Supplementary Information: Routing entanglement in the quantum internet

Pant et al.

SUPPLEMENTARY NOTE 1. ENTANGLEMENT GENERATION RATE AS A FUNCTION OF X,Y

The rate plots in Fig. 3 in the main paper are only in the $X = Y$ direction. Fig. 1 plots the rate as a function of X, Y which captures the behavior in all directions. The differences between different curves is similar across different directions, but there is an enhancement of the rate along the $X = Y$ direction because of the presence of multiple paths of similar length between Alice and Bob. This is visible in the contour plot Fig. 1(d).

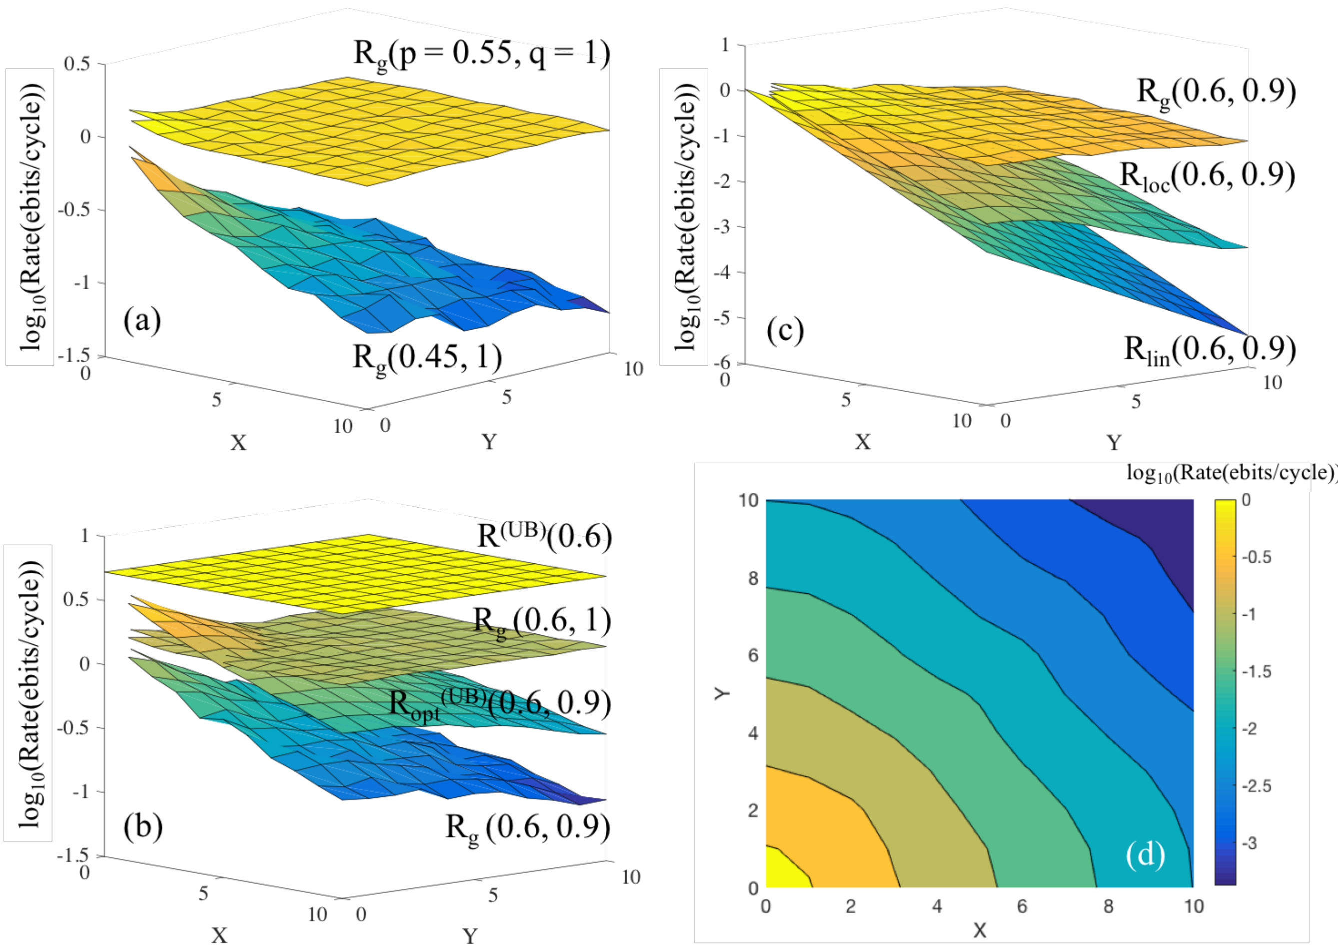

Supplementary Figure 1. Entanglement generation rate as a function of the Alice-Bob separation along X and Y (on a square grid) and(p, q); (a) $R_{\mathrm{g}}(p, q)$ is the rate attained by a global-knowledge-based protocol we propose where each node, in each time step, knows whether any link in the entire network succeeded or failed to establish entanglement. For the case of $q = 1$, R_{g} is distance independent when p is greater than the bond percolation threshold (0.5 for the square lattice) and decays exponentially if it is below the threshold. (b) $R^{(\mathrm{UB})}(0.6)$ is the distance-independent Pirandola rate upper bound for $p = 0.6$, achieving which requires perfect quantum processing at repeater nodes. $R_{\mathrm{g}}(0.6, 1)$ is also distance independent, and within a factor 3.6 of $R^{(\mathrm{UB})}(0.6)$. With $q < 1$, e.g., $R_g(0.6, 0.9)$, the rate decays exponentially with distance. $R^{(\mathrm{UB})}_{\mathrm{opt}}$ is an upper bound on the rate attainable with global-knowledge by any protocol. (c) R_{loc} is attained by a protocol we propose where each node, in each time step, only needs to know the link state of neighboring edges. The rate-distance scaling exponent of R_{loc} is clearly worse than R_{g}, but is significantly superior to that of a linear repeater chain along the shortest path, R_{lin}, demonstrating multi-path routing advantage even with local link-state knowledge. (d) Contour plot of the entanglement generation rate with the local rule when $p = 0.6$ and $q = 0.9$. Although the Alice to Bob distance along the network links is $X + Y$, there is a noticeable enhancement in the rate along the $X = Y$ direction because of more Alice-Bob paths of similar length.

