

584 **A Hyperparameters for PNA dataset.**

585 In this section we provide the hyperparameters used for the different models on the PNA multitask  
 586 benchmark. We train all models for 2000 steps and with 3 layers. The remaining hyperparameters  
 587 for hidden size of each layer, learning rate, number of message passing steps (only valid for MPNN  
 588 models), number of rotation matrices and same example frequency (when relevant) are provided in  
 589 Table 6.

Table 6: Training hyperparameters for PNA dataset

Model	#Hidden size	L. rate	#MP steps	#Rotation matrices	#Same Examples
GAT	64	$10^{-4}$	-	-	-
GCN	64	$10^{-4}$	-	-	-
DGN	256	$10^{-3}$	-	-	-
MPNN	256	$10^{-3}$	2	-	-
ER GNN	128	$10^{-3}$	2	-	-
ER (node) embed.	64	$10^{-3}$	1	-	-
ER (edge) embed.	256	$10^{-3}$	2	-	-
ER (edge) embed.	256	$10^{-3}$	2	-	-
All ER features	256	$10^{-4}$	2	23	9
HT + ER (rand rot)	512	$10^{-4}$	2	23	4

590 **B Details of MPNN Framework**

591 As discussed previously, the architectures used in the experiments conform to the MPNN framework,  
 592 which allows affinity measures to be added as additional node and edge features. We describe the  
 593 details here for completeness.

594 Assume that our input graph,  $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ , has node features  $\mathbf{x}_u \in \mathbb{R}^n$ , edge features  $\mathbf{x}_{uv} \in \mathbb{R}^m$   
 595 and graph-level features  $\mathbf{x}_{\mathcal{G}} \in \mathbb{R}^l$ , for nodes  $u, v \in \mathcal{V}$  and edges  $(u, v) \in \mathcal{E}$ . We provide encoders  
 596  $f_n : \mathbb{R}^n \rightarrow \mathbb{R}^k$ ,  $f_e : \mathbb{R}^m \rightarrow \mathbb{R}^k$  and  $f_g : \mathbb{R}^l \rightarrow \mathbb{R}^k$  that transform these inputs into a latent space:

$$\mathbf{h}_u^{(0)} = f_n(\mathbf{x}_u) \quad \mathbf{h}_{uv}^{(0)} = f_e(\mathbf{x}_{uv}) \quad \mathbf{h}_{\mathcal{G}}^{(0)} = f_g(\mathbf{x}_{\mathcal{G}}) \quad (3)$$

597 Our MPNN then performs several message passing steps:

$$\mathbf{H}^{(t+1)} = P_{t+1}(\mathbf{H}^{(t)}) \quad (4)$$

598 where  $\mathbf{H}^{(t)} = \left( \left\{ \mathbf{h}_u^{(t)} \right\}_{u \in \mathcal{V}}, \left\{ \mathbf{h}_{uv}^{(t)} \right\}_{(u,v) \in \mathcal{E}}, \mathbf{h}_{\mathcal{G}}^{(t)} \right)$  contains all of the latents at a particular process-  
 599 ing step  $t \geq 0$ .

600 This process is iterated for  $T$  steps, recovering final latents  $\mathbf{H}^{(T)}$ . These can then be *decoded* into  
 601 node-, edge-, and graph-level predictions (as required), using analogous decoder functions  $g_n, g_e$  and  
 602  $g_g$ :

$$\mathbf{y}_u = g_n(\mathbf{h}_u^{(T)}) \quad \mathbf{y}_{uv} = g_e(\mathbf{h}_{uv}^{(T)}) \quad \mathbf{y}_{\mathcal{G}} = g_g(\mathbf{h}_{\mathcal{G}}^{(T)}) \quad (5)$$

603 Generally,  $f$  and  $g$  are simple MLPs, whereas we use the MPNN update rule for  $P$ . It computes  
 604 message vectors,  $\mathbf{m}_{uv}^{(t)}$ , to be sent across the edge  $(u, v)$ , and then aggregates them in the receiver  
 605 nodes as follows:

$$\mathbf{m}_{uv}^{(t+1)} = \psi_{t+1} \left( \mathbf{h}_u^{(t)}, \mathbf{h}_v^{(t)}, \mathbf{h}_{uv}^{(0)} \right), \quad \mathbf{h}_u^{(t+1)} = \phi_{t+1} \left( \mathbf{h}_u^{(t)}, \sum_{u \in \mathcal{N}_v} \mathbf{m}_{vu}^{(t+1)} \right) \quad (6)$$

606 The message function  $\psi_{t+1}$  and the update function  $\phi_{t+1}$  are both MLPs. All of our models have  
 607 been implemented using the jraph library [16].

