

MIT Open Access Articles

Short and Wide Network Paths

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Marla, Lavanya, Varshney, Lav R, Shah, Devavrat, Prakash, Nirmal A and Gale, Michael E. 2022. "Short and Wide Network Paths." IEEE Transactions on Network Science and Engineering, 9 (2).

As Published: 10.1109/TNSE.2021.3123311

Publisher: Institute of Electrical and Electronics Engineers (IEEE)

Persistent URL: <https://hdl.handle.net/1721.1/143876>

Version: Original manuscript: author's manuscript prior to formal peer review

Terms of use: Creative Commons Attribution-Noncommercial-Share Alike



Short and Wide Network Paths

Lavanya Marla¹, Lav R. Varshney¹, Devavrat Shah², Nirmal A. Prakash¹, and
Michael E. Gale¹

Abstract

Network flow is a powerful mathematical framework to systematically explore the relationship between structure and function in biological, social, and technological networks. We introduce a new pipelining model of flow through networks where commodities must be transported over single paths rather than split over several paths and recombined. We show this notion of pipelined network flow is optimized using network paths that are both short and wide, and develop efficient algorithms to compute such paths for given pairs of nodes and for all-pairs. Short and wide paths are characterized for many real-world networks.

To further demonstrate the utility of this network characterization, we develop novel information-theoretic lower bounds on computation speed in nervous systems due to limitations from anatomical connectivity and physical noise. For the nematode *Caenorhabditis elegans*, we find these bounds are predictive of biological timescales of behavior. Further, we find the particular *C. elegans* connectome is globally less efficient for information flow than random networks, but the hub-and-spoke architecture of functional subcircuits is optimal under constraint on number of synapses. This suggests functional subcircuits are a primary organizational principle of this small invertebrate nervous system.

I. INTRODUCTION

In studying complex systems via the interconnection of their elements, network science has emerged in the last two decades as an insightful approach for understanding collective behavior in brains, societies, and physical infrastructures. Common network science analysis techniques draw on dynamical systems theory [1], [2] and many universal properties of disparate networks have been found [3]. Other prevalent analysis techniques are based on network flow [4].

¹ University of Illinois at Urbana-Champaign

² Massachusetts Institute of Technology

We take the flow perspective and introduce a novel notion of network flow that arises in many biological, social, and technological networks, yet has not previously been studied. Consider a network where a commodity to be transmitted from a source node to a destination node can be split into pieces in time and sent in a pipelined fashion over (possibly) several hops using as many time slots as needed. The commodity, however, must go over a single route rather than being split over several routes to be recombined by the destination [5], [6]. This is in contrast to the maximum capacity problem [7], [8] where flow between two nodes may use as many different routes as needed. This “circuit-switched” model with a single route is prevalent in systems without the ability to split and recombine, e.g. signal flow in simple neuronal networks, message flow in social networks, and the flow of train cars in railroad networks with few engines.

We will see that the optimal paths for the pipelined network flow problem must not only be short in terms of number of hops but also wide in terms of the bottleneck edge in the route. That is, to maximize flow requires finding the single best route between the two nodes: the route that minimizes the weight of the maximum-weight edge in the route and yet is short in path length. Finding all-pairs shortest paths in weighted networks can be accomplished in polynomial time using Floyd’s dynamic programming algorithm [9]. This is an optimization problem in a metric space. Finding widest paths in weighted networks can be accomplished by taking paths in a maximum spanning tree [5], [6], also in polynomial time. This is an optimization problem in an ultrametric (non-Archimedean) space [10]. As far as we can tell, our problem of finding paths that minimize the width-length product between two nodes has remained unstudied in the literature. We develop efficient algorithms for finding short and wide paths between two given nodes, as well as for all pairs. As part of the development, we prove correctness and also characterize the computational complexity as polynomial time. Depth-first search strategies that enumerate all simple paths between two nodes [11], would check a factorial number of paths in the worst case.

Efficient algorithms enable us to characterize the all-pairs distribution of short and wide paths for many complex networks. Note that traditional notions of network diameter and average path length are studied extensively in network science [2], but the all-pairs geodesic distance distribution of unweighted graphs of fairly arbitrary topology is also starting to be of interest [12], [13], essentially building on results for Erdős-Rényi random graphs [14], [15]. As far as we know, this distance distribution for weighted graphs remains unstudied, as does the distribution

of our short-and-wide path lengths. In studying the short-and-wide path length distribution, we do not see universality across networks.

To demonstrate the detailed structure-function insights that pipelined flow gives, we consider neuronal networks. Indeed with advances in experimental connectomics producing wiring diagrams of many neuronal networks, there is growing interest in informational systems theories to provide insight [16]–[18]. For concreteness, we focus on hermaphrodite nematode *Caenorhabditis elegans*, which is a standard model organism in biology [19]–[21] and has exactly 302 neurons [22]. We consider three scientific questions in asking whether information transmission through the nervous system is a bottleneck that limits behavior. (Neural efficiency hypotheses of intelligence also argue information flows better in the nervous systems of bright individuals [23], [24].)

Question I.1. *Do neuronal circuits allow behaviors to happen as quickly as possible under information flow limitations imposed by synaptic noise properties and neuronal connectivity patterns?*

Question I.2. *Are information flow properties of neuronal networks significantly different from random graphs drawn from ensembles that match other network functionals? That is, are networks non-random [25] in allowing information flow that is faster or slower than other networks?*

Question I.3. *Does the synaptic microarchitecture of functional subcircuits optimize information flow under constraint on number of synapses?*

Since the exact computations performed by the nervous system are unclear, we use general information-theoretic methods to lower bound optimal computational performance of a given neural circuit in terms of its physical noise and connectivity structure. This approach is inspired by information-theoretic limits in distributed computing and control [26], [27]. If the performance of a neural circuit is close to the lower bound, then it is operating close to optimally.

We specifically consider gap junctions in *C. elegans*, where neurons are directly electrically connected to each other through pores in their membranes. There can be more than one gap junction connecting two given neurons. We, for the first time, model and compute the Shannon channel capacity of gap junctions. Channel capacity used together with the network topology of the system in the short-and-wide path computation, and with an estimate of the informational

requirements to perform computations, yields a bound on the minimum time to perform biologically plausible computations. Remarkably, when this lower bound is applied to *C. elegans*, the result is rather close to behaviorally-observed timescales. This suggests nematodes may be operating close to the behavioral limits imposed by physical properties of their nervous system.

In asking whether the network is non-random [25] in allowing behavior that is faster or slower than other networks, we surprisingly find that the complete *C. elegans* connectome has greater distances and is therefore slower in supporting global information flow than random networks. Contrarily, we prove that the hub-and-spoke architecture of *C. elegans* functional subcircuits [28], [29] optimizes computation speed under constraint on number of synapses. As such, global information flow may not be a relevant criterion for neurobiology, at least for a small invertebrate system like *C. elegans*. Rather, functional subcircuits may be a primary organizational principle.

II. PIPELINING MODEL OF INFORMATION FLOW

Consider a network where a commodity is to be sent from a source to a destination in a manner that can be split into pieces in time and sent in a pipelined fashion over (possibly) several hops using as many time slots as needed. The commodity must go over a single route rather than being split over several routes to be recombined by the destination. As an example, in *C. elegans*, each neuron is identified by name and is different from any other neuron [22]; computational specialization may arise from neuronal specialization, which in turn may require specific paths for specific information.

In this model, maximizing flow requires finding the single best route between the two nodes: the route that minimizes the weight of the maximum-weight edge in the route and yet is short in path length. In the context of network behavior, note that since we adopt the short-and-wide path view of information flow rather than the maximum capacity view [7], [8], bounds on computation speed will be governed by an appropriate notion of graph diameter rather than by notions of graph conductance [27]. Since diameter provides weaker bounds than graph conductance, this is without loss of generality. Our notion of graph diameter is defined in the next section.