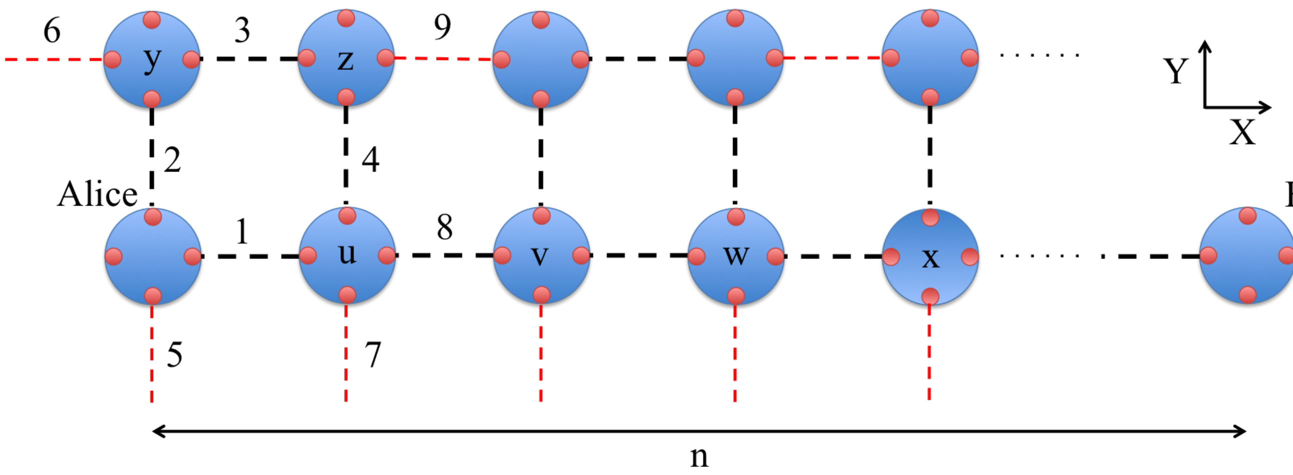

Supplementary Figure 2. Network used to prove the lower bound on entanglement generation rate with our local routing rule which shows that scaling of the rate with Alice-Bob manhattan distance for our rule is better than the scaling of the rate along a linear repeater chain along the shortest path between Alice-Bob.

SUPPLEMENTARY NOTE 2. MULTIPATH RATE ADVANTAGE

A. Analytical lower bound on the rate achieved by the local routing rule

In this subsection, we derive an analytical lower bound on the entanglement generation rate attained by our local routing rule (using the L^2 norm as the distance metric), with the objective of demonstrating multi-path routing advantage, i.e., the rate-vs.-distance scaling attained by our local rule is strictly better than that attained by a linear repeater chain along the shortest path between Alice and Bob.

Consider routing entanglement between Alice and Bob located at (X, Y) and $(X+n, Y)$ respectively, i.e., n hops apart along the X dimension of the square lattice. We will evaluate a lower bound on R_{loc} by only evaluating the rate contributions from paths in which all the (external) links belong to the set of black dashed links shown in Fig. 2. The choice of internal links made at repeater nodes proceed as usual per our local rule. As a result, there are instances in which our local rule routes entanglement through paths comprising not just the black links, resulting in flows that do not contribute to our rate lower bound.

We will refer to Fig. 2 for the ensuing discussion. Recall that external links succeed (are 'up') with probability p and fail (are 'down') with probability $1-p$, whereas internal links succeed with probability q. Consider $P(A \leftrightarrow v)$, the probability that there is a path between Alice A and repeater v that uses only black links. $P(A \leftrightarrow v)$ includes the probability of making the required internal links to create a path between A and v, but not the probability of any internal links at the end points A or v. It is easy to see that in any given time step, there can be no more than one edge-disjoint path between A and v along the black dashed links, since link 8 must be part of the path. Let l $\left(\widetilde{l}\right)$ be the event that the external link l is up (down). Further, note that at any given time step, of all the possible (0, 1 or 2) internal links attempted by our local rule at a repeater node, only one internal link, if successful, contributes to $A \leftrightarrow v$. Let $l-m$ be the event that the internal link attempted at a repeater node to connect external links l and m is successful. If links 1 and 8 are both up, node u attempts to connect those two links based on our local rule, regardless of the other links. If links 2, 3, 4 and 8 are up, but 1, 5, 6, 7 and 9 are down, u attempts to connect 4 and 8, z attempts to connect 3 and 4 and y attempts to connect 2 and 3. Considering these two possibilities, we have

$$\begin{aligned} P(A \leftrightarrow v) &> \Pr(1, 8, \text{1-8}) \\ &\quad + \Pr(2, 3, 4, 8, \widetilde{1}, \widetilde{5}, \widetilde{6}, \widetilde{7}, \widetilde{9}, \text{2-3}, \text{3-4}, \text{4-8}) \\ &= \left[p + p^3(1-p)^5 q^2\right] pq \\ &= p' pq, \end{aligned} \tag{1}$$

where $p' = p + p^3(1-p)^5 q^2 > p$.

