TUDelft

Delft University of Technology

Multi-Node Quantum Networks with Diamond Qubits

Pompili, M.

Important note
To cite this publication, please use the final published version (if applicable).
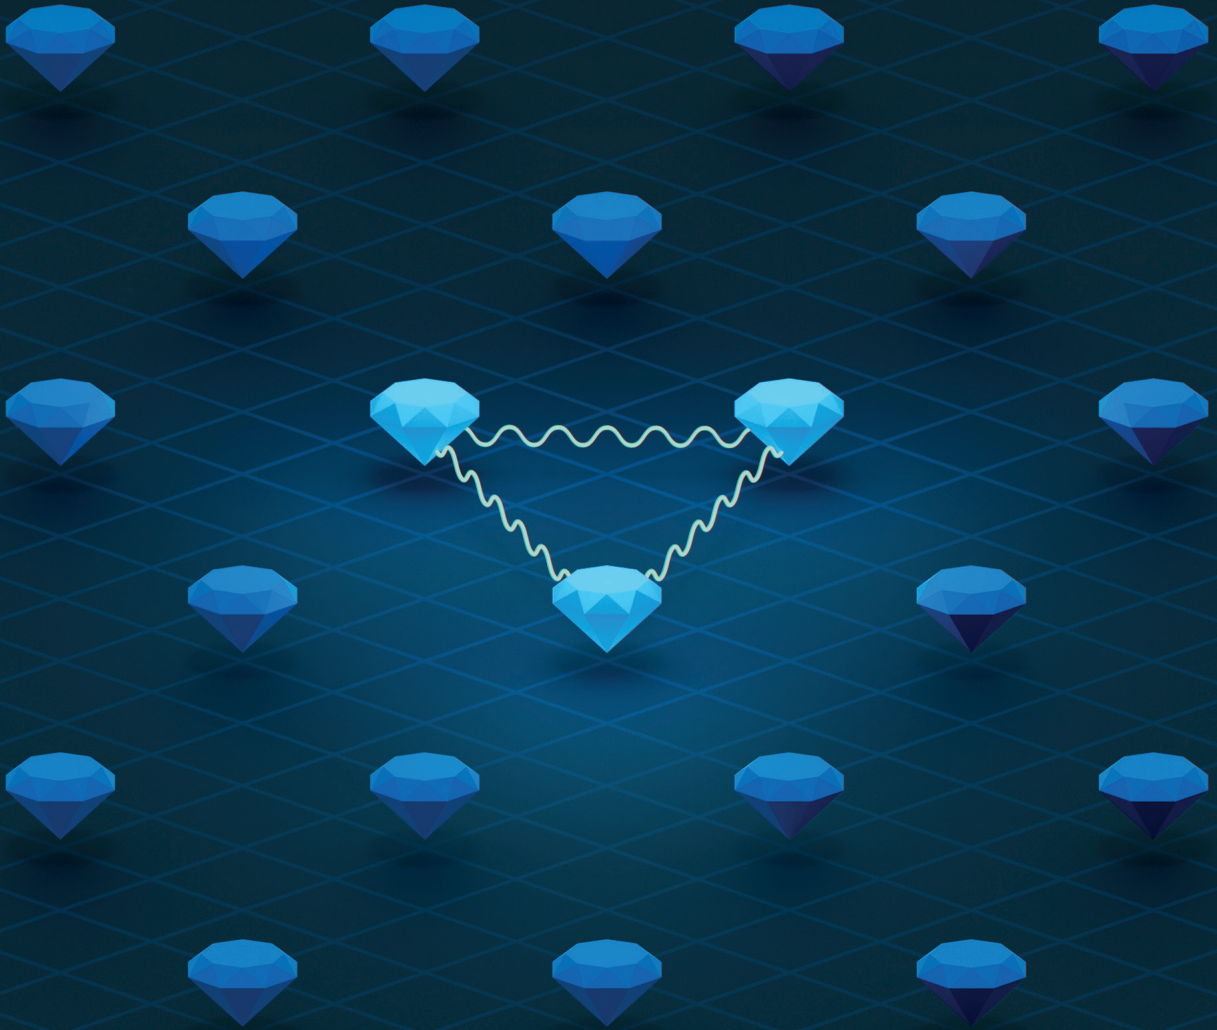Please check the document version above.

MULTI-NODE QUANTUM NETWORKS WITH DIAMOND QUBITS

MATTEO POMPILI

Multi-Node Quantum Networks with Diamond Qubits

Multi-Node Quantum Networks
with Diamond Qubits

Dissertation

for the purpose of obtaining the degree of doctor
at Delft University of Technology,
by the authority of the Rector Magnificus, prof. dr. ir. T. H. J. J. van der Hagen,
chair of the Board for Doctorates,
to be defended publicly on
Wednesday, 15 December 2021 at 12:30 o'clock

by

Matteo POMPILI

Master of Science in Physics,
La Sapienza Università di Roma, Italy,
born in Rome, Italy.

This dissertation has been approved by the promotors.

Composition of the doctoral committee:

Rector Magnificus,	chairperson
Prof. dr. ir. R. Hanson,	Delft University of Technology, promotor
Prof. dr. S. D. C. Wehner,	Delft University of Technology, promotor

Independent members:

Prof. dr. L. DiCarlo,	Delft University of Technology
Prof. dr. S. Gröblacher,	Delft University of Technology
Prof. dr. T. E. Northup,	University of Innsbruck, Austria
Prof. dr. A. S. Sørensen,	University of Copenhagen, Denmark
Prof. dr. R. Van Meter,	Keio University, Japan

Isn't it enough to see that a garden is beautiful
without having to believe that there are
fairies at the bottom of it too?

Douglas Adams

Contents

Summary

The Internet has revolutionized the way we live. It has enabled applications far beyond what it was originally built for, and it will continue to exceed our expectations for the future.

Quantum computers, and the network that will connect them—the Quantum Internet—are likely going to follow the same path. Differently from normal computers, quantum computers can share a property called *entanglement*, which allows the qubits (quantum bits) to be connected at a much more fundamental level. This property enables a range of new applications that span secure communication to enhanced metrology.

Over the past fifteen years, significant progress has been made in connecting rudimentary quantum network nodes via long-distance entanglement. Several quantum platforms have demonstrated entanglement generation between two physically separated qubits. In this thesis we take a significant step forward, both in terms of experimental complexity achieved, and in the abstraction of said complexity for future developments. Moving past two-node experiments required a fundamental redesign of our experimental apparatus, as well as developing the capabilities to control simultaneously the additional node.

The first result of this thesis is building a three-node entanglement-based quantum network. We demonstrated distribution of Greenberger-Horne-Zeilinger states over the network, as well as a building block for larger networks: entanglement swapping. Differently from previous multi-node demonstrations, which relied on post-selection, our network is able to perform the entanglement distribution in a *heralded* fashion: a signal will notify the users that the protocol was successful, and that the state is ready to be used.

The second result builds on the first, by adding control over a fifth qubit, improving the quality of the entanglement, and introducing a novel repetitive readout technique, to achieve quantum teleportation of a qubit from the third node to the first—nodes that do not share a direct entanglement channel.

The third and final result is the demonstration of entanglement delivery using a quantum network stack. The Internet is built using a plethora of physical platforms: optical fibers, Ethernet cables, Wi-Fi, satellite signals etc. To abstract their functionality, and make applications work regardless of the underlying platform, a layered approach was developed in the 1970s (the Internet protocol). Taking inspiration from classical network stacks, we demonstrate the first two layers of a quantum network stack, the *physical layer* (where the qubits, lasers and signal generators live), and the *link layer*, which abstracts the concepts of qubit and entanglement generation such that they can be used by applications at the higher-layers, hiding the complexity of the quantum platform being used.

Samenvatting

Het internet heeft een revolutie teweeggebracht in de manier waarop we leven. Het heeft toepassingen mogelijk gemaakt die veel verder gaan dan waarvoor het oorspronkelijk was gebouwd, en het zal onze verwachtingen voor de toekomst blijven overtreffen.

Quantumcomputers en het netwerk dat ze zal verbinden, het Quantum Internet, zullen waarschijnlijk hetzelfde pad volgen. Anders dan normale computers, kunnen quantumcomputers verstrengelde toestanden delen, waardoor de qubits (quantumbits) op een veel fundamenteler niveau met elkaar kunnen worden verbonden. Deze eigenschap maakt een reeks nieuwe toepassingen mogelijk, van veilige communicatie tot verbeterde metrologie.

In de afgelopen vijftien jaar is er aanzienlijke vooruitgang geboekt bij het verbinden van rudimentaire quantumnetwerkknooppunten via langeafstandsverstrengeling. Verschillende quantumplatforms hebben verstrengeling tussen twee fysiek gescheiden qubits aangetoond. In dit proefschrift zetten we een belangrijke stap voorwaarts, zowel in termen van de bereikte experimentele complexiteit, als in de abstractie van deze complexiteit voor toekomstige ontwikkelingen. Om verder te gaan dan experimenten met twee knooppunten, was een fundamenteel herontwerp van ons experimentele apparaat nodig, evenals het ontwikkelen van de mogelijkheden om tegelijkertijd het extra knooppunt te besturen.

Het eerste resultaat van dit proefschrift is het bouwen van een op verstrengeling gebaseerd quantumnetwerk bestaande uit drie knooppunten. We demonstreerden de distributie van Greenberger-Horne-Zeilinger-staten over het netwerk, evenals een bouwsteen voor grotere netwerken: verstrengeling doorgeven. In tegenstelling tot eerdere demonstraties met meerdere knooppunten, die gebaseerd waren op dataselectie achteraf, is ons netwerk in staat om de verstrengelingsdistributie uit te voeren met een successignaal, een signaal zal de gebruikers informeren dat het protocol geslaagd is en dat de toestand klaar is voor gebruik.

Het tweede resultaat bouwt voort op het eerste, door controle over een vijfde qubit toe te voegen, de kwaliteit van de verstrengeling te verbeteren en een nieuwe repetitieve uitleestechniek te introduceren, om quantumteleportatie van een qubit van het derde knooppunt naar het eerste knooppunt te bereiken, waarbij deze knooppunten geen directe verbinding met elkaar hebben.

Het derde en laatste resultaat is de demonstratie van het genereren van verstrengeling met behulp van een quantumnetwerkstack. Het internet is gebouwd met behulp van een overvloed aan fysieke platforms, optische glasvezels, Ethernet-kabels, Wi-Fi, satellietsignalen enzovoorts. Om hun functionaliteit te abstraheren en applicaties te laten werken ongeacht het onderliggende platform, werd in de jaren zeventig een gelaagde benadering ontwikkeld (het Internet Protocol). Geïnspireerd door klassieke netwerkstacks, demonstreren we de eerste twee lagen van een quantumnetwerkstack, de *physical layer* (waar de qubits, lasers en signaalgeneratoren toe behoren), en de *link layer*, die het concept van een qubit en het genereren van verstrengeling beheert, zodat ze kunnen worden gebruikt

door toepassingen in de hogere lagen, waardoor de complexiteit van het gebruikte quan-
tumplatform wordt verborgen voor de eindgebruiker.

Sommario

Internet ha rivoluzionato il modo in cui viviamo. Ha consentito applicazioni ben oltre ciò per cui è stato originariamente creata e continuerà a superare le nostre aspettative per il futuro.

I computer quantistici e la rete che li collegherà, la Quantum Internet, probabilmente seguiranno lo stesso percorso. Diversamente dai normali computer, i computer quantistici possono condividere una proprietà chiamata *entanglement*, che consente di collegare i qubit (bit quantistici) a un livello molto più fondamentale. Questa proprietà consente una gamma di nuove applicazioni che spaziano dalla crittografia alla metrologia di precisione.

Negli ultimi quindici anni, sono stati compiuti importanti progressi nella connessione, tramite entanglement, dei primi rudimentali nodi di una rete quantistica. Diverse piattaforme quantistiche hanno dimostrato la generazione di entanglement tra due qubit fisicamente separati. In questa tesi si compie un significativo passo avanti, sia in termini di complessità sperimentale raggiunta, sia nell'astrazione di tale complessità per sviluppi futuri. Andare oltre gli esperimenti a due nodi ha richiesto una riprogettazione fondamentale del nostro apparato sperimentale, oltre a sviluppare le capacità per controllare contemporaneamente il nodo aggiuntivo.

Il primo risultato di questa tesi è la costruzione di una rete quantistica a tre nodi basata sull'entanglement. Abbiamo dimostrato la distribuzione di stati Greenberger-Horne-Zeilinger sulla rete, nonché un elemento costitutivo per reti più grandi: lo scambio di entanglement (entanglement swapping). A differenza delle precedenti dimostrazioni multinodo, che si basavano sulla post-selezione, la nostra rete è in grado di eseguire la distribuzione dell'entanglement in modo *annunciato*: un segnale avviserà gli utenti che il protocollo ha avuto successo e che lo stato è pronto per essere utilizzato.

Il secondo risultato si basa sul primo, aggiungendo il controllo su un quinto qubit, migliorando la qualità dell'entanglement e introducendo una nuova tecnica di lettura ripetitiva, per ottenere il teletrasporto quantistico di un qubit dal terzo nodo al primo—nodi che non condividono un canale di entanglement diretto.

Il terzo e ultimo risultato è la dimostrazione della generazione di entanglement utilizzando una stack di rete quantistica. L'Internet è costruita utilizzando una varietà di piattaforme fisiche: fibre ottiche, cavi Ethernet, Wi-Fi, segnali satellitari, ecc. Per astrarre la loro funzionalità e far funzionare le applicazioni indipendentemente dalla piattaforma sottostante, negli anni '70 è stato sviluppato un approccio a più livelli (il protocollo Internet). Prendendo ispirazione dalla stack di rete classica, dimostriamo i primi due livelli di una stack di rete quantistica, il *livello fisico* (dove risiedono i qubit, i laser e i generatori di segnale) e il *livello di collegamento*, che astrae i concetti di generazione di qubit e di entanglement in modo che possano essere utilizzati dalle applicazioni ai livelli superiori, nascondendo le complessità della piattaforma quantistica utilizzata.

1

1

Introduction

Part of this introduction is taken from a post I wrote for the QuTech Blog in 2018, still in my first year of PhD. While this is not a conventional PhD thesis introduction, I hope it conveys the excitement that I have had during my PhD research, and that I still have for quantum technologies.

1.1 A Quantum Internet made of Diamonds

We are constantly connected to Internet. With our computers, our smartphones, our cars, our fridges (mine is not, yet, but you get the idea). In its very first days, the Internet was a very rudimentary, yet revolutionary, connection between computers [1]. It enabled one computer on the network to send messages to any other computer on the network, whether it was directly connected to it (that is, with a cable) or not. Some of the computers on the network acted as routing nodes for the information, so that it could get directed toward the destination. In 1969 there were four nodes on the then-called ARPANET. By 1973 there were ten times as many. In 1981 the number of connected computers was more than 200. This year (2018)[1] the number of devices capable of connecting to the Internet was 8.4 billion (with a b!) [3]. Computers on their own are already great, but there is a whole range of applications that, without a network infrastructure, would be inaccessible. Do you see where I am going?

If you are not new to this blog you will probably already know what a quantum computer is, and in the last post [4] Stephan Philips showed how we could make one. If I had to shrink to a handful of words, the reason why we want to make them and how is: quantum computers will exploit the weird laws of micro(nano!)scopical objects to solve some problems way faster than any future normal computer. They will do so by encoding the information in quantum systems, which will therefore be quantum information. A Quantum Internet is a network capable of routing this quantum information between quantum computers [5]. We can already foresee some nice applications for this quantum network, like establishing an inviolable secure communication link between any two nodes [6], connecting far-apart telescopes to take ultra-sharp images of stars and galaxies [7], or linking

---

[1]The current forecast, at the time of writing this thesis, is 29.3 billion devices by 2023 [2].
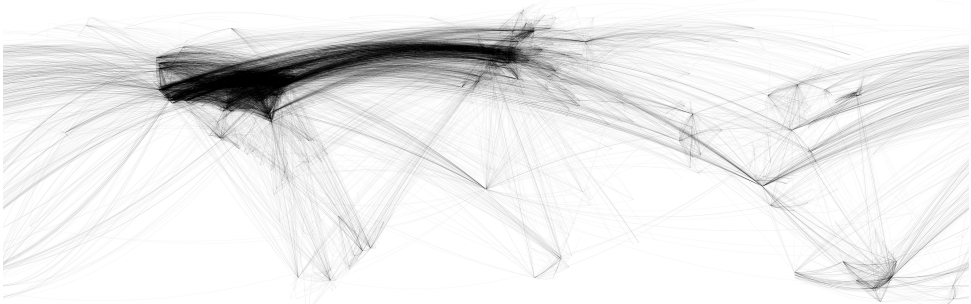
Figure 1.1: "Internet Map" by Chris Harrison: The intensity of edge contrast reflects the number of connections between two points. No country borders or geographic features are shown—the only thing plotted is data. Credit: Chris Harrison, Carnegie Mellon University, USA.

small quantum computers into a huge powerful one (a bit like cloud computing, but much stronger) [8]. In all likelihood, the best way to use these new technologies will come over time, with applications we cannot anticipate. For example, in the early 1960s the invention of the laser was welcomed as "a solution looking for a problem" [9]. Now we have lasers everywhere! I think that the same thing will happen for quantum computers and their internet.

One of the quirks of quantum information is that it cannot be copied. When you try to copy it, you irreversibly destroy the original information. This is, under the hood, what makes quantum communication so secure; but on the other hand, it could also make sharing quantum information really hard! If we had to rely on a single quantum system (a photon for example, a particle of light) to travel undisturbed across the globe, we could as well stop this now. Fortunately, quantum mechanics offers us a solution to this problem: teleportation. It's not like Star Trek, we can't teleport people, but we can teleport quantum information. We can transfer the information stored in a single atom in Amsterdam to an electron in London, without reading the atom, without knowing what the information is and, crucially, without making it travel the distance "physically". Of course, this does not come for free; we have to pay a price, and that price is entanglement.

Quantum mechanics predicts the possibility of this rather weird phenomenon. If you take two separate quantum objects, say two nanometer-sized M&M's, and you make them interact in some particular way, you can make them entangled: the two M&M's lose their individuality and can only be described as one of the parts of a two M&M's system. Let's say that our nano-M&M's can only be of either two colors: red and blue. A non-entangled scenario would be, for example, if the first was red and the second blue. Each M&M's has its own color, its own identity. Now, let's take the two M&M's (one red and one blue), shuffle them a little a bit, just to lose track of which one is which, and then send them, one to you and one to me. When you observe the color of your M&M's, you immediately get to know also the color of mine! While this is interesting, this is not quantum entanglement. This is called correlation, and it is not quantum at all. The two M&M's had their own color the whole time, we simply did not know what it was.

When you entangle the two M&M's, they actually lose their own color! You can make

1
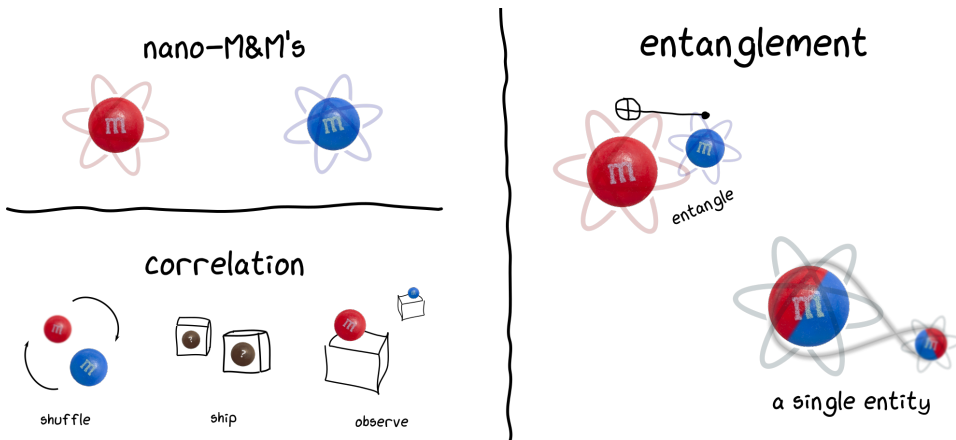
Figure 1.2: Illustration of the difference between classical correlation and entanglement. In the classical case (correlation) the two M&M's had their own color throughout the thought experiment. We simply did not know what it was. In the case of entanglement, the M&M's form a single system, and talking about their individual color does not make sense anymore.

an entangled two-M&M's system, in which you know that the M&M's will have different colors when you will observe them, but until then their color will not be assigned. This effect is so weird that even great scientists believed it was too weird to be true [10]. Now we have the tools to prove with experiments that the effect is indeed real [11–13], and we can exploit it to build new technologies, such as a quantum internet. The idea is to share entangled objects between the nodes of the network and protect their entanglement from the noisy non-quantum environment in which we live, such that when we need to send a quantum bit of information, we can spend the connection of the entangled objects to teleport the (qu)bit. But how can we share entanglement on this network?

This is where our diamonds finally find a place. Diamonds are crystals made of carbon atoms arranged in a very compact way. Sometimes the crystal may have some defects, like an intruder atom (say nitrogen) or a missing carbon atom (what we call a vacancy). If we are lucky enough, these two defects happen one next to the other. Such a system is called an NV center (nitrogen-vacancy). NV centers are one of the most promising candidates to act as nodes of a quantum network. A node is made of three ingredients: a processor to handle information, a memory to temporary store it, and a link to the other nodes.

We use the spin of a pair of electrons localized around the NV center as the processor of our node. We can read its state using lasers and manipulate it using microwave signals. The information stored in this spin has a short lifetime: the system loses memory of the information we store in it too quickly to use it as a reliable memory. Luckily enough, Nature provides us with a strong quantum memory not far apart from the NV center. About 99% of the carbon atoms in nature are ^{12}C which is spin-less; it does not have a spin. Most of the remaining carbons are ^{13}C, which has a spin (due to the additional neutron in the nucleus). We can talk to these ^{13}C atoms in the diamond thanks to their spin-spin interaction with our electrons in the NV center. Since the ^{13}C spin "feels" the electronic spin, we can manipulate the latter to perform operations on the first. The last

1
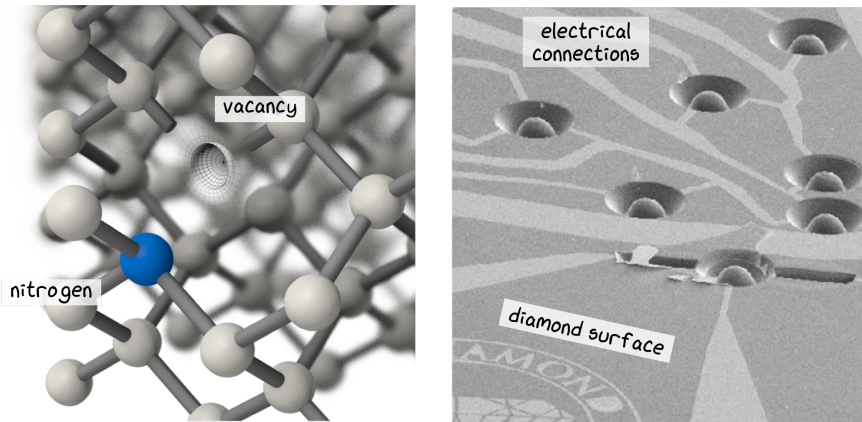
Figure 1.3: (**Left**) Illustration of the lattice of a diamond with an NV center, credit: James Hedberg. (**Right**) Scanning electron microscope image of one of our diamonds. In each of those domes there is an NV center. The electrical connections are used to deliver microwave pulses and other signals.

ingredient of the node is the ability to link it to other nodes. We do this by making the NV center emit a photon that is entangled with the electronic spin (like the M&M's). A second NV center, in a second diamond, in a different place, does the same. Then, by making the two photons interact, we can transfer the entanglement that they have with the electronic spins (photon-spin entanglement), to entanglement between the two electronic spins [14]. This entanglement can then be used for quantum network applications.

We are living a new quantum revolution. We will not be just spectators of quantum mechanics, we will use it as a technology. In a couple of decades we will have access to quantum computers connected through a quantum internet, to design drugs, to optimize airports, to play video games and who knows for what else.
Aren't you excited? I certainly am.

1.2 Thesis overview
This thesis is structured in the following way,

Chapter 2 introduces the NV center in diamond as a platform for quantum network nodes. We provide theoretical background, a practical description of the required control methods, and a detailed account of the hardware and software necessary to operate a diamond-based quantum network node.

In Chapter 3, three independent NV center quantum network nodes are used to demonstrate the first entanglement-based multi-node quantum network. We achieve heralded distribution of Greenberger-Horne-Zeilinger states on the network as well as any-to-any connectivity via entanglement swapping, a fundamental building block for future quantum networks.

Chapter 4 extends on the results of the previous chapter, controlling an additional memory

qubit, implementing a novel repetitive readout scheme and increasing the fidelity of the elementary links, to achieve qubit teleportation between non-neighboring nodes.

In Chapter 5 we take a concrete step towards a programmable quantum network, by implementing and demonstrating platform-independent entanglement distribution via a quantum network stack, composed of a physical layer and a link layer.

Chapter 6 summarizes the results presented in the thesis and gives an outlook towards the development of larger entanglement-based quantum networks.

References

[1] R. Van Meter, *Quantum networking*, Networks and telecommunications series (John Wiley & Sons, 2014).

[2] *Cisco Annual Internet Report - Cisco Annual Internet Report (2018–2023) White Paper*, (2021), https://www.cisco.com/c/en/us/solutions/collateral/executive-perspectives/annual-internet-report/white-paper-c11-741490.html.

[3] A. Nordrum, *Popular Internet of Things Forecast of 50 Billion Devices by 2020 Is Outdated*, (2016), https://spectrum.ieee.org/popular-internet-of-things-forecast-of-50-billion-devices-by-2020-is-outdated.

[4] S. Philips, *Making quantum computers with spin qubits*, (2018), https://blog.qutech.nl/2018/04/20/making-quantum-computers-with-spin-qubits/.

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

[6] A. Ekert and R. Renner, *The ultimate physical limits of privacy*, Nature **507**, 443 (2014).

[7] D. Gottesman, T. Jennewein, and S. Croke, *Longer-Baseline Telescopes Using Quantum Repeaters*, Physical Review Letters **109**, 070503 (2012).

[8] L. Jiang, J. M. Taylor, A. S. Sørensen, and M. D. Lukin, *Distributed quantum computation based on small quantum registers*, Physical Review A **76**, 062323 (2007).

[9] J. Hecht, *Beam: The Race to Make the Laser*, Optics and Photonics News **16**, 24 (2005).

[10] A. Einstein, B. Podolsky, and N. Rosen, *Can Quantum-Mechanical Description of Physical Reality Be Considered Complete?* Physical Review **47**, 777 (1935).

[11] B. Hensen, H. Bernien, A. E. Dréau, A. Reiserer, N. Kalb, M. S. Blok, J. Ruitenberg, R. F. L. Vermeulen, R. N. Schouten, C. Abellán, W. Amaya, V. Pruneri, M. W. Mitchell, M. Markham, D. J. Twitchen, D. Elkouss, S. Wehner, T. H. Taminiau, and R. Hanson, *Loophole-free Bell inequality violation using electron spins separated by 1.3 kilometres*, Nature **526**, 682 (2015).

[12] M. Giustina, M. A. Versteegh, S. Wengerowsky, J. Handsteiner, A. Hochrainer, K. Phelan, F. Steinlechner, J. Kofler, J.-Å. Larsson, C. Abellán, W. Amaya, V. Pruneri, M. W. Mitchell, J. Beyer, T. Gerrits, A. E. Lita, L. K. Shalm, S. W. Nam, T. Scheidl, R. Ursin,

1

B. Wittmann, and A. Zeilinger, *Significant-Loophole-Free Test of Bell's Theorem with Entangled Photons,* Physical Review Letters **115**, 250401 (2015).

[13] L. K. Shalm, E. Meyer-Scott, B. G. Christensen, P. Bierhorst, M. A. Wayne, M. J. Stevens, T. Gerrits, S. Glancy, D. R. Hamel, M. S. Allman, K. J. Coakley, S. D. Dyer, C. Hodge, A. E. Lita, V. B. Verma, C. Lambrocco, E. Tortorici, A. L. Migdall, Y. Zhang, D. R. Kumor, W. H. Farr, F. Marsili, M. D. Shaw, J. A. Stern, C. Abellán, W. Amaya, V. Pruneri, T. Jennewein, M. W. Mitchell, P. G. Kwiat, J. C. Bienfang, R. P. Mirin, E. Knill, and S. W. Nam, *Strong Loophole-Free Test of Local Realism,* Physical Review Letters **115**, 250402 (2015).

[14] H. Bernien, B. Hensen, W. Pfaff, G. Koolstra, M. S. Blok, L. Robledo, T. H. Taminiau, M. Markham, D. J. Twitchen, L. Childress, and R. Hanson, *Heralded entanglement between solid-state qubits separated by three metres,* Nature **497**, 86 (2013).

2

Control methods for NV center quantum nodes

This chapter provides the theoretical background and the experimental methods that are used throughout the thesis. The requirements of a quantum network end-node are listed in Section 2.1. Section 2.2 introduces the nitrogen-vacancy center in diamond, which is the platform for all our experiments, and Section 2.3 outlines the characteristics of our quantum devices. Section 2.4 discusses the control of the communication qubit, which can be used to establish remote entanglement, examined in Section 2.5. Section 2.6 describes the control of memory qubits. Finally, Section 2.7 summarizes the aspects of the experimental setups which are common to the following chapters.

A s the capabilities of the diamond devices have improved over the years, we have felt the need to update the nomenclature we have used in the past for various components of our system, with the intention to accentuate the shift from physics experiments to quantum technology demonstrations, making the concepts discussed here accessible to researchers with diverse background. While some topics remain strictly in the domain of experimental quantum physicists (such as the description of the experimental apparatus), most are also suited for computer scientists and quantum information theoreticians.

In particular, we will often talk in terms of *communication qubits* and *memory qubits*, as opposed to electronic and nuclear spins; *preparation*, *reset* and *measurement*, as opposed to repumping, spin-pumping and readout. Finally, when discussing long-distance entanglement, we will focus on *flying qubits*, as opposed to photons.

2.1 A quantum network node

A future quantum Internet [1] will likely be based on several quantum platforms and technologies, that will take advantage of their specific strengths, and coordinate with each other to perform complex network applications [2]. For example, satellite-based entangled photon sources could allow quantum key distribution over global distances [3] without incurring in the losses of optical fibers. Rare-earth-doped crystals can be used as multiplexed quantum memories to store photons for long-distance quantum repeaters over telecom fibers [4]. The access point of a general-purpose quantum network (quantum network end-node [5]), the one at which a user might perform some application, has different requirements than a quantum repeater or a quantum-key-distribution node, which are part of the network but are not directly exposed to the end-user. The end-node requires:

1. A communication qubit that can be entangled with a flying qubit. Remote end-nodes can then be entangled by means of their flying qubits.

2. Individually addressable memory qubits to store previously generated entangled states.

3. Universal quantum control over the communication qubit and the memory qubits.

4. Real-time communication and feed-forward capabilities to perform advanced quantum network applications.

The experiments in this thesis focus on the realization of quantum networks of end-nodes with diamond-based devices, which fulfill all the above-mentioned requirements.

2.2 The NV center in diamond

The nitrogen vacancy (NV) center in diamond is a point defect consisting of a substitutional nitrogen atom and an adjacent missing carbon atom—a vacancy (fig. 2.1A). In its neutral charge configuration, NV^0, the five electrons from the dangling bonds of the nitrogen atom and the three carbon atoms form an electronic spin-½ system. If an additional electron is captured from the environment (e.g. a nearby charge trap), the negatively charged state is formed, NV^-. This second state has attractive quantum properties that enabled
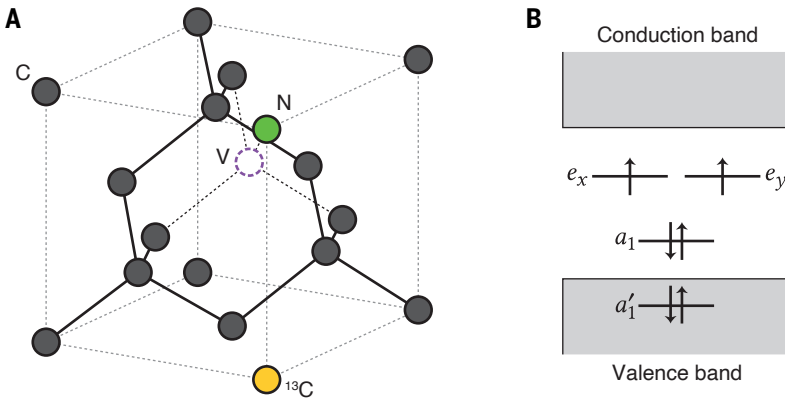
Figure 2.1: The NV center in diamond. (**A**) Schematic representation of the NV center and the surrounding diamond lattice. The ^{12}C (gray) has a natural abundance of ≈98.9 %, while the ^{13}C isotope (yellow) has ≈1.1 %. The nitrogen atom (green) and the neighboring vacancy (V) form the point defect. (**B**) Electronic ground state occupancy of the molecular orbitals of NV$^-$ and their energies with respect to the diamond band-gap. The labeling of the molecular orbitals follows the irreducible representations of the C_{3v} symmetry group of the NV center: A_1, A_2, and E. Adapted from Pfaff and Bernien [6, 7].

the experiments presented in this thesis; when not specified, the negative charge state is implied.

Fig. 2.1B depicts the molecular orbitals of NV$^-$, constructed by linear combination of the atomic orbitals sp^3 from the nitrogen atom and the three carbon atoms surrounding the vacancy [8]. Both the ground state and the first electronic excited states have energy levels well within the band-gap of diamond. This property grants the NV$^-$ properties similar to that of trapped ion, since it prevents electron loss to the conduction band, while being hosted in a solid-state platform.

The six electrons of NV$^-$ form a spin-1 system which in its ground state occupies an orbital-singlet spin-triplet state (fig. 2.2A-B); the communication qubit subspace is defined in two of the three spin states (see Section 2.4). The ground state couples optically to an orbital-doublet spin-triplet excited state, which can decay back into the ground state either directly or via the spin-singlet states. The optical transitions in the spin-triplet states can happen resonantly at 637 nm in the ZPL (zero-phonon line), or with the additional emission of a phonon in the PSB (phonon sideband). As will be discussed later, PSB photons are not suitable for entanglement generation and the fraction of photons that are emitted in the ZPL is ≈2.55 % [9].

The nitrogen atom has either a nuclear spin $I = 1$, in the case of the more common ^{14}N (99.3 % natural abundance), or $I = 1/2$ for ^{15}N (0.7 %). In the following chapters, only ^{14}N NV center are used. While the nuclear spin of the nitrogen atom has been used as qubit in several experimental demonstrations [10, 11], the interaction strength of ≈2 MHz with the NV center electronic spin hinders its use as memory qubit during quantum network applications (see Section 2.6).

Approximately 1.1 % of the carbon atoms (in non-isotopically-purified diamond) are ^{13}C, which carries a nuclear spin $I = 1/2$. The coupling of these spins with the NV cen-

2

ter electronic spin is position dependent, therefore centers with appropriate interactions strengths can be selected, which are suitable for quantum network applications (more in Section 2.6).

The spin-selective optical transitions that form the NV⁻ ZPL, are only resolved at cryogenic temperatures [12, 13]. They enable high-fidelity measurement of the quantum state in a single shot (single-shot read-out) and spin-photon entanglement, which are both required for the quantum network applications presented in this thesis. All our experiments are performed in closed-cycle cryostats at $T \approx 5$ K.

2.3 Device characteristics

The diamond devices used in this thesis are fabricated from high-purity type-IIa (lowest concentration of nitrogen impurities) chemical-vapor-deposition diamond cut along the $\langle 111 \rangle$ crystal orientation (grown by Element Six). Here we give a quick outline of the device fabrication process; for a more detailed treatment see Refs. [14, 15].

Naturally occurring NV centers are identified and initially characterized in a room-temperature setup with green (off-resonant) excitation and PSB collection. We select defects aligned with the $\langle 111 \rangle$ crystal axis (via polarization selection measurements) and without strongly coupled ^{13}C spins in the surroundings (coupling < 1 MHz), as these would not perform as well as memory qubits. Solid immersion lenses (SILs) are milled in the diamond via focused ion beam. The SIL improves the photon collection efficiency by reducing reflections at the diamond-air (or diamond-vacuum) interface. To further reduce the reflections, an Al_2O_3 anti-reflection coating is added via atomic layer deposition to the diamond surface. Gold electrodes (to deliver DC signals) and a gold stripline (to deliver MW pulses as well as DC voltages) are patterned via electron-beam lithography. The sample is finally mounted on a printed circuit board (PCB) and the electrical pads on the diamond sample are wired to connectors on the PCB.

2.4 Communication qubit

A communication qubit can be entangled with a flying qubit (a photon / photonic mode); this enables entangling the communication qubits of spatially separated quantum nodes by interaction and measurement of the flying qubits. The communication qubit of an NV center quantum network node is defined via the electronic spin. The ground-state spin Hamiltonian is (neglecting external electric field and strain contributions)

$$H_{GS}/\hbar = DS_z^2 + \gamma_e(S_x B_x + S_y B_y + S_z B_z), \tag{2.1}$$

where $\hbar$ is the reduced Planck constant, $D \approx 2\pi \times 2.88$ GHz is the zero-field splitting between the $m_S = 0$ and $m_S = \pm 1$ states, $\gamma_e \approx 2\pi \times 2.802$ MHz/G is the gyromagnetic ratio, S_i are the spin-1 Pauli matrices and B_i is the magnetic field along $\hat{i}$; the NV center symmetry axis is defined as $\hat{z}$. The zero-field splitting D is mostly due to spin-spin interactions in the $^3A_2$ state [8]. The second term is the Zeeman effect, which lifts the degeneracy of the $m_S = \pm 1$ states when an external magnetic field is applied. For the experiments in this thesis, magnetic fields aligned to the $\hat{z}$ axis ranging from ≈ 100 G to ≈ 2000 G are applied to the NV centers using permanent magnets. For the full ground-state Hamiltonian (including its dependence on the host nitrogen spin and electric field), see Ref. [8].

The communication qubit is defined as $|0\rangle \equiv |m_S = 0\rangle$ and either $|1\rangle \equiv |m_S = -1\rangle$ or $|1\rangle \equiv |m_S = +1\rangle$. The choice between $m_S = \pm 1$ is due to essentially two factors: first, if one of the two states is too far in energy to be controlled with the available electronics (too high of a transition frequency), the choice is forced onto the other one; second, if both states are experimentally addressable, one of the two will usually have better performance (gate and measurement fidelities).
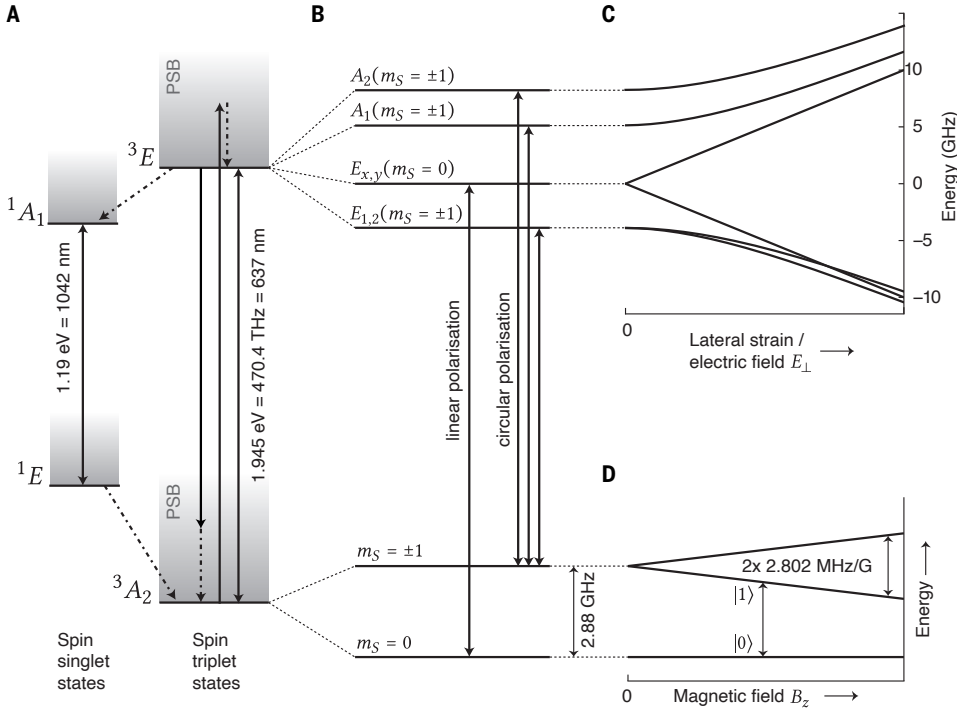
Figure 2.2: Electronic energy structure of NV⁻. (**A**) Optical transitions between the ground states $^3A_2$ and the excited states ^3E can occur resonantly in the ZPL (zero-phonon line) or off-resonantly in the PSB (phonon sideband), with the additional emission of a phonon (dashed-dotted lines). Additionally, ^3E can decay via the singlet states, with spin-dependent decay rates and times [16, 17]. (**B**) A zero-field splitting of ≈2.88 GHz separates the $m_S = 0$ state to the degenerate states $m_S = \pm 1$. At cryogenic temperatures, the excited state ^3E reveals its fine structure [12, 13]. The $E_{x,y}$ and $E_{1,2}$ levels are doubly-degenerate. Transitions from the ground to the excited states are spin selective. (**C**) The energies of the ^3E levels depend on the applied electric field and lateral strain, which have the same effect to first order [18, 19], in contrast to the ground state energies which are relatively unaffected [8]. By applying an electric field across the NV center, it is possible to tune the frequency of the optical transitions for several GHz, as is required for photon indistinguishability during long-distance entanglement generation. (**D**) The $m_S = \pm 1$ degeneracy in the $^3A_2$ ground state can be lifted via Zeeman splitting by applying a magnetic field parallel to the NV center axis; it is then possible to define a qubit subspace using two of the three levels. Adapted from Pfaff, Bernien and Kalb [6, 7, 15].

2.4.1 Reset

The communication qubit is prepared into the $|0\rangle \equiv |m_S = 0\rangle$ state by shining a laser that couples the $|m_S = \pm 1\rangle$ states to the $|A_1\rangle$, $|A_2\rangle$ or $|E_{1,2}\rangle$ excited states (fig. 2.2B). Once in the

2

excited state, the system can either decay back to the $|m_S = \pm1\rangle$ states, decay directly to the $|m_S = 0\rangle$ state (via mixing in the excited state) or decay via the spin-singlet states (fig. 2.2B). The combination of these decay channels and their relative probabilities makes it possible to pump the population into the $|0\rangle$ state with fidelities in excess of 99.5 % [13, 17].

At low magnetic fields ($B_z \lesssim 400$ G) a single laser can be used to couple both $|m_S = \pm1\rangle$ ground states to the $|E_{1,2}\rangle$ excited states. At higher magnetic fields ($B_z = 1890$ G, for node Bob in chapters 3 and 4 and Server in chapter 5), the excited states become closer to being spin eigenstates ($|E_+\rangle$ and $|E_-\rangle$) and the optical transitions $|m_S = +1\rangle \leftrightarrow |E_+\rangle$ and $|m_S = -1\rangle \leftrightarrow |E_-\rangle$ become too far apart to be addressed with a single laser (we found a line splitting of $\approx$250 kHz/G). The higher magnetic field is beneficial for the memory qubits (discussed later) but requires either a second laser or an efficient sideband generation (for example with an electro-optic modulator) for fast and high-fidelity reset.

2.4.2 Measurement

Similar to the reset, the measurement procedure consists in addressing one of the optical transitions with a laser. By shining a weak ($\lesssim$1 nW) pulse on either the $|0\rangle \leftrightarrow |E_x\rangle$ or $|0\rangle \leftrightarrow |E_y\rangle$ transition, the NV center will produce fluorescence only if it is in the $|0\rangle$ state. The choice of $|E_x\rangle$ vs $|E_y\rangle$ comes down to essentially two factors: the first is whether it is possible to tune two NV centers in resonance, as it is likely that only one of the two transitions will shift in the correct direction; the second is the cyclicity of the transitions (which translates into measurement fidelity) given the applied strain or electric field. A good cyclicity of the transition allows the NV center to emit several photons, enabling measurement in a single shot [13]. The detection of at least one photon is regarded as $|0\rangle$ measurement outcome. If no photons are detected, the result of the measurement is assigned to $|1\rangle$. The ability to assign the measurement outcome in a single-shot (i.e. not having to average over several repetitions to get an outcome) is a fundamental advantage in quantum network applications, since it allows us to perform further operations (such as feed-forward gates) based on the outcome of the measurement.

There are two performance metrics that summarize the quality (fidelity) of the measurement procedure:

$$F_0 \equiv \langle 0|M_0|0\rangle \quad \text{and} \quad F_1 \equiv \langle 1|M_1|1\rangle, \tag{2.2}$$

which are respectively the probability to measure outcome 0 (1) given that the qubit was in state $|0\rangle$ ($|1\rangle$). $M_{0/1}$ are the POVMs (positive operator-valued measures) that describe the measurement operation. For the NV centers used in this thesis, $0.92 \lesssim F_0 \lesssim 0.94$, $0.985 \lesssim F_1 \lesssim 0.995$. The number usually reported is the average of the two, which represent the measurement fidelity for an unknown state. Appendix A (page 141) details the procedure of correcting experimental results for the known error in the measurement procedure.