A. Distance and Effective Diameter

Consider the following standard definitions of graph distance for undirected, weighted graphs.

Definition II.1. Let $G = (V, E)$ be a weighted graph. Then the geodesic distance between nodes $s, t \in V$ is denoted $d_G(s, t)$ and is the number of edges connecting s and t in the path with the smallest number of hops between them. If there is no path connecting the two nodes, then $d_G(s, t) = \infty$.

Definition II.2. Let $G = (V, E)$ be a weighted graph. Then the weighted distance between nodes $s, t \in V$ is denoted $d_W(s, t)$ and is the total weight of edges connecting s and t , in the path with the smallest total weight between them. If there is no path connecting the two nodes, then $d_W(s, t) = \infty$.

Another notion of distance arises from the pipelining model of flow. We want a path between two nodes that has a small number of hops but is also such that the weight of the maximum-weight edge is small; we measure path length weighted by this bottleneck weight.

Definition II.3. Let $G = (V, E)$ be a weighted graph. Then the bottleneck distance between nodes $s, t \in V$ is denoted $d_B(s, t)$ and is the number of edges connecting s and t , scaled by the weight of the maximum-weight edge, in the path with the smallest total scaled weight between them. If there is no path connecting the two nodes, then $d_B(s, t) = \infty$.

Proposition II.1. If weights of all actual edges are 1 or less, geodesic distance upper bounds the bottleneck distance:

$$d_B(s, t) \leq d_G(s, t). \quad (1)$$

Proof: Consider the path between s and t that governs $d_G(s, t)$. Since the weights of all actual edges are 1 or less, the maximum-weight edge weight is 1 or less. Hence the bottleneck weight of this path must be less than or equal to the number of edges connecting s and t in that path. Since by definition $d_B(s, t)$ minimizes the bottleneck weight among paths connecting s and t , $d_B(s, t) \leq d_G(s, t)$. ■

Proposition II.2. Weighted distance lower bounds the bottleneck distance:

$$d_B(s, t) \geq d_W(s, t). \quad (2)$$

Proof: Consider the path between s and t that governs $d_B(s, t)$. Let w_0 be the maximum-

weight edge weight and h the number of edges in it. Since other weights in the path are only less than w_0 , the total weight of this path is less than hw_0 , but by definition $d_W(s, t)$ is less than or equal to this quantity. Hence $d_B(s, t) \geq d_W(s, t)$. ■

It is convenient for the sequel to write these distances as constrained optimization problems. We first define a set of constraints on the decision variables x_{ij} that indicate how different edges are used in optimal paths, and N indicates neighborhood.

$$\sum_{j:(s,j) \in E} x_{ij} = 1 \quad (3)$$

$$\sum_{j:(i,j) \in E} x_{ij} - \sum_{j:(j,i) \in E} x_{ji} = 0 \quad \text{for all } i \in N \setminus \{s, t\} \quad (4)$$

$$\sum_{i:(i,t) \in E} x_{it} = 1 \quad (5)$$

$$x_{ij} \in \{0, 1\} \quad \forall (i, j) \in E \quad (6)$$

These constraints simply enforce the movement of one unit of flow from node s to node t and maintain flow balance at all other nodes. Integrality constraints ensure that exactly one path is chosen. These constraints are the same as for standard shortest path problems [4] and therefore the constraint matrix is totally unimodular. The distance expressions use the notation w_{ij} for the edge weight between nodes i and j .

$$d_G(s, t) = \min \sum_{(i,j) \in E} x_{ij} \quad \text{such that (3)–(6) hold.} \quad (7)$$

$$d_W(s, t) = \min \sum_{(i,j) \in E} w_{ij} x_{ij} \quad \text{such that (3)–(6) hold.} \quad (8)$$

$$d_B(s, t) = \min \left\{ \max_{(i,j) \in E} \{w_{ij} x_{ij}\} \sum_{(i,j) \in E} x_{ij} \right\} \quad \text{such that (3)–(6) hold.} \quad (9)$$

Notice the objective functions are nonlinear; yet we develop efficient algorithms in Section III.

Any of these distance functions can be used to define the all-pairs distance distribution of a network, which is just the empirical distribution of distances among all $\binom{n}{2}$ pairs of vertices, for a graph of size $|V| = n$. These distance distributions have various moments and order statistics,

such as the average path length and the diameter.

Definition II.4. *The graph diameter is*

$$D = \max_{s,t \in V} d(s, t). \quad (10)$$

We also define a notion of effective diameter where node pairs that are outliers in the all-pairs distance distribution do not enter into the calculation. Recall that the quantile function corresponding to the cumulative distribution function (cdf) $F(\cdot)$ is $Q(p) = \inf\{x \in \mathbb{R} | p \leq F(x)\}$ for a probability value $0 < p < 1$.

Definition II.5. *For a network of size n , let $F(x)$ be the empirical cdf of the distances of all $\binom{n}{2}$ distinct node pairs. Then the effective diameter is:*

$$D_e = Q(0.95). \quad (11)$$

This definition is more stringent than others in the literature [30]. Of course, $D_e \leq D$. Moving forward, we use effective diameter rather than diameter since it characterizes when most of a commodity would have reached its destination. Thresholds other than 0.95 can be easily defined.

B. Unique Property of Bottleneck Paths

In this subsection, we discuss a property of short-and-wide paths that is distinct from geodesic and weighted paths and that has algorithmic importance.

A key attribute of shortest paths exploited by many shortest path algorithms is the *optimal substructure property*, that all subpaths of shortest paths are shortest paths [4]. This follows from metric structure, but for the ultrametric space induced by short-and-wide paths, the property does not hold. We prove this through a counterexample.

Consider the network in Figure 1, with the network weights w_{ij} shown on the edges. For source node A and destination node H , the geodesic, weighted, and short-and-wide paths are, respectively: $A - E - G - H$ (3 units), $A - B - D - F - H$ (1.367 units), and $A - B - D - F - H$ (or $A - B - C - F - H$, 2 units). Instead, if we compute the same paths with A as the source and K as the destination, path $A - E - G - H - I - J - K$ is the geodesic path (6 units) $A - B - D - F - H - I - J - K$ is the weighted path (3.367 units) and path $A - E - G - H - I - J - K$




Fig. 1: Network that demonstrates the optimal substructure violation for short and wide paths.

is the short-and-wide path. Notice if we compute the short-and-wide path from A to K , this does not guarantee that its sub-path up to node H is optimal for source A and destination H .

This implies short-and-wide paths do not form trees, unlike geodesic or weighted paths. Hence, classic tree-based shortest path algorithms [4] must be modified significantly for this setting, as we now show.

III. EFFICIENT ALGORITHMS FOR COMPUTING SHORT AND WIDE PATHS

In this section, we present algorithms for computing short-and-wide paths.

A. One-to-All Bottleneck Distance Algorithm

We present Algorithm 1 to compute the bottleneck distance. The algorithm maintains label sets at each node $j \in V$, denoted by $L_p(j) = [p, l_p^1(j), l_p^2(j), l_p^3(j), pred_p(j), pp_p(j)]$, where the label sets include: (i) p , the index of the label set, (ii) $l_p^k(j)$, the value of label k in label set p at node j , $k = 1, 2, 3$; where: (a) $l_p^1(j)$ tracks the number of edges traversed from s until the current node, (b) $l_p^2(j)$ tracks the maximum width along a path from the origin until the current node, (c) $l_p^3(j) (= l_p^1(j) \times l_p^2(j))$ tracks the product of the maximum width and the number of edges traversed from s until the current node j ; (iii) $pred_p(j)$: predecessor node for label set p at node j , (iv) $pp_p(j)$: index of predecessor's label set for label p at node j . We also let $np(j)$ denote the number of non-dominated label sets for node j .

