , we get that the total weight of spikes fired towards u in block
 T_{i-1} is at least $L \cdot b(u)/8$. Due to Observation 28, we get

$$L \cdot w(\mathbf{S}(A_{0,i-1}(u)), u) + w(\mathbf{W}(A_{0,i-1}(u)), u) \geq \frac{L}{8} \cdot b(u).$$

By the definition of $A'_{0,j}$ and the induction assumption, it holds that

$$A_{0,i-1} \subseteq \bigcup_{j \leq i-1} A'_{0,j} \subseteq \bigcup_{i' \leq i-1} A'_{1,i'}.$$

We now consider the simulation with $x_0 = 1$ and the latency function ℓ_1 , and partition
 all the rounds until τ_i into k blocks of L rounds (except perhaps the last one). Formally, for
 every $j \leq k - 2$, let $B_j = [jL, jL + (L - 1)]$ and let $B_{k-1} = [(k - 1)L, \tau_{i-1} + 15]$. Denote by
 $\mathbf{S}(B_j)$ and $\mathbf{W}(B_j)$ the strong and weak (respectively) incoming neighbors of u that fire in
 some round of B_j . Using these notations, we can write

$$\sum_{j=0}^{k-1} L \cdot w(\mathbf{S}(B_j), u) + w(\mathbf{W}(B_j), u) \geq \frac{L}{8} \cdot b(u).$$

Case 2.1: Most of the weight is in the last block. We first assume that

$$L \cdot w(\mathbf{S}(B_{k-1}), u) + w(\mathbf{W}(B_{k-1}), u) \geq \frac{L}{16} \cdot b(u).$$

1089 Consider the algorithm that defines ℓ_1 , and recall that $\mathcal{R}_{\tau_{i'}}(u)$ is the set of pending spikes
 1090 that were not yet scheduled when the algorithm considered the important round τ_i . The
 1091 interesting case is when u did not fire in any round of B_{k-1} . In such a case, all the spikes
 1092 generated towards u in the rounds of B_{k-1} were added to the pending list of $\mathcal{R}_{\tau_i}(u)$. Note
 1093 that each strong neuron $v \in \mathbf{S}(B_{k-1})$ fires at least 16 spikes in B_{k-1} , since $\tau_i - \tau_{i-1} = 16$.
 1094 Furthermore, each $v \in \mathbf{W}(B_{k-1})$ fires at least one spike in B_{k-1} . Moreover, the gap between
 1095 any $\tau_{i'} \in B_{k-1}$ and τ_i is at most L rounds, so they do not exceed the maximal latency in τ_i .
 1096 Altogether, we get that

$$1097 \quad w(\mathcal{R}_{\tau_i}(u)) \geq 16 \cdot w(\mathbf{S}(B_{k-1}), u) + w(\mathbf{W}(B_{k-1}), u) \geq \tag{3}$$

$$1098 \quad \frac{16}{L} \cdot (L \cdot w(\mathbf{S}(B_{k-1}), u) + w(\mathbf{W}(B_{k-1}), u)) \geq b(u) .$$

1099

23:28 The Computational Cost of Asynchronous Neural Communication

1100 Therefore u fires in τ_i and $u \in A_{1,i'}$ for some $i' \leq i$ as desired.

Case 2.2: Most of the weight is in the first $k - 1$ blocks. It remains to consider the complementary case where

$$\sum_{j=0}^{k-2} L \cdot w(\mathbf{S}(B_j), u) + w(\mathbf{W}(B_j), u) \geq \frac{L}{16} \cdot b(u).$$

1101 Since $i < L^2/1024$ and each block B_j for $j \leq k - 2$ consists of $L/32$ important rounds, we
 1102 have $k \leq \frac{L^2/1024}{L/32} = \frac{L}{32}$. Therefore, by an averaging argument there exists B_j for $j \leq k - 2$
 1103 satisfying that:

$$1104 \quad L \cdot w(\mathbf{S}(B_j), u) + w(\mathbf{W}(B_j), u) \geq 2 \cdot b(u). \quad (4)$$

1106 First observe that every strong neuron $s \in \mathbf{S}(B_j)$ fires for at least $L/2$ rounds in this block.
 1107 The reason is that there is a gap of $L/2$ rounds between the last important rounds of B_j
 1108 and the round where the inhibiting spike from x arrives. During this time interval every
 1109 strong neuron in $\mathbf{S}(B_j)$ keeps on firing. Now, assume that u does not fire in any round of
 1110 B_j , and denote the first important round of B_{j+1} by $\tau_{i'}$. Again, consider the algorithm
 1111 that defines ℓ_1 . Since u did not fire in any round of the block B_j , all the spikes that are
 1112 fired towards u in B_j are in the residual set $\mathcal{R}_{\tau_{i'}}(u)$. Therefore by Eq. (4), we get that
 1113 $w(\mathcal{R}_{\tau_{i'}}(u)) \geq (L/2) \cdot w(\mathbf{S}(B_j), u) + w(\mathbf{W}(B_j), u) \geq b(u)$, and u fires in $\tau_{i'}$. Therefore, we get
 1114 that u fires either in some important round of B_j or in $\tau_{i'}$. In both cases there is a round
 1115 $\tau_{i''}$ with $i'' \leq i$ such that $u \in A_{1,i''}$. This implies $u \in A'_{1,i''}$ for $i'' \leq i$, and the condition
 1116 holds. \blacktriangleleft

1117 Finally, it remains to prove Claim 29.

1118 **Proof of Claim 29.** Recall that $\mathbf{S}(A_{0,i-1}(u))$ and $\mathbf{W}(A_{0,i-1}(u))$ are the strong and weak
 1119 (respectively) incoming neighbors of u that fire in block T_{i-1} . If $w(\mathbf{S}(A_{0,i-1}), u) \geq b(u)/8$,
 1120 then by Observation 28 the total spike weight fired in block $i - 1$ is at least $L \cdot b(u)/8$, and
 1121 we are done. Therefore, it remains to consider the case where $w(\mathbf{S}(A_{0,i-1}), u) < b(u)/8$
 1122 and $w(\mathbf{W}(A_{0,i-1})) < L \cdot b(u)/8$. We will show that in this case, there is a way to schedule
 1123 all spikes fired towards u in block T_{i-1} to arrive in rounds $[(i - 1)L + 2, iL]$, such that u
 1124 does not get activate in any of these rounds. By the definition of ℓ_0 , we get that u does
 1125 not get activated in any of the rounds $[(i - 1)L + 2, iL]$, and in particular $u \notin A'_{0,i}$, thus a
 1126 contradiction.

1127 First observe that $b(u) > 0$ since u did not fire in round $(i - 1)L$ (as $u \notin A_{0,i-1}$) and
 1128 it did not receive any negative spike in that round (as all negative spikes arrive in the last
 1129 rounds of the blocks). We next show that all the spikes generated in block T_{i-1} can be
 1130 scheduled in rounds $[(i - 1)L + 2, iL]$ without making u fire in any of these rounds. Since
 1131 the scheduling algorithm of ℓ_0 works in this manner, we will get a contradiction to the fact
 1132 that $u \in A_{0,i}$.

1133 **Scheduling spikes from weak neighbors.** Let $F_{\mathbf{W}} = \{\langle v, u, (i - 1)L \rangle \mid v \in \mathbf{W}(A_{0,i-1})\}$
 1134 be the spikes of weak neighbors fired in the block T_{i-1} . Recall that by Observation 28, these
 1135 weak neurons fire only in the first round. Since these spikes are fired in round $(i - 1)L$,
 1136 they can arrive in any of the rounds $[(i - 1)L + 2, iL]$. As the total weight of the weak
 1137 spikes is at most $Lb(u)/8$, we show that we can schedule them in a greedy manner into
 1138 at most $L/2 - 2$ rounds while keeping the total weight in each such round to strictly less
 1139

1140 than $b(u)$. We traverse the weak spikes one by one, and start throwing them into rounds
 1141 in $[(i-1)L+2, iL-1]$. We add a spike to round τ as long as the total weight of weak
 1142 spikes scheduled to it is at most $b(u)/2$. If the addition of the next weak spike raises the
 1143 weight to above $b(u)$ it is deferred to the next round $\tau+1$. Let τ' be the last round to
 1144 which the weak spikes are scheduled. Since in each $\tau \in [(i-1)L+2, \tau'-1]$ the total weight
 1145 of weak spikes is at least $b(u)/2$, we get that $\tau' \leq (i-1)L+L/4+3 \leq (i-1)L+L/2$ as desired.

1146

1147 **Scheduling spikes from strong neighbors.** We next turn to show that also the strong
 1148 spikes can be scheduled in a balanced manner in the remaining $L/2$ slots of the block T_i without
 1149 activating the neuron u . Let $F_{\mathbf{S}} = \{\langle v, u, \tau \rangle \mid v \in \mathbf{S}(A_{0,i-1}), \tau \in T_{0,i-1}\}$ be the spikes of
 1150 strong neighbors fired in block T_{i-1} . For a spike $\langle v, u, \tau \rangle \in F_{\mathbf{S}}$ with $\tau \leq (i-1)L + (L/2 - 1)$,
 1151 schedule $\langle v, u, \tau \rangle$ to arrive in round $\tau + L/2 + 1$. For $\langle v, u, \tau \rangle \in F_{\mathbf{S}}$ with $\tau \geq (i-1)L + L/2$,
 1152 schedule $\langle v, u, \tau \rangle$ to arrive in round $\tau + 1$. In this way, due to Observation 28, u receive
 1153 two spikes from each $v \in \mathbf{W}(A_{0,i-1})$ in each round $\tau \in [(i-1)L + L/2 + 1, iL]$. Since
 1154 $w(\mathbf{S}(A_{0,i-1})) < b(u)/8$, we get that the total weight of spikes that u receives in each of
 1155 these rounds is less than $b(u)/4$, and therefore u does not get activated. Overall, all spikes
 1156 generated in the block T_{i-1} are scheduled by ℓ_0 without activating the neuron u in any of
 1157 the rounds $[(i-1)L+2, iL]$, contradiction to the fact that $u \in A_{i,0}$. The claim follows. ◀

1158 **B** Missing Proofs for the Positive Results

1159 **B.1** Synchronization of Boolean Gates

1160 **Proof of Observation 13.** The network is as follows: connect each input neuron x_i to the
 1161 output neuron z by an edge of weight $w(x_i, z) = 1$, and let the bias of z be $b(z) = 1$. First
 1162 note that if all input neurons x_i did not fire in round 0, then $\text{pot}(z, \tau) = -1$ for all τ , and z
 1163 will not fire. If a neuron x_i fires in round τ , since the latency of each edge is at most L , there
 1164 is a round $\tau' \in [\tau + 1, \tau + L]$ in which the spike from x_i arrives to z . Thus, in round τ' , the
 1165 weighted incoming sum to z is at least 1, therefore $\text{pot}(z, \tau') \geq 0$ and z fires in round τ' . ◀

1166 **The Complete Proof of Lemma 14:** We analyze the correctness of the network NOT_{sync} .
 1167 We begin by proving the following auxiliary claim.

1168 ▷ **Claim 30.** If all the intermediate neurons v_0, \dots, v_L fire starting round τ for at least
 1169 $L(L+1)$ rounds, then there exists a round $\tau' \in [\tau + 1, \tau + L(L+1)]$ in which the output
 1170 neuron z fires (i.e., regardless of the latencies of the edges).

1171 **Proof.** For every $i \in \{0, 1, \dots, L(L+1)\}$, denote the L -length interval $T_i = [\tau + i \cdot L, \tau + (i+1)L - 1]$. In addition, define $\tilde{T} = T_0 \cup \dots \cup T_{L+1}$. Let q_i be the number of spikes that were
 1172 fired in the interval of T_i but received by z in the next interval T_{i+1} . Note that since the
 1173 maximum edge latency is L , in the worst case the spikes of interval T_i must arrive to z by
 1174 the end of the next interval T_{i+1} . We next prove by induction on i that either z fires by the
 1175 end of the interval T_i , or $q_{i+1} \geq i \cdot L$. For the base of the induction, consider $i = 0$. If z did
 1176 not fire in some round during T_0 , we claim that $q_1 \geq L$. Since all the $L+1$ neurons fire
 1177 in every round during the interval, overall $L(L+1)$ many spikes were fired. By the fact
 1178 that z did not fire during T_0 , we have that in each of these rounds, it received at most L
 1179 spikes. This implies that z received at most L^2 many spikes during T_0 , and therefore at least
 1180 $q_1 \geq L$ many spikes will be received by z in the interval T_1 . Assume that the claim holds
 1181 up to $i-1$ and consider the i^{th} interval. If z fired by the end of the i^{th} interval T_i , we are
 1182 done. Otherwise, by induction assumption for $i-1$, we have that $q_i \geq i \cdot L$. In addition,

23:30 The Computational Cost of Asynchronous Neural Communication

1184 all these q_i spikes must be received at z during the interval T_i . Then, in interval T_i we
 1185 again have a total of $L(L+1)$ fresh spikes by the neurons v_0, \dots, v_L . This creates a total of
 1186 $L(L+1) + i \cdot L$ spikes. As z did not fire in T_i , it received at most L^2 many spikes, leaving at
 1187 least $q_{i+1} \geq L(L+1) + i \cdot L - L^2 \geq (i+1)L$ spikes for the next interval. This completes the
 1188 proof of the induction step.

