Deterministic delivery of remote entanglement on a quantum network

Humphreys, Peter C.; Kalb, Norbert; Morits, Jaco P.J.; Schouten, Raymond N.; Vermeulen, Raymond F.L.; Twitchen, Daniel J.; Markham, Matthew; Hanson, Ronald

DOI
10.1038/s41586-018-0200-5

Publication date
2018

Document Version
Accepted author manuscript

Published in
Nature

Deterministic delivery of remote entanglement on a quantum network

Peter C. Humphreys,[1,*] Norbert Kalb,[1,*] Jaco P. J. Morits,[1] Raymond N. Schouten,[1] Raymond F. L. Vermeulen,[1] Daniel J. Twitchen,[2] Matthew Markham,[2] and Ronald Hanson[1]

[1] *QuTech & Kavli Institute of Nanoscience, Delft University of Technology, PO Box 5046, 2600 GA Delft, The Netherlands*

[2] *Element Six Innovation, Fermi Avenue, Didcot, Oxfordshire OX11 0QE, U.K.*

Large-scale quantum networks promise to enable secure communication, distributed quantum computing, enhanced sensing and fundamental tests of quantum mechanics through the distribution of entanglement across nodes [1–7]. Moving beyond current two-node networks [8–13] requires the rate of entanglement generation between nodes to exceed their decoherence rates such that entanglement can be generated faster than it is lost. Beyond this critical threshold, intrinsically probabilistic entangling protocols can be subsumed into a powerful building block that deterministically provides remote entangled links at pre-specified times. Here we surpass this threshold using diamond spin qubit nodes separated by 2 metres. We realise a fully heralded single-photon entanglement protocol that achieves entangling rates of up to 39 Hz, three orders of magnitude higher than previously demonstrated two-photon protocols on this platform [14]. At the same time, we suppress the decoherence rate of remote entangled states to 5 Hz by dynamical decoupling. By combining these results with efficient charge-state control and mitigation of spectral diffusion, we are able to deterministically deliver a fresh remote state with average entanglement fidelity exceeding 0.5 at every clock cycle of ∼100 ms without any pre- or post-selection. These results demonstrate a key building block for extended quantum networks and open the door to entanglement distribution across multiple remote nodes.

* These authors contributed equally to this work.

The power of future quantum networks will derive from entanglement that is shared between the network nodes. Two critical parameters for the performance of such networks are the entanglement generation rate r_{ent} between nodes and the entangled-state decoherence rate r_{dec}. Their ratio, that we term the quantum link efficiency $\eta_{\text{link}} = r_{\text{ent}}/r_{\text{dec}}$ [8, 15], quantifies how effectively entangled states can be preserved over the timescales necessary to generate them. Alternatively, from a complementary perspective, the link efficiency determines the average number of entangled states that can be created within one entangled state lifetime. Achieving a link efficiency of unity therefore represents a critical threshold beyond which entanglement can be generated faster than it is lost. Exceeding this threshold is central to allowing multiple entangled links to be created and maintained simultaneously, as required for the distribution of many-body quantum states across a network [6, 15].

Consider an elementary entanglement delivery protocol that delivers states at predetermined times. This can be achieved by making multiple attempts to generate entanglement, and then protecting successfully generated entangled states from decoherence until the required delivery time (Fig. 1a, steps 1, 2 & 3). If we try to generate entanglement for a period t_{ent}, the cumulative probability of success will be $p_{\text{succ}} = 1 - e^{-r_{\text{ent}} t_{\text{ent}}}$. For a given p_{succ}, the average fidelity F_{succ} of the successfully generated states is solely determined by the quantum link efficiency η_{link} (derivation is given in methods). We plot F_{succ} versus p_{succ} for several values of η_{link} in Fig 1b.

This protocol allows entangled states to be delivered at specified times, but with a finite probability of success. By delivering an unentangled state (state fidelity $F_{\text{unent}} \leq \frac{1}{2}$) in cycles in which all entanglement generation attempts failed, the protocol can be cast into a fully deterministic black-box (Fig 1a, step 4). The states output from such a black-box will have a fidelity with a Bell state of

$$F_{\text{det}} = p_{\text{succ}} F_{\text{succ}} + (1 - p_{\text{succ}}) F_{\text{unent}}. \tag{1}$$

The maximum achievable fidelity $F_{\text{det}}^{\text{max}}$ of this deterministic state delivery protocol, found by optimising p_{succ}, is also only determined by the quantum link efficiency η_{link}. For $F_{\text{unent}} = \frac{1}{4}$ (fully mixed state), we find (see Fig 1c):

$$F_{\text{det}}^{\text{max}} = \frac{1}{4}\left(1 + 3\eta_{\text{link}}^{\frac{1}{1-\eta_{\text{link}}}}\right). \tag{2}$$

Beyond the threshold $\eta_{\text{link}} \gtrsim 0.83$, there exists a combination of p_{succ} and F_{succ} high enough to compensate for cycles in which entanglement is not heralded, allowing for the deterministic delivery of states that are on-average entangled ($F_{\text{det}}^{\text{max}} \geq \frac{1}{2}$). Demonstrating deterministic entanglement delivery therefore presents a critical benchmark of a network's performance, certifying that the network quantum link efficiency is of order unity or higher. Furthermore, the ability to specify in advance the time at which entangled states are delivered may assist in designing multi-step quantum information tasks such as entanglement routing [16, 17].

To date, this threshold has remained out of reach for solid-state quantum networks. Quantum dots have demonstrated kHz entanglement rates r_{ent}, but tens of MHz decoherence rates r_{dec} limit their achieved quantum link efficiencies to $\eta_{\text{link}} \sim 10^{-4}$ [18, 19]. Nitrogen vacancy (NV) centres, point-defects in diamond with a long-lived electron spin and bright optical transitions, have demonstrated entanglement rates r_{ent} of tens of mHz [10, 14] and, in separate experiments, decoherence rates r_{dec} of order 1 Hz [20], which would together give $\eta_{\text{link}} \sim 10^{-2}$.

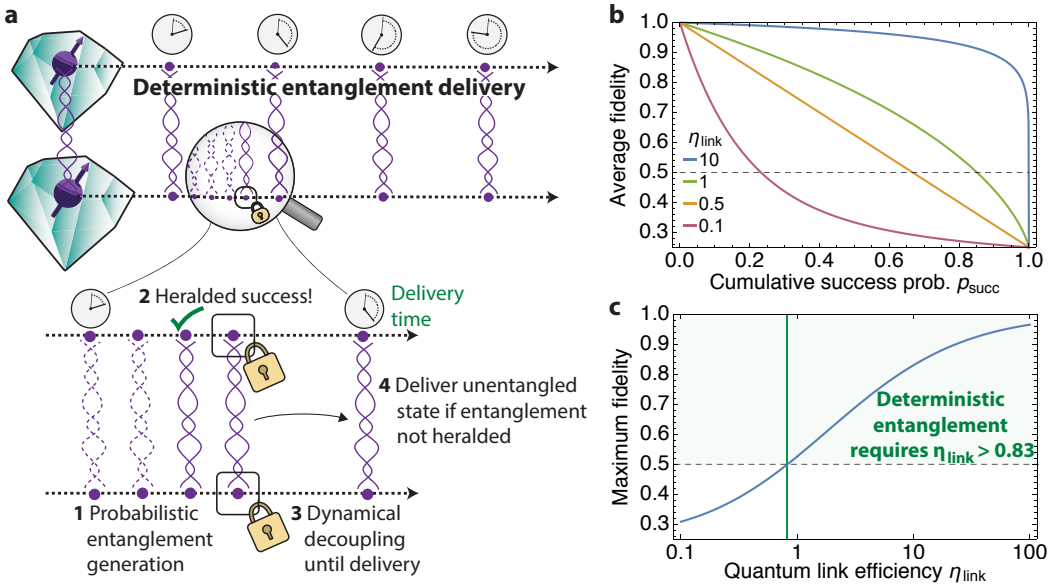

FIG. 1. **Deterministic remote entanglement delivery. a**, Deterministic entanglement delivery guarantees the output of states with average entanglement fidelity exceeding 0.5 at pre-specified times. In our protocol, underlying this deterministic delivery is a probabilistic but heralded entanglement process. Repeated entangling attempts are made and then, upon heralded success, the entangled state is protected from decoherence until the specified delivery time. If no attempt at entanglement generation succeeds within one cycle, an unentangled state must be delivered. **b**, For the underlying entanglement generation and state preservation protocol (steps 1, 2 & 3 in (a)), the effectiveness of the trade-off between the average fidelity of the delivered entangled state and the success probability is determined by the quantum link efficiency η_{link}. **c**, Maximum fidelity of deterministically delivered states as a function of η_{link}. A critical threshold $\eta_{\mathrm{link}} \gtrsim 0.83$ must be surpassed in order to deliver an on-average entangled state at every cycle.

Here we achieve η_{link} well in excess of unity by realising an alternative entanglement protocol for NV centres in which we directly use the state heralded by the detection of a single photon (Fig. 2) [21, 22]. The rate for such single-photon protocols scales linearly with losses, which, in comparison with previously used two-photon-mediated protocols [9, 14], provides a dramatic advantage in typical remote entanglement settings. Recent experiments have highlighted the potential of such single-photon protocols by generating local entanglement [23, 24], and remote entanglement in post-selection [18, 19]. By realising a single-photon protocol in a fully heralded fashion and protecting entanglement through dynamical decoupling, we achieve the deterministic delivery of remote entangled states on a ∼10 Hz clock.

Our experiment employs NV centres residing in independently operated cryostat setups separated by 2 metres (further experimental details are given in methods). We use qubits formed by two of the NV centre ground-state spin sub-levels ($|\uparrow\rangle \equiv |m_s = 0\rangle, |\downarrow\rangle \equiv |m_s = -1\rangle$). Single-photon entanglement generation (Fig. 2a) proceeds by first initialising each node in $|\uparrow\rangle$ by optical pumping [25], followed by a coherent rotation using a microwave