When controlling memory qubits, which will be discussed later in this chapter, it is important to minimize the probability that the communication qubit will change state after the measurement is performed (due to the limited cyclicity of the optical transition used). Such a change in state would result in dephasing noise on the memory qubit. For this reason, during the communication qubit measurement, we monitor the PSB detector for clicks: as soon as a photon is detected, we turn off the measurement laser (reaction time $\approx$5 μs, limited by speed of the DAC module used). To further reduce the probability that the communication qubit will change state after the measurement process is over—

Sequence 2.1: MEASUREMENT OF T_1 CONSTANT

```
1  for i in experimental_repetitions:
2      for time in wait_times:
3          for state in [0, 1]:
4              reset()  # Prepare the qubit in |0⟩.
5              if state is 1:
6                  pi_pulse()  # Rotate to |1⟩ with a π pulse.
7              free_evolution(time)  # Wait "time".
8              if state is 1:
9                  pi_pulse()  # Rotate back to |0⟩ with a π pulse.
10             measure()  # Measure the qubit in the computational basis.
```

2

again, when controlling a memory qubit—we perform a slower measurement ≈100 μs, using ≈100 pW of laser power; this way, the residual time the measurement laser is on will have a smaller effect on the communication qubit. A faster DAC module would further reduce this effect.

2.4.3 Universal control
Arbitrary single-qubit gates on the communication qubit are realized using microwave (MW) pulses at the frequency that is resonant with the $|0\rangle \leftrightarrow |1\rangle$ transition (typically in the 1 GHz to 4 GHz range for magnetic fields below 500 G). The MW pulses are delivered via the on-chip strip-line, producing an oscillating magnetic field that induces transitions between the spin states. The different rotation angles are achieved by varying the duration, shape, and amplitude of the pulse, while different rotation axes on the equatorial plane are selected with the phase of the pulse, which is controlled via the IQ modulation of the MW source. The duration of a pulse, with the equipment in use, is typically ≈100 ns. See Section 3.5.7 (page 54) for details about the pulse shape used and the compensation of distortions induced by the transmission line.

2.4.4 Longitudinal relaxation (T_1)
The terms T_1 and T_2 are borrowed from the NMR field (nuclear magnetic resonance) and, while connected to their original meaning, have slightly different denotations when discussing relaxation times for qubits.

   While in NMR T_1 and T_2 are quantities averaged over many spins (how quickly the net magnetization vector decays for a given sample), for a qubit they are characteristic times of its decay dynamics (a single spin, in the case of the NV center), therefore they are averages of many repetitions of the same experiment.

   In particular, T_1 is the decay time of the qubit population when initialized in one of the computational basis states (eigenstates of the ground-state Hamiltonian, equation (2.1)). The experimental sequence[1] is outlined in Sequence 2.1. The population in the initial state ($|0\rangle/|1\rangle$) relaxes into a mixture of the ground-state spin-states mostly because of residual control fields applied to the NV center when they should instead be switched off: the finite extinction ratio of the MW signal generators and of the optical modulators used for the

---

[1]Experimental sequences are described here using a Python-like pseudocode. In practice, the experiment logic is usually split across multiple devices, such as PCs, microcontrollers, waveform generators and FPGAs. See Section 2.7 for details on the experimental setup.

laser fields prevents them from being completely turned off. This results in an exponential decay of the initial population with characteristic time T_1. Provided a sufficient extinction ratio of the above-mentioned control fields, T_1 approaching 1 hour has been shown on our platform [20]. In the experiments presented in this thesis, T_1 is not a limiting factor, as most sequences are executed in <1 s. It is worth noting that when applying a decoupling sequence to eigenstates, small gate errors will accumulate, leading to an effective depolarizing noise which is greater than what is measured by T_1. This additional noise source should be taken into account when simulating the decoupling of an unknown qubit state.

Other factors contributing to longitudinal relaxation, such as two-photon Raman processes or Orbach-type processes, are negligible at our operating temperature of $\approx 5\,\mathrm{K}$ [21, 22].

2.4.5 Coherence time ($T_{\mathbf{coh}}$)

While T_1 represents the characteristic decay time of the computational basis states, the decay dynamic of superposition states, such as $|x\rangle = (|0\rangle + |1\rangle)/\sqrt{2}$, is caused by different processes and therefore characterized by different constants. In NMR, the spins in the sample will precess at different frequencies due to the local variations in the magnetic field. When preparing the spins in a superposition state with a $\pi/2$ pulse, such as $(|\uparrow\rangle + |\downarrow\rangle)/\sqrt{2}$, each spin will acquire over time a different phase proportional to its precession frequency. The net magnetization vector will, as a consequence, reduce in magnitude as the spins get out of phase (so-called dephasing).

It is possible to recover some lost coherence using decoupling pulses (or echo pulses) which periodically flip the spins, counteracting the effect of different precession frequencies. Several time constants are defined depending on what kind of experiment is performed: The first is the characteristic time $T_2^\star$ of the decay of the net magnetization vector when no decoupling pulses are applied (a measurement known as free induction decay). The second, T_2, is the time constant of the decay in the presence of a decoupling pulse that re-phases the spins in the sample (Hahn echo experiment). Applying more decoupling pulses can lead to a coherence time, T_{coh}, much longer than $T_2^\star$ and T_2. As in the case of T_1, similar constants are defined in the case of a qubit. The NV center communication qubit experiences dephasing due to interactions with the surrounding spin environment (nuclear spins of the $^{13}\mathrm{C}$ atoms and nitrogen atom of the NV center itself). $T_2^\star$ and T_{coh} are then defined as the decay constants extracted via the measurements in Sequence 2.2.

For NV centers in non-isotopically-purified devices, $1\,\mu\mathrm{s} \lesssim T_2^\star \lesssim 10\,\mu\mathrm{s}$ for the communication qubit, which in practice requires the use of decoupling sequences for applications with more than a few gates. T_{coh}, the coherence time, will depend on the quality of the decoupling sequence applied. A higher decoupling pulse rate will generally improve T_{coh}, because it will shield the qubit from higher frequency interactions [23]; the finite fidelity of the decoupling pulses though, will result in error accumulation when the decoupling sequence requires thousands of pulses. For the experiments presented in this thesis, coherence times $T_{\mathrm{coh}} \approx 100\,\mathrm{ms}$ were routinely achieved for the communication qubit using $\approx 10\,000$ decoupling pulses.

Other sources of noise that limit T_{coh} are electrical noise on the microwave stripline which could induce unwanted rotations—when applied over a prolonged time—and projective noise caused by limited extinction of the laser fields.

Sequence 2.2: MEASUREMENT OF $T_2^\star$ AND T_{coh} CONSTANTS

```
1  # T₂* measurement, or free induction decay.
2  for i in experimental_repetitions:
3      for time in wait_times:
4          for state in [x, y, -x, -y]:
5              reset()  # Prepare the qubit in |0⟩.
6              pi2_pulse(state)  # Rotate to |state⟩ with a π/2 pulse.
7              free_evolution(time)  # Wait "time".
8              -pi2_pulse(state)  # Rotate back to |0⟩ with a −π/2 pulse.
9              measure()  # Measure the qubit in the computational basis.
10
11 # T_coh measurement, coherence time.
12 for i in experimental_repetitions:
13     for time in wait_times:
14         for state in [x, y, -x, -y]:
15             reset()  # Prepare the qubit in |0⟩.
16             pi2_pulse(state)  # Rotate to |state⟩ with a π/2 pulse.
17             decoupling_sequence(time)  # Apply a decoupling sequence of "time" duration.
18             -pi2_pulse(state)  # Rotate back to |0⟩ with a −π/2 pulse.
19             measure()  # Measure the qubit in the computational basis.
```

Remarkably, a coherence time longer than 1 second has been reported on our platform with a device similar to ours [20]; to achieve such extended T_{coh}, one needs to carefully characterize the nuclear spin environment to avoid unwanted interactions with surrounding spins during the decoupling sequence.

2.4.6 Preparation

In the last paragraphs, we have assumed that the NV center is in the NV$^-$ charge state and that the lasers used for reset and measurement are resonant with their respective optical transitions. Unfortunately that is not always the case: the NV center could lose one of the six electrons and go into the NV^0 charge state (via a two-photon absorption process [24, 25]), or spectral diffusion due to changes in the local charge environment could shift the optical transition away from the laser frequencies. Therefore, before starting to use the NV center's qubits, one needs to make sure that it is in the right charge state and in resonance with all the relevant optical transitions. This is accomplished with a procedure called a charge and resonance check (CR check).

The CR check (outlined in Sequence 2.3) consists in shining both the reset (spin pumping) and measurement (readout) lasers for a fixed amount of time, usually ≈50 µs. If the NV center is in the correct charge state (NV$^-$), and the lasers are on resonance with the transitions used, a considerable amount of fluorescence can be collected in the PSB: with a few nW on both lasers, emission rates ≈10 MHz can be reached, resulting in tens of photons collected during the procedure (including photon loss). If the number of counts exceeds a preset threshold, the CR check is successful and the NV center is ready to be used. If the threshold is not met, three scenarios are possible: 1) the NV center is in resonance with the lasers and in NV$^-$, but the randomness in the number of photons detected made it such that the threshold was not met; 2) the NV center is in the correct charge state, but the lasers are not properly on resonance, therefore the fluorescence is not as strong as it should be; 3) the NV center is in the NV^0 state, which does not produce a strong signal when excited with 637 nm lasers. One can distinguish case 3) from the other two

Sequence 2.3: CHARGE AND RESONANCE CHECK

```python
1   def CR_check(cr_check_time, resonance_threshold, repump_threshold):
2       # Perform a charge and resonance check.
3       # Return True if CR check succeeded, False otherwise.
4
5       turn_on_measure_laser()
6       turn_on_reset_laser()
7
8       counts = collect_photons_for(cr_check_time)
9
10      turn_off_measure_laser()
11      turn_off_reset_laser()
12
13      if (counts > resonance_threshold):
14          return True   # The NV center is ready to be used.
15      else:
16          if (counts < repump_threshold):
17              # We are very likely in the NV^0 charge state, apply a pulse
18              # of either green or yellow laser to go back to NV^-.
19              apply_repump_laser_pulse()
20          # The CR check was not successful and should be repeated. The experiment cannot start yet.
21          return False
22
23  # CR check usage example.
24  for i in experimental_repetitions:
25      repeat:
26          preparation_result = CR_check(50e-6, 30, 1)
27      until (preparation_result is True)  # Repeat the CR check until it succeeds.
28      perform_actual_experiment(repetition=i)
```

using an additional threshold: NV^0 will emit less than one photon per CR check in general, much less than an off-resonant NV^- (provided the laser frequencies are not too far off the required optical transitions).

There are two procedures to convert NV^0 into the required NV^- state (a procedure called repumping, not to be confused with spin-pumping, which is instead communication qubit reset). The first is to shine a strong and short laser pulse (tens of µW for tens of µs) with an off-resonant green laser (510 to 540 nm). After the green laser pulse, the probability to find the NV center in NV^- is ≈75 % [24]. The green laser pulse also induces spectral diffusion, which can help to bring the NV^- in resonance with the lasers when the CR check is not successful. The second is to use a yellow laser (575 nm) to address the NV^0 ZPL, which deterministically ionizes the NV center to NV^- with a weak and long laser pulse (tens of nW for ≈300 µs). This technique, while being more advanced and technologically demanding, has the benefit of being compatible with DC Stark tuning of the NV center [10, 26], which is required to tune the optical transitions for entanglement generation. The CR check procedure will be repeated until successful, therefore selecting only NV center configurations that are in the correct charge state and in resonance with the required lasers.

2.5 Remote entanglement generation

While in a classical network the fundamental task is delivering a package of information (packet) from one node to the other, quantum network applications require the distribution of entangled states across the nodes [2].

NV centers in diamond have been used to establish remote entanglement across network nodes over the past decade [10, 26–29]. Photons emitted in the ZPL—i.e. without the concurrent emission of a phonon—can be used as flying qubits to interface with remote network nodes. The emission of a phonon, that is lost to the environment, leaves the spin-photon state in a classically-correlated but non-entangled state (indeed the correlation with the spin state is used to perform qubit measurement).

The optical fibers connecting the remote nodes will inevitably induce photon loss. Probabilistic protocols that use the detection of a photon as heralding signal, counteract the photon loss at the expense of success rate.

While the experiments presented in this thesis only employ a single-photon protocol, it is useful to first discuss the two-photon protocol, which is easier to realize experimentally, at the expense of a reduced probability of success.

2.5.1 Two photon protocol

The protocol, proposed by Barrett and Kok [30], starts by resetting each communication qubit in the $|0\rangle$ state. A $\pi/2$ pulse around the $\hat{y}$ axis on each qubit brings them to

$$|\psi\rangle_{A/B} = \frac{|0\rangle + |1\rangle}{\sqrt{2}}, \tag{2.3}$$

where A and B stand for Alice and Bob, the name of the two quantum nodes. A short (≈ 1 ns) optical pulse on either the $|0\rangle \rightarrow |E_x\rangle$ or $|0\rangle \rightarrow |E_y\rangle$ transition, selectively brings the NV center to the optically excited state; from there, it will spontaneously decay, emitting a photon. The communication qubit is now entangled with the presence/absence of a photon in the fluorescence of the NV center (the flying qubit):

$$|\psi\rangle_{A/B} = \frac{|0\rangle_c |1\rangle_\gamma + |1\rangle_c |0\rangle_\gamma}{\sqrt{2}}. \tag{2.4}$$

Here $|0/1\rangle_c$ are the communication qubit states and $|0/1\rangle_\gamma$ are the photon number states. We now apply a π pulse to the communication qubit, followed by a second optical excitation pulse:

$$|\psi\rangle_{A/B} = \frac{|1\rangle_c |1\rangle_\gamma |0\rangle_{\gamma'} + |0\rangle_c |0\rangle_\gamma |1\rangle_{\gamma'}}{\sqrt{2}}, \tag{2.5}$$

By defining the flying qubit states as early $|e\rangle = |1\rangle_\gamma |0\rangle_{\gamma'}$ (if a photon is present in the first time bin and not in the second one) and late $|l\rangle = |0\rangle_\gamma |1\rangle_{\gamma'}$ (if a photon is present in the second time bin and not in the first one), the state can now be written with the time-bin encoding as

$$|\psi\rangle_{A/B} = \frac{|1\rangle_c |e\rangle + |0\rangle_c |l\rangle}{\sqrt{2}}. \tag{2.6}$$

The two optical modes are now brought on a balanced beam splitter that removes the which-path information: on the other side of the beam splitter, it is not possible to know whether a photon came from Alice's side or Bob's side. By selecting events in which one photon is detected in the early bin and one photon is detected in the late bin, the two

communication qubits are projected into the entangled state

$$|\Psi^\pm\rangle = \frac{|0\rangle_A |1\rangle_B \pm |1\rangle_A |0\rangle_B}{\sqrt{2}},\qquad(2.7)$$

where the $\pm$ sign depends on whether the two photons were detected on the same detector or on different ones (the beam splitter has a π phase difference between the two output arms [31]). The protocol can be also be described as an entanglement swapping procedure on the flying qubits: the beam splitter entangles the two flying qubits, which are then projected into one of the four possible outcomes. Since the flying qubits were entangled with the communication qubits, projecting the first in an entangled state, will leave the second also in an entangled state.

One of the experimental requirements of this protocol is that the photons emitted by the two NV centers should be indistinguishable, i.e. should have the same polarization, the same temporal mode and the same spatial mode. Polarization and spatial mode are guaranteed by collecting the ZPL photons in polarization maintaining (PM) single mode fibers and interfering them on a fiber beam splitter (also known as fiber coupler). For the temporal mode, photons in the same time-bin from two different nodes should arrive simultaneously at the beam splitter, and they should have the same frequency.

The lasers used to excite the NV centers need to be frequency-stabilized. In Ref. [26], where the lasers (637 nm) were at a distance of 1.3 km, light from the first laser was sent through an optical fiber and interfered with light from the second laser. A feedback loop would then lock the frequency of the two lasers. Over longer distances, one could lock the lasers to cavities, and then lock the cavities to each other using a second frequency that is compatible with standard telecom fiber.

The probability of success of this protocol is $p_{det}^2/2$, where p_{det} is the probability to detect a photon after the beam splitter after an optical excitation. It includes the probability to emit in the ZPL, the collection efficiency out of the diamond, the losses in the time, frequency and spatial filtering, the propagation losses in the optical fibers to the beam splitter and finally the efficiency of the single photon detectors. With bulk samples (like the ones used in this thesis), p_{det} is on the order of 10^{-3} to 10^{-4}, excluding the losses due to propagation in the optical fibers. For lab-scale experiments (tens of meters of fibers at most), those losses are negligible, but for long-distance experiments, such as Ref. [26], they are critical (637 nm light propagates with ≈ 10 dB/km loss in single mode optical fiber). Given that an entanglement attempt takes $\approx 5\,\mu s$, the two photon protocol limits the entanglement generation rate on our devices (and in lab-scale experiments) to $\approx 10\,$mHz (one heralded state every few minutes), not including experimental overhead (such as, for example, CR checking).

The Barret and Kok protocol has been used to generate heralded entanglement between remote NV centers in Refs. [10, 26, 27]. Other platforms for quantum network end-nodes also use two photon protocols to generate entanglement, for example trapped ions [32, 33], and neutral atoms [34, 35].

2.5.2 Single photon protocol
While the Barret and Kok protocol allows for, in principle, very high fidelity entangled states (there are no protocol-induced errors, all the noise sources could in principle be

mitigated, see Ref. [26]), the generation rate scales with the square of the detection probability, which makes it extremely sensitive to photon losses.

The protocol proposed by Cabrillo *et al.* [36] and by Bose *et al.* [37] uses the detection of a single photon to herald an entangled state, dramatically increasing the generation rate at the expense of a reduced fidelity and higher technological complexity.

The protocol starts with the reset of the communication qubits to the $|0\rangle$ state. A rotation around the $\hat{y}$ axis then brings each qubit into the superposition state

$$|\psi\rangle_{A/B} = \sqrt{\alpha_{A/B}}|0\rangle + \sqrt{1 - \alpha_{A/B}}|1\rangle, \tag{2.8}$$

with α (the population in the $|0\rangle$ state) in general different for the two setups. The rotation angle required to achieve the state is $\vartheta(\alpha) = 2\cos^{-1}(\sqrt{\alpha})$. An optical excitation pulse (the same as with the Barret and Kok two photon protocol), creates an entangled state of the communication qubit and the presence/absence of a photon in the fluorescence of the NV center,

$$|\psi\rangle_{A/B} = \sqrt{\alpha_{A/B}}|0\rangle_c|1\rangle_\gamma + \sqrt{1 - \alpha_{A/B}}|1\rangle_c|0\rangle_\gamma. \tag{2.9}$$

Different from the two-photon protocol, this is all we need from the NV center to generate our entangled state: by directing the flying qubits toward a central beam splitter and letting them interfere, we remove the which-path information from the quantum state, entangling the flying qubits. By measuring one photon on the other side of the beam splitter we herald a state between the communication qubits that is close (more details later) to the EPR pair

$$|\Psi^\pm\rangle = \frac{|0\rangle_A|1\rangle_B \pm e^{i\Delta\theta}|1\rangle_A|0\rangle_B}{\sqrt{2}}, \tag{2.10}$$

where the ± sign depends on which of the two detectors clicked, and the $\Delta\theta$ phase between the two states depends on the optical path difference that the two flying qubits have to travel before interfering on the beam splitter.

Access to this entangled state relies on the ability to know and stabilize the optical phase. Fluctuations in the optical setup would otherwise randomize $\Delta\theta$ multiple times a second, washing out the quantum correlations from the heralded state: the generated state would change from shot to shot, making it unusable. Active phase stabilization for entanglement generation was first demonstrated by Stockill *et al.* with quantum dots [38], and later realized for NV center network nodes in Ref. [29]. The added experimental requirements of optical phase stabilization make single-photon protocol much more challenging than two-photon protocols. Chapter 3 describes a novel stabilization scheme that is used to phase-stabilize the three-node quantum network, which improves on the feedback bandwidth of Ref. [29] by almost three orders of magnitude. The faster feedback allows for the generation of entangled states using optical fibers that, connecting setups in different laboratories, are subject to much higher phase noise.

The probability of success per attempt, assuming balanced probabilities of detection, high losses ($p_{\text{det}}^2 \ll p_{\text{det}}$, and equal α for Alice and Bob, is $2\alpha p_{\text{det}}$, bringing the generation rate into the $10\,\text{Hz}$ regime, three orders of magnitude faster than with the two-photon protocol on similar devices [29].

Apart from the technological challenge of phase stabilization, the single-photon protocol suffers from an additional entangled state infidelity that scales as $F = 1 - \alpha$. This is due

to the possibility that one of the two photons emitted by the nodes is not detected, either because it is lost, or because of the non-number-resolving detectors employed. More details about modeling of the generated states and the noise sources that affect this protocol are in Chapters 3 and 4.

2.6 Memory qubits

Each quantum network node requires at least one communication qubit to establish entanglement with a remote node. Once entanglement is generated, it needs to be stored in additional qubits—we call these memory qubits.

The nuclear spins of ^{13}C atoms ($I = 1/2$) in the surroundings of defect centers have been used in pioneering experiments as memory qubits, both to perform quantum computation and to store long-distance entanglement once generated [11, 28, 39–41]. The nuclear spin of the host nitrogen ($I = 1$ for ^{14}N) in the NV center has also been used as memory qubit, see Refs. [6, 11]. In this thesis, we will only focus on ^{13}C memory qubits, since they offer better resilience to the entanglement generation procedure [17, 28, 42] (more on this below).

Each ^{13}C atom in the surrounding of the NV center is coupled to the electronic spin via an always-on hyperfine interaction. The Hamiltonian of each ^{13}C nuclear spin is therefore a sum of an external-magnetic-field-dependent term, and a hyperfine interaction term between the nuclear spin and the electronic spin of the NV center (composed of an isotropic part, the Fermi contact interaction, and a non-isotropic part, the dipole-dipole interaction, see Ref. [43] for details). If the hyperfine interaction is small compared to the external magnetic field applied (secular approximation), the ^{13}C nuclear spin Hamiltonian in the communication qubit rotating frame is

$$H/\hbar = \omega_L I_z + A_\parallel S_z I_z + A_\perp S_z I_x, \tag{2.11}$$

with $\omega_L = \gamma_{13C} B_z$ the Larmor frequency of the nuclear spin, $\gamma_{13C} = 2\pi \times 1.071 \, \text{kHz/G}$ the gyromagnetic ratio for ^{13}C, I_i are the spin-1/2 Pauli matrices for the nuclear spin, and $A_\parallel$ and $A_\perp$ are the position-dependent hyperfine couplings between the electronic spin and the ^{13}C nuclear spin. It is possible to rearrange the Hamiltonian and reveal the electronic-spin dependent dynamics of the ^{13}C nuclear spins:

$$H = H_0 |0\rangle\langle 0| + H_1 |1\rangle\langle 1|,$$
$$H_0/\hbar = \omega_L I_z, \tag{2.12}$$
$$H_1/\hbar = (\omega_L \pm A_\parallel) I_z \pm A_\perp I_x,$$

where $|0\rangle\langle 0|$ and $|1\rangle\langle 1|$ are the communication qubit (electronic spin) projectors. The $\pm$ sign in the Hamiltonian depends on the communication qubit $|1\rangle$ state definition ($m_S = \pm 1 \rightarrow \pm$).

When the communication qubit is in $|0\rangle$, the ^{13}C nuclear spin experiences precession due to the applied external magnetic field. If the communication qubit is in $|1\rangle$, the nuclear spin precession frequency and precession axis are modified; the new precession frequency is $\tilde{\omega}_L = \sqrt{(\omega_L \pm A_\parallel)^2 + A_\perp^2}$. By carefully designing decoupling sequences for the communication qubit, it is possible to apply a time-averaged Hamiltonian to the memory qubit that applies arbitrary communication-qubit-controlled and -uncontrolled gates. Alternatively,

Sequence 2.4: NV CENTER FINGERPRINT MEASUREMENT

```
1  for i in experimental_repetitions:
2      for tau in interpulse_delays:
3          reset()  # Prepare the qubit in |0⟩.
4          pi2_pulse()  # Rotate to |+⟩ with a π/2 pulse.
5          # Apply a decoupling sequence with "tau" interpulse delay and "N" π pulses.
6          decoupling_sequence(tau, N)
7          -pi2_pulse()  # Rotate back to |0⟩ with a -π/2 pulse.
8          measure()  # Measure the qubit in the computational basis.
```

it is possible to interleave the decoupling sequence with direct drive of the individual memory qubit transition with radio-frequency pulses, as demonstrated in Ref. [11].

2.6.1 Universal control

The decoupling sequences used throughout this thesis for the communication qubit are of the form $(-\tau-\pi-2\tau-\pi-\tau-)^{N/2}$, with N the number of applied microwave π pulses and τ the refocusing time. In particular, we use the XY8 sequence, which alternates π pulses around the X axis and the Y axis in the following way: XY XY YX YX; this sequence is robust against imperfect calibration of the pulses [44], for example consistently over- or under-rotations.

During the decoupling operation, the memory qubit experiences an evolution that depends, in general, on the initial state of the communication qubit. In particular, the memory qubit will undergo a rotation around an axis that (might) depend on the communication qubit's initial state, and an angle that depends on the number of π pulses applied. For details on the Hamiltonian that the memory qubit experiences during the decoupling sequence, see Ref. [39].

By measuring the coherence left in the communication qubit after preparing it in a superposition state and applying a decoupling sequence, it is possible to probe the environment of the NV center and characterize the memory qubits available: sharp drops in coherence at periodic values of τ represent resonance conditions at which the communication qubit can perform conditional gates with a specific memory qubit (see fig. 2.3). Since each nuclear spin environment is unique to the NV center used (because of the random position of the ^{13}C atoms in the diamond lattice) such a measurement—outlined in Sequence 2.4—is called an NV center fingerprint.

2.6.2 Lifetime

Due to their always-on interaction with the communication qubit, the memory qubits experience different noise mechanisms depending on the activity of the quantum network node. Using dynamical decoupling sequences, coherence times >10 s have been shown for ^{13}C-nuclear-spin memory qubits when the network node is idling and the communication qubit is kept in $|1\rangle$ [11].

When the network node is busy generating remote entanglement, the memory qubits are subject mostly to dephasing noise that arises from imperfect control of the communication qubit [17, 42]. Depending on the state of the communication qubit, a memory qubit will precess at one of two different frequencies (see equation (2.12)). To operate the
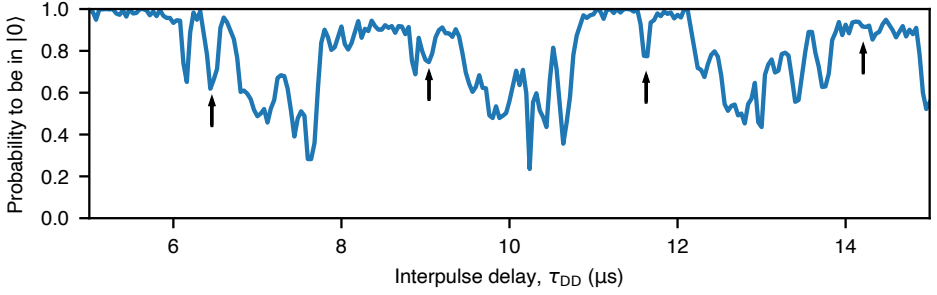
Figure 2.3: NV center fingerprint measurement. Result of performing Sequence 2.4 with one of our NV centers. The decoupling sequence is composed of 16 inversion pulses (two XY8 blocks). The arrows point at the expected resonances of a single memory qubit with parallel hyperfine coupling $A_{\parallel} = 2\pi \times 31.5\,\text{kHz}$. Once a memory qubit frequency has been identified, one can proceed with the optimization of the conditional and unconditional gates (by varying the number of pulses applied, and the resonance selected), for details on the calibration procedure see Refs. [15, 39, 45].

memory qubit, we keep track of the phase acquired by the nuclear spin at all times. This is relatively straightforward in deterministic sequences, such as a decoupling sequence on the communication qubit, where the time spent at each frequency can be known in advance. When the communication qubit operation is non-deterministic, instead, the memory qubit will experience a dephasing proportional to the time the communication qubit spends in an unknown state.

During each entanglement attempt, the communication qubit is reset in $|0\rangle$ using a fast ($\approx 1.5\,\mu\text{s}$) and strong ($\approx 1\,\mu\text{W}$) laser pulse on the spin-pumping transition. Depending on the time after which a collapse to $|0\rangle$ would eventually happen, and on the overall effectiveness of the reset process, the memory qubit will acquire a different phase. Since the acquired phase changes from entanglement attempt to entanglement attempt in a stochastic way, the effect on the memory qubit is an effective dephasing.

After the microwave α pulse—that creates the communication qubit superposition—the memory qubit will start to acquire a communication-qubit-dependent phase. This is the same mechanism that allows us to entangle the memory qubit with the communication qubit (and therefore to perform controlled gates). Upon (likely) failure of the entanglement attempt, the communication qubit will be reset to $|0\rangle$ (in the next entanglement attempt) and the memory qubit will have acquired an unknown phase. To mitigate this effect, it is possible to add a communication qubit decoupling pulse during the entanglement generation attempt that balances the time spent by the memory qubit under the two different Hamiltonians before the communication qubit is reset in the next attempt. Since this additional decoupling pulse has finite fidelity, it introduces an additional dephasing mechanism for the memory qubit during network activity. Numerical simulations for the relative and absolute effects of the communication qubit reset and the decoupling pulse on the dephasing of memory qubits are in Ref. [17].

Since the noise induced by the network activity is orders of magnitude higher than that without entanglement generation, a common metric (used also throughout this thesis) for the lifetime of quantum network memory qubits is $N_{1/e}$, the number of entanglement

attempts after which the Bloch vector length of the memory qubit has decreased to $1/e$ of its initial value. A state-of-the-art experiment performed in our group, prior to the work presented here, showed $N_{1/e} \approx 270$ attempts for weakly coupled ^{13}C nuclear spins ($\Delta\omega = |\omega_L - \tilde{\omega}_L| \approx 25$ kHz). Using decoherence-protected subspaces, $N_{1/e} \approx 1000$ have been shown with our platform [46], albeit with reduced memory qubit control fidelity. This number should be compared to the average number of attempts required to establish entanglement with current diamond devices using a single-photon protocol, $N_{\text{ent}} \approx 20\,000$.

Successful storage of a quantum state in the memory qubit during entanglement generation currently relies on protocol timeouts: if entanglement has not been generated in a pre-set number of attempts, the protocol is restarted. In the experiments presented in Chapter 3, following the proposal of Kalb *et al.* [17] to increase the magnetic field applied to the NV center, we report an $N_{1/e} \approx 1800$ and show that the lifetime is in that case limited by $T_2^\star$ of the memory qubit rather than the entanglement generation procedure.

We further improved the memory qubit lifetime in the experiments of Chapter 4, where using a decoupling pulse on the memory qubit after the successful entanglement generation, we achieve $N_{1/e} \approx 5300$. More complex decoupling sequences, such as those using direct driving of the memory qubits—demonstrated in Ref. [11]—might bring the memory qubit lifetime well past N_{ent}, opening the way to deterministic quantum network protocols that do not rely on timeouts.

2.7 Experimental setup

An NV center-based quantum network end-node is composed of several interconnected components. The main one is of course the diamond device itself, as it has been discussed previously. To make the quantum device function, several other components are required: the cryogenic system to keep the device cold; the optics to deliver and collect light to and from the NV center; the electronics for, among other things, qubit control; last but not least, the software that allows to control all the experimental devices and coordinate the experiments.

2.7.1 Cryogenics

Our diamond devices are installed in commercially-available 4 K closed cycles cryostats, either Montana Instruments s50 or attocube attoDRY 800. The cryogenic temperature is required to resolve the fine structure of the ^3E electronic excited state (fig. 2.2), and therefore to perform spin-selective excitation of the NV center for both single-shot measurement and reset of the communication qubit and spin-photon entanglement (entanglement of the communication qubit with a flying qubit for remote entanglement generation). A room-temperature microscope objective (Olympus MPLFLN100x) is brought in proximity of the device (working distance 1 mm) for confocal microscopy on the NV center.

In the case of the Montana Instruments s50, we modify the cryostat housing in-house and connect an external room-temperature vacuum housing that will contain a 3-axis piezo stage (Physik Instrumente P-615). The microscope objective is mounted on the piezo stage, allowing for sub-nm scanning resolution with approximately $10 \times 10 \times 10$ µm range. The device is mounted on an in-house machined gold-plated copper sample-holder (diamond surface in the zx plane). Neodymium permanent magnets (supermagnete.de) are

2
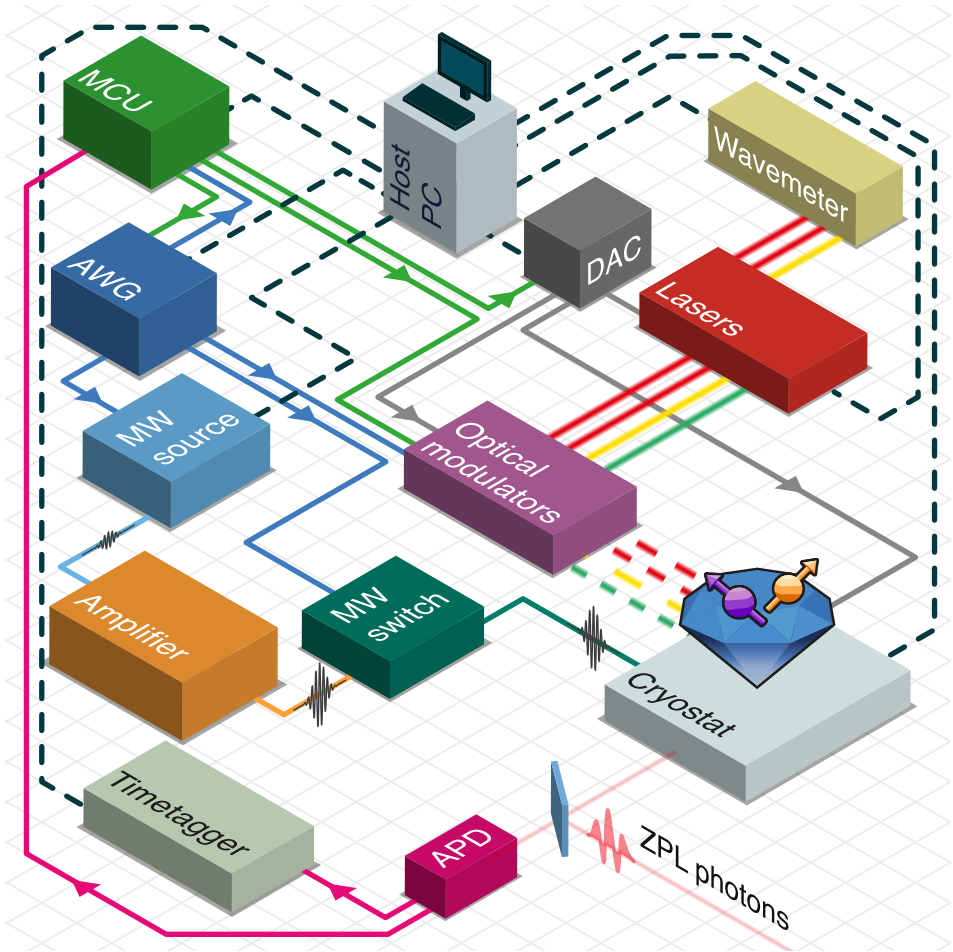
Figure 2.4: Software and hardware connections for the quantum network node. A host PC is connected to several experimental devices (dashed dark-blue lines), both to exchange data in real-time (wavemeter, DAC, timetagger, cryostation) and to upload programs (MCU and AWG). The MCU, which is responsible for microsecond-control of the setup, communicates with the AWG, in charge of the nanosecond-control, either via simple digital triggers (chapters 3 and 4) or additionally using a parallel digital interface (chapter 5). Both the MCU and the AWG can control the optical modulators (AOMs and EOM) to apply light pulses to the NV center. The AWG controls the MW source via IQ modulation. The GHz-range pulses from the MW source are amplified and, after passing through a MW switch that increases on-off ratio, are delivered to the diamond device in the cryostat. PSB photons from the NV center are detected with an APD; the events are both recorded by a timetagger connected to the host PC (for analysis and debugging), and monitored in real-time by the MCU to perform qubit measurements. The ZPL photons, separated from the PSB with a dichroic mirror, are sent to the central beam splitter to generate remote entanglement. The lasers that require frequency stabilization (qubit measurement and reset, NV center preparation) are also sent to a wavemeter; the host PC runs software-based PID loops that stabilize the frequency via a DAC module (occasionally MCU DAC channels are also used). A DC signal generated by the DAC module (controlled and modulated by the MCU) is sent to the NV center to tune the frequency of the optical transitions.

mounted on the back of the sample-holder to provide a constant magnetic field.

For the attocube attoDRY 800, we mount the sample on an attocube positioning stage, composed of three stepper modules (ANPx101 and ANPx101, mm travel range) and a 3-axis scanner module for sub-nm positioning (ANSxyz100). On top of the stack, an in-house built adapter plate holds the permanent magnet. A thermal connection plate (ATC100) ensures good thermal connection between the cryostat cold plate and the sample, which is mounted on the very top of the stack horizontally (diamond surface in the xy plane). The microscope objective is mounted on the cryostat housing and approaches the sample from the top, once the cryostat is closed.

The systems are actively stabilized at ≈5 K using a heating element to minimize temperature fluctuations when applying high-duty-cycle microwave sequences (such as entanglement generation). Temperature fluctuations, because of our design, induce changes in the temperature of the permanent magnet, which in turn leads to changes in the magnetic field.

If properly vacuum-tight, the systems can stay cold without significant ice buildup for several months. If ice is detected (reduced fluorescence or increased spectral diffusion) a partial warm-up, pumping, and cool-down usually is enough to return to a working point. Occasional compressor or main unit software glitches can result in unwanted system warm-up. It is good to take precautions if sensitive piezo devices are at cold temperature, because the change in temperature (without grounding the piezos) can lead to permanent damage to the stack. In our case, a monitoring loop checks the status of the cryostat, and in case a malfunction is detected, it grounds all cryogenic piezos.

2.7.2 Optics

The optics in the quantum network nodes play essentially three roles: delivering laser light to the NV center, collecting single photons from the NV center, and phase stabilization for entanglement generation.

When possible, we use in-fiber optical components, since they require less maintenance than their free-space counterparts. Throughout our setups, optical fibers with free-space interface (such as for in-coupling or out-coupling light) have FC/APC (ferrule connector, angled physical contact) termination, that minimizes back reflection both in the fiber and in free-space. For fiber-to-fiber connection (butt connection) of PM fibers, we find that one should prefer either splicing, or FC/PC (non-angled) termination. Using FC/APC termination significantly reduces the PER (polarization extinction ratio) of the combined fiber patch. Using a PM splicer (Ericsson FSU 995 PM), we routinely achieve ≤0.1 dB loss, 30 dB PER splices.

The lasers used in the node range in wavelength from 515 nm to 637 nm. For lasers that need to be frequency stabilized (all except the green), a 1 to 10% tap is sent to a wavemeter (HighFinesse WS8 / WS6) and stabilized with a PC-controlled PID loop (≈1 MHz long-term stability).

We use AOM (acousto-optic modulators, G&H Fiber-Q) to generate laser pulses (rise-time ≈20 ns). In some cases (yellow-repump modulation, phase stabilization frequency offset) we dynamically choose the frequency of the AOM driving to change the frequency shift applied to the laser (see section 3.5). To generate the short laser pulses (1 ns) used for entanglement generation, we use an in-fiber electro-optic amplitude modulator (EOM,

Jenoptik AM635). The various laser beams are combined in free-space using beam splitters and dichroic mirrors. The combined beam is coupled into a PM460-HP fiber (Thorlabs) that provides good transmission for all the wavelengths used. The fiber is then launched in free-space inside a closed box (for darkness) that includes the cryostat, the PSB detectors, the ZPL collection and the phase stabilization optics. For details on the layout, see Fig. 3.6 on page 57. Both cryostat models have anti-reflection-coated wedged windows—to deliver and collect light—reducing in-path reflections.

Fluorescence in the PSB is separated, from ZPL and reflected light, with a dichroic mirror (Semrock, 635 nm), collected into a multimode fiber (to increase long-term alignment stability) and measured into a single photon detector (Laser Components COUNT). The multimode fiber has anti-reflection coating on both ends to increase collection and detection efficiency.

ZPL photons are separated from reflected excitation light using a band-pass filter (Semrock, 5 nm) and polarization rejection (Thorlabs PBSW-633). A deformable mirror (Boston Micromachines Multi-DM) is also in the ZPL optical path, to correct for aberrations in the optical setup and increase the efficiency of collection of the single photons in a single mode PM fiber (PM-630HP). The single mode fibers that carry the ZPL photons from two nodes are combined on a in-fiber 50 : 50 beam splitter (Evanescent Optics). One of two input arms of the beam splitter has a piezo fiber-stretcher incorporated (see section 3.5) that is used for phase stabilization. The ZPL photons are detected by SNSPDs (superconducting nanowire single photon detectors) optimized for 637 nm, with >95 % detection efficiency and a few dark counts per minute (PhotonSpot).

2.7.3 Electronics

Controlling a diamond-based quantum-network end-node requires electronics that operate at many different frequencies, time scales and logical capabilities. A single device able to perform all the required tasks is not available yet on the market, therefore we use several devices (each with their strengths and weaknesses) that interact with each other to cover the span of required signals.

Low-noise high-voltage DC signals are provided by an in-house built DAC and amplifier (QuTech SPI D5a + S1h, 18 bits, ±90 V). These are connected to either the gates on the diamond device, or they are used to bias the microwave stripline.

Microwave signals are generated with an IQ modulated source (R&S SGS100A) that produces signals up to 12 GHz. The output of the source is brought to the required power level of ≈45 dBm with an amplifier (AR 40S1G4), goes through a in-house built MW switch (based on the HMC8038, Analog Devices) and is then connected to the microwave stripline of the device. The output of the device is terminated on 50 Ω after being attenuated to reduce back-reflections.

An arbitrary waveform generator (AWG) is responsible for nanosecond-accuracy signal generation and simple experimental logic. In Chapters 3 and 4 we used the Tektronix 5014 to generate MW pulses envelopes (which are used for IQ modulation of the MW source) as well as driving signals for the AOMs and for the EOM. Additionally, the AWG produced digital signals to coordinate and communicate with other equipment. The 5014 can have simple logic embedded in the uploaded waveforms, such as wait for trigger before playing, repeat a sequence a certain number of times, jump to another element. An

additional microcontroller, discussed below, communicates with the 5014 to achieve the complex sequences demonstrated in Chapters 3 and 4. The limited programming function- alities (and limited communication complexity with the microcontroller) of the 5014 would have made the quantum network stack demonstration presented in Chapter 5 possibly un- feasible (and definitely not scalable). The AWGs in the network nodes were upgraded to Zurich Instruments HDAWG8. These, apart from a higher timing resolution and faster reaction time, expose a programmable FPGA and a digital communication interface that enables complex sequences and information to be shared in real-time.

A time-deterministic microcontroller (MCU) is fundamental to ensure that the required operations (such as turning on a laser, or counting photons) are executed at the right time and for the correct duration. We use the Jäger ADwin-Pro II T12 on each node. The AD- win is an event-based programmable microcontroller with 1 GHz clock rate, that can be connected to several modules (DAC, ADC, DIO, counters, SPI, ...) and exchange informa- tion across them. Most of the network node logic (for which μs accuracy is sufficient) is executed by the microcontroller—NV center preparation, reset, measurement, synchro- nization with another node, feedbacks, feedforward communication and so on. For tasks that require higher timing resolution (such as MW pulses or entanglement generation), the microcontroller triggers the AWG and waits for a confirmation or fail signal to come back before proceeding to the next step.

A complex programmable logic device (CPLD, less complex but similar functionalities as an FPGA) is used to herald entanglement (Altera MAX V 5M570ZF256C5N). By perform- ing a suitable logic/counting operation with the ZPL photon arrival and a time reference provided by the AWG, it detects whether a photon has arrived in the specified window and produces the heralding signal used by all the other devices.

For calibration, analysis and debugging purposes, we use time-taggers to store the arrival times of photons both in the PSB (PicoQuant TimeHarp, 1 ns or 250 ps time bin) and in the ZPL (PicoQuant HydraHarp, 1 ps time bin). These are used, for example, when calibrating the amplitude of the optical excitation pulses, when making sure the photons from two different nodes arrive at the same time at the central beam splitter, and when analyzing data to only select events that were heralded in a specific time window.

2.7.4 Software

During the course of this thesis, we have changed the software framework used to operate the quantum network node. In particular, in the experiments of Chapters 3 and 4, the software was written for a relatively outdated measuring framework called qtlab, which was developed in 2008 for QuTech, and it was based on Python 2.7. It allows one to control instruments connected to the PC, acquire and save data, display plots and create GUIs (graphical user interfaces). For the quantum network stack demonstration of Chapter 5, a new framework was developed by QuTech called QMI, built on more modern technologies, cross-platform, multithreaded and with network transparency. For both qtlab and QMI, the host PC runs Python scripts to interface with instruments. Via the host PC, we program the other devices that are used during the experiments.