Definition III.1. A label set $L_p(j)$ is strictly dominated by label set $L_q(j)$ at node j , if $l_p^1(j) > l_q^1(j)$ and $l_p^2(j) > l_q^2(j)$ (and consequently, $l_p^3(j) > l_q^3(j)$). Label set p is dominated by label set q at node j , if either: (a) $l_p^1(j) > l_q^1(j)$ and $l_p^2(j) = l_q^2(j)$ or (b) $l_p^1(j) = l_q^1(j)$ and $l_p^2(j) > l_q^2(j)$

(and again consequently, $l_p^3(j) > l_q^3(j)$). A label set that is not strictly dominated or dominated, is non-dominated.

Remark III.1. *Observe that each node j can have at most W non-dominated labels, where W is the maximum number of discrete values of weights w_{ij} , over all edges (i, j) in the network. This is a consequence of the fact $l_p^1(j)$ and $l_p^2(j)$ are the only two quantities being tracked (and $l_p^3(j)$ is the product of $l_p^1(j)$ and $l_p^2(j)$); and the possibilities for non-dominated labels are if $l_p^1(j) < l_q^1(j)$ and $l_p^2(j) > l_q^2(j)$ (or vice versa) in labels p and q of node j . Also note that all non-dominated labels must be maintained since the optimal substructure property does not hold, and each such label must be propagated to the downstream nodes because it may dominate after propagation. Since l_p^1 is always an integer, the number of discrete weights upper bounds the possible number of non-dominated labels that can exist at each node.*

Lemma III.1. *Along a given $a - b$ path, the value of the labels on the path monotonically (but not necessarily strictly monotonically) increases.*

Proof: Note that the values in labels l_p^1 and l_p^2 are always non-negative because the labels measure the number of edges and the max width thus far, respectively. The UpdateLabels operation defined below can only increase label values, and thus, on a given $a - b$ path (including if a is the same as b), the label values only increase. ■

Lemma III.2. *The $s - t$ path corresponding to bottleneck distance will not contain any cycles.*

Proof: Proof is by contradiction. Suppose the $s - t$ path with the bottleneck distance contains a cycle, that is the path is $s - \dots - u - u_1 \dots u_k - u - \dots - t$, with cycle $u - u_1 \dots u_k - u$. However, by Lemma III.1, we know that as we travel along $u - u_1 \dots u_k - u$, the value of $d_B(s, t)$ only increases. Therefore the path $s - \dots - u - \dots - t$ that excludes the cycle will have a smaller value of $d_B(s, t)$, contradicting that $s - \dots - u - u_1 \dots u_k - u - \dots - t$ is the shortest path. ■

Theorem III.1. *When Algorithm 1 terminates, all nodes i have the bottleneck distance from s to i as labels.*

Proof: The correctness of the algorithm follows because the label sets keep track of $l_p^1(i), l_p^2(i)$,

and $l_p^3(i)$ for each node i and each label p corresponding to that node i . At each iteration, each non-dominated label is explored, in increasing order of $l_p^3(i)$. That is, each label set from each edge out of the selected node i is propagated downstream during the UpdateLabels step. Due to the ConsolidateLabels step at each node in each iteration, a maximum of W labels can be present at each node (Remark III.1). All non-dominated labels are explored before termination, and due to Lemma III.1 and III.2, all possible paths are explored, resulting in the short-and-wide path. ■

Having established correctness, we also analyze the computational complexity, in terms of the number of nodes n , number of edges m , and the number of possible discrete edge weights W .

Lemma III.3. *The running time of Algorithm 1 is $O(W^2m \log n)$, which is pseudopolynomial.*

Proof: The algorithm is structured like Dijkstra's algorithm which has complexity $O(m \log n)$ [4], but during each iteration, since the UpdateLabels step propagates from each label out of each edge and at each node, a consolidation step must be performed. With a discrete number of weights W over all edges in the network, the possible labels at each node is upper bounded by W . Hence, the order of the algorithm is $O(W^2m \log n)$. ■

B. All-Pairs Bottleneck Distance Algorithm

Now we consider the computation for all-pairs, rather than separately computing for each source-destination pair. This is Algorithm 2.

Theorem III.2. *When Algorithm 2 terminates, all edges (i, j) have bottleneck distance from i to j as labels.*

Proof: Note that the short-and-wide path between any nodes i and j will not contain any cycles, and therefore have at most $n - 1$ edges. In the k th iteration, we consider adding node k to each edge along the path connecting i and j , from each label set at i to each label set at j . This algorithm is equivalent to enumerating all possible combinations of the label sets $[i, k]$ and $[k, j]$. If the short-and-wide path between i and j contains node k , then the condition $d[i, j] > \max\{l_2[i, k], l_2[k, j]\} \times (l_1[i, k] + l_1[k, j])$ will be violated in iteration k , and k is added to one of the edges on the path. If there is an improving or non-dominated label that can be added to a node, it is added even if k is not on the bottleneck path between i and j . If the

Algorithm 1 One-to-All Bottleneck Distance

```

1: Initialize:
    $S = \emptyset, S' = \emptyset$ 
    $l_1^k(s) = 0, \forall k = 1, 2, 3$ 
    $S = S \cup \{s\}, S' = S' \cup \{l_p^k(s)\} \forall k = 1, 2, 3, p = 1, 2$ 
    $l_p^k(j) = \infty; \forall k = 1, 2, 3, p = 1, 2; np(j) = 0, p_p(j) = -1$  and
    $pp(j) = -1 \forall j \in N \setminus \{s\}$ 
2:  $V' = V' \cup \{l_p^k(j)\} \forall j \in V, \forall k = 1, 2, 3$ 
3: while  $S \neq V$  and  $S' \neq V'$  do
4:   for each node  $j \in V \setminus S$  do
5:     for each non-dominated label set  $L_p(j)$  (see Definition III.1) do
6:        $T := T \cup \{L_p(j)\}$ 
7:     end for
8:   end for
9:   Order label sets in  $T$  in increasing order of  $l_p^3(j), \forall p, \forall j$ , breaking ties arbitrarily. Find
   label set  $L_p(i) \in T$  such that  $l_p^3(i) = \min_{L \in V' \setminus S'} l_q^3(j)$ 
10:   $T = T \setminus \{L_p(i)\}, S' = S' \cup \{L_p(i)\}, S := S \cup \{i\}$ 
11:  UpdateLabels( $L_p(i)$ ):
12:  for each edge  $(i, j) \in E$  do
13:    for each label set  $q = 1, \dots, np(j)$  at node  $j$  do
14:      if  $l_q^1(j) > l_p^1(i) + 1 \wedge l_q^2(j) > \max\{l_p^2(i), w_{ij}\}$  then
15:         $pred_q(j) = i, pp_q(j) = p$ 
16:         $l_q^1(j) = l_p^1(i) + 1$ 
17:         $l_q^2(j) = \max\{l_p^2(i), w_{ij}\}$ 
18:         $l_q^3(j) = l_q^1(j) * l_q^2(j)$ 
19:      else if  $(l_q^1(j) \leq l_p^1(i) + 1 \text{ and } l_q^2(j) > \max\{l_p^2(i), w_{ij}\}) \vee (l_q^1(j) > l_p^1(i) + 1 \text{ and } l_q^2(j) \leq \max\{l_p^2(i), w_{ij}\})$  then
20:        Create new temporary label  $L_{q'}(j)$  at  $j$  with  $pred_{q'}(j) = i, pp_{q'}(j) = p, l_{q'}^1 = l_p^1(i) + 1, l_{q'}^2(j) = \max\{l_p^2(i), w_{ij}\}, l_{q'}^3(j) = l_{q'}^1(j) * l_{q'}^2(j)$ 
21:      end if
22:    end for
23:    ConsolidateLabels( $j$ ):
24:    for all (including temporary) labels  $p$  at  $i$  do
25:      Delete all dominated labels (Definition III.1). Also combine labels  $p$  and  $q$  with
       $l_q^1(j) = l_p^1(j)$  and  $l_q^2(j) = l_p^2(j)$ . Temporary label made permanent if non-dominated. Update
       $np(j)$ .
26:    end for
27:  end for
28: end while