$P(v \leftrightarrow x)$ is the probability that there is a path between v and x that uses only black links (the probability of internal link successes at the end points v and x are not included). $P(v \leftrightarrow x)$ and $P(A \leftrightarrow v)$ are not independent events because they both involve link 9. $P(v \leftrightarrow x|A \leftrightarrow v)$ is the probability that there exists a path along black dashed lines between v and X given that a path along black dashed lines exists between A and v. We now show that $P(v \leftrightarrow x|A \leftrightarrow v) > P(v \leftrightarrow x)$.

$$\begin{aligned}
P(v \leftrightarrow x|A \leftrightarrow v) &= P(v \leftrightarrow x|A \leftrightarrow v, 9)\Pr(9|A \leftrightarrow v) + \\
&\quad P\left(v \leftrightarrow x|A \leftrightarrow v, \widetilde{9}\right) \Pr\left(\widetilde{9}|A \leftrightarrow v\right) \\
&= P(v \leftrightarrow x|9)\Pr(9|A \leftrightarrow v) + \\
&\quad P\left(v \leftrightarrow x|\widetilde{9}\right) \Pr\left(\widetilde{9}|A \leftrightarrow v\right) \\
&= P(v \leftrightarrow x|9)\left(1 - \Pr\left(\widetilde{9}|A \leftrightarrow v\right)\right) + \\
&\quad P\left(v \leftrightarrow x|\widetilde{9}\right) \Pr\left(\widetilde{9}|A \leftrightarrow v\right) \\
&= \Pr\left(\widetilde{9}|A \leftrightarrow v\right) \times \\
&\quad \left(P\left(v \leftrightarrow x|\widetilde{9}\right) - P(v \leftrightarrow x|9)\right) \\
&\quad + P(v \leftrightarrow x|9)
\end{aligned} \tag{2}$$

where $P(v \leftrightarrow x|A \leftrightarrow v, 9) = P(v \leftrightarrow x|9)$ and $P\left(v \leftrightarrow x|A \leftrightarrow v, \widetilde{9}\right) = P\left(v \leftrightarrow x|\widetilde{9}\right)$ because link 9 being up or down is the only probabilistic event that influences both $P(A \leftrightarrow v)$ and $P(v \leftrightarrow x)$. Further,

$$\begin{aligned}
P(v \leftrightarrow x) &= P(v \leftrightarrow x|9)\Pr(9) + P\left(v \leftrightarrow x|\widetilde{9}\right) \Pr\left(\widetilde{9}\right) \\
&= P(v \leftrightarrow x|9)\left(1 - \Pr\left(\widetilde{9}\right)\right) + \\
&\quad P\left(v \leftrightarrow x|\widetilde{9}\right) \Pr\left(\widetilde{9}\right) \\
&= \Pr\left(\widetilde{9}\right)\left(P\left(v \leftrightarrow x|\widetilde{9}\right) - P(v \leftrightarrow x|9)\right) \\
&\quad + P(v \leftrightarrow x|9).
\end{aligned} \tag{3}$$

Comparing 2 and 3, $\Pr\left(\widetilde{9}|A \leftrightarrow v\right) = \Pr\left(\widetilde{9}\right) \Pr\left(A \leftrightarrow v|\widetilde{9}\right) / \Pr\left(A \leftrightarrow v\right) > \Pr\left(\widetilde{9}\right)$ because $\Pr\left(A \leftrightarrow v|\widetilde{9}\right) > \Pr\left(A \leftrightarrow v\right)$ following equation 1. Similarly, $\left(P\left(v \leftrightarrow x|\widetilde{9}\right) - P(v \leftrightarrow x|9)\right) > 0$. Hence, $P(v \leftrightarrow x|A \leftrightarrow v) > P(v \leftrightarrow x)$.