The ADwin is programmed with a Basic-like language (called ADbasic). The program is compiled by the host PC, uploaded to the ADwin and triggered by the PC. The structure of an ADwin program is as follows: a first INIT section is executed when the program

starts and is generally used to initialize variables or prepare equipment (such as turning off all lasers). An event timer is pre-configured to execute the section EVENT periodically. In our case, the event timer is 1 µs. This means that the ADwin will execute the same piece of code every microsecond, deterministically. To perform multi-cycle operations, we build a finite state machine in the EVENT section that describes what to do for each possible state and what the state should be in the next cycle.

In the case of the Tektronix 5014 AWG, the program is essentially a text file with a list of waveforms to play, with additionally the options *wait for trigger*, *repeat (N times)*, *jump to* and *go to*. For the Zurich Instruments HDAWG8, alongside the list of waveforms, a C-like program can be uploaded that specifies the sequence (and logic) with which the waveforms should be played. This allows for counting, choosing pulses and making decisions directly from the AWG, instead of having to go back to the MCU when complex logic is required.

Finally, a small VHDL program for the CPLD implements the heralding scheme.

References

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

[2] R. Van Meter, *Quantum networking*, Networks and telecommunications series (John Wiley & Sons, 2014).

[3] S.-K. Liao, W.-Q. Cai, J. Handsteiner, B. Liu, J. Yin, L. Zhang, D. Rauch, M. Fink, J.-G. Ren, W.-Y. Liu, Y. Li, Q. Shen, Y. Cao, F.-Z. Li, J.-F. Wang, Y.-M. Huang, L. Deng, T. Xi, L. Ma, T. Hu, L. Li, N.-L. Liu, F. Koidl, P. Wang, Y.-A. Chen, X.-B. Wang, M. Steindorfer, G. Kirchner, C.-Y. Lu, R. Shu, R. Ursin, T. Scheidl, C.-Z. Peng, J.-Y. Wang, A. Zeilinger, and J.-W. Pan, *Satellite-Relayed Intercontinental Quantum Network,* Physical Review Letters **120**, 030501 (2018).

[4] D. Lago-Rivera, S. Grandi, J. V. Rakonjac, A. Seri, and H. de Riedmatten, *Telecom-heralded entanglement between multimode solid-state quantum memories,* Nature **594**, 37 (2021).

[5] S. Wehner, D. Elkouss, and R. Hanson, *Quantum internet: A vision for the road ahead,* Science **362** (2018), 10.1126/science.aam9288.

[6] W. Pfaff, *Quantum measurement and entanglement of spin Quantum bits*, PhD Thesis, Delft University of Technology (2013).

[7] H. Bernien, *Control, measurement and Entanglement of Remote Quantum Spin Registers in Diamond*, PhD Thesis, Delft University of Technology (2014).

[8] M. W. Doherty, N. B. Manson, P. Delaney, F. Jelezko, J. Wrachtrup, and L. C. L. Hollenberg, *The nitrogen-vacancy colour centre in diamond,* Physics Reports The nitrogen-vacancy colour centre in diamond, **528**, 1 (2013).

[9] D. Riedel, I. Söllner, B. J. Shields, S. Starosielec, P. Appel, E. Neu, P. Maletinsky, and R. J. Warburton, *Deterministic Enhancement of Coherent Photon Generation from a Nitrogen-Vacancy Center in Ultrapure Diamond,* Physical Review X **7**, 031040 (2017).

[10] W. Pfaff, B. J. Hensen, H. Bernien, S. B. v. Dam, M. S. Blok, T. H. Taminiau, M. J. Tiggelman, R. N. Schouten, M. Markham, D. J. Twitchen, and R. Hanson, *Unconditional quantum teleportation between distant solid-state quantum bits,* Science **345**, 532 (2014).

[11] C. Bradley, J. Randall, M. Abobeih, R. Berrevoets, M. Degen, M. Bakker, M. Markham, D. Twitchen, and T. Taminiau, *A Ten-Qubit Solid-State Spin Register with Quantum Memory up to One Minute,* Physical Review X **9**, 031045 (2019).

[12] P. Tamarat, N. B. Manson, J. P. Harrison, R. L. McMurtrie, A. Nizovtsev, C. Santori, R. G. Beausoleil, P. Neumann, T. Gaebel, F. Jelezko, P. Hemmer, and J. Wrachtrup, *Spin-flip and spin-conserving optical transitions of the nitrogen-vacancy centre in diamond,* New Journal of Physics **10**, 045004 (2008).

[13] L. Robledo, L. Childress, H. Bernien, B. Hensen, P. F. A. Alkemade, and R. Hanson, *High-fidelity projective read-out of a solid-state spin quantum register,* Nature **477**, 574 (2011).

[14] B. Hensen, *Quantum Nonlocality with Spins in Diamond*, PhD Thesis, Delft University of Technology (2016).

[15] N. Kalb, *Diamond-based quantum networks with multi-qubit nodes*, PhD Thesis, Delft University of Technology (2018).

[16] M. Goldman, A. Sipahigil, M. Doherty, N. Yao, S. Bennett, M. Markham, D. Twitchen, N. Manson, A. Kubanek, and M. Lukin, *Phonon-Induced Population Dynamics and Intersystem Crossing in Nitrogen-Vacancy Centers,* Physical Review Letters **114**, 145502 (2015).

[17] N. Kalb, P. C. Humphreys, J. J. Slim, and R. Hanson, *Dephasing mechanisms of diamond-based nuclear-spin memories for quantum networks,* Physical Review A **97**, 062330 (2018).

[18] P. Tamarat, T. Gaebel, J. R. Rabeau, M. Khan, A. D. Greentree, H. Wilson, L. C. L. Hollenberg, S. Prawer, P. Hemmer, F. Jelezko, and J. Wrachtrup, *Stark Shift Control of Single Optical Centers in Diamond,* Physical Review Letters **97**, 083002 (2006).

[19] L. C. Bassett, F. J. Heremans, C. G. Yale, B. B. Buckley, and D. D. Awschalom, *Electrical Tuning of Single Nitrogen-Vacancy Center Optical Transitions Enhanced by Photoinduced Fields,* Physical Review Letters **107**, 266403 (2011).

[20] M. H. Abobeih, J. Cramer, M. A. Bakker, N. Kalb, M. Markham, D. J. Twitchen, and T. H. Taminiau, *One-second coherence for a single electron spin coupled to a multi-qubit nuclear-spin environment,* Nature Communications **9**, 2552 (2018).

[21] S. Takahashi, R. Hanson, J. van Tol, M. S. Sherwin, and D. D. Awschalom, *Quenching Spin Decoherence in Diamond through Spin Bath Polarization,* Physical Review Letters **101**, 047601 (2008).

2

[22] A. Jarmola, V. M. Acosta, K. Jensen, S. Chemerisov, and D. Budker, *Temperature- and Magnetic-Field-Dependent Longitudinal Spin Relaxation in Nitrogen-Vacancy Ensembles in Diamond,* Physical Review Letters **108**, 197601 (2012).

[23] G. d. Lange, Z. H. Wang, D. Ristè, V. V. Dobrovitski, and R. Hanson, *Universal Dynamical Decoupling of a Single Solid-State Spin from a Spin Bath,* Science **330**, 60 (2010).

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

[25] S. Baier, C. Bradley, T. Middelburg, V. Dobrovitski, T. Taminiau, and R. Hanson, *Orbital and Spin Dynamics of Single Neutrally-Charged Nitrogen-Vacancy Centers in Diamond,* Physical Review Letters **125**, 193601 (2020).

[26] B. Hensen, H. Bernien, A. E. Dréau, A. Reiserer, N. Kalb, M. S. Blok, J. Ruitenberg, R. F. L. Vermeulen, R. N. Schouten, C. Abellán, W. Amaya, V. Pruneri, M. W. Mitchell, M. Markham, D. J. Twitchen, D. Elkouss, S. Wehner, T. H. Taminiau, and R. Hanson, *Loophole-free Bell inequality violation using electron spins separated by 1.3 kilometres,* Nature **526**, 682 (2015).

[27] H. Bernien, B. Hensen, W. Pfaff, G. Koolstra, M. S. Blok, L. Robledo, T. H. Taminiau, M. Markham, D. J. Twitchen, L. Childress, and R. Hanson, *Heralded entanglement between solid-state qubits separated by three metres,* Nature **497**, 86 (2013).

[28] N. Kalb, A. A. Reiserer, P. C. Humphreys, J. J. W. Bakermans, S. J. Kamerling, N. H. Nickerson, S. C. Benjamin, D. J. Twitchen, M. Markham, and R. Hanson, *Entanglement distillation between solid-state quantum network nodes,* Science **356**, 928 (2017).

[29] P. C. Humphreys, N. Kalb, J. P. J. Morits, R. N. Schouten, R. F. L. Vermeulen, D. J. Twitchen, M. Markham, and R. Hanson, *Deterministic delivery of remote entanglement on a quantum network,* Nature **558**, 268 (2018).

[30] S. D. Barrett and P. Kok, *Efficient high-fidelity quantum computation using matter qubits and linear optics,* Physical Review A **71**, 060310 (2005).

[31] H. Fearn and R. Loudon, *Quantum theory of the lossless beam splitter,* Optics Communications **64**, 485 (1987).

[32] D. L. Moehring, P. Maunz, S. Olmschenk, K. C. Younge, D. N. Matsukevich, L.-M. Duan, and C. Monroe, *Entanglement of single-atom quantum bits at a distance,* Nature **449**, 68 (2007).

[33] L. Stephenson, D. Nadlinger, B. Nichol, S. An, P. Drmota, T. Ballance, K. Thirumalai, J. Goodwin, D. Lucas, and C. Ballance, *High-Rate, High-Fidelity Entanglement of Qubits Across an Elementary Quantum Network,* Physical Review Letters **124**, 110501 (2020).

[34] S. Ritter, C. Nölleke, C. Hahn, A. Reiserer, A. Neuzner, M. Uphoff, M. Mücke, E. Figueroa, J. Bochmann, and G. Rempe, *An elementary quantum network of single atoms in optical cavities,* Nature **484**, 195 (2012).

[35] J. Hofmann, M. Krug, N. Ortegel, L. Gérard, M. Weber, W. Rosenfeld, and H. Weinfurter, *Heralded Entanglement Between Widely Separated Atoms,* Science **337**, 72 (2012).

[36] C. Cabrillo, J. I. Cirac, P. García-Fernández, and P. Zoller, *Creation of entangled states of distant atoms by interference,* Physical Review A **59**, 1025 (1999).

[37] S. Bose, P. L. Knight, M. B. Plenio, and V. Vedral, *Proposal for Teleportation of an Atomic State via Cavity Decay,* Physical Review Letters **83**, 5158 (1999).

[38] R. Stockill, M. Stanley, L. Huthmacher, E. Clarke, M. Hugues, A. Miller, C. Matthiesen, C. Le Gall, and M. Atatüre, *Phase-Tuned Entangled State Generation between Distant Spin Qubits,* Physical Review Letters **119**, 010503 (2017).

[39] T. H. Taminiau, J. Cramer, T. van der Sar, V. V. Dobrovitski, and R. Hanson, *Universal control and error correction in multi-qubit spin registers in diamond,* Nature Nanotechnology **9**, 171 (2014).

[40] J. Cramer, N. Kalb, M. A. Rol, B. Hensen, M. S. Blok, M. Markham, D. J. Twitchen, R. Hanson, and T. H. Taminiau, *Repeated quantum error correction on a continuously encoded qubit by real-time feedback,* Nature Communications **7**, 11526 (2016).

[41] C. Nguyen, D. Sukachev, M. Bhaskar, B. Machielse, D. Levonian, E. Knall, P. Stroganov, R. Riedinger, H. Park, M. Lončar, and M. Lukin, *Quantum Network Nodes Based on Diamond Qubits with an Efficient Nanophotonic Interface,* Physical Review Letters **123**, 183602 (2019).

[42] M. Pompili, S. L. N. Hermans, S. Baier, H. K. C. Beukers, P. C. Humphreys, R. N. Schouten, R. F. L. Vermeulen, M. J. Tiggelman, L. d. S. Martins, B. Dirkse, S. Wehner, and R. Hanson, *Realization of a multinode quantum network of remote solid-state qubits,* Science **372**, 259 (2021).

[43] A. P. Nizovtsev, S. Y. Kilin, A. L. Pushkarchuk, V. A. Pushkarchuk, S. A. Kuten, O. A. Zhikol, S. Schmitt, T. Unden, and F. Jelezko, *Non-flipping 13C spins near an NV center in diamond: hyperfine and spatial characteristics by density functional theory simulation of the C510[NV]H252 cluster,* New Journal of Physics **20**, 023022 (2018).

[44] T. Gullion, D. B. Baker, and M. S. Conradi, *New, compensated Carr-Purcell sequences,* Journal of Magnetic Resonance (1969) **89**, 479 (1990).

[45] J. Cramer, *Quantum error correction with spins in diamond,* PhD Thesis, Delft University of Technology (2016).

[46] A. Reiserer, N. Kalb, M. S. Blok, K. J. van Bemmelen, T. H. Taminiau, R. Hanson, D. J. Twitchen, and M. Markham, *Robust Quantum-Network Memory Using Decoherence-Protected Subspaces of Nuclear Spins,* Physical Review X **6**, 021040 (2016).

2

3

A multinode quantum network of remote solid-state qubits

M. Pompili*, S. L. N. Hermans*, S. Baier*, H. K. C. Beukers,
P. C. Humphreys, R. N. Schouten, R. F. L. Vermeulen,
M. J. Tiggelman, L. dos Santos Martins, B. Dirkse, S. Wehner &
R. Hanson

The distribution of entangled states across the nodes of a future quantum internet will unlock fundamentally new technologies. Here, we report on the realization of a three-node entanglement-based quantum network. We combine remote quantum nodes based on diamond communication qubits into a scalable phase-stabilized architecture, supplemented with a robust memory qubit and local quantum logic. In addition, we achieve real-time communication and feed-forward gate operations across the network. We demonstrate two quantum network protocols without postselection: the distribution of genuine multipartite entangled states across the three nodes and entanglement swapping through an intermediary node. Our work establishes a key platform for exploring, testing, and developing multinode quantum network protocols and a quantum network control stack.

F uture quantum networks sharing entanglement across multiple nodes [1, 2] will enable a range of applications such as secure communication, distributed quantum computing, enhanced sensing, and fundamental tests of quantum mechanics [3–8]. Efforts in the past decade have focused on realizing the building blocks of such a network: quantum nodes capable of establishing remote entangled links as well as locally storing, processing, and reading out quantum information.

Entanglement generation through optical channels between a pair of individually controlled qubits has been demonstrated with trapped ions and atoms [9–12], diamond nitrogen-vacancy (NV) centers [13, 14], and quantum dots [15, 16]. In addition, a number of quantum network primitives have been explored on these elementary two-node links, including nonlocal quantum gates [17, 18] and entanglement distillation [19]. Moving these qubit platforms beyond two-node experiments has so far remained an outstanding challenge owing to the combination of several demanding requirements. Multiple high-performance quantum nodes are needed that include a communication qubit with an optical interface as well as an efficient memory qubit for storage and processing. Additionally, the individual entanglement links need to be embedded into a multinode quantum network, requiring a scalable architecture and multinode control protocols.

Here, we report on the realization and integration of all elements of a multinode quantum network: optically mediated entanglement links connected through an extensible architecture, local memory qubit and quantum logic, and real-time heralding and feed-forward operations. We demonstrate the full operation of the multinode network by running two key quantum network protocols. First, we establish Greenberger-Horne-Zeilinger (GHZ) entangled states across the three nodes. Such distributed genuine multipartite entangled states are a key ingredient for many network applications [2] such as anonymous transmission [20], secret sharing [21], leader election [22], and clock stabilization [8]. Second, we perform entanglement swapping through an intermediary node, which is the central protocol for entanglement routing on a quantum network enabling any-to-any connectivity [23, 24]. Owing to efficient coherence protection on all qubits, combined with real-time feed-forward operations, these protocols are realized in a heralded fashion, delivering the final states ready for further use. This capability of heralding successful completion of quantum protocols is critical for scalability; its demonstration here presents a key advance from earlier experiments using photons [25] and quantum memories [26].

Our network is composed of three spatially separated quantum nodes (fig. 3.1, A and B), labeled Alice, Bob, and Charlie. Each node consists of an NV center electronic spin as a communication qubit. In addition, the middle node Bob uses a carbon-13 nuclear spin as a memory qubit. Initialization and single-shot readout of the communication qubits are performed through resonant optical excitation and measurement of state-dependent fluorescence [14]. Universal quantum logic on the electronic-nuclear register is achieved through tailored microwave pulses delivered on chip (section 3.5). The nodes are connected through an optical fiber network for the quantum signals, as well as classical communication channels for synchronizing the control operations and relaying heralding signals (see below).

Remote entanglement generation hinges on indistinguishability between emitted photons. For NV centers in high-purity low-strain diamond devices, the optical transition
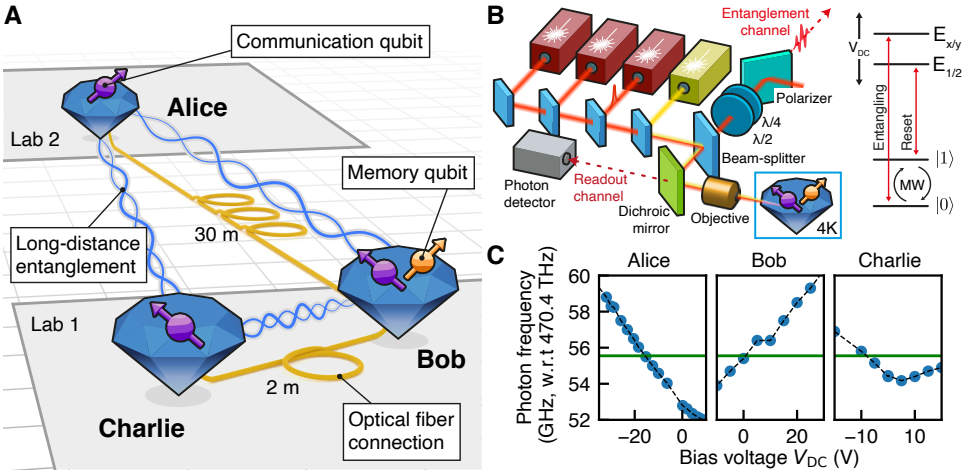
Figure 3.1: The three-node quantum network. (**A**) Layout of the network. Three nodes, labeled Alice, Bob, and Charlie, are located in two separate labs. Each node contains an NV center communication qubit (purple). At Bob, an additional nuclear spin qubit (orange) is used in the presented experiments. Fiber connections between the nodes (lengths indicated) enable remote entanglement generation on the links Alice-Bob and Bob-Charlie, which, combined with local quantum logic, allow for entanglement to be shared between all nodes (wiggly lines). (**B**) On the left is a simplified schematic of the optical setup at each node (see fig. 3.6, table 3.1, and section 3.5 for additional details). On the right is a diagram of the relevant levels of the electronic spin qubit, showing optical transitions for remote entanglement generation and readout ("entangling"), qubit reset ("reset"), and resonant microwaves ("MW") for qubit control (see figs. 3.7 and 3.8 for additional details). The memory qubit at Bob is initialized, controlled, and read out via the electronic qubit (fig. 3.9). Optical transition frequencies are tuned via the dc bias voltages (V_{DC}). $\lambda/2$ ($\lambda/4$) is a half-waveplate (quarter-waveplate); $E_{x/y}$ and $E_{1/2}$ are electronic excited states. (**C**) Tuning of the optical "entangling" transition at each of the three nodes. The solid line is the working point, 470.455 55 THz; the dashed line is a guide to the eye. w. r. t. , with respect to.

frequencies show relatively minor variations (few GHz). We remove the remaining offsets by using dc Stark tuning at each node with bias fields generated on chip (fig. 3.1C). We are thus able to bring the relevant optical transitions of all three nodes to the same frequency, which we choose to be the zero-bias frequency of Bob.

3.1 Establishing remote entanglement in a network architecture

To generate remote entanglement between a pair of nodes (i.e., one elementary link), a single-photon protocol is used [27, 28] (fig. 3.2A). The communication qubits of the nodes are each prepared in a superposition state $|\alpha\rangle = \sqrt{\alpha}|0\rangle + \sqrt{1-\alpha}|1\rangle$. At each node, pulsed optical excitation, which is resonant only for the $|0\rangle$ state, and subsequent photon emission deterministically create an entangled state between the communication qubit and the presence-absence of a photon (the flying qubit). The photonic modes from the two nodes are then interfered on a beam splitter, removing the which-path information. The beam splitter closes an effective interferometer formed by the optical excitation and collection paths. Detection of a single photon after the beam splitter heralds the state $|\psi^{\pm}\rangle \approx (|01\rangle \pm e^{i\Delta\theta}|10\rangle)/\sqrt{2}$ between the two communication qubits, where the $\pm$ sign de-

pends on which of the two detectors clicked and $\Delta\theta$ is the optical phase difference between the two arms of the effective interferometer (section 3.5). Experimentally, this phase difference is set to a known value by stabilizing the full optical path using a feedback loop [14, 16]. This scheme yields states at maximum fidelity $1 - \alpha$ at a rate $\approx 2\alpha p_{\text{det}} r_{\text{attempt}}$, with p_{det} the probability that an emitted photon is detected and r_{attempt} the entanglement attempt rate.

Scaling this entangling scheme to multiple nodes requires each elementary link to be phase-stabilized independently (fig. 3.2B), posing a number of new challenges. The different links, and even different segments of the same link, will generally be subject to diverse noise levels and spectra. Additionally, the optical power levels used are vastly different, from microwatts for the excitation path to attowatts for the single-photon heralding station, requiring different detector technologies for optimal signal detection. We solve these challenges with a hybrid phase-stabilization scheme that is scalable to an arbitrary number of nodes. We decompose the effective interferometer for each link into three independently addressable interferometers and stabilize each separately (see fig. 3.2C for the Alice-Bob link; the link Bob-Charlie is phase-stabilized in an analogous and symmetric way; see figs. 3.10 and 3.13).

First, each node has its own local stabilization that uses unbalanced heterodyne phase detection (fig. 3.2C, left). In comparison to the previous homodyne stabilization method [14], this enables us to obtain a higher bandwidth phase signal from the small part of the excitation light that is reflected from the diamond surface ($\approx 1\,\%$) by boosting it with a strong reference-light beam at a known frequency offset. Moreover, this scheme allows for optimal rejection of the reflected excitation light by polarization selection, thus preventing excitation light from entering the single-photon path toward the heralding detectors and creating false entanglement heralding events. The measured phase signals are fed back on piezoelectric-mounted mirrors to stabilize the local interferometers.

Second, the global part of the effective interferometer (fig. 3.2C, right) is stabilized by single-photon-level homodyne phase detection with feedback on a fiber stretcher: A small fraction of the strong reference-light beam is directed into the single-photon path, and the interference is measured using the same detectors used for entanglement generation.

This architecture provides scalability in the number of nodes and a higher feedback bandwidth compared with our previous implementation on a single link (fig. 3.14; see section 3.5 for details). In our current implementation, the central node – Bob – has combining optics to merge the signals coming from Alice and Charlie, so that the single-photon detectors can be shared by the two links.

Crucially, this architecture enables the successive generation of entanglement on the two elementary links as required for network protocols exploiting multinode entanglement. We benchmark its performance by running entanglement generation on both elementary links within a single experimental sequence (fig. 3.2D).
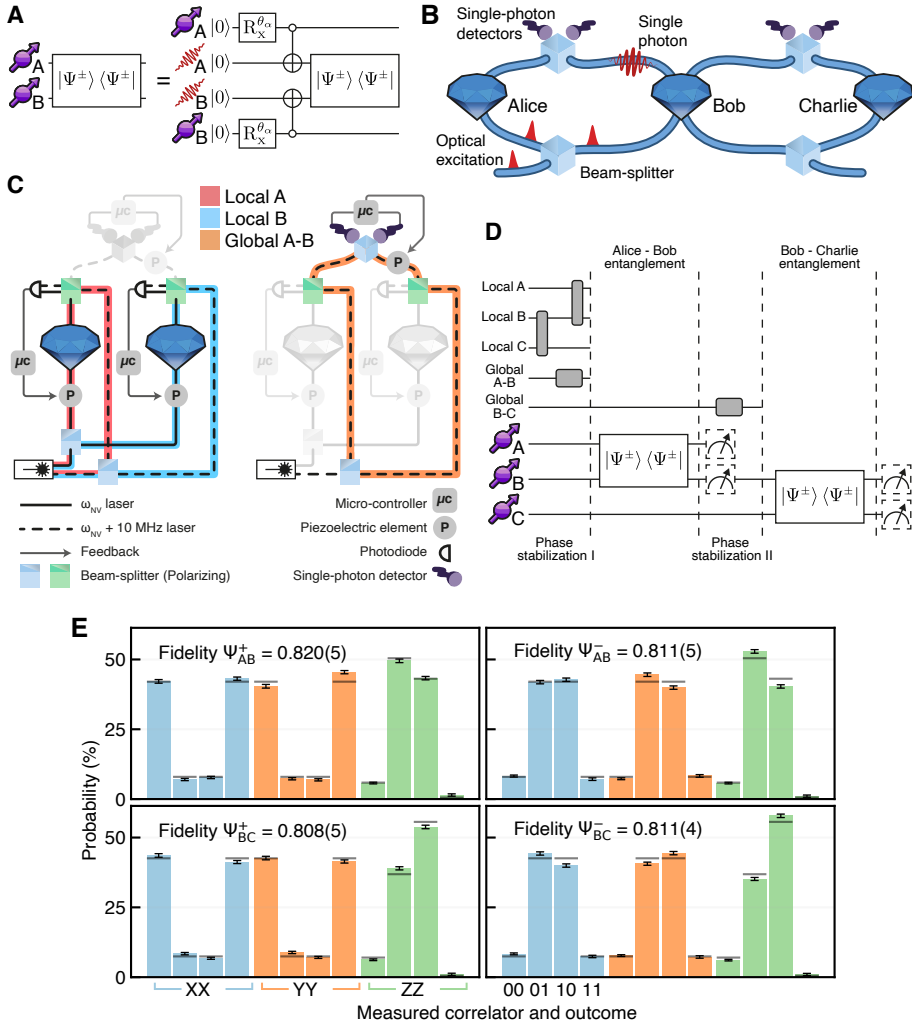
Figure 3.2: Establishing remote entanglement in a network architecture. (**A**) Circuit diagram of the single-photon entanglement protocol, where $R_X^{\theta_\alpha}$ is a rotation around the x axis with angle θ_α, $\theta_\alpha = 2\cos^{-1}(\sqrt{\alpha})$. (**B**) Sketch of three quantum network nodes in line configuration, showing the two effective interferometers. (**C**) Phase stabilization diagram of the Alice-Bob link, highlighting the local interferometers (left) and the global interferometer (right). See Section 3.5 for further details. (**D**) Experimental sequence to generate Bell pairs on both Alice-Bob (A-B) and Bob-Charlie (B-C) links. Dashed boxes display measurements used in (E). (**E**) Correlation measurements on entangled states on A-B (top) and B-C (bottom) links. The left plots correspond to $|\Psi^-\rangle$ states; the right plots correspond to $|\Psi^+\rangle$ states. Shown are observed probabilities for outcomes (from left to right) 00, 01, 10, and 11 for correlation measurements in the bases XX (blue), YY (orange), and ZZ (green). Gray bars depict values from the theoretical model. Error bars indicate one standard deviation.

We achieve fidelities of the entangled Bell states exceeding 0.8 for both links (fig. 3.2E), on par with the highest fidelity reported for this protocol for a single link [14]. For the same fidelity, the entangling rates are slightly higher than in Ref. [14] (9 and 7 Hz for

links Alice-Bob and Bob-Charlie, respectively), despite the additional channel loss from connecting the two links. The main sources of infidelity are the probability α that both nodes emit a photon, remaining optical phase uncertainty, and double excitation during the optical pulse (see table 3.2 and section 3.5). A detailed physical model that includes known error sources is used here and below for comparison to the experimental data (section 3.5); predictions by the model are indicated by the gray bars in the correlation and fidelity plots.

3.2 Memory qubit performance and real-time feed-forward operations

To distribute entangled states across multiple nodes, generated entangled states must be stored in additional qubits while new entanglement links are created. Carbon-13 nuclear spins are excellent candidates for such memory qubits, thanks to their long coherence times, controllability, and isolation from the control drives on the electronic qubit [29]. Recent work [30] indicated that their storage fidelity under network activity is mainly limited by dephasing errors resulting from the coupling to the electronic spin that is randomized on failed entanglement generation. It was suggested that the memory robustness to such errors may be further improved by operating under an increased applied magnetic field. Here, we use a magnetic field of 189 mT for our central node, as opposed to ≈40 mT used in past experiments [19, 30].
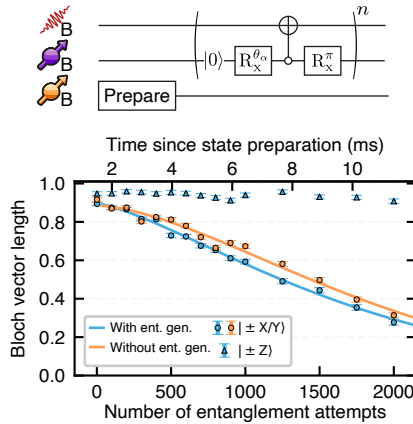
Figure 3.3: Memory qubit coherence under network activity. (**Top**) Circuit diagram displaying the experimental sequence, where n is the number of entanglement attempts. (**Bottom**) Blue represents the measured Bloch vector length of memory qubit eigenstates (triangles) and superposition states (circles) versus entanglement attempts, for $\alpha = 0.05$. Orange represents measured superposition decay versus time in the absence of entanglement attempts. Solid lines are fits, yielding decay constants of $N_{1/e} = 1843 \pm 32$ (2042 ± 36) with (without) entanglement generation attempts (see table 3.3 and section 3.5 for additional details).

This higher field puts much stricter demands on the relative field stability in order to not affect the qubit frequencies; we achieve an order of magnitude reduction in field fluctuations by actively stabilizing the temperature of the sample holder, which in turn

stabilizes the permanent magnet inside the cryostat (section 3.5). Additionally, the higher magnetic field splits the two optical transitions used for electronic spin initialization, hindering fast qubit resets; the addition of a second initialization laser, frequency locked to the first one with an offset of 480 MHz, enables us to maintain high-fidelity (> 0.99) and fast (few microsecond) resets (section 3.5).

We measure the fidelity of stored states on Bob's memory qubit for a varying number of entanglement generation attempts (fig. 3.3). The two eigenstates ($\pm Z$) do not show appreciable decay as we increase the number of entanglement generation attempts, as expected from the pure dephasing nature of the process [30]. The superposition states degrade with an average decay constant of $N_{1/e} \approx 1800$ attempts. To gain insight into the contribution of network activity to this decay, we repeat these measurements in the absence of entanglement attempts, in which case dephasing of the memory qubit is mainly due to uncontrolled interactions with nearby nuclear spins. We find this intrinsic dephasing time to be $T_2^\star = 11.6(2)$ ms, equivalent to the duration of ≈ 2000 entanglement generation attempts. We conclude that the intrinsic dephasing accounts for most of the decay observed under network activity, indicating the desired robustness. For the experiments discussed below, we use a timeout of 450 attempts before the sequence is restarted, as a balance between optimizing entanglement generation rate and fidelity of the stored state.

Executing protocols over quantum networks requires real-time feed-forward operations among the various nodes: measurement outcomes at the heralding station or at nodes need to be translated into quantum gates on other nodes. We implement an asynchronous bidirectional serial communication scheme between microcontrollers at the nodes, enabling both the required timing synchronization of the nodes and the exchange of feed-forward information for the quantum network protocols (section 3.5). Furthermore, we integrate the feed-forward operations with local dynamical decoupling protocols that actively protect the communication qubits from decoherence. The resulting methods enable us to run multinode protocols in a heralded fashion: *Flag* signals indicate in real time the successful execution of (sub)protocols and generation of desired states that are then available for further use, thus critically enhancing the efficiency and removing the need for any postselection.

3.3 Demonstration of multinode network protocols

We now turn to the full operation of the three-node network that combines the different elements discussed above. We perform two canonical network protocols: the distribution of genuine multipartite entanglement and entanglement swapping to two non–nearest-neighbor nodes.

In both protocols, the sequence depicted in Fig. 3.4A is used to establish a remote entangled state on each of the two links. This sequence starts with a preparation step (depicted only in fig. 3.15) that synchronizes the microcontrollers of the nodes and makes sure that the NV centers in each node are in the desired charge state and in resonance with all the relevant lasers. After initialization of the memory qubit, the first entangled state is prepared on the link Alice-Bob. We interleave blocks of entanglement generation attempts with phase-stabilization cycles. Once Alice-Bob entanglement is heralded, Alice's entangled qubit is subject to a dynamical decoupling sequence while awaiting further communication from the other nodes. At Bob, deterministic quantum logic is used to swap

the other half of the entangled state to the memory qubit.

The second part of the phase stabilization is then executed, followed by the generation of remote entanglement between the communication qubits of Bob and Charlie. In case of a timeout (no success within the preset number of attempts), the full protocol is restarted. In case of success, a dynamical decoupling sequence is started on Charlie's communication qubit analogous to the protocol on Alice. At Bob, a Z-rotation is applied to the memory qubit to compensate for the acquired phase that depends linearly on the (a priori unknown) number of entanglement attempts. This gate is implemented through an XY4 decoupling sequence on the communication qubit, with a length set in real time by the microcontroller based on which entanglement attempt was successful (section 3.5). After this step, the two links each share an entangled state ready for further processing: one between the communication qubit at Alice and the memory qubit at Bob and one between the communication qubits of Bob and Charlie.

The first protocol we perform is the generation of a multipartite entangled GHZ state across the three nodes. The circuit diagram describing our protocol is depicted in Fig. 3.4B. We first entangle the two qubits at Bob, followed by measurement of the communication qubit in a suitably chosen basis. The remaining three qubits are thereby projected into one of four possible GHZ-like states, which are all equivalent up to a basis rotation. The specific basis rotation depends both on the measurement outcome at Bob and on which Bell states ($|\Psi^+\rangle$ or $|\Psi^-\rangle$) were generated in the first part of the sequence, which in turn depends on which two photon detectors heralded the remote entangled states. These outcomes are communicated and processed in real time and the corresponding feed-forward operations are applied at Charlie. As a result, the protocol is able to achieve delivery of the same GHZ state $|\text{GHZ}\rangle_{\text{ABC}} = (|000\rangle + |111\rangle)/\sqrt{2}$, irrespective of the intermediate outcomes. Here, we choose to herald only on Bob reporting the $\|0\rangle$ readout outcome, because the asymmetry in the communication qubit readout fidelities renders this outcome more faithful (section 3.5). Additionally, this choice automatically filters out events in which the NV center of Bob was in the incorrect charge state or off resonance (occurrence $\approx 10\,\%$ in this experiment; see section 3.5). With this heralding choice, the protocol delivers GHZ states at a rate of about 1/(90s).

We extract the fidelity to the ideal GHZ state from correlation measurements by using $F = (1 + \langle IZZ\rangle + \langle ZIZ\rangle + \langle ZZI\rangle + \langle XXX\rangle - \langle XYY\rangle - \langle YXY\rangle - \langle YYX\rangle)/8$ and find $F = 0.538(18)$ (fig. 3.4C). The state fidelity above 0.5 certifies the presence of genuine multipartite entanglement distributed across the three nodes [31].

In this experiment, the fidelities of the entangled states on the elementary links bound the fidelity of the heralded GHZ state to about 0.66. Other relevant error sources are the dephasing of the memory qubit and accumulation of small quantum gate errors (see table 3.4). We emphasize that, contrary to earlier demonstrations of distributed GHZ states with photonic qubits [25] and ensemble-based memories [26] that relied on postselection, we achieve heralded GHZ state generation: a real-time heralding signal indicates the reliable delivery of the states.
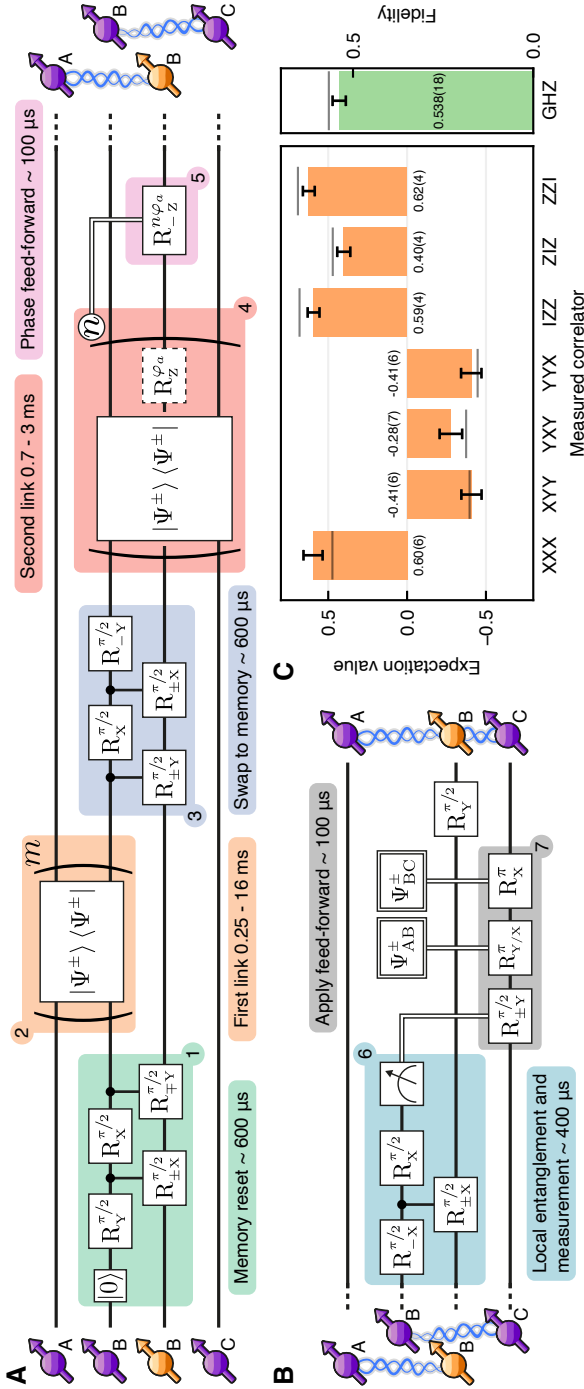
Figure 3.4: Distribution of genuine multipartite entanglement across the quantum network. (A) Circuit diagram displaying the experimental sequence used to establish entanglement on both elementary links. (B) Circuit diagram displaying the experimental sequence for distributing a three-partite GHZ state across the three nodes. (C) Outcomes of correlation measurements and the resulting fidelity of the heralded GHZ state, demonstrating genuine multipartite entanglement. Gray bars depict values from the theoretical model. Error bars indicate one standard deviation.

3

The second protocol, illustrated in fig. 3.5A, demonstrates entanglement swapping of the two direct links into an entangled state of the outer two nodes. Once entanglement is established on the two links as described above, the central part of the entanglement swapping is executed: Bob, the central node, performs a Bell state measurement (BSM) on its two qubits. One way to read this protocol is that the BSM induces teleportation of the state stored on Bob's memory qubit to Charlie, by consuming the entangled state shared by Bob's communication qubit and Charlie. Because the state teleported to Charlie was Bob's share of an entangled state with Alice, the teleportation establishes direct entanglement between Alice and Charlie.
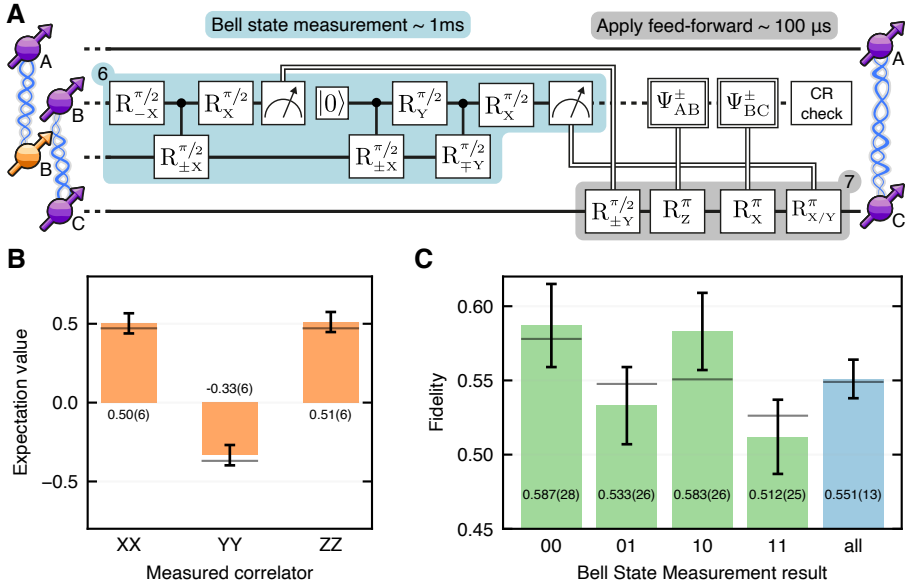
Figure 3.5: Entanglement swapping on a multinode quantum network. (**A**) Circuit diagram displaying the experimental sequence for entanglement swapping, yielding an entangled state shared between the two nonconnected nodes. (**B**) Outcomes of correlation measurements on the heralded entangled state shared between Alice and Charlie for the selected Bell-state measurement outcome (see main text). (**C**) State fidelities for different outcomes of Bob's Bell-state measurement (green) and the state fidelity averaged over all outcomes (blue). In (B) and (C), gray bars depict values from the theoretical model, and error bars indicate one standard deviation.

After the BSM is completed, we perform a charge and resonance (CR) check on Bob to prevent heralding on events in which the NV center of Bob was in the incorrect charge state or off resonance. We note that this CR check was not used in the heralding procedure of the GHZ generation protocol because its current implementation induces decoherence on Bob's memory qubit, which is part of the final GHZ state to be delivered. To complete the entanglement swapping, feed-forward operations are performed at Charlie to account in real time for the different measurement outcomes, analogous to the previous protocol, resulting in the delivery of the Bell state $|\Phi^+\rangle_{AC} = (|00\rangle + |11\rangle)/\sqrt{2}$.

We assess the performance of the entanglement swapping by measuring three two-node correlators on the generated Bell state shared by Alice and Charlie. Because the

BSM is performed with local quantum logic and single-shot readout, it is (except for the CR check step) a deterministic operation. However, given the asymmetry in the readout errors as discussed above, the fidelity of the final state will depend on the readout outcomes. Fig. 3.5B shows the results of the correlation measurements on the delivered state for heralding on Bob obtaining twice the outcome $|0\rangle$, yielding a state fidelity of $F = 0.587(28)$. Fig. 3.5C compares the state fidelities across the different BSM outcomes, displaying the expected lower fidelities for outcomes of $|1\rangle$ and an average fidelity over all outcomes of $F = 0.551(13)$. The combined heralding rate is $1/(40\text{s})$. The sources of infidelity are similar to the ones discussed above (see table 3.5). This experiment constitutes the first demonstration of entanglement swapping from previously stored remote entangled states, enabled by the network's ability to asynchronously establish heralded elementary entanglement links, to store these entangled states, and then to efficiently consume them to teleport entanglement to distant nodes.

3.4 Conclusion and outlook

We have demonstrated the realization of a multinode quantum network. We achieved multipartite entanglement distribution across the three nodes and any-to-any connectivity through entanglement swapping. It is noteworthy that the data acquisition for the network protocols has been performed fully remotely because of the COVID-19 pandemic, highlighting the versatility and stability of our architecture. Near-term advances in the capabilities and performance of the network will be driven by further reducing the infidelities of the elementary links (section 3.5), by adding new subprotocols such as control methods [29], decoupling sequences [30], and repetitive readout [32] for the nuclear spin qubits; by improved photonic interfaces to enhance the entangling rates [33–35]; and by improved control over the charge state of the NV center [36].

Our results open the door to exploring advanced multinode protocols and larger entangled states, for instance, by extending the local registers at the nodes. We note that a fully controlled 10-qubit register has recently been demonstrated on a similar device [29]. Furthermore, the network provides a powerful platform for developing and testing higher-level quantum network control layers [37–39], such as the recently proposed link layer protocol for quantum networks [40]. Quantum frequency conversion of the NV photons [41] can be used to interface the network nodes with deployed telecom fiber, paving the way to near-term quantum network tests over metropolitan distances. Finally, we expect the methods developed here to provide guidance for similar platforms reaching the same level of maturity in the future [42–45].

3.5 Supplementary information

3.5.1 Experimental setup

Our experiments are performed on three quantum network nodes. Each node houses a Nitrogen-Vacancy (NV) center in a high-purity type-IIa chemical-vapor-deposition diamond cut along the $\langle 111 \rangle$ crystal orientation (Element Six). All three samples have a natural abundance of carbon isotopes. Fabrication of solid immersion lenses and an anti-reflection coating on the diamond samples enhances the photon-collection efficiencies from the NV centers. The samples are housed in home-built cryogenic confocal microscope setups at 4 K. Experimental equipment used for each node is summarized in Table 3.1. In the following we use the letters A, B and C to identify nodes Alice, Bob and Charlie. The numbers 0, 1 refer to the computational basis $|0\rangle, |1\rangle$. Node A is in a different laboratory than nodes B and C, 7 m away. Nodes B and C are on the same optical table, approximately 2 m apart, see also Fig. 3.1A of the main text. The optical fiber that connects A with B is 30 m long, while the one that connect B to C is 2 m long.