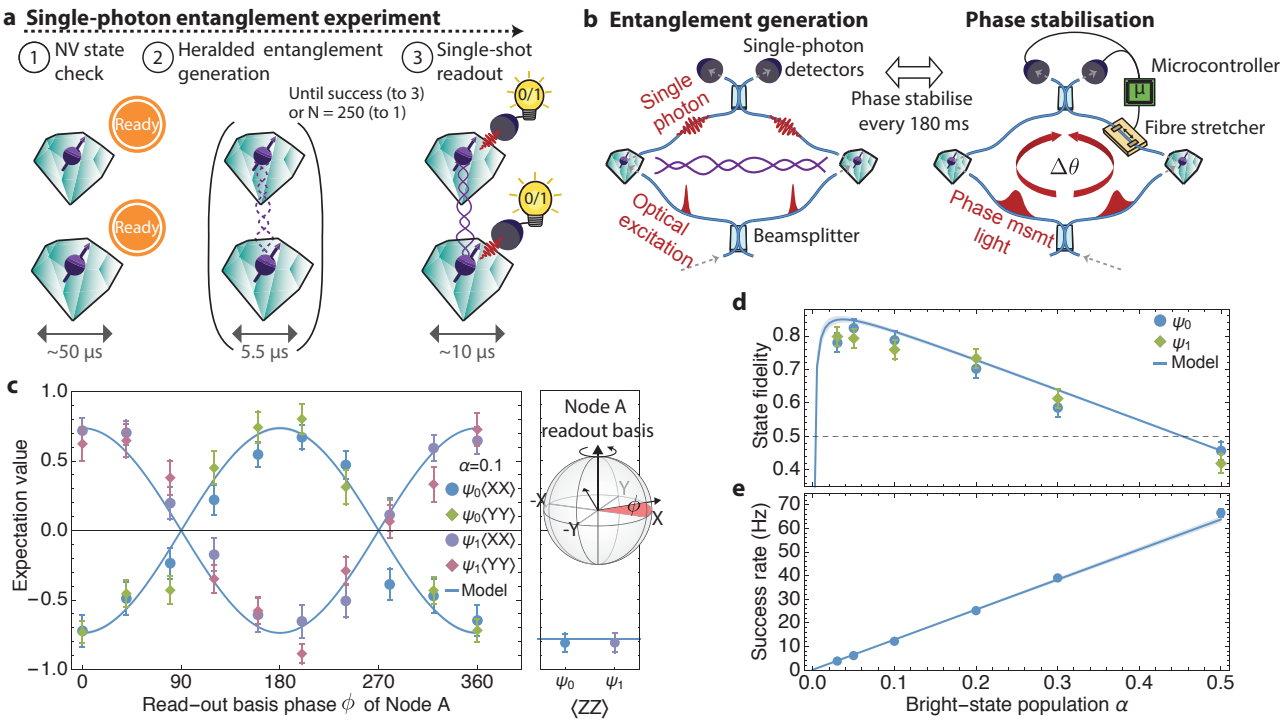

FIG. 2. **Benchmarking single-photon entanglement generation.** **a**, Experimental protocol: ① Before entanglement generation, an NV state check verifies that the NV centre is in the correct charge state (NV^-) and resonant with the excitation laser (discussed further in methods). This is repeated until the check passes. ② Entanglement generation is attempted until success is heralded, in which case we continue to readout. If 250 attempts have been made without success, we revert to step 1. ③ Upon heralded success, the spin states are read out in a chosen basis by using microwaves to rotate the state followed by single-shot readout. **b**, The optical phase difference $\Delta\theta$ acquired in an interferometer formed by the two nodes must be known. For the data reported in this figure, we stabilise the phase difference every 180 ms. **c**, Measured $\langle XX\rangle$ and $\langle YY\rangle$ correlations (left panel) for $\psi_{0/1}$ (where 0/1 denotes the heralding detector) and $\alpha = 0.1$ as readout basis is swept at node A. Right panel shows the measured $\langle ZZ\rangle$ correlations. **d**, Fidelity of the heralded states with a Bell state and **e**, entanglement generation success rate, for different values of α. For all plots, solid lines (with shaded 1 standard deviation (s.d.) statistical uncertainties) give the predictions of our model solely based on independently determined parameters (see methods). The data error bars represent 1 s.d.

pulse [26] to create the state

$$|NV\rangle = \sqrt{\alpha}\,|\uparrow\rangle + \sqrt{1-\alpha}\,|\downarrow\rangle\,. \tag{3}$$

We then apply resonant laser light to selectively excite the 'bright' state $|\uparrow\rangle$ to an excited state, which rapidly decays radiatively back to the ground state by emitting a single photon. This entangles the spin state of the NV with the presence $|1\rangle$ or absence $|0\rangle$ of a photon in the emitted optical mode:

$$|NV, \text{optical mode}\rangle = \sqrt{\alpha}\,|\uparrow\rangle\,|1\rangle + \sqrt{1-\alpha}\,|\downarrow\rangle\,|0\rangle\,. \tag{4}$$

Emitted photons are transmitted to a central station at which a beamsplitter is used to remove their which-path information. Successful detection of a photon at this station indicates that at least one of the NVs is in the bright state $|\uparrow\rangle$ and therefore heralds the creation of a spin-spin entangled state. However, given the detection of one photon, the conditional probability that the other NV is also in the state $|\uparrow\rangle$, but the photon it emitted was lost, is given by $p = \alpha$ (in the limit that the photon detection efficiency $p_{\text{det}} \ll 1$). This degrades the heralded state from a maximally-entangled Bell state $|\psi\rangle = \frac{1}{\sqrt{2}}(|\uparrow\downarrow\rangle + |\downarrow\uparrow\rangle)$ to

$$\rho_{NV,NV} = (1-\alpha)\,|\psi\rangle\langle\psi| + \alpha\,|\uparrow\uparrow\rangle\langle\uparrow\uparrow|\,. \tag{5}$$

The probability of successfully heralding entanglement is given by $2\,p_{\text{det}}\alpha$. The state fidelity $F = 1-\alpha$ can therefore be directly traded off against the entanglement rate. The corresponding success probability of a two-photon protocol is given by $\frac{1}{2}p_{\text{det}}^2$; for a given acceptable infidelity α, single-photon protocols will thus provide a rate increase of $4\,\alpha/p_{\text{det}}$. For example, for our system's $p_{\text{det}} \sim 4\times10^{-4}$, if a 10% infidelity is acceptable, the rate can be increased by three orders of magnitude over two-photon protocols.

The primary challenge in implementing single-photon entanglement is that the resulting entangled state depends on the optical phase acquired by the laser pulses used to create spin-photon entanglement at each node, as well as the phase acquired by the emitted single photons as they propagate (Fig 2b). The experimental setup therefore acts as an interferometer from the point at which the optical pulses are split to the point at which the emitted optical modes interfere. For a total optical phase difference of $\Delta\theta$, the entangled state created is given by

$$\left|\psi_{0/1}(\Delta\theta)\right\rangle = |\uparrow\downarrow\rangle \pm e^{\mathrm{i}\Delta\theta}\,|\downarrow\uparrow\rangle\,, \tag{6}$$

where 0/1 (with corresponding $\pm$ phase factor) denotes which detector at the central station detected the incident photon. This optical phase difference must be known in order to ensure that entangled states are available for further use.

We overcome this entangled-state phase sensitivity by interleaving periods of optical-phase stabilisation with our entanglement generation. During phase stabilisation we input bright laser light at the same frequency as the NV excitation light and detect the light reflected from the diamond substrate using the same detectors that are used to herald entanglement. The measured optical phase, estimated from the detected counts, is used to adjust the phase back to our desired value using a piezoelectric fibre stretcher. We achieve an average steady-state phase stability of 14.3(3)°, limited by the mechanical oscillations of the optical elements in our experimental setup (discussed further in methods and Extended Data Fig. 6).

To demonstrate the controlled generation of entangled states, we run the single-photon entangling protocol with a bright-state population of $\alpha = 0.1$. After entanglement is heralded, we apply basis rotations and single-shot state readout [25] at each node to measure $\langle\sigma_i^A\sigma_j^B\rangle$ correlations between the nodes, where the standard Pauli matrices will be referred to here in the shorthand $\sigma_X, \sigma_Y, \sigma_Z = X, Y, Z$. We observe strong correlations both for $\langle XX\rangle$ and $\langle YY\rangle$, and, when sweeping the readout basis for node A, oscillations of these coherences as expected from the desired entangled state (Fig. 2c, left panel). In combination with the measured $\langle ZZ\rangle$ correlations (Fig. 2c, right panel), this unambiguously proves the establishment of entanglement between our nodes.

We explore the tradeoff between the entangled state fidelity and the entanglement rate by measuring $\langle XX\rangle$, $\langle YY\rangle$ and $\langle ZZ\rangle$ correlations for a range of different initial bright-state

populations α. Using these correlations, we calculate the fidelity of the heralded state to the desired maximally entangled Bell state for each value of α (Fig. 2d), along with the measured success rate (Fig. 2e). As predicted, the fidelity increases with decreasing α as the weight of the unentangled state $|\uparrow\uparrow\rangle\langle\uparrow\uparrow|$ diminishes (Eqn. 5). For small α, the fidelity saturates because the detector dark-count rate becomes comparable to the detection rate.

Choosing α to maximise fidelity, we find that our protocol allows us to generate entanglement with a fidelity of 0.81(2) at a rate of $r_{\text{ent}} = 6$ Hz (for $\alpha = 0.05$). Alternatively, by trading the entanglement fidelity for rate, we can generate entanglement at $r_{\text{ent}} = 39$ Hz with an associated fidelity of 0.60(2) ($\alpha = 0.3$). This represents a two orders of magnitude increase in the entangling rate over all previous NV experiments [10] and a three orders of magnitude increase in rates over two-photon protocols under the same conditions [14].

Compared to the maximum theoretical fidelity for $\alpha = 0.05$ of 0.95, the states we generate have a 3% reduction in fidelity due to residual photon distinguishability, 4% from double excitation, 3% from detector dark counts, and 2% from optical-phase uncertainty (discussed further in methods).

In order to reach a sufficient link efficiency η_{link} to allow for deterministic entanglement delivery, the single-photon protocol must be combined with robust protection of the generated remote entangled states. To achieve this, we carefully shielded our NVs from external noise sources including residual laser light and microwave amplifier noise, leaving as the dominant noise the slowly-fluctuating magnetic field induced by the surrounding nuclear spin bath.

We mitigate this quasi-static noise by implementing dynamical decoupling with XY8 pulse sequences (Fig. 3a, see methods and Extended Data Fig. 9). The fixed delay between microwave pulses in these sequences is optimised for each node [27]. Varying the number of decoupling pulses allows us to protect the spins for different durations. This dynamical decoupling extends the coherence time of Node A and B from a T_2^* of $\sim 5\,\mu$s to 290(20) ms and 680(70) ms respectively, as shown in Fig. 3b. The difference in coherence times for the two nodes is attributed to differing nuclear spin environments and microwave pulse fidelities.

To investigate the preservation of remote entangled states, we incorporate dynamical decoupling for varying time durations after successful single-photon entanglement generation (Fig. 3c). We find an entangled state coherence time of 200(10) ms (decoherence rate r_{dec} of 5.0(3) Hz). The observed entangled-state fidelities closely match the predictions of our model, which is solely based on independently determined parameters (see methods and Extended Data Table 1). In particular, the decoherence of the remote entangled state is fully explained by the combination of the individual decoherence rates of the individual nodes.

The combination of dynamical decoupling with the single-photon entanglement protocol achieves a quantum link efficiency of $\eta_{\text{link}} \sim 8$ (comparable to the published state-of-the-art in ion traps, $\eta_{\text{link}} \sim 5$ [8]), pushing the NV-based platform well beyond the critical threshold of $\eta_{\text{link}} \gtrsim 0.83$.