608 We incorporate edge-based affinity features (e.g., effective resistances and hitting times) in  $f_e$   
 609 and node-based affinity features (e.g., resistive embeddings) in  $f_n$ . Note that node-based affinity

610 features may also naturally be incorporated as edge features by concatenating the node features at the  
611 endpoints.

612 Occasionally, the dataset in question will be easy to overfit with the most general form of message  
613 function (see (6)). In these cases, we resort to assuming that  $\psi$  factorises into an *attention mechanism*:

$$\mathbf{m}_{uv}^{(t+1)} = a_{t+1} \left( \mathbf{h}_u^{(t)}, \mathbf{h}_v^{(t)}, \mathbf{h}_{uv}^{(0)} \right) \psi_{t+1} \left( \mathbf{h}_u^{(t)} \right) \quad (7)$$

614 where the attention function  $a$  is scalar-valued. We will refer to this particular MPNN baseline as a  
615 graph attention network (GAT) [44].

## 616 C Omitted Proofs

617 **Lemma 3.2.** *For any pair of nodes  $u, v$ , we have  $\|\mathbf{r}_u - \mathbf{r}_v\|_2^2 = \text{Res}(u, v)$ .*

*Proof.*

$$\begin{aligned} \|\mathbf{r}_u - \mathbf{r}_v\|_2^2 &= \|C^{1/2} B L_G^{-1} (\mathbf{1}_u - \mathbf{1}_v)\|_2^2 \\ &= (\mathbf{1}_u - \mathbf{1}_v)^T L^\dagger (B^T C B) L^\dagger (\mathbf{1}_u - \mathbf{1}_v) \\ &= (\mathbf{1}_u - \mathbf{1}_v)^T L^\dagger L L^\dagger (\mathbf{1}_u - \mathbf{1}_v) \\ &= (\mathbf{1}_u - \mathbf{1}_v)^T L^\dagger (\mathbf{1}_u - \mathbf{1}_v) = \text{Res}(u, v). \quad \square \end{aligned}$$

618 **Corollary 4.2.** *For any fixed vectors  $\alpha, \beta \in \mathbb{R}^n$ , if we let  $X := \sum_i \alpha_i x_i$ ,  $\hat{X} := \sum_i \alpha_i \hat{x}_i$  and  
619 similarly  $Y := \sum_i \beta_i x_i$ ,  $\hat{Y} := \sum_i \beta_i \hat{x}_i$ ; then:*

$$\left| \langle X, Y \rangle - \langle \hat{X}, \hat{Y} \rangle \right| \leq \frac{\epsilon}{2} (\|X\|^2 + \|Y\|^2).$$

620 *Proof.* Since  $\langle X, Y \rangle = \frac{1}{4} (\|X + Y\|^2 - \|X - Y\|^2)$ , we can bound  $A = |\langle X, Y \rangle - \langle \hat{X}, \hat{Y} \rangle|$  from  
621 above as:

$$\begin{aligned} A &= \left| \frac{1}{4} \left( \|X + Y\|^2 - \|\hat{X} + \hat{Y}\|^2 - \|X - Y\|^2 + \|\hat{X} - \hat{Y}\|^2 \right) \right| \\ &\leq \frac{1}{4} \left( \left| \|X + Y\|^2 - \|\hat{X} + \hat{Y}\|^2 \right| + \left| \|X - Y\|^2 - \|\hat{X} - \hat{Y}\|^2 \right| \right) \\ &\leq \frac{1}{4} (\epsilon \cdot \|X + Y\|^2 + \epsilon \cdot \|X - Y\|^2) \quad (8) \\ &= \frac{\epsilon}{4} (\|X + Y\|^2 + \|X - Y\|^2) \\ &= \frac{\epsilon}{2} (\|X\|^2 + \|Y\|^2), \end{aligned}$$

622 where (8) follows from Lemma 4.1 with probability  $1 - o(1)$ , by our choice of  $k$  (as Lemma 4.1  
623 guarantees that each of  $\|\hat{X} + \hat{Y}\|^2 = (1 \pm \epsilon)\|X + Y\|^2$  and  $\|\hat{X} - \hat{Y}\|^2 = (1 \pm \epsilon)\|X - Y\|^2$  holds  
624 with probability  $1 - o(1)$ , and one can take a union bound over the two events).  $\square$

625 **Lemma 4.3.**  $H_{u,v} = 2M \langle \mathbf{r}_v - \mathbf{r}_u, \mathbf{r}_v - \mathbf{p} \rangle$  where  $\mathbf{p} := \sum_u \pi_u \mathbf{r}_u$ .