Fig. 3.6 depicts the optics used to deliver and collect light to each sample. For phonon-sideband (PSB) detection, a dichroic mirror (Semrock) and an additional long-pass filter (Semrock) are used to block reflections of the excitation lasers. Photon emission is detected via an avalanche photo-diode (APD, Laser components, quantum efficiency approximately 80 %), with a total collection efficiency of approximately 10 %. For zero-phonon line (ZPL) detection, we isolate the single photons first with a narrow bandpass filter (5 nm, Semrock), then by blocking the reflected excitation light via two polarising beam-splitters (Thorlabs and Semrock). Spatial mode shaping via a deformable mirror (Boston Micromachines) enhances coupling to a polarization-maintaining single-mode fiber. The optical signals from each node are combined on in-fiber polarization-maintaining beam-splitters (Evanescent Optics). The final beam-splitter (where the single photons interfere) has an integrated fiber stretcher used for optical phase stabilization. Finally, the single photons are detected on superconducting nanowire single photons detectors (Photon Spot). They are optimized for 637 nm, have a detection efficiency >95 % and a dark count rate <1 Hz.

The level structures of the three nodes are depicted in Fig. 3.7. Each structure depends on local strain, electric fields and the applied magnetic field B. For nodes A and B the magnetic field is created with a permanent neodymium magnet inside the cryostat, which is located close to the sample and attached to the sample holder. The magnetic field is fine-tuned to be along the symmetry axis of the NV center using permanent neodymium magnets outside the cryostat. Node C has only a single permanent magnet outside the cryostat.

For optical excitation we set the laser frequencies (red arrows in Fig. 3.7) to the corresponding $^3A_2$ to ^3E transition. Spin-selective excitation of ZPL transitions ($\lambda = 637.25$ nm, $\omega = 2\pi \times 470.45$ THz) enables qubit readout ("Entangling" in Fig. 1B of the main text, $m_s = 0 \longleftrightarrow E_{x/y}$) and qubit reset via optical spin-pumping ("Reset" in Fig. 1B of the main text, $m_s = \pm 1 \longleftrightarrow E_{1,2}$). While at low field a single laser is sufficient to address both qubit reset transitions, in case of node B, which operates at 189 mT, we find a reset transitions splitting of 480 MHz. An additional laser is implemented in order to drive both reset transitions efficiently.

In order to tune the readout transitions of each NV center into resonance we employ the DC Stark effect via DC-biasing the strip-line that is used to deliver microwave (MW)

signals. The feedback sequence is analogous to the one used in Ref. [46]. Node B operates at 0 V tuning (it is grounded), and it uses non resonant charge reset with a green laser (515 nm). We observe small day to day drift in the readout frequency of node B that we attribute to slow ice build-up on the sample as the transition frequency can be brought back to its original value by a warm-up cool-down cycle. Nodes A and C, which use resonant charge reset with a yellow laser (575 nm), are brought into resonance with node B before starting a measurement.

The memory qubit of node B is the nuclear spin of a ^{13}C atom in the proximity of the NV center. Its electronic-spin-dependent precession frequencies are $\omega_0 = 2\pi \times 2025$ kHz and $\omega_{-1} = 2\pi \times 2056$ kHz, resulting in parallel hyperfine coupling of $A_\parallel \approx 2\pi \times 30$ kHz. The nuclear spin is controlled using dynamical decoupling sequences [19, 29]. The conditional $\pi/2$-rotations on the nuclear spin are performed with 56 decoupling pulses with an inter-pulse delay of $2\tau = 2 \times 2.818$ μs. Gate sequences to readout the memory qubit via the communication qubit are summarized in Fig. 3.9.

For synchronization purposes, the micro-controllers at each node (Jäger ADwin-Pro II T12) share a common 1 MHz clock.

To increase the on-off ratio of the AOM RF drivers, and therefore reduce unwanted light leakage, we use home-built fast (150 ns rise-time) RF switches, based on the HMC8038 (Analog Devices), to disconnect the RF drivers from the AOMs when no power should be delivered.

We terminate the MW delivery line on each cryostat with a home-built MW envelope detector, that allows us to see on an oscilloscope the microwave pulses being delivered to each sample. We use this for debugging purposes.

Parts that are not mentioned in the description above are the same as in Refs. [14, 19, 46].

3.5.2 Model of the generated states

Model and sources of error

The Python code to model all the generated states and to produce the figures in the main text can be found at [47]. The Jupyter notebooks that generate the figures make direct use of that code to plot the simulated states. The communication qubits of nodes A, B and C are encoded in the NV center electronic spin states $|0/1\rangle_A \equiv |m_S = 0/ + 1\rangle$, $|0/1\rangle_B \equiv |m_S = 0/ - 1\rangle$ and $|0/1\rangle_C \equiv |m_S = 0/ - 1\rangle$. The memory qubit of node B is encoded in the nuclear spin state of the addressed ^{13}C atom, $|0/1\rangle \equiv |m_I = \pm\frac{1}{2}\rangle$.

Regarding the generation of Bell states on the Alice-Bob and Bob-Charlie links, we extend the model presented in Ref. [14] to allow for different values of the parameters α in the two nodes. We find that to obtain maximum state fidelity the condition $\alpha_A p_A^{\text{det}} \approx \alpha_B p_B^{\text{det}}$ must hold, where $\alpha_{A,B}$ are the populations of the $|0\rangle$ state of each node and $p_{A,B}^{\text{det}}$ is the probability of detecting a photon emitted by the respective node in the detection window. The state that is heralded by the protocol is the following (assuming $p^{\text{det}} \ll 1$):

$$\rho_{AB}^{\pm} = \frac{1}{p_{\text{tot}}} \begin{pmatrix} p_{00} & 0 & 0 & 0 \\ 0 & p_{01} & \pm\sqrt{V p_{01} p_{10}} & 0 \\ 0 & \pm\sqrt{V p_{01} p_{10}} & p_{10} & 0 \\ 0 & 0 & 0 & p_{11} \end{pmatrix}, \tag{3.1}$$

$$p_{00} = \alpha_A \alpha_B (p_A^{\text{det}} + p_B^{\text{det}} + 2p_{\text{dc}}), \tag{3.2}$$

$$p_{01} = \alpha_A (1 - \alpha_B)(p_A^{\text{det}} + 2p_{\text{dc}}), \tag{3.3}$$

$$p_{10} = \alpha_B (1 - \alpha_A)(p_B^{\text{det}} + 2p_{\text{dc}}), \tag{3.4}$$

$$p_{11} = 2(1 - \alpha_A)(1 - \alpha_B)p_{\text{dc}}, \tag{3.5}$$

$$p_{\text{tot}} = p_{00} + p_{01} + p_{10} + p_{11} \tag{3.6}$$

where V is the visibility of the two-photon quantum interference, $p_{\text{dc}} \ll 1$ is the probability of detecting a dark count (or in general a non-NV photon) in the detection window, the $\pm$ sign depends on which detector clicked. The off-diagonal terms neglect the contribution due to the dark counts with respect to the contribution due to $p_{A,B}^{\text{det}}$, i.e. we assume $p_{\text{dc}} \ll p_{A,B}^{\text{det}}$.

If one assumes $p_{\text{det}} = p_A^{\text{det}} = p_B^{\text{det}}$, $p_{\text{dc}} = 0$, $\alpha = \alpha_A = \alpha_B$ and $V = 1$, then the fidelity of ρ_{AB} with the closest Bell state is $F = 1 - \alpha$, and the generation rate is $r_{AB} = 2\,\alpha\,p_{\text{det}}\,r_{\text{attempt}}$, with r_{attempt} the attempt rate.

Additional sources of infidelity are uncertainty in the phase of the entangled state and double excitation. See Ref. [14] for details on how they are modeled. We summarize in Table 3.2 the infidelity arising from the aforementioned sources, which is reasonably in agreement with the measured state fidelities. The Bell states between Alice and Bob were generated with $\alpha_A, \alpha_B = 0.07, 0.05$, while the ones between Bob and Charlie with $\alpha_B, \alpha_C = 0.05, 0.10$. These values have been chosen as a trade-off between protocol success rate and fidelity.

To model the states generated in the two demonstrated protocols (GHZ state between Alice Bob and Charlie, and Bell state between Alice and Charlie) we take into account:

- the Bell states generated between Alice and Bob and between Bob and Charlie,

- the dephasing of the nuclear spin during the entanglement generation between Bob and Charlie,

- depolarising noise on the nuclear spin that combines initialisation, swap and readout error,

- communication qubit readout errors at Bob that would generate a wrong feed-forward operation at Charlie,

- the depolarising noise on the communication qubits of Alice and Charlie during their dynamical decoupling sequences,

- the possibility that Alice and/or Charlie are in the wrong charge state (NV^0) at the end of the sequence.

Table 3.4 and Table 3.5 summarize the error budget of the generated states. In the case of the entanglement swapping we also report the expected infidelity when accepting any Bell state measurement (BSM) result.

In the protocol demonstrating GHZ distribution across the three nodes we chose to herald on readout outcome "0" only, as explained in the main text. Our model predicts that GHZ states heralded on measurement outcome "1" would have had an additional 3 % infidelity (due to the asymmetry in the electron readout infidelities).

Future improvements to the remote entanglement links

The remote entanglement fidelity can be increased in the near term by improving phase stabilization at Alice to a similar level as Bob and Charlie, by lowering the double-excitation probability through a reduction of the optical excitation pulse width, and by using similar devices with higher collection efficiency [46] allowing for operation at smaller α. Additionally, an improvement in fidelity of approximately $\eta_{\mathrm{PSB}}(\alpha + p_{2\mathrm{e}})$ (where η_{PSB} is the probability to collect a photon in the phonon sideband emission and $p_{2\mathrm{e}}$ is the probability of double excitation during the optical pulse) can be obtained by rejecting heralding events for which simultaneously a photon was detected in the phonon sideband emission channel on one of the nodes. In the same way, this filtering will reduce the errors due to double excitation during the optical pulse as well as errors due to heralding on dark counts. This rejection could be implemented in real-time using an additional FPGA on each node. The combination of these improvements would bring the entanglement fidelity above 0.90.

Going beyond the current hardware, the entanglement generation rates may be increased by up to two orders of magnitude by enhancing the collection of coherent NV photons through the use of optical cavities [33, 35]. Alternatively, the development of other color centers in optical waveguides and/or cavities may bring a similar improvement in rates [42–45].

3.5.3 Phase stabilization

Inherent to an implementation where active phase stabilization is interleaved with free evolution time, there is a trade-off between phase stability (or fidelity of the entangled state) and the free evolution time (which is the time used for entanglement generation). The more often the system is stabilized, the higher the bandwidth of the stabilization and the lower the final uncertainty in $\Delta\theta$ will be.

Our previous implementation used a single homodyne phase detection scheme [14]. While that method allows for the stabilization of the phase of the entangled state, there are several aspects that can be improved; The small fraction of excitation light that is reflected from the diamond surface is partially coupled in the single-photon detection path. By measuring the interference signal after the beam-splitter at the heralding station it is possible to obtain the phase of the interferometer. But leaking some of the reflected excitation light into the single-photon path for phase stabilization purposes increases the chance that, during entanglement generation, some of the reflected excitation light will be detected and mistakenly herald an entangled state. To counteract this effect, the amount of leaked light was somewhat minimized by polarization selection (but never completely, since some light is needed to detect a phase signal) and long integration times (24 ms) were used during phase detection, reducing the phase stabilization bandwidth. Furthermore, exposing the NV center to a relatively long and strong laser pulse makes it more susceptible to spectral jumps and ionization.

To solve these challenges we devised and implemented a new phase stabilization scheme that combines higher bandwith and optimal rejection of the excitation light from the single-photon paths, while maintaining robustness against power level fluctuations and scalability to a higher number of nodes.

Phase detection methods

In a homodyne phase detection scheme the light has the same frequency in both arms of the interferometer. Depending on the optical phase difference $\Delta\theta$, light will constructively or destructively interfere on the output ports of the beam-splitter. Assuming common polarization and perfectly overlapping spatial modes, the intensity $I_{3,4}$ in the output ports is

$$I_{3,4} = I_1 + I_2 \pm 2\sqrt{I_1 I_2}\cos\Delta\theta. \tag{3.7}$$

For known input levels $I_{1,2}$, $\Delta\theta$ can be calculated from the difference in intensity in the output ports of the beams-splitter. Fluctuations in the intensity of the input signals will lead to an error in the phase measurement, except for the case $\cos\Delta\theta = 0$ which gives $I_3 = I_4$ independent of the input intensity.

In a heterodyne phase detection scheme the light has different frequencies in the two arms of the interferometer. Again, assuming common polarization and perfectly overlapping modes, the light will interfere in the output ports resulting in a signal with amplitude

$$I_{3,4} = I_1 + I_2 \pm \sqrt{I_1 I_2}\big(\cos((\omega_1 - \omega_2)t - \Delta\theta) + \cos((\omega_1 + \omega_2)t + \Delta\theta)\big) \tag{3.8}$$

where $\omega_{1,2}$ are the angular frequencies of the light. When we pick a relatively small frequency difference, $(\omega_1 - \omega_2)/2\pi \approx 10\,\mathrm{MHz}$, we can ignore the last term in Eq. 3.8 and the resulting 10 MHz beat signal can be measured with a photodiode and efficiently filtered from the DC background signal (the last term of Eq. 3.8 will have a frequency in the optical domain and will not be picked up by the photodiode due to the limited bandwidth). The phase of this beat signal corresponds to the optical phase difference in the two paths. Since the phase information is not translated to the amplitude of the beat signal, fluctuations in the input intensity will not cause an error in the measurement. Moreover, this method is very suitable to measure small signals: if the signal is very small in one of the arms, the amplitude of the beat signal can be increased by increasing the intensity in the other arm.

Splitting the interferometer in parts

In the experiments with three quantum nodes we have two effective interferometers that share part of their optical paths. We split the interferometers into six parts, see Figures 3.10, 3.11. In total there are four local interferometers and two global interferometers, where the local interferometer comprises the excitation path and free space optical path close to the cryostat of each node and the global interferometer includes the fibers connecting the nodes to the central beam-splitter. With the measured phase, an error signal is computed and feedback is applied to the optical path, either with a mirror on a piezoelectric element or a fiber stretcher.

The two global interferometers, using homodyne phase detection, stabilize the optical path to the beam-splitter and single photon detectors used for entanglement heralding. Since the detectors are shared for the two entanglement links, the optical phase measurement for the two global interferometers has to be multiplexed in time. The local interferometers are stabilized using heterodyne phase detection. The excitation light (the same we use for the optical excitation pulse that generates spin-photon entanglement) is reflected off the diamond surface and since it has (close to) orthogonal polarization with the

NV centers emitted photons it can be separated from the single photons using a polarizing beam-splitter (PBS). Afterwards, the weak reflected excitation pulse interferes with a strong laser pulse from the other arm with the frequency offset. The beat signal is measured with a photodiode and the optical phase difference is extracted using an electronic reference signal. The middle node has two local interferometers, one for each link. When all separate interferometers are stabilized, the paths of the excitation light and the single photons used for entanglement heralding will be phase stable.

Technical description of the local interferometer
For all the local interferometers we use a heterodyne phase detection scheme. A diagram of the optics and electronics is plotted in Fig. 3.12. For each entanglement link (Alice-Bob and Bob-Charlie) the phase and excitation light are provided by the outer nodes (Alice and Charlie). To generate the known 10 MHz frequency offset between the light paths, we take advantage of the acousto-optic modulators (AOMs) we use to generate light pulses. By driving two AOMs at respectively 200 MHz and 210 MHz, we establish the required frequency difference between the light paths. Part of the RF signals used to drive the AOMs are tapped off and combined in a mixer to obtain an electronic reference signal. The light from the AOMs is launched in a free space path with several optical elements. The first polarizing beam-splitter (PBS) ensures the phase light to be linearly polarized. The second PBS separates the reflected excitation light from the single photons. At this point the phase-reference light and the reflected excitation light have orthogonal polarization. The waveplates in front of the third beam-splitter rotate their polarization such that they can be interfered on the third PBS. This interference leads to a beating signal that can be detected with the photodiode. Consequently, the beating signal is filtered, amplified and, together with the electronic reference signal, used as inputs for the phase detector (Mini-Circuits ZRDP-1+). The output of the phase detector is filtered and impedance matched to an analog to digital converter (ADC) input of the micro-controller, the ADwin.

Timings
The phase stabilization requires synchronization between the different nodes. Nodes A and C provide the phase and excitation light, but all nodes measure the phase of at least one local interferometer. Some of the detectors used for the phase measurements are shared among different interferometers, so not all measurements can be done at the same time. Figure 3.13 shows how the various phase stabilization cycles are interleaved with entanglement generation time.

The choice of free evolution time is governed by the noise sources in the different parts of the system. The local interferometer of node A experiences noise with high frequency components (compared to the other intereferometers) hence the free running time must be short enough to achieve the necessary feedback bandwith. The duration of the preparation part of the experiment, which includes charge and resonance checks, as well as synchronization steps between nodes, can vary from approximately 50 μs to a few seconds. When the phase is completely scrambled due to a too long free running time, it is not possible to reach the set-point in a single feedback round. For this reason we start with multiple rounds of phase stabilization without any free evolution time in between.

Phase stability

To characterize the performance of the phase stabilization we look at three different aspects: the free evolution of the phase without any stabilization, the frequency spectrum of the noise and the distribution of the phase while actively stabilizing. All the results for the six interferometers are plotted in Figure 3.14 (see Fig. 3.11 for the labeling). The differences in performance can be explained by the noise sources present in our experimental lay-out. We identify two main sources of noise: the relatively noisy fiber connection between nodes A and B and the positioning stages of each node. The three nodes are built in two separates rooms and we use optical fibers (30 m) to connect node A to node B. All nodes have a microscope objective for optically accessing the diamond samples. On node B and C this microscope objective is mounted on a piezo-electric stage. For node A the design is different: here the sample is mounted on a piezo-stack and the microscope objective is fixed. All these piezo-electric stages are susceptible to the vibrations generated by the operation of the cryostats.

The sample stage of node A cause relatively-strong high-frequency (>500 Hz) noise; the microscope objective stage of nodes B and C cause lower-frequency noise and the optical fiber connection between nodes A and B causes relatively-strong low-frequency components. In the experimental sequence we interleave experimental time with rounds of phase stabilization. With the used timings (see Fig. 3.13) we are able to stabilize frequencies ≤ 500 Hz. Due to its relatively high-frequency components, the noise of the local interferometer of node A is the limiting factor in terms of phase stability of the overall apparatus. We expect that fixing the sample to the cold-finger of the cryostat, and only moving the microscope objective (like we do on nodes B and C) will allow us to lower the phase noise on node A in the future.

Entangled-state phase drifts

While the phase stabilization scheme allows us to access the entangled state generated by the single photon protocol by fixing the phase $\Delta\theta$, we observe that the phase of the generated entangled state undergoes small drifts on a timescale of hours. That is, even though all the interferometers are stabilized to the same value, the phase of the entangled state will slowly drift by ≈10 deg/h. We hypothesize that these drifts are due to the relative position of the microscope objective and the NV center: while the light used for phase stabilization is reflected off the diamond surface, the NV-emitted photons are generated inside the diamond. Small changes in distance and angle of the microscope objective would not lead to observable differences in the fluorescence measurement we use for position optimization, but may slightly alter the path the photons have to travel. To solve this challenge, after every position optimization (≈ once every hour), we re-calibrate the phase of the generated entangled state (≈ 5 minute measurement per link). More robust positioning systems (both for the sample and the microscope objective) may reduce the phase drifts and alleviate the need for entangled-phase re-calibration

3.5.4 Single-shot readout correction

We correct tomography-related single-shot readouts (SSROs) for known error in order to obtain a reliable estimate of the actual generated states.[1]

Single qubit case

For a single qubit:

$$\vec{m} = \hat{R}\,\vec{p}, \tag{3.9}$$

where $\vec{p} = (p_0, p_1)^T$ is the (column) vector of expected populations, $\vec{m} = (m_0, m_1)^T$ is the (column) vector of measured populations, and

$$\hat{R} = \begin{pmatrix} r_{00} & r_{01} \\ r_{10} & r_{11} \end{pmatrix} = \begin{pmatrix} F_0 & 1 - F_1 \\ 1 - F_0 & F_1 \end{pmatrix}$$

is the SSRO operator that connects the two. For example:

$$m_0 = F_0 p_0 + (1 - F_1)p_1,$$

i.e. the measured population in $|0\rangle$ is given by the correctly assigned population in $|0\rangle$ plus the incorrectly assigned population in $|1\rangle$. From Eq. 3.9 it follows that:

$$\vec{p} = \hat{R}^{-1}\,\vec{m}, \tag{3.10}$$

which is what we use in practice to apply the readout correction. This allow us to obtain the vector of expected populations given the measured populations and the SSRO error operator. Experimentally we cannot directly measure $\vec{m}$. We measure events in which the communication qubit is either in $|0\rangle$ or in $|1\rangle$. We repeat this process N times, obtaining N_0 times the outcome $|0\rangle$ and N_1 times the outcome $|1\rangle$. From this we estimate the measured populations $\vec{m}$:

$$m_0 = N_0/N,\ m_1 = N_1/N \tag{3.11}$$

The probability distribution of the number of events N_0 is a Binomial distribution with expected value $N m_0$ and variance $N m_0(1 - m_0)$. From this it is possible to calculate the experimental value and uncertainty for m_0 (and m_1):

$$m_0 = N_0/N \tag{3.12}$$

$$m_1 = 1 - m_0 \tag{3.13}$$

$$\sigma_{m_0} = \sigma_{m_1} = \sqrt{\frac{m_0}{N}(1 - m_0)} \tag{3.14}$$

The covariance between m_0 and m_1 is:

$$\mathrm{Cov}(m_0, m_1) = -\frac{m_0(1 - m_0)}{N} \tag{3.15}$$

---

[1] The method described here is what was used for the analysis of the data in the publication (Ref. [48]). In Appendix A of this thesis (page 141) a new method is described which is to be preferred to the one described here.

Once Eq. 3.10 has been calculated it is possible to evaluate the expectation value of $\vec{p}$ and its uncertainty. In the one qubit scenario it is easy to invert the expression analytically:

$$p_0 = \frac{F_1 \, m_0 + (F_1 - 1) \, m_1}{F_0 + F_1 - 1} = \frac{F_1 + m_0 - 1}{F_0 + F_1 - 1} \tag{3.16}$$

$$p_1 = \frac{(F_0 - 1) \, m_0 + F_0 \, m_1}{F_0 + F_1 - 1} = \frac{F_0 - m_0}{F_0 + F_1 - 1} = 1 - p_0 \tag{3.17}$$

$$\sigma_{p_0} = \sigma_{p_1} = \frac{\sigma_{m_0}}{F_0 + F_1 - 1} \tag{3.18}$$

and it is straightforward to propagate uncertainties in $F_{0/1}$ to p_0 and p_1.

Two and three qubit case

For two (and more) qubits, the measurement outcomes will be distributed according to a Multinomial distribution (as opposed to a Binomial). While the expectation values of $p_0, ..., p_i$ can still be computed analytically relatively straightforwardly, their uncertainties need to take into account the non-trivial covariances in the m_i. Additionally, taking into account uncertainties in the $F_{0/1}$ makes the error propagation even more tedious. We therefore use a Monte Carlo simulation that takes into account the Multinomial distribution as well as the $F_{0/1}$ of each setup to estimate uncertainties on the correlation measurements and the state fidelities, without having to assume normality of the data. The code to run the Monte Carlo simulation is included in the Jupyter notebooks that produce the figures of the main text [47].

3.5.5 Phase feed-forward on the memory qubit

The nuclear spin memory qubit of Bob precesses at a frequency that depends on the spin state of the electronic spin (the communication qubit). Throughout the experimental sequence we keep track of the phase acquired by the nuclear spin to be able to readout and apply gates in the correct bases. While most operations are deterministic in time (nuclear spin initialisation, gates on the electronic spin, etc.) and the phase evolution of the nuclear spin can be calculated in advance, entanglement generation is a probabilistic process. This means that it is not known in advance how long the entanglement operation (number of entanglement attempts) is going to take, and therefore how much phase the nuclear spin is going to acquire. To solve this challenge, we implement a phase feed-forward mechanism that applies a Z-rotation to the nuclear spin that cancels this acquired phase. Since the used Arbitrary Waveform Generator (AWG) only has limited real-time programming capability, we implement this mechanism via a real-time interaction between our node micro-controller (ADwin) and the AWG. Once entanglement is heralded between Bob and Charlie, the AWG of Bob jumps out of the entanglement generation subroutine and starts an XY4 decoupling sequence on the communication qubit. During this XY4, the AWG interacts with the ADwin of Bob (which has recorded how much phase the nuclear spin has acquired during the entanglement operation) to select, via a binary decision tree, the time in between microwave pulses. The binary decision tree allows us to vary the (additional) duration of the XY4 element in steps of 2 ns up to 512 ns, which is more than a 2π precession for the nuclear spin ($\tau_L = 490$ ns, feed-forward resolution $\approx 1.5°$). Regardless of

the inter-pulse time selected, the communication qubit will be decoupled. Consequently, the needed additional phase to re-phase the nuclear spin can be conveniently set via the length of the XY4 sequence. We pre-compile the timings that the ADwin will communicate to the AWG to reduce the computational load on the ADwin. We anticipate that an AWG with an integrated programmable FPGA will be able to completely take over the task of phase tracking without need for interaction with the node micro-controller, reducing experimental overhead.

3.5.6 Memory qubit lifetime and increased magnetic field

An important resource in our experiments is the ability to store entanglement in the nuclear spin memory qubit of Bob while performing further operations on the node. While we have implemented methods to keep track and actively compensate for the phase acquired during entanglement generation (see previous section), additional dephasing may occur. The major source of nuclear spin dephasing during entanglement generation was found to be [30] failed electronic spin control (initialization errors or MW pulse errors).

An entanglement attempt can be broken into the following pieces: communication qubit reset (via optical pumping), MW pulse that creates the communication qubit superposition (named in the following the α pulse), optical excitation pulse that creates the spin-photon entanglement, and a decoupling MW π pulse. The time τ between the α and the de-coupling pulse is chosen such that it equals the time between the decoupling pulse and the average reset time in the subsequent entanglement attempt (see Ref. [30] for details). This ensures that regardless of its initial state, the communication qubit spends an equal amount of time in the $|0\rangle$ and $|1\rangle$ states. However, an error in the MW π pulse will result in an unknown acquired phase on the nuclear spin and lead to dephasing. Previous work [30] suggested that such dephasing can be mitigated when working at a higher magnetic field, which allows for a shorter spacing between subsequent MW pulses.

In order to work at higher fields we have installed a stronger permanent magnet inside the cryostat of Bob reaching a field of 189 mT at the location of the NV center. At such fields, temperature fluctuations of the magnet, mainly due to the MW pulses applied to the sample, can result in a significant change of the magnetic field amplitude. Hence, we stabilize the sample holder via an active feedback loop, ensuring a stable temperature of the permanent magnet. We reach a stability of 1 μT, which results in a maximum variation of the nuclear spin precession frequency of ≈ 10 Hz, one order of magnitude below the dephasing rate due to interactions with other spins in its environment.

These improvements allow us to shorten the interpulse spacing to 942 ns, limited by the waiting time after the optical excitation pulse that we need to include in order to allow the AWG to respond in real time to a successful entanglement attempt and jump out of the entangling sequence. As Fig. 3.3 of the main text shows, a similar nuclear memory lifetime is observed when applying entanglement attempts or when idling. This shows that the lifetime of the memory qubit, in our magnetic field regime, is mainly limited by natural dephasing and not by electronic spin control errors. We fit the two decays with the following function:

$$f(N) = A \exp\left(-\left(\frac{N}{N_{1/e}}\right)^n\right),$$

(3.19)

with N the number of entanglement generation attempts, $N_{1/e}$ the N at which the Bloch

vector length has decayed to $1/e$ of its initial value A, and n the exponent of the decay. The results of the fit are reported in Table 3.3. For the results *Without ent. gen.* the entanglement generation attempt is replaced by the equivalent free evolution time.

3.5.7 Microwave pulse fidelity

Errors in the MW pulses can limit the control of the communication qubit as well as induce decoherence on the nuclear spin memory qubit [30]. We use Hermite MW pulse envelopes [29] to perform rotations of the communication qubit spin:

$$h(t) = \left(1 - p\left(\frac{t}{T}\right)^2\right) e^{-\left(\frac{t}{T}\right)^2},$$

(3.20)

where p affects the shape of the pulse and T changes the length of the pulse. The pulses get distorted by the transmission line before they get to the sample. We apply a linear pre-distortion in frequency domain to compensate part of the error via the following IQ signals:

$$I = a \cdot h(t)$$

(3.21)

$$Q = ab\frac{t}{\pi T^2}\left(p + 1 - p\left(\frac{t}{T}\right)^2\right) e^{-\left(\frac{t}{T}\right)^2},$$

(3.22)

where a is the amplitude of the pulse and b is the skewness (slope) of the pre-distortion in frequency domain.

The MW π-pulses are calibrated by initializing the qubit in the $|0\rangle$ state, applying an odd number of consecutive pulses and reading out the final state. If the pulses were perfect one would measure $|1\rangle$ as outcome. The effect of the skewness on the pulse fidelity is investigated with a two dimensional scan; evaluating the fidelity for pulses with different amplitudes (a) and skewness (b). Figure 3.8 shows an example of such a scan, where it is clear one can calibrate a and b almost independently. We find that different set-ups require different levels of pre-distortion b, ranging from 10^{-11} to 10^{-8}. We estimate that the current errors of our MW pulses are between 0.1 % and 1 % for all the three nodes.

3.5.8 Classical communication

The three nodes can share information in several ways. The slowest method is based on Python socket interfaces between the measurement computers that allow us to share necessary values and information at a rate of approximately 10 Hz; this method is used for example to frequency lock the lasers, to coordinate calibrations on all nodes from a single computer and to share and record environmental data such as the temperature in the different laboratories. The second, and fastest, method is a direct connection between the micro-controller and the AWGs. This enables the triggering of all the AWGs from a single node, reducing jitter on the output waveforms. The third method is implemented on the micro-controllers and is used for the feed-forward operations across the nodes. Each micro-controller has one input and one output communication port (physically it is a normal digital input-output coaxial port). Bob, which receives signals from both Alice and Charlie, has a digital summing box (OR gate) at its input port, that combines the signals

coming from the other two nodes. We designed the experimental sequence such that it is clear who sent a specific message depending on when it arrives. Messages are sent over an off-the-shelf coaxial cable, using a serial communication scheme, with an average bit interval of 60 ns (the shortest the micro-controller can achieve). At the input port of each micro-controller, a fast edge detection (100 MHz) stores changes in the signal level (and the time at which they occur). It is therefore possible to reconstruct what pattern (i.e. message) was sent from one node to the other, directly on the micro-controller. Sending a message takes up to 300 ns (we send up to 5 bits at a time). Receiving and decoding take up to 2 µs combined. A flowchart of the communication steps between the micro-controllers used in the networks protocols demonstrated in the main text is shown in fig. 3.15 .

3.5.9 Feed-forward operations between nodes

We implement the feed-forward operations needed for our experimental protocols by combining the classical communication just discussed with a real-time pulse selection sequence by the micro-controller on the AWG. For both network protocols demonstrated we need to apply gates on the communication qubit of Charlie conditional on measurement outcomes at Bob. Once Bob has performed the required readout operations (on the communication qubit for the GHZ state generation or on both qubits for the entanglement swapping) it combines the readout results with the Bell states generation outcomes (i.e. which detectors clicked in the A-B and B-C entanglement generation) to obtain one of four possible feed-forward messages. Combining this information on Bob is an optimization of our communication resources; we could, alternatively, send the bits of information one by one to Charlie and combine the information there should that be a requirement of the protocol (for example in a blind quantum computation scenario). At this point a *FAIL* message could also be sent from Bob to all the nodes in order to abort the whole sequence, for example if the Bell State Measurement result is not the one that gives high-fidelity (see main text). We choose to not send *FAIL* messages and instead continue with the protocol to be able to assess the protocol performance for the less faithful Bell State Measurement outcomes (see Fig. 5C of the main text). In the meantime, Charlie has been applying an XY8 decoupling sequence to the communication qubit to protect its coherence while Bob performed the readout operations. Once Charlie receives the feed-forward information, its micro-controller starts a decision-tree sequence with its AWG to select the required microwave pulse-sequence. This decision tree is incorporated into an XY8 block of the AWG, such that the slow response time of the AWG (1 µs per bit of information) does not affect the coherence of the communication qubit. The microwave pulse sequence selected via the decision tree is appended to the aforementioned XY8 block. After the feed-forward operations are performed, the delivery of the states by the network protocol is completed. Finally, the delivered states are analyzed using a readout sequence (composed of an optional basis rotation and state readout).

3.5.10 Data acquisition and calibrations

The data supporting the protocol demonstrations in the main text (Figures 4C, 5B, 5C) was gathered in the month of October 2020. Due to the restrictions imposed by the COVID19 pandemic, we operated the setups remotely (from home) and went to the laboratories only when something needed in-situ intervention (like a broken power-supply).

The data has been collected in blocks of approximately 1 h, interleaved by calibration routines of approximately 20 min. For the GHZ state generation protocol we set the target number of data points at 2000. For the entanglement swapping we set the target number of data points at 4000. We stopped the experiment once the measurement block was completed in which the target number of data points was surpassed.

For the GHZ state generation we acquired 55 blocks over 10 days (effective measurement time $\approx$ 50 h), obtaining 2028 events, equivalent to a rate of $r_{\mathrm{GHZ}} \approx (90\,\mathrm{s})^{-1}$.

For the entanglement swapping demonstration we acquired 53 blocks over 7 days (effective measurement time $\approx$ 45 h), obtaining 853 events with BSM result "00", equivalent to a rate of $r_{\mathrm{swapping}} \approx (3\,\mathrm{minutes})^{-1}$. The other BSM results were: "01": 1030 events, "10": 1004 events, "11": 1168 events. The ratio of events between the BSM results matches the readout characteristics of node B: measured (expected) share of the events, $0.21 : 0.25 : 0.25 : 0.29$ ($0.23 : 0.25 : 0.25 : 0.27$). Combining all the BSM results we obtained a total of 4055 events, equivalent to a rate $r'_{\mathrm{swapping}} \approx (40\,\mathrm{s})^{-1}$.

Every three measurement blocks we performed a fidelity check on the entangled states between Alice-Bob and Bob-Charlie at the target α (total duration 20 min). These fidelity checks, combined over the GHZ and Entanglement Swapping datasets, are used for Fig. 3.2E of the main text. We performed a total of 58 fidelity checks, that combined generated: 24 197 Ψ_{AB}^{+} events, 25 057 Ψ_{AB}^{-} events, 26 383 Ψ_{BC}^{+} events and 27 459 Ψ_{BC}^{-} events. The asymmetry in the number of events between the Ψ^{+} and the Ψ^{-} states is due in part to the beam-splitter having a non ideal splitting ratio ($0.493 : 0.507$), to a slight difference in detector efficiencies ($\approx 1\,\%$) and to the brightnesses ($\alpha\,p_{\mathrm{det}}$) of the two setups involved not being completely balanced. The asymmetry between the number of events for the Ψ_{AB} and the Ψ_{BC} states is due to the different probability for node B to be in the wrong charge state (NV^{0}) at the end of the sequence for the two links. To obtain a reliable estimate of the fidelities of the Bell states, we discard events in which a CR (charge and resonance) check performed after readout gives a negative result. We remark that, as mentioned in the main text, we do not perform such an operation for the network protocols demonstrations, which are free from any post-selection. For the GHZ state generation, by heralding only on the $|0\rangle$ readout outcome of the communication qubit of Bob, we automatically reject events in which the NV center of Bob was either in the wrong charge state or off resonant. For the Entanglement Swapping demonstration, we perform a CR check after the Bell state measurement is performed on node B, and we herald success of the whole protocol only if this final CR check gives a positive result. We find that the test gives a positive result in approximately 90 % of the cases.

3.5.11 Experimental monitoring

Analogous to what reported in section J of the SI of Ref. [46], we implement checks while the experiment is running to ensure that the nodes are performing as expected. If one of the checks does not pass, we mark all future data to be disregarded (until the check is passed) and / or pause the experiment to perform further calibrations. Following is a list of all the checks that we use to mark future data to be disregarded (if they don't pass):

- Check that the measured phase of each interferometer is below 50° before the last piezo feedback is performed.

- Check that the number of photons collected during the qubit reset by optical pumping part of the entanglement generation sequence, averaged over the preceding second, is above a pre-set threshold. If the check does not pass within a matter of seconds, we pause the experiment and scan the laser frequency to find back the qubit reset transition frequency.

- Check that the number of photons collected during the spin-photon entanglement part of the entanglement generation sequence, averaged over the preceding second, is above a pre-set threshold. If the check does not pass within a matter of seconds, we pause the experiment and scan the bias voltage of the setup.

3
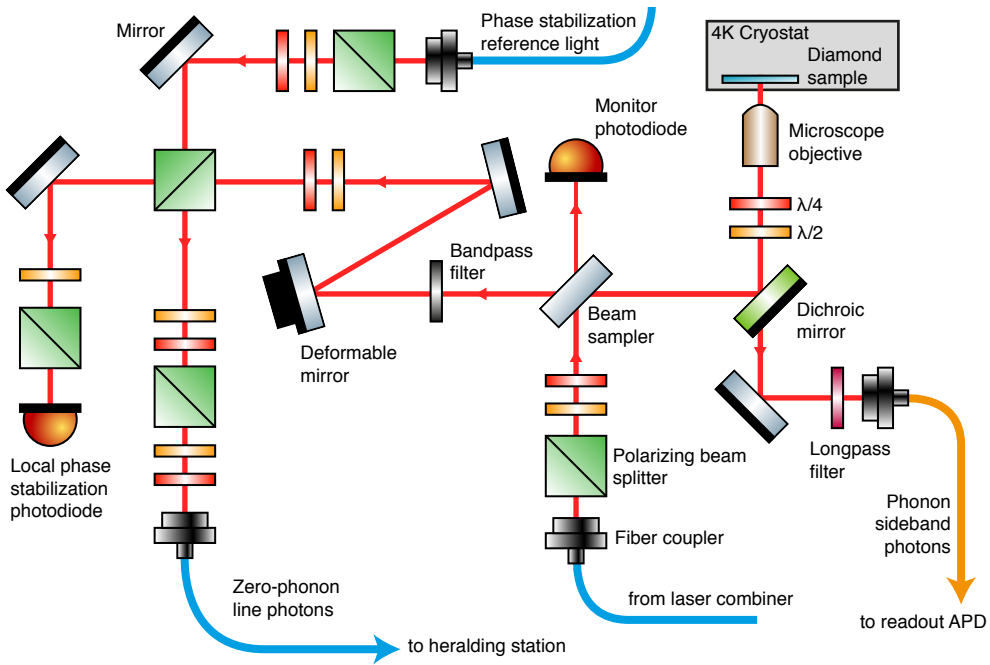
Figure 3.6: Schematic of the optics used for each node. The red lines indicate the optical path used both by the laser beams and the single photons. Blue fibers are single-mode polarization maintaining fiber. The orange fiber is a multi-mode fiber. The laser combiner (not depicted) combines, via beam-splitters and dichroic mirrors, the various laser beams and couples them into the single-mode fiber shown in the diagram. The laser combiner also includes a piezoelectric-mounted mirror that is used for the local phase stabilization feedback. The monitor photodiode records the 90 % of excitation light that goes through the beam-sampler (and that would otherwise be discarded). We monitor this signal on a digital oscilloscope connected to the measurement computers for debugging purposes.
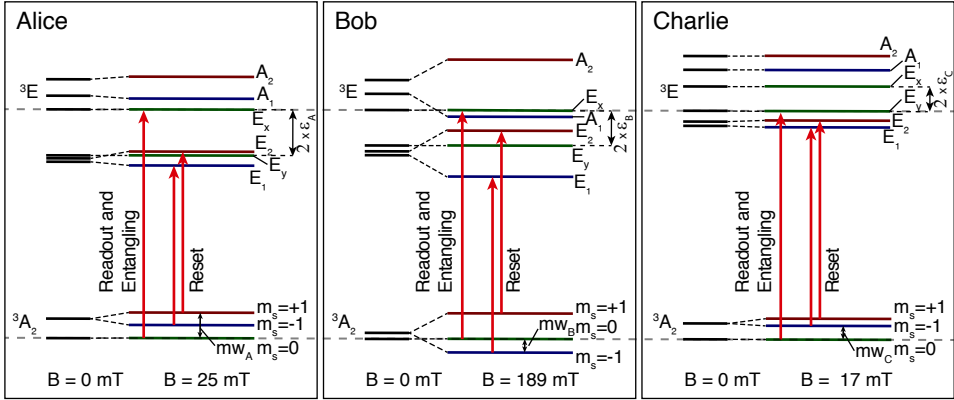
Figure 3.7: Level structure for the three NV centers. The optical transitions used within this work are indicated by red solid arrows. DC stark tuning brings all readout transitions to the same frequency, ensuring that the photons generated via the optical excitation pulse are indistinguishable. The spin state $m_S = 0/ + 1/ - 1$ of each level is indicated by color (green/red/blue). The communication qubits of nodes A, B and C are encoded in the NV center electronic spin states $|0/1\rangle_A \equiv |m_S = 0/ + 1\rangle$, $|0/1\rangle_B \equiv |m_S = 0/ - 1\rangle$ and $|0/1\rangle_C \equiv |m_S = 0/ - 1\rangle$. The memory qubit of node B is encoded in the nuclear spin state of the addressed ^{13}C atom, $|0/1\rangle \equiv |m_I = \pm\frac{1}{2}\rangle$.
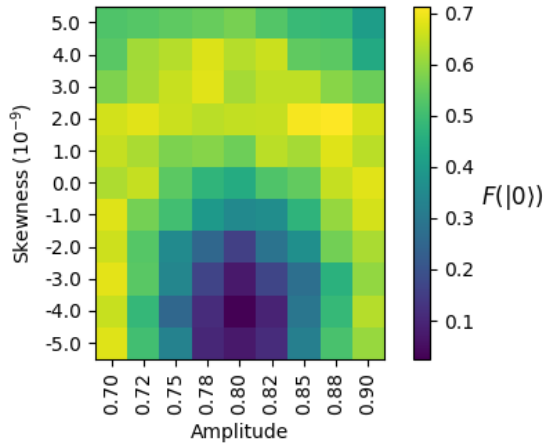
Figure 3.8: Calibrating the pre-distorted microwave (MW) Hermite pulses. The π pulses are calibrated by applying 11 sequential pulses: the probability of being in $|0\rangle$ at the end of the sequence is measured for different amplitudes (*a*) and skewness (*b*) of the Hermite pulse. The linear frequency pre-distortion allows us to achieve lower errors for the MW pulses.
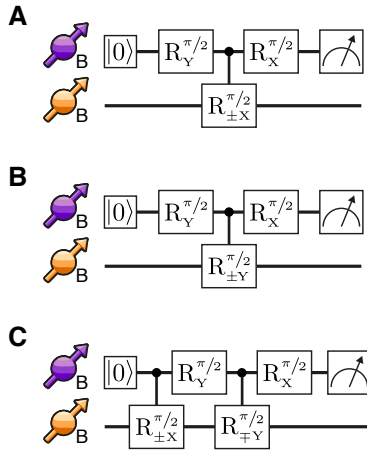
Figure 3.9: Memory qubit readout sequences. (A-C) Readout sequences of the nuclear spin memory qubit expectation values for $\hat{X}, \hat{Y}, \hat{Z}$ via the communication qubit. The controlled rotations are to be read as follows: $R_{\pm X}^{\pi/2}$ is a rotation of the memory qubit around the X axis with an angle of $\pi/2$ if the communication qubit is in $|0\rangle$, and with an angle of $-\pi/2$ if the communication qubit is in $|1\rangle$.
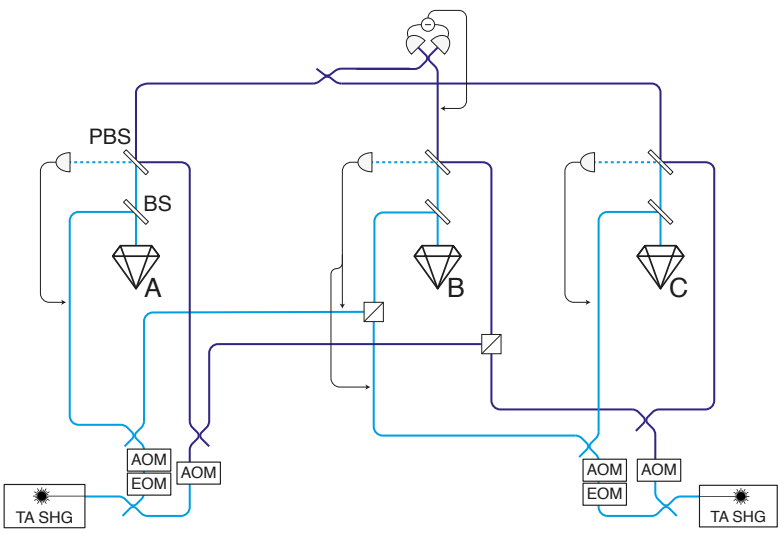
Figure 3.10: Diagram of the entire layout. Shown are the paths used by the excitation laser (solid light-blue lines) and the phase light (solid dark-blue lines), which has a frequency offset of $\approx 10\,\text{MHz}$ with respect to the excitation laser. The frequency offset is generated using different frequency modulation set-points for the acousto-optic modulators (AOM) in the excitation and phase path respectively.