```

Algorithm 2 All-to-All Bottleneck Distances

```

1: Initialize:
    $l^1[i, j] = \emptyset \forall i, j \in N$ 
    $l^2[i, j] = \emptyset \forall i, j \in N$ 
   Let  $l_1^1[i, j] = 1$  if nodes  $i$  and  $j$  are adjacent
   Let  $l_1^2[i, j]$  be the weight between nodes  $i$  and  $j$  if they are connected
    $np[i, j] = l_1^1[i, j] \forall i, j \in N$ 
2: for all nodes  $k \in N$  do
3:   for all nodes  $i \in N$  do
4:     for all nodes  $j \in N$  do
5:       Node Insertion on label sets  $L[i, k]$  and  $L[k, j]$ 
6:       Maximize Labels on label sets  $L[i, j]$  and  $L'[i, j]$ 
7:     end for
8:   end for
9: end for

```

Algorithm 3 Node Insertion

```

1: Initialize:
    $p = np[i, k]; q = np[k, j]$ 
2: while  $L_p[i, k]$  and  $L_q[k, j]$  exist do
3:   append to front  $l_p^1[i, k] + l_q^1[k, j]$  to  $l'^1[i, j]$ 
4:   append to front  $\max\{l_p^2[i, k], l_q^2[k, j]\}$  to  $l'^2[i, j]$ 
5:   if  $l_p^2[i, k] = l_q^2[k, i]$  then  $p = p - 1; q = q - 1$ 
6:   end if
7:   if  $l_p^2[i, k] > l_q^2[k, i]$  then  $p = p - 1$ 
8:   end if
9:   if  $l_p^2[i, k] < l_q^2[k, i]$  then  $q = q - 1$ 
10:  end if
11: end while

```

bottleneck path does not contain node k , then the condition will not be violated and the current best path and distance are retained, leading to correctness. ■

After establishing correctness, we characterize computational complexity in terms of the number of nodes n , number of edges m , and the number of discrete edge weights W .

Theorem III.3. *The worst-case runtime for Algorithm 2 is pseudopolynomial $O(n^3W^2)$.*

Proof: Note that the structure of the algorithm is similar to the Floyd-Warshall algorithm for shortest paths, which is $O(n^3)$ [4]. Rather than one comparison inside of this, the bottleneck distance algorithm has two sub-algorithms, Algorithm 3 and Algorithm 4, to compare all of the

Algorithm 4 Maximize Labels

```

1: Initialize:
    $bestL^1 = \infty$ 
    $p, q = 1$ 
    $np[i, j] = 0$ 
   Note  $L[i, j]$  and  $L'[i, j]$  are sorted ascending by  $l^2$ 
2: while  $L_p[i, k]$  and  $L'_q[i, j]$  exist do
3:   if  $l_p^2[i, j] < l_q^2[i, j]$  then
4:     if  $l_p^1[i, j] < bestL^1$  then
5:       append  $L_p[i, j]$  to  $L''[i, j]$ 
6:        $bestL^1 = l_p^1[i, j]$ 
7:     end if
8:      $p = p + 1$ 
9:   else if  $l_p^2[i, j] > l_q^2[i, j]$  then
10:    if  $l_q^1[i, j] < bestL^1$  then
11:      append  $L'_q[i, j]$  to  $L''[i, j]$ 
12:       $bestL^1 = l_q^1[i, j]$ 
13:    end if
14:     $q = q + 1$ 
15:  else if  $l_p^2[i, j] = l_q^2[i, j]$  then
16:    if  $l_p^1[i, j] \leq l_q^1[i, j] \wedge l_p^1[i, j] < bestL^1$  then
17:      append  $L_p[i, j]$  to  $L''[i, j]$ 
18:       $bestL^1 = l_p^1[i, j]$ 
19:    else if  $l_p^1[i, j] > l_q^1[i, j] \wedge l_q^1[i, j] < bestL^1$  then
20:      append  $L'_q[i, j]$  to  $L''[i, j]$ 
21:       $bestL^1 = l_q^1[i, j]$ 
22:    end if
23:     $p = p + 1$ 
24:     $q = q + 1$ 
25:  end if
26: end while
27: for each remaining node in  $L[i, j]$  or  $L'[i, j]$  do
28:   if  $l^1[i, j] < bestL^1$  then append to  $L''[i, j]$ 
29:   end if
30: end for
31:  $L[i, j] = L''[i, j]$ 
32: Delete  $L''[i, j]$  and  $L'[i, j]$ 
33: update  $np[i, j]$ 

```

label sets. Each sub-algorithm at worst iterates through two label sets in parallel with size W if there are W levels of weights based on Remark III.1 (or maximum size m). Therefore their runtimes are $O(W^2)$ (or $O(m^2)$ if all edges have different weights). Combining this with the whole algorithm gives the total runtime as $O(n^3W^2)$ (or worst-case of $O(n^3m^2)$). ■

IV. ALL-PAIRS BOTTLENECK DISTANCE DISTRIBUTION IN COMPLEX NETWORKS


With efficient algorithms in place, we can compute the all-pairs bottleneck distance distribution for several undirected, weighted, real-world networks drawn from the Index of Complex Networks [31] database, which is publicly available. We choose networks that span a variety of systems, including transportation networks, biological networks, and social networks that may naturally support pipelined flow. Table I details the size, type, and sources of each network.

TABLE I: Characteristics of Real-World Networks


Name	Nodes	Edges	Type	Source
US airports	1574	28236	Transportation	US airport networks (2010) [31]
Mumbai bus routes	2266	3042	Transportation	India bus routes (2016) [31]
Chennai bus routes	1009	1610	Transportation	India bus routes (2016) [31]
Author collaborations	475	625	Social	Social Networks authors (2008) [31]
Free-ranging dogs	108	1296	Social	Wilson-Aggarwal dogs [31]
Game of Thrones	107	353	Social	Game of Thrones coappearances [31]
Resting state fMRI network	638	18625	Biological	Human brain functional coactivations [31]
Human brain coactivation	638	18625	Biological	Human brain functional coactivations [31]

In Figure 2 we plot the survival functions for the all-to-all geodesic, weighted, and bottleneck distances d_G , d_W , and d_B . For each network, we find that the inequalities $d_W \leq d_B \leq d_G$ hold, as required by Propositions II.1 and II.2. If all edge weights were 1, all distance metrics are equivalent, i.e., $d_W = d_B = d_G$. For example, note that in the Authors' collaboration network (Figure 2d), the values of d_B and d_G very nearly coincide because the maximum weight along nearly every path is 1. However, because a significant number of weights are far from 1, the values of d_W and d_B do not coincide. Additionally, we observe that the geodesic, bottleneck, and weighted distances diverge as the weights are distributed away from 1.