These innovations enable the design a deterministic entanglement delivery protocol that guarantees the delivery of entangled states at specified intervals, without any post-selection of results or pre-selection based on the nodes being in appropriate conditions (Fig. 4a). Phase stabilisation occurs at the start of each cycle, after which there is a preset period before an entangled state must be delivered. This window must therefore include all NV state checks (necessary to mitigate spectral diffusion via feedback control and verify the charge-state and resonance conditions [9]), entanglement generation attempts and dynamical decoupling

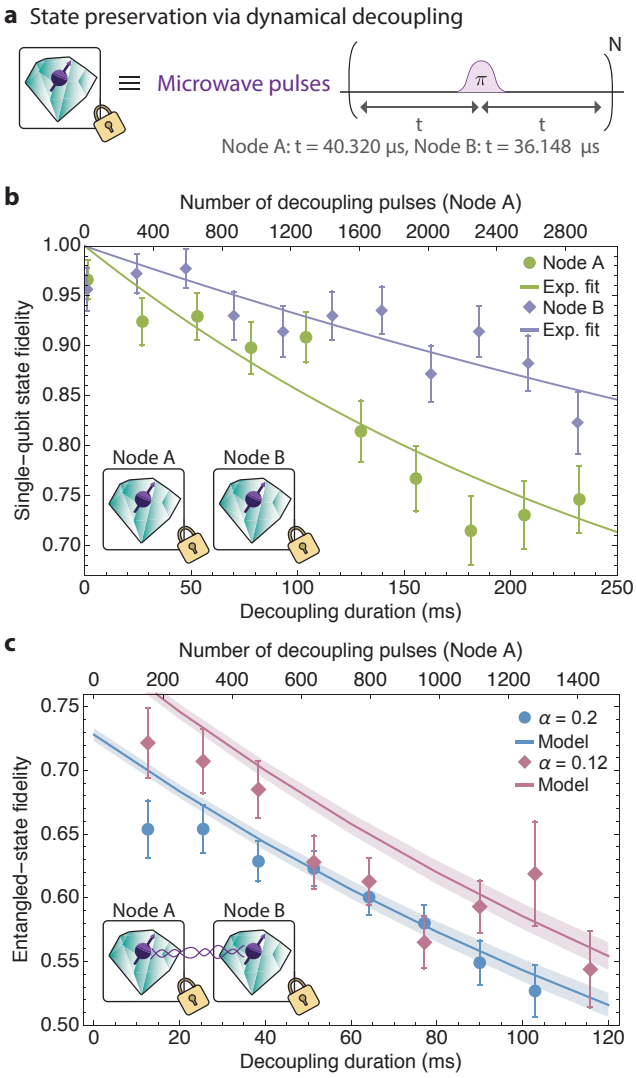

FIG. 3. **Coherence protection of remote entangled states. a**, Dynamical decoupling protects the state of the NV spins from quasi-static environmental noise. **b**, Fidelity with the initial state for dynamical decoupling of the state $\frac{1}{\sqrt{2}}(|\uparrow\rangle+|\downarrow\rangle)$ at each of our nodes. Solid lines show exponential fits with coherence times of 290(20) ms and 680(70) ms for nodes A and B respectively. **c**, Dynamical decoupling of entangled states created using the single-photon entanglement protocol for $\alpha = 0.12$ & 0.2. Solid lines (with shaded 1 s.d. statistical uncertainties) show the predictions of our model (discussed further in methods) based on the data in (b), from which the entangled state coherence time is expected to be $\tau = 200(10)$ ms. For both plots, data error bars represent 1 s.d.

necessary to deliver an entangled state. Fast conditional logic is used to dynamically adapt the experimental sequence based on the detection of a heralding signal [9, 10, 28, 29]. Further details on the experimental implementation are given in methods and Extended Data Fig. 1.

We run our deterministic entanglement delivery protocol at two values of α (0.2 & 0.12) and for delivery rates ranging from 7-12 Hz. We divide the experiment into runs of 1500 cycles (i.e. 1500 deterministic state deliveries), for a total data set of 42000 cycles.

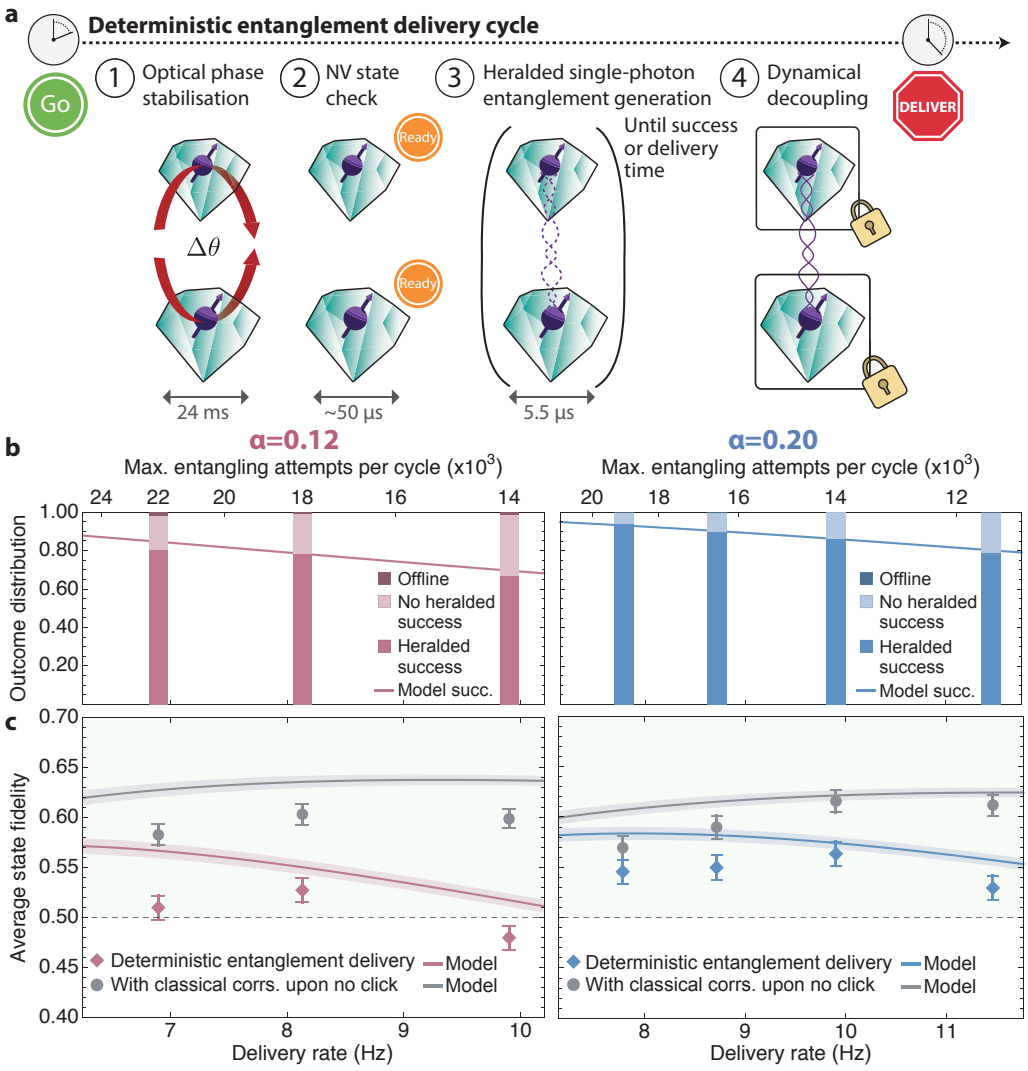

FIG. 4. **Deterministic entanglement delivery.** **a**, Each entanglement cycle combines: ① Optical phase stabilisation. ② NV state checks, repeated until a threshold number of photons are detected at each node. ③ Attempts at probabilistic entanglement generation (Fig. 2). ④ Upon heralded entanglement success, the state is protected by dynamical decoupling until the delivery time. **b**, Distribution of deterministic entanglement delivery outcomes for $\alpha = 0.12$ & 0.2 and different delivery rates. Shown is the fraction of cycles in which a herald photon is detected (heralded success), in which no herald is detected (no heralded success), and in which the NV state checks for at least one of the NV centres fail repeatedly for the whole cycle (offline, often too small to be visible in the plot). The lines give the success rates predicted by our model. **c**, Average fidelity of deterministically-delivered entangled states for $\alpha = 0.12$ & 0.2 and different delivery rates (diamonds). Also plotted is the average fidelity if classically-correlated states were delivered for cycles in which no success event is heralded (circles). The associated lines (with shaded 1 s.d. statistical uncertainties) plot the corresponding predictions of our model (discussed further in methods). For all plots, the data error bars represent 1 s.d.

We first confirm that heralded entanglement occurs with the expected probabilities (Fig. 4b) by determining the fraction of cycles in which entanglement is heralded, in which no entangling attempts succeed, and in which entanglement attempts do not occur at all as the NV state check never succeeds. In order to establish reliable and useful quantum networks, it is important that entangled states can be delivered with high confidence over long periods. The nodes must therefore not be offline due, for example, to uncompensated drifts in the resonant frequencies of the optical transitions. We therefore do not stop the experiment from running once it starts and include any such offline cycles in our datasets. Their negligible contribution (0.8% of cycles) confirms the high robustness of our experimental platform and the effectiveness of our NV frequency and charge-state control (discussed further in methods).

For each value of α and for each pre-set delivery interval, we determine the average fidelity of the deterministically delivered states by measuring their $\langle XX \rangle$, $\langle YY \rangle$ and $\langle ZZ \rangle$ correlations (Fig. 4c). We find that for $\alpha = 0.2$ and a rate of 9.9 Hz, we are able to create states with a fidelity of 0.56(1), proving successful deterministic entanglement delivery.

Our model (solid lines in Fig. 4c) effectively captures the trends of the deterministic entanglement delivery data. However, the observed state fidelities are slightly lower than the predicted ones, hinting at sources of decoherence that are not included in our model (discussed further in methods and Extended Data Fig. 4). Identifying these potential sources will be the subject of future work.

During cycles in which entanglement is not successfully heralded, the spin states are nonetheless delivered and read out. In these case, we deliver the state that the NVs are left in after a failed entanglement attempt, which has a low fidelity with the desired Bell state (e.g. $F_{\text{unent}} = 0.04$ for $\alpha = 0.2$). While this stringent test highlights the robust nature of our protocol, we could instead deliver a mixed state ($F_{\text{unent}} = \frac{1}{4}$) or a classically-correlated state ($F_{\text{unent}} = \frac{1}{2}$) when a successful event is not heralded. The resulting fidelities for our experimental data if classically-correlated states were delivered are also plotted in Fig. 4b (grey circles). In this case we would be able to deliver entangled states deterministically with fidelities of 0.62(1) at a rate of 9.9 Hz.

The deterministic entanglement delivery between remote NV centres demonstrated here is enabled by a quantum link efficiency exceeding unity. Straightforward modifications to our experiment are expected to further increase our quantum link efficiency. Refinements to the classical experimental control will allow us to reduce the entanglement attempt duration from 5.5 μs to below 2 μs, which would more than double the entangling rate. Furthermore, the entangled state coherence time could be significantly improved by exploiting long-lived nuclear spin quantum memories [10, 30, 31]. We anticipate that this will allow for link efficiencies in excess of 100 in the near term. Further improvements to the photon detection efficiency (including enhancement of zero-phonon line emission) [32, 33] would lead to an additional increase of at least an order of magnitude.

In combination with recent progress on robust storage of quantum states during remote entangling operations [10, 34], the techniques reported here reveal a direct path to the creation of many-body quantum states distributed over multiple quantum network nodes. Moreover, given the demonstrated potential for phase stabilisation in optical fibre over tens of kilometre distances [22], our results open up the prospect of entanglement-based quantum

networks at metropolitan scales.

[1] Kimble, H. J. The quantum internet. *Nature* **453**, 1023–1030 (2008).

[2] Broadbent, A., Fitzsimons, J. & Kashefi, E. Universal Blind Quantum Computation. In *50th Annual IEEE Symposium on Foundations of Computer Science* (2009).