3
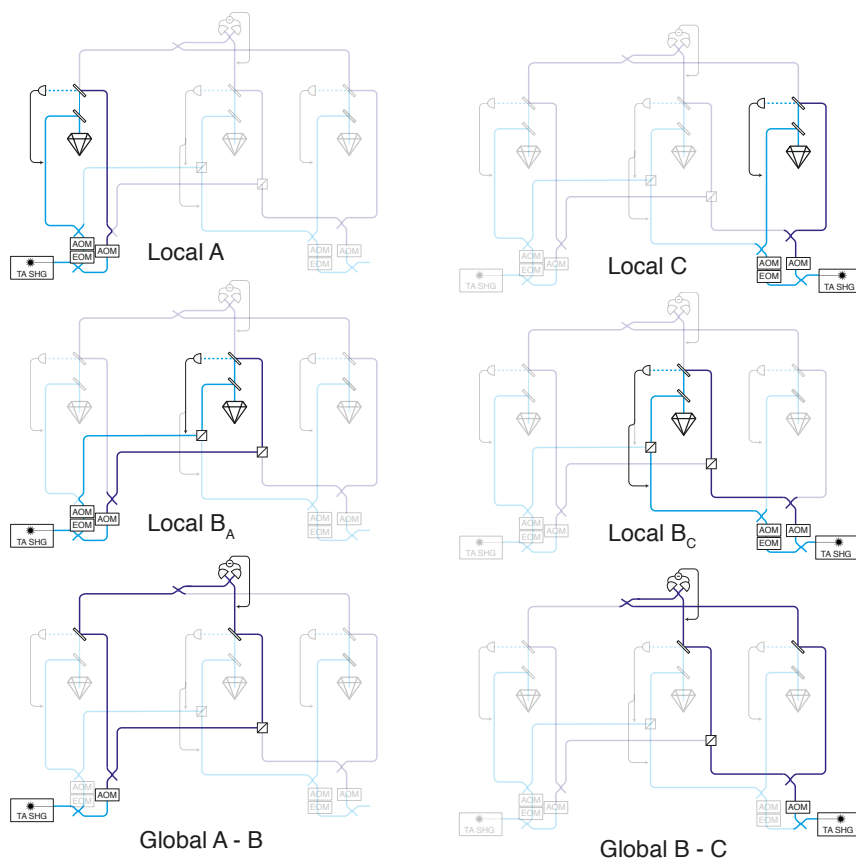
Figure 3.11: Diagrams of the six interferometers in which the optical set-up is divided. For the local interferometers, the heterodyne beat signal (dashed light-blue lines) is measured, compared to an electronic reference signal and feedback is applied to the optical path via piezo-electric mounted mirrors. For the global interferometers, the interference is measured by the single photon detectors. The detectors are shared for the two entanglement links using a beam-splitter that combines the photons from Alice with photons from Charlie. A feedback signal is applied to a fiber stretcher which is also shared by the two global interferometers.
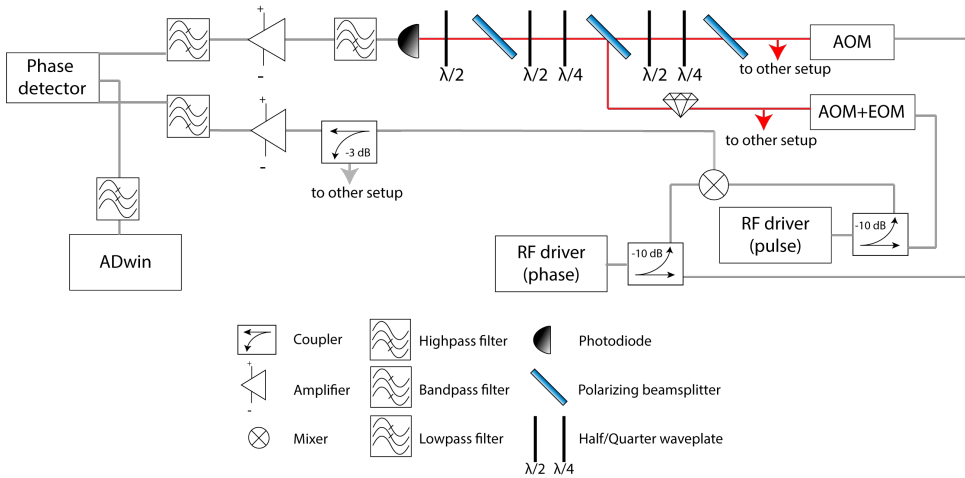
Figure 3.12: Diagram of the electronics and optics for a local interferometer using a heterodyne phase detection scheme. Both the electronic reference signal and the excitation light are shared with another setup.
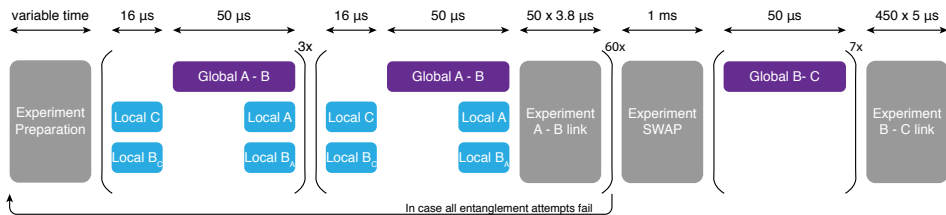
Figure 3.13: Overview of the timings related to the phase stabilization. Experimental time (gray blocks) is interleaved with phase stabilization cycles, which include a phase measurement and a feedback. The subscript to B indicates which light is used, either from setup A or C. The local phase stabilization of A and B_A and the global phase stabilization A-B can be performed at the same time since they use the same light sources.

3
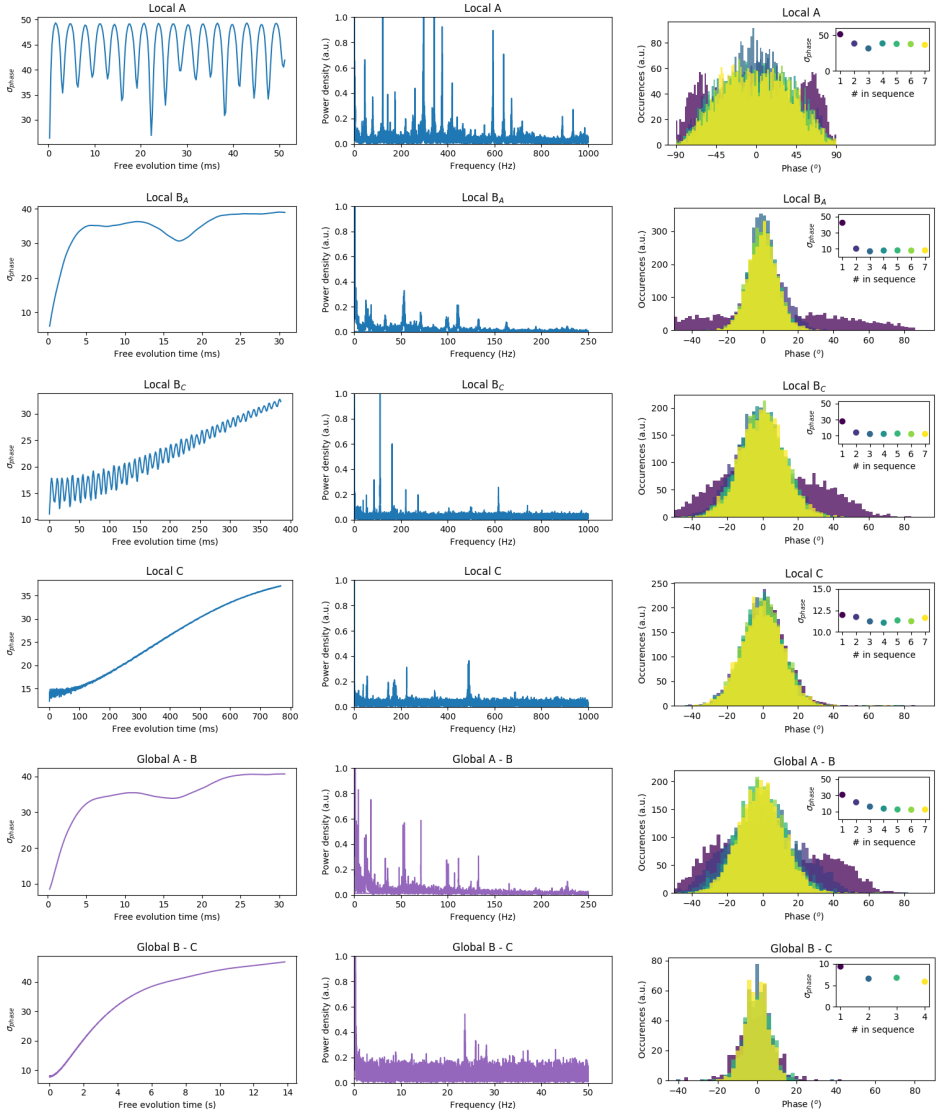
Figure 3.14: Characterization of the phase stabilization of all six interferometers. (**Left**) Standard deviation of the measured phase while changing the free evolution time. (**Center**) Frequency spectrum of the measured noise. (**Right**) Phase distribution for the different rounds of phase stabilization. (**Insets**) Standard deviation of the phase per stabilization round.

3
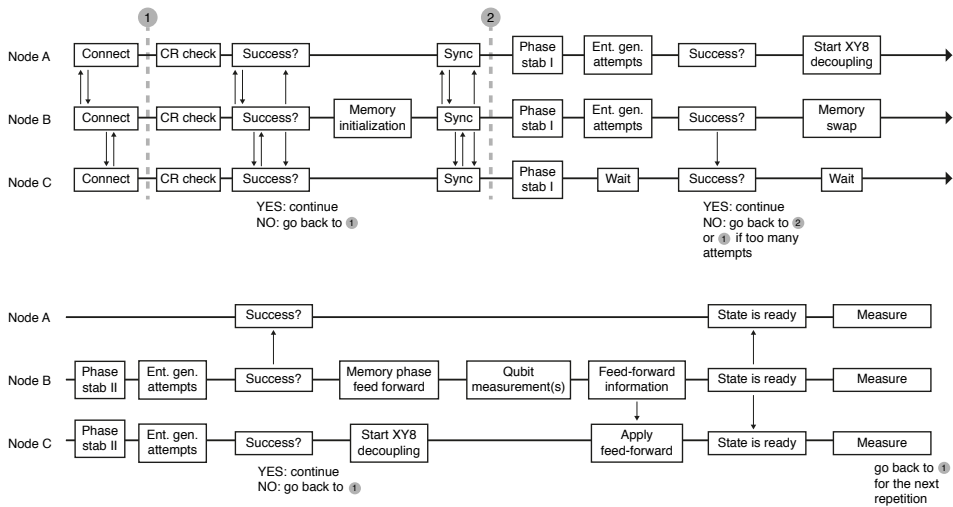
Figure 3.15: Flowchart of the network protocols demonstrations. The micro-controllers of the three nodes exchange information to synchronize the experimental sequence and apply feed-forward operations. Vertical lines between the nodes represent communication steps.

Table 3.1: Experimental equipment used in the three nodes. Non-listed equipment is unchanged from previous experiments and identical for the three setups.

	Alice	Bob	Charlie
Cryostat	attocube attoDRY800	Montana Instruments Cryostation S50	Montana Instruments Cryostation S50
Positioner	Sample on attocube xyz stack	Microscope objective on PI P-615	Microscope objective on PI P-615
Micro-controller	Jäger ADwin-Pro II T12	Jäger ADwin-Pro II T12	Jäger ADwin-Pro II T12
Arbitrary Waveform Generator (AWG)	Tektronix AWG5014	Tektronix AWG5014C	Tektronix AWG5014C
MW source	R&S SGS100A SGMA - up to 6 GHz	R&S SGS100A SGMA - up to 12.75 GHz	R&S SGS100A SGMA - up to 6 GHz
MW amplifier	AR 40S1G4	AR 25S1G4A	AR 40S1G4
Entangling and qubit readout laser	Toptica TA-SHG pro 637 nm	Toptica DL pro 637 nm	Toptica TA-SHG pro 637 nm
Qubit reset laser	Toptica DL pro 637 nm	1 - Toptica TA-SHG pro 637 nm 2 - New Focus Velocity 637 nm	New Focus Velocity 637 nm
Charge reset laser	Toptica DL-SHG pro 575 nm	Cobolt 515 nm	Toptica DL-SHG pro 575 nm
EOM	Jenoptik AM635	Jenoptik AM635	Jenoptik AM635
Deformable mirror	Boston Micromachines 12x12	Boston Micromachines 12x12	Boston Micromachines 12x12

3

Table 3.2: Error budget of the generated Bell states and experimental parameters. The error due to the probability that both nodes emit a photon is related to the values of α (see section 3.5.2) and is therefore intrinsic to the protocol. The infidelity contribution for each of the other errors is estimated as if that error were the only other error present, this way one can easily compare the relative effect of the different infidelity sources. When combined we take into account all the errors at the same time.

Source of infidelity	Expected state infidelity	
	Alice - Bob	Bob - Charlie
Probability that both nodes emit a photon	6.1×10^{-2}	8.0×10^{-2}
Phase uncertainty	6.0×10^{-2}	1.5×10^{-2}
Double excitation	5.5×10^{-2}	7.0×10^{-2}
Photon distinguishability	2.4×10^{-2}	2.3×10^{-2}
Non-NV and dark counts	5×10^{-3}	5×10^{-3}
Combined	0.191	0.186
Measured Ψ^+ infidelity	0.180(5)	0.192(5)
Measured Ψ^- infidelity	0.189(5)	0.189(4)
Experimental parameters	Alice - Bob	Bob - Charlie
p_{det}	$p_A^{\text{det}} = 3.6 \times 10^{-4}$ $p_B^{\text{det}} = 4.4 \times 10^{-4}$	$p_B^{\text{det}} = 4.2 \times 10^{-4}$ $p_C^{\text{det}} = 3.0 \times 10^{-4}$
α	$\alpha_A = 0.07$ $\alpha_B = 0.05$	$\alpha_B = 0.05$ $\alpha_C = 0.10$
p_{dc}	1.5×10^{-7}	1.5×10^{-7}
Visibility V	0.90	0.90
Phase uncertainty	30°	15°
Entanglement attempt duration	3.8 μs	5.0 μs
Probability of double excitation	0.06	0.08

Table 3.3: Fit results for the curves displayed in Fig. 3.3 with and without entanglement generation (Ent. Gen.). See section 3.5.6 for details on the fitting function.

	With Ent. Gen.	Without Ent. Gen.
$N_{1/e}$	1843(32)	2042(36)
n	1.37(5)	1.61(6)
A	0.895(6)	0.885(6)

Table 3.4: Error budget of the generated GHZ states. For each source of infidelity considered we estimate two quantities: the infidelity induced on the state as if it were the only source of error present; the improvement in fidelity if that error were to be removed while all other errors remain present. When combined we take into account all the errors at the same time.

Source of infidelity	Expected infidelity if only source present	Expected improvement once removed
Ψ_{AB} state infidelity	0.191	0.120
Ψ_{BC} state infidelity	0.186	0.122
Memory qubit depolarising noise	8.3×10^{-2}	4.4×10^{-2}
Memory qubit dephasing noise	2.8×10^{-2}	1.5×10^{-2}
Dynamical decoupling of A and C	3.7×10^{-2}	1.9×10^{-2}
At least one node is in NV^0	1.6×10^{-2}	8×10^{-3}
Feed-forward errors	6×10^{-3}	3×10^{-3}
Ψ_{AB} and Ψ_{BC} combined	0.337	0.275
Combined	0.433	
Measured GHZ infidelity	0.462(18)	

Table 3.5: Error budget of the generated Alice-Charlie states. For each source of error considered we estimate two quantities: the infidelity induced on the state as if it were the only source of error present; the improvement in fidelity if that error were to be removed while all other errors remain present. When combined we take into account all the errors at the same time. Errors reported for different Bell state measurement (BSM) results.

Source of infidelity	Expected infidelity if only source present	Expected improvement once removed
Ψ_{AB} state infidelity	0.191	0.115
Ψ_{BC} state infidelity	0.186	0.109
Memory qubit depolarising noise	8.2×10^{-2}	4.0×10^{-2}
Memory qubit dephasing noise	2.8×10^{-2}	1.2×10^{-2}
Dynamical decoupling of A and C	3.7×10^{-2}	1.7×10^{-2}
At least one node is in NV^0	1.6×10^{-2}	6×10^{-3}
Feed-forward errors (00 BSM result)	1.3×10^{-2}	6×10^{-3}
Feed-forward errors (any BSM result)	7.5×10^{-2}	3.4×10^{-2}
Combined (00 BSM result)	0.422	
Combined (any BSM result)	0.451	
Measured Φ_{AC} infidelity (00 BSM result)	0.413(28)	
Measured Φ_{AC} infidelity (any BSM result)	0.449(13)	

References

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

[2] S. Wehner, D. Elkouss, and R. Hanson, *Quantum internet: A vision for the road ahead,* Science **362** (2018), 10.1126/science.aam9288.

[3] L. Jiang, J. M. Taylor, A. S. Sørensen, and M. D. Lukin, *Distributed quantum computation based on small quantum registers,* Physical Review A **76**, 062323 (2007).

[4] A. Broadbent, J. Fitzsimons, and E. Kashefi, *Universal Blind Quantum Computation,* in *2009 50th Annual IEEE Symposium on Foundations of Computer Science* (2009) pp. 517–526.

[5] D. Gottesman, T. Jennewein, and S. Croke, *Longer-Baseline Telescopes Using Quantum Repeaters,* Physical Review Letters **109**, 070503 (2012).

[6] A. Ekert and R. Renner, *The ultimate physical limits of privacy,* Nature **507**, 443 (2014).

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

[8] P. Kómár, E. M. Kessler, M. Bishof, L. Jiang, A. S. Sørensen, J. Ye, and M. D. Lukin, *A quantum network of clocks,* Nature Physics **10**, 582 (2014).

[9] D. L. Moehring, P. Maunz, S. Olmschenk, K. C. Younge, D. N. Matsukevich, L.-M. Duan, and C. Monroe, *Entanglement of single-atom quantum bits at a distance,* Nature **449**, 68 (2007).

[10] S. Ritter, C. Nölleke, C. Hahn, A. Reiserer, A. Neuzner, M. Uphoff, M. Mücke, E. Figueroa, J. Bochmann, and G. Rempe, *An elementary quantum network of single atoms in optical cavities,* Nature **484**, 195 (2012).

[11] J. Hofmann, M. Krug, N. Ortegel, L. Gérard, M. Weber, W. Rosenfeld, and H. Weinfurter, *Heralded Entanglement Between Widely Separated Atoms,* Science **337**, 72 (2012).

[12] L. Stephenson, D. Nadlinger, B. Nichol, S. An, P. Drmota, T. Ballance, K. Thirumalai, J. Goodwin, D. Lucas, and C. Ballance, *High-Rate, High-Fidelity Entanglement of Qubits Across an Elementary Quantum Network,* Physical Review Letters **124**, 110501 (2020).

[13] H. Bernien, B. Hensen, W. Pfaff, G. Koolstra, M. S. Blok, L. Robledo, T. H. Taminiau, M. Markham, D. J. Twitchen, L. Childress, and R. Hanson, *Heralded entanglement between solid-state qubits separated by three metres,* Nature **497**, 86 (2013).

[14] P. C. Humphreys, N. Kalb, J. P. J. Morits, R. N. Schouten, R. F. L. Vermeulen, D. J. Twitchen, M. Markham, and R. Hanson, *Deterministic delivery of remote entanglement on a quantum network,* Nature **558**, 268 (2018).

3

[15] A. Delteil, Z. Sun, W.-b. Gao, E. Togan, S. Faelt, and A. Imamoğlu, *Generation of heralded entanglement between distant hole spins,* Nature Physics **12**, 218 (2016).

[16] R. Stockill, M. Stanley, L. Huthmacher, E. Clarke, M. Hugues, A. Miller, C. Matthiesen, C. Le Gall, and M. Atatüre, *Phase-Tuned Entangled State Generation between Distant Spin Qubits,* Physical Review Letters **119**, 010503 (2017).

[17] P. Maunz, S. Olmschenk, D. Hayes, D. N. Matsukevich, L.-M. Duan, and C. Monroe, *Heralded Quantum Gate between Remote Quantum Memories,* Physical Review Letters **102**, 250502 (2009).

[18] S. Daiss, S. Langenfeld, S. Welte, E. Distante, P. Thomas, L. Hartung, O. Morin, and G. Rempe, *A quantum-logic gate between distant quantum-network modules,* Science **371**, 614 (2021).

[19] N. Kalb, A. A. Reiserer, P. C. Humphreys, J. J. W. Bakermans, S. J. Kamerling, N. H. Nickerson, S. C. Benjamin, D. J. Twitchen, M. Markham, and R. Hanson, *Entanglement distillation between solid-state quantum network nodes,* Science **356**, 928 (2017).

[20] M. Christandl and S. Wehner, *Quantum Anonymous Transmissions,* in *Advances in Cryptology - ASIACRYPT 2005*, Lecture Notes in Computer Science, edited by B. Roy (Springer, Berlin, Heidelberg, 2005) pp. 217–235.

[21] M. Hillery, V. Bužek, and A. Berthiaume, *Quantum secret sharing,* Physical Review A **59**, 1829 (1999).

[22] A. Ambainis, H. Buhrman, Y. Dodis, and H. Rohrig, *Multiparty quantum coin flipping,* in *Proceedings. 19th IEEE Annual Conference on Computational Complexity, 2004.* (2004) pp. 250–259.

[23] H.-J. Briegel, W. Dür, J. I. Cirac, and P. Zoller, *Quantum Repeaters: The Role of Imperfect Local Operations in Quantum Communication,* Physical Review Letters **81**, 5932 (1998).

[24] M. Pant, H. Krovi, D. Towsley, L. Tassiulas, L. Jiang, P. Basu, D. Englund, and S. Guha, *Routing entanglement in the quantum internet,* npj Quantum Information **5**, 1 (2019).

[25] D. Bouwmeester, J.-W. Pan, M. Daniell, H. Weinfurter, and A. Zeilinger, *Observation of Three-Photon Greenberger-Horne-Zeilinger Entanglement,* Physical Review Letters **82**, 1345 (1999).

[26] B. Jing, X.-J. Wang, Y. Yu, P.-F. Sun, Y. Jiang, S.-J. Yang, W.-H. Jiang, X.-Y. Luo, J. Zhang, X. Jiang, X.-H. Bao, and J.-W. Pan, *Entanglement of three quantum memories via interference of three single photons,* Nature Photonics **13**, 210 (2019).

[27] C. Cabrillo, J. I. Cirac, P. García-Fernández, and P. Zoller, *Creation of entangled states of distant atoms by interference,* Physical Review A **59**, 1025 (1999).

[28] S. Bose, P. L. Knight, M. B. Plenio, and V. Vedral, *Proposal for Teleportation of an Atomic State via Cavity Decay,* Physical Review Letters **83**, 5158 (1999).

3

[29] C. Bradley, J. Randall, M. Abobeih, R. Berrevoets, M. Degen, M. Bakker, M. Markham, D. Twitchen, and T. Taminiau, *A Ten-Qubit Solid-State Spin Register with Quantum Memory up to One Minute,* Physical Review X **9**, 031045 (2019).

[30] N. Kalb, P. C. Humphreys, J. J. Slim, and R. Hanson, *Dephasing mechanisms of diamond-based nuclear-spin memories for quantum networks,* Physical Review A **97**, 062330 (2018).

[31] O. Gühne and G. Tóth, *Entanglement detection,* Physics Reports **474**, 1 (2009).

[32] L. Jiang, J. S. Hodges, J. R. Maze, P. Maurer, J. M. Taylor, D. G. Cory, P. R. Hemmer, R. L. Walsworth, A. Yacoby, A. S. Zibrov, and M. D. Lukin, *Repetitive Readout of a Single Electronic Spin via Quantum Logic with Nuclear Spin Ancillae,* Science **326**, 267 (2009).

[33] D. Riedel, I. Söllner, B. J. Shields, S. Starosielec, P. Appel, E. Neu, P. Maletinsky, and R. J. Warburton, *Deterministic Enhancement of Coherent Photon Generation from a Nitrogen-Vacancy Center in Ultrapure Diamond,* Physical Review X **7**, 031040 (2017).

[34] E. Janitz, M. K. Bhaskar, and L. Childress, *Cavity quantum electrodynamics with color centers in diamond,* Optica **7**, 1232 (2020).

[35] M. Ruf, M. Weaver, S. van Dam, and R. Hanson, *Resonant Excitation and Purcell Enhancement of Coherent Nitrogen-Vacancy Centers Coupled to a Fabry-Perot Microcavity,* Physical Review Applied **15**, 024049 (2021).

[36] S. Baier, C. Bradley, T. Middelburg, V. Dobrovitski, T. Taminiau, and R. Hanson, *Orbital and Spin Dynamics of Single Neutrally-Charged Nitrogen-Vacancy Centers in Diamond,* Physical Review Letters **125**, 193601 (2020).

[37] R. Van Meter, *Quantum networking*, Networks and telecommunications series (John Wiley & Sons, 2014).

[38] A. Pirker and W. Dür, *A quantum network stack and protocols for reliable entanglement-based networks,* New Journal of Physics **21**, 033003 (2019).

[39] W. Kozlowski and S. Wehner, *Towards Large-Scale Quantum Networks,* in *Proceedings of the Sixth Annual ACM International Conference on Nanoscale Computing and Communication*, NANOCOM '19 (Association for Computing Machinery, New York, NY, USA, 2019) pp. 1–7.

[40] A. Dahlberg, M. Skrzypczyk, T. Coopmans, L. Wubben, F. Rozpędek, M. Pompili, A. Stolk, P. Pawełczak, R. Knegjens, J. de Oliveira Filho, R. Hanson, and S. Wehner, *A link layer protocol for quantum networks,* in *Proceedings of the ACM Special Interest Group on Data Communication*, SIGCOMM '19 (Association for Computing Machinery, New York, NY, USA, 2019) pp. 159–173.

[41] A. Tchebotareva, S. L. Hermans, P. C. Humphreys, D. Voigt, P. J. Harmsma, L. K. Cheng, A. L. Verlaan, N. Dijkhuizen, W. de Jong, A. Dréau, and R. Hanson, *Entanglement between a Diamond Spin Qubit and a Photonic Time-Bin Qubit at Telecom Wavelength,* Physical Review Letters **123**, 063601 (2019).

3

[42] B. C. Rose, D. Huang, Z.-H. Zhang, P. Stevenson, A. M. Tyryshkin, S. Sangtawesin, S. Srinivasan, L. Loudin, M. L. Markham, A. M. Edmonds, D. J. Twitchen, S. A. Lyon, and N. P. d. Leon, *Observation of an environmentally insensitive solid-state spin defect in diamond,* Science **361**, 60 (2018).

[43] C. Nguyen, D. Sukachev, M. Bhaskar, B. Machielse, D. Levonian, E. Knall, P. Stroganov, R. Riedinger, H. Park, M. Lončar, and M. Lukin, *Quantum Network Nodes Based on Diamond Qubits with an Efficient Nanophotonic Interface,* Physical Review Letters **123**, 183602 (2019).

[44] M. E. Trusheim, B. Pingault, N. H. Wan, M. Gündoğan, L. De Santis, R. Debroux, D. Gangloff, C. Purser, K. C. Chen, M. Walsh, J. J. Rose, J. N. Becker, B. Lienhard, E. Bersin, I. Paradeisanos, G. Wang, D. Lyzwa, A. R.-P. Montblanch, G. Malladi, H. Bakhru, A. C. Ferrari, I. A. Walmsley, M. Atatüre, and D. Englund, *Transform-Limited Photons From a Coherent Tin-Vacancy Spin in Diamond,* Physical Review Letters **124**, 023602 (2020).

[45] N. T. Son, C. P. Anderson, A. Bourassa, K. C. Miao, C. Babin, M. Widmann, M. Niethammer, J. Ul Hassan, N. Morioka, I. G. Ivanov, F. Kaiser, J. Wrachtrup, and D. D. Awschalom, *Developing silicon carbide for quantum spintronics,* Applied Physics Letters **116**, 190501 (2020).

[46] B. Hensen, H. Bernien, A. E. Dréau, A. Reiserer, N. Kalb, M. S. Blok, J. Ruitenberg, R. F. L. Vermeulen, R. N. Schouten, C. Abellán, W. Amaya, V. Pruneri, M. W. Mitchell, M. Markham, D. J. Twitchen, D. Elkouss, S. Wehner, T. H. Taminiau, and R. Hanson, *Loophole-free Bell inequality violation using electron spins separated by 1.3 kilometres,* Nature **526**, 682 (2015).

[47] M. Pompili, S. Hermans, S. Baier, H. Beukers, P. Humphreys, R. Schouten, R. Vermeulen, M. Tiggelman, L. d. S. Martins, B. Dirkse, S. Wehner, and R. Hanson, *Data and software supporting "Realization of a multi-node quantum network of remote solid-state qubits",* (2021), 10.4121/13600589.v1.

[48] M. Pompili, S. L. N. Hermans, S. Baier, H. K. C. Beukers, P. C. Humphreys, R. N. Schouten, R. F. L. Vermeulen, M. J. Tiggelman, L. d. S. Martins, B. Dirkse, S. Wehner, and R. Hanson, *Realization of a multinode quantum network of remote solid-state qubits,* Science **372**, 259 (2021).

4

Qubit teleportation between non-neighboring nodes

S. L. N. Hermans*, M. Pompili*, H. K. C. Beukers, S. Baier, J. Borregaard & R. Hanson

Future quantum internet applications will derive their power from the ability to share quantum information across the network. Quantum teleportation allows for the reliable transfer of quantum information between distant nodes, even in the presence of highly lossy network connections. While many experimental demonstrations have been performed on different quantum network platforms, moving beyond directly connected nodes has so far been hindered by the demanding requirements on the pre-shared remote entanglement, joint qubit readout and coherence times. Here we realize quantum teleportation between remote, non-neighboring nodes in a quantum network. The network employs three optically connected nodes based on solid-state spin qubits. The teleporter is prepared by establishing remote entanglement on the two links, followed by entanglement swapping on the middle node and storage in a memory qubit. We demonstrate that once successful preparation of the teleporter is heralded, arbitrary qubit states can be teleported with fidelity above the classical bound, even with unit efficiency. These results are enabled by key innovations in the qubit readout procedure, active memory qubit protection during entanglement generation and tailored heralding that reduces remote entanglement infidelities. Our work demonstrates a prime building block for future quantum networks and opens the door to exploring teleportation-based multi-node protocols and applications.

Quantum teleportation is the central routine for reliably sending qubits across lossy network links [1] as well as a key primitive of quantum network protocols and applications [2–4]. Using a teleporter in the form of a pre-shared entangled state between the sending node and the receiving node, the quantum information is transferred by performing a joint Bell-state measurement on the sender's part of the entangled state and the qubit state to be teleported, followed by a gate operation on the receiving node conditioned on the Bell-state measurement outcome [1]. Since the quantum information is not transmitted by a physical carrier, the protocol is insensitive to loss in the connecting photonic channels and on intermediate nodes. A deterministic Bell-state measurement combined with real-time feed-forward enables unconditional teleportation, in which state transfer is achieved each time a qubit state is inserted into the teleporter.

Pioneering explorations of quantum teleportation protocols were performed using photonic states [5–7]. Following the development of quantum network nodes with stationary qubits, remote qubit teleportation was realized between trapped ions [8], trapped atoms [9, 10], diamond NV centers [11] and memory nodes based on atomic ensembles [12].

While future quantum network applications will widely employ teleportation between non-connected nodes in the network, the demanding set of requirements on the pre-shared entanglement, the Bell-state measurement and the coherence times for enabling real-time feed-forward has so far prevented the realization of teleportation beyond directly connected stationary network nodes.

Here, we overcome these challenges by a set of key innovations and achieve qubit teleportation between non-neighboring network nodes (see fig. 4.1a). Our quantum network consists of three nodes in a line configuration, Alice, Bob and Charlie. Each node contains a nitrogen-vacancy (NV) center in diamond. Using the NV electronic spin as the communication qubit we are able to generate remote entanglement between each pair of neighboring nodes. In addition, Bob and Charlie each employ a nearby ^{13}C nuclear spin as a memory qubit. The steps of the teleportation protocol are shown in Fig. 4.1b. To prepare the teleporter we use an entanglement swapping protocol mediated by Bob to establish entanglement between Alice and Charlie. Once successful preparation of teleporter is heralded, the input qubit state is prepared on Charlie and finally teleported to Alice.
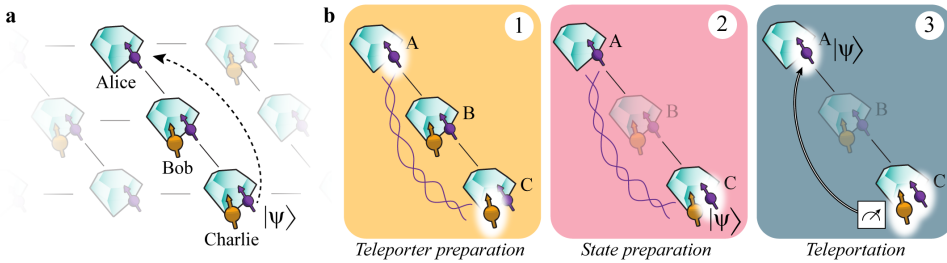
Figure 4.1: Teleporting a qubit between non-neighboring nodes of a quantum network. (a) Three network nodes, Alice (A), Bob (B) and Charlie (C) are connected via optical fiber links (lines) in a line configuration. Each setup has a communication qubit (purple) that enables entanglement generation with its neighboring node. Additionally, Bob and Charlie contain a memory qubit (yellow). (b) The steps of the teleportation protocol: (1) We prepare the teleporter by establishing entanglement between Alice and Charlie using an entanglement swapping protocol on Bob, followed by swapping the state at Charlie to the memory qubit. (2) The qubit state to be teleported is prepared on the communication qubit on Charlie. (3) A Bell-state measurement is performed on Charlie's qubits and the outcome is communicated to Alice over a classical channel. Dependent on this outcome, Alice applies a quantum gate to obtain the teleported qubit state.

4.1 Entanglement fidelity of the network links

A key parameter for quantum teleportation is the fidelity of the pre-shared entangled state between Alice and Charlie. As we generate this state by entanglement swapping, its fidelity is upper bounded by the errors on the individual links. Therefore, mitigating error sources on the individual links is critical. Our network generates entanglement between neighboring nodes using a single-photon protocol [13, 14] in an optical-phase-stabilized architecture [15]. The building block of this protocol is a qubit-photon entangled state created at each node. To generate this entangled state we initialize the communication qubit in a superposition state $|\psi\rangle = \sqrt{\alpha}|0\rangle + \sqrt{1-\alpha}|1\rangle$ and apply a state-selective optical pulse that transfers the population from $|0\rangle$ to an optically excited state. Following spontaneous emission, the qubit state is entangled with photon number (0 or 1 photon). We perform this protocol on both nodes and interfere the resonant photonic states on a beam splitter (Fig. 4.2a). Detection of a single photon in one of the beam splitter output ports ideally heralds the generation of an entangled state $|\psi\rangle = (|01\rangle \pm |10\rangle)/\sqrt{2}$, where the ± phase is set by which detector clicked. Fig. 4.2b displays the joint outcomes of qubit measurements in the computational basis after entanglement is heralded, showing the expected correlations.

The infidelity of the generated state has three main contributions: double $|0\rangle$ state occupancy, double optical excitation and finite distinguishability of the photons [15, 16] (see section 4.6). In the case of double $|0\rangle$ state occupancy (which occurs with probability α), both communication qubits are in the $|0\rangle$ state and have emitted a photon. Detection of one of these photons leads to false heralding of an entangled state. The second effect, double excitation, is due to the finite length of the optical pulse compared to the emitter's optical lifetime. There is a finite chance that the communication qubit emits a photon during this pulse, is subsequently re-excited during the remainder of the pulse and then emits another photon, resulting in the qubit state being entangled with two photons. Detection or loss of the first photon destroys the coherence of the qubit-photon entangled state and detection of the second photon can then falsely herald the generation of an entangled

state.

Crucially, false heralding events due to double $|0\rangle$ state occupancy and double excitation are both accompanied by an extra emitted photon. Therefore, detection of this additional photon allows for unambiguous identification of such events and thus for real-time rejection of the corresponding false heralding signals. We implement this rejection scheme by monitoring the off-resonant phonon-side band (PSB) detection path on both setups during and after the optical excitation (see fig. 4.2a).

To investigate the effect of this scheme, we generate entanglement on the individual links and extract the entanglement heralding events for which the PSB monitoring flagged the presence of an additional photon. For these events, we again analyze the corresponding qubit measurements in the computational basis (Fig. 4.2c).

We identify two separate regimes: one during the optical pulse (purple) and one after the optical pulse (yellow). When a photon is detected on Alice's (Bob's) PSB detector during the optical pulse, we see that the outcome 01 (10) is most probable (purple data in Fig. 4.2c) showing that only one setup was in the $|0\rangle$ state and thus that both detected photons originated from Alice (Bob). The detection of PSB photons during the optical pulse thus primarily flags double excitation errors. In contrast, when a photon is detected after the optical pulse in either Alice's or Bob's PSB detector, the outcome 00 is most probable (yellow data in Fig. 4.2c), indicating that both setups were in the $|0\rangle$ state and both emitted one photon. PSB photon detection after the optical pulse thus flags the double $|0\rangle$ state occupancy error. We find similar results to Fig. 4.2c for the entangled states generated on the Bob-Charlie link, see section 4.6. The improvement in fidelity from rejecting these false heralding events in our experiment is set by the combined probability of occurrence ($\approx 9\,\%$, see section 4.6) multiplied by the probability to flag them (given here by the total PSB photon detection efficiency of $\approx 10\,\%$).

The third main source of infidelity, the finite distinguishability, can arise from frequency detunings between the emitted photons [17]. While most of these detunings are eliminated upfront by the charge-resonance (CR) check before the start of the protocol (see section 4.6), the communication qubits may still be subject to a small amount of spectral diffusion. In our single-photon protocol, this leads to dephasing that is stronger for photons that are detected later relative to the optical pulse. By shortening our detection window, we can increase the fidelity of the entangled state at the expense of a lower entangling rate. For the experiments below (unless mentioned differently) we use a detection window length of 15 ns. Fig. 4.2d summarizes the measured improvements on the individual links. For the teleporter, we estimate that their combined effect is an increase in Alice-Charlie entangled state fidelity by $\approx 3\,\%$. This increase is instrumental in pushing the teleportation fidelity above the classical bound.
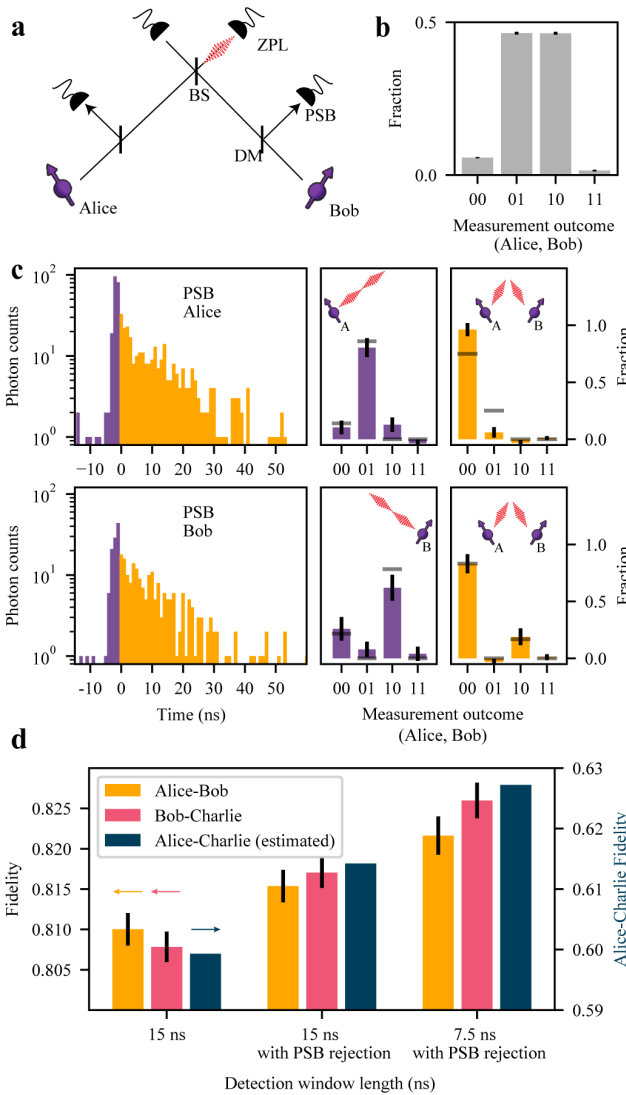
Figure 4.2: High-fidelity entangled network links. (a) Simplified schematic of the optical link used for generating entanglement between neighboring nodes. Photons emitted by the communication qubits are filtered by a dichroic mirror (DM) to separate the resonant (zero-phonon line, ZPL) photons (3 % of emission) from the off-resonant (phonon-side band, PSB) photons (97 % of emission). The resonant photons are sent to the beam splitter (BS); detection of a single photon at one of the ZPL detectors heralds successful generation of an entangled state between the two nodes. (b) Measured correlations of the communication qubits in the computational basis, conditioned on a heralding event on the ZPL detectors. (c) (left) Histograms of the PSB photon detection times on Alice (top) or Bob (bottom), conditioned on a simultaneous ZPL detection in the same entanglement generation attempt. Gray lines show expected correlations based on a quantum-optical model (see section 4.6). (d) Measured fidelity of the network links, without PSB rejection (left), with PSB rejection (middle) and with PSB rejection plus shortened detection window (right). The dark blue bars indicate the corresponding expected fidelity on Alice-Charlie after entanglement swapping for each case (see section 4.6).

4.2 Memory qubit coherence

In the preparation of the teleporter it is crucial that the first entangled state between Alice and Bob is reliably preserved on the memory qubit while the second link between Bob and Charlie is being generated. For this reason, we abort the sequence and start over when the second entangled state is not heralded within a fixed number of attempts, the timeout.

The ^{13}C memory qubits can be controlled with high fidelity via the communication qubit while they can be efficiently decoupled when no interaction is desired. Recent work showed that in a magnetic field of 189 mT entanglement generation attempts with the communication qubit do not limit the memory dephasing time $T_2^\star$ [15], opening the door to significantly extending the memory preservation time with active coherence protection from the spin bath [18]. We realize this protection by integrating a decoupling π-pulse on the memory qubit into the experimental sequence that follows a heralding event, while ensuring that all phases that are picked up due to the probabilistic nature of the remote entangling process are compensated in real time (Fig. 4.3a).

In Fig. 4.3b we check the performance of this sequence by storing a superposition state on the memory qubit and measuring the Bloch vector length. We compare the results for the sequence with and without the decoupling π-pulse, and with and without entanglement attempts. We observe that without the decoupling pulse the decay of the Bloch vector length is not altered by the entanglement attempts, in line with previous findings [15]. In contrast, when we apply the decoupling pulse the decay is slowed down by more than a factor of 6, yielding a $N_{1/e}$ decay constant of ≈5300 entanglement attempts, the highest number reported to date for diamond devices. In addition, we observe a difference in the shape of the decay between the cases with and without entangling attempts, indicating that intrinsic decoherence is no longer the only limiting error source. The improved memory coherence enables us to use a timeout of 1000 entangling attempts, more than double that of Ref. [15], which doubles the entanglement swapping rate.

4.3 Memory qubit readout

High-fidelity memory qubit readout is required both in the preparation of the teleporter (at Bob) and during the teleportation protocol itself (at Charlie). The memory qubit is read out by mapping its state onto the communication qubit using quantum logic followed by single-shot readout of the communication qubit using state-dependent optical excitation and detection [19]. Due to limited photon collection efficiency (≈10 %) and finite cyclicity of the optical transition (≈99 %), the communication qubit readout fidelity is different for $|0\rangle$ and $|1\rangle$. As a result, for random initial states the probability that the correct state was assigned is significantly larger if one or more photons were detected (assigned outcome 0) than if no photons were detected (assigned outcome 1) [20]. In previous work we circumvented this issue by conditioning on obtaining the outcome 0 [15]. However, this approach scales unfavorably, as it forces the protocol to prematurely abort with probability >50 % at each memory qubit readout. Therefore, to access more complex protocols with multiple memory qubit readouts, near-deterministic readout schemes are required.

We resolve this challenge by introducing a basis-alternating repetitive readout for the memory qubit (see fig. 4.3c). The key point of this readout strategy is, in contrast to earlier work [21], to alternately map the computational basis states of the memory qubit

to the communication qubit state $|0\rangle$. Fig. 4.3d shows the readout fidelities of the n-th readout repetition for the two initial states for the memory qubit on Bob (for Charlie, see section 4.6). We clearly observe the expected alternating pattern due to the asymmetry of the communication qubit readout fidelities. Importantly, the readout fidelity decays only by ≈1 % per readout, showing that the readout is mostly non-demolition and multiple readouts are possible without losing the state. We model the readout procedure using measured parameters (see section 4.6) and plot the model's predictions as dashed lines in Fig. 4.3d-f.