Previous investigations of the geodesic distance distribution of random graphs had suggested good fits by Weibull, gamma, lognormal, and generalized three-parameter gamma distributions [12] as well as basic generative models to explain these distributions. We consider the same parametric families to understand if short-and-wide path length distributions are well-described by such parametric forms. For each network in Table I, we fit the survival functions from Figure 2




(a) US airports 2010




(b) Mumbai bus network




(c) Chennai bus network




(d) Authors' collaboration network




(e) Free ranging dogs social network



(f) Game of Thrones coappearances



(g) Resting State fMRI network



(h) Human brain coactivation network

Fig. 2: Survival function for the empirical all-pairs lengths.

and find goodness-of-fit for several different parametric families using the Kolmogorov-Smirnov test. No single parametric family was best for all networks, but the gamma distribution provided reasonable fits for several networks, as shown in Table II which provides the fitted distribution parameters, along with the χ^2 and p -values from the test. As far as we can tell, there is no universality in the bottleneck distance distribution across networks.

TABLE II: Gamma Distribution Fits of Bottleneck Distance

Name	Shape	Location	Scale	χ^2	p -value
US airports	0.27	0	0.64	5.80	0.63
Chennai bus routes	0.54	0	0.22	1.03	0.06
Mumbai bus routes	0.41	0	0.23	21.25	0.095
Author collaborations	0.83	0	0.32	8.19	0.22
Free-ranging dogs	0.67	0	0.40	3.56	0.61
Game of Thrones	0.55	0	0.18	4.71	0.008
Resting state fMRI network	0.61	0	0.19	4.91	0.23
Human brain coactivation	368.63	-7.47	0.02	8.756	0.85

V. C. *elegans* NEURONAL NETWORK: GLOBAL AND LOCAL FLOW

We turn attention specifically to the *C. elegans* gap junction network, to investigate how information flow limits behavior. A bound on the Shannon capacity [bits/sec] of a single gap junction is developed in the Appendix. Although there is no reason to believe capacity-achieving codes are used in neural signaling, Shannon capacity provides bounds on the information rate for any signaling scheme.

The topology of the gap junction network has been characterized in some detail in our prior work [22]. The somatic network consists of a giant component comprising 248 neurons, two small connected components, and several isolated neurons. Within the giant component, the average geodesic distance between two neurons is 4.52. Since this characteristic path length is similar to that of a random graph and since the clustering coefficient is large with respect to a random graph, the network is said to be a small-world network. Moreover, the *C. elegans* network overall has good expander properties [32, App. C].

Here we compute the effective diameter of the giant component of the gap junction network, with respect to the bottleneck distance. We also find upper and lower bounds. Figure 3 shows the survival function of the empirical all-pairs geodesic, weighted, and bottleneck distances. As can be observed, the effective diameter for bottleneck distance is between 6 and 7. The upper and lower bounds are close to one another since, as shown in Figure 4, the minimax width of paths




Fig. 3: Survival function for the empirical all-pairs distance distributions of the *C. elegans* gap junction neuronal network giant component. The weighted and bottleneck distances are listed in terms of inverse gap junctions.




Fig. 4: Survival function for the empirical all-pairs bottleneck width of the *C. elegans* gap junction neuronal network giant component, without considering path length. The width is listed in terms of inverse gap junctions.

in the *C. elegans* network when ignoring path length is almost always one inverse gap junction rather than smaller. Taking path lengths into account only increases the bottleneck width.

A. Limits in Computation Speed

Having characterized the channel capacity of links and the topology of neuronal connectivity, we now develop an information-theoretic model of computation, which in turn yields a limit on computation speed derived from information flow bounds.

Consider the chemosensing problem faced by an organism like *C. elegans*. It has 700 different types of chemoreceptors [33] and must take behavioral actions based on the chemical properties

of its environment [28], [34]. Suppose that differentiating between 700 chemicals requires 10 bits of information, which we call the message volume $\log M$. To perform an action, the neurons must reach consensus among themselves based on the sensory neuron signals [28]. This may, in principle, need transport of information between all parts of the neuronal network. We make the following natural assumption: the consensus time is bounded by the amount of time it takes to transport 10 bits of information across the network (maximum over any two pair of neurons).

The intuition behind this assumption is the following, which follows from the sparsity of sensory response in *C. elegans* [35]. Suppose one part of the neuronal network has “strong” sensing information about an event in the environment, e.g. worm near a chemical, but the other part has little or no such sensing information. Then effectively all information is communicated from one part of the network to the other. Accounting for such instances naturally justifies the above assumption. Note that the actual computational procedures used by *C. elegans* may require several sweeps of signals through the organism, but for bounding purposes, we assume that one sweep is enough to spread the requisite information.

In effect, the time for transportation of information across the network and hence to reach consensus is bounded below by

$$t = \frac{D_e \log M}{C} = \frac{7 \cdot 10}{1700} = 0.041\text{s}.$$

This bound uses the geodesic effective diameter and assumes bandwidth 1700 Hz or equivalently a refractory period of 1/1.7 ms (see Appendix). There is evidence suggesting that the *C. elegans* refractory period is likely to be near 1 ms instead. In that case, the above bound of 41 ms would become 70 ms. The bound of 41 ms or 70 ms applies to the whole giant component.

There may be smaller subcircuits within the neuronal network responsible for specific functional reactions, within the giant component. If consensus is to be reached only in those functional subcircuits, we should utilize their diameter in place of effective diameter 7. As explained in Section V-B, they may have extremal diameters of 2. Then the information propagation time would be bounded by 12 ms (or 20 ms under 1 ms refractory period).

Using circuit-theoretic techniques, the predicted timescale of operation of functional circuits was between 20 ms and 83 ms [22]. This is clearly an excellent match to the range predicted by our bounds: 12 ms to 41 ms (or 20 ms to 70 ms under the 1 ms refractory period). This is

rather surprising given that our technique is only attempting to derive fundamental lower bounds using anatomical information.

Now we compare our lower-bound predictions to experimentally observed behavioral times. This is a true test of our methods in bounding the propagation time for decision-making information. The behavioral switch times in response to chemical gradients as fast as 200 ms have been observed in the literature [36, Supp. Fig. 1]. Since the time required for motor action like turning around must also be taken account—the worm can straighten itself in viscous fluids within 6 to 20 ms [37, Fig. 4B]—the lower-bound results are in agreement.

Collectively, these agreements with the model calculation bounds suggest that information propagation is likely to be a primary bottleneck in the behavioral decision making of *C. elegans*.

B. Hub-and-Spoke Architecture

As mentioned earlier, there are smaller subcircuits of the neuronal network that are responsible for certain functional reactions. For such subcircuits, the hub-and-spoke architecture has optimality properties. The basic premise is that the diameter of such a subcircuit should be small for computational speed; a hub-and-spoke network structure provides the smallest possible diameter of 2 with the constraint on the number of edges as well as connectivity requirement. Formally, we state the following easy fact.

Proposition V.1. *Given a connected graph G of $n \geq 3$ nodes and $n - 1$ edges, the smallest possible diameter is 2 and is achieved by the hub-and-spoke structure.*

Proof: Clearly for a connected network with $n \geq 3$ nodes and $n - 1$ edges, it is a tree (no cycles). Further if there are $n \geq 3$ nodes and $n - 1$ edges, there must be a pair of nodes not connected to each other through an edge. Since the graph is connected, they must be at least 2 hops apart. That is, the diameter of such a graph must be at least 2. A hub-and-spoke network, by construction has diameter 2. ■

Indeed, as shown in Figure 5, certain known functional subcircuits in the *C. elegans* neuronal network do indeed follow the hub-and-spoke architecture (or nearly so) [28], [29]. Note that other arguments also suggest the benefits to neuronal networks of small diameter [38], but not the optimality of hub-and-spoke architectures.