[3] Jiang, L. *et al.* Quantum repeater with encoding. *Phys. Rev. A* **79**, 032325 (2009).

[4] Ekert, A. & Renner, R. The ultimate physical limits of privacy. *Nature* **507**, 443–447 (2014).

[5] Gottesman, D., Jennewein, T. & Croke, S. Longer-Baseline Telescopes Using Quantum Repeaters. *Phys. Rev. Lett.* **109**, 070503 (2012).

[6] Nickerson, N. H., Fitzsimons, J. F. & Benjamin, S. C. Freely Scalable Quantum Technologies Using Cells of 5-to-50 Qubits with Very Lossy and Noisy Photonic Links. *Phys. Rev. X* **4**, 041041 (2014).

[7] Kómár, P. *et al.* A quantum network of clocks. *Nat. Phys.* **10**, 582–587 (2014).

[8] Hucul, D. *et al.* Modular entanglement of atomic qubits using photons and phonons. *Nat. Phys.* **11**, 37–42 (2015).

[9] Hensen, B. *et al.* Loophole-free Bell inequality violation using electron spins separated by 1.3 kilometres. *Nature* **526**, 682–686 (2015).

[10] Kalb, N. *et al.* Entanglement distillation between solid-state quantum network nodes. *Science* **356**, 928–932 (2017).

[11] Reiserer, A. & Rempe, G. Cavity-based quantum networks with single atoms and optical photons. *Rev. Mod. Phys.* **87**, 1379–1418 (2015).

[12] Hofmann, J. *et al.* Heralded entanglement between widely separated atoms. *Science* **337**, 72–75 (2012).

[13] Northup, T. & Blatt, R. Quantum information transfer using photons. *Nat. Photonics* **8** (2014).

[14] Pfaff, W. *et al.* Unconditional quantum teleportation between distant solid-state quantum bits. *Science* **345**, 532–535 (2014).

[15] Monroe, C. *et al.* Large-scale modular quantum-computer architecture with atomic memory and photonic interconnects. *Phys. Rev. A* **89**, 1–16 (2014).

[16] Pant, M. *et al.* Routing entanglement in the quantum internet ArXiv 1708.07142 (2017).

[17] Schoute, E., Mancinska, L., Islam, T., Kerenidis, I. & Wehner, S. Shortcuts to quantum network routing ArXiv 1610.05238 (2016).

[18] Stockill, R. *et al.* Phase-Tuned Entangled State Generation between Distant Spin Qubits. *Phys. Rev. Lett.* **119**, 010503 (2017).

[19] Delteil, A. *et al.* Generation of heralded entanglement between distant hole spins. *Nat. Phys.* **12**, 218–223 (2016).

[20] Bar-Gill, N., Pham, L. M., Jarmola, A., Budker, D. & Walsworth, R. L. Solid-state electronic spin coherence time approaching one second. *Nat. Commun.* **4**, 1743 (2013).

[21] Cabrillo, C., Cirac, J. I., García-Fernández, P. & Zoller, P. Creation of entangled states of distant atoms by interference. *Phys. Rev. A* **59**, 1025–1033 (1999).

[22] Minář, J., de Riedmatten, H., Simon, C., Zbinden, H. & Gisin, N. Phase-noise measurements in long-fiber interferometers for quantum-repeater applications. *Phys. Rev. A* **77**, 052325 (2008).

[23] Casabone, B. *et al.* Heralded Entanglement of Two Ions in an Optical Cavity. *Phys. Rev. Lett.* **111**, 100505 (2013).

[24] Sipahigil, A. *et al.* An integrated diamond nanophotonics platform for quantum-optical networks. *Science* **354**, 847–850 (2016).

[25] Robledo, L. *et al.* High-fidelity projective read-out of a solid-state spin quantum register. *Nature* **477**, 574–578 (2011).

[26] Fuchs, G. D., Dobrovitski, V. V., Toyli, D. M., Heremans, F. J. & Awschalom, D. D. Gigahertz dynamics of a strongly driven single quantum spin. *Science* **326**, 1520–1522 (2009).

[27] Abobeih, M. H. *et al.* One-second coherence for a single electron spin coupled to a multi-qubit nuclear-spin environment ArXiv 1801.01196 (2018).

[28] Chou, C.-W. *et al.* Functional quantum nodes for entanglement distribution over scalable quantum networks. *Science* **316**, 1316–1320 (2007).

[29] Matsukevich, D. N. *et al.* Deterministic single photons via conditional quantum evolution. *Phys. Rev. Lett.* **97**, 013601 (2006).

[30] Maurer, P. C. *et al.* Room-Temperature Quantum Bit Memory Exceeding One Second. *Science* **336**, 1283–1286 (2012).

[31] Yang, S. *et al.* High-fidelity transfer and storage of photon states in a single nuclear spin. *Nature Photonics* **10**, 507 (2016).

[32] Riedel, D. *et al.* Deterministic enhancement of coherent photon generation from a nitrogen-vacancy center in ultrapure diamond. *Phys. Rev. X* **7**, 031040 (2017).

[33] Wan, N. H. *et al.* Efficient Extraction of Light from a Nitrogen-Vacancy Center in a Diamond Parabolic Reflector ArXiv 1711.01704 (2017).

[34] Reiserer, A. *et al.* Robust Quantum-Network Memory Using Decoherence-Protected Subspaces of Nuclear Spins. *Physical Review X* **6**, 021040 (2016).

ACKNOWLEDGEMENTS

We thank Suzanne van Dam, Mohamed Abobeih, Tim Taminiau, Filip Rozpędek, Kenneth Goodenough and Stephanie Wehner for helpful discussions. We acknowledge support from the Netherlands Organisation for Scientific Research (NWO) through a VICI grant and the European Research Council through a Starting Grant and a Synergy Grant.

AUTHOR CONTRIBUTIONS

P.C.H., N.K. and J.P.J.M. prepared the experimental apparatus. P.C.H. and N.K. carried out the experiments. R.F.L.V. and R.N.S. conceived the microwave switch circuit. P.C.H. analysed the data and wrote the manuscript with input from N.K. and R.H. The diamond substrates were grown by D.J.T. and M.M. The project was supervised by R.H.

AUTHOR INFORMATION

Correspondence should be addressed to R.H. (r.hanson@tudelft.nl). The authors declare no competing interests.

DATA AVAILABILITY

The datasets generated and analysed during the current study are available from the corresponding author on reasonable request.

METHODS

Derivation of deterministically delivered entangled state fidelity as a function of quantum link efficiency

We assume an entanglement generation rate r_{ent} and an entangled state decoherence rate r_{dec}. If the rate at which entanglement attempts occur is much faster than r_{ent} (i.e there is a low probability of success), we can approximate entanglement generation as a continuous process. In this case, the probability density for successfully generating entanglement at a time t after beginning our attempts is given by $p_{\text{ent}}(t) = r_{\text{ent}}e^{-r_{\text{ent}}t}$. The corresponding cumulative probability of success is $p_{\text{succ}}(t) = 1 - e^{-r_{\text{ent}}t}$.

Once we succeed at creating entanglement, the state will decohere until the time at which we deliver it. For single-qubit depolarising noise at each site, the fidelity of the resulting state after storage for a time t is given by

$$F(t) = \frac{1}{4} + \frac{3}{4}e^{-r_{\text{dec}}t}. \tag{7}$$

If we deliver our entangled state at time $t_{\text{ent}} = \beta/r_{\text{dec}}$ (where β is simply used to parameterise the time in terms of the decoherence rate), the average fidelity of the delivered state (given a success occurred) is therefore

$$\begin{aligned} F_{\text{succ}} &= \frac{1}{p_{\text{succ}}(t_{\text{ent}})}\int_0^{t_{\text{ent}}} p_{\text{ent}}(t)F(t_{\text{ent}}-t)\mathrm{d}t \\ &= \frac{1}{p_{\text{succ}}(t_{\text{ent}})}\int_0^{t_{\text{ent}}} r_{\text{ent}}e^{-r_{\text{ent}}t}(\frac{1}{4} + \frac{3}{4}e^{-r_{\text{dec}}(t_{\text{ent}}-t)})\mathrm{d}t. \\ &= \frac{3e^{-\beta}\eta_{\text{link}} + (1-4\eta_{\text{link}})e^{-\eta_{\text{link}}\beta} + \eta_{\text{link}} - 1}{4(\eta_{\text{link}}-1)p_{\text{succ}}(t_{\text{ent}})} \end{aligned} \tag{8}$$

We note that $p_{\text{succ}}(t_{\text{ent}}) = 1 - e^{-\beta\eta_{\text{link}}}$, and therefore $\beta = -\ln(1 - p_{\text{succ}}(t_{\text{ent}}))/\eta_{\text{link}}$. Using this, along with the shorthand $p_{\text{succ}} = p_{\text{succ}}(t_{\text{ent}})$, we find that

$$F_{\text{succ}} = \frac{3\eta_{\text{link}} + p_{\text{succ}} - 3\eta_{\text{link}}(1-p_{\text{succ}})^{1/\eta_{\text{link}}} - 4\eta_{\text{link}}p_{\text{succ}}}{4p_{\text{succ}}(1-\eta_{\text{link}})}. \tag{9}$$

As discussed in the main text, we can choose to draw a black box around this process, delivering an unentangled state (state fidelity $F_{\text{unent}} \leq \frac{1}{2}$) for cycles in which no attempt at entanglement generation succeeds such that a state is always delivered. Reproducing Eq. 1 of the main text, this means that the states output from this black-box will have a fidelity with a Bell state of

$$F_{\text{det}} = p_{\text{succ}}F_{\text{succ}} + (1-p_{\text{succ}})F_{\text{unent}}, \tag{10}$$

where F_{succ} is as given above. The maximum achievable fidelity when outputting a fully mixed state ($F_{\text{unent}} = \frac{1}{4}$) upon failure is found by optimising F_{succ} for a given quantum link efficiency η_{link}:

$$F_{\text{det}}^{\text{max}} = \frac{1}{4}\left(1 + 3(\eta_{\text{link}})^{\frac{1}{1-\eta_{\text{link}}}}\right). \tag{11}$$

Note that the full state of a quantum system can only be determined experimentally using an ensemble of identical states. This means that, in the absence of information about which deterministic entanglement delivery cycles have a heralded success, the only accurate description of the output of such a black-box system is that a statistical mixture is deterministically output at each cycle.

Experiment design

We use chemical-vapour-deposition homoepitaxially grown diamonds of type IIa with a natural abundance of carbon isotopes. Both diamonds have been cut along the $\langle 111 \rangle$ crystal axis and were grown by Element Six. They are situated in home-built confocal microscope set-ups within closed-cycle cryostats (4 K, Montana Instruments) separated by two meters. We use fast microwave switches to shield both NV centres from microwave amplifier noise and therefore increase the coherence times dramatically (Node A uses Qorvo TGS2355-SM and Node B uses Analog Devices HMC544). All other parts of the set-up and sample details have been described in the Supplementary Information of Refs. [9, 10].

One cycle of the deterministic entanglement protocol consists of optical phase stabilisation (described further below), charge-resonance checks to ensure that both NVs are in the appropriate charge state and on-resonance [25], heralded single-photon entanglement generation, and finally dynamical decoupling to protect the state of the NVs from their environment until the delivery time. The experimental sequences used in each step of this protocol (and also the single-photon entanglement generation experiment) are detailed in Extended Data Fig. 1.