Next, we assign the state using the first readout and continue the sequence only when the consecutive readouts are consistent with the first readout. The subsequent readouts therefore add confidence to the assignment in the case of consistent outcomes, while cases of inconsistent outcomes (which have a higher chance of indicating an incorrect assignment) are filtered out. In Fig. 4.3e we plot the readout fidelity resulting from this strategy for up to five readouts, with the corresponding rejected fraction due to inconsistent outcomes plotted in Fig. 4.3f. We observe that using two readouts already eliminates most of the asymmetry, reducing the average infidelity from ≈6 % to below 1 %. At this point, the remaining observed infidelity mainly results from cases where the memory qubit was flipped during the first readout block. While adding further readout blocks does not lead to significant improvements in fidelity, each two additional readouts cut the amount of consistent outcomes by ≈10 %, due to the communication qubit readout infidelities and gate errors. For the experiments reported below (unless mentioned differently) we use two readout repetitions to benefit from a high average readout fidelity (Bob: 99.2(4) %, Charlie: 98.1(4) %) and a high probability to continue the sequence (Bob and Charlie: ≈88 %).
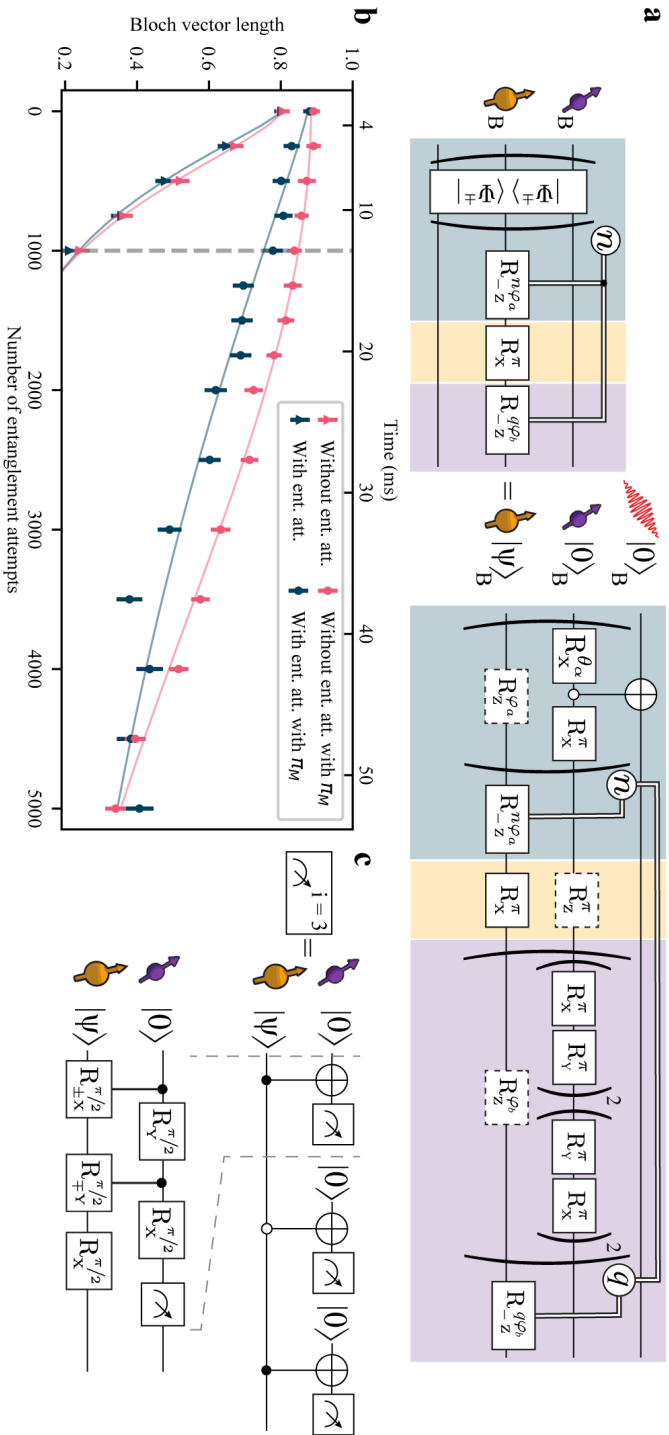
Figure 4.3: Memory qubit coherence and readout. (a) Gate sequence on Bob for entanglement generation with the communication qubit while preserving states stored on the memory qubit. Entanglement generation attempts are repeated until success or a predetermined timeout. Upon success, a phase feed-forward is applied to maintain the correct reference frame of the memory qubit [15], followed by a decoupling pulse on the memory qubit. The decoupling π_M pulse causes a Z-rotation on the communication qubit. Afterwards, we rephase the memory qubit for the same amount of time as it took to herald entanglement while applying an XY8 decoupling sequence on the communication qubit, and we end with another phase feed-forward on the memory qubit to compensate for any phase picked up during this decoupling. **(b)** Bloch vector length of a superposition state stored on the memory qubit for different number of entanglement attempts or a time-equivalent wait element. In the case of no decoupling (no π_M) on the memory qubit, the gates in the yellow shaded box in (a) are left out. The gray dashed line indicates the chosen timeout of 1000 entanglement attempts. **(c)** Gate sequence for the basis-alternating repetitive readout of the memory qubit.
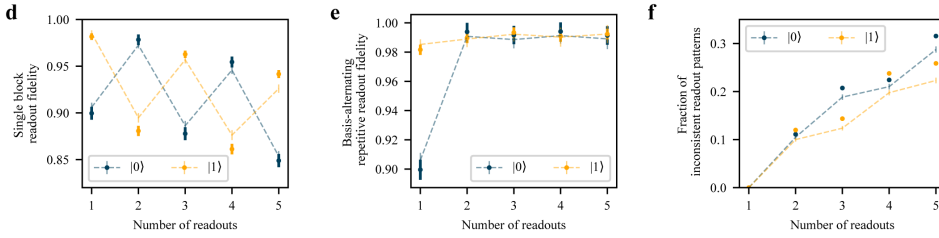
Figure 4.3 (continued): High-fidelity entangled network links. (**d**) Readout fidelity for each readout repetition, for state $|0\rangle$ and $|1\rangle$. (**e**) Readout fidelity of the basis-alternating repetitive readout scheme for different number of readout repetitions. (**f**) Fraction of inconsistent readout patterns for different number of readout repetitions. In (d-f) the dashed lines show a numerical model using measured parameters.

4.4 Teleporting qubit states from Charlie to Alice

With all innovations described above implemented, we perform the protocol as shown in Fig. 4.4a. First, we generate entanglement between Alice and Bob and store Bob's part of the entangled state on the memory qubit using a compiled SWAP operation. Second, we generate entanglement between Bob and Charlie, while preserving the first entangled state on the memory qubit with the pulse sequence as described in Fig. 4.3a. Next, we perform a Bell-state measurement on Bob followed by a CR check. We continue the sequence if the communication qubit readout yields outcome 0, the memory qubit readout gives a consistent outcome pattern, and the CR check is passed. At Charlie, we perform a quantum gate that depends on the outcome of the Bell-state measurement and on which detectors clicked during the two-node entanglement generation. Next, we swap the entangled state to the memory qubit. At this point the teleporter is ready and Alice and Charlie share an entangled state with an estimated fidelity of 0.61.

Subsequently, we generate the qubit state to be teleported, $|\psi\rangle$, on Charlie's communication qubit and run the teleportation protocol. First, a Bell-state measurement is performed on the communication and memory qubits at Charlie. With the exception of unconditional teleportation (discussed below), we only continue the sequence when we obtain a 0 outcome on the communication qubit, when we have a consistent readout pattern on the memory qubit and when Charlie passes the CR check. The outcomes of the Bell-state measurement are sent to Alice and by applying the corresponding gate operation we obtain $|\psi\rangle$ on Alice's side.

We teleport the six cardinal states ($\pm X, \pm Y, \pm Z$), which form an unbiased set [22], and measure the fidelity of the teleported states to the ideally prepared state (Fig. 4.4b). We find an average teleported state fidelity of $F = 0.702(11)$ at an experimental rate of 1/(117 s). This value exceeds the classical bound of 2/3 by more than three standard deviations, thereby proving the quantum nature of the protocol. We note that this value provides a lower bound to the true teleportation fidelity, as the measured fidelity is lowered by errors in the preparation of the qubit states at Charlie (estimated to be 0.5 %, see section 4.6).

The differences in fidelity between the teleported states arise from an interplay of errors in different parts of the protocol that either affect all three axes (depolarizing errors) or only two axes (dephasing errors). These differences are qualitatively reproduced by our

model (gray bars in Fig. 4.4b). In Fig. 4.4c we plot the teleportation fidelity for each possible outcome of the Bell-state measurement. Due to the basis-alternating repetitive readout, the dependence on the second bit (from the memory qubit readout) is small, whereas for the first bit (communication qubit readout) the best teleported state fidelity is achieved for outcome 0 due to the asymmetric readout fidelities. We also analyze the case in which no feed-forward is applied at Alice section 4.6; as expected, the average state fidelity reduces to a value consistent with a fully mixed state (fidelity $F = 0.501(7)$), emphasizing the critical role of the feed-forward in the teleportation protocol.

Finally, we demonstrate that the network can achieve unconditional teleportation between Alice and Charlie. Unconditional teleportation requires that, following preparation of the teleporter by establishing the remote entangled state, the protocol runs deterministically (each qubit state prepared at Charlie ends up at Alice) while surpassing the classical fidelity bound. We thus require that the Bell-state measurement at Charlie and the subsequent feed-forward operations are performed deterministically. To this end, we revise the protocol at Charlie to accept both communication qubit outcomes, use all memory qubit readout patterns including the inconsistent ones and disregard the outcome of the CR check after the Bell-state measurement. Using this fully deterministic Bell-state measurement lowers the average teleportation fidelity by a few percents (Fig. 4.4d). At the same time, shortening the detection windows of the two-node entanglement generation is expected to yield an improvement in the fidelity, as discussed above. We find indeed that the average unconditional teleportation fidelity increases with shorter window lengths, reaching $F = 0.688(10)$ for a length of 7.5 ns and a rate of 1/(100 s). The current quantum network is thus able to perform teleportation beyond the classical bound, even under the strict condition that every state inserted into the teleporter be transferred.

Figure 4.4: Qubit teleportation between non-neighboring network nodes. (**a**) Circuit diagram of the teleportation protocol using notation defined in Fig. 4.3. See section 4.6 for the full circuit diagram. (**b**) Teleported state fidelities for the six cardinal states and their average. The gray lines show the expected fidelities from simulations. The dashed lines in (b-d) represents the classical bound of 2/3. (**c**) Average teleported state fidelity for the different outcomes of the Bell-state measurement on Charlie. The right-most bar shows the resulting fidelity when no feed forward operation on Alice would be applied. (**d**) Average state fidelity for a conditional and unconditional teleportation, for different detection window lengths of the two-node entanglement generation processes. The blue bordered data point is the same point as shown in (b).

4.5 Conclusion and outlook

In this work we have realized unconditional qubit teleportation between non-neighboring nodes in a quantum network. The innovations introduced here on memory qubit readout and protection during entanglement generation, as well as the real-time rejection of false heralding signals, will be instrumental in exploring more complex protocols [2–4, 23, 24]. Also, these methods can be readily transferred to other platforms such as the group-IV color centers in diamond, the vacancy-related qubits in SiC and single rare-earth ions in solids [25–31].

The development of an improved optical interface for the communication qubit [32] will increase both the teleportation protocol rate and fidelity. Because of the improved memory qubit performance reported here, the network already operates close to the threshold where nodes can reliably deliver a remote entangled state while preserving previously stored quantum states in their memory qubits. With further improvements, for instance by integrating multi-pulse memory decoupling sequences [18] into the entanglement generation, demonstration of deterministic qubit teleportation may come within reach. In that case, the network is able to teleport a qubit state with unit efficiency at any given time, removing the need for heralding successful preparation of the teleporter and opening the door to exploring applications that call the teleportation routine multiple times.

Finally, by implementing a recently proposed link layer protocol [33], qubit teleportation and applications making use of the teleportation primitive may be executed and tested on the network through platform-independent control software, an important prerequisite for a large-scale future network.

4.6 Supplementary information

4.6.1 Full gate circuit

Our quantum network consists of three nodes, Alice, Bob and Charlie. In the experiment, we will teleport a qubit from Charlie to Alice, two non-neighboring nodes. The full gate circuit is shown in Fig. 4.5. Prior to the sequence, we do a Charge-Resonance (CR) check on each node to ensure that the communication qubits are in the correct charge state (NV⁻) and on resonance with the control lasers. Once all the nodes have passed this check, we do a first round of optical phase stabilization of the interferometers, which enables the entanglement generation using the single click protocol [13–16]. After these preparation steps, the sequence is triggered on all setups.

On Bob, we initialize the memory qubit into $|0\rangle$ using the communication qubit [19]. Next, we generate entanglement between the communication qubits of Alice and Bob. When entanglement is heralded, we perform a SWAP operation to store Bob's part of the entangled state on the memory qubit.

We continue with a second round of phase stabilization (not shown in the circuit) and generate entanglement between the communication qubits of Bob and Charlie. Each entanglement attempt slightly decoheres the memory qubit, therefore we limit the number of attempts by a timeout. If we do not succeed within the timeout, we abort the sequence and start over.

During entanglement generation, the memory qubit of Bob picks up an average phase $n\varphi_a$ dependent on the number of entanglement attempts n. Due to the probabilistic nature of the entanglement generation process, we do not know which attempt will be successful, therefore this phase is unknown at the start of the sequence. To maintain the correct reference frame of the memory qubit this phase needs to be corrected in real-time before any other gate can be applied to the memory qubit. We perform this real-time correction by changing the time between pulses on the communication qubit [15]. After the phase correction, the decoupling pulse is applied to the memory qubit via the communication qubit. The back-action of this gate causes a Z-rotation on the communication qubit. To rephase the memory qubit, we wait for the same amount of time as it took to herald the second entangled state while decoupling the communication qubit. This imprints a phase $q\varphi_b$ on the memory qubit, which we compensate in an analogous way.

Bob now shares two entangled states; his memory qubit is entangled with Alice and his communication qubit with Charlie. To establish an entangled state between Alice and Charlie we perform a Bell-state measurement on the two qubits of Bob. To do so, we entangle the communication and memory qubits and do a measurement on the communication qubit. We map its state onto the communication qubit and measure the communication qubit. In the basis-alternating repetitive readout, we repeat the measurement sequence twice. During the first readout we map the $|0\rangle$ state to the $|0\rangle$ state of the communication qubit, and in the second readout we map $|1\rangle$ to $|0\rangle$. The first outcome is used to assign the state and the second outcomes serves as a check. By continuing the sequence only when we measure consistent patterns (for instance $(m1, m2) = (1, 0)$) we increase our average readout fidelity. After the readout procedure, we perform a CR check on Bob to filter out any event where Bob was in the wrong charge state.

Bob communicates to Charlie which gate operation should be done to obtain the correct entangled state. Which operation is required is determined by the outcomes of the

Bell-state measurement on Bob and by which detector heralded the individual links. Charlie performs the feed-forward gate operation and subsequently stores its part of the entangled state on the memory qubit using a SWAP gate. At this point in the sequence the teleporter is ready.

To prepare the state that is to be teleported, we initialize the communication qubit at Charlie and perform the desired qubit rotation.

To teleport the qubit, we perform a Bell-state measurement on the qubits of Charlie. Locally, we entangle the communication qubit with the memory qubit. We readout the communication qubit and use the basis-alternating repetitive readout for the memory qubit. Additionally, we do a CR check on Charlie. Charlie communicates the results of the Bell-state measurement to Alice, and Alice performs a feed-forward operation to obtain the teleported state.

To verify the teleported state, we measure the state of Alice in the corresponding basis. To prevent any bias in the tomography we measure in both directions, e.g. when we teleport $|+Z\rangle$ we measure both along +Z and -Z axes.

4.6.2 Experimental setup

The basics of the experimental setup are described in [15]. In the current experiment, Charlie has access to a ^{13}C nuclear spin that acts as a memory qubit. The parameters used for the memory qubits of Bob and Charlie can be found in Table 4.1. Additionally, we have set up a classical communication channel between Charlie and Alice such that Charlie can directly send the results of the Bell-state measurement to Alice.

4.6.3 Tailored heralding of the remote entangled states

In the main text we describe several noise mechanisms that reduce the remote two-node entangled state fidelity. Two of these noise mechanisms, double $|0\rangle$ occupancy and double optical excitation, are accompanied by the emission of an extra photon. This extra photon can be detected using the local phonon-side band (PSB) detectors. By monitoring the PSB detectors, we can real-time-reject false heralding events.

In Figs. 4.6 and 4.7, we plot the histograms of the detection times of the PSB photons conditioned on a simultaneous heralding (zero-phonon line, ZPL) photon detection in the same entanglement generation attempt, for the Alice-Bob and Bob-Charlie entangled link respectively. The correlations are measured in the computational (or Z) basis, and in the X and Y basis. In the computational basis we see the behavior dependent on the detection time of the PSB photon as described in the main text together with the simulations (gray bars). In the X and Y basis, all outcomes are equally probable, and the quantum correlations are washed out.

From the data collected, we can extract the probability to detect these additionally emitted PSB photons. We assume the dark counts of the detectors to be negligible, the PSB detections during the pulse to be fully dominated by the double optical excitation error, and the PSB detections after the pulse to be only caused by double $|0\rangle$ occupancy. By correcting for the PSB detection efficiency, we can estimate the probability for double $|0\rangle$ occupancy and double optical excitation errors. The results are given in Table 4.2. The double $|0\rangle$ state error is expected to occur with probability α. The extracted numbers correspond well to the parameter values we use during remote entanglement generation

($\alpha_{\text{Alice}} = 0.07, \alpha_{\text{Bob}} = 0.05, \alpha_{\text{Charlie}} = 0.10$). The probability for the double optical excitation to occur depends on the shape and the amplitude of the optical excitation pulse, and differs per node.

Numerical model

We compare our PSB detection data (previous section) to a numerical model. We model the NV center as a three level system with two stable ground states $|0\rangle, |1\rangle$ and one excited state $|e\rangle$. The optical $|0\rangle \leftrightarrow |e\rangle$ transition is driven by a resonant laser pulse and is assumed to be a closed transition. The Hamiltonian describing the dynamics of the system in a suitable rotating frame is

$$\hat{H} = \Omega(t)|e\rangle\langle 0| + \Omega^*(t)|0\rangle\langle e|, \tag{4.1}$$

where $\Omega(t)$ describes the (time-dependent) driving of the optical transition. From the excited state, the NV can spontaneously emit a photon and decay to $|0\rangle$. Without specifying the particular mode this photon is emitted in, we simply model such an emission with a Lindblad jump operator of the form $\hat{L}_1 = \sqrt{\gamma}|0, 1_p\rangle\langle e|$. Here γ is the rate of spontaneous emission, $|0, 1_p\rangle$ denotes the state where the NV is in state $|0\rangle$ and one photon was emitted, and we use the convention that when not explicitly stated, there is no emitted photon i.e. $|e\rangle$ denotes the NV in state $|e\rangle$ with zero emitted photons.

To account for double emission errors in the entanglement scheme, we expand the model by letting states $|0, 1_p\rangle, |e, 1_p\rangle$ be coupled by a similar Hamiltonian as in equation (4.1) with the same coupling $\Omega(t)$. Double emission is then captured by a Lindblad jump operator $\hat{L}_2 = \sqrt{\gamma}|0, 2_p\rangle\langle 1_p, e|$. For the specific excitation pulses used in the experiment, we can then numerically solve the Master equation of the system in a basis of $\{|0\rangle, |e\rangle, |0, 1_p\rangle, |e, 1_p\rangle, |0, 2_p\rangle\}$ to obtain the probability of zero (P_0), one (P_1), or two (P_2) photons being emitted from the system ($P_0 + P_1 + P_2 = 1$). Note that in this model, we neglect the probability of emitting more than two photons from the NV.

Assuming an initial state $\sqrt{\alpha}|0\rangle + \sqrt{1-\alpha}|1\rangle$ of the NV center, the state after the optical excitation is then modeled as

$$|\psi\rangle = \sqrt{\alpha}\left(\sqrt{P_0}|0\rangle + \sqrt{P_1}|0, 1_p\rangle + \sqrt{P_2}|0, 2_p\rangle\right) + \sqrt{1-\alpha}|1\rangle. \tag{4.2}$$

The emitted photons are either PSB (= 97%) or ZPL (= 3%) photons. We model this by performing a standard beam splitter transformation on the photonic modes. Letting $\hat{a}^\dagger$ be the creation operator of a photon ($|1_p\rangle = \hat{a}^\dagger|0_p\rangle$), we make the transformation $\hat{a}^\dagger \rightarrow \sqrt{P_z}\hat{a}_z^\dagger + \sqrt{1-P_z}\hat{a}_b^\dagger$, where $\hat{a}_z^\dagger$ ($\hat{a}_b^\dagger$) is the creation operator of a ZPL (PSB) photon and $P_z = 3\%$. Consequently, $|1_p\rangle \rightarrow \sqrt{P_z}|1_z\rangle + \sqrt{1-P_z}|1_b\rangle$, where $|1_z\rangle$ ($|1_b\rangle$) is an emitted ZPL (PSB) photon.

The photons can be emitted either inside or outside the detection time window, i.e. the time interval in which detected photons are accepted. This time interval is in general

different for the PSB and ZPL photons. This results in the following transformations:

$$|1_z\rangle \quad \rightarrow \quad \sqrt{P_{dz,1}}|1_{d,z}\rangle + \sqrt{1-P_{dz,1}}|1_{nd,z}\rangle \tag{4.3}$$

$$|1_b\rangle \quad \rightarrow \quad \sqrt{P_{db,1}}|1_{d,b}\rangle + \sqrt{1-P_{db,1}}|1_{nd,b}\rangle \tag{4.4}$$

$$|2_z\rangle \quad \rightarrow \quad \sqrt{P_{dz,2}}|2_{d,z}\rangle + \sqrt{P_{dz,3}}|1_{d,z}\rangle|1_{nd,z}\rangle + \sqrt{1-P_{dz,2}-P_{dz,3}}|2_{nd,z}\rangle \tag{4.5}$$

$$|2_b\rangle \quad \rightarrow \quad \sqrt{P_{db,2}}|2_{d,b}\rangle + \sqrt{P_{db,3}}|1_{d,b}\rangle|1_{nd,b}\rangle + \sqrt{1-P_{db,2}-P_{db,3}}|2_{nd,b}\rangle \tag{4.6}$$

$$|1_z\rangle|1_b\rangle \quad \rightarrow \quad \sqrt{P_{dzb,1}}|1_{d,z}\rangle|1_{d,b}\rangle + \sqrt{P_{dzb,2}}|1_{nd,z}\rangle|1_{d,b}\rangle \tag{4.7}$$
$$+ \sqrt{P_{dzb,3}}|1_{d,z}\rangle|1_{nd,b}\rangle + \sqrt{1-P_{dz,2}-P_{dz,2}-P_{dz,3}}|1_{nd,z}\rangle|1_{nd,b}\rangle.$$

The probabilities $P_{dz,1}, P_{db,1},...$ are defined in Table 4.3 and are found through the numerical simulation described above.

Finally, we model transmission loss with standard beam splitter transformations acting on the photon modes emitted in the detection window. Letting $\hat{a}_{d,z}^\dagger$ ($\hat{a}_{d,b}^\dagger$) be the creation operator of a ZPL (PSB) photon emitted in the detection time window, we make the transformations

$$\hat{a}_{d,z}^\dagger \rightarrow \sqrt{\eta_z}\hat{a}_{d,z}^\dagger + \sqrt{1-\eta_z}\hat{a}_{nd,z}^\dagger \tag{4.8}$$

$$\hat{a}_{d,z}^\dagger \rightarrow \sqrt{\eta_b}\hat{a}_{d,b}^\dagger + \sqrt{1-\eta_z}\hat{a}_{nd,b}^\dagger. \tag{4.9}$$

where η_z is the total transmission efficiency from the NV to the central beam splitter while η_b is the total transmission and detection efficiency of the PSB photons. The operators $\hat{a}_{nd,z}^\dagger$ and $\hat{a}_{nd,b}^\dagger$ describe the lost/undetected modes. Tracing over the undetected modes, the output state of a single NV can be written as

$$\rho_\psi = \rho_0 \otimes |0\rangle\langle0|_{d,b} + \rho_1 \otimes |1\rangle\langle1|_{d,b} + \rho_2 \otimes |2\rangle\langle2|_{d,b}, \tag{4.10}$$

where we have neglected any coherence between the photonic PSB modes since these are accompanied by undetected non-radiative decay (phonon emission). The unnormalized density matrices ρ_0, ρ_1, and ρ_2 describe the state of the NV center communication qubit and the ZPL photons emitted in the time window of the ZPL detectors and transmitted to the central beam splitter. In the limit $\eta_z \ll 1$, we can neglect terms of $|2_{d,z}\rangle$ and these density matrices will all be of the form

$$\rho_j = \sum_{i=1}^{4} |\phi_{i,j}\rangle\langle\phi_{i,j}|, \tag{4.11}$$

where $|\phi_{i,j}\rangle = (a_{i,j}|1\rangle + b_{i,j}|0\rangle)|0_z\rangle + c_{i,j}|0\rangle|1_{d,z}\rangle$ and $j = 0, 1, 2$. In Equation (4.11), i refers to the different number of lost undetected photons

$\quad i = 1,$ zero photons being lost

$\quad i = 2,$ one ZPL photon being lost

$\quad i = 3,$ one PSB photon being lost

$\quad i = 4,$ two photons being lost, either two ZPL, two PSB or one ZPL and one PSB

and j to the number of detected PSB photons. We note that all $a_{i,1}$ and $a_{i,2}$ will be zero since ρ_1 and ρ_2 are accompanied by PSB photons (see equation (4.10)) meaning that the NV was in state $|0\rangle$. Furthermore, the only non-zero term in ρ_2 will be $b_{i,2}$ since two PSB photons were emitted, meaning that no ZPL photon was emitted since we neglect higher order emissions.

The only term in equation (4.10) from which remote spin-spin entanglement between two NVs can be created is $\rho_0 \otimes |0\rangle\langle 0|_{d,b}$ since this does not have any detected PSB photons. However, PSB and ZPL photons that were emitted but not detected will still decrease the entangled state fidelity. Such events are responsible for the contributions of $|\phi_{2,0}\rangle, |\phi_{3,0}\rangle$ and $|\phi_{4,0}\rangle$ in ρ_0. The only term where no PSB photons were emitted and no ZPL photons were undetected is $|\phi_{1,0}\rangle = \sqrt{1-\alpha}|1\rangle|0_{zpl}\rangle + \sqrt{\alpha}|0\rangle(\sqrt{P_0}|0_{zpl}\rangle + \sqrt{P_1 P_{zpl} P_{d,zpl}}|1_{d,zpl}\rangle)$.

The combined state from the two NV centers before the central beam splitter is $\rho_\psi \otimes \tilde{\rho}_\psi$, where $\tilde{\rho}_\psi$ (the state of the second NV) is of the same form as in equation (4.10) but including that parameters such as initial rotation (α), driving strength (Ω) and transmission efficiencies (η_z, η_b) can be different for the two centers. Furthermore, we include a phase difference between the two paths to the central beam splitter. The central beam splitter is modeled as a perfect 50:50 beam splitter and the finite detection efficiency of the output detectors is assumed to be equal and can be directly included in the transmission efficiencies (η_z) while dark counts are negligible in the experiment and not included. Finally, we include non-perfect visibility between the ZPL photons by reducing the coherence between the output modes of the beam splitter by a factor v. This visibility is estimated from experimental data and can e.g. originate from slightly off-resonant driving of the NV centers.

4.6.4 Memory qubit coherence Bob

We use the sequence described in Fig. 4.3a of the main text to preserve the state of the memory qubit during entanglement attempts. To characterize the decoupling sequence, we compare it to the sequence where we do not apply the decoupling pulse on the memory qubit and/or the sequence where we idle instead of performing entanglement attempts. We characterize the coherence of the memory qubit by storing the six cardinal states. We average the results for the eigenstates ($|0\rangle, |1\rangle$) and superposition states ($|\pm X\rangle$ and $|\pm Y\rangle$). In Fig. 4.8 we plot the Bloch vector length $b = \sqrt{b_x^2 + b_y^2 + b_z^2}$ with b_i the Bloch vector component in direction i.

Over the measured range, the eigenstates show little decay. The decay of the superposition states is fitted with the function $f(x) = Ae^{-(x/N_{1/e})^n}$. The fitted parameters can be found in Table 4.4.

The use of the decoupling pulse π_M on the memory qubit increases the $N_{1/e}$ by more than a factor 6. Moreover, the initial Bloch vector length A is higher with the π_M pulse. This is mainly explained by the second round of phase stabilization [15] in between swapping the state onto the memory qubit and starting the entanglement generation process. The phase stabilization takes $\approx 350\,\mu s$ and during this time the memory qubit is subject to intrinsic $T_2^\star$ dephasing, which can be efficiently decoupled using the π_M pulse.

4.6.5 Communication qubit coherence

In various parts of the protocol, we decouple the communication qubits from the spin bath environment to extend their coherence time. On Alice, we start the decoupling when the first entangled link is established and stop when the results of the Bell-state measurement to teleport the state are sent by Charlie. On Bob, we decouple the communication qubit when the memory qubit is being re-phased. On Charlie, the communication qubit is decoupled from the point that entanglement with Bob is heralded up to the point where Bob has finished the Bell-state measurement, performed the CR check and has communicated the results. All these decoupling times are dependent on how many entanglement attempts are needed to generate the entangled link between Bob and Charlie.

We characterize the average state fidelities for different decoupling times, see Fig. 4.9. We investigate eigenstates and superposition states separately. We fit the fidelity with the function $f(t) = Ae^{-(t/\tau_{coh})^n} + 0.5$. The fitted parameters are summarized in Table 4.5. For each setup, the minimum and maximum used decoupling times are indicated by the shaded regions in Fig. 4.9. The left-most border is the decoupling time when the first entanglement attempt on Bob and Charlie would be successful, the right-most border when the last attempt before the timeout of 1000 attempts would herald the entangled state.

4.6.6 Basis-alternating repetitive readout

In the main text we discuss the basis-alternating repetitive readout and the results on Bob's memory qubit are shown in Fig. 4.3. Here we show the results for Charlie's memory qubit. We assign the state using the first readout and only accept the result when the consecutive readouts give a consistent pattern. The results for two different initial states of the memory qubit are plotted in Fig. 4.10. We model the expected performance with a Monte Carlo simulation which takes into account the electron readout fidelities, the initial state populations and gate errors, see Ref. [34]. In the case of unconditional teleportation, the state is assigned using the first readout and is accepted regardless of the second readout result.

4.6.7 Teleportation results

The numerical values of the data displayed in Fig. 4.4b-c in the main text can be found in Table 4.6 and Table 4.7, respectively.

4.6.8 Data acquisition and experimental rates

For the data acquisition, we interleaved blocks of measurements with calibrations. We collected the data in blocks of ≈1 hour. In total, we have acquired 79 blocks of data, and we measured 2272 events ($|+X\rangle$ 382, $|-X\rangle$ 385, $|+Y\rangle$ 385, $|-Y\rangle$ 378, $|+Z\rangle$ 375, $|-Z\rangle$ 367) for the conditional teleportation over a time span of 21 days. We can determine the experimental rate including all overhead (such as CR checks, communication time and phase stabilization) by dividing the number of measured data points by the total measurement time. In Fig. 4.11 we plot the experimental rate for both the conditional and unconditional teleportation sequence. In the case of the unconditional teleportation, we accept all Bell-state measurement outcomes on Charlie and therefore the experimental rate is higher. For shorter detection windows during the two-node entanglement, the success probability per attempt is smaller, and thus the experimental rate is lower.

4.6.9 Model of the teleported state

A detailed model of the teleported state can be found at Ref. [34]. The model comprises elements from [15] and is further extended for the teleportation protocol. We take the following noise sources into account

- imperfect Bell states between Alice and Bob, and between Bob and Charlie,

- dephasing of the memory qubit of Bob during entanglement generation between Bob and Charlie,

- depolarizing noise on the memory qubits of Bob and Charlie, due to imperfect initialization and swap gates,

- readout errors on the communication qubits of Bob and Charlie and readouts errors on the memory qubits of Bob and Charlie when using the basis-alternating readout scheme which result in incorrect feed-forward gate operations after the Bell-state measurements,

- depolarizing noise on Alice during the decoupling sequence,

- ionization probability on Alice.

An overview of the input parameters and the effect of the different error sources is given in Tables 4.8 and 4.9.

4.6.10 Effect of the three key innovations on the teleported state fidelity and experimental rate

We assess the effect of each innovation on the teleportation protocol. First, we estimate the average state fidelity and experimental rate with a set of baseline parameters based on the performance in [15]. We use a timeout of 1000 entanglement attempts for the second link (between Bob and Charlie) before aborting the protocol and starting over. In both Bell-state measurements, we continue the sequence for the outcomes "00" and "01" (communication qubit, memory qubit), or abort and start over (in the case of conditional teleportation). Then we incrementally add (1) the basis-alternating repetitive readout scheme for the memory qubits, (2) the improved memory qubit coherence and (3) the tailored heralding scheme of the remote entanglement generation. The results are summarized in Table 4.10.

4.6.11 Estimated fidelity of state to be teleported

The state to be teleported is prepared on the communication qubit of Charlie. Errors in the preparation originate from imperfect initialization and imperfect MW pulses, which are estimated to be $p_{\text{init}} = 1.2 \times 10^{-3}$ and $p_{\text{MW}} = 8 \times 10^{-3}$ [35]. Averaged over the six cardinal states, we estimate the state preparation fidelity to be ≈ 0.995.

4.6.12 Calculation of teleported state fidelity without feed-forward operation

In Fig. 4.4c in the main text we show the fidelity of the teleported state in case no feed-forward operations would have been applied on Alice. To extract this data we follow the

Table 4.1: Memory qubit characteristics. In each setup we use a magnetic field with strength B_z aligned to the NV axis. The nuclear spin precession frequencies ($\omega_{m_S=0}$ and $\omega_{m_S=-1}$) depends on the electron spin state. From the frequency difference the parallel component $A_\parallel$ of the hyperfine interaction can be determined. Conditional (unconditional) pulses are applied by doing N_{con} (N_{unc}) pulses on the electron spin with an inter-pulse delay of τ_{con} (τ_{unc}).

Setup	B_z	$\omega_{m_S=0}$	$\omega_{m_S=-1}$	$A_\parallel$
Bob	189 mT	$2\pi \times$ 2025 kHz	$2\pi \times$ 2056 kHz	$2\pi \times$ 30 kHz
Charlie	16.5 mT	$2\pi \times$ 177 kHz	$2\pi \times$ 240 kHz	$2\pi \times$ 63 kHz

Setup	τ_{con}	N_{con}	τ_{unc}	N_{unc}
Bob	2.818 μs	54	4.165 μs	144
Charlie	6.003 μs	56	11.996 μs	30

Table 4.2: Estimated probabilities for the double optical excitation error and the double $|0\rangle$ occupancy error per node.

Node	Double optical excitation probability	Double $|0\rangle$ occupancy probability
Alice	4.1(5)	7.6(4)
Bob (with Alice)	2.6(3)	4.9(3)
Bob (with Charlie)	6.9(8)	4.7(8)
Charlie	5.7(4)	9.4(4)

same method as in [11]. We perform classical bit flips on the measurement outcomes to counteract the effect of the feed-forward gate operations (as if the gate was not applied) for each Bell-state measurement outcome. We do this for all six cardinal states and compute the average fidelity. We assume the errors of the gate in the feed-forward operations to be small.
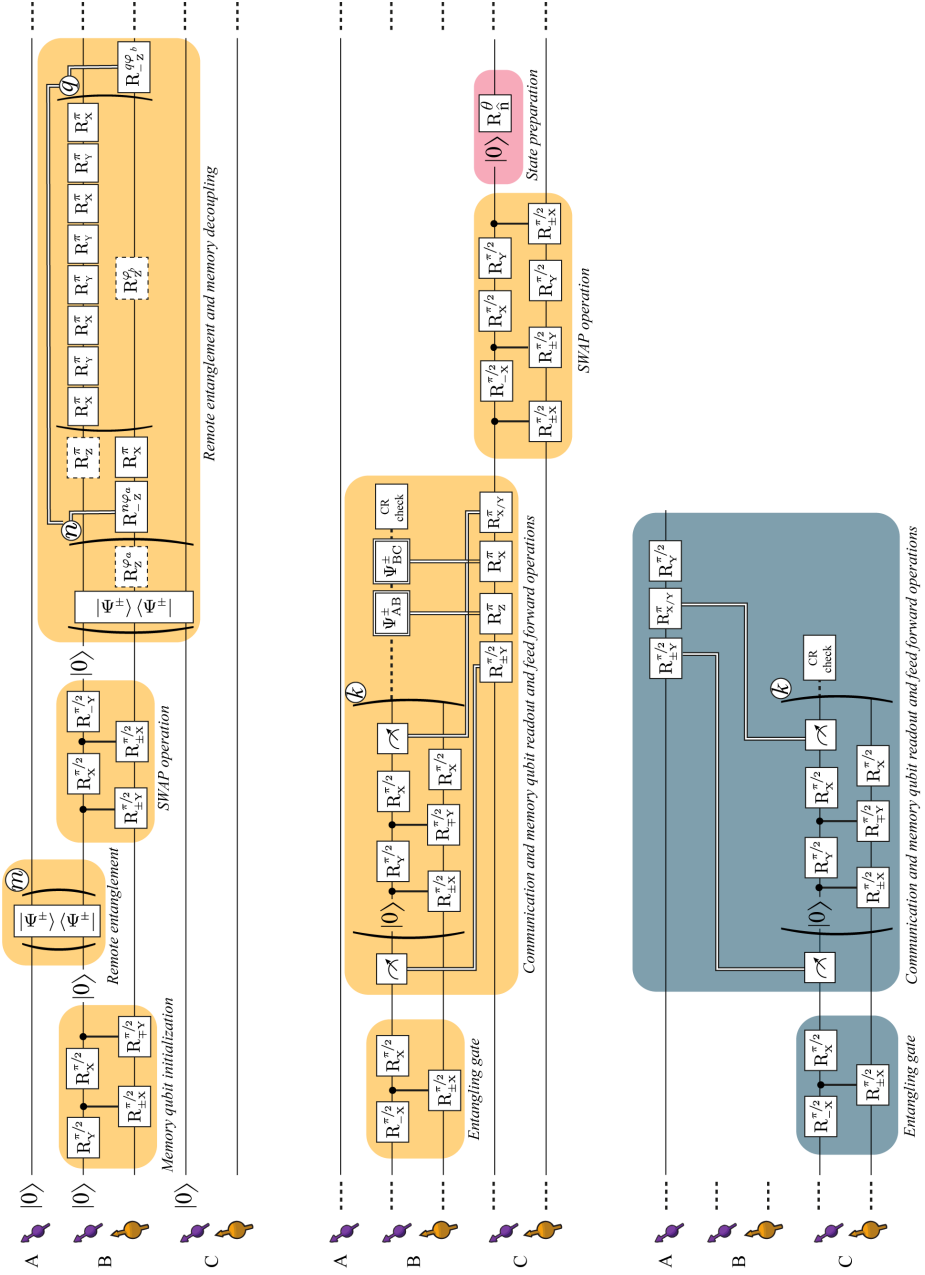
Figure 4.5: Full gate circuit for the teleportation protocol. See text for the description of each element.

4
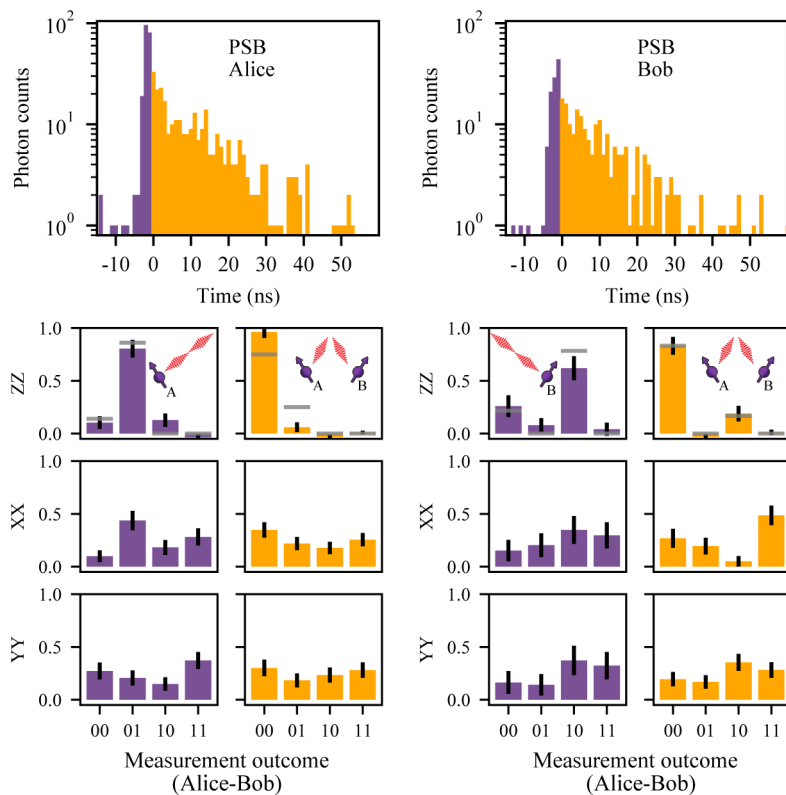
Figure 4.6: (Top) Histograms of the detected PSB photons conditioned on a simultaneous ZPL detection in the entanglement generation attempt, for Alice(left) and Bob(right). (Bottom) Corresponding measured correlations in all bases.
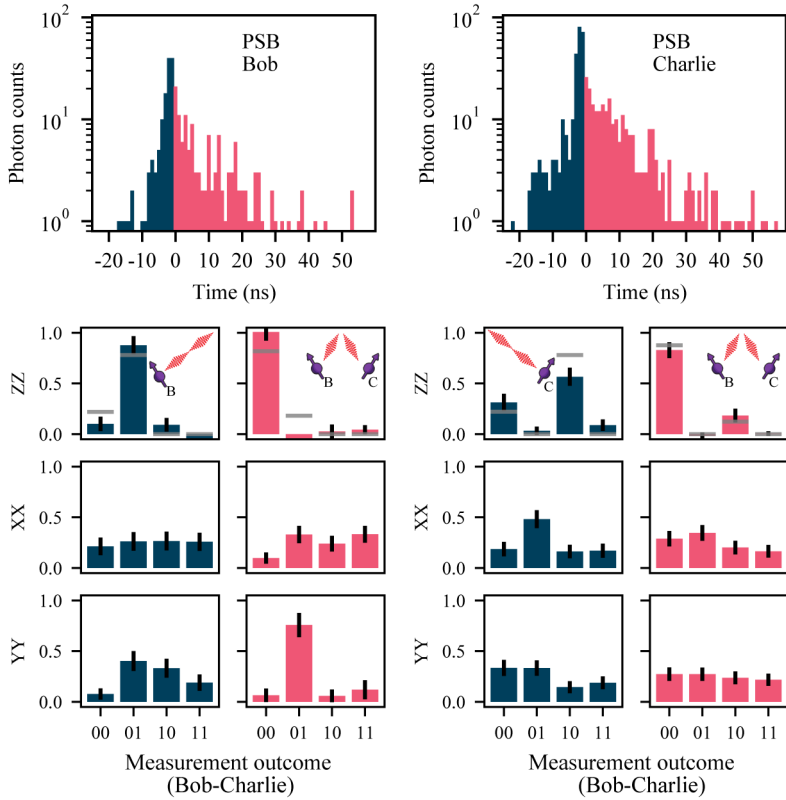
Figure 4.7: (Top) Histograms of the detected PSB photons conditioned on a simultaneous ZPL detection in the entanglement generation attempt, for Bob(left) and Charlie(right). (Bottom) Corresponding measured correlations in all bases.

Table 4.3: Explanation of the parameters used in the numerical simulation of the entanglement generation protocol.

Parameter	Description
γ	Spontaneous emission rate of the excited state.
$\Omega(t)$	Optical driving strength.
α	Initial population of the $|0\rangle$ state.
P_0	Probability of emitting 0 photons (ZPL or PSB).
P_1	Probability of emitting 1 photons (ZPL or PSB).
P_2	Probability of emitting 2 photons (ZPL or PSB or both).
P_z	Probability that an emitted photon is a ZPL photon.
$P_{dz,1}$	Probability that a ZPL photon is within the ZPL detection window, conditioned on a single ZPL photon being emitted.
$P_{db,1}$	Probability that a PSB photon is within the PSB detection window, conditioned on a single PSB photon being emitted.
$P_{dz,2}$	Probability that 2 ZPL photons are within the ZPL detection window, conditioned on two ZPL photons being emitted.
$P_{dz,3}$	Probability that one ZPL photons is within the ZPL detection window and one is not, conditioned on two ZPL photons being emitted.
$P_{db,2}$	Probability that 2 PSB photons are within the PSB detection window, conditioned on two PSB photon being emitted.
$P_{dz,3}$	Probability that one PSB photons is within the PSB detection window and one is not, conditioned on two PSB photons being emitted.
$P_{dzb,1}$	Probability that a ZPL photon is within the ZPL detection window and a PSB photon is within the PSB detection window, conditioned on one ZPL and one PSB photon being emitted.
$P_{dzb,2}$	Probability that a ZPL photon is not within the ZPL detection window and a PSB photon is within the PSB detection window, conditioned on one ZPL and one PSB photon being emitted.
$P_{dzb,3}$	Probability that a ZPL photon is within the ZPL detection window and a PSB photon is not within the PSB detection window, conditioned on one ZPL and one PSB photon being emitte.
η_z	Total transmission and detection efficiency of ZPL photons.
η_p	Total transmission and detection efficiency of PSB photons.