1189 Overall, for $i = L$, we have that either z fired by the end of the interval T_L , or that
 1190 $q_{L+1} \geq L(L+1)$. In the latter case, since all these spikes must arrive during the last interval
 1191 T_{L+1} , by the pigeonhole principle there must be a round in this interval in which z received
 1192 at least $L+1$ spikes and fire. This completes the proof of the claim. ◀

1193 **Proof of Lemma 14.** Due to Claim 30, it remains to show that if x did not fire in round
 1194 0, then there must be a starting round τ , in which all the neuron v_0, \dots, v_L fire for at least
 1195 $L(L+1)$ rounds.

1196 If x did not fire, then neither the memory neuron m nor the reset neuron r fire during
 1197 the execution. Since we assume that v^* fires in round 0, it must hold that all these
 1198 neurons fire starting some round $\tau \in [5L^2, 5L^3 + 2L]$. This holds due to the self-loops on
 1199 the neurons v_0, \dots, v_L , and the chain of length $5L^2$. By Claim 30, there exists a round
 1200 $\tau' \in [\tau + 1, \tau + L(L+1)]$ in which z fires.

1201 We next show that if x fired in round 0, then z would not fire in any round. The key
 1202 observation is as follows:

1203 \triangleright **Observation 31.** In order for z to fire in some round τ , it must receive spikes from *at least*
 1204 two different v_i neurons.

1205 **Proof.** To see this, note that since the maximum edge latency is L , in round τ , z can receive
 1206 spikes only from the L previous rounds $\tau - L, \dots, \tau - 1$. In particular, a single neuron can
 1207 be accounted for at most L many spikes received by z in a given round. Finally, since the
 1208 bias of z is $L+1$, and all edge weights are 1, we conclude that z must receive spikes from at
 1209 least two neurons in order to fire. ◀

1210 When x fires in round 0, the memory neuron m fires from round $\tau_m \in [1, L+1]$ ahead, due
 1211 to its self-loop. Hence, starting round $\tau_r \in [2, 2L+2]$, the reset neuron r starts firing at
 1212 least *once* in every interval of $2L$ rounds. Recall that each v_i gets a negative spike from the
 1213 inhibitor r and positive spike from the neuron $c_{5iL} \in C$. We next show that each neuron v_i
 1214 gets inhibited at least L rounds *before* the activation of the neurons v_{i+1} . As a result, at any
 1215 point of time, there will be no two neurons v_i and v_j such that z received both of their spikes
 1216 in the same round. By induction on i , the first intermediate neuron v_0 has an incoming
 1217 edge from $c_0 = v^*$, and thus it begins to fire in some round $\tau' \in [0, L]$. Due to the negative
 1218 edge from the reset neuron r , it stops firing before round $3L+2$. Since v_1 has an incoming
 1219 edge from c_{5L} , it starts firing only after round $5L+1$, and therefore z starts receiving spikes
 1220 from v_1 only starting round $5L+2$. Assume the claim is correct for neurons v_0, \dots, v_{i-1} and
 1221 consider neuron v_i . If the neuron v_i starts firing in round τ_i , by round $\tau_i + 2L$ it is inhibited
 1222 by r . Because v_i starts to fire after receiving a spike from $c_{5 \cdot i \cdot L}$ and v_{i+1} starts firing after
 1223 receiving a spike from $c_{5(i+1)L}$, neuron v_{i+1} begins to fire only after round $\tau_i + 4L$, at least
 1224 L rounds after v_i is inhibited.

1225 Hence, z cannot receive input from two different neurons v_i, v_j at the same round, and
 1226 the claim follows by combining Observation 31. Finally, the next observation plays a role in
 1227 the subsequent constructions.

1228 \triangleright **Observation 32.** The correctness still holds even if the chain starts to fire at some round
 1229 $\tau > 0$. In this case the output neuron z fires in some round $t \in [\tau + 1, \tau + \Theta(L^3)]$.

1230 **Proof of Observation 32.** The correctness of the observation follows from the fact that the
 1231 input neuron x activates the memory neuron m , that keeps on firing (i.e., presenting the
 1232 state of x) due to its self-loop. Thus all arguments in Lemma 14 still hold in case the chain
 1233 starts to fire in any later round. ◀

1234 B.2 Synchronization of a Boolean Circuit, Proof of Lemma 15


1235 **The Construction.** Given a Boolean circuit \mathcal{A} of OR / NOT gates g_1, \dots, g_m of depth d ,
 1236 we describe a construction of an analogous neural network \mathcal{N} with a similar execution. For
 1237 every g_i , let $\text{Sync}(g_i)$ be the synchronized sub-network of the gate g_i . Specifically, for a NOT
 1238 gate (resp., OR) g_i , the sub-network $\text{Sync}(g_i)$ is taken from Lemma 14 (resp., Observation 13).
 1239 Recall, that for a NOT gate g_i , its synchronized sub-network $\text{Sync}(g_i)$ contains a chain of
 1240 neurons where the head of the chain c_0 will be denoted hereafter by v_i^* . The network \mathcal{N}
 1241 consists the following components:

- 1242 1. Input neurons x_1, \dots, x_n , and output neurons z_1, \dots, z_k , that serve as the input and the
 1243 output for the network \mathcal{N} .
- 1244 2. A chain $C = [c_0, \dots, c_q]$ containing $q + 1 = \alpha d L^3 + 1$ neurons, where α is a constant
 1245 satisfying that $\alpha L^3 \geq 5L^3 + L$. For every $i \geq 0$, the neuron c_i has bias $b(c_i) = 1$.
 1246 Moreover, for every $i \geq 1$ the neuron c_i has a positive incoming edge from c_{i-1} with weight
 1247 $w(c_{i-1}, c_i) = 1$. Our simulation starts with neuron c_0 firing.
- 1248 3. A $\text{Sync}(g_i)$ network (using Lemma 14 and Observation 13 respectively) for every gate g_i .
 1249 The connections between these components are as follows:
 - 1250 1. For every gate g_i in the first layer, the input for its synchronized sub-network $\text{Sync}(g_i)$ is
 1251 given by $x_{i,1}, \dots, x_{i,k_i}$, namely, the input bits of the gate g_i in the circuit \mathcal{A} .
 - 1252 2. For every gate g_i in layer $j \geq 2$, denote by $g_{i,1}, \dots, g_{i,k_i}$ the input of the gate g_i in
 1253 the circuit \mathcal{A} . In the network \mathcal{N} , the input to the sub-network $\text{Sync}(g_i)$ are the output
 1254 neurons of the sub-networks $\text{Sync}(g_{i,1}), \dots, \text{Sync}(g_{i,k_i})$.
 - 1255 3. The output gates of the network \mathcal{N} are the output neurons of the sub-networks
 1256 $\text{Sync}(o_1), \dots, \text{Sync}(o_k)$, where o_1, \dots, o_k are the output gates of the circuit \mathcal{A} .
 - 1257 4. Finally, the synchronized sub-networks of the NOT gates are connected to the chain C as
 1258 follows. For each NOT gate g_i in every layer j , the $(j\alpha L^3)^{\text{th}}$ neuron $c_{j\alpha L^3}$ in the chain
 1259 has an outgoing edge to v_i^* with weight 1 (where v_i^* is the head of the internal chain in
 1260 $\text{Sync}(g_i)$), since the bias of v_i^* is 1, a spike from $c_{j\alpha L^3}$ makes v_i^* fire.

1261 Figure 3 illustrates the construction for a circuit with 4 NOT and OR gates of depth 3. We
 1262 note that one can shave an L -factor in the size and time overhead of lemma 15, by reusing
 1263 the synchronization chain for all the Boolean gates in the network. For clarity of explanation,
 1264 we defer this improvement to the full paper.

1265 **Correctness.** Let V be the total set of neurons in \mathcal{N} , and let $\ell : V \times V \times \mathbb{N} \rightarrow [1, L]$
 1266 be a fixed (arbitrary) nice latency function. First note that in the global chain C , each of
 1267 the neurons fires once, in a sequential manner. Recall, that we assume that the starter c_0
 1268 fires in round $\tau_0 = 0$. For every $j \in \{1, \dots, d\}$, let τ_j be the round in which $c_{j\alpha L^3}$ fires (i.e.,
 1269 the spike from $c_{j\alpha L^3-1}$ is received at $c_{j\alpha L^3}$ in round $\tau_j - 1$).

1271 For every gate g_i in the circuit \mathcal{A} in layer $j \geq 1$, denote by $\text{out}(g_i, \mathcal{A})$ the final state of g_i
 1272 after receiving its inputs in the circuit \mathcal{A} . In addition, let q_i be the output neuron in the
 1273 sub-network $\text{Sync}(g_i)$, and let $\sigma_t(q_i, \mathcal{N})$ be the state of the neuron in round t when simulating
 1274 the network \mathcal{N} . Our goal is to show that for every g_i , its corresponding output q_i in the
 1275 network \mathcal{N} , has the same "output" as g_i in the circuit \mathcal{A} .



■ **Figure 3** The transformation of the circuit on the left with 4 inputs and 3 layers. For each gate we add the corresponding synchronized sub-network, where we connect the input and output neurons of the sub-network according to the original circuit. In addition we introduce a global chain that activates the sub-networks in each layer after the previous layers have already finished the computation. The first neuron in the global chain is set to be the starter neuron which fires in the beginning of the simulation.

1276 \triangleright **Claim 33.** For every layer $j \in \{1, \dots, d\}$ and every gate g_i in layer j of circuit \mathcal{A} , it holds
 1277 that: (i) If $out(g_i, \mathcal{A}) = 0$, then $\sigma_t(q_i, \mathcal{N}) = 0$ for every t , and (ii) If $out(g_i, \mathcal{A}) = 1$, then
 1278 there exists $t \in [\tau_{j-1} + 1, \tau_j]$ such that $\sigma_t(q_i, \mathcal{N}) = 1$.

1279 **Proof.** We prove by induction on the layer j . For $j = 1$, recall that the input neurons
 1280 $x_{i,1}, \dots, x_{i,k_i}$ of the sub-network $\text{Sync}(g_i)$ are the input neurons of the network \mathcal{N} . Therefore,
 1281 in round 0 in the simulation of \mathcal{N} , the sub-network $\text{Sync}(g_i)$ has the same input as gate g_i
 1282 in the circuit \mathcal{A} . Assume first that g_i is a NOT gate. Then the spike of the starter neuron
 1283 c_0 arrived at the head chain v_i^* by round L . Combining with Observation 32 we get that
 1284 if $out(g_i, \mathcal{A}) = 1$ then $\sigma_t(q_i, \mathcal{N}) = 1$ for some $t \in [L, L + 5L^3] \subseteq [1, \alpha L^3]$. In addition, if
 1285 $out(g_i, \mathcal{A}) = 0$ then $\sigma_t(q_i, \mathcal{N}) = 0$ for every t . The case where g_i is an OR gate is even simpler
 1286 and follows by Observation 13. Since the path from c_0 to $c_{\alpha L^3}$ in the chain C is of length
 1287 αL^3 , we have that $\tau_1 \geq \alpha L^3$. Therefore $[1, \alpha L^3] \subseteq [\tau_0 + 1, \tau_1]$, and the claim holds for $j = 1$.

1288 For the induction step, let $j \geq 2$, and assume correctness up to layer $j - 1$. We now
 1289 prove the claim for layer j . Let g_i be a gate in layer j . By Observation 32, the important
 1290 thing to take care of regarding a NOT gate g is to make sure that its inputs have the correct
 1291 states (i.e., as the corresponding states in \mathcal{A}) by the time that the head of the chain v_i^* in
 1292 $\text{Sync}(g)$ has received the spike from $c_{(j-1)\alpha L^3}$. Denote by $q_{i,1}, \dots, q_{i,k_i}$ the output neurons
 1293 of the sub-networks $\text{Sync}(g_{i,1}), \dots, \text{Sync}(g_{i,k_i})$. By the induction assumption, for each $g_{i,h}$, if
 1294 $out(g_{i,h}, \mathcal{A}) = 0$ then $\sigma_t(q_{i,h}) = 0$ for every t , and otherwise $\sigma_t(q_{i,h}) = 1$ for some $t \leq \tau_{j-1}$.
 1295 Since the neurons $q_{i,1}, \dots, q_{i,k_i}$ are the input neurons of g , it holds that the sub-network
 1296 $\text{Sync}(g)$ gets the same input as the input of g in \mathcal{A} , by round τ_{j-1} of the simulation of \mathcal{N} .