After delivery, the state of each NV is measured in a chosen basis. We use spin-selective optical read-out of the NV electron spin in a single shot via the optical E_x transition on both nodes [25]. We measure single-shot read-out fidelities of 0.959(3) (0.950(3)) for the bright $|m_s = 0\rangle \equiv |\uparrow\rangle$ ground-state and 0.995(1) (0.996(1)) for the dark $|m_s = -1\rangle \equiv |\downarrow\rangle$ state on Node A (Node B). These values are subsequently used to correct for read-out errors of the electron spins in state tomography measurements.

Experiment control and communication logic

Extended Data Fig. 2 gives the decision trees and control logic for the ADwin microprocessors (Jaeger ADwin Pro II) that control the experiments. These microcontrollers are responsible for controlling all other experimental hardware and also communicate with each other to synchronise the experiment.

Herald photon detection window

We use a combination of polarization and temporal filtering to separate the excitation pulse from photons emitted by the NV. This necessitates a compromise between collecting as

much of the emission light as possible, while ensuring that contamination from the pulse is minimised. In this experiment, we choose a temporal filter window (Supp. Fig. 3) such that the pulse (assumed to have a Gaussian profile) is suppressed to the level of the detector dark counts by the beginning of the window. The end of the window at ~ 30 ns after the pulse is chosen so that, for all of the data sets taken, the rate of detected NV photons is greater than ten times the dark count rate at all points within the window. We use a complex programmable logic device (CPLD) to apply this temporal filtering during our experiment and herald the successful generation of an entangled state in real-time.

Theoretical model of deterministic entanglement delivery

We developed a detailed model to determine the expected performance of the deterministic entanglement delivery experiment, based on the independently measured parameters given in Extended Data Table 1.

Once the set-ups are determined to be ready, the core entanglement sequence begins with single-photon entanglement generation. This proceeds by first initialising each node in $|\uparrow\rangle$, followed by a coherent rotation using a microwave pulse to create the state

$$|NV\rangle = \sqrt{\alpha}\,|\uparrow\rangle + \sqrt{1-\alpha}\,|\downarrow\rangle\,. \tag{12}$$

Resonant excitation of the NV nodes excites only the 'bright' $|\uparrow\rangle$ level to an excited state, which rapidly decays radiatively back to the ground state by emitting a single photon. This entangles the state of the NV with the presence $|1\rangle$ or absence $|0\rangle$ of a photon in the emitted optical mode:

$$|NV, \text{optical mode}\rangle = \sqrt{\alpha}\,|\uparrow\rangle\,|1\rangle + \sqrt{1-\alpha}\,|\downarrow\rangle\,|0\rangle\,. \tag{13}$$

The photons emitted by each NV are transmitted to a central station at which a beamsplitter is used to remove their which-path information. Successful detection of a photon at this station indicates that at least one of the NVs is in the bright state $|\uparrow\rangle$ and therefore heralds the creation of a spin-spin entangled state. This entangled state, expressed as $|NV_{\text{Node A}}, NV_{\text{Node B}}\rangle$, is given (in unnormalised form) by

$$\rho_{\text{sc}} = \left|\psi^{\pm}\right\rangle\left\langle\psi^{\pm}\right| + p_{\uparrow\uparrow}\,|\uparrow\uparrow\rangle\langle\uparrow\uparrow| + p_{\downarrow\downarrow}\,|\downarrow\downarrow\rangle\langle\downarrow\downarrow|\,, \tag{14}$$

where

$$\left|\psi^{\pm}\right\rangle\left\langle\psi^{\pm}\right| = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & p_{\uparrow\downarrow} & \pm\sqrt{V\,p_{\uparrow\downarrow}\,p_{\downarrow\uparrow}} & 0 \\ 0 & \pm\sqrt{V\,p_{\uparrow\downarrow}\,p_{\downarrow\uparrow}} & p_{\downarrow\uparrow} & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}. \tag{15}$$

This state is parametrized by

$$\begin{aligned} p_{\uparrow\uparrow} &= \alpha^2\left((1-p_{\text{dc}})^2(p^A_{\text{det}}(1-p^B_{\text{det}}) + p^B_{\text{det}}(1-p^A_{\text{det}})\right) \\ &\quad + 2(1-p_{\text{dc}})\,p_{\text{dc}}(1-p^A_{\text{det}})(1-p^B_{\text{det}})\big) \\ p_{\uparrow\downarrow} &= \alpha\,(1-\alpha)\left((1-p_{\text{dc}})^2\,p^A_{\text{det}} + 2\,p_{\text{dc}}(1-p_{\text{dc}})(1-p^A_{\text{det}})\right) \\ p_{\downarrow\uparrow} &= \alpha\,(1-\alpha)\left((1-p_{\text{dc}})^2\,p^B_{\text{det}} + 2\,p_{\text{dc}}(1-p_{\text{dc}})(1-p^B_{\text{det}})\right) \\ p_{\downarrow\downarrow} &= 2\,(1-\alpha)^2\,p_{\text{dc}}\,(1-p_{\text{dc}}) \end{aligned} \tag{16}$$

where V gives the visibility of two-photon interference, p_{dc} gives the dark count probability per detector (given by the product of the dark count rate ν_{dark} and the 25 ns detection window length), and p^A_{det} (p^B_{det}) gives the probability of detecting a photon emitted by Node A (B). In the limit of $p_{\mathrm{det}} \ll 1$, for balanced detection probabilities $p_{\mathrm{det}} = p^A_{\mathrm{det}} = p^B_{\mathrm{det}}$ and no other imperfections, this tends to the expression given in the main text: $\rho_{NV,NV} = (1-\alpha)\,|\psi\rangle\langle\psi| + \alpha\,|\uparrow\uparrow\rangle\langle\uparrow\uparrow|$.

The corresponding probability of successfully heralding entanglement is given by

$$p_{\mathrm{herald}} = (1-p_{\mathrm{dc}})(\alpha(p^A_{\mathrm{det}} + p^B_{\mathrm{det}} - 2p^A_{\mathrm{det}}p^B_{\mathrm{det}}\alpha) + p_{\mathrm{dc}}(2 - 3(p^A_{\mathrm{det}} + p^B_{\mathrm{det}})\alpha + 4p^A_{\mathrm{det}}p^B_{\mathrm{det}}\alpha^2)) \quad (17)$$

The modelled success rate (plotted in main text Fig. 2e) is calculated by dividing the entangling attempt duration (5.5 μs) by p_{herald}.

We model double excitation (discussed further below) by applying a Pauli Z transformation to each of the NV states with probability $p_{2\mathrm{ph}}/2$. Phase instability is modelled similarly by applying a Pauli Z transformation to one of the states with probability $\frac{1}{2}(1 - e^{-\frac{1}{2}((\nu_{\mathrm{Int}}t_p)^2+\sigma^2_{\mathrm{Int}})})$, where t_p denotes the time elapsed since phase stabilisation.

Finally, we model the impact of dynamical decoupling by assuming that it acts as a depolarising channel for each qubit [27]. We therefore apply single-qubit depolarising errors with a probability determined by the measured dynamical-decoupling coherence times. For decoupling for a total time duration of t_d, the total probability of a depolarising error (i.e. the application of a Pauli X, Y or Z transformation with an equal probability) is given by $\frac{3}{4}(1 - e^{-t_d/T_2})$.

This model, based only on independently determined parameters (Extended Data Table 1), effectively captures the trends of our deterministic entanglement generation data (main text Fig. 4). However, we find that its predictions are slightly offset from these experimental measurements, suggesting that it does not include a small source of infidelity that is present in the experimental data. One potential origin of this discrepancy could be the up to two orders-of-magnitude more attempts at generating entanglement after NV-state verification made here as compared to previous experiments [9, 10]. Any additional sources of infidelity that may occur over this period (for example, due to the passive charge-state stabilisation process, discussed further below) are not included in the model. A detailed study of these potential imperfections is outside the scope of this present work. Nonetheless, as an estimate of the order of this effect, we find that a small systematic correction of 3% to the heralded entangled state fidelity is sufficient to effectively match our model to the data, as shown in Extended Data Fig. 4.

Passive charge-state stabilization of individual NV centres

The negatively charged NV centre (NV^-) can be ionized under optical illumination via a two-photon absorption process [35]. Due to the different level structure of the neutral charge state NV^0, the NV will remain dark if such an ionization event occurs during one of our entangling attempts. Ionization thus hampers the performance of our deterministic entangling protocol by diminishing the success rate and delivery of a separable state upon success. Previous experiments with NV centres that worked in the regime of probabilistically generated, yet heralded, remote entanglement overcame NV-ionization by frequent charge-state verification between protocols and actively converting the NV centre back to NV^- by interleaved resonant excitation of the optical transitions of NV^0 [36].

Such active stabilization protocols would require additional logical overhead in our scenario where entanglement is generated deterministically. Instead we passively stabilize the charge-state during our entangling sequence by shining in an additional weak laser beam that is resonant with the optical transition of NV^0 (Extended Data Fig. 5). This provides negligible disturbance to the spin initialization fidelity of NV^- while bringing the NV centre back into NV^- if it was converted to NV^0. We additionally identify that the optical reset beam (duration 1.5 μs) is the main cause of ionization in our system and carefully balance the power of both beams such that the spin state is still well initialized and that ionization is a negligible process for our deterministic entangling protocol (up to 15000 entangling attempts). Note that reducing the applied power further by elongating the spin reset duration would decrease the entanglement rate and limit our quantum link efficiency.

Extended Data Fig. 5 depicts the basic element that, in repetition, forms our sequence to probe the ionization rate. We use simultaneous charge and spin reset beams followed by a single microwave π rotation that brings the NV into $|1\rangle$ and thus guarantees optical excitation during the next round. The NV is then readout after a final optical reinitialisation into the bright state $|0\rangle$. By increasing the number of sequence repetitions, we observe a decay of the final readout fidelity that is associated with the ionization rate. By increasing the optical intensity of the charge-state reset beam we obtain a negligible decay as a function of sequence repetitions, therefore allowing us to overcome ionization in our deterministic entangling protocol. Note that the illumination strength of the charge-reset beam is weak enough to avoid inducing noticeable spectral diffusion of the NV emission as our measured entangled states are consistent with a high degree of indistinguishability for both NV emission profiles (discussed further below).

Optical phase stabilisation

The single-photon entanglement experiment requires that optical phase of an effective interferometer between the two nodes is known, as shown in Fig. 2 of the main text. The optical phase difference between the paths of this interferometer must be known in order to ensure that entangled states are available for further use. This is achieved by interleaving periods of optical-phase stabilisation with our entanglement generation.

For phase stabilisation we input bright laser light at the same frequency but orthogonally polarised to the light used for excitation of the NVs. The orthogonal polarisation is chosen because we use a crossed-polariser to filter out the excitation light from the NV emission. Using orthogonally polarised light for phase stabilisation therefore allows us to collect more light reflected from the diamond substrate. Before doing this, we verified that there is no measurable difference in the relative phase of the two polarisations within our interferometer.