Table 4.4: Fitted parameters for the memory coherence decay of the superposition states.

	A	$N_{1/e}$	n
With ent. att. with π_M	0.875(15)	5327(319)	1.13(11)
With ent. att. without π_M	0.806(19)	848(39)	1.21(9)
Without ent. att. with π_M	0.884(11)	5239(163)	1.94(16)
Without ent. att. without π_M	0.807(19)	880(34)	1.37(10)
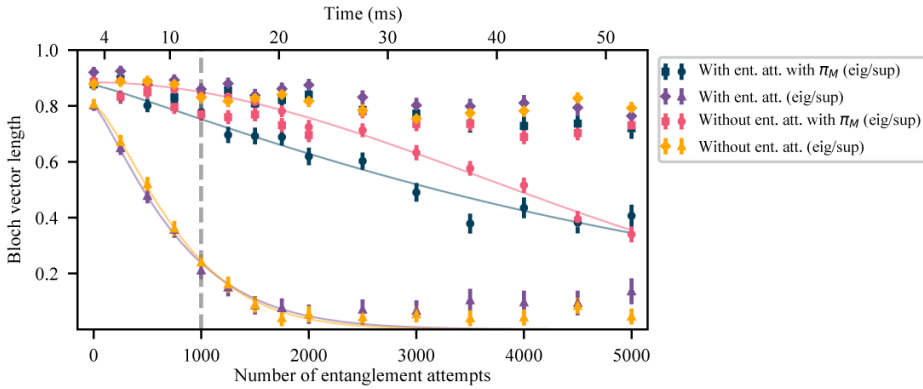
Figure 4.8: Coherence of Bob's memory qubit for superposition states (triangles and circles) and eigenstates (squares and diamonds). We perform the sequence as described in the main text with and without the decoupling pulse π_M on the memory qubit, the dark blue and purple points respectively. Additionally, we perform the sequence with a wait time instead of entanglement attempts with (pink points) and without the decoupling pulse (yellow points). The gray dashed line indicates the timeout of the entanglement generation process used in the teleportation protocol.
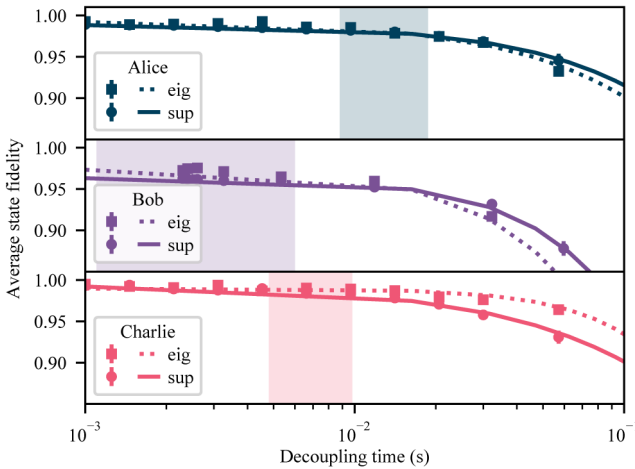
Figure 4.9: Decoupling of the communication qubits. The average state fidelity is plotted for different decoupling times for each setup. The shaded area represent the decoupling times used in the teleportation protocol.

Table 4.5: Fitted parameters for average state fidelity state during communication qubit decoupling.

		A	τ_{coh}(s)	n
Alice	Eigenstate	0.4930(13)	0.459(12)	1.04(3)
	Superposition	0.4889(18)	0.540(20)	1.07(5)
Bob	Eigenstate	0.4738(11)	0.130(3)	1.41(4)
	Superposition	0.4634(15)	0.177(6)	1.47(6)
Charlie	Eigenstate	0.4897(9)	0.357(7)	1.67(6)
	Superposition	0.4936(19)	0.560(20)	0.92(4)
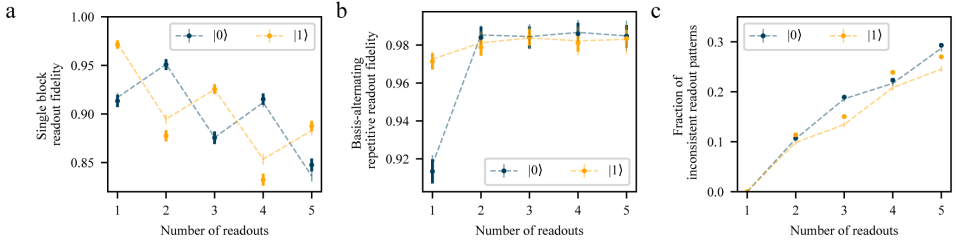
Figure 4.10: Basis-alternating repetitive (BAR) readout results for Charlie's memory qubit. **a** Measured fraction of memory qubit states that were assigned "0" per readout block, for initialization in $|0\rangle$ and in $|1\rangle$. **b** Readout fidelity of the basis-alternating repetitive readout scheme for different number of readout blocks. **c** Fraction of inconsistent readout patterns for different number of readout blocks. The dashed lines represent a numerical model using measured parameters.

Table 4.6: Numerical values of the data displayed in Fig. 4.4b.

	Teleported state fidelity
X	0.760(24)
-X	0.745(25)
Y	0.656(27)
-Y	0.651(27)
Z	0.731(26)
-Z	0.671(27)
Average	0.702(11)

Table 4.7: Numerical values of the data displayed in Fig. 4.4c.

Bell-state measurement outcome (memory qubit, communication qubit)	Average teleported state fidelity
00	0.707(15)
01	0.696(14)
10	0.698(15)
11	0.671(14)
No feed forward	0.501(7)
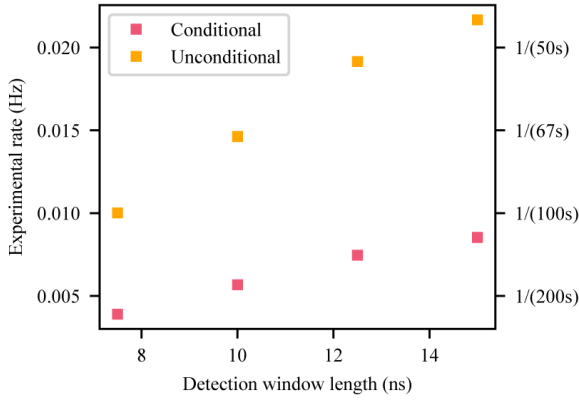
Figure 4.11: Experimental rates of the conditional and unconditional teleportation protocol for different detection window lengths in the two-node entanglement generation.

Table 4.8: Overview of parameters used in the simulations for the two-node entangled states. The error due to the non-zero bright state populations is a result of the single click protocol. For the other error sources, we compute the estimated infidelity as if it was the only error source present apart from the protocol error. This allows easy comparison between the different error sources.

	Parameter AB	Parameter BC	Infidelity Ψ_{AB}	Infidelity Ψ_{BC}
Detection window length	15 ns	15 ns		
Detection probability setup 1	3.4×10^{-4}	4.3×10^{-4}		
Detection probability setup 2	5.1×10^{-4}	2.4×10^{-4}		
Average detection probability PSB	0.1	0.12		
Bright state populations (α_1, α_2)	(0.07, 0.05)	(0.05, 0.1)	5.5×10^{-2}	6.7×10^{-2}
Dark count rate	10 Hz	10 Hz	5.1×10^{-3}	5.3×10^{-3}
Visibility	0.90	0.90	2.4×10^{-2}	2.4×10^{-2}
Average double excitation probability	0.06	0.08	5.5×10^{-2}	7.1×10^{-2}
Optical phase uncertainty	21^o	12^o	3.1×10^{-2}	1.0×10^{-2}
All error sources combined			0.16	0.17

Table 4.9: Overview of parameters used in the simulations for the average teleported state fidelity in case of a conditional Bell state measurement on Charlie. For each error source, we compute the estimated infidelity as if it was the only error source present, apart from the single click protocol errors of the two-node entangled states. This allows easy comparison between the different error sources.

4

	Parameter	Infidelity
Ionization probability Alice	0.7%	0.6×10^{-2}
Depolarizing noise Alice	0.04	1.7×10^{-2}
Depolarizing noise memory qubit Bob	0.12	5.0×10^{-2}
Dephasing noise memory qubit Bob ($N_{1/e}, n$)	(5300, 1.1)	2.1×10^{-2}
Depolarizing noise memory qubit Charlie	0.14	5.9×10^{-2}
Readout fidelities memory qubit Bob ($|0\rangle, |1\rangle$)	(0.99, 0.99)	0.6×10^{-2}
Readout fidelities communication qubit Bob ($|0\rangle, |1\rangle$)	(0.93, 0.995)	0.3×10^{-2}
Readout fidelities memory qubit Charlie ($|0\rangle, |1\rangle$)	(0.98, 0.98)	1.1×10^{-2}
Readout fidelities communication qubit Charlie ($|0\rangle, |1\rangle$)	(0.92, 0.99)	0.6×10^{-2}
Two-node entangled states combined		0.192
All error sources combined		0.305

Table 4.10: Simulated effect of the innovations on the teleported state fidelity and experimental rate.

	Fidelity	Rate (Hz)
Baseline parameters using timeout = 1000, BSM outcomes (communication qubit, memory qubit) = "00" or "01"	0.666	1/(53s)
With basis-alternating repetitive readout	0.679	1/(73s)
With improved memory coherence	0.687	1/(73s)
With tailored heralding scheme	0.695	1/(74s)

References

[1] C. H. Bennett, G. Brassard, C. Crépeau, R. Jozsa, A. Peres, and W. K. Wootters, *Teleporting an unknown quantum state via dual classical and Einstein-Podolsky-Rosen channels,* Physical Review Letters **70**, 1895 (1993).

[2] S. Wehner, D. Elkouss, and R. Hanson, *Quantum internet: A vision for the road ahead,* Science **362** (2018), 10.1126/science.aam9288.

[3] M. Ben-Or, C. Crepeau, D. Gottesman, A. Hassidim, and A. Smith, *Secure Multiparty Quantum Computation with (Only) a Strict Honest Majority,* in *2006 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS'06)* (2006) pp. 249–260.

[4] A. S. Arora, J. Roland, and S. Weis, *Quantum weak coin flipping,* in *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2019 (Association for Computing Machinery, New York, NY, USA, 2019) pp. 205–216.

[5] D. Bouwmeester, J.-W. Pan, K. Mattle, M. Eibl, H. Weinfurter, and A. Zeilinger, *Experimental quantum teleportation,* Nature **390**, 575 (1997).

[6] D. Boschi, S. Branca, F. De Martini, L. Hardy, and S. Popescu, *Experimental Realization of Teleporting an Unknown Pure Quantum State via Dual Classical and Einstein-Podolsky-Rosen Channels,* Physical Review Letters **80**, 1121 (1998).

[7] A. Furusawa, J. L. Sørensen, S. L. Braunstein, C. A. Fuchs, H. J. Kimble, and E. S. Polzik, *Unconditional Quantum Teleportation,* Science (1998).

[8] S. Olmschenk, D. N. Matsukevich, P. Maunz, D. Hayes, L.-M. Duan, and C. Monroe, *Quantum Teleportation Between Distant Matter Qubits,* Science (2009).

[9] C. Nölleke, A. Neuzner, A. Reiserer, C. Hahn, G. Rempe, and S. Ritter, *Efficient Teleportation Between Remote Single-Atom Quantum Memories,* Physical Review Letters **110**, 140403 (2013).

[10] S. Langenfeld, S. Welte, L. Hartung, S. Daiss, P. Thomas, O. Morin, E. Distante, and G. Rempe, *Quantum Teleportation between Remote Qubit Memories with Only a Single Photon as a Resource,* Physical Review Letters **126**, 130502 (2021).

[11] W. Pfaff, B. J. Hensen, H. Bernien, S. B. v. Dam, M. S. Blok, T. H. Taminiau, M. J. Tiggelman, R. N. Schouten, M. Markham, D. J. Twitchen, and R. Hanson, *Unconditional quantum teleportation between distant solid-state quantum bits,* Science **345**, 532 (2014).

[12] X.-H. Bao, X.-F. Xu, C.-M. Li, Z.-S. Yuan, C.-Y. Lu, and J.-W. Pan, *Quantum teleportation between remote atomic-ensemble quantum memories,* Proceedings of the National Academy of Sciences **109**, 20347 (2012).

[13] C. Cabrillo, J. I. Cirac, P. García-Fernández, and P. Zoller, *Creation of entangled states of distant atoms by interference,* Physical Review A **59**, 1025 (1999).

4

[14] S. Bose, P. L. Knight, M. B. Plenio, and V. Vedral, *Proposal for Teleportation of an Atomic State via Cavity Decay,* Physical Review Letters **83**, 5158 (1999).

[15] M. Pompili, S. L. N. Hermans, S. Baier, H. K. C. Beukers, P. C. Humphreys, R. N. Schouten, R. F. L. Vermeulen, M. J. Tiggelman, L. d. S. Martins, B. Dirkse, S. Wehner, and R. Hanson, *Realization of a multinode quantum network of remote solid-state qubits,* Science **372**, 259 (2021).

[16] P. C. Humphreys, N. Kalb, J. P. J. Morits, R. N. Schouten, R. F. L. Vermeulen, D. J. Twitchen, M. Markham, and R. Hanson, *Deterministic delivery of remote entanglement on a quantum network,* Nature **558**, 268 (2018).

[17] T. Legero, T. Wilk, A. Kuhn, and G. Rempe, *Time-resolved two-photon quantum interference,* Applied Physics B **77**, 797 (2003).

[18] C. Bradley, J. Randall, M. Abobeih, R. Berrevoets, M. Degen, M. Bakker, M. Markham, D. Twitchen, and T. Taminiau, *A Ten-Qubit Solid-State Spin Register with Quantum Memory up to One Minute,* Physical Review X **9**, 031045 (2019).

[19] J. Cramer, N. Kalb, M. A. Rol, B. Hensen, M. S. Blok, M. Markham, D. J. Twitchen, R. Hanson, and T. H. Taminiau, *Repeated quantum error correction on a continuously encoded qubit by real-time feedback,* Nature Communications **7**, 11526 (2016).

[20] L. Robledo, L. Childress, H. Bernien, B. Hensen, P. F. A. Alkemade, and R. Hanson, *High-fidelity projective read-out of a solid-state spin quantum register,* Nature **477**, 574 (2011).

[21] L. Jiang, J. S. Hodges, J. R. Maze, P. Maurer, J. M. Taylor, D. G. Cory, P. R. Hemmer, R. L. Walsworth, A. Yacoby, A. S. Zibrov, and M. D. Lukin, *Repetitive Readout of a Single Electronic Spin via Quantum Logic with Nuclear Spin Ancillae,* Science **326**, 267 (2009).

[22] S. J. van Enk, N. Lütkenhaus, and H. J. Kimble, *Experimental procedures for entanglement verification,* Physical Review A **75**, 052318 (2007).

[23] R. Van Meter, *Quantum networking*, Networks and telecommunications series (John Wiley & Sons, 2014).

[24] A. Broadbent, J. Fitzsimons, and E. Kashefi, *Universal Blind Quantum Computation,* in *2009 50th Annual IEEE Symposium on Foundations of Computer Science* (2009) pp. 517–526.

[25] B. C. Rose, D. Huang, Z.-H. Zhang, P. Stevenson, A. M. Tyryshkin, S. Sangtawesin, S. Srinivasan, L. Loudin, M. L. Markham, A. M. Edmonds, D. J. Twitchen, S. A. Lyon, and N. P. d. Leon, *Observation of an environmentally insensitive solid-state spin defect in diamond,* Science **361**, 60 (2018).

[26] C. Nguyen, D. Sukachev, M. Bhaskar, B. Machielse, D. Levonian, E. Knall, P. Stroganov, R. Riedinger, H. Park, M. Lončar, and M. Lukin, *Quantum Network Nodes Based on Diamond Qubits with an Efficient Nanophotonic Interface,* Physical Review Letters **123**, 183602 (2019).

[27] M. E. Trusheim, B. Pingault, N. H. Wan, M. Gündoğan, L. De Santis, R. Debroux, D. Gangloff, C. Purser, K. C. Chen, M. Walsh, J. J. Rose, J. N. Becker, B. Lienhard, E. Bersin, I. Paradeisanos, G. Wang, D. Lyzwa, A. R.-P. Montblanch, G. Malladi, H. Bakhru, A. C. Ferrari, I. A. Walmsley, M. Atatüre, and D. Englund, *Transform-Limited Photons From a Coherent Tin-Vacancy Spin in Diamond,* Physical Review Letters **124**, 023602 (2020).

[28] N. T. Son, C. P. Anderson, A. Bourassa, K. C. Miao, C. Babin, M. Widmann, M. Niethammer, J. Ul Hassan, N. Morioka, I. G. Ivanov, F. Kaiser, J. Wrachtrup, and D. D. Awschalom, *Developing silicon carbide for quantum spintronics,* Applied Physics Letters **116**, 190501 (2020).

[29] D. M. Lukin, M. A. Guidry, and J. Vučković, *Integrated Quantum Photonics with Silicon Carbide: Challenges and Prospects,* PRX Quantum **1**, 020102 (2020).

[30] J. M. Kindem, A. Ruskuc, J. G. Bartholomew, J. Rochman, Y. Q. Huan, and A. Faraon, *Control and single-shot readout of an ion embedded in a nanophotonic cavity,* Nature **580**, 201 (2020).

[31] S. Chen, M. Raha, C. M. Phenicie, S. Ourari, and J. D. Thompson, *Parallel single-shot measurement and coherent control of solid-state spins below the diffraction limit,* Science **370**, 592 (2020).

[32] M. Ruf, N. H. Wan, H. Choi, D. Englund, and R. Hanson, *Quantum networks based on color centers in diamond,* Journal of Applied Physics **130**, 070901 (2021).

[33] A. Dahlberg, M. Skrzypczyk, T. Coopmans, L. Wubben, F. Rozpędek, M. Pompili, A. Stolk, P. Pawełczak, R. Knegjens, J. de Oliveira Filho, R. Hanson, and S. Wehner, *A link layer protocol for quantum networks,* in *Proceedings of the ACM Special Interest Group on Data Communication*, SIGCOMM '19 (Association for Computing Machinery, New York, NY, USA, 2019) pp. 159–173.

[34] S. Hermans, M. Pompili, H. Beukers, S. Baier, J. Borregaard, and R. Hanson, *Data and software supporting "Qubit teleportation between non-neighboring nodes in a quantum network",* Available after peer review is complete (2021), 10.4121/16645969.v1.

[35] N. Kalb, P. C. Humphreys, J. J. Slim, and R. Hanson, *Dephasing mechanisms of diamond-based nuclear-spin memories for quantum networks,* Physical Review A **97**, 062330 (2018).

4

5

Entanglement delivery with a quantum network stack

**M. Pompili*, C. Delle Donne*, I. te Raa, B. van der Vecht,
M. Skrzypczyk, G. Ferreira, L. de Kluijver, A. J. Stolk,
S. L. N. Hermans, P. Pawełczak, W. Kozlowski & R. Hanson,
S. Wehner**

Scaling current quantum communication demonstrations to a large-scale quantum network will require not only advancements in quantum hardware capabilities, but also robust control of such devices to bridge the gap to user demand. Moreover, the abstraction of tasks and services offered by the quantum network should enable platform-independent applications to be executed without knowledge of the underlying physical implementation. Here we experimentally demonstrate, using remote solid-state quantum network nodes, a link layer and a physical layer protocol for entanglement-based quantum networks. The link layer abstracts the physical-layer entanglement attempts into a robust, platform-independent entanglement delivery service. The system is used to run full state tomography of the delivered entangled states, as well as preparation of a remote qubit state on a server by its client. Our results mark a clear transition from physics experiments to quantum communication systems, which will enable the development and testing of components of future quantum networks.

B y sharing entangled states over large distances, the future Quantum Internet [1, 2] can unlock new possibilities in secure communication [3], distributed and blind quantum computation [4, 5], and metrology [6, 7]. Fundamental primitives for entanglement-based quantum networks have been demonstrated across several physical platforms, including trapped ions [8, 9], neutral atoms [10, 11], diamond color centers [12–15], and quantum dots [16, 17]. To scale up such physics experiments to intermediate-scale quantum networks, researchers have been investigating how to enclose the complex nature of quantum entanglement generation into more robust abstractions [18–24].

A natural way to render a complex system scalable, is to design its architecture as a stack of layers that go from specialized physical medium protocols to more general services. Inspired by the TCP/IP protocol stack commonly employed in classical networks, similar stacks have been proposed for quantum networks [19–21], like the one depicted in Fig. 5.1. These recent efforts, along with proposals for resource scheduling and routing techniques (e.g. [25–31]), pave the way for larger-scale quantum networks.

In this work we experimentally demonstrate—for the first time—a link layer protocol for entanglement-based quantum networks. The link layer abstracts the generation of entangled states between two physically separated solid-state qubits into a robust and platform-independent service. An application can request entangled states from the link layer and then, in addition, apply local quantum operations on the entangled qubits in real-time. Using the link layer, we perform full state tomography of the generated states and achieve remote state preparation—a building block for blind quantum computation—as well as measuring the latency of the entanglement generation service.
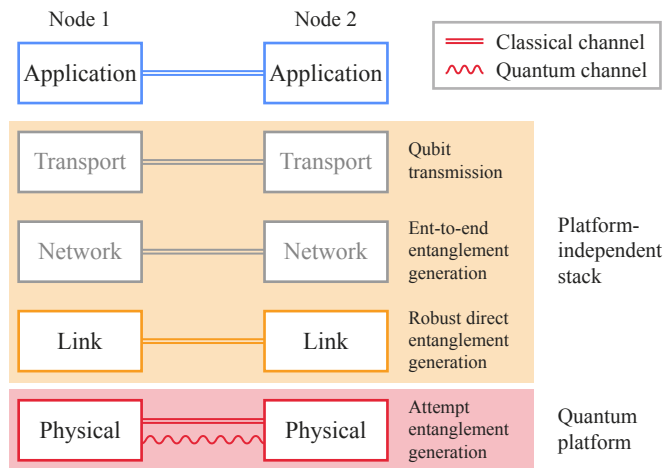
Figure 5.1: Quantum network stack architecture. At the bottom of the stack, the physical layer (red), which is highly quantum platform-dependent, is tasked with attempting entanglement generation. The link layer (yellow) uses the functionality provided by the physical layer to provide a platform-independent and robust entanglement generation service between neighboring nodes to the higher layers. Network and transport layer (not implemented in this work, grayed out) will support end-to-end connectivity and qubit transmission. Applications (blue) use the services offered by the stack to perform quantum networking tasks. Based on Dahlberg *et al.* [19].

To evaluate correct operation and performance of our system, we measure (a) the fidelity of the generated states and (b) the latency incurred by link layer and physical layer when generating entangled pairs. For both fidelity and latency, we find that our system performs with marginal overhead with respect to previous non-platform-independent experiments. We also identify the sources of the additional overhead incurred, and propose improvements for future realizations.

5.1 Quantum Link Layer Protocol

Remote entanglement generation constitutes a fundamental building block of quantum networking. However, for a user to be able to integrate it into more complex quantum networking applications and protocols, the entanglement generation service must also be: (a) robust, meaning that the user should not have to deal with entanglement failures and retries, and that an entanglement request should result in the delivery of an entangled pair; (b) quantum platform-independent, in order for the user to be able to request entanglement without having to understand the inner workings of the underlying physical implementation; (c) on-demand, such that the user can request and consume entanglement as part of a larger quantum communication application. Robust, platform-independent, on-demand entanglement generation must figure as one of the basic services offered by a system running on a quantum network node. In other words, establishing a reliable quantum link between two directly connected nodes is the task of the first layer above the physical layer in a quantum networking protocol stack, as portrayed in Fig. 5.1. Following the TCP/IP (Internet protocol suite) model nomenclature, we refer to this layer as the *link layer*. We remark that, in the framework of a multi-node network, a quantum network stack should also feature a *network layer* (called *internet layer* in the TCP/IP model) to establish links between non-adjacent nodes, and optionally a *transport layer* to encapsulate qubit transmission into a service [19–21] (as shown in fig. 5.1).

5.1.1 Link Layer Service

The service provided by a link layer protocol for quantum networks should expose a few configuration parameters to its user. To ensure a platform-independent interaction with the link layer, such parameters should be common to all possible implementations of the quantum physical device. In this work, we implement a revised version of the link layer protocol proposed—but not implemented—in Ref. [19], with the following service description. The interface exposed by the link layer should allow the higher layer to specify: (a) *Remote node ID*, an identifier of the remote node to produce entanglement with (in case the requesting node has multiple neighbors); (b) *Number of entangled pairs*, to allow for the creation of several pairs with one request; (c) *Minimum fidelity*, an indication of the desired minimum fidelity for the produced pairs; (d) *Delivery type*, whether to keep the produced pair for future use (type *K*), measure it directly after creation (type *M*), or measure the local qubit immediately and instruct the remote node to keep its own for future use (type *R*, used for remote state preparation); (e) *Measurement basis*, the basis to use when measuring *M*- or *R*-type entangled pairs; (f) *Request timeout*, to indicate a time limit for the processing of the request. After submitting an entanglement generation request, the user should expect the link layer to coordinate with the remote node and to handle

entanglement generation attempts and retries until all the desired pairs are produced (or until the timeout has expired). When completing an entanglement generation request, the link layer should then report to the above layer the following: (a) *Produced Bell state*, the result of entanglement generation; (b) *Measurement outcome*, in case of *M*- or *R*-type entanglement requests; (c) *Entanglement ID*, to uniquely identify an entangled pair consistently across source and destination of the request.

5.1.2 A Quantum Link Layer Protocol

A design of a quantum link layer protocol that offers the above service is the *quantum entanglement generation protocol* (QEGP) proposed by Dahlberg *et al.* [19]. As originally designed, this protocol relies on the underlying quantum physical layer protocol to achieve accurate timing synchronization with its remote peer and to detect inconsistencies between the local state and the state of the remote counterpart. To satisfy such requirements, QEGP is accompanied by a quantum physical layer protocol, called *midpoint heralding protocol* (MHP), designed to support QEGP on heralded entanglement-based quantum links.

Entanglement requests and agreement. QEGP exposes an interface for its user to submit *entanglement requests*. An entanglement request can specify all the aforementioned configuration parameters (remote node ID, number of entangled pairs, minimum fidelity, request type, measurement basis), and an additional set of parameters which can be used to determine the priority of the request. In the theoretical protocol proposed in [19], agreement on the requests between the nodes was achieved using a distributed queue protocol (DQP) which added the incoming requests to a joint queue. The distributed queue, managed by the node designated as primary, ensures that both nodes schedule pending entanglement requests in the same order. Moreover, QEGP attaches a timestamp to each request in the distributed queue, so that both nodes can process the same entanglement request simultaneously.

Time synchronization. Time-scheduling entanglement generation requests is necessary for the two neighboring nodes to trigger entanglement generation at the same time, and avoid wasting entanglement attempts. QEGP relies on MHP to maintain and distribute a synchronized clock, which QEGP itself uses to schedule entanglement requests. The granularity of such a clock is only marginally important, but its consistency across the two neighboring nodes is paramount to make sure that entanglement attempts are triggered simultaneously on the two ends.

Mismatch verification. One of the main responsibilities of MHP is to verify that both nodes involved in entanglement generation are servicing the same QEGP request at the same time, which the protocol achieves by sending an auxiliary classical message to the heralding station when the physical device sends the flying qubit. The heralding station can thus verify that the messages fetched by the two MHP peers are consistent and correspond to the same QEGP request.

QEGP challenges. We identify three main challenges that would be faced when deploying QEGP on a large-scale quantum network, while suggesting an alternative solution for each of these. (C1) Using a link-local protocol (DQP) to schedule entanglement requests, albeit sufficient for a single-link network, becomes challenging in larger networks, given that a node might be connected to more than just one peer. In such scenarios, the scheduling of entanglement requests can instead be deferred to a centralized scheduling

entity, one which has more comprehensive knowledge of the entire (sub)network [32]. (C2) Entrusting the triggering of entanglement attempts to QEGP would impose very stringent real-time constraints on the system where QEGP itself is deployed—even microsecond-level latencies on either side of the link can result in out-of-sync (thus wasteful) entanglement attempts. While Dahlberg *et al.* [19] identify this problem as well, the original MHP protocol assumes that both QEGP peers issue an entanglement command to the physical layer at the same clock cycle. In this scheme, MHP initiates an entanglement attempt regardless of the state of the remote counterpart. We believe that fine-grained entanglement attempt synchronization should pertain to the physical layer only, building on the assumption that the real-time controllers deployed at the physical layer of each node are anyway highly synchronized [15]. (C3) Checking for request mismatches at the heralding station requires the latter to be capable of performing such checks in real-time. Given that the two neighboring MHP protocols have to anyway synchronize before attempting entanglement, we suggest that, as an alternative approach, consistency checks be performed at the nodes themselves, rather than at the heralding station, just before entering the entanglement attempt routine.

5.1.3 Revised Protocol

To address the present QEGP and MHP challenges with the proposed solutions, we have made some modifications to the original design of the two protocols. In particular, we adopted a centralized request scheduling mechanism [32] to tackle challenge (C1), we delegated the ultimate triggering of entanglement attempts to MHP as a solution to challenge (C2), and we assigned request mismatch verification to the MHP protocol running on each node, rather than to the heralding station, to address challenge (C3).

    Centralized request scheduling. To avoid using a link-local protocol (DQP) to schedule entanglement requests, our version of QEGP defers request scheduling to a *centralized request scheduler*, whereby a node's entanglement generation schedule is computed on the basis of the whole network's needs. Delegating network scheduling jobs to centralized entities is, albeit not the only alternative, a common paradigm of classical networks, and especially of software-defined networking (SDN)—a concept that has been recently investigated in the context of quantum networking [22, 23]. In our system, the centralized scheduler produces a time-division multiple access (TDMA) network schedule—one for each node in the network—where each time bin is reserved for a certain class of entanglement generation requests [32]. A class of requests may comprise, for instance, all requests coming from the same application and asking for the same fidelity of the entangled states. While reserving time bins may be redundant in a single-link network, integrating a centralized scheduling mechanism early on into the link layer protocol will facilitate future developments.

    MHP synchronization and timeout. Although centralized request scheduling makes the synchronization of QEGP peers easier, precise triggering of entanglement attempts should still be entrusted to the component of the system where time is the most deterministic—in our case, the physical layer protocol MHP. In contrast to Ref. [19], once MHP fetches an entanglement instruction from QEGP, the protocol announces itself as ready to its remote peer, and waits for the latter to do so as well. After this synchronization step succeeds, the two MHP peers can instruct the underlying hardware to trigger an entanglement at-

tempt at a precise point in time. If, instead, one of the two MHP peers does not receive announcements from its remote counterpart within a set timeout, it can conclude that the latter is not ready, or temporarily not responsive, and can thus return control to QEGP without wasting entanglement attempts. This MHP synchronization step is also useful for the two sides to verify that they are processing the same QEGP request, and thus catch mismatches.

The MHP synchronization routine inherently incurs some overhead, which is also larger on longer links. We mitigate this overhead by batching entanglement attempts— that is, the physical layer attempts entanglement multiple times after synchronization before reporting back to the link layer. The maximum number of attempts per batch is a purely physical-layer parameter, and it has no relation with the link layer entanglement request timeout parameter described in Ref. [19]—although batches should be small enough for the link layer timeout to make sense.

5.1.4 Implementation

The original design of the QEGP and MHP protocols, as well as our revision, specifies the conceptual interaction between the two protocols and the service exposed to a higher layer in the system, but does not impose particular constraints on how to implement link layer and physical layer, how to realize the physical interface between them, and how to configure things such as the centralized request scheduler and the entanglement attempt procedure. Fig. 5.2 gives an overview of the architecture of our quantum network nodes. We briefly describe our most relevant implementation choices here and in Section 5.2.

Application processing. At the application layer, user programs—written in Python using a dedicated *software development kit* [33]—are processed by a rudimentary compilation stage, which translates abstract quantum networking applications into gates and operations supported by our specific quantum physical platform. Such gates and operations are expressed in a low-level assembly-like language for quantum networking applications called *NetQASM* [34]. As part of our software stack, we also include an *instruction processor*, conceptually placed above the link layer, which is in charge of dispatching entanglement requests to QEGP and other application instructions to the physical layer directly.

Interface. Ref. [19] did not provide a specification of the interface to be exposed by the physical layer. We designed this interface such that the physical layer can accept *commands* from the higher layer, specifically: (a) qubit initialization (INI), (b) qubit measurement (MSR), (c) single-qubit gate (SQG), (d) entanglement attempt (ENT, or ENM for M- or R-type requests), (e) pre-measurement gates selection (PMG, to specify in which basis to measure the qubit for M- or R-type requests). For each command, the physical layer reports back an *outcome*, which indicates whether the command was executed correctly, and can bear the result of a qubit measurement and the Bell state produced after a successful entanglement attempt. Our software stack also comprises a *hardware abstraction layer* (HAL) that sits below QEGP and the instruction processor. The HAL encodes and serializes commands and outcomes, and is thus used to interface with the device controller.

TDMA network schedule. Designing a full-blown centralized request scheduler is a challenge in and of its own, outside the scope of this work. Instead of implementing such a scheduler, we compute static TDMA network schedules [32] and install them manually on the two network nodes upon initialization. TDMA schedules for our simple single-link
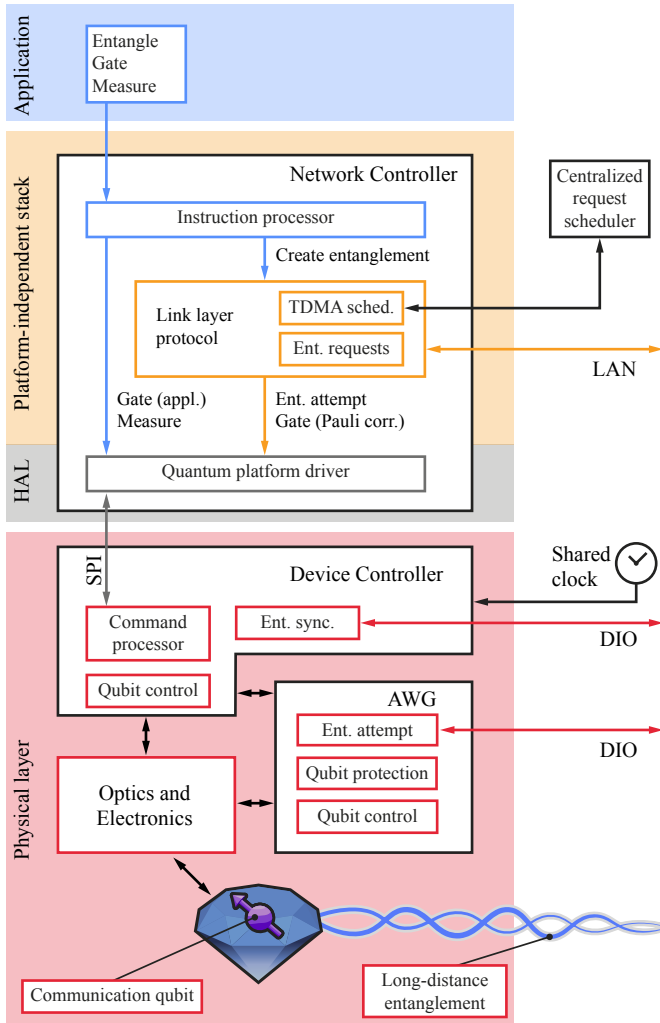
Figure 5.2: Quantum network node architecture. From top to bottom: At the application layer, a simple platform-independent routine is sent to the network controller. The network controller implements the platform-independent stack—in this work only the link layer protocol—and a hardware abstraction layer (HAL) to interface with the physical layer's device controller. An instruction processor dispatches instructions either directly to the physical layer, or to the link layer protocol in case a remote entangled state is requested by the application. The link layer schedules entanglement requests and synchronizes with the remote node (on a local area network, LAN) using a time-division multiple access (TDMA) schedule computed by a centralized scheduler (external). At the physical layer, the device controller fetches commands from—and replies with outcomes to—the network controller. Driven by a clock shared with the neighboring node, it performs hard-real-time synchronization for entanglement generation using a digital input/output (DIO) interface. By controlling the optical and electronic components (among which an arbitrary waveform generator, AWG), the device controller can perform universal quantum control of the communication qubit in real-time, as well as attempt long-distance entanglement generation with the neighboring node.

experiments are quite trivial (see section 5.5.1), as the network resources of a node are not contended by multiple links.

Entanglement attempts. Producing entanglement on a link can take several attempts. To minimize the number of ENT commands fetched by MHP from QEGP, as well as to mitigate the MHP synchronization overhead incurred after each entanglement command, we batch entanglement attempts at the MHP layer, such that synchronization and outcome reporting only happens once per batch of attempts.

Delivered entangled states. In our first iteration, we implemented QEGP such that it always delivers $|\Phi^+\rangle$ Bell states to the higher layer. This means that, when the physical layer produces a different Bell state, QEGP (on the node where the entanglement request originates) issues a single-qubit gate—a Pauli correction—to transform the entangled pair into the $|\Phi^+\rangle$ state[1]. A future version of QEGP could allow the user to request any Bell state, and could extract the Pauli correction from QEGP so that the application itself can decide, depending on the use case, whether to apply the correction or not.

Mismatch verification. As per our design specification, MHP should also be responsible for verifying that the entanglement commands coming from the two QEGP peers belong to the same request. We did not implement this feature yet because, in our simple quantum network, we do not expect losses on the classical channel used by the two MHP parties to communicate—a lossy classical channel would be the primary source of inconsistencies at the MHP layer [19]. However, we believe that this verification step will prove very useful in real-world networks where classical channels do not behave as predictably.

Deployment. We implemented QEGP as a software module in a system that also includes the instruction processor and the hardware abstraction layer. QEGP, the instruction processor and the hardware abstraction layer, forming the *network controller*, are implemented as a C/C++ standalone runtime developed on top of FreeRTOS, a real-time operating system for embedded platforms [35]. The runtime and the underlying operating system are deployed on a dedicated Avnet MicroZed—an off-the-shelf platform based on the Zynq-7000 SoC, which hosts two ARM Cortex-A9 processing cores, of which only one is used, clocked at 667 MHz. QEGP connects to its remote peer via TCP over a Gigabit Ethernet interface. The interface to the physical layer is realized through a 12.5 MHz SPI connection. The user application is sent from a general-purpose 4-core desktop machine running Linux, which connects to the instruction processor through the same Gigabit Ethernet interface that QEGP uses to communicate with its peer.

5.2 Physical Layer Control in Real-Time

In this section, we outline the design and operation of the physical layer, which executes the commands issued by the higher layers on the quantum hardware and handles time-critical synchronization between the quantum network nodes. The physical layer of a quantum network, as opposed to the apparatus of a physics experiment, needs to be able to execute commands coming from the layer above in real-time. Additionally, when performing the requested operations, it needs to leave the quantum device in a state that is compatible with future commands (for example, as discussed below, it should protect

---

[1]We abbreviate the four two-qubit maximally entangled Bell states as $|\Phi^\pm\rangle = (|00\rangle \pm |11\rangle)/\sqrt{2}$ and $|\Psi^\pm\rangle = (|01\rangle \pm |10\rangle)/\sqrt{2}$.

qubits from decoherence while it awaits further instructions). Finally, if a request cannot be met (e.g. the local quantum hardware is not ready, the remote quantum hardware is not available, etc.), the physical layer should notify the link layer of the issue without interrupting its service.

Our quantum network is composed of two independent nodes based on diamond NV centers physically separated by ≈2 m (see fig. 5.2 for the architecture of one node, and fig. 5.6 for details on the connections between the two nodes). We will refer to the two nodes as *client* and *server*, noting that this is only a logical separation useful to describe the case studies—the two nodes have the exact same capabilities. On each node, we implement the logic of the physical layer in a state-machine-based algorithm deployed on a time-deterministic microcontroller, the *device controller* (Jäger ADwin Pro II, based on Zynq-7000 SoC, dual-core ARM Cortex-A9, clocked at 1 GHz). Additionally, each node uses an arbitrary waveform generator (AWG, Zurich Instruments HDAWG8, 2.4 GSa/s, 300 MHz sequencer) for nanosecond-resolution tasks, such as fast optical and electrical pulses; the use of such a user-programmable FPGA-based AWG, as opposed to a more traditional upload-and-play instrument (such as the ones used in Ref. [15]), enables the real-time control of our quantum device.

5.2.1 Single node operation

On our quantum platform, before a node is available to execute commands, it needs to perform a qubit readiness procedure called *charge and resonance check* (CR check). This ensures that the qubit system is in the correct charge state and that the necessary lasers are resonant with their respective optical transitions. Other quantum platforms might have a similar preparation step, such as loading and cooling for atoms and ions [9, 10]. Once the CR check is successful, the device controller can fetch a command from the network controller. Depending on the nature of the command, the device controller might need to coordinate with other equipment in the node or synchronize with the device controller of the other node.

For qubit initialization and measurement commands (INI and MSR), the device controller shines the appropriate laser for a pre-defined duration (INI≈100 μs, MSR≈10 μs). Both operations are deterministic and carried out entirely by the device controller.

Single qubit gates (SQG) require the coordination of the device controller and the AWG. For our communication qubits, they consist of generating an electrical pulse with the AWG (duration ≈100 ns), which is then multiplied to the qubit frequency (≈2 GHz), amplified and finally delivered to the quantum device. The link layer can request rotations in steps of $\pi/16$ around the X, Y or Z axis of the Bloch sphere (here we implement only X and Y rotations, Z rotations will be implemented in the near future, see section 5.5.2). When a new gate is requested by the link layer, the device controller at the physical layer informs the AWG of the gate request via a parallel 32-bit DIO interface. The AWG will then select one of the 64 pre-compiled waveforms, play it, and notify the device controller that the gate has been executed. The device controller will in turn notify the network controller of the successful operation.

After the rotation has been performed, our qubit—if left idling—would lose coherence in ≈5 μs. A coherence time exceeding 1 s has been reported on our platform [36] using decoupling sequences (periodic rotations of the qubit that shield it from environmental

noise). By interleaving decoupling sequences and gates, one can perform extended quantum computations [37]. These long sequences of pulses have in the past been calculated and optimized offline (on a PC), then uploaded to an AWG, and finally executed on the quantum devices with minimal interaction capabilities (mostly binary branching trees, see [15]). In our case, it is impossible to pre-calculate these sequences, since we cannot know in advance which gates are going to be requested by the link layer. To solve this challenge, we implement a *qubit protection* module on the AWG, that interleaves decoupling sequences with the requested gates in real-time. As soon as the first gate in a sequence is requested, the AWG starts a decoupling sequence on the qubit. Then, it periodically checks if a new gate has been requested, and if so, it plays it at the right time in the decoupling sequence. The AWG will continue the qubit protection routine until the device controller will ask for it to stop (e.g. to perform a measurement). This technique allows us to execute universal qubit control without prior knowledge of the sequence to be played, and—crucially—in real-time.

5.2.2 Entanglement generation

Differently from the commands previously discussed, attempting entanglement generation (ENT) requires tight timing synchronization between the device controllers—and AWGs—of the two nodes. In our implementation, the two device controllers share a common 1 MHz clock as well as a DIO connection to exchange synchronization messages (see Ref. [15]). When the device controllers are booted, they synchronize an internal cycle counter that is used for time-keeping, and is shared, at each node, with their respective network controllers to provide timing information to the link layer and the higher layers. Over larger distances, one could use well-established protocols to achieve sub-nanosecond, synchronized, GPS-disciplined common clocks [38].