1297 Now, assume that g_i is a NOT gate. Then, by round $\tau_{j-1} + L$, the head of the chain v_i^*
 1298 has received the spike from $c_{(j-1)\alpha L^3}$. Combining with Observation 32, when $out(q_i, \mathcal{N}) = 1$
 1299 , it holds that $\sigma_t(q_i, \mathcal{N}) = 1$ for some $t \in [\tau_{j-1} + 1, \tau_{j-1} + L + 5L^3]$. In addition, when
 1300 $out(q_i, \mathcal{N}) = 0$ then $\sigma_t(q_i, \mathcal{N}) = 0$ for every t . Again, since the path from $c_{(j-1)\alpha L^3}$ to $c_{j\alpha L^3}$
 1301 in the chain C is of length αL^3 , we have that $\tau_j \geq \tau_{j-1} + L + 5L^3$, and the claim follows.
 1302 The case where g_i is an OR gate follows in a similar way by Observation 13. \blacktriangleleft

1303 Lemma 15 follows by using Claim 33 with $j = d$, and noting that each output neuron z_i in
 1304 \mathcal{N} is the output neuron $q_{i'}$ for some sub-network $\text{Sync}(g_{i'})$ where $g_{i'}$ is a gate in layer d of \mathcal{A} .
 1305 This completes the correctness and the bound on the time overhead. We finally bound the
 1306 size of the network. The network \mathcal{N} consists of a chain of $O(dL^3)$ neurons, and a $\text{Sync}(g_i)$
 1307 sub-network of size $O(L^2)$ for each gate g_i in \mathcal{A} . Therefore, there are overall $O(dL^3 + mL^2)$
 1308 auxiliary neurons.

1309 B.3 Synchronization of a Single Deterministic Threshold Gate

1310 We now turn to consider the synchronized implementation of a single deterministic threshold
 1311 gate and prove Lemma 16.

1312 Thanks to a result of [30], we can assume without loss of generality that the weights and
 1313 bias values can be represented using binary vectors of length $\lceil \Delta \log \Delta \rceil$. Hastad [10] also
 1314 showed that this bound is tight. In addition, we can also assume without loss of generality
 1315 that $b(z) \geq 0$. The key part is to implement the single threshold gate by a Boolean circuit.
 1316 This requires small adaptations from existing results in the area, specifically we will use the
 1317 following known facts.

1318 \triangleright **Fact 34.** [31, 38][Iterated Addition] Given two input binary vectors $\bar{x} = [x_1, \dots, x_\Delta]$
 1319 and $\bar{y} = [y_1, \dots, y_\Delta]$, there exists a Boolean circuit with $\text{poly}(\Delta)$ gates and $O(1)$ depth that
 1320 outputs the binary representation of $\text{dec}(\bar{x}) + \text{dec}(\bar{y})$.

1321 \blacktriangleright **Corollary 35.** [Multiple Iterated Addition] Given Δ input binary vectors $\bar{x}_1, \dots, \bar{x}_\Delta$ where
 1322 $\bar{x}_i \in \{0, 1\}^m$ for some integer $m \geq 1$, there exists a Boolean circuit with $\text{poly}(\Delta, m)$ gates
 1323 and $O(\log \Delta)$ depth that outputs the binary representation of $\sum_i \text{dec}(\bar{x}_i)$.

1324 \triangleright **Observation 36 (Comparison).** Given two input binary vectors $\bar{x} = [x_1, \dots, x_\Delta]$ and
 1325 $\bar{y} = [y_1, \dots, y_\Delta]$, there exists a Boolean circuit with $\text{poly}(\Delta)$ gates and $O(1)$ depth that
 1326 outputs 1 iff $\text{dec}(\bar{x}) \geq \text{dec}(\bar{y})$.

1327 We are now ready to implement a threshold gate by a small depth Boolean circuit of
 1328 polynomial size. This lemma explains the dependency in the largest in-degree Δ of the final
 1329 synchronized solution.

1330 \blacktriangleright **Lemma 37.** Given a threshold gate g with Boolean inputs x_1, \dots, x_Δ with weights
 1331 w_1, \dots, w_Δ , and an output neuron z with bias $b(z)$, there exists a Boolean circuit that
 1332 computes g (i.e., outputs 1 iff $\sum w_i \cdot x_i \geq b(z)$) using $\text{poly}(\Delta)$ gates and depth $O(\log \Delta)$.

1333 **Proof.** Each input x_i is connected to $\ell = \lceil \Delta \cdot \log \Delta \rceil$ neurons $w_{i,1}, \dots, w_{i,\ell}$, where the edge
 1334 weight $w(x_i, w_{i,j})$ is 1 if the j^{th} -bit in w_i is 1 and 0 otherwise. We set the bias values to
 1335 be $b(w_{i,j}) = 1$. Thus, the outgoing edge weights of x_i encode the binary representation
 1336 of the weight w_i . As a result, once x_i fires in round τ , after at most L rounds, $w_{i,j}$ fires
 1337 iff the j^{th} bit in the representation of w_i is 1. As we will see, those $\Delta^2 \cdot \log \Delta$ neurons
 1338 $w_{1,1}, \dots, w_{\Delta,\ell}$ will serve as the input layer to the circuit. In addition, we also represent the
 1339 bias of z using ℓ neurons b_1, \dots, b_ℓ that encode the binary representation of $b(z)$: the bias
 1340 of $b_j = 1$ if the j^{th} bit in $b(z)$ is 0, and $b_j = -1$ otherwise. Let $\bar{x}_{pos} = \{x_i \mid w_i \geq 0\}$
 1341 and $\bar{x}_{neg} = \{x_i \mid w_i < 0\}$. In the same manner, let $W_{pos} = \sum \{w_i \mid w_i \geq 0\}$ and
 1342 $W_{neg} = \sum \{|w_i| \mid w_i < 0\}$. We will use the Multiple Iterated Addition circuit of Corollary
 1343 35 to compute the binary representation of W_{pos} and $W_{neg} + b(z)$. Finally, we use the
 1344 Comparison circuit of Observation 36 to compare those values, such that the output will be
 1345 1 iff $W_{pos} \geq W_{neg} + b(z)$, hence computing the function of the threshold gate. \blacktriangleleft

1346 The final synchronous implementation of g is obtained by applying Lemma 15 on \mathcal{C} , i.e.,
 1347 $\text{Sync}(g) \leftarrow \text{Sync}(\mathcal{C})$. The construction uses a total $O(\log \Delta \cdot L^3 + \text{poly}(\Delta) \cdot L^2)$ auxiliary
 1348 neurons, and computation time of $O(L^4 \cdot \log \Delta)$ rounds. This completes the proof of Lemma 16.

1349 B.4 Probabilistic Threshold Gate

1350 B.4.1 Description of the Boolean Circuit

1351 The construction of the boolean circuit \mathcal{A} approximating a probabilistic threshold gate is
 1352 achieved using two main steps. First we sample an almost uniform random variable, then we
 1353 use the sampled value in order to approximate a sample from the Logistic distribution.

1354
 1355 **Step 1: Sampling from the Almost Uniform Distribution.** We introduce $k =$
 1356 $4 \log(1/\epsilon)$ uniformly random gates, denoted as r_1, \dots, r_k . Hence, $\text{dec}(\bar{r})$ encodes an integer
 1357 number that is uniformly sampled between 0 and $\lceil (1/\epsilon)^4 \rceil$. In addition, we introduce
 1358 k input bits (with fixed value) a_1, \dots, a_k such that $\text{dec}(\bar{a}) = \lceil (1/\epsilon)^4 \rceil$. Thus, the value
 1359 $r' = \text{dec}(\bar{r})/\text{dec}(\bar{a})$ is sampled uniformly at random from the set $\{0, \epsilon^4, 2\epsilon^4, 3\epsilon^4, \dots, 1\}$.

1360
 1361 **Step 2: Sampling from the Almost Logistic Distribution.** Next, we transform
 1362 the sample r' from Step 1 into a sample from an almost Logistic distribution. This is done by
 1363 using the method of inverse transform sampling. In our context, for a sample r u.a.r in $[0, 1]$,
 1364 the value $b + \ln(r/(1-r))$ is a sample from the Logistic distribution with mean b and scale
 1365 1. To compute the expression $b + \ln(r'/(1-r'))$ using a Boolean circuit, we approximate
 1366 the $\ln(x)$ function (up to $\pm \text{poly}(\epsilon)$) using the first $O(\log 1/\epsilon)$ terms of the Taylor expansion
 1367 around a point x_0 where $0 \leq x_0 - x \leq 1/2$.

1368 **► Definition 38 (ϵ -Approximation of the $\ln(x)$ Function).** Given $x > 0$ and a positive integer
 1369 k , let $\widehat{\ln}_k(x)$ be the \ln -approximation of x obtained by computing the first k terms of the
 1370 Taylor expansion around a point x_0 , where $0 \leq x_0 - x \leq 1/2$. When k is clear from the
 1371 content we may omit it and simply write $\widehat{\ln}(x)$.

1372 The task of sampling from the (almost) Logistic distribution then boils into computing
 1373 $f(r') = \widehat{\ln}_k(r'/(1-r'))$ with $k = \lceil 4 \log 1/\epsilon \rceil$. We first use a Boolean circuit to distinguish
 1374 between the case where $r' \leq 1-r'$, and the complementary case. Using the vectors \bar{r} and \bar{a} , this
 1375 can be done using integer operations and comparison as $r'/(1-r') = \text{dec}(\bar{r})/(\text{dec}(\bar{a}) - \text{dec}(\bar{r}))$.
 1376 When $r' > 1-r'$, we calculate $f(-r')$, and then either add or subtracts it from the bias b
 1377 respectively.

1378 In what follows, assume that $r' \leq 1-r'$, and therefore $r'/(1-r') \in [0, 1]$. To pick
 1379 the point x_0 around which the Taylor approximation is expanded, we let $x_0 = 1/2$ when
 1380 $r'/(1-r') \leq 1/2$, and $x_0 = 1$ otherwise. This latter condition can also be easily checked
 1381 with a Boolean circuit.

1382 To finally be able to compute the function $\widehat{\ln}_k(x)$ using a Boolean circuit, we must ensure
 1383 that all our operations are applied on *integers*. Therefore, instead of computing $f(r')$, we will
 1384 be actually computing $q \cdot f(r')$ for some large enough constant q that guarantees that $q \cdot f(r')$
 1385 is an integer. Specifically, letting $q = (\text{dec}(\bar{a}) - \text{dec}(\bar{r}))^k$ does the job as the function $\widehat{\ln}_k(x)$ is
 1386 a polynomial of degree k . This factor of q would not affect the correctness of the computation
 1387 as it will be canceled out later on. Using the circuit for iterated addition [38, 31] and fast
 1388 multipliers [9], we can compute $q \cdot f(r')$ using only integer addition and multiplication. The
 1389 output of the final Boolean circuit is \bar{y} where $\text{dec}(\bar{y}) = q \cdot b + q \cdot f(r')$. In the analysis section,
 1390 we show that $\text{dec}(\bar{y})/q$ is sampled from a distribution that is $\text{poly}(\epsilon)$ -close to the Logistic

1391 distribution with mean b .

1392

1393 **Putting all Together: The Output Circuit.** Let w_1, \dots, w_Δ be the weights of the
 1394 probabilistic threshold gate g . To cancel out the multiplication of q in the output bias
 1395 value from the previous step, we multiply all the incoming weights by q as well. We can
 1396 then use the construction of Lemma 16 for a deterministic threshold gate with weights
 1397 $w'_1 = q \cdot w_1, \dots, w'_\Delta = q \cdot w_\Delta$ and bias $b'' = \text{dec}(\bar{y})$. This completes the description of the
 1398 construction.

1399 B.4.2 Analysis and Proof of Lemma 17

1400 We now turn to prove Lemma 17 and start with several auxiliary claims.

▷ **Claim 39.** Let $r_1, r_2 \in [0, 1]$ such that $|r_1 - r_2| \leq \epsilon^2$ and $\epsilon \leq r_1 \leq 1 - \epsilon$, then

$$|\ln(r_1/(1 - r_1)) - \ln(r_2/(1 - r_2))| \leq 2\epsilon .$$

1401 **Proof.** By the definition of r_1 and r_2 we get the following inequalities:

$$\begin{aligned} 1402 \quad |\ln(r_1/(1 - r_1)) - \ln(r_2/(1 - r_2))| &= |\ln(r_1) - \ln(1 - r_1) - \ln(r_2) + \ln(1 - r_2)| \\ 1403 &\leq \left| \ln\left(\frac{r_1 + \epsilon^2}{r_1}\right) \right| + \left| \ln\left(\frac{1 - r_1 + \epsilon^2}{1 - r_1}\right) \right| \\ 1404 &\leq |\ln(1 + \epsilon^2/r_1)| + |\ln(1 + \epsilon^2/(1 - r_1))| \\ 1405 &\leq 2\ln(1 + \epsilon) \leq 2 \cdot \epsilon , \end{aligned}$$

1406 where the last inequality is due to the Taylor expansion of $\ln(1 + x)$ around 0. ◀

1407 Recall that given $x > 0$ and an integer $k > 0$, $\widehat{\ln}_k(x)$ is the \ln -approximation of x
 1408 obtained by computing the first k terms of the Taylor expansion around a point x_0 where
 1409 $0 \leq x_0 - x \leq 1/2$.