Measurements of the phase drift (Extended Data Fig. 6a) show a slow drift on second time scales, but several strong resonances at hundreds of hertz (Extended Data Fig. 6b). These resonances are thought to be from mechanical elements in the path of the beam, including the microscope objective mount. As we were unable to completely suppress these resonances in the current set-ups, we need to measure the phase over a complete oscillation to estimate the mean phase reliably. The phase must therefore be measured for approximately 10 ms.

We calculate an estimate of the phase from the counts detected at the heralding single-photon detectors. This estimate is used to adjust the phase back to our desired value using a homebuilt piezoelectric fibre stretcher and a proportional-integral-derivative (PID) routine within our ADwin micro-controller. We find that it takes between 2-3 PID cycles to

optimally stabilise the phase. We stabilise the phase for 3 cycles during the single-photon entanglement experiment and for 2 cycles during the deterministic entanglement experiment. This difference is because phase stabilisation occurs during every cycle of the deterministic entanglement delivery experiment (~ 100 ms), while it only occurs every 180 ms during the single-photon entanglement experiment and therefore the phase drifts slightly less after one experimental cycle.

We achieve an average steady-state phase stability of $14.3(1)^\circ$, as measured by calibration routines spaced throughout the measurement of our data set (Extended Data Fig. 6c,d). This stability is limited by the previously identified mechanical oscillations of the optical elements in our experimental set-up. The standard deviation of the phase averaged over a 10 ms period during active stabilisation is $4.8(1)^\circ$.

We note that optical phase stabilisation is also likely to be feasible for long-distance network links. Using long-wavelength off-resonant light for phase measurements would allow for continuous stabilisation during entanglement attempts with negligible impact upon the NV state. An experimental study [22] has shown that two network nodes separated by 36 km over a commercial fibre network would still allow for interference visibilities of 99%. For longer distances, it would also be possible to passively track the phase at the time of entanglement delivery and feed this information back to the nodes in which the state is stored, requiring only a coherence time longer than the communication time.

Two-photon quantum interference

The quality of photon-mediated heralded entanglement between two emitters hinges on the indistinguishability of their emitted photons. We probe this indistinguishability by interfering emitted single photons on a beam splitter and measuring the number of events in which single-photon detectors connected to the output ports of the beamsplitter both detect a photon. For completely indistinguishable single photons, Hong-Ou-Mandel interference ensures that both photons always exit from the same port of the beamsplitter, and therefore no coincident events should be detected.

Our TPQI experiment proceeds by exciting each emitter with a series of well separated optical excitation pulses (separated by $1\,\mu$s). We collect statistics on coincidence events in which one detector registers a photon after one excitation pulse, and then the other detector registers a photon after a later excitation pulse. For an infinite pulse train, the number of coincidence events detected for each number of pulses between the detection events should be constant. However, for a finite pulse train, there are some pulses for a given pulse separation for which there is no partner excitation pulse and therefore no coincident events will be detected. This leads to a linearly decreasing number of coincidence events as a function of pulse difference (Extended Data Fig. 7a).

We use a linear fit to the coincidence events to infer the number of coincidences that would be detected from the same pulse (pulse difference of zero), if fully distinguishable single photons were input (Extended Data Fig. 7b). Because these are nonetheless single photons, a counting argument shows that, for balanced emission probabilities from each emitter, this expected number of events is given by half of the value of the linear fit at zero pulse difference.

The ratio r between the measured number of coincident events within the same pulse and the expected number of events for fully distinguishable photons is related to the single-photon wave function overlap $V = |\langle\psi_a|\,\psi_b\rangle|^2$ by $V = (1 - r)$ (again for balanced emission

probabilities from each emitter). Incorporating the effect of the known imbalance in emission probabilities in our experiment, we find $V = 0.90(2)$.

Dephasing of entangled states due to double excitation

An optical rabi pulse is used to excite the NV nodes to a higher lying level via a spin-conserving transition. The NV subsequently decays back down to its original level through spontaneous emission, thereby entangling the spin state of the NV and the emitted optical mode. For optical rabi pulses of finite duration, there is a chance that the NV will spontaneously emit a photon during the optical pulse and be re-excited before the end of the pulse. the first emitted photon will be lost to the environment, as it is impossible to distinguish from the excitation light. However, if the subsequent emitted photon is detected in this double excitation process, this will falsely herald entanglement. We measured the width of our optical pulse (Extended Data Fig. 8) and used a quantum-jump based simulation to calculate the corresponding double-excitation probability. Given that the NV emitted a photon within the detection window, the probability that double excitation occurred is $p_{2\mathrm{ph}} = 0.04$.

State storage via dynamical decoupling

The coherence time of NV centres is limited by interactions with other magnetic impurities. In our samples the dominant source of magnetic field noise is the surrounding bath of slowly fluctuating ^{13}C nuclear spins (natural abundance of 1.1 %) resulting in typical coherence times of 5 μs. We use dynamical decoupling XY8 sequences of the form $(t - \pi_X - 2t - \pi_Y - 2t - \pi_X - 2t - \pi_Y - 2t - \pi_Y - 2t - \pi_X - 2t - \pi_Y - 2t - \pi_X - t)^{N/8}$ to elongate the coherence times of both NV centres (see Fig. 3 main text), with microwave inversion pulses π, the waiting time t and the number of pulses N (see also Extended Data Fig. 1.4). A given decoupling duration is obtained by arbitrary combinations of t and N. We find the optimal combination for a targeted protection duration of ~ 100 ms by varying t for a fixed $N = 1024$. We specifically choose $N = 1024$ as the introduced infidelity from inversion pulse errors is still moderate for both nodes.

Extended Data Fig. 9 shows the results of our decoupling optimization procedure. We prepare the NV in a balanced superposition and choose waiting times that are integer multiples of the inverse ^{13}C-nuclear-spin Larmor frequency ν_L to avoid coupling with the nuclear spin bath (Node A: $\nu_L = 443.342$ kHz; Node B: $\nu_L = 442.442$ kHz). Following the techniques of Abobeih et al. [27], we further avoid coupling to other magnetic noise sources that result in loss of NV coherence by picking five waiting times with a total variation of 16 ns for each multiple of the inverse Larmor frequency. The data (grey) are then sorted for the waiting time with the best state preservation quality (blue) at each multiple, giving the minimal NV coherence decay for this number of inversion pulses. We then proceed to pick the waiting time that guarantees a low number of inversion pulses while still providing high-quality state

protection (red).

---

[35] Aslam, N., Waldherr, G., Neumann, P., Jelezko, F. & Wrachtrup, J. Photo-induced ionization dynamics of the nitrogen vacancy defect in diamond investigated by single-shot charge state detection. *New J. Phys.* **15**, 013064 (2013).

[36] Pfaff, W. *et al.* Unconditional quantum teleportation between distant solid-state quantum bits. *Science* **345**, 532–535 (2014).

[37] Minář, J., de Riedmatten, H., Simon, C., Zbinden, H. & Gisin, N. Phase-noise measurements in long-fiber interferometers for quantum-repeater applications. *Phys. Rev. A* **77**, 052325 (2008).

	Node A	Node B	Description
T_2 (ms)	290(20)	680(70)	Dephasing time of the electron spin state.
T_1 (s)	> 1	> 1	Relaxation time of electron spin eigenstates.
p_{det} (10^{-4})	2.8(1)	4.2(1)	Probability to detect a ZPL photon after a single excitation.
p_{ionize}	$\leq 10^{-6}$	$\leq 10^{-6}$	Probability of passive charge-state control failure per entangling attempt (detailed in methods).
t (μs)	40.320	36.148	Optimized inter-pulse delay for state storage.
F_0	0.959(3)	0.950(3)	Fidelity of the electron read-out for $\|\uparrow\rangle$.
$F_{\pm 1}$	0.995(1)	0.996(1)	Fidelity of the electron read-out for $\|\downarrow\rangle$.
V	0.90(2)		Visibility of the two-photon quantum interference (detailed in methods).
p_{2ph}	0.04		Estimated probability of double excitation during the optical π-pulse (detailed in methods).
ν_{dark} (Hz)	20		Dark count rate per detection channel.
σ_{Int}	14.3(1)°		Initial uncertainty of the interferometric drift (detailed in methods).
ν_{Int} ($/s$)	$\sim 20°$		Estimated drift rate of the free running interferometer (detailed in methods).

Extended Data Table 1. Independently measured experimental parameters for the performance of the nodes used in our experiment.