Fig. 5: Gap junction connectivity of a known chemosensory subcircuit [28]. The worm is almost left-right symmetric; the circuits on both sides are shown. The circuit on the left side of the worm (a) is a hub-and-spoke circuit whereas the circuit on the right side (b) is nearly so.

C. Comparison to Random Networks

Now we wish to study whether the bottleneck diameter of the *C. elegans* gap junction network is more than, less than, or similar to the bottleneck diameter of random graphs that have certain other network functionals fixed.

To evaluate the nonrandomness of the bottleneck diameter of the *C. elegans* network giant component, we compare it with the same quantity expected in random networks. We start with a weighted version of the Erdős-Rényi random network ensemble because it is a basic ensemble. Constructing the topology requires a single parameter, the probability of a connection between two neurons. There are 514 gap junction connections over 279 somatic neurons in *C. elegans*, and so we choose the probability of connection as $0.0133 = 2 \times 514/279/278$. After fixing the topology, we choose the multiplicity of the connections by sampling randomly according to the *C. elegans* multiplicity distribution [22, Fig. 3(B)], which is well-modeled as a power-law with parameter 2.76. Note that in general the giant component for such a construction will be much larger than that of *C. elegans*.

Figure 6 shows the survival function of the empirical all-pairs geodesic, weighted, and bottleneck distances of one hundred random networks. A random example is highlighted. As can be observed, the effective diameter bottleneck distance is roughly 6, significantly less than that for the *C. elegans* network.

Now we consider a degree-matched weighted ensemble of random networks. In such a random network, the degree distribution matches the degree distribution of the gap junction network; the




Fig. 6: Survival function for the empirical all-pairs distance distributions of 100 Erdős-Rényi random network giant components; a random example is highlighted in colored lines. The weighted and bottleneck distances are in terms of inverse gap junctions.

degree of a neuron is the number of neurons with which it makes a gap junction. Such a random ensemble is created using a numerical rewiring procedure to generate samples [39], [40]. Upon fixing the topology, the multiplicity of connections is sampled as for the Erdős-Rényi ensemble. Note that in general the giant component for such a construction will be much larger than that of *C. elegans*.

Figure 7 shows the survival function of the empirical all-pairs geodesic, weighted, and bottleneck distances of one hundred random networks. A random example is highlighted. As can be observed, the effective diameter for bottleneck distance is just below 5, quite significantly less than that for the *C. elegans* network. Comparing Figures 6 and 7 to Figure 3, note that the results hold for many defining thresholds for effective diameter, not just 0.95.

These results reveal a key nonrandom feature in synaptic connectivity of the *C. elegans* gap junction network, but perhaps contrary to expectation. The network has a non-randomly *worse* bottleneck diameter compared to basic random graph ensembles. It enables globally *slower* behavioral speed than similar random networks. In contrast, Section V-B had found that at the micro-level of small functional sub-circuits, the *C. elegans* gap junction network has several hub-and-spoke structures [28], [29], which are actually optimal from an information flow perspective. Thus, these results lend greater nuance to efficient flow hypotheses in neuroscience.




Fig. 7: Survival function for the empirical all-pairs distance distributions of 100 degree-matched random network giant components; a random example is highlighted in colored lines. The weighted and bottleneck distances are in terms of inverse gap junctions.

VI. CONCLUSION

We have modeled pipelined network flow, as arises in several biological, social, and technological networks and developed algorithms to efficiently find short-and-wide paths that optimize this novel network flow. This new model of flow also specifically provides a new approach to understand the limits of information propagation and behavioral speed supported by neuronal networks. Within the field of network information theory itself, it is of interest to study the novel notion of bottleneck distance in detail from a theoretical perspective. Since we did not universal scaling laws across several real-world networks, it is also of interest to analytically characterize its distribution in random ensembles such as Watts-Strogatz small worlds, Barabási-Albert scale-free networks, Kronecker random graphs, or random geometric graphs.

Beyond our general study, we also specifically considered circuit neuroscience and connectomics, where the overarching goal is try to understand how an animal’s behavior arises from the operations of its nervous system. The nervous system must transport information from one part to another, whether engaged in communication, computation, control, or maintenance. This paper proposes a way to characterize information flow through the nervous system from detailed properties of anatomical connectivity data and to use this characterization to make lower bound statements on the behavioral timescales of animals.

The efficacy of these techniques was demonstrated by explaining the communication bottle-

necks in the gap junction network of the nematode *C. elegans*. Remarkably, the timescale lower bounds are predictive of behavioral timescales.

In considering the possibility of changing the network topology itself, we discovered that the network has much worse bottleneck distance than similar random graphs (whether Erdős-Rényi or degree-matched). The network does not seem to be optimized for global information flow. On the other hand, we noted the prominence of hub-and-spoke functional subcircuits in the *C. elegans* gap junction network and proved their optimality for information flow under number of gap junctions constraints. In terms of neural organization, this suggests that smaller subcircuits within the larger neuronal network are responsible for specific functions, and these should have fast information flow (to quickly achieve the computational objective of that circuit, such as chemotaxis). Behavioral speed of the global network may not be biologically relevant.

As more and more connectomes are uncovered and more details of the biophysical properties of synapses are determined, these information flow techniques may provide a general methodology to understand how physical constraints lead to informational and thereby behavioral limits in nervous systems.

Authors' Contributions: LM developed and implemented efficient algorithms for computing bottleneck distance, proved their correctness and characterized their complexity, participated in the design of the study, and carried out analysis; LRV conceived of the study, designed the study, coordinated the study, proved certain results, participated in data analysis, and drafted the manuscript; DS conceived of the pipelining model of network flow and proved certain theoretical results; NAP and MEG implemented algorithms and participated in data analysis. All authors gave final approval for publication and agree to be held accountable for the work performed therein.

REFERENCES

- [1] D. J. Watts and S. H. Strogatz, "Collective dynamics of 'small-world' networks," *Nature*, vol. 393, no. 6684, pp. 440–442, Jun. 1998.
- [2] M. E. J. Newman, *Networks: An Introduction*. Oxford University Press, 2010.
- [3] A.-L. Barabási and R. Albert, "Emergence of scaling in random networks," *Science*, vol. 286, no. 5439, pp. 509–512, Oct. 1999.
- [4] R. K. Ahuja, T. L. Magnanti, and J. B. Orlin, *Network Flows: Theory, Algorithms, and Applications*. Pearson, 1993.
- [5] M. Pollack, "The maximum capacity through a network," *Oper. Res.*, vol. 8, no. 5, pp. 733–736, Sept.-Oct. 1960.
- [6] T. C. Hu, "The maximum capacity route problem," *Oper. Res.*, vol. 9, no. 6, pp. 898–900, Nov.-Dec. 1961.