1410 ▷ **Claim 40.** Fix $r_1, r_2 \in [0, 1]$ such that $|r_1 - r_2| \leq \epsilon^2$ and $\epsilon \leq r_1 \leq 1 - \epsilon$, denote
 1411 $\widehat{b}_1 = b + \ln(r_1/(1 - r_1))$ and $\widehat{b}_2 = b + \widehat{\ln}(r_2/(1 - r_2))$. Then, $|\widehat{b}_1 - \widehat{b}_2| \leq 3\epsilon$.

1412 **Proof.** Fix $x \in (0, 1)$. Since $\widehat{\ln}(x)$ is obtained by using the first k terms in the Taylor
 1413 expansion of $\ln(x)$ around x_0 , we have that $|\ln(x) - \widehat{\ln}(x)| = \frac{1}{x_0^k} \cdot \frac{(x - x_0)^k}{k} \cdot \eta^k$, where
 1414 $\eta \in [x, x_0]$. Since $x_0 \geq x$, also $x_0 \geq \eta$. As $|x - x_0| \leq 1/2$, we get that $|\ln(x) - \widehat{\ln}(x)| \leq (1/2)^k$.
 1415 By plugging $k = \Theta(\log 1/\epsilon)$, we have that $|\ln(x) - \widehat{\ln}(x)| \leq \epsilon$ for every x .

1416 Thus, combining with Claim 39 we conclude the following:

$$\begin{aligned} 1417 \quad |\widehat{b}_1 - \widehat{b}_2| &= |b + \ln(r_1/(1 - r_1)) - b - \widehat{\ln}(r_2/(1 - r_2))| & (5) \\ 1418 &\leq |\ln(r_1/(1 - r_1)) - \ln(r_2/(1 - r_2)) + \epsilon| \\ 1419 &\leq \epsilon + |\ln(r_1/(1 - r_1)) - \ln(r_2/(1 - r_2))| \leq 3 \cdot \epsilon . \end{aligned}$$

1420 ◀

1421 ▷ **Claim 41.** Consider two threshold gates g_1, g_2 with the same weighted sum and bias values
 1422 $b_1 \leq b_2$ such that $b_2 - b_1 \leq \epsilon$. Then $|\Pr[g_1 = 1] - \Pr[g_2 = 1]| \leq \sqrt{\epsilon}$.

23:36 The Computational Cost of Asynchronous Neural Communication

1423 **Proof.** Let W be the weighted incoming sum to both g_1 and g_2 . The probability that g_1
 1424 outputs 1 is $1/(1 + e^{-(W-b_1)})$, and the probability that g_2 outputs 1 is $1/(1 + e^{-(W-b_2)})$.
 1425 The following holds:

$$\begin{aligned}
 1426 \quad \Pr[g_2 = 1] &= 1/(1 + e^{-(W-b_2)}) \geq 1/(1 + e^{-(W-b_1-\epsilon)}) = 1/(1 + e^\epsilon e^{-(W-b_1)}) \\
 1427 &\geq \frac{1}{e^\epsilon \cdot (1 + e^{-(W-b_1)})} \geq \frac{1}{(1 + \sqrt{\epsilon}) \cdot (1 + e^{-(W-b_1)})} \\
 1428 &= (1 - \sqrt{\epsilon})(1/(1 + e^{-(W-b_1)})) \geq 1/(1 + e^{-(W-b_1)}) - \sqrt{\epsilon} = \Pr[g_1 = 1] - \sqrt{\epsilon}.
 \end{aligned}$$

In the third inequality we use the fact that $e < (1 + \sqrt{\epsilon})^{\frac{1}{\epsilon}}$ and thus $e^\epsilon < 1 + \sqrt{\epsilon}$. On the other hand, since $b_2 \geq b_1$ it holds that

$$\Pr[g_2 = 1] = 1/(1 + e^{-(W-b_2)}) \leq 1/(1 + e^{-(W-b_1)}) = \Pr[g_1 = 1].$$

1429 Hence, we conclude that $|\Pr[g_2 = 1] - \Pr[g_1 = 1]| \leq \sqrt{\epsilon}$ as required. \blacktriangleleft

1430 **Analysis of Step 1.** In the first step of the construction, since each uniformly random gate r_i
 1431 is 1 with probability $1/2$, the value $\text{dec}(\bar{r})$ is a uniform sample in $\{0, 1, \dots, (1/\epsilon)^4\}$. Therefore,
 1432 $r' = \text{dec}(\bar{r})/\text{dec}(\bar{a}) = \epsilon^4 \cdot \text{dec}(\bar{r})$ is sampled uniformly at random from $\{0, \epsilon^4, 2\epsilon^4, 3\epsilon^4, \dots, 1\}$.
 1433 By a simple coupling argument, sampling r' is equivalent to the process of sampling a uniform
 1434 random variable $r_1 \in [0, 1]$ and rounding it to the closest value of the form $i \cdot \epsilon^4$ for some
 1435 integer i . In this manner, these two samples have an additive distance of at most ϵ^4 .

1436
 1437 **Analysis of Step 2.** Denote the probability z outputs 1 by q , and the probability u
 1438 outputs 1 by p . Recall that g is the probabilistic gate and g' is the output gate of the Boolean
 1439 circuit that approximates g .

1440 In the second step, we compute $\text{dec}(\bar{y}) = q \cdot (b + f(r'))$ where $q = (\text{dec}(\bar{a}) - \text{dec}(\bar{r}))^k$ and
 1441 $f(r') = \widehat{\ln}(r'/(1 - r'))$. Then g' outputs 1 iff $\text{dec}(\bar{y}) \leq W \cdot q$, or simply iff $b + f(r') \leq W$.
 1442 Given that $r' \in [2\epsilon^2, 1 - 2\epsilon^2]$ by Claim 40, $b' = b + f(r')$ satisfies that $|b^* - b'| \leq 3\epsilon^2$ where b^*
 1443 is a true sample from the Logistic distribution with mean b . Therefore, the following holds.

$$\begin{aligned}
 1444 \quad \Pr[g' = 1 \mid r' \in [2\epsilon^2, 1 - 2\epsilon^2]] &= \Pr[W \geq b' \mid r' \in [2\epsilon^2, 1 - 2\epsilon^2]] \\
 1445 &\leq \Pr[W + 3\epsilon^2 \geq b^* \mid r' \in [2\epsilon^2, 1 - 2\epsilon^2]] \\
 1446 &= 1/(1 + e^{-(W-b+3\epsilon^2)}),
 \end{aligned}$$

and in addition

$$\Pr[g' = 1] \geq \Pr[W - 3\epsilon^2 \geq b^* \mid r' \in [2\epsilon^2, 1 - 2\epsilon^2]] = 1/(1 + e^{-(W-b-3\epsilon^2)}).$$

1447 Recall that $\Pr[g = 1] = \frac{1}{1 + e^{-(W-b)}}$. By claim 41 we conclude that $|\Pr[g' = 1] - \Pr[g = 1]| \leq 3\epsilon$.

We note that $r' \in [2\epsilon^2, 1 - 2\epsilon^2]$ with probability at least $1 - 4\epsilon^2$. Hence, we conclude that:

$$\Pr[g' = 1] \leq \Pr[g' = 1 \mid r' \in [2\epsilon^2, 1 - 2\epsilon^2]] + 4\epsilon^2 \leq \Pr[g = 1] + 3\epsilon + 4\epsilon^2 = p + \Theta(\epsilon),$$

1448 and on the other hand:

$$\begin{aligned}
 1449 \quad \Pr[g' = 1] &\geq (1 - 4\epsilon^2) \Pr[g' = 1 \mid r' \in [2\epsilon^2, 1 - 2\epsilon^2]] \\
 1450 &\geq (1 - 4\epsilon^2)(\Pr[g = 1] - 3\epsilon) \geq \Pr[g = 1] - \Theta(\epsilon).
 \end{aligned}$$

1451 Thus, $|\Pr[g = 1] - \Pr[g' = 1]| = O(\epsilon)$ as required.

1452
 1453 **Complexity.** We assume that the bias and weights of the given probabilistic threshold
 1454 gate g are polynomial in $1/\epsilon$. We first claim that with high probability of $1 - \Theta(\epsilon)$, the
 1455 approximate bias sampled from the almost Logistic distribution is also bounded by $\text{poly}(1/\epsilon)$.

1456 ▷ **Claim 42.** Given that $|\mu| = \text{poly}(1/\epsilon)$, for a random variable x drawn from the logistic
 1457 distribution with mean μ it holds that $|x| = \text{poly}(1/\epsilon)$ with probability greater than $1 - \epsilon$.

Proof. By the definition of the Logistic CNF function it holds that

$$\Pr[x > 2 \ln(1/\epsilon) + \mu] = 1 - \frac{1}{1 + e^{-2 \ln(1/\epsilon) - \mu + \mu}} = \frac{\epsilon^2}{1 + \epsilon^2} < \epsilon^2 .$$

On the other hand

$$\Pr[x \leq -2 \ln(1/\epsilon) + \mu] = \frac{1}{1 + e^{2 \ln(1/\epsilon) - \mu + \mu}} = \frac{1}{1 + 1/\epsilon^2} < \epsilon^2 .$$

1458

1459 Thus we can assume from now on that all integer numbers can be representing using
 1460 $O(\text{poly}(1/\epsilon))$ bits. Using circuits for fast integers multiplication as described in [9] and
 1461 iterated addition [38, 31], there exists a Boolean circuit computing $W \cdot q$ as well as $b \cdot q$ using
 1462 $\text{poly}(\Delta, \log(1/\epsilon))$ gates and $\text{poly}(\log \Delta, \log(1/\epsilon))$ depth. When computing the polynomial
 1463 $q \cdot \ln(\frac{r'}{1-r'})$ (of total degree $2k$), calculating each term requires $O(\log k)$ multiplicity operations.
 1464 Since we have k summands, in total we use $k \cdot \log k$ multiplicity operations, each requires
 1465 $O(k \cdot \log k \cdot 2^{O(\log^* k)})$ gates (and depth), and $\log k$ addition operations. The comparison
 1466 circuits uses $\text{poly}(\log 1/\epsilon)$ gates and depth, and the final threshold gate circuit requires
 1467 $\text{poly}(\Delta, \log 1/\epsilon)$ gates and depth $\text{poly}(\log \Delta, \log 1/\epsilon)$. We conclude that the Boolean circuit
 1468 has $\text{poly}(\Delta, \log(1/\epsilon))$ gates and depth of $\text{poly}(\log \Delta, \log(1/\epsilon))$.

1469 B.4.3 Synchronizing a Probabilistic Threshold Gate

1470 In order to construct a synchronized neural network computing the Boolean Circuit described
 1471 in Lemma 17, we use the construction for synchronized Boolean circuits as described in
 1472 Lemma 15. We are left with describing the implementation of the random bits \bar{r} and the
 1473 constant bits \bar{a} .

- 1474 ■ In order to represent \bar{a} , we introduce k neurons a_1, \dots, a_k . If the i^{th} bit in the binary
 1475 representation of $\text{dec}(\bar{a}) = (1/\epsilon)^4$ equals 1 we set the bias of a_i to be $b(a_i) = -1$ and
 1476 otherwise we set the bias to be $b(a_i) = 1$. As a result, the neurons that represent the bits
 1477 that are 1 in the binary representation fire on every round, and the other neurons idle
 1478 thought the execution.
- 1479 ■ In order to represent \bar{r} we introduce k *spiking* neurons r_1, \dots, r_k . For the computation
 1480 to succeed, we need to sample each random variable r_i only once. Therefore, each neuron
 1481 r_i has a very large bias $b(r_i) = \text{poly}(1/\epsilon)$ and an incoming edge from the starter neuron
 1482 s with weight $w(s, r_i) = b(r_i)$. As a result, as long as r_i did not receive a spike from s ,
 1483 with high probability it does not fire. On the other hand, when neuron r_i receives a spike
 1484 from the starter neuron v^* it fires with probability $1/2$.

1485 B.5 The Complete Synchronization Scheme

1486 Finally, we describe the synchronizer for a given neural network and prove Theorem 4. We
 1487 start by describing the construction for a network of deterministic threshold gates. The
 1488 adaptation to a network of spiking neurons is quite straightforward as discussed in the
 1489 end of the section. The construction has two parts: a global pulse generator that can
 1490 be used to synchronize many networks, and an adaptation of the given network \mathcal{N} into a
 1491 network $\text{Sync}(\mathcal{N})$, see Figure 2. The *pulse generator* is implemented by a directed cycle

1492 $PG = [c_0, \dots, c_k]$ of length $k = O(L^4 \log \Delta)$. All neurons in PG have bias $b(c_i) = 1$. In
 1493 addition, for every $i \geq 1$ neuron c_i has an incoming edge from neuron c_{i-1} with weight
 1494 $w(c_{i-1}, c_i) = 1$, and the first neuron c_0 has an incoming edge from the last neuron c_k with
 1495 weight $w(c_0, c_k) = 1$. The last neuron of the chain c_k will declare the end of each phase. We
 1496 assume throughout that the simulation starts by a spike of the starter $v^* = c_0$.