From Fig. 2, we can see that in order to get a path along black dashed lines from A to x, there must be a path along black dashed lines from A to v and from v to x, and the internal link at v must succeed. Therefore,

$$\begin{aligned}
P(A \leftrightarrow x) &= P(A \leftrightarrow v)qP(v \leftrightarrow x|A \leftrightarrow v) \\
&> P(A \leftrightarrow v)P(v \leftrightarrow x)q \\
&= \left(P(A \leftrightarrow v)\right)^2 q \\
&= (p'pq)^2 q,
\end{aligned} \tag{4}$$

where we use symmetry between $A \leftrightarrow v$ and $v \leftrightarrow x$ in the third line. Repeating this for all repeaters between Alice and Bob, it is easy to see that

$$\begin{aligned}
R_{\mathrm{loc}} > P(A \leftrightarrow B) &> p'^{\lceil n/2 \rceil} p^{\lfloor n/2 \rfloor} q^{n-1} \\
&\geq \left(\sqrt{p'p}\right)^n q^{n-1} \\
&= \left[\left(\sqrt{p'p}\right) q\right]^n q^{-1} \\
&= (pq)^{\beta n} q^{-1},
\end{aligned} \tag{5}$$

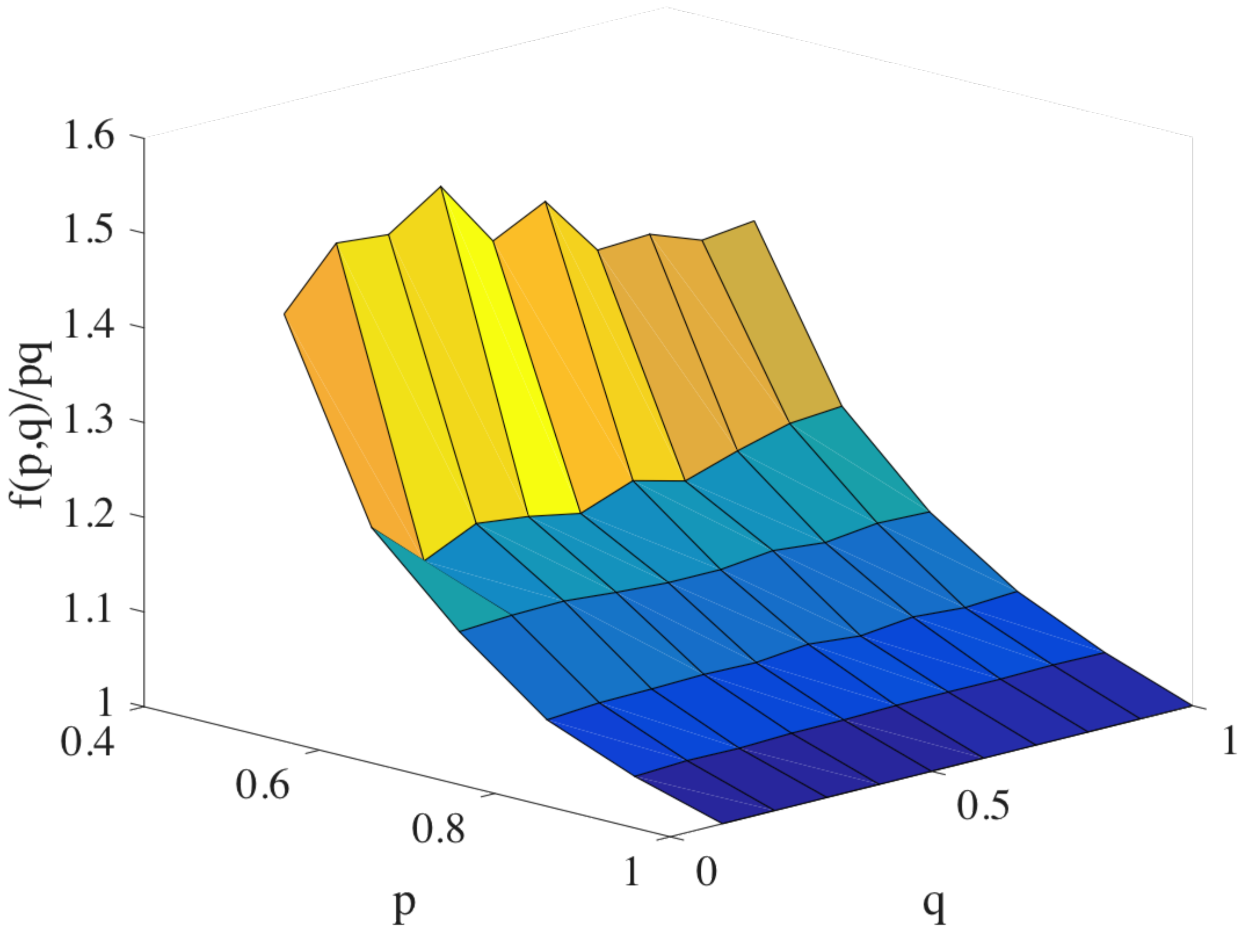

Supplementary Figure 3. $f(p,q)/pq$ quantifies the improvement in the scaling of $R_{\text{loc}}(p,q)$ with respect to $R_{\text{lin}}(p,q)$ with respect to the Alice-Bob Manhattan distance, n. $f(p,q)/pq$ increases as p is reduced in $[1,p_c]$ but changing q has a negligible effect.

where $\lceil n/2 \rceil$ is the smallest integer greater than or equal to $n/2$ and $\lfloor n/2 \rfloor$ is the largest integer smaller than or equal to $n/2$. The second inequality uses the fact that $p' > p$ and $n > 0$. $\beta = \log\left[\left(\sqrt{p'p}\right)q\right]/\log[pq] < 1$ because $p < p' < 1$ and $q < 1$.