- [7] L. R. Ford, Jr. and D. R. Fulkerson, “Maximal flow through a network,” *Can. J. Math.*, vol. 8, pp. 399–404, 1956.
- [8] P. Elias, A. Feinstein, and C. E. Shannon, “A note on the maximum flow through a network,” *IRE Trans. Inf. Theory*, vol. IT-2, no. 4, pp. 117–119, Dec. 1956.
- [9] R. W. Floyd, “Algorithm 97: Shortest path,” *Commun. ACM*, vol. 5, no. 6, p. 345, Jun. 1962.
- [10] R. Rammal, G. Toulouse, and M. A. Virasoro, “Ultrametricity for physicists,” *Rev. Mod. Phys.*, vol. 58, no. 3, pp. 765–788, July-Sept. 1986.
- [11] P. E. Black, “All simple paths,” in *Dictionary of Algorithms and Data Structures*, V. Pieterse and P. E. Black, Eds. National Institute of Standards and Technology, 2008.
- [12] C. Bauckhage, K. Kersting, and F. Hadji, “Parameterizing the distance distribution of undirected networks,” in *Proc. 31st Annu. Conf. Uncertainty in Artificial Intelligence (UAI’15)*, Jul. 2015.
- [13] S. Melnik and J. P. Gleeson, “Simple and accurate analytical calculation of shortest path lengths,” arXiv:1604.05521v1 [physics.soc-ph], Apr. 2016.
- [14] V. D. Blondel, J.-L. Guillaume, J. M. Hendrickx, and R. M. Jungers, “Distance distribution in random graphs and application to network exploration,” *Phys. Rev. E*, vol. 76, no. 6, p. 066101, Dec. 2007.
- [15] E. Katzav, M. Nitzan, D. ben Avraham, P. L. Krapivsky, R. Kühn, N. Ross, and O. Biham, “Analytical results for the distribution of shortest path lengths in random networks,” *Europhys. Lett.*, vol. 111, no. 2, p. 26006, Jul. 2015.
- [16] O. Sporns, G. Tononi, and R. Kötter, “The human connectome: A structural description of the human brain,” *PLoS Comput. Biol.*, vol. 1, no. 4, pp. 0245–0251, Sep. 2005.
- [17] H. S. Seung, “Neuroscience: Towards functional connectomics,” *Nature*, vol. 471, no. 7337, pp. 170–172, Mar. 2011.
- [18] S. Seung, *Connectome: How the Brain’s Wiring Makes Us Who We Are*. Boston: Houghton Mifflin Harcourt, 2012.
- [19] C. Rockland, “The nematode as a model complex system: A program of research,” M.I.T. Laboratory for Information and Decision Systems, Working Paper WP-1865, Apr. 1989.
- [20] M. de Bono and A. V. Maricq, “Neuronal substrates of complex behaviors in *C. elegans*,” *Annu. Rev. Neurosci.*, vol. 28, pp. 451–501, Jul. 2005.
- [21] P. Sengupta and A. D. T. Samuel, “*Caenorhabditis elegans*: A model system for systems neuroscience,” *Curr. Opin. Neurobiol.*, vol. 19, no. 6, pp. 637–643, Dec. 2009.
- [22] L. R. Varshney, B. L. Chen, E. Paniagua, D. H. Hall, and D. B. Chklovskii, “Structural properties of the *Caenorhabditis elegans* neuronal network,” *PLoS Comput. Biol.*, vol. 7, no. 2, p. e1001066, Feb. 2011.
- [23] I. J. Deary, L. Penke, and W. Johnson, “The neuroscience of human intelligence differences,” *Nat. Rev. Neurosci.*, vol. 11, no. 3, pp. 201–211, Mar. 2010.
- [24] T.-W. Lee, Y.-T. Wu, Y. W.-Y. Yu, H.-C. Wu, and T.-J. Chen, “A smarter brain is associated with stronger neural interaction in healthy young females: A resting EEG coherence study,” *Intelligence*, vol. 40, no. 1, pp. 38–48, Jan.-Feb. 2012.
- [25] S. Song, P. J. Sjöström, M. Reigl, S. Nelson, and D. B. Chklovskii, “Highly nonrandom features of synaptic connectivity in local cortical circuits,” *PLoS Biol.*, vol. 3, no. 3, pp. 0507–0519, Mar. 2005.
- [26] N. C. Martins and M. A. Dahleh, “Feedback control in the presence of noisy channels: “Bode-like” fundamental limitations of performance,” *IEEE Trans. Autom. Control*, vol. 53, no. 7, pp. 1604–1615, Aug. 2008.
- [27] O. Ayaso, D. Shah, and M. A. Dahleh, “Information theoretic bounds for distributed computation over networks of point-to-point channels,” *IEEE Trans. Inf. Theory*, vol. 56, no. 12, pp. 6020–6039, Dec. 2010.
- [28] E. Z. Macosko, N. Pokala, E. H. Feinberg, S. H. Chalasani, R. A. Butcher, J. Clardy, and C. I. Bargmann, “A hub-and-spoke

- circuit drives pheromone attraction and social behaviour in *C. elegans*,” *Nature*, vol. 458, no. 7242, pp. 1171–1176, Apr. 2009.
- [29] I. Rabinowitch, M. Chatzigeorgiou, and W. R. Schafer, “A gap junction circuit enhances processing of coincident mechanosensory inputs,” *Curr. Biol.*, vol. 23, no. 11, pp. 963–967, Jun. 2013.
- [30] J. Leskovec, J. Kleinberg, and C. Faloutsos, “Graph evolution: Densification and shrinking diameters,” *ACM Trans. Knowl. Discovery Data*, vol. 1, no. 1, p. 2, Mar. 2007.
- [31] A. Clauset, E. Tucker, and M. Sainz, “The Colorado index of complex networks,” <https://icon.colorado.edu>, 2016.
- [32] L. R. Varshney, B. L. Chen, E. Paniagua, D. H. Hall, and D. B. Chklovskii, “Structural properties of the *Caenorhabditis elegans* neuronal network,” arXiv:0907.2373v4 [q-bio], Jun. 2010.
- [33] A. J. Whittaker and P. W. Sternberg, “Sensory processing by neural circuits in *Caenorhabditis elegans*,” *Curr. Opin. Neurobiol.*, vol. 14, no. 4, pp. 450–456, Aug. 2004.
- [34] S. H. Chalasani, N. Chronis, M. Tsunozaki, J. M. Gray, D. Ramot, M. B. Goodman, and C. I. Bargmann, “Dissecting a circuit for olfactory behaviour in *Caenorhabditis elegans*,” *Nature*, vol. 450, no. 7166, pp. 63–70, Nov. 2007.
- [35] A. Zaslaver, I. Liani, O. Shtangel, S. Ginzburg, L. Yee, and P. W. Sternberg, “Hierarchical sparse coding in the sensory system of *Caenorhabditis elegans*,” *Proc. Natl. Acad. Sci. U.S.A.*, vol. 112, no. 4, pp. 1185–1189, Jan. 2015.
- [36] D. R. Albrecht and C. I. Bargmann, “High-content behavioral analysis of *Caenorhabditis elegans* in precise spatiotemporal chemical environments,” *Nat. Methods*, vol. 8, no. 7, pp. 599–605, Jul. 2011.
- [37] C. Fang-Yen, M. Wyart, J. Xie, R. Kawai, T. Kodger, S. Chen, Q. Wen, and A. D. T. Samuel, “Biomechanical analysis of gait adaptation in the nematode *Caenorhabditis elegans*,” *Proc. Natl. Acad. Sci. U.S.A.*, vol. 107, no. 47, pp. 20 323–20 328, Nov. 2010.
- [38] M. Kaiser and C. C. Hilgetag, “Nonoptimal component placement, but short processing paths, due to long-distance projections in neural systems,” *PLoS Comput. Biol.*, vol. 2, no. 7, p. e95, Jul. 2006.
- [39] S. Maslov and K. Sneppen, “Specificity and stability in topology of protein networks,” *Science*, vol. 296, no. 5569, pp. 910–913, May 2002.
- [40] M. Reigl, U. Alon, and D. B. Chklovskii, “Search for computational modules in the *C. elegans* brain,” *BMC Biol.*, vol. 2, p. 25, Dec. 2004.
- [41] F. Rieke, D. Warland, R. de Ruyter van Steveninck, and W. Bialek, *Spikes: Exploring the Neural Code*. Cambridge, MA: MIT Press, 1997.
- [42] A. Borst and F. E. Theunissen, “Information theory and neural coding,” *Nat. Neurosci.*, vol. 2, no. 11, pp. 947–957, Nov. 1999.
- [43] B. W. Connors and M. A. Long, “Electrical synapses in the mammalian brain,” *Annu. Rev. Neurosci.*, vol. 27, pp. 393–418, Jul. 2004.
- [44] T. C. Ferrée and S. R. Lockery, “Computational rules for chemotaxis in the nematode *C. elegans*,” *J. Comput. Neurosci.*, vol. 6, no. 3, pp. 263–277, May 1999.
- [45] A. Manwani and C. Koch, “Detecting and estimating signals in noisy cable structures, I: Neuronal noise sources,” *Neural Comput.*, vol. 11, no. 8, pp. 1797–1829, Nov. 1999.
- [46] J. B. Johnson, “Thermal agitation of electricity in conductors,” *Phys. Rev.*, vol. 32, no. 1, pp. 97–109, Jul. 1928.
- [47] H. Nyquist, “Thermal agitation of electric charge in conductors,” *Phys. Rev.*, vol. 32, no. 1, pp. 110–113, Jul. 1928.
- [48] S. R. Lockery and M. B. Goodman, “The quest for action potentials in *C. elegans* neurons hits a plateau,” *Nat. Neurosci.*, vol. 12, no. 4, pp. 377–378, Apr. 2009.