1497

1498 **Modifications to the Network $\text{Sync}(\mathcal{N})$.** The input layer and output layer in $\text{Sync}(\mathcal{N})$
 1499 are *exactly* as in \mathcal{N} . We will now focus on the set of *auxiliary* neurons V in \mathcal{N} . The network
 1500 $\text{Sync}(\mathcal{N})$ contains the vertices V of the original network \mathcal{N} , and in addition, for each neuron
 1501 $v_i \in V$ we add the following components to the network:

- 1502 ■ A synchronized sub-network $\text{Sync}(v_i)$ using Lemma 16 implementing the threshold gate
 1503 defined by neuron v_i . The input neurons to the sub-network $\text{Sync}(v_i)$ are the incoming
 1504 neighbors of v_i in \mathcal{N} . The first neuron v_i^* in the internal chain of the sub-network $\text{Sync}(v_i)$
 1505 has an incoming edge from the L^{th} neuron of PG cycle, namely c_L with weight 1 and
 1506 bias $b(v_i^*) = 1$. Denote the output of the sub-network $\text{Sync}(v_i)$ by v_i^{out} .
- 1507 ■ An AND module AND_i whose output neuron is v_i . This module is implemented by a
 1508 circuit of OR_{sync} and NOT_{sync} gates with three layers (using simple De-morgan rule). The
 1509 AND_i module receives input from the neuron v_i^{out} and from the $(\alpha L^4)^{\text{th}}$ neuron in PG,
 1510 $c_{\alpha L^4}$ where α is a large enough constant. The internal chains of the AND_i circuit receive
 1511 input from neuron c_β in PG where $\beta = \alpha L^4 + L$, making sure the circuit begins the
 1512 execution *after* receiving all its inputs¹⁰.

1513 **Modifications to the Circuit Synchronization of Sec. 4.1.** So far, we handled the
 1514 synchronization of circuits. In order to handle general networks (e.g., that contains self-loops
 1515 and recurrent edges), we need to apply small adaptations to the synchronized sub-networks of
 1516 Sec. 4.1. Specifically, unlike circuits, in the execution of a network, certain neurons (or gates)
 1517 might be activated several times. To be able to re-use the sync. sub-networks throughout
 1518 the execution, we need to reset the states kept by their self-loops.

1519 We therefore adapt the construction of the sync. sub-network presented in Section 4.1
 1520 to reset themselves at the end of their computation. For each NOT_{sync} gate, we augment
 1521 its internal chain by $3 \cdot L^2$ neurons, and the last neuron of this chain is connected to an
 1522 inhibitor neuron v_r . The inhibitor v_r has outgoing edges of weight $-\infty$ to all neurons in the
 1523 sub-network. Due to Claim 30, it holds that the inhibition by v_r (i.e., the round in which
 1524 v_r fires) occurs *after* the output neuron has already fired. Observe that the timing of the
 1525 inhibition by v_r is set in a way that guarantees that all gates in the sub-network will be idle
 1526 from that point on (i.e., there will be no delayed spikes that arrive after this inhibition). For
 1527 the sub-network OR_{sync} which do not contain self-loops, no adaptation is needed.

1528 B.6 Correctness

1529 Throughout, we fix a synchronous execution Π_{sync} and an asynchronous execution Π_{async} .
 1530 For every neuron v and phase p , define the beginning of phase p of v in the asynchronous
 1531 execution ($r(v, p)$) as the round in which the p^{th} spike of c_0 is fired. I.e., the p^{th} phase of v
 1532 is the time interval $[r(v, p), r(v, p + 1))$. For every round p , let $V_{\text{sync}}^+(p)$ be the set of neurons
 1533 that fire in round p in Π_{sync} (i.e., the neurons with positive entries in σ_p). Similarly, let
 1534 $V_{\text{async}}^+(p)$ be the set of neurons that fire during phase p .

¹⁰We say that the circuit receives its input, if every gate in the first layer has received the signals from its incoming input.

1535 ► **Lemma 43.** *The networks $\text{Sync}(\mathcal{N})$ and \mathcal{N} have similar executions.*

1536 In order to show the networks $\text{Sync}(\mathcal{N})$ and \mathcal{N} have similar executions, we show by induction
 1537 on round (resp., phase) p that $V_{\text{sync}}^+(p) = V_{\text{async}}^+(p)$. For $p = 1$, let $V_{\text{sync}}^+(0)$ be the neurons
 1538 that fired at the beginning of the simulation in round 0. We will show that every neuron
 1539 $v_i \in V$ fires in phase 1 iff $v_i \in V_{\text{sync}}^+(1)$. For $v_i \in V$, the spikes from its incoming neighbors
 1540 in $V_{\text{sync}}^+(0)$ reach the sub-network $\text{Sync}(v_i)$ by round L . The global chain in $\text{Sync}(v_i)$ is then
 1541 activated by the neuron c_L in some round $\tau \in [L, L^2]$. Therefore, by Lemma 16 there exists
 1542 a constant γ such that v_i^{out} fires in some round $\tau_i \in [2, \tau + \gamma \cdot \log \Delta \cdot L^4]$ iff the output of the
 1543 threshold function corresponding to v_i is 1, meaning that $v_i \in V_{\text{sync}}^+(1)$. We next note that
 1544 the first layer of the sub-network AND_i consists of two NOT_{sync} sub-networks with input
 1545 from $c_{\alpha \cdot L^4}$ and v_i^{out} . Hence, by Observation 32 as long as AND_i receives the information
 1546 from $c_{\alpha \cdot L^4}$ and v_i^{out} before the activation of the global chain of the network AND_i in some
 1547 round τ^* its output neuron fires by round $\tau^* + O(L^4)$ iff both v_i^{out} and $c_{\alpha \cdot L^4}$ fired.

1548 The global chain of AND_i is activated by neuron c_β for $\beta = \alpha \cdot L^4 + L$ and therefore is
 1549 indeed activated *after* AND_i receives the spike from $c_{\alpha \cdot L^4}$. In addition, we choose α such that
 1550 $\alpha L^4 > L^2 + \gamma \cdot \log \Delta \cdot L^4$. Therefore the neuron c_β fires *after* round $\tau_i + L$, i.e. after AND_i
 1551 received the spike from v_i^{out} as well. We conclude that v_i fires in some round $\tau'' \in [\beta, O(L^4)]$
 1552 iff v_i^{out} fires in round τ_i . We choose k to be large enough to make sure that c_k fires after
 1553 round τ'' and therefore all neurons in $V^+(1)$ fired during the first phase.

1554 Next, we assume that $V_{\text{sync}}^+(p) = V_{\text{async}}^+(p)$ and consider phase $p + 1$. Let τ_p be the round
 1555 that c_0 fired at the beginning of phase p and let τ_{p+1} be the round in which c_0 fired at
 1556 the beginning of phase $p + 1$. In addition, we denote the round in which $c_{\alpha \cdot L^4}$ fired during
 1557 phase p by τ_α . By the induction assumption, neuron v_i fires between round τ_p and round
 1558 τ_{p+1} iff $v_i \in V_{\text{sync}}^+(p)$. Moreover, since the activation of the sub-network AND_i is performed
 1559 by neuron c_β , every $v_i \in V_{\text{sync}}^+(p)$ fires after round τ_α . We choose α to be large enough
 1560 such that by round τ_α , all sub-networks $\text{Sync}(v_i)$ have been reset due to the modification
 1561 in the circuit synchronization. Hence, for neuron $v_i \in V$, the spikes from its incoming
 1562 neighbors in $V_{\text{async}}^+(p)$ reach $\text{Sync}(v_i)$ after the sub-network has already been reset. Thus,
 1563 when the global chain of the sub-network $\text{Sync}(v_i)$ is activated by the neuron c_L in round
 1564 $\tau_L \in [\tau_{p+1} + L, \tau_{p+1} + L^2]$, the sub-network $\text{Sync}(v_i)$ received spikes from the incoming
 1565 neighbors of v_i in $V_{\text{async}}^+(p)$. Combining with Lemma 16 we conclude that v_i^{out} fires in round
 1566 $\tau_i \in [\tau_{p+1} + L, \tau_{p+1} + L^2 + \gamma \cdot \log \Delta \cdot L^4]$ iff $v_i \in V_{\text{sync}}^+(p + 1)$. Thus, when neuron c_β fires in
 1567 phase $p + 1$, the sub-network AND_i has received the spikes from both v_i^{out} and $c_{\alpha \cdot L^4}$. Since
 1568 the global chain of AND_i is activated by the neuron c_β , we conclude that v_i fires in some
 1569 round $\tau^* \in [\tau_{p+1} + \beta, \tau_{p+1} + \Theta(L^4)]$, iff $v_i \in V_{\text{sync}}^+(p + 1)$. Choosing k to be large enough, τ^*
 1570 occurs before c_k fires and ends the phase.

1571
 1572 **Synchronization of a Spiking Neural Network.** We next explain the adaptation
 1573 of the construction given a network of spiking neurons \mathcal{N} . Let n be the number of auxiliary
 1574 neurons in \mathcal{N} and let t be the number of rounds. Each spiking neuron implemented by a
 1575 probabilistic threshold gate can be made synchronized using Cor. 19 where we use an error
 1576 parameter of $\epsilon = 1/\text{poly}(n, t)$. Thus, The network $\text{Sync}(\mathcal{N})$ consists of $\text{poly}(\Delta, \log n \cdot t) \cdot L^4 \cdot n$
 1577 auxiliary neurons and uses $\text{poly}(\log \Delta, \log n \cdot t) \cdot L^5$ rounds.

1578 To compare the simulation of the given spiking neural network \mathcal{N} and the synchronized
 1579 network $\text{Sync}(\mathcal{N})$, we fix the randomness used by \mathcal{N} throughout the simulation and use these
 1580 coins when simulated the network $\text{Sync}(\mathcal{N})$. For neuron $v \in V$ and round $\tau \geq 1$, by Cor. 19,
 1581 with probability at least $1 - 1/\text{poly}(n \cdot t)$ it holds that $v \in V_{\text{sync}}^+(\tau)$ iff $v \in V_{\text{async}}^+(\tau)$. By
 1582 applying the union bound over all n neurons and t rounds of the simulation, we conclude

1583 that with high probability \mathcal{N} and $\text{Sync}(\mathcal{N})$ have similar executions.

1584 **C Synchronization in the Node-Delay Model**

1585 **C.1 Network Dynamics in the Node Delay Setting**

Network evolution proceeds in *seconds*, namely, a sufficiently small time unit. For a given integer $T \geq 1$, the dynamics is specified by a node-delay function $t : V \rightarrow \mathbb{N}_{\leq T}$ interpreted as follows: the round duration on each neuron v consists of $t(v)$ seconds. Specifically, the i^{th} round of v is defined by the time interval $R_i(v) = [(i-1)t(v) + 1, i \cdot t(v)]$ for every $i \geq 1$. All spikes are assumed to arrive with a delay of a single second¹¹. For the neuron v and integer i , the set of spikes received at v during its i^{th} round is given by

$$A(v, i) = \{(u, j \cdot t(u)) \mid j \cdot t(u) + 1 \in R_i(v)\}.$$

1586 The state of v in its i -round (i.e., at the second $i \cdot t(v)$) is given by:

$$1587 \quad \text{pot}(v, i) = \sum_{(u, j \cdot t(u)) \in A(v, i)} w(u, v) \cdot \sigma_j(v) - b(v) \quad \text{and} \quad \sigma_i(v) = 1 \text{ iff } \text{pot}(v, i) \geq 0. \quad (6)$$

1588 If v is a probabilistic threshold gate then it fires in second $i \cdot t(v)$ with probability $p(v, i) =$
1589 $\frac{1}{1 + e^{-\text{pot}(v, i)}}$.

1590 **► Definition 44 (The T -bounded Node-Delay Setting).** *We are given a network \mathcal{N} and an*
1591 *integer T . It is assumed the network contains a special neuron, the starter, that fires in*
1592 *the first round of the simulation. The dynamic is determined by a node-delay function*
1593 *$t : V \rightarrow \mathbb{N}_{\leq T}$. This function t can be chosen arbitrarily.*

1594 **► Definition 45 (Computation of a Boolean Function in the T -bounded Node-Delay Setting).**
1595 *Let $f : \{0, 1\}^n \rightarrow \{0, 1\}^k$ be a Boolean function. A network \mathcal{N} with n input neurons x_1, \dots, x_n*
1596 *and k output neurons z_1, \dots, z_k computes f in this setting if for every T -bounded function*
1597 *$t : V \rightarrow \mathbb{N}_{\leq T}$ and for every fixed possible assignment to the input neurons b_1, \dots, b_n the*
1598 *following holds: (i) If $f_i(b_1, \dots, b_n) = 1$, then there exists a round in which z_i fires, where*
1599 *$f_i(\cdot)$ is the i^{th} bit in the output of f . (ii) If $f_i(b_1, \dots, b_n) = 0$ then z_i does not fire throughout*
1600 *the entire execution.*

1601 **Synchronizers for the Node-Delay.** A synchronizer ν is an algorithm that gets as input
1602 a network \mathcal{N} and integer T , and outputs a network $\mathcal{N}' = \text{sync}_V(\mathcal{N}, T)$ that contains all the
1603 neurons of \mathcal{N} , plus additional auxiliary neurons. One of the auxiliary neurons in \mathcal{N}' is a
1604 *starter* neuron that fires in the *first* round of the simulation. The network \mathcal{N}' works in the
1605 asynchronous setting and should have *similar execution* to \mathcal{N} in the sense that for every
1606 neuron $v \in V(\mathcal{N})$, the firing pattern of v in the asynchronous network should be similar to
1607 the one in the synchronous network. The output network \mathcal{N}' simulates each round of the
1608 network \mathcal{N} by a phase.