626 *Proof.* Consider the following expression of hitting times in terms of commute times by [43].

$$H_{u,v} = \frac{1}{2} \left[ K_{u,v} + \sum_i \pi_i (K_{v,i} - K_{u,i}) \right]. \quad (9)$$

627 Dividing both sides of eq. (9) and using the relation  $K_{u,v} = 2M \text{Res}(u, v)$ , we see that:

$$\begin{aligned} \frac{H_{u,v}}{2M} &= \frac{1}{2} \left[ \text{Res}(u, v) + \sum_i \pi_i (\text{Res}(v, i) - \text{Res}(u, i)) \right] \\ &= \frac{1}{2} \left[ \|\mathbf{r}_u - \mathbf{r}_v\|^2 + \sum_i \pi_i (\|\mathbf{r}_v - \mathbf{r}_i\|^2 - \|\mathbf{r}_u - \mathbf{r}_i\|^2) \right]. \quad (10) \end{aligned}$$

628 Let’s focus on the inner summation. After expanding out the squared norms, we see that:

$$\begin{aligned}
\sum_i \pi_i (\|\mathbf{r}_v - \mathbf{r}_i\|^2 - \|\mathbf{r}_u - \mathbf{r}_i\|^2) &= \sum_i \pi_i (\|\mathbf{r}_v\|^2 - \|\mathbf{r}_u\|^2) \\
&\quad - 2 \sum_i \pi_i \langle \mathbf{r}_v - \mathbf{r}_u, \mathbf{r}_i \rangle \\
&= (\|\mathbf{r}_v\|^2 - \|\mathbf{r}_u\|^2) \\
&\quad - 2 \langle \mathbf{r}_v - \mathbf{r}_u, \sum_i \pi_i \mathbf{r}_i \rangle \\
&= (\|\mathbf{r}_v\|^2 - \|\mathbf{r}_u\|^2) - 2 \langle \mathbf{r}_v - \mathbf{r}_u, \mathbf{p} \rangle.
\end{aligned}$$

629 Substituting this back into eq. (10), we can express  $\frac{1}{2M} H_{u,v}$  as:

$$\begin{aligned}
\frac{1}{2} (\|\mathbf{r}_v - \mathbf{r}_u\|^2 + \|\mathbf{r}_v\|^2 - \|\mathbf{r}_u\|^2 - 2 \langle \mathbf{r}_v - \mathbf{r}_u, \mathbf{p} \rangle) \\
= \|\mathbf{r}_v\|^2 - \langle \mathbf{r}_u, \mathbf{r}_v \rangle - \langle \mathbf{r}_v - \mathbf{r}_u, \mathbf{p} \rangle = \langle \mathbf{r}_v - \mathbf{r}_u, \mathbf{r}_v - \mathbf{p} \rangle. \quad \square
\end{aligned}$$

630 **Lemma 4.4.**  $|\widehat{H}_{u,v} - H_{u,v}| \leq 3\epsilon H_{\max}$ .

631 *Proof.* Using Lemma 4.3, we see that

$$\begin{aligned}
|\widehat{H}_{u,v} - H_{u,v}| &= 2M |\langle \widehat{\mathbf{r}}_v - \widehat{\mathbf{r}}_u, \widehat{\mathbf{r}}_v - \widehat{\mathbf{p}} \rangle - \langle \mathbf{r}_v - \mathbf{r}_u, \mathbf{r}_v - \mathbf{p} \rangle| \\
&\leq \epsilon M (\|\mathbf{r}_v - \mathbf{r}_u\|^2 + \|\mathbf{r}_v - \mathbf{p}\|^2) \\
&\leq 3\epsilon H_{\max},
\end{aligned}$$

632 where we used Corollary 4.2 in the first inequality and Definition 3.4 in the last inequality.  $\square$

## 633 D Comparison: Effective Resistances vs. Shortest Path Distances

634 Given that effective resistance (ER) captures times associated with random walks in a graph, it is  
635 tempting to ask how effective resistances compare to shortest path distances (SPDs) between nodes in  
636 a graph. Indeed, for some simple graphs, e.g., trees, shortest path distances and effective resistances  
637 turn out to be identical. However, in general, effective resistances and shortest path distances behave  
638 quite differently.