- [49] T. B. Achacoso and W. S. Yamamoto, *AY's Neuroanatomy of C. elegans for Computation*. CRC Press, 1992.
- [50] N. Chayat and S. Shamai (Shitz), "Bounds on the information rate of intertransition-time-restricted binary signaling over an AWGN channel," *IEEE Trans. Inf. Theory*, vol. 45, no. 6, pp. 1992–2006, Sep. 1999.
- [51] L. R. Varshney, P. J. Sjöström, and D. B. Chklovskii, "Optimal information storage in noisy synapses under resource constraints," *Neuron*, vol. 52, no. 3, pp. 409–423, Nov. 2006.

APPENDIX: GAP JUNCTION LINKS

Although nearly all information-theoretic investigation of synaptic transmission has focused on electrochemical synapses [41], [42], purely electrical gap junctions are also rather ubiquitous, not only in *C. elegans* but also in the mammalian brain [43]. Among the 282 somatic neurons in the *C. elegans* connectome, there are 890 gap junctions [22]. As such, it is important to provide a mathematical model of signal flow through gap junctions, along with the noise that perturbs signals.

A. Thermal Noise

A gap junction is a hollow protein that allows electrical current to flow between cells, and thereby allows signaling in both directions. In previous modeling efforts, it has been determined that gap junctions can essentially just be modeled as resistors with thermal noise that is additive white Gaussian (AWGN) [44]. Noise due to stochastic chemical effects or due to random background synaptic activity [45] need not be considered.

The electrical conductance of a *C. elegans* gap junction is 200 pS [22]; a resistor with conductance 200 pS corresponds to a resistance of 5000 M Ω . In order to compute the root mean square (RMS) thermal noise voltage v_n , we use the Johnson-Nyquist formula [46], [47]:

$$v_n = \sqrt{4k_B T R \Delta f}, \quad (12)$$

where k_B is Boltzmann's constant 1.38×10^{-23} J/K, T is the absolute temperature in K, R is the resistance in Ω , and Δf is the bandwidth in Hz over which the noise is measured. We assume room temperature (298 K) and for reasons that will become evident in the sequel, we take the bandwidth to be 1700 Hz. Then

$$\begin{aligned} v_n &= \sqrt{4 \cdot 1.38 \times 10^{-23} \cdot 298 \cdot 5 \times 10^9 \cdot 1700} \\ &= 3.74 \times 10^{-4} \text{ V}. \end{aligned} \quad (13)$$

We consider the thermal noise RMS voltage of a *C. elegans* gap junction to be this 0.374 mV value.

B. Plateau Potentials

Although one might hope that the signaling scheme that can be used over a gap junction is unrestricted, there are biophysical constraints that impose limits. We describe these signaling limitations for *C. elegans*.

Signaling arises from regenerative events in neurons. Electrophysiologists have described four main types of regenerative events: action potentials, graded potentials, intrinsic oscillations, and plateau potentials. Plateau potentials are prolonged all-or-none depolarizations that can be triggered and terminated by brief positive- and negative-current pulses, respectively [48]. Plateau potentials are the biological equivalent of Schmitt triggers [48].

It is thought that plateau potentials may be used by many neurons in *C. elegans*, and that they may arise through synaptic interaction [48]. RMD neurons in *C. elegans* have two stable resting potentials, one near -70 mV and one near -35 mV [48], and these values are thought to hold across the nervous system. These two levels can be thought of as the two possible input levels to a gap junction channel:

$$v_0 = -70 \times 10^{-3} \text{ V}, \quad (14)$$

$$v_1 = -35 \times 10^{-3} \text{ V}, \quad (15)$$

as part of a random telegraph signal.

Switching between the two levels cannot happen arbitrarily quickly, as there is an absolute refractory period between regenerative events due to biochemical constraints. Unfortunately, the absolute refractory period for *C. elegans* is not known due to the difficulty in performing the requisite electrophysiology experiments [email communication with S. R. Lockery (14 Dec. 2010) and later confirmation].

The fastest impulse potentials observed in neurons in common reference are, perhaps, in the Renshaw cells of the mammalian spinal motor system, and have been reported as high as 1700 per second [49, p. 47]. Moving forward, we use this as a bound, however the typical value for the absolute refractory period across the animal kingdom is 1 ms; *C. elegans* may be even slower. This is where the noise bandwidth value of 1700 Hz also comes from.

C. Capacity

Having determined the noise distribution and the signaling constraints, we aim to find the capacity of this continuous-time, intertransition-time-restricted, binary-input AWGN channel. This capacity computation problem has been studied by Chayat and Shamai [50], assuming no timing jitter. Rather than using those precise results, we take a simplified approach through the discretization of time. We assume slotted pulse amplitude modulation, which is robust to any timing jitter that may be present in *C. elegans* signal propagation.

In particular, we consider a discrete-time channel with 1700 channel usages per second, two binary input levels of -35 and -70 , and AWGN noise with standard deviation 0.374 . As can be noted, the signal-to-noise ratio is rather high, $17.5^2/0.374^2 = 2.2 \times 10^3$, and so the capacity will be approximately one bit per channel usage.

Just to be sure, we compute this more precisely. The capacity is

$$C = h(Y) - \frac{1}{2} \log 2\pi e, \quad (16)$$

where $h(Y)$ is the differential entropy of the distribution:

$$p(y) = \frac{1}{2} \left[\frac{1}{\sqrt{2\pi}} \exp \left\{ \frac{-(y - \sqrt{\text{SNR}})^2}{2} \right\} \right] + \frac{1}{2} \left[\frac{1}{\sqrt{2\pi}} \exp \left\{ \frac{-(y + \sqrt{\text{SNR}})^2}{2} \right\} \right], \quad (17)$$

with $\text{SNR} = 2.2 \times 10^3$, see e.g. [51]. Performing the calculation demonstrates the rate loss below 1 bit per channel usage due to noise is negligible. Thus we assume that the capacity of a *C. elegans* gap junction is 1700 bits per second, or equivalently 5.9×10^{-4} seconds per bit.

Synaptic connection between two neurons may contain more than one gap junction.¹ Although it is difficult to maintain electrical separation between individual gap junctions, for the purposes of this paper we assume that each gap junction can act independently. Hence the channel capacity of parallel gap junction links is simply taken to be the number of gap junctions between the two neurons multiplied by the capacity of an individual gap junction.

¹The mean number of gap junctions between two connected neurons is 1.73; see [22, Fig. 3(b)] for the distribution, which is well-modeled by a power law with exponent 2.76.