1609 **► Definition 46 (Phases).** *We partition the execution of \mathcal{N}' into phases $1, 2, \dots$, using a*
1610 *function $r : V(\mathcal{N}) \times \mathbb{N} \rightarrow \mathbb{N}$ that defines the beginning of phase p . Hence, the p^{th} phase is*
1611 *the round interval $[r(v, p), r(v, p + 1))$.*

¹¹As discussed in the introduction, this model can be generalized to support both edge-delays and node-delays, to isolate the node-delay effect we assume that all edges have latency of 1.

1612 ► **Definition 47** (Similar Executions (Deterministic Networks)). *The synchronous execution*
 1613 *Π of a deterministic network \mathcal{N} is specified by a list of states $\Pi = \{\sigma_1, \dots\}$ where each σ_i*
 1614 *is a binary vector describing the firing status of the neurons in round i . The asynchronous*
 1615 *execution of the network $\mathcal{N}' = \text{sync}_V(\mathcal{N}, T)$ with a node-delay function $t : V \rightarrow \mathbb{N}_{\leq T}$ denoted*
 1616 *by $\Pi'(t)$ is defined analogously only when applying the asynchronous dynamic. The execution*
 1617 *$\Pi'(t)$ is divided into phases according the a function $r : V(\mathcal{N}) \times \mathbb{N} \rightarrow \mathbb{N}$.*

1618 *The network \mathcal{N} and the pair $\langle \mathcal{N}', t \rangle$ have a **similar execution** if $V(\mathcal{N}) \subseteq V(\mathcal{N}')$, and*
 1619 *in addition, a neuron $v \in V(\mathcal{N})$ fires in round p in the execution Π iff v fires during phase p*
 1620 *in $\Pi'(t)$.*

1621 *The networks \mathcal{N} and \mathcal{N}' are **similar** if \mathcal{N} and $\langle \mathcal{N}', t \rangle$ have a similar execution for every*
 1622 *node-delay function t .*

1623 As for the edge-delay model, the extension for randomized networks is made by fixing
 1624 the random bits in the simulation of the input network.

1625 C.2 Reduction to the Edge-Delay Model: A Simulation Result

1626 Given a neural network \mathcal{N} and an integer parameter T , our goal is to construct a network
 1627 $\mathcal{N}_R = \text{sync}_V(\mathcal{N}, T)$ in the T -bounded node-delay model that behaves similarly to \mathcal{N} , i.e.,
 1628 that \mathcal{N} and \mathcal{N}_R are similar according to Definition 47.

1629 Given the network \mathcal{N} and the delay bound T , we start by computing the network
 1630 $\mathcal{N}_{\mathcal{L}} = \text{sync}_E(\mathcal{N}, L)$ with $L = 5T^2$. The desired $\mathcal{N}_R = \text{sync}_V(\mathcal{N}_T)$ is obtain by changing some
 1631 of the edge weights in $\mathcal{N}_{\mathcal{L}}$. Our proof of correctness is based on similarity between a network
 1632 in the node-delay model and a network in the edge-delay model.

1633 **Similarity between the networks \mathcal{N}_R and $\mathcal{N}_{\mathcal{L}}$.** Fix integer parameters T, L . Given an
 1634 edge-delay network $\mathcal{N}_{\mathcal{L}}$, a latency function $\ell : E(\mathcal{N}_{\mathcal{L}}) \times \mathbb{N} \rightarrow \mathbb{N}_{\leq L}$, a node-delay network \mathcal{N}_R
 1635 on the same neuron set and a node-delay function $t : V(\mathcal{N}_R) \rightarrow \mathbb{N}_{\leq T}$, we want to define
 1636 similarity between the simulations $\langle \mathcal{N}_{\mathcal{L}}, \ell \rangle$ and $\langle \mathcal{N}_R, t \rangle$, where both simulations use the same
 1637 initial configuration.
 1638

This notion of similarity is based on defining different time scales in each of the simulations. Specifically, for every $i \geq 1$ and neuron $u \in V$ the time window $R_i(u)$ will be the time that u collects spikes for its round i in the simulation of $\langle \mathcal{N}_R, t \rangle$. Moreover, for every $i \geq 0$ the time window $L_i(u)$ correspond to the firing period of round i of u in the simulation of $\langle \mathcal{N}_{\mathcal{L}}, \ell \rangle$, where

$$R_i(u) = [(i-1) \cdot t(u) + 1, i \cdot t(u)] \text{ and } L_i(u) = [i \cdot T \cdot t(u), i \cdot T \cdot t(u) + (T \cdot t(u) - 1)].$$

Furthermore, for every second τ_R in the simulation of $\langle \mathcal{N}_R, t \rangle$ we will have the corresponding block $B_{\tau_R} = [\tau_R \cdot T, \tau_R \cdot T + (T - 1)]$ in the simulation of $\langle \mathcal{N}_{\mathcal{L}}, \ell \rangle$. For the simulation of $\langle \mathcal{N}_{\mathcal{L}}, \ell \rangle$ define for every neuron u and $i \geq 0$:

$$\sigma_i(u, \mathcal{N}_{\mathcal{L}}) = \begin{cases} 1 & u \text{ is strong and } u \text{ fires in every } \tau_{\mathcal{L}} \in L_i(u) \\ 1 & u \text{ is weak and } u \text{ fires in } \tau_{\mathcal{L}} \in L_i(u) \text{ only for } \tau_{\mathcal{L}} = i \cdot T \cdot t(u) \\ 0 & u \text{ never fires in } L_i(u) \\ \emptyset & \text{Otherwise.} \end{cases}$$

For the simulation of $\langle \mathcal{N}_R, t \rangle$ define for every neuron u and $i \geq 0$:

$$\sigma_i(u, \mathcal{N}_R) = \begin{cases} 1 & u \text{ fires in round } i \text{ of } u \text{ (i.e. in the second } i \cdot t(u)) \\ 0 & \text{Otherwise.} \end{cases}$$

1639 ► **Definition 48.** The simulations $\langle \mathcal{N}_R, t \rangle, \langle \mathcal{N}_L, \ell \rangle$ are similar, denoted as $\langle \mathcal{N}_R, t \rangle \sim \langle \mathcal{N}_L, \ell \rangle$,
 1640 if for every neuron u and $i \geq 0$ it holds that $\sigma_i(u, \mathcal{N}_L) = \sigma_i(u, \mathcal{N}_R)$.

1641 A network \mathcal{N}_L in the L -bounded edge-delay model and a network \mathcal{N}_R in the T -bounded node-
 1642 delay model are similar, denoted by $\mathcal{N}_L \sim \mathcal{N}_R$, if for every node-delay function $t : V(\mathcal{N}_R) \rightarrow$
 1643 $\mathbb{N}_{\leq T}$ there exists a latency function $\ell : E(\mathcal{N}_L) \times \mathbb{N} \rightarrow \mathbb{N}_{\leq L}$ such that $\langle \mathcal{N}_R, t \rangle \sim \langle \mathcal{N}_L, \ell \rangle$.

1644 The key simulation lemma used in the synchronization scheme is as follow:

1645 ► **Lemma 49.** Given a network \mathcal{N}_L in the L -bounded edge delay model such that:

- 1646 1. $b(u) > 0$ for every neuron u .
- 1647 2. Every weak neuron v has no self-loop.
- 1648 3. There is no edge from a strong neuron to a strong neuron.
- 1649 4. Every negative edge has weight $-\infty$.
- 1650 5. For every neuron u , either any excitatory incoming neighbor of u is weak, or any excitatory
 1651 incoming neighbor of u is strong.
- 1652 6. Let v be a strong incoming neighbor of a neuron u , and let f be an inhibitor. Then if f
 1653 has an edge to v , it also has an edge to u .

1654 Then there exists a network \mathcal{N}_R in the T -bounded node-delay model with $T \leq \sqrt{L/5}$ with
 1655 $V(\mathcal{N}_R) = V(\mathcal{N}_L)$ such that \mathcal{N}_R and \mathcal{N}_L are similar.

Defining the node-delay network \mathcal{N}_R . The network \mathcal{N}_R is exactly as \mathcal{N}_L , up to small
 adaption of the weights. Denote by $w_L : V \rightarrow \mathbb{R}$ the weight function of the network \mathcal{N}_L .
 Define the weight function w_R of \mathcal{N}_R as

$$w_R(v, u) = \begin{cases} T \cdot w_L(v, u) & v \neq u, v \text{ is strong,} \\ w_L(v, u) & \text{Otherwise.} \end{cases}$$

1656

1657

1658 **Correctness.** We will show that \mathcal{N}_L and \mathcal{N}_R are similar. Fix a node-delay function
 1659 $t : V \rightarrow \mathbb{N}_{\leq T}$. First, we define the corresponding latency function ℓ and prove it is valid, i.e.
 1660 that ℓ is nice and $\ell(v, u, \tau) \in [1, L]$ for every neurons v, u and round τ . Then, we restate
 1661 Lemma 49 in order to prove its correctness by induction on the round.

1662

1663 **Definition of the latency function ℓ .** First, set the latency of self-spikes to be of
 1664 value 1. For a neuron u , we say that u is *weak-incoming* if any excitatory incoming neighbor
 1665 of u is weak, and we say that u is *strong-incoming* if any excitatory incoming neighbor of u is
 1666 strong. Note that by property 5, every neuron u is either weak-incoming or strong-incoming.
 1667 For a strong-incoming neuron u , an inhibitor v and $\tau_L \geq 0$, set $\ell(v, u, \tau_L) = 2T^2 + 1$. Now
 1668 consider the remaining spikes, which are either positive spikes, or spikes to a weak-incoming
 1669 neuron u .

1670 For every $\tau_L \geq 0$ define the latency value for the spike event $\langle v, u, \tau_L \rangle$ as follows. Let j
 1671 be an integer satisfying that $\tau_L \in L_j(v)$, and let i be such that $j \cdot t(v) + 1 \in R_i(u)$, hence
 1672 $(v, j \cdot t(v)) \in A(u, i)$.

1673 If v is weak, then for $\tau_L = j \cdot T \cdot t(v)$ set $\ell(v, u, \tau_L) = i \cdot T \cdot t(u) - \tau_L$. That is, the
 1674 spike $\langle v, u, \tau_L \rangle$ is scheduled to arrive in the first round of $L_i(u)$. For $\tau_L > j \cdot T \cdot t(v)$, set
 1675 $\ell(v, u, \tau_L) = 1$. Otherwise, if v is strong, consider the following argument. For every second
 1676 τ_L in the edge-latency simulation, let τ_R be the second in the node-delay simulation such
 1677 that $\tau_L \in B_{\tau_R}$.

1678 ■ **Case (I):** there exists a second in $[\tau_R + 1, \tau_R + 2T]$ such that u fires in the node-delay
 1679 simulation, let τ'_R be the first such second. Set $\ell(v, u, \tau_L) = \tau'_R \cdot T - \tau_L$, that is schedule
 1680 $\langle v, u, \tau_L \rangle$ to arrive in round $\tau'_R \cdot T$.

1681 ■ **Case (II):** case I does not apply, and there is an inhibitor f which is an incoming
 1682 neighbor of u , and a second $\tau'_R \in [\tau_R - T, \tau_R + 2T]$ such that f fires in τ'_R in the node-
 1683 delay simulation. Then for such τ'_R , set $\ell(v, u, \tau_{\mathcal{L}}) = \tau'_R \cdot T + (2T^2 + 1) - \tau_R$, that is
 1684 schedule $\langle v, u, \tau_{\mathcal{L}} \rangle$ to arrive in round $\tau'_R \cdot T + (2T^2 + 1)$.

1685 ■ **Case (III):** neither case (I) nor case (II) apply. Set $\ell(v, u, \tau_{\mathcal{L}}) = 1$.

1686 The intuition is that for a positive spike in the edge-delay simulation, we look for a round
 1687 such that u is supposed to fire in the next $2T^2$ rounds. If we cannot find one, we want to
 1688 send the spike to a round that we know it will not activate u . This is a round in which u
 1689 receives a negative spike (since negative spikes are of weight $-\infty$). If such round also does
 1690 not exist, it implies that the total weight of positive incoming neighbors of u that fired in
 1691 round $\tau_{\mathcal{L}}$ is low, and we can schedule all these spikes to arrive together in $\tau_{\mathcal{L}} + 1$ without
 1692 activating u . We next show that ℓ is valid.

1693 ▷ **Claim 50.** ℓ is a valid latency function for $\mathcal{N}_{\mathcal{L}}$.