639 Nevertheless, it is tempting to ask how effective resistance features compare to SPD features in  
640 GNNs, especially as there have been a number of recent model architectures that make use of SPD  
641 features (e.g., Graphormer [49], Position-Aware GNNs [52], DE-GNN [29]). We first note that  
642 the most natural direct comparison of our ER-based MPNNs with SPD-based networks does not  
643 quite make sense. The reason is that the analogous comparison would be to determine the effect of  
644 replace ERs with SPDs as features in our MPNNs. However, since our networks only use ER features  
645 along edges of the given graph, the corresponding SPD features would then be trivial (as the SPD  
646 between two nodes directly connected by an edge in the graph is 1, resulting in a constant feature on  
647 every edge)!

648 As a result, graph learning architectures that use SPDs typically either (a.) use a densely-connected  
649 network (e.g., Graphormer [49], which uses a densely-connected attention mechanism) that incurs  
650  $O(n^2)$  overhead, or (b.) pick a small set of *anchor nodes* or *landmark nodes* to which SPDs from  
651 all other nodes are computed and incorporated as node features (e.g., Position-Aware GNNs [52],  
652 DE-GNN [29]). We stress that the former approach generally modifies the graph (by connecting  
653 all pairs of nodes) and therefore does not fall within the standard MPNN approach, while the latter  
654 includes architectures that fall within the MPNN paradigm.

655 Furthermore, we note that DE-GNNs are arguably one of the closest proposals to ours, as they  
656 compute distance-encoded features. These features can be at least as powerful as our proposed  
657 affinity-based features *if* polynomially many powers of the adjacency matrix are used. However,  
658 for all but the smallest graphs, using this many powers will be impractical—in fact, [29] only use  
659 powers of  $A$  up to 3, which would not be able to reliably approximate affinity-based features. We also

660 observe that the DE-GNN paper is concerned with learning representations of small sets of nodes  
 661 (e.g., node-, link-, and triangle-prediction) and does not show how to handle graph prediction tasks,  
 662 which the authors mention as possible future work. This makes a direct comparison of our methods  
 663 with DE-GNNs difficult.

## 664 D.1 Empirical Results

665 In an effort to empirically compare the expressivity of ER features with that of SPD features, we once  
 666 again perform experiments on the PNA dataset, picking the following baselines that make use of SPD  
 667 features:

- 668 • The first baseline is roughly an MPNN with *Graphormer-based features*. More precisely, it  
 669 is a densely-connected MPNN with SPDs *from the original graph* as edge features. In order  
 670 to retain the structure of the original graph, we also use additional edge features to indicate  
 671 whether or not an edge in the dense (complete) graph is a true edge of the original graph.  
 672 We also explore the use of the *centrality encoding* (in-degree and out-degree embeddings)  
 673 from Graphormer as additional node features.
- 674 • The second baseline is the Position-Aware GNN (P-GNN), which makes use of “anchor sets”  
 675 of nodes and encodes distances to these nodes.

676 The results of these baselines are shown in Table 7. In particular, we note that our ER-based MPNNs  
 677 outperform all aforementioned baselines.

Table 7: Results on the PNA dataset for MPNNs with Graphormer-based features (yellow) as well as SPD-based P-GNNs (orange). Here, CE refers to the *centrality encoding*, which is incorporated in the relevant MPNNs as additional node features. Similarly, SPD refers to *shortest path distance* features — in the relevant MPNNs, shortest path distances between all pairs of nodes in the graph are incorporated as edge features, along with an additional edge feature indicating whether an edge exists in the input graph. Therefore, the MPNN baselines are all variants of the same model with additional node/edge features. Similarly, P-GNN [52] uses SPD features with respect to a set of chosen *anchor nodes*. The average score metric is, as before, the average of the  $\log(MSE)$  metric over all six tasks, as in Table 1.

Model	Average score
*MPNN + CE	-2.728
*MPNN (dense) + SPD	-2.157
*MPNN (dense) + CE + SPD	-2.107
*P-GNN	-2.650
<b>MPNN w/ resistive (edge) embeddings</b>	<b>-2.789</b>
<b>MPNN w/ all affinity measure features</b>	<b>-3.106</b>

677

## 678 D.2 Theory: ER vs. SPD

679 In addition to experimental results, we would like to provide some theory for why effective resistances  
 680 can capture structure in GNNs that SPDs are unable to.