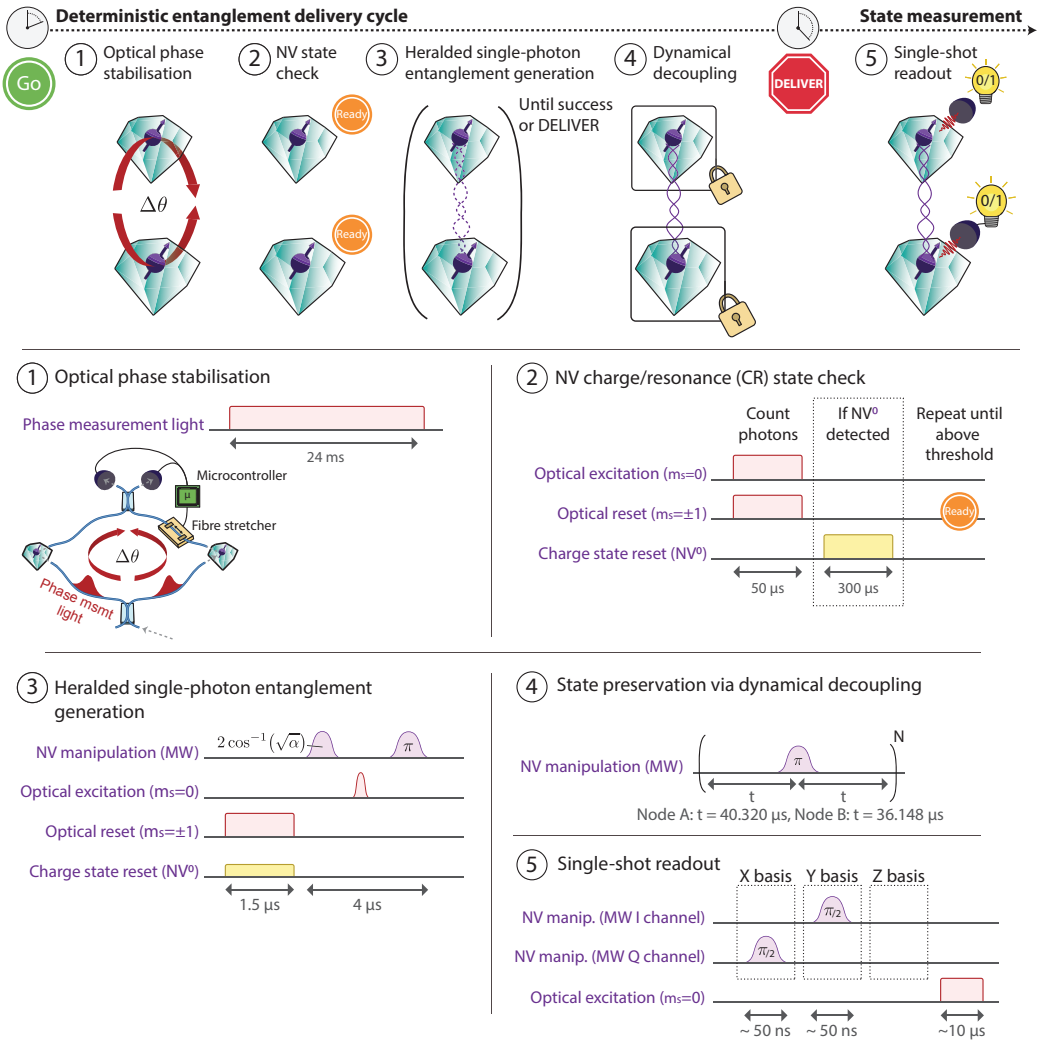

Extended Data Fig. 1. **Deterministic entanglement delivery sequences** Pulse sequences for each step of the deterministic entanglement delivery protocol. These sequences are also employed in the single-photon entanglement generation experiment. ① **Optical phase stabilisation**: Bright light is input to measure and stabilise the interferometer (see methods). Note that the duration is different for the single-photon entanglement experiment. ② **NV state check**: By shining in two lasers that are together resonant with transitions from all of the ground states, the NV will fluoresce regardless of its ground-state occupation. By counting photons emitted by the NV we are able to verify that both NVs are in the desired charge state NV^- and that they are on resonance with the applied lasers. The NV centre is deemed to be on resonance if the number of photons detected during the CR check surpasses a certain threshold. If no photons are detected, the NV is assumed to be in the NV^0 state and a resonant laser is applied to reset it to NV^-. ③ **Heralded single-photon entanglement generation:** Entanglement generation proceeds by optically repumping the spins to $|\uparrow\rangle$ (including passive charge-state stabilisation, see methods) before a microwave (MW) pulse is used to create the desired bright-state population α at each node. A resonant excitation pulse then generates spin-photon entanglement. A subsequent MW π pulse is used to ensure that the NV state is refocused before the next stage should success be heralded. ④ **Dynamical decoupling**: MW pulses are used to implement dynamical decoupling (see methods). ⑤ **Single-shot readout**: The NV nodes can be readout in arbitrary bases in a single shot. If required, a MW pulse is applied to rotate the qubit state before a resonant laser is applied. Fluorescence photons from the NV are detected if the NV is in the state $|\uparrow\rangle$.

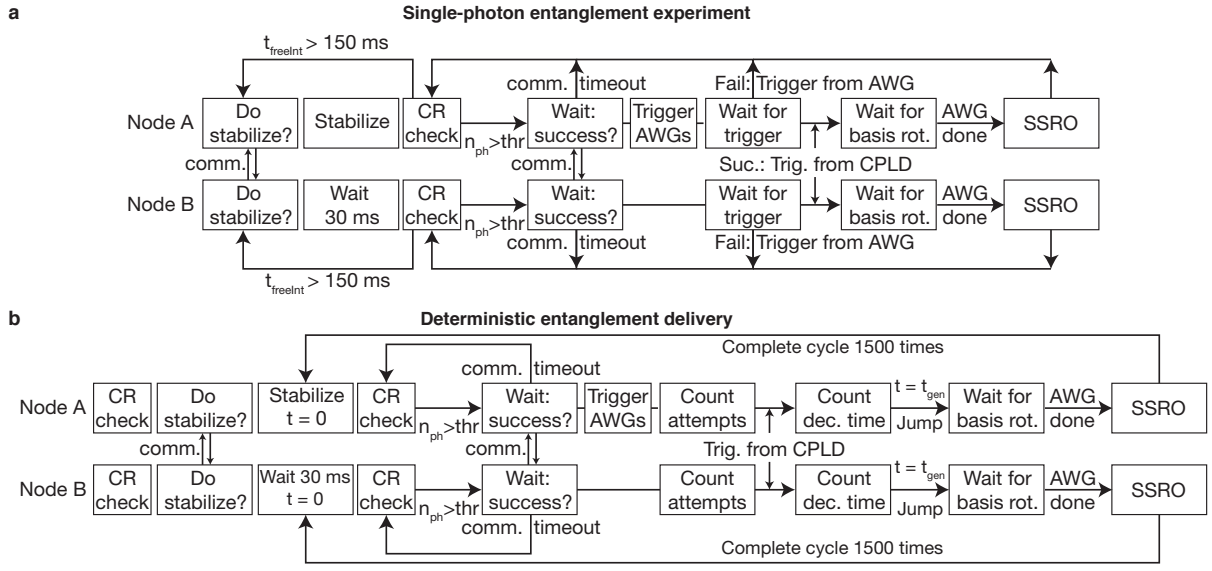

Extended Data Fig. 2. **Flowchart of the experimental sequences.** Shown are the decision trees of the ADwin microprocessors (Jaeger ADwin Pro II) that create the overarching measurement and control loops for network nodes A and B. Both nodes use arbitrary waveform generators for microwave and laser pulse sequencing (AWG, Tektronix AWG5014C). We additionally use a complex programmable logic device (CPLD) to herald the successful generation of an entangled state in real-time (described further in methods). **a**, Decision tree when benchmarking the entangled state. **b**, Deterministic entanglement delivery. Here the ADwins keep track of the elapsed time since the end of the phase stabilization ($t = 0$). **CR check**: As explained in Extended Data Fig. 1, the NV centre is deemed to be on resonance with the excitation lasers if the number of photons detected during the CR check surpasses a certain threshold n_{thr}. The CR check is repeated until this occurs. **comm. & comm. timeout:** Both ADwins exchange classical communication, such as success of the CR check, via a three-step-handshake. If one ADwin waits longer than 1 ms for a response from its counterpart the communication times out and we return to the previous logical step (see arrow). **Count attempts**: count the number of entangling attempts N until $N = N_{\mathrm{max}}$. **Count dec. time**: track the elapsed time since phase stabilization. If the elapsed time equals the prespecified state-generation time t_{gen} then trigger the AWG such that the local readout sequences are executed. **Wait for basis rot.**: ADwins wait for a trigger input from the AWG (AWG done) which heralds that the last MW-rotation before optical readout has been completed. **Trigger AWGs**: The ADwin of Node A triggers the AWGs of both nodes to initiate the microwave and entangling sequences. We use a single ADwin as trigger source to avoid timing jitter between both generated sequences. **SSRO**: Optical single-shot read-out. **Success: Trigger from CPLD / Fail: Trigger from AWG**: During entanglement generation, the CPLD communicates successful detection of a photon to the nodes. During the single-photon entanglement benchmarking experiment, the AWG at each node flags failure of the round after 250 entangling attempts. **Do stabilize?**: The ADwins communicate that phase stabilization will be the next step in the experimental sequence. The Node A ADwin then proceeds with the phase stabilization while the ADwin of Node B waits until the phase stabilisation has finished. The deterministic entangling sequence is run a total of 1500 times (500 times per read-out basis) before a new round is called in which starts again with the verification of resonant conditions for both NVs.

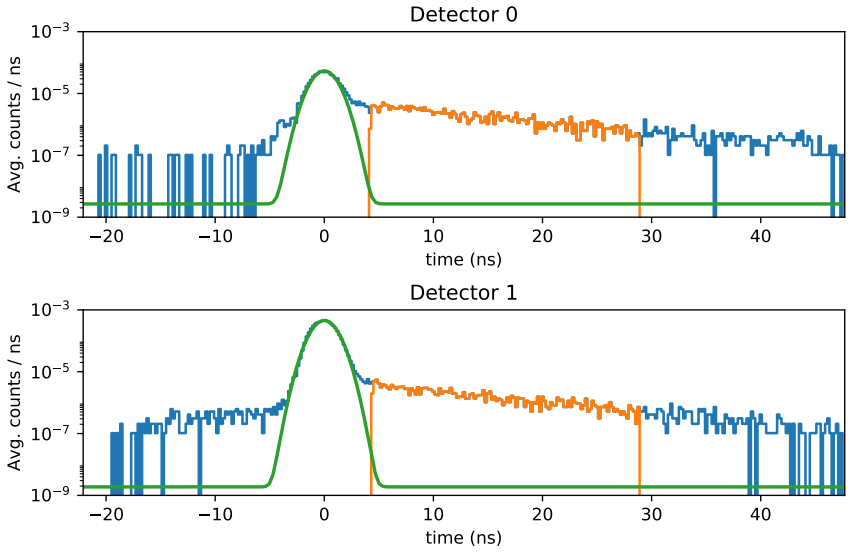

Extended Data Fig. 3. **Temporal filtering of photons.** Histogram of the times at which photons are detected at each single-photon detector (blue) during a deterministic entanglement delivery experiment with bright state population $\alpha = 0.12$. The orange histograms show the photons that were detected within the temporal filter window and therefore were counted as valid entanglement events. The green line shows a gaussian fit to the pulse with a FWHM of 2.26 as measured in Extended Data Fig. 8. This is used to estimate the contribution of residual pulse photons within the filter window.

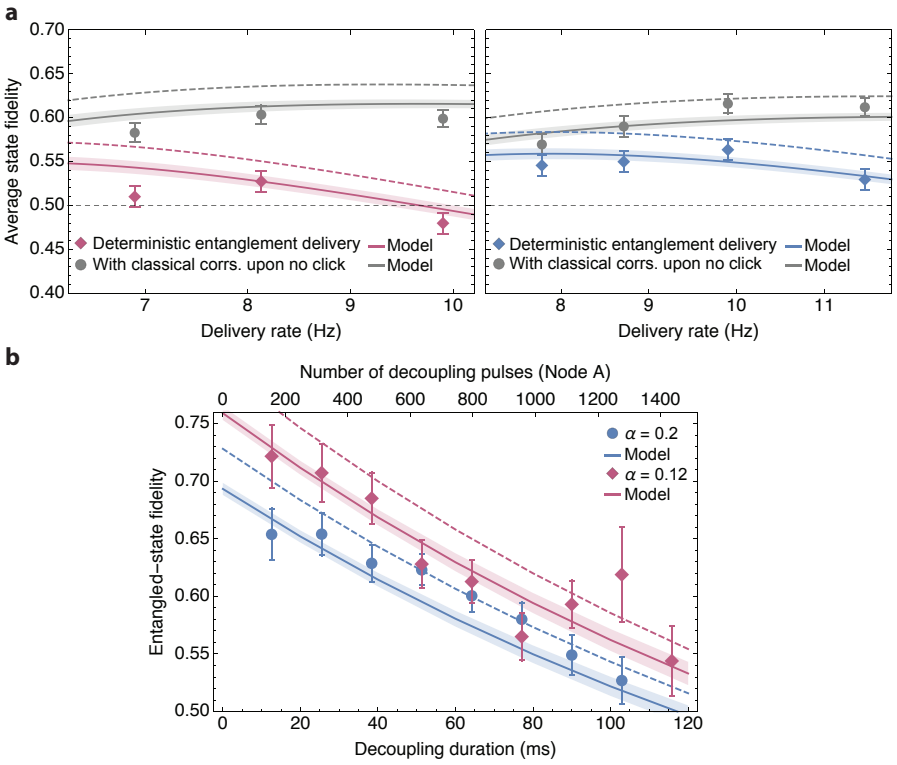

Extended Data Fig. 4. **Comparison of experimental model and data.** Including a 3% source of infidelity in our model (which otherwise consists completely of independently determined parameters) is sufficient to account for the offset observed between this model and our experimental data. **a**, The modified model, plotted with experimental data reproduced from main text Fig. 4. Dashed lines show the model given in the main text (without this infidelity parameter). **b**, This infidelity also applies to the model shown in main text Fig. 3, as an equally large number of entanglement repetitions was employed in generating this data. Error bars for data and shaded model uncertainties are 1 s.d.