1694 **Proof.** First, since all self-spikes have latency value 1, ℓ is nice. For a negative spike $\langle v, u, \tau_{\mathcal{L}} \rangle$
 1695 such that u is strong-incoming, it holds that $\ell(v, u, \tau) = 2T^2 + 1 < L$. Therefore we are
 1696 left to show validity for positive spikes, and for negative spikes that are fired towards a
 1697 weak-incoming neuron. Consider a spike $\langle v, u, \tau_{\mathcal{L}} \rangle$, and let j be an integer satisfying that
 1698 $\tau_{\mathcal{L}} \in L_j(v)$. Furthermore, let i be an integer such that $j \cdot t(v) + 1 \in R_i(u)$.

1699 Next, assume that v is weak. We distinguish between two cases depending whether $\tau_{\mathcal{L}}$ is
 1700 the first round in the block or not. For $\tau_{\mathcal{L}} = i \cdot T \cdot t(v)$ we have $\ell(v, u, \tau_{\mathcal{L}}) = i \cdot T \cdot t(u) - \tau_{\mathcal{L}}$.
 1701 Recall that $R_i(u) = [(i-1) \cdot t(u) + 1, i \cdot t(u)]$, thus $j \cdot t(v) + 1 \leq i \cdot t(u)$, and $\ell(v, u, \tau_{\mathcal{L}}) =$
 1702 $T \cdot (i \cdot t(u) - j \cdot t(v)) \geq T \geq 1$. Furthermore $j \cdot t(v) + 1 \geq (i-1) \cdot t(u) + 1$, hence
 1703 $i \cdot t(u) - j \cdot t(v) \leq t(u) \leq T$, and $\ell(v, u, \tau_{\mathcal{L}}) \leq T \cdot (i \cdot t(u) - j \cdot t(v)) \leq L$. Otherwise, i.e. for
 1704 $\tau_{\mathcal{L}} \geq i \cdot T \cdot t(v)$, it holds that $\ell(v, u, \tau_{\mathcal{L}}) = 1$, and thus $\ell(v, u, \tau_{\mathcal{L}}) \in [1, L]$.

1705 It remains to consider the case where v is strong. Let τ_R be the second in the node-
 1706 delay simulation such that $\tau_{\mathcal{L}} \in B_{\tau_R}$. Consider the definition of ℓ for a spike $\langle v, u, \tau_{\mathcal{L}} \rangle$.
 1707 In case (I), we have $\ell(v, u, \tau_{\mathcal{L}}) = \tau'_R \cdot T - \tau_{\mathcal{L}}$, and since $\tau'_R \in [\tau_R + 1, \tau_R + 2T]$ it holds
 1708 that $1 \leq \tau'_R \cdot T - \tau_{\mathcal{L}} \leq 2T^2 < L$. In case (II), since $\tau'_R \in [\tau_R - T, \tau_R + 2T]$, we have
 1709 that $\ell(v, u, \tau_{\mathcal{L}}) = \tau'_R \cdot T + (2T^2 + 1) - \tau_{\mathcal{L}} \in [1, 5T^2]$. Finally, in case (III) we simply have
 1710 $\ell(v, u, \tau_{\mathcal{L}}) = 1$. Hence, in all cases it holds that $\ell(v, u, \tau_{\mathcal{L}}) \in [1, L]$. ◀

1711 In order to show that $\langle \mathcal{N}_R, t \rangle \sim \langle \mathcal{N}_{\mathcal{L}}, \ell \rangle$, we restate the condition for similarity in the following
 1712 lemma. We then prove the lemma by induction on the round $\tau_{\mathcal{L}}$.

1713 ► **Lemma 51 (Restating Lemma 49).** *For every round $\tau_{\mathcal{L}} \geq 0$ of the simulation $\langle \mathcal{N}_{\mathcal{L}}, \ell \rangle$ and*
 1714 *for every neuron u , let i be such that $\tau_{\mathcal{L}} \in L_i(u)$. Then:*

- 1715 1. *If $\sigma_i(u, \mathcal{N}_R) = 1$:*
 - 1716 ■ *If $\tau_{\mathcal{L}} = i \cdot T \cdot t(u)$ then u fires in $\tau_{\mathcal{L}}$.*
 - 1717 ■ *If $\tau_{\mathcal{L}} > i \cdot T \cdot t(u)$ then u fires iff u is strong.*
- 1718 2. *If $\sigma_i(u, \mathcal{N}_R) = 0$ then u does not fire in $\tau_{\mathcal{L}}$.*

1719 For the base case $\tau_{\mathcal{L}} = 0$, the correctness follows the fact that both simulations have the
 1720 same starting configuration. Now, let $\tau_{\mathcal{L}} \geq 1$ and assume correctness for every $\tau'_{\mathcal{L}} \leq \tau_{\mathcal{L}} - 1$.
 1721 Fix a neuron u and let i be an integer such that $\tau_{\mathcal{L}} \in L_i(u)$. We start with a useful auxiliary
 1722 claim.

1723 ▷ **Claim 52.** Let u be a weak-incoming neuron, v an incoming neighbor of u , and $\tau'_{\mathcal{L}} \geq 0$.
 1724 Furthermore, let j be such that $\tau'_{\mathcal{L}} \in L_j(v)$, and i such that $j \cdot t(v) + 1 \in R_i(u)$. Then the
 1725 spike $\langle v, u, \tau'_{\mathcal{L}} \rangle$ occurs and arrives to u in round $\tau_{\mathcal{L}} = i \cdot T \cdot t(u)$ in the simulation $\langle \mathcal{N}_{\mathcal{L}}, \ell \rangle$ iff

1726 $\tau'_\mathcal{L} = j \cdot t(v)$ and the spike $\langle v, u, j \cdot t(v) \rangle$ occurs and arrives to u in $R_i(u)$ in the simulation
1727 $\langle \mathcal{N}_R, t \rangle$.

1728 **Proof of Claim 52.** Since u is weak-incoming v is weak, then by the induction assumption
1729 for $\tau'_\mathcal{L}$ and the definition of ℓ , the spike event $\langle v, u, \tau' \rangle$ occurs and arrives in round $\tau_\mathcal{L}$ iff there
1730 exists j such that $\tau'_\mathcal{L} = j \cdot T \cdot t(v)$ and $\sigma_j(v, \mathcal{N}_R) = 1$. This happens iff in the simulation of
1731 $\langle \mathcal{N}_R, t \rangle$ the spike event $\langle v, u, j \cdot t(v) \rangle$ occurs and arrives to u in $R_i(u)$. ◀

1732 We split the proof of Lemma 51 into two cases.

1733 **Case 1: u is weak-incoming.** Assume $\tau_\mathcal{L} = i \cdot T \cdot t(u)$, we want to show that u fires in round
1734 $\tau_\mathcal{L}$ iff $\sigma_i(u, \mathcal{N}_R) = 1$. By Claim 52, we get that the mapping $\langle v, u, j \cdot T \cdot t(v) \rangle \mapsto \langle v, u, j \cdot t(v) \rangle$
1735 is a bijection between the set of non self-spikes that u receives in $\tau_\mathcal{L}$ in the simulation $\langle \mathcal{N}_\mathcal{L}, \ell \rangle$
1736 and the set of non-self spikes that u receives in $R_i(u)$ in the simulation $\langle \mathcal{N}_R, t \rangle$. As for
1737 self-spikes, note that if u is weak it has no self-loop. If u is strong, then by the induction
1738 assumption u fires in $\tau_\mathcal{L} - 1$ iff $\sigma_{i-1}(u, \mathcal{N}_R)$. Thus, u receives the self-spike $\langle u, u, \tau_\mathcal{L} - 1 \rangle$ in
1739 $\tau_\mathcal{L}$ iff it receives the self-spike $\langle u, u, (i-1) \cdot T \cdot t(u) \rangle$ in $R_i(u)$. Since $w_\mathcal{L}(v, u) = w_R(v, u)$
1740 for every weak neuron v and for $v = u$, we get that the total spike weight that u receives
1741 in $\tau_\mathcal{L}$ equals to the total spike weight it receives in $R_i(u)$. Thus, u fires in round $\tau_\mathcal{L}$ iff
1742 $\sigma_i(u, \mathcal{N}_R) = 1$.

1743 Now, assume $\tau_\mathcal{L} > i \cdot T \cdot t(u)$ and that either v is weak, or v is strong and $\sigma_i(u, \mathcal{N}_R) = 0$.
1744 We want to show that u does not fire. Note that if v is weak then it has no self-loop, and if
1745 v is strong and $\sigma_i(u, \mathcal{N}_R) = 0$ then by the induction assumption for $\tau_\mathcal{L} - 1$, u does not fire
1746 in $\tau_\mathcal{L} - 1$. Thus, in both cases u does not receive a self-spike in $\tau_\mathcal{L}$. Furthermore, u has no
1747 strong neighbors, therefore by Claim 52 u does not receive any positive spikes from incoming
1748 neighbors. Since $b(u) > 0$, u does not fire in $\tau_\mathcal{L}$.

1749 Finally, assume $\tau_\mathcal{L} > i \cdot T \cdot t(u)$, and assume u is strong and $\sigma_i(u, \mathcal{N}_R) = 1$. We want
1750 to show that u fires. Note that by Claim 52, u does not receive a negative spike in $\tau_\mathcal{L}$.
1751 Furthermore, since $\sigma_i(u, \mathcal{N}_R) = 1$ by the induction assumption for $\tau_\mathcal{L} - 1$, u fires in $\tau_\mathcal{L} - 1$
1752 and therefore u receives a self-spike in $\tau_\mathcal{L}$. Since $w_\mathcal{L}(u, u) \geq b(u)$, u fires in $\tau_\mathcal{L}$.

1753 **Case 2: u is strong-incoming.** By the properties of $\mathcal{N}_\mathcal{L}$ there is no edge between strong
1754 neurons, and weak neurons have no self-loop. Hence u is weak and has no self-loop. We
1755 handle separately the following sub-cases:

1756 **Case 2.1: $\sigma_i(u, \mathcal{N}_R) = 1$ and $\tau_\mathcal{L} = i \cdot T \cdot t(u)$.** We want to show that u fires in $\tau_\mathcal{L}$. Let
1757 $\langle v, u, j \cdot t(v) \rangle$ be a positive spike in the simulation $\langle \mathcal{N}_R, t \rangle$ that arrives to u in $R_i(u)$, and
1758 let $\tau'_\mathcal{L}$ be one of the T rounds $[j \cdot T \cdot t(v), j \cdot T \cdot t(v) + (T-1)]$. Since v is strong then by
1759 the induction assumption for $\tau'_\mathcal{L}$ v fires in $\tau'_\mathcal{L}$, and therefore the spike event $\langle v, u, \tau'_\mathcal{L} \rangle$ occurs
1760 in the simulation $\langle \mathcal{N}_\mathcal{L}, \ell \rangle$. We now show that $\langle v, u, \tau'_\mathcal{L} \rangle$ arrives to u in $\tau_\mathcal{L}$, according to the
1761 definition of ℓ for spikes from strong neurons.

1762 Since $\sigma_i(u, \mathcal{N}_R) = 1$, u fires in the second $r_R = i \cdot t(u)$ in the simulation $\langle \mathcal{N}_R, t \rangle$.
1763 Note that $j \cdot t(v) + 1 \in R_i(u)$ implies that $i \cdot t(u) - j \cdot t(v) \leq T$. Hence in particular
1764 $i \cdot t(u) \in [j \cdot t(v) + 1, j \cdot t(v) + 2T]$. Let $r'_R \in [j \cdot t(v) + 1, j \cdot t(v) + 2T]$ with $r'_R < i \cdot t(u)$.
1765 Note that $r'_R \in R_i(u)$, therefore r'_R is not an end of a round of u . Hence u does not fire in
1766 r'_R . Therefore the second $r_R = i \cdot t(u)$ is the first second in $[j \cdot t(v) + 1, j \cdot t(v) + 2T]$ that u
1767 fires, and due to the definition of ℓ the spike $\langle v, u, \tau'_\mathcal{L} \rangle$ arrives in round $\tau_\mathcal{L}$.

1768 Now, let f be an inhibitory incoming neighbor of u . By the definition of ℓ , a spike from
1769 f to u can arrive only in a round of the form $\tau'_R \cdot T + T^2 + 1$ for some second τ'_R , which is
1770 not a multiplicity of T . Note that $\tau_\mathcal{L} = i \cdot T \cdot t(u)$ is a multiplicity of T . Thus u does not
1771 receive a negative spike in $\tau_\mathcal{L}$.