Therefore, since $R_{\text{loc}} > (pq)^{\beta n}\, q^{-1}$ with $\beta < 1$ and $R_{\text{lin}} = (pq)^n\, q^{-1}$, the exponent in the scaling with n is smaller in R_{loc} compared to R_{lin}, i.e. the rate-vs.-distance scaling is better with multi-path routing. Using a similar reasoning, it is easy to see that the same is true even when Alice and Bob are at located at different Y coordinates. It should be noted that the lower bound we derive here is not meant to be tight (see Section SUPPLEMENTARY NOTE 2 B for a full numerical evaluation of the exponents for R_{loc} and R_{lin}). The only purpose of this subsection was to prove that the rate-vs.-distance scaling for entanglement routing strictly benefits from multi-path routing.

B. Numerical Evaluation

The goal this subsection is to quantify the improvement in the rate-vs.-distance exponent achieved by our local rule over that of a linear chain along the shortest path, for all possible pairs of values of p and q. Fig. 1(c) shows this improvement, i.e., that of $R_{\text{loc}}(p,q)$ compared to $R_{\text{lin}}(p,q)$, for $p = 0.6$ and $q = 0.9$. Clearly, $R_{\text{lin}}(p,q) = (pq)^n(p)/q \sim (1/q)[pq]^n$, where n is the Manhattan distance between Alice and Bob. We have numerically verified that $R_{\text{loc}}(p,q) \sim g(p,q)[f(p,q)]^n$ for n large. We hence quantify the rate improvement by numerically evaluating the ratio $f(p,q)/(pq)$ exhaustively for all $(p,q) \in [0,1] \times [0,1]$, using Monte Carlo simulations. The results are shown in Fig. 3, for configurations of Alice and Bob located along 45° with respect to the grid axes. We see that $f(p,q)/pq$