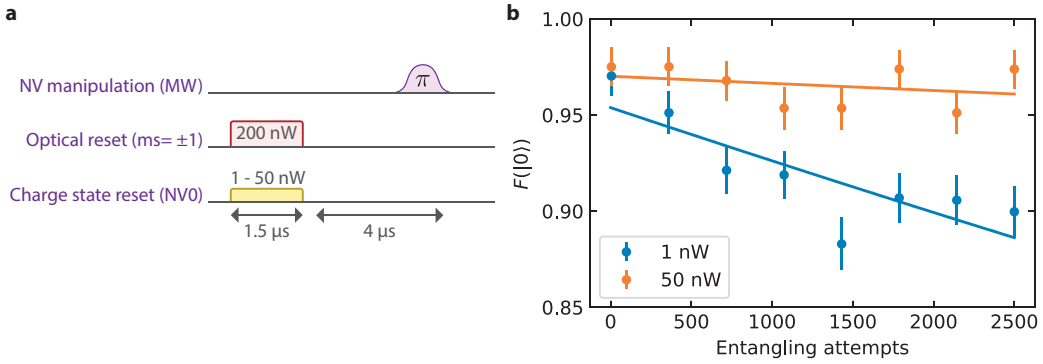

Extended Data Fig. 5. **Verifying passive charge-state stabilization into NV$^-$.** **a**, Elementary sequence to probe the NV ionization rate. **b**, Applying our sequence many times results in a decay of the NV readout fidelity due to ionization. By exploring the ionization rate for different charge-reset powers we find an optimal regime in which the spin initialization of NV^- is not affected by the additional blue-detuned beam and ionization is effectively mitigated over thousands of trials.

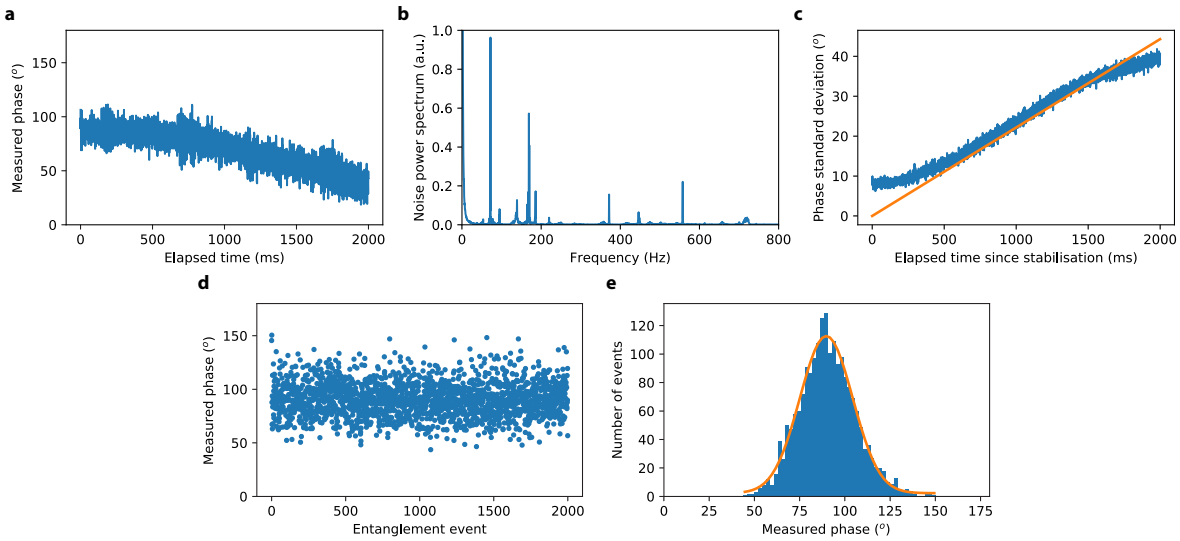

Extended Data Fig. 6. **Optical phase stabilisation.** Single-photon entanglement requires that optical phase of an effective interferometer between the two nodes is known. **a**, A typical trace of the interferometer optical phase as it is passively tracked for two seconds. **b**, Power spectrum of the optical phase signal showing peaks thought to be due to mechanical resonances of components in the setup. **c**, Active phase stabilisation is used to correct for phase drifts. Here the phase is stabilised and then the interferometer is allowed to passively drift for two seconds. Plotted is the standard deviation of the phase as a function of elapsed time for a data set of 100 of these measurements. The orange line shows a linear fit, used to estimate the rate of phase drift $\nu_{\mathrm{Int}} \sim 20°/s$. **d**, Here the phase is repeatedly actively stabilised every 180 ms. Entanglement generation occurs during the periods in between stabilisation. The interferometer phase is measured directly after each successful heralded entanglement event. **e**, Histogram of the measured post-entanglement optical phases. Also plotted in orange is a Gaussian fit with a standard deviation fixed to the average measured standard deviation for all entanglement data taken, $\sigma_{\mathrm{Int}} = 14.3(1)°$. Error bars for data are 1 s.d.

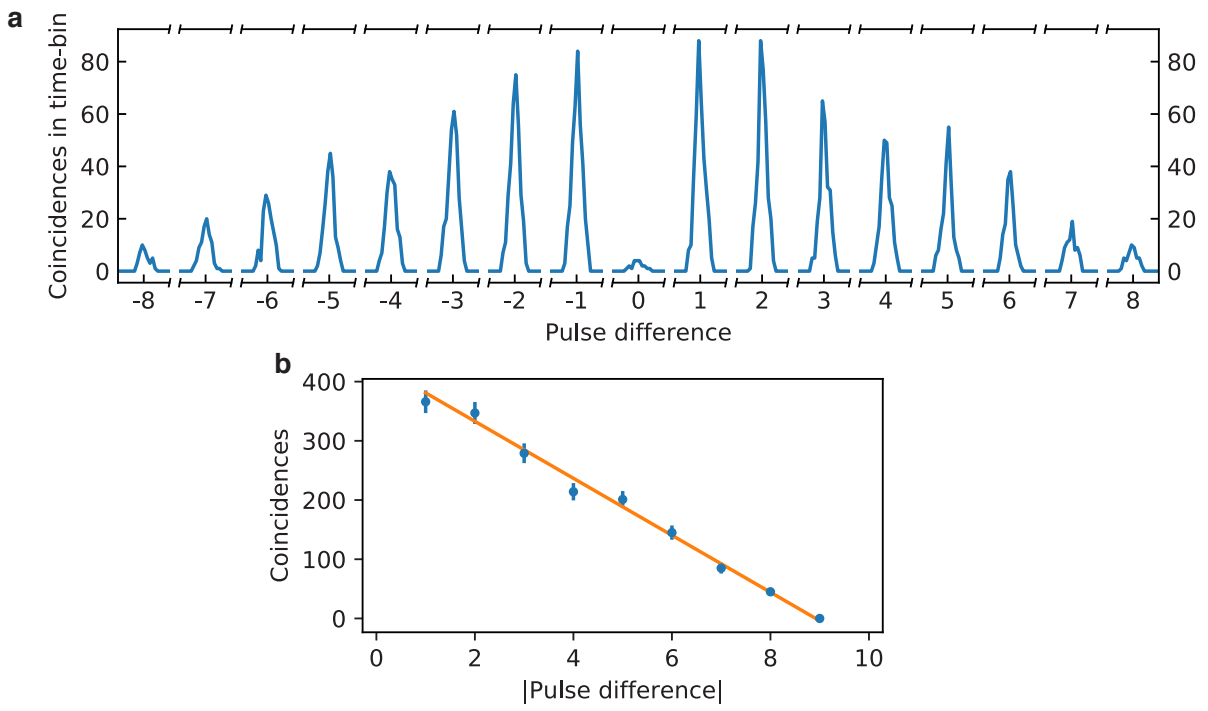

Extended Data Fig. 7. **Two-photon quantum interference. a**, Histogram for coincident events measured by two single-photon detectors in a two-photon quantum interference experiment, measured by cross-referencing photon detection events from a pulse train of 10 optical π-pulses that excite both emitters. Hong-Ou-Mandel interference of simultaneously coinciding photons ideally results in vanishing coincidence events within a single excitation round. The time difference between individual excitation rounds is 1 μs. Histograms of coincidence counts are shown with a bin-size of 5 ns. **b**, Total number of coincidences as a function of the number of pulses separating the two detection events. We extrapolate the measured coincidences to infer the expected coincidences for distinguishable photons at zero pulse difference by fitting a linear regression (orange). Using this to normalise the 22 observed coincidences for zero pulse difference allows us to estimate the TPQI visibility $V = 0.90(2)$. Error bars for data are 1 s.d.

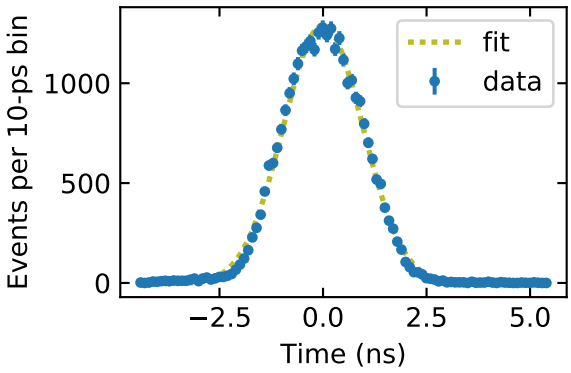

Extended Data Fig. 8. **Width of the optical π-pulse**. We find a full-width half maximum of 2.26 ns. This measurement is necessary to compute the dual-excitation probability shown in Extended Data Table 1, given a radiative lifetime of 12 ns. Error bars for data are 1 s.d.

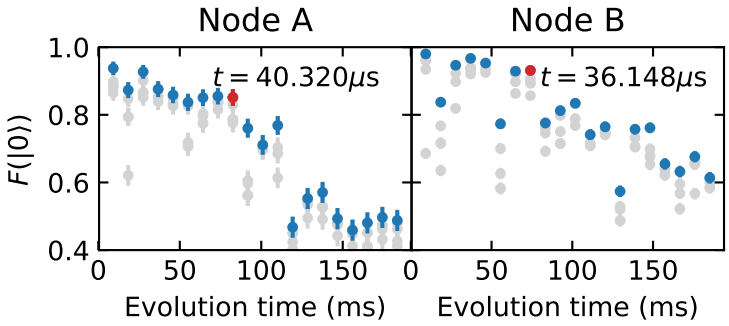

Extended Data Fig. 9. **Determining the optimal inter-pulse delay for state storage and 1024 inversion pulses.** We initialize a superposition state on the NV electron spin, preserve it via dynamical decoupling and finally perform optical readout after another $\pi/2$ pulse. We probe the coherence of the NV by varying the inter-pulse delay t in steps of the Larmor period $1/\nu_L \approx 2.25\,\mu$s and additionally shifting the delay in steps of 4 ns for a total of five data points per Larmor period (grey data). For each multiple of the Larmor period we pick the best, i.e. most preserving, inter-pulse delay (blue data). We determine the optimal delay t by selecting an inter-pulse delay that provides sufficient state preservation, i.e. $\sim$ 100 ms, for a moderate number of pulses (red data point and text inset in both panels). Left: Node A. Right: Node B. Error bars are 1 s.d.