1772 We get that in round $\tau_\mathcal{L}$, u receives only positive spikes in $\tau_\mathcal{L}$, and for every positive spike

1773 $\langle v, u, j \cdot t(v) \rangle$ that arrives to u in $R_i(u)$ and every $\tau'_L \in [j \cdot T \cdot t(v), j \cdot T \cdot t(v) + (T-1)]$, u receives a
 1774 spike $\langle v, u, \tau'_L \rangle$. Since $w_R(v, u) = T \cdot w_L(v, u)$ for every strong v and $[j \cdot T \cdot t(v), j \cdot T \cdot t(v) + (T-1)]$
 1775 contains T rounds, we get that the total spike weight that u receives in τ_L is at least the
 1776 total spike weight it receives in $R_i(u)$ in the node-delay simulation. Since $\sigma_i(u, \mathcal{N}_R) = 1$, u
 1777 receives in $R_i(u)$ a spike weight of at least $b(u)$, which implies the same for round τ_L in the
 1778 edge-delay simulation. Thus u fires in round τ_L .

1779 **Case 2.2:** $\sigma_i(u, \mathcal{N}_R) = 0$ or $\tau_L > i \cdot T \cdot t(u)$. We want to show that u does not fire in τ_L .
 1780 Towards contradiction, assume that it does. First note that if u receives no positive spikes
 1781 in τ_L , then since $b(u) > 0$ u does not fire in τ_L . Otherwise, let $\langle v, u, \tau'_L \rangle$ be a positive spike
 1782 that arrives to u in round τ_L . Recall that since v is strong, there are three cases for defining
 1783 the latency value of $\langle v, u, \tau'_L \rangle$.

1784 We will now show that $\langle v, u, \tau'_L \rangle$ belongs to case (II). It does not belong to case (I), since
 1785 it does not hold that $\sigma_i(u) = 1$ and $\tau_L = i \cdot T$. If we are in case (II), then there exists an
 1786 inhibitor f which is connected to u that fired in second τ'_R in the node-delay simulation that
 1787 arrived in τ_L , i.e. such that $\tau_L = \tau'_R \cdot T + (T^2 + 1)$. By the induction assumption for $\tau'_R \cdot T$,
 1788 u fires in round $\tau'_R \cdot T$ in the edge-delay simulation, and since u is strong-incoming then by
 1789 the definition of ℓ the spike $\langle f, u, \tau'_R \cdot T \rangle$ arrives to u in round $\tau'_R \cdot T + (T^2 + 1) = \tau_L$. Since
 1790 negative spikes are of weight $-\infty$, u does not fire in τ_L . Therefore, $\langle v, u, \tau'_L \rangle$ belongs to case
 1791 (III).

1792 By the definition of case (III), $\langle v, u, \tau'_L \rangle$ was generated in round $\tau'_L = \tau_L - 1$. Let j be
 1793 such that $\tau_L - 1 \in L_j(v)$, and let τ_R such that $\tau_L - 1 \in B_{\tau_R}$. Furthermore, let v be an
 1794 excitatory incoming neighbor that fires in $\tau_L - 1$, let j be such that $\tau_L - 1 \in L_j(v)$, and
 1795 let τ_R such that $\tau_L - 1 \in B_{\tau_R}$. Our goal is to show that v fires in $[\tau_R + 1, \tau_R + T]$ in the
 1796 node-delay simulation, by showing that it receives enough positive spikes from its neighbors
 1797 in this interval.

Let f be an inhibitor that has an edge to v . By the network properties f also has an
 edge to u , and since we are not in case (II) in the definition of ℓ , f does not fire in the
 interval $[\tau_R - T, \tau_L + 2T]$ in the node-delay simulation. This implies that v does not receive
 a negative spike in $[\tau_R - T + 1, \tau_R + 2T + 1]$. Notice that $\tau_R \in [j \cdot t(v), (j + 1) \cdot t(v) - 1]$, and
 since $t(v) \leq T$ we get

$$R_{j+1}(v) = [j \cdot t(v) + 1, (j + 1) \cdot t(v)] \subseteq [\tau_R - T + 1, \tau_R + 2T + 1].$$

1798 Therefore, v does not receive a negative spike in $R_{j+1}(v)$.

1799 By the induction assumption for $\tau_L - 1$ we have $\sigma_j(v, \mathcal{N}_R) = 1$. Together with the fact that
 1800 v is strong and receives no negative spikes in $R_{j+1}(v)$, we get that $\sigma_{j+1}(v, \mathcal{N}_R) = 1$, i.e. v
 1801 fires in the node-delay simulation in the second $(j + 1) \cdot t(v)$. This implies that u receives a
 1802 spike from v in $(j + 1) \cdot t(v) + 1$, which is inside the interval $[\tau_R + 1, \tau_R + T]$. If so, let W
 1803 the total weight of the incoming neighbors of u that fired in round $\tau_L - 1$. Since u fires in
 1804 round τ_L , it holds that $W \geq b(u)$. We will show this implies that u fires in some round in
 1805 $[\tau_R + 1, \tau_R + 2T]$, which contradicts the fact that none of the arriving spikes belong to case I.

We showed that for every neuron v that fires in $\tau_L - 1$ in the edge-latency simulation, u
 receive a spike from v in some round $\tau'_R \in [\tau_R + 1, \tau_R + T]$ in the node-delay simulation. By
 the definition of w_R it holds that $w_R(v, u) = T \cdot w_L(v, u)$, and therefore we get

$$\sum_{\tau'_R = \tau_R + 1}^{\tau_R + T} W_{\tau'_R} \geq T \cdot W.$$

1806 By an averaging argument there is a second $\tau'_R \in [\tau_R + 1, \tau_R + T]$ with $W_{\tau'_R} \geq (T \cdot W) / T = W$.

1807 Let i' be an integer such that $\tau'_R \in R_{i'}(u)$. Therefore u receive in $R_{i'}(u)$ a total positive
 1808 spike weight of at least $W \geq b(u)$. Furthermore, since no spike belongs to case C.2, u
 1809 do not receive a negative spike in $R_{i'}(u) \subseteq [\tau_R + 1, \tau_R + 2T]$. Thus, u fires in the second
 1810 $i' \cdot t(u) \in [\tau_R + 1, \tau_R + 2T]$, a contradiction.

1811 C.3 The Complete Synchronization Scheme

1812 We are now ready to complete the proof of Theorem 5. We consider a neural network \mathcal{N}
 1813 and an integer parameter T . Set $L = 5T^2$ and let $\mathcal{N}_{\mathcal{L}} = \text{sync}_E(\mathcal{N}, L)$ be the synchronized
 1814 network of \mathcal{N} in the L -bounded node-delay model. We will now show that $\mathcal{N}_{\mathcal{L}}$ satisfies the
 1815 properties in the conditions of Lemma 49.

1816

1817 **Showing that $\mathcal{N}_{\mathcal{L}}$ satisfies the properties of Lemma 49.**

1818 Note that by the definition of the edge-delay synchronization scheme given in Section 4.3,
 1819 every neuron $u \in \mathcal{N}_{\mathcal{L}}$ is contained in one of the following modules: (1) an OR_{sync} or NOT_{sync}
 1820 subnetwork (Section 4.1), (2) a chain of a threshold gate which is implemented as a boolean
 1821 circuit subnetwork (Section B.3), or (3) the chain of the global pulse generator (Section 4.3).
 1822 By the definitions of these modules, properties 1 and 2 hold. Moreover, together with the fact
 1823 that edges between the modules connect only weak excitatory neurons, we also get property
 1824 3. Furthermore, note that the only inhibitors in the network are r and v_r neurons (which is
 1825 later added in 4.3) in the NOT_{sync} module, and all their edges have weight $-\infty$. Therefore,
 1826 property 4 is satisfied.

1827 The remaining properties 5 and 6 are relevant only for strong neurons. Therefore, consider
 1828 the NOT_{sync} module (Lemma 14), which is the only module that contains strong neurons.
 1829 By the module definition, there are two possible types of strong neurons: (i) the memory
 1830 neuron m , that is only connected to the reset neuron r ; and (ii) intermediate neuron v_i , that
 1831 is only connected to the output neuron z . In case (i), v has only one incoming inhibitor,
 1832 which is the neuron v_r that resets the whole network after it finishes. Thus v_r also has an
 1833 edge to r . In case (ii), v has two incoming inhibitors, v_r and r , which both have an edge to
 1834 z . Therefore property 6 holds. Furthermore, both r and z have no edges from weak neurons.
 1835 Hence, property 5 holds.

1836 Indeed, $\mathcal{N}_{\mathcal{L}}$ satisfies the conditions of Lemma 49, and therefore there exists a network
 1837 \mathcal{N}_R in the T -bounded node delay model which is similar to $\mathcal{N}_{\mathcal{L}}$. We are left to show the
 1838 transitivity of similarity, i.e. that if $\mathcal{N} \sim \mathcal{N}_{\mathcal{L}}$ and $\mathcal{N}_{\mathcal{L}} \sim \mathcal{N}_R$, also $\mathcal{N} \sim \mathcal{N}_R$.

1839

1840 **Showing transitivity of similarity.**

1841 Let t be a node-delay function for \mathcal{N}_R . First, by the similarities of the networks we get
 1842 $V(\mathcal{N}) = V(\mathcal{N}_{\mathcal{L}}) = V(\mathcal{N}_R)$. Moreover, by the definition of $\mathcal{N}_{\mathcal{L}} \sim \mathcal{N}_R$ there exists a latency
 1843 function ℓ for $\mathcal{N}_{\mathcal{L}}$ such that $\langle \mathcal{N}_{\mathcal{L}}, \ell \rangle \sim \langle \mathcal{N}_R, t \rangle$. Let Π be the execution of \mathcal{N} , $\Pi_{\mathcal{L}}$ be the
 1844 execution of $\langle \mathcal{N}, \ell \rangle$, and Π_R be the execution of $\langle \mathcal{N}, t \rangle$. Let the interval $[r_{\mathcal{L}}(v, p), r_{\mathcal{L}}(v, p + 1))$
 1845 be the p^{th} phase of $\Pi_{\mathcal{L}}$, and define $[r_R(v, p), r_R(v, p + 1))$ as the p^{th} phase of Π_R , where the
 1846 definition of $r_R(v, p)$ is as follows. Let $L_p^*(v)$ be the earliest block $L_i(v)$ whose first round τ_p^*
 1847 is contained in phase p of $\Pi_{\mathcal{L}}$, then $r_R(v, p) = \tau_p^*/T$. We wish to prove the following claim.

1848 \triangleright **Claim 53.** For every neuron v and $p \geq 0$, v fires in round p of Π iff v fires in phase p of
 1849 Π_R .

1850 First, note that by the construction of $\mathcal{N}_{\mathcal{L}} = \text{sync}_{\mathcal{L}}(\mathcal{N}, \mathcal{L})$, every neuron $v \in V(\mathcal{N})$ can fire
 1851 only after the chain neuron $c_{\alpha L^4 + L}$ fires. Since $\alpha L^4 > t(v) \cdot T$ this implies that v does not
 1852 fire in the first $t(v) \cdot T$ rounds of each phase in $\Pi_{\mathcal{L}}$. We prove the two directions of the claim.

- 1853 ■ Assume neuron v fires in round p in Π . Because $\mathcal{N} \sim \langle \mathcal{N}_{\mathcal{L}}, \ell \rangle$ there is a round $\tau_{\mathcal{L}}$ in
 1854 phase p of $\Pi_{\mathcal{L}}$ where v fires. Since v does not fire in the first $t(v) \cdot T$ rounds of each
 1855 phase we have $\tau_{\mathcal{L}} \geq r_{\mathcal{L}}(v, p) + t(v) \cdot T$. Since $L_j(v)$ consists of $t(v) \cdot T$ rounds, the first
 1856 round of $L_j(v)$ is in phase p . Therefore, $j \cdot t(v) \cdot T \geq \tau^*$, and therefore $j \cdot t(v) \geq r_R(v, p)$.
 1857 Furthermore we have that $j \cdot t(v)$ is not in phase $p + 1$. Hence also $j \cdot t(v) < r_R(v, p + 1)$,
 1858 i.e. $j \cdot t(v)$ is in phase p of Π . Due to the similarity $\langle \mathcal{N}_{\mathcal{L}}, \ell \rangle \sim \langle \mathcal{N}_R, t \rangle$, since v fires in
 1859 $L_j(v)$ it also fires in $j \cdot t(v)$. Hence v fires in phase p of Π_R .
- 1860 ■ Assume that v fires in phase p in Π_R . Assume this happens in round τ_R , then $\tau_R \geq \tau_p^*/T$.
 1861 Thus $j \cdot T \cdot t(v) \geq \tau_p^* \geq r_{\mathcal{L}}(v, i)$. Furthermore, $j \cdot t(v) < \tau_p^*/T$ implies that round $j \cdot T \cdot t(v)$
 1862 was before phase $p + 1$ of $\Pi_{\mathcal{L}}$. Therefore $j \cdot T \cdot t(v)$ is in phase p of $\Pi_{\mathcal{L}}$. By the similarity
 1863 $\langle \mathcal{N}_{\mathcal{L}}, \ell \rangle \sim \langle \mathcal{N}_R, t \rangle$ we have that v fires in round $j \cdot T \cdot t(v)$ in $\Pi_{\mathcal{L}}$. Hence v fires in phase p
 1864 of $\Pi_{\mathcal{L}}$. By the similarity $\mathcal{N} \sim \langle \mathcal{N}_{\mathcal{L}}, \ell \rangle$ we get that v fires in round p of Π .