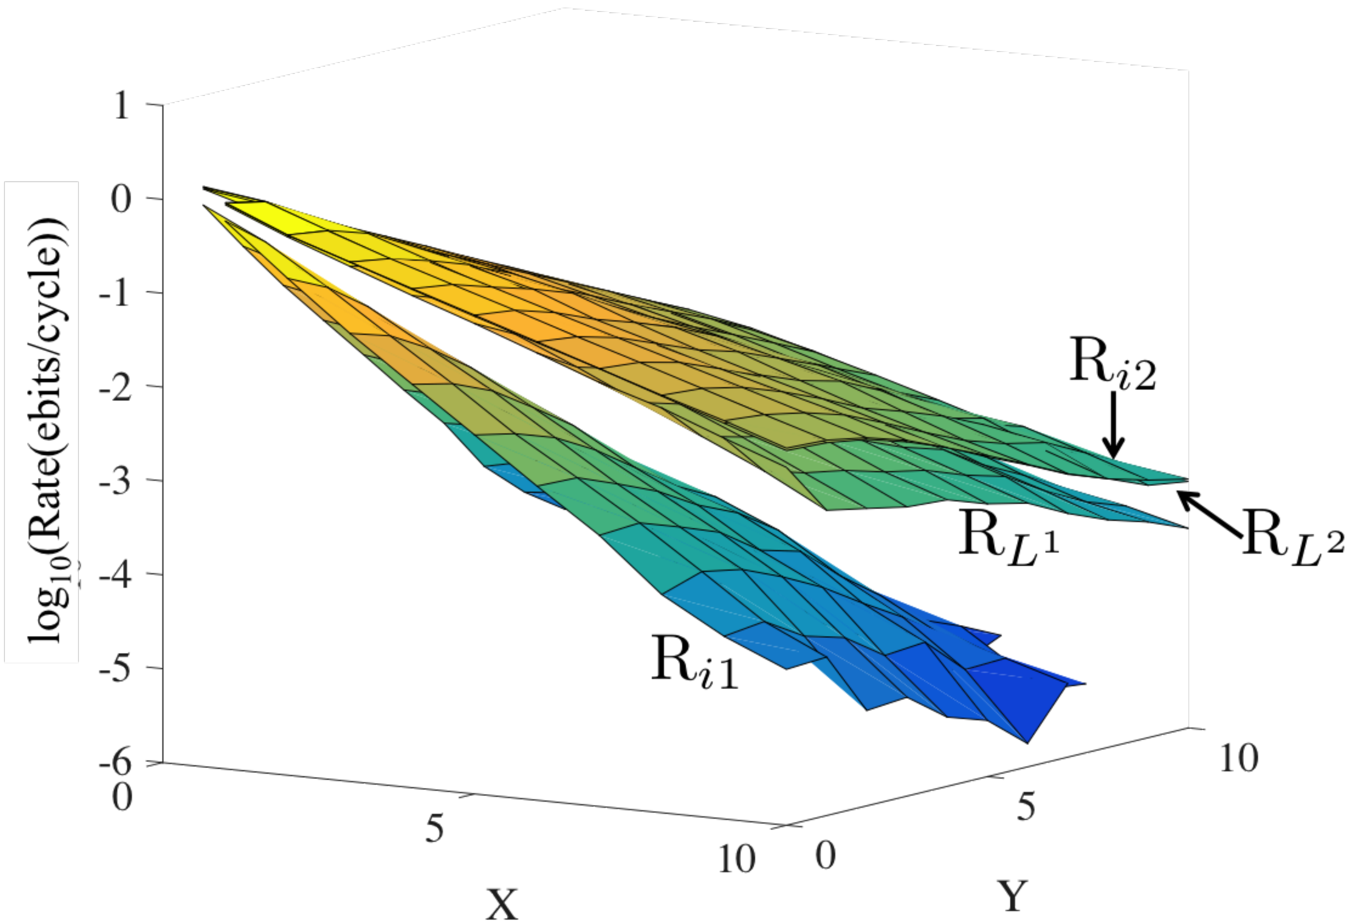

Supplementary Figure 4. Entanglement generation rates with different distance metrics. R_{L^1} and R_{L^2} are evaluated using the L_1 and L_2 norms respectively. The distance metric for R_{i1} (iteration 1) is calculated using R_{L^1}, and R_{i2} (iteration 2) is calculated using R_{i1}. R_{i2} and R_{L^1} are nearly indistinguishable as they coincide.

increases as p decreases in $[p_c, 1]$, but changing q has a negligible effect on this ratio.