681 We will call an initialization function  $u \mapsto \mathbf{h}_u^{(0)}$  on nodes of a graph *node-based* if it assigns values  
 682 that are independent of the edges of the graph. Such an initialization is, however, allowed to depend  
 683 on node identities (e.g., for the single-source shortest path problem from a source  $s$ , one might find it  
 684 natural to define  $\mathbf{h}_s^{(0)} = 0$  and  $\mathbf{h}_u^{(0)} = +\infty$  for all  $u \neq s$ ).

685 Consider the task of computing “single-source effective resistances,” i.e., the effective resistance  
 686 from a particular node to every other node. We show that a GNN with a limited number of message  
 687 passing steps cannot possibly learn single-source effective resistances, even to nearby nodes.

688 **Theorem D.1.** *Suppose we fix  $k > 0$ . Then, given any node-based initialization function  $\mathbf{h}_u^{(0)}$ , it*  
 689 *is impossible for a GNN to compute single-source effective resistances from a given node  $w$  to any*  
 690 *nodes within a  $k$ -hop neighborhood.*

691 *More specifically, for any update rule*

$$\begin{aligned} \mathbf{m}_{uv}^{(t+1)} &= \psi_{t+1} \left( \mathbf{h}_u^{(t)}, \mathbf{h}_v^{(t)}, f_e(\mathbf{x}_{uv}) \right) \\ \mathbf{h}_u^{(t+1)} &= \phi_{t+1} \left( \mathbf{h}_u^{(t)}, f(\{\mathbf{m}_{uv} : v \in \mathcal{N}(u)\}) \right), \end{aligned} \quad (11)$$

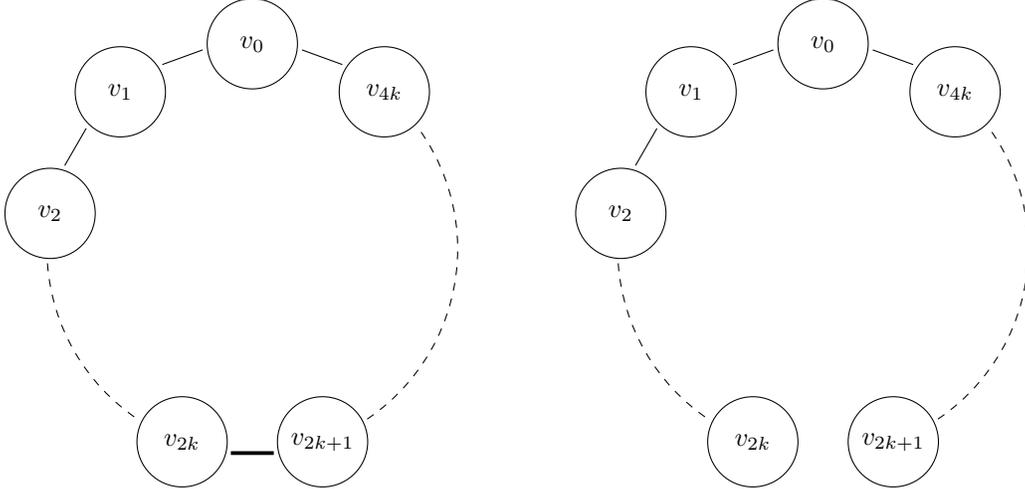
692 *there exists a graph  $G = (V, E)$  and  $u \in V$  such that after  $k$  rounds of message passing,  $h_v^{(k)} \neq$*   
 693  *$\text{Res}(u, v)$  for some  $v \neq u$  within a  $k$ -hop neighborhood of  $u$ .*

694 *On the other hand, there exists an initialization with respect to which  $k$  rounds of message passing*  
 695 *will compute the correct shortest path distances to all nodes within  $k$ -hop neighborhood.*

696 Note that the assumption on the initialization function in the above theorem is reasonable because  
 697 enabling the use of arbitrary, unrestricted functions would allow for the possibility of precomputing  
 698 effective resistances in the graph and trivially incorporating them as node features, which would  
 699 defeat the purpose of computing them using message-passing.