SUPPLEMENTARY NOTE 3. DISTANCE METRIC FOR THE LOCAL ROUTING RULE USING L^1 NORM AND RECURSION

Our entanglement routing protocol with local link-state information uses the 'distance' of neighboring repeater stations from Alice and Bob to decide which memories at a repeater should undergo entanglement swap attempts. The results presented in the paper use the L^2 norm as the distance metric. While the L^2 norm can be easily calculated for the square grid, it may not be easily generalizable for other (e.g., non-planar) topologies. Further, even though we do not prove the rate optimality of our local link-state routing protocol, given a network topology, it is not clear whether or not the L^2 norm is the optimal distance metric to be used in our protocol.

In order to adapt our algorithm for arbitrary network topologies, and also to find a near-optimal distance metric for our algorithm, we employ the following numerical recursive method. Our evaluation begins with calculating $R_{L^1}(\mathbf{n}_1, \mathbf{n}_2)$, the entanglement generation rate achieved when our local rule is used to route entanglement between nodes $\mathbf{n}_1$ and $\mathbf{n}_2$, using the L^1 norm as the distance metric. In Fig. 4, we plot $R_{L^1}(\mathbf{n}_1, \mathbf{n}_2)$ as a function of (X, Y), where X and Y are the distance (in hops) between $\mathbf{n}_1$ and $\mathbf{n}_2$ along the horizontal and vertical dimensions of the square grid, respectively. The rate-distance scaling exponent for R_{L^1} is worse than that of R_{L^2}, the rate attained by our protocol, using the L^2 norm as the distance metric. Next, for every repeater node $\mathbf{n}$, we define distances d_A and

d_B to Alice **A** and Bob **B** respectively, with respect to the following new distance metric (let us name this metric $i1$): $d_A := 1/\mathrm{R}_{L^1}(\mathbf{n}, \mathbf{A})$ and $d_B := 1/\mathrm{R}_{L^1}(\mathbf{n}, \mathbf{B})$. We then calculate $\mathrm{R}_{i1}(\mathbf{n}_1, \mathbf{n}_2)$, the entanglement generation rate achieved when our local rule is used with the $i1$ distance metric to route entanglement between every pair of nodes $\mathbf{n}_1$ and $\mathbf{n}_2$. In Fig. 4, we plot $\mathrm{R}_{i1}(\mathbf{n}_1, \mathbf{n}_2)$ as a function of (X, Y). We see that the rate-distance scaling achieved by R_{i1} is even lower than that of R_{L^1}. However, when we go through the second iteration of the algorithm—i.e., define distance metric $i2$, under which $d_A = 1/\mathrm{R}_{i1}(\mathbf{n}, \mathbf{A})$ and $d_B = 1/\mathrm{R}_{i1}(\mathbf{n}, \mathbf{B})$, and use our local rule to evaluate $\mathrm{R}_{i2}(\mathbf{n}_1, \mathbf{n}_2)$ as a function of (X, Y)—we find that the resulting rate R_{i2} is almost the same (visually indistinguishable in the plot) as R_{L^2}, the rate we obtained directly when using the L^2 norm as the distance metric. This suggests that: (a) for the square grid (and presumably for any planar network topology) the L^2 norm metric might be near-optimal for use within our local rule, and that (b) for any given network topology, one could potentially pre-compute the optimal distance metric by a recursive strategy on the given topology using the L^1 norm as the starting point. However, there are instances where our local rule does not give the rate-optimal local routing rule. As an example, when $p = 1$ and $q = 1$, it is possible to find four disjoint paths without any link-state knowledge (the links are all deterministic) and the optimal rate is four ebits/cycle for any location of Alice and Bob. However, the fact that we are trying to route every flow through the best possible path without any coordination between different flows leads to collisions, which results in a rate that is below the optimal rate of four ebits/cycle. Finding the rate-optimal local routing rule across different parameter values is left for future research.