700 *Proof.* Consider the following set of graphs, each on  $4k + 1$  nodes:

Figure 2: Both of the above graphs are on  $4k + 1$  vertices, labeled  $v_0, v_1, \dots, v_{4k}$ . The only difference is a single edge, i.e., the graph on the left has an edge between  $v_{2k}$  and  $v_{2k+1}$ , while the one on the right does not have this edge.



701 Let  $V = \{v_0, v_1, \dots, v_{4k}\}$ . The first graph  $G = (V, E)$  is a cycle, while the second graph  $G' =$   
 702  $(V, E')$  is a path, obtained by removing a single edge from the first graph (namely, the one between  
 703  $v_k$  and  $v_{k+1}$ ). Suppose the edge weights are all 1 in the above graphs.

704 Let  $w = v_0$  be the source and let  $\{\mathbf{h}_v^{(0)} : v \in V\}$  be a “local” node feature initializa-  
 705 tion. Note that for any GNN (i.e., update and aggregation rules in (11), add the formal update  
 706 rule somewhere), the computation tree after  $k$  rounds of message passing is identical for nodes  
 707  $v_0, v_1, \dots, v_k, v_{3k+1}, v_{3k+2}, \dots, v_{4k}$  (i.e., the nodes within the  $k$ -hop neighborhood of  $v_0$ ) in both  
 708  $G$  and  $G'$ . This is because the only difference between  $G$  and  $G'$  is the existence of the edge  
 709 between  $v_{2k}$  and  $v_{2k+1}$ , and this edge is beyond a  $k$ -hop neighborhood centered at any one of the  
 710 aforementioned nodes. Therefore, we will necessarily have that  $\mathbf{h}_{v_i}^{(k)}$  is identical in both  $G$  and  $G'$  for  
 711  $i = 1, \dots, k, 3k + 1, 3k + 2, \dots, 4k$ .

712 However, it is easy to calculate the effective resistances in both graphs. In  $G$ , we have  $\text{Res}_G(v_0, v_i) =$   
 713  $\frac{i(4k+1-i)}{4k+1}$ , while in  $G'$ , we have  $\text{Res}_{G'}(v_0, v_i) = \min\{i, 4k + 1 - i\}$ . Therefore,  $\text{Res}_G(v_0, v_i) \neq$   
 714  $\text{Res}_{G'}(v_0, v_i)$  for all  $i = 1, 2, \dots, k, 3k + 1, 3k + 2, \dots, 4k$ .

715 It follows that for any  $i = 1, 2, \dots, k, 3k + 1, 3k + 2, \dots, 4k$ , the execution of  $k$  message passing  
 716 steps of a GNN cannot result in  $\mathbf{h}_{v_i}^{(k)} = \text{Res}(v_0, v_i)$  for both  $G$  and  $G'$ , which proves the first claim  
 717 of the theorem.

718 For the second part (regarding single-source shortest paths), observe that single-source shortest path  
719 distances can, indeed, be realized via aggregation and update rules for a message passing network. In  
720 particular, for  $k$  rounds of message passing, it is possible to learn shortest path distances of all nodes  
721 within a  $k$ -hop neighborhood. Specifically, for a source  $w$ , we can use the following setup: Take  
722  $\mathbf{h}_w = 0$  and  $\mathbf{h}_u = \infty$  for all  $u \neq w$ . Moreover, for any edge  $(u, v)$ , let the edge feature  $\mathbf{x}_{uv} \in \mathbb{R}$   
723 simply be the weight of  $(u, v)$  in the graph. Then, take the update rule (11) with  $f_e, \psi_{t+1}$  as identity  
724 functions and

$$\begin{aligned}
f_e(\mathbf{x}_{uv}) &= \mathbf{x}_{uv} \\
\psi_{t+1}(\mathbf{h}_u^{(t)}, \mathbf{h}_v^{(t)}, f_e(\mathbf{x}_{uv})) &= \mathbf{h}_u^{(t)} + \mathbf{x}_{uv} \\
f(S) &= \min_S \{s \in S\} \\
\phi_{t+1}(a, b) &= \min\{a, b\}.
\end{aligned}$$

725 It is clear that the above update rule simply simulates the execution of an iteration of the Bellman-Ford  
726 algorithm. Therefore,  $k$  message passing steps will simulate  $k$  iterations of Bellman-Ford, resulting  
727 in correct shortest path distances from the source  $w$  for every node within a  $k$ -hop neighborhood.  $\square$