When a device controller fetches an ENT command, it starts a three-way handshake procedure with the device controller of the other node. If the other node has also fetched an ENT command, they will synchronize and proceed with the entanglement generation procedure. If one of the two nodes is not available (e.g. it is still trying to pass the CR check) the other node will time out, after 0.5 ms, and return an *entanglement synchronization failure* (ENT_SYNC_FAIL) to its link layer. The duration of the timeout is chosen such that is comparable with the average time taken by a node to pass the charge and resonance check (if correctly on resonance). This is to avoid unnecessary interactions between physical layer and link layer. After the entanglement synchronization step, the device controllers proceed with an optical phase stabilization cycle [15], and then the AWGs are triggered to attempt entanglement generation. In our implementation, one device controller (the server's) triggers both AWGs to achieve sub-nanosecond jitter between the two AWGs (see section 5.5.3 for a discussion on longer distance implementation). Each entanglement attempt lasts 3.8 µs, and includes fast qubit initialization, communication-qubit to flying-qubit entanglement, and probabilistic entanglement swapping of the flying qubits [15]. The AWGs attempt entanglement up to 1000 times before timing out and reporting an *entanglement failure* (ENT_FAIL). Longer batches of entanglement attempts would increase the probability that one of the nodes goes into the unwanted charge state (and therefore cannot produce entanglement, see section 5.5.7). While in principle possible, we did not implement, in this first realization, the charge stabilization mechanism proposed in

Ref. [14] that would allow for significantly longer batches of entanglement attempts.
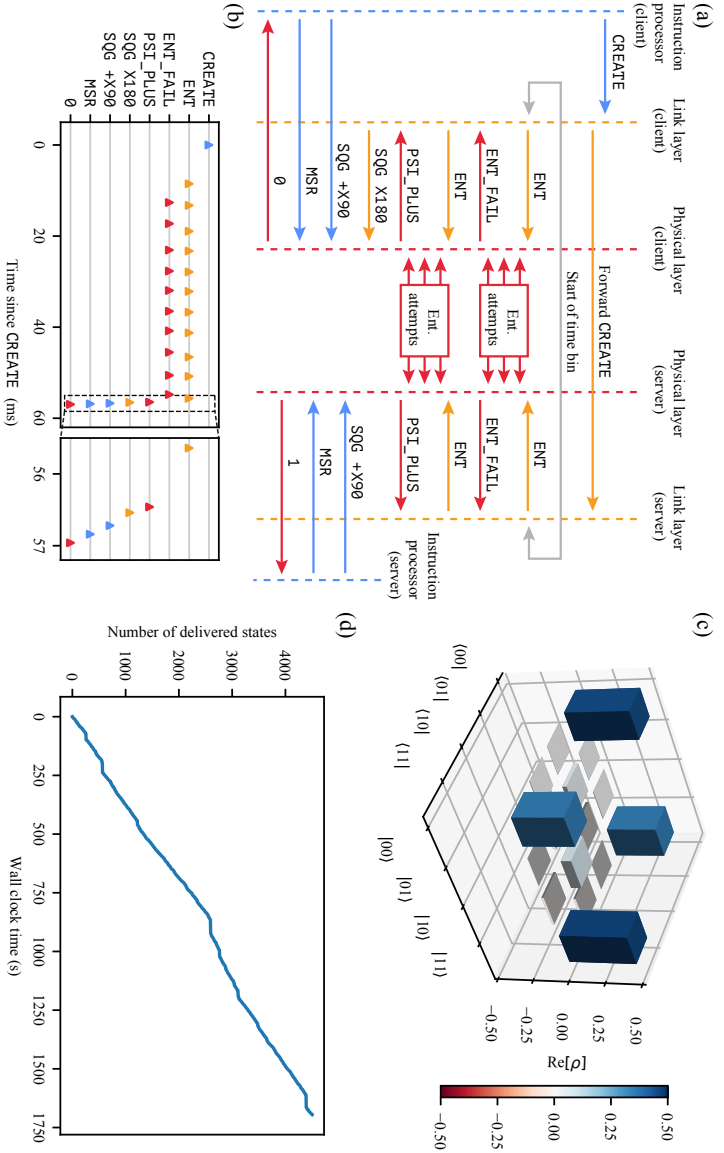
If an entanglement generation attempt is successful (probability $\approx 5 \times 10^{-5}$), the communication qubits of the two nodes will be projected into an entangled state (either $|\Psi^+\rangle$ or $|\Psi^-\rangle$, depending on which detector clicked at the heralding station). To herald success of the entanglement attempt, a CPLD (Complex Programmable Logic Device, Altera MAX V 5M570ZF256C5N) sends a fast digital signal to both AWGs and device controllers, to prevent a new entanglement attempt from being played (which would destroy the generated entangled state). When the heralding signal is detected, the AWGs enter the qubit protection routine and wait for further instructions from the device controllers, which in turn notify the link layer of the successful entanglement generation, as well as which state was generated.

To satisfy *M*- or *R*-type entanglement requests, the link layer can instruct the physical layer to apply an immediate measurement to the entangled qubit by means of an ENM command. Up until heralding of the entangled state, the physical layer operates as it does for the ENT command. When the state is ready, it proceeds immediately with a sequence of single qubit gates (as prescribed by an earlier PMG command) and a qubit measurement. The result of the measurement, together with which entangled state was generated, is communicated to the link layer. It is worth noting that the two nodes could fetch different types of requests and still generate entanglement. In fact, this will be used later in the remote state preparation application.

5.3 Evaluation

To demonstrate and benchmark the capabilities of the link layer protocol, the physical layer, and of our system as a whole, we execute—on our two-node network—three quantum networking applications, all having a similar structure: the client asks for an entangled pair with the server, which QEGP delivers in the $|\Phi^+\rangle$ Bell state, and then both client and server measure their end of the pair in a certain basis. First, we perform full quantum state tomography of the delivered entangled states. Second, we request and characterize entangled states of varying fidelity. Third, we execute remote preparation of qubit states on the server by the client. For all three applications, we study the quality of the entangled pairs delivered by our system. Additionally, we use the second application to assess the latency incurred by our link layer, and to compare it to the overall entanglement generation latency, including that of the physical layer. Crucially, the three applications are executed back-to-back on the quantum network, without any software or hardware changes to the system—the only difference being the quantum-platform-independent application sent to the instruction processor (see 5.5.4).

The sequence diagram in Fig. 5.3a exemplifies the general flow between system components during the execution of an application. At first, the instruction processor issues a request to create entanglement to link layer (CREATE). Then, the client's link layer forwards the request with the server's counterpart (Forward CREATE). The request is processed as soon as the designated time bin in the TDMA schedule starts, at which point the first entanglement command (ENT) is fetched by physical layer. After an entangled state is produced successfully (PSI_PLUS), the link layer of the client issues, if needed, a Pauli correction (π rotation around the X axis, SQG X180) to deliver the pair in the $|\Phi^+\rangle$ state. Finally, the instruction processor issues a gate ($\pi/2$ rotation around the X axis, SQG X90)

Figure 5.3: Full state tomography with the quantum network stack. (a) Sequence diagram of the communication steps across the network stack and the two nodes to perform one repetition of the tomography application (in particular, measurement of the $\langle YY \rangle$ correlator). The coloring follows that of Fig. 5.1. CREATE: entanglement request, ENT: entanglement attempts request, ENT_FAIL: failed the batch of entanglement attempts, PSI_PLUS: successful entanglement attempt with generated state Ψ^+, SQG: single-qubit gate, X180: 180° rotation around X axis, MSR: qubit measurement, 0/1: qubit measurement outcome. See Table 5.1 for a complete list of commands. Note that the client's link layer protocol requests a X180 gate after entanglement generation to deliver the $|\Phi^+\rangle$ Bell state to the higher layer. **(b)** Example time trace of (a) for the client. Several batches of entanglement attempts are required before an entangled state is heralded. On the right, a zoomed-in part of the trace (corresponding to the dashed box in the left plot). **(c)** Reconstructed density matrix of the states delivered by the link layer. Only the real part is plotted (imaginary elements are all ≈0, see main text). We estimate a fidelity F with $|\Phi^+\rangle$ of $F = 0.793(9)$. **(d)** Total number of delivered states over time. The occasional pauses in entanglement delivery (plateaus) are due to the client's NV center becoming off-resonant with the relevant lasers (see section 5.5.7). Differences in slope are due to changes in resonance conditions that increase the time necessary to pass the charge and resonance check.

5

and a measurement (MSR) to read out the entangled qubit in a certain basis, and receives an outcome from the physical layer (0). Fig. 5.3b illustrates the actual latencies between these interactions in one iteration of the full state tomography application.

For all our experiments, we configured TDMA time bins to be of 20 ms. In a larger network, the duration of time bins should be calibrated according to the average time it takes, on a certain link, to produce an entangled pair of a certain fidelity [32]. By doing so, one can maximize network usage and thus reduce qubit decoherence on longer end-to-end paths. However, in our single-link network, the duration of time bins only influences the frequency at which new entanglement requests are processed. Our time bin duration accommodates up to four batches of 1000 entanglement attempts.

5.3.1 Full Quantum State Tomography

The first application consists in generating entangled states at the highest *minimum fidelity* currently available on our physical setup (0.80), and measuring the two entangled qubits in varying bases to learn their joint quantum state. We measure all 9 two-node correlators ($\langle XX \rangle, \langle XY \rangle, ..., \langle ZZ \rangle$) as well as all their ± variations ($\langle +X + X \rangle, \langle +X - X \rangle$, etc.) to minimize the bias due to measurement errors. For each of the $9 \times 4 = 36$ combinations, we measure 125 data points, for a total of 4500 entangled states generated and measured. Sequence 5.1 in Section 5.5.4 contains a pseudocode description of the application.

The collected measurement outcomes are then analyzed using QInfer [39], in particular the Monte Carlo method described in Ref. [40] for Bayesian estimation of density matrices from tomographic measurements. The reconstructed density matrix is displayed in Fig. 5.3c (only the real part is shown in the figure) and its values and uncertainties are

$$\mathrm{Re}[\rho] = \begin{pmatrix} 0.443(4) & 0.001(6) & 0.003(4) & 0.343(6) \\ 0.001(6) & 0.029(7) & -0.019(3) & -0.001(4) \\ 0.003(4) & -0.019(3) & 0.071(5) & -0.015(5) \\ 0.343(6) & -0.001(4) & -0.015(5) & 0.457(8) \end{pmatrix},$$

$$\mathrm{Im}[\rho] = \begin{pmatrix} 0 & -0.003(4) & -0.015(4) & 0.021(5) \\ 0.003(4) & 0 & 0.022(9) & -0.001(6) \\ 0.015(4) & -0.022(9) & 0 & -0.001(4) \\ -0.021(5) & 0.001(6) & 0.001(4) & 0 \end{pmatrix}.$$

Here $\rho_{ij,mn} = \langle ij| \, \rho \, |mn \rangle$, with i, m (j, n) being the client (server) qubit states in the computational basis. The uncertainty on each element of the density matrix is calculated as the standard deviation of that element over the probability distribution approximated by the Monte Carlo reconstruction algorithm (probability distribution approximated by 1×10^5 Monte Carlo particles [40]). It is then possible to estimate the fidelity of the delivered entangled states with respect to the maximally entangled Bell state, which we find to be $F = 0.793(9)$. The measured fidelity is slightly lower ($\approx 2\,\%$) than what measured in Ref. [15] without the use of the QEGP abstraction (and the whole network controller where QEGP runs). This discrepancy could be due to the additional physical-layer decoupling sequences required for real-time operation ($\approx 300\,\mu s$) and the additional single-qubit gate issued by the link layer to always deliver $|\Phi^+\rangle$ (see 5.5.5).

It is to be noted that, in order to obtain the most faithful estimate of the generated state (see section 5.5.6 for details), the measured expectation values are corrected, in post-

processing, to remove known tomography errors of both client and server [41], and events in which at least one physical device was in the incorrect charge state.

Finally, we show, in Fig. 5.3d, that our system can sustain a fairly stable entanglement delivery rate over ≈30 min of data acquisition—plateaus and changes in slope can be attributed to varying conditions of resonance between the NV centers and the relevant lasers (see section 5.5.7).

5.3.2 Latency vs Fidelity

The QEGP interface allows its user to request entangled pairs at various minimum fidelities. For physical reasons, higher fidelities will result in lower entanglement generation rates [14, 17]. The trade-off between fidelity and throughput is particularly interesting in a scenario where some applications might require high-fidelity entangled pairs and are willing to wait a longer time, while others might prefer lower-fidelity states but higher rates [19]. Clearly, for the link layer to offer a range of fidelities to choose from, the underlying physical layer must support such a range. We benchmark the capabilities of the link layer and of the physical layer to deliver states at various fidelities in a single application by measuring the $\langle XX \rangle$, $\langle YY \rangle$ and $\langle ZZ \rangle$ correlators (and their ± variations, as we did above, for a total of $3 \times 4 = 12$ correlators) for seven different target fidelities, (0.50, 0.55, 0.60, 0.65, 0.70, 0.75, 0.80). We generate 1500 entangled states per fidelity, for a total of 10 500 delivered states (see Sequence 5.2 in section 5.5.4). With this case study, we analyze both the resulting fidelity and the system's latency for different requested fidelities.

The results for measured fidelity versus requested fidelity are shown in Fig. 5.4a. It is worth noting that the application iterates over the range of fidelities in real-time, and thus the physical layer is prepared to deliver any of them at any point. We calibrate the physical layer to deliver states of slightly higher fidelity than the requested ones (0.03 more), since entanglement requests specify the *minimum* desired fidelity. The measured fidelities are—within measurement uncertainty—always matching or exceeding the requested minimum ones (the dashed gray line in fig. 5.4a is the $y = x$ diagonal). As in the previous application, measurement outcomes are post-processed to eliminate tomography errors and events in which the physical devices were in the incorrect charge state (we refer to the latter as *charge state correction*). For arbitrary applications that use the delivered entangled states for something other than statistical measurements, applying the second correction directly at the link layer might prove challenging, since the information concerning whether to discard an entangled pair is only available at the physical layer *after* the entangled state is delivered to the link layer (when the next CR check is performed). However, a mechanism to identify *bad* entangled pairs retroactively at the link layer—like the expiry functionality included in the original design of QEGP [19]—could be used to discard entangled states after they have been delivered by the physical layer. For completeness, we also report, again in Fig. 5.4a, the measured fidelity when the wrong charge state correction is not applied.

For each requested fidelity we also measure the entanglement generation latency [19], defined as the time between the issuing of the CREATE request to the link layer, until the successful entanglement outcome reported by the physical layer (refer to fig. 5.3a for a diagram of the events in between these two). Fig. 5.4b shows the measured average latency, grouped by requested fidelity and broken down into the various sources of latency. When
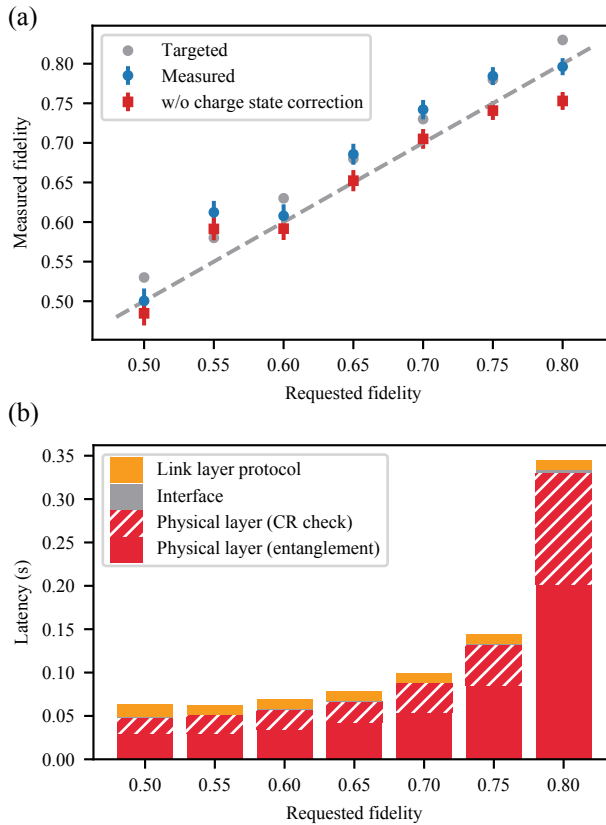
Figure 5.4: Performance of the entanglement delivery service. **(a)** Measured fidelity of the states delivered by the link layer for varying requested fidelity. Targeted fidelity at the physical layer is 0.03 higher than the link layer protocol's *minimum fidelity* request. When not correcting for wrong charge state events, fidelity is reduced by a few percents (see section 5.5.6). Error bars represent 1 s.d. **(b)** Average latency of the entanglement delivery per requested fidelity, broken down into sources of latency. Entanglement generation and charge and resonance check at the physical layer are the largest sources of latency (at higher fidelities, more entanglement attempts are required before success). Running the link layer protocol introduces a small but measurable overhead ($\approx 10\,\mathrm{ms}$) to the entanglement generation procedure, which does not depend on the requested fidelity, and that could be mitigated by requesting multiple entangled states in a single instruction. The communication delays between quantum network controller and quantum device controller (Interface) introduce negligible overall latency.

calculating the average latencies, we have ignored entanglement requests that required more than 10 s to be fulfilled. These high-latency requests correspond to the horizontal plateaus of Fig. 5.3d (see section 5.5.7 for details). The main contribution to the total latency comes from the entanglement generation process at the physical layer, followed by the NV center preparation time (CR check). Both latency values are consistent with the expected number of entanglement attempts required by the single-photon entanglement protocol employed at the physical layer [14]. The link layer protocol adds, on average, ≈10 ms of extra latency to all requests, regardless of their fidelity. This is due partly to the synchronization of the CREATE request between the two nodes (i.e. a simple TCP message), but mostly to the nodes having to wait for the next time bin in the network schedule to start (the larger the time bins, the larger the worst-case waiting time, see section 5.5.1). We remark that, by requesting multiple entangled states in a single CREATE, one can distribute this overhead over many generated pairs, to the point where it becomes negligible. While our applications did not issue multi-pair CREATE requests, this would be the more natural choice for real applications, and would result in better utilization of the allocated time bins. Finally, the overhead incurred by the interface between microcontrollers is rather small (barely visible in fig. 5.4b), but could however be further reduced by integrating device controller and network controller into a single device. It is worth mentioning that, in our simple scenario in which each entanglement request is only submitted to QEGP after the previous one completes, and thus the request queue never grows larger than one element, throughput happens to be almost exactly the same as the inverse of latency, and hence it is not reported here.

Overall, we observe that the extra entanglement generation latency incurred when deploying an abstraction layer (QEGP) on top of the physical layer, while not too modest, is only a small part of the whole, particularly at higher fidelities. Nevertheless, optimizing the length of TDMA time bins could result in an even smaller overhead (see section 5.5.1).

5.3.3 Remote State Preparation

One of the use cases of the QEGP service is to prepare quantum states on a remote quantum server [19]. Remote state preparation is a fundamental step to execute a blind quantum computation application [5], whereby a client quantum computer with limited resources can run quantum applications on a powerful remote quantum server using the many qubits the server has, while keeping the performed computation private.

Remote state preparation is different from the previous two cases in that the client can measure its end of the entangled pair as soon as the pair is generated, while the server has to keep its qubit alive waiting for further instructions. For such a scenario, the client can make use of QEGP's service to issue R-type entanglement requests, so that the local end of the entangled pair can be measured (in a certain basis) as soon as it is generated, while the server's qubit can be protected for later usage. An R-type entanglement request results in an ENM command on the client and an ENT command on the server. For this type of requests (as well as for M-type ones), since the local end of the pair is measured immediately, the client's QEGP can skip the Pauli correction used to always deliver $|\Phi^+\rangle$, and can instead apply a classical correction to the received measurement outcome (refer to section 5.5.5).

To showcase this feature of QEGP we use the client node to prepare the six cardinal states on the server ($|\pm x\rangle$, $|\pm y\rangle$, $|0\rangle$ and $|1\rangle$) by having the client measure its share of the
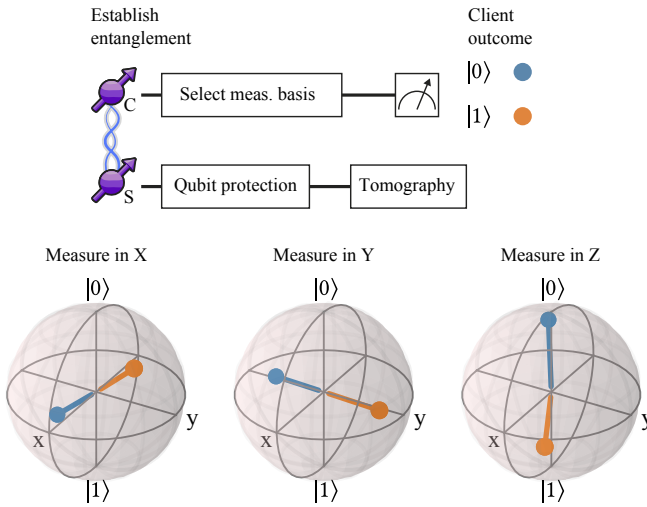
Figure 5.5: Tomography of states prepared on the server by the remote client. For each chosen measurement axis of the client (X, Y, Z), and for each obtained measurement outcome at the client ($|0\rangle$, $|1\rangle$), a different state is prepared on the server. Plotted on the Bloch spheres are the results of the tomography on the server's qubit. Uncertainties on each coordinate are ≈ 0.05 (see section 5.5.6). We find an average preparation fidelity of $F = 0.853(8)$.

5

entangled state in the six cardinal bases. We then let the server measure the prepared states—again in the six cardinal bases—to perform tomography. For each client measurement basis, and for each server tomography basis, we deliver 125 entangled states at a requested fidelity of 0.80, for a total of $6 \times 6 \times 125 = 4500$ remote state preparations (see Sequence 5.3 in section 5.5.4). The results are presented in Fig. 5.5, which displays the tomography of the prepared states on the server, for the three different measurement axes of the client and the two possible measurement outcomes of the client. The prepared states are affected by the measurement error of the client ($F_0 = 0.928(3)$, $F_1 = 0.996(1)$): an error in the measurement of the client's qubit results in an incorrect identification of the state prepared on the server. By alternating between positive and negative readout orientations, we make sure that the errors affect all prepared states equally, instead of biasing the result. We note that we exclude, once again, events in which at least one of the two devices was in the wrong charge state, and we correct for the known tomography error on the server (results without corrections are in section 5.5.6). Overall, we find an average remote state preparation fidelity of $F = 0.853(8)$. The asymmetry in the fidelity of the $|0\rangle$ and $|1\rangle$ states is caused by the asymmetry in the populations $\langle 01| \rho |01\rangle$ vs $\langle 10| \rho |10\rangle$ of the delivered entangled state, which in turn is due to the double $|0\rangle$ occupancy error of the single-photon protocol used to generate entanglement [14, 15].

5.4 Conclusions

In summary, we have demonstrated the operation of a link layer and a physical layer for entanglement-based quantum networks. The link layer abstracts the entanglement

generation procedure provided by the physical layer—implemented here with two NV center-based quantum network nodes—into a robust platform-independent service that can be used to run quantum networking applications. We performed full quantum state tomography of the states delivered by the link layer, tested its ability to deliver states at different fidelities in real-time, and verified remote state preparation of a qubit from the client on the server, a fundamental step towards blind quantum computation [5]. We have shown that our implementation of link and physical layers can deliver entangled states at the fidelity requested by the user, despite some marginal inefficiencies—some of which can be addressed in a future version of the protocols (e.g. avoiding Pauli corrections unless necessary). We have also quantified the additional latency incurred by deploying the link layer protocol on top of the physical layer. Although not detrimental, the extra overhead is still noticeable, but can also be scaled down by optimizing the scheduling of entanglement generation requests. We also acknowledge that scheduling a quantum node's resources is still an open problem [32, 42, 43], and that the simple TDMA approach taken here might be a suboptimal choice in more advanced quantum networks.

The adoption of the techniques presented here (which are not specific to our diamond devices) by other quantum network platforms [9, 10, 17, 44–47] will boost the development towards large-scale and heterogeneous quantum networks. Real-time control of memory qubits, as well as the availability of multi-node networks and dynamic network schedules, will enable demonstrations of the higher layers of the network stack [48], which in turn will open the door to end-to-end connectivity on a platform-independent quantum network.

5.5 Supplementary information

5.5.1 Pre-computed TDMA schedule

As mentioned in Section 5.1, TDMA network schedules are redundant in a one-link network. Time-binning network activity also forces nodes to only process entanglement requests at the beginning of a time division, thus introducing latency and idle time. Particularly, longer time bins potentially result in entanglement requests to wait longer to be processed. However, an application asking for multiple entangled pairs with just one request would experience smaller average latencies, as all pairs—but the first one—would be generated in close succession.

In our experiments, TDMA schedules are just a constant division of 20 ms time bins, each of which is reserved to the only application running. We chose the duration of the time-bin—somewhat arbitrarily, given the small effect on our experiments—to be equal to 1000 communication cycles between the device controller and the network controller (20 ms = 1000 × 20 µs).

5.5.2 Single qubit gates implemented

At the physical layer, we implement real-time rotations around the X and Y axes of the qubit Bloch sphere, using a resolution of $\pi/16$ =11.25°. That is, the link layer can request any rotation that is a multiple of $\pi/16$ around either the X or Y axis. The different rotations are performed using Hermite-shaped pulses (as described in Ref. [15]) of calibrated amplitude. The choice of X(Y) rotation axis is implemented using the I(Q) channel of the microwave vector source.

While supported at the link layer, our physical layer currently does not implement Z axis rotations. Such rotations around the Z axis could be implemented by virtual rotations of the Bloch sphere: a π pulse around the Z axis is equivalent to multiplying future I and Q voltages by –1. By keeping track of the accumulated Z rotations, and by adjusting I and Q mixing accordingly, one can perform effective Z rotations with very high resolution and virtually no infidelity. The AWGs currently in use have the required capabilities, and the implementation of said Z gates is planned for the near future.

5.5.3 Clock sharing and AWG triggering over longer distances

One of the technical challenges of realizing a large scale quantum network is synchronizing equipment at the physical layer across nodes. The synchronization is required to generate entanglement—the photons from the two nodes need to arrive at the same time at the heralding station (compared to their duration, 12 ns for NV centers in bulk diamond samples); failing to do so would reduce (or even remove) their indistinguishability, which is required to establish long-distance entanglement [15]. Our two nodes are located in a single laboratory, on the same optical table, approximately 2 m apart. This allows for some simplifications, for the purpose of demonstrating entanglement delivery using a network stack, which would not be possible over longer distances. Specifically:

1. We use a single laser—the client's—to excite both nodes, as in Ref. [15]. Over longer distances, one would need to phase-lock the excitation lasers at the two nodes to ensure phase-stability of the entangled states.

2. The Device Controllers (ADwin Pro II microcontroller) are triggered every 1 μs by the same signal generator, advancing the state machine algorithm that implements the physical layer. This ensures that the two microcontrollers have a common shared clock. Over longer distances, one could use existing protocols (and commercially available hardware) to obtain a shared clock [38], and use that to trigger the microcontrollers.

3. The two AWGs need to be triggered to play entanglement attempts. In our implementation, one device controller—the server's—triggers both AWGs. This ensures that the triggering delay between the two AWGs is constant, and we can therefore calibrate it out. Triggering the AWGs with two independent microcontrollers would result in jitter (realistically on the order of nanoseconds). Over larger distances, one could derive—from the shared clock—a periodic trigger signal that is gated by the microcontroller at each node. In this way the microcontroller can decide whether the AWG will be triggered on the next cycle, but the accuracy of the trigger's timing will be derived from the shared clock between the nodes, rather than from the microprocessor.

4. The phase stabilization scheme we use, developed in Ref. [15], is designed to work at a single optical frequency (in our case, the 637 nm emission frequency of the NV center). Over longer distances, conversion of the NV center photons to the telecom band will be necessary to overcome photon loss. The phase stabilization scheme will therefore need to be adapted to new optical frequencies used.

For reference, our client (server) is based on node Charlie (Bob) of the multi-node quantum network presented in Ref. [15].

5.5.4 Applications

Following are pseudocode sequences that describe the applications executed via the quantum network stack. Their purpose is to outline how the applications were executed (sweep order), and the differences between the three applications.

Sequence 5.1: Full quantum state tomography.

```
 1  # The common list of correlators that are going to be measured (client, server).
 2  correlator_list = [
 3      (-X, -X), (-X, -Y), (-X, -Z), (-X, +X), (-X, +Y), (-X, +Z),
 4      (-Y, -X), (-Y, -Y), (-Y, -Z), (-Y, +X), (-Y, +Y), (-Y, +Z),
 5      (-Z, -X), (-Z, -Y), (-Z, -Z), (-Z, +X), (-Z, +Y), (-Z, +Z),
 6      (+X, -X), (+X, -Y), (+X, -Z), (+X, +X), (+X, +Y), (+X, +Z),
 7      (+Y, -X), (+Y, -Y), (+Y, -Z), (+Y, +X), (+Y, +Y), (+Y, +Z),
 8      (+Z, -X), (+Z, -Y), (+Z, -Z), (+Z, +X), (+Z, +Y), (+Z, +Z)]
 9
10  def client_application():
11      for rep in range(125):
12          for corr in correlator_list:
13              client_basis = corr[0]
14              # Establish the entangled state.
15              client_qubit = create_ent(
16                  with=Server,
17                  req_type=Keep,
18                  min_fidelity=0.8)
19              # Perform the required rotation.
20              client_qubit.rotate_basis(client_basis)
21              # Measure the qubit, store the result.
22              outcomes[rep, corr] = client_qubit.measure()
23
24  def server_application():
25      for rep in range(125):
26          for corr in correlator_list:
27              server_basis = corr[1]
28              # Establish the entangled state.
29              server_qubit = receive_ent(
30                  with=Client,
31                  req_type=Keep,
32                  min_fidelity=0.8)
33              # Perform the required rotation.
34              server_qubit.rotate_basis(server_basis)
35              # Measure the qubit, store the result.
36              outcomes[rep, corr] = server_qubit.measure()
```

5

Sequence 5.2: Delivery of entangled states at varying fidelity and rate.

```python
# The common list of correlators that are going to be measured (client, server).
correlator_list = [
    (-X, -X), (-X, +X), (-Y, -Y), (-Y, +Y), (-Z, -Z), (-Z, +Z),
    (+X, -X), (+X, +X), (+Y, -Y), (+Y, +Y), (+Z, -Z), (+Z, +Z)]

# The target fidelities to generate.
fidelity_list = [0.50, 0.55, 0.60, 0.65, 0.70, 0.75, 0.80]

def client_application():
    for rep in range(125):
      for fid in fidelity_list:
        for corr in correlator_list:
          client_basis = corr[0]
          # Establish the entangled state.
          client_qubit = create_ent(
                  with=Server,
                  req_type=Keep,
                  min_fidelity=fid)
          # Perform the required rotation.
          client_qubit.rotate_basis(client_basis)
          # Measure the qubit, store the result.
          outcomes[rep, fid, corr] = client_qubit.measure()

def server_application():
    for rep in range(125):
      for fid in fidelity_list:
        for corr in correlator_list:
          server_basis = corr[1]
          # Establish the entangled state.
          server_qubit = receive_ent(
                  with=Client,
                  req_type=Keep,
                  min_fidelity=fid)
          # Perform the required rotation.
          server_qubit.rotate_basis(server_basis)
          # Measure the qubit, store the result.
          outcomes[rep, fid, corr] = server_qubit.measure()
```

Sequence 5.3: Remote preparation of a qubit on the server.

```
1  # The common list of correlators that are going to be measured (client, server).
2  correlator_list = [
3      (-X, -X), (-X, -Y), (-X, -Z), (-X, +X), (-X, +Y), (-X, +Z),
4      (-Y, -X), (-Y, -Y), (-Y, -Z), (-Y, +X), (-Y, +Y), (-Y, +Z),
5      (-Z, -X), (-Z, -Y), (-Z, -Z), (-Z, +X), (-Z, +Y), (-Z, +Z),
6      (+X, -X), (+X, -Y), (+X, -Z), (+X, +X), (+X, +Y), (+X, +Z),
7      (+Y, -X), (+Y, -Y), (+Y, -Z), (+Y, +X), (+Y, +Y), (+Y, +Z),
8      (+Z, -X), (+Z, -Y), (+Z, -Z), (+Z, +X), (+Z, +Y), (+Z, +Z)]
9
10 def client_application():
11     for rep in range(125):
12         for corr in correlator_list:
13             client_basis = corr[0]
14             # Establish the entangled state and measure in the specified basis.
15             outcomes[rep, corr] = create_ent(
16                 with=Server,
17                 req_type=RemoteStatePreparaion,
18                 measurement_basis=client_basis,
19                 min_fidelity=0.8)
20
21 def server_application():
22     for rep in range(125):
23         for corr in correlator_list:
24             server_basis = corr[1]
25             # Establish the entangled state.
26             server_qubit = receive_ent(
27                 with=Client,
28                 req_type=RemoteStatePreparation,
29                 min_fidelity=0.8)
30             # Perform the required rotation.
31             server_qubit.rotate_basis(server_basis)
32             # Measure the qubit, store the result.
33             outcomes[rep, corr] = server_qubit.measure()
```

5.5.5 Post-measurement Pauli correction

The physical layer, depending on the specific quantum platform, will deliver in general one of the four possible Bell states. With the single photon protocol we employ, the physical layer can produce either $|\Psi^+\rangle = (|01\rangle + |10\rangle)/\sqrt{2}$ or $|\Psi^-\rangle = (|01\rangle - |1\rangle)/\sqrt{2}$, see Refs. [14, 15].

We choose to offer the generation of $|\Phi^+\rangle = (|00\rangle + |11\rangle)/\sqrt{2}$ as the link layer service (the choice of the specific state is arbitrary). In principle, one could also make the generated state a parameter of the link layer request. To deliver the desired Bell state, the link layer applies, for K-type requests (entangle and keep), a Pauli correction, by requesting a π rotation around either the X or Y axis on the client's qubit. For M-type (entangle and measure) and R-type (remote state preparation) requests, the physical layer performs a measurement in a basis prespecified—using the PMG command—and reports the measurement outcome, as well as which state was generated, to the link layer. The link layer can, depending on the generated state and on the chosen measurement basis, apply a classical bit flip on the client's outcome to obtain the correct measurement statistics. In particular, the link layer flips the measurement outcome of the client in the following cases:

- State delivered $|\Psi^+\rangle$, measurement basis $\pm Y$;

- State delivered $|\Psi^-\rangle$, measurement basis $\pm X$;

- State delivered $|\Psi^+\rangle$ or $|\Psi^-\rangle$, measurement basis $\pm Z$.

5.5.6 Results with and without corrections

The data presented in the main text is corrected for known measurement errors, and events in which at least one of the two devices was in the wrong charge state are removed (the CR check following the delivery of entanglement reports zero counts). While it is useful to correct for such errors in order to obtain the most faithful reconstruction of the delivered states, these errors cannot always be avoided in a real network scenario. For completeness, we report here the same results as in Section 5.3, first without any corrections applied, and then with only the measurement error correction applied. All the results, the raw datasets, and the software to analyze them, are available at Ref. [49].

Full Quantum State Tomography

The events in which the two devices generated 0 photon counts in the following CR check were 37 for the client and 380 for the server (out of the 4500 total). When combined, (client or server in the wrong charge state), we obtain 417 events (in zero events both client and server were in the wrong charge state). Without any corrections (tomography errors or wrong charge state), we obtain the following density matrix (which has a fidelity with the target Bell state F=0.681(16)):

$$\text{Re}[\rho] = \begin{pmatrix} 0.397(9) & 0.011(9) & 0.001(7) & 0.256(14) \\ 0.011(9) & 0.058(14) & -0.005(13) & -0.007(9) \\ 0.001(7) & -0.005(13) & 0.092(12) & -0.027(13) \\ 0.256(14) & -0.007(9) & -0.027(13) & 0.452(9) \end{pmatrix},$$

$$\text{Im}[\rho] = \begin{pmatrix} 0 & 0.000(18) & -0.029(9) & 0.036(9) \\ -0.000(18) & 0 & 0.010(12) & -0.002(8) \\ 0.029(9) & -0.010(12) & 0 & -0.000(8) \\ -0.036(9) & 0.002(8) & 0.000(8) & 0 \end{pmatrix}$$

Only applying tomography error correction (but not removal of wrong charge state events) yields the following density matrix (fidelity F=0.764(13)):

$$\text{Re}[\rho] = \begin{pmatrix} 0.413(7) & 0.002(11) & -0.003(3) & 0.313(8) \\ 0.002(11) & 0.026(9) & -0.015(6) & -0.018(11) \\ -0.003(3) & -0.015(6) & 0.092(7) & -0.017(8) \\ 0.313(8) & -0.018(11) & -0.017(8) & 0.469(6) \end{pmatrix},$$

$$\text{Im}[\rho] = \begin{pmatrix} 0 & 0.003(5) & -0.015(6) & 0.036(6) \\ -0.003(5) & 0 & 0.014(9) & -0.007(7) \\ 0.015(6) & -0.014(9) & 0 & -0.003(5) \\ -0.036(6) & 0.007(7) & 0.003(5) & 0 \end{pmatrix}$$

Fidelity vs Rate

The events in which the two devices generated 0 photon counts in the following CR check were 74 for the client and 709 for the server (out of the 10 500 total). When combined, (client or server in the wrong charge state), we obtain 781 events (there were two events in which both client and server were in the wrong charge state). Without any corrections (tomography errors or wrong charge state), we obtain the following delivered fidelities:

0.454(18), 0.540(18), 0.548(17), 0.596(17), 0.640(16), 0.674(16), 0.679(15). Only applying tomography error correction (but not removal of wrong charge state events) yields the following fidelities: 0.485(15), 0.591(14), 0.592(14), 0.652(13), 0.705(13), 0.741(12), 0.753(11).

Remote State Preparation

As mentioned in the main text, for the remote state preparation analysis, we only apply the tomography error correction for the server, while remove wrong charge state events of both the server and the client. The events in which the two devices generated 0 photon counts in the following CR check were 29 for the client and 365 for the server (out of the 4500 total). When combined, (client or server in the wrong charge state), we obtain 394 events (there were zero events in which both client and server were in the wrong charge state). Following are the numerical values that result in the plot in the main text (average fidelity F=0.853(8)):

Client	Server			
	$\langle X \rangle$	$\langle Y \rangle$	$\langle Z \rangle$	Fidelity
Measured $|x\rangle$	0.634(48)	−0.123(61)	−0.004(59)	0.817(24)
Measured $|y\rangle$	−0.028(58)	−0.650(45)	0.005(61)	0.825(23)
Measured $|0\rangle$	−0.081(65)	−0.083(66)	0.849(31)	0.925(16)
Measured $|-x\rangle$	−0.645(43)	0.135(59)	0.030(63)	0.823(22)
Measured $|-y\rangle$	0.026(65)	0.719(40)	−0.013(61)	0.860(20)
Measured $|1\rangle$	0.032(58)	−0.069(58)	−0.736(39)	0.868(19)

Without any corrections (tomography errors or wrong charge state), we obtain the following prepared states, with average fidelity F=0.807(10):

Client	Server			
	$\langle X \rangle$	$\langle Y \rangle$	$\langle Z \rangle$	Fidelity
Measured $|x\rangle$	0.534(55)	−0.090(62)	0.009(62)	0.767(27)
Measured $|y\rangle$	0.024(60)	−0.582(51)	−0.013(62)	0.791(26)
Measured $|0\rangle$	−0.073(69)	−0.072(69)	0.786(42)	0.893(21)
Measured $|-x\rangle$	−0.552(49)	0.143(61)	0.055(63)	0.776(24)
Measured $|-y\rangle$	0.052(64)	0.623(47)	−0.018(62)	0.811(23)
Measured $|1\rangle$	0.030(57)	−0.028(55)	−0.606(46)	0.803(23)

When only applying tomography error correction, we find an average preparation fidelity F=0.829(9):

Client	Server			Fidelity
	$\langle X \rangle$	$\langle Y \rangle$	$\langle Z \rangle$	
Measured $|x\rangle$	0.573(49)	−0.096(59)	0.010(58)	0.786(24)
Measured $|y\rangle$	0.025(56)	−0.624(45)	−0.014(59)	0.812(23)
Measured $|0\rangle$	−0.078(64)	−0.077(65)	0.843(32)	0.921(16)
Measured $|-x\rangle$	−0.592(44)	0.153(57)	0.059(59)	0.796(22)
Measured $|-y\rangle$	0.056(61)	0.667(41)	−0.020(59)	0.834(20)
Measured $|1\rangle$	0.032(54)	−0.030(53)	−0.650(40)	0.825(20)

5.5.7 NV center resonance control

The two quantum network nodes use different techniques to control the resonance of their NV centers (see Ref. [15] for implementations details). The server uses an off-resonant charge randomization strategy: when its NV center is not on resonance (it does not pass the charge and resonance check), it can apply an off-resonant (green, 515 nm) laser pulse to shuffle the charge environment and probabilistically recover the correct charge and resonance state. The server cannot get *stuck* in a non-resonance state: in a few tens of failed CR checks and green laser pulses (overall less than 1 ms) the NV center will be in resonance again.

The client, which needs to be tuned in resonance with the other node, uses a resonant strategy. When in the wrong charge state (zero counts during CR check), it applies a resonant laser pulse (yellow, 575 nm, NV^0 zero-phonon line) to go back to NV^-. To bring NV^- in resonance with the necessary lasers, it adjusts a biasing voltage applied to the diamond sample, which shifts the resonance frequencies. This process is mostly automated. However, occasional human intervention is still required when the resonance frequencies of the NV center shift too far—for example due to a charge in the vicinity of the NV center changing position in the lattice—for the automatic mechanism to find its way back. The horizontal steps in Fig. 5.3d are due to the jumps in the client's NV optical transitions, which then require manual optimization of the laser frequencies and/or the diamond biasing voltage—depending on the magnitude of the frequency shift, it requires tens of seconds to a few minutes to recover the optimal resonance condition.

Table 5.1: List of possible outcomes of the physical layer to link layer requests.

Link layer request	Physical layer outcome	Description
`INI`	`SUCCESS`	Qubit initialization is always successful.
`MSR`	`SUCCESS_0`	Measurement outcome is $|0\rangle$.
	`SUCCESS_1`	Measurement outcome is $|1\rangle$.
`SQG`	`SUCCESS`	Single qubit gates are always successful.
`PMG`	`SUCCESS`	Updating the pre-measurement gate information is always successful.
`ENT`	`SUCCESS_PSI_PLUS`	Entanglement generation was successful, the state generated was $|\Psi^+\rangle$.
	`SUCCESS_PSI_MINUS`	Entanglement generation was successful, the state generated was $|\Psi^-\rangle$.
`ENM`	`SUCCESS_PSI_PLUS_0`	Entanglement generation was successful, the state generated was $|\Psi^+\rangle$, and the measurement outcome was $|0\rangle$.
	`SUCCESS_PSI_PLUS_1`	Entanglement generation was successful, the state generated was $|\Psi^+\rangle$, and the measurement outcome was $|1\rangle$.
	`SUCCESS_PSI_MINUS_0`	Entanglement generation was successful, the state generated was $|\Psi^-\rangle$, and the measurement outcome was $|0\rangle$.
	`SUCCESS_PSI_MINUS_1`	Entanglement generation was successful, the state generated was $|\Psi^-\rangle$, and the measurement outcome was $|1\rangle$.
`ENT, ENM`	`ENT_FAILURE`	Entanglement generation was attempted and failed.
	`ENT_SYNC_FAILURE`	Entanglement generation was not attempted because the synchronization step failed (other node is busy).
Any request	`HARDWARE_FAILURE`	The node has experienced a hardware problem and cannot fulfill requests.

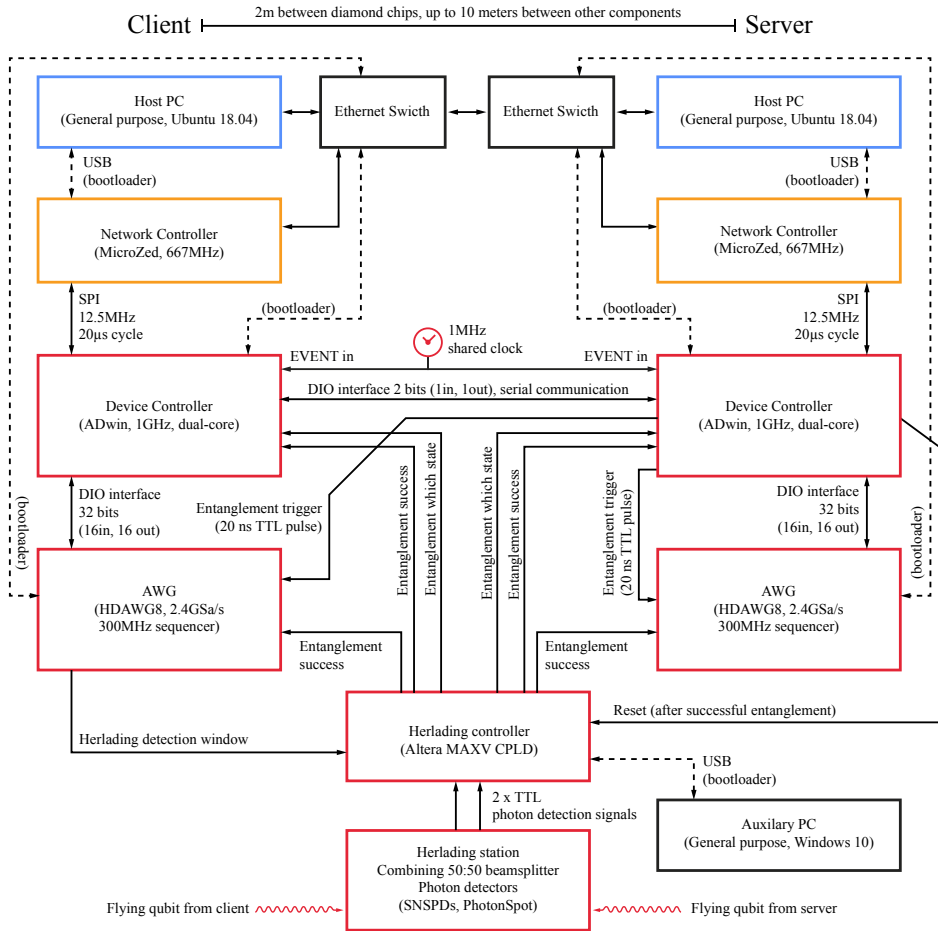**Figure 5.6:** Schematic of the connections between client and server and among the various components of each quantum network node. The dashed lines represent connections used when flashing devices (boot loading), and they are not used during the real-time operation of the network stack. Not shown are additional optical and electronic components used to control the qubits, see Ref. [15] for details on the equipment.

References

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

[2] S. Wehner, D. Elkouss, and R. Hanson, *Quantum internet: A vision for the road ahead,* Science **362** (2018), 10.1126/science.aam9288.

[3] A. Ekert and R. Renner, *The ultimate physical limits of privacy,* Nature **507**, 443 (2014).

[4] L. Jiang, J. M. Taylor, A. S. Sørensen, and M. D. Lukin, *Distributed quantum computation based on small quantum registers,* Physical Review A **76**, 062323 (2007).

[5] A. Broadbent, J. Fitzsimons, and E. Kashefi, *Universal Blind Quantum Computation,* in *2009 50th Annual IEEE Symposium on Foundations of Computer Science* (2009) pp. 517−526.

[6] D. Gottesman, T. Jennewein, and S. Croke, *Longer-Baseline Telescopes Using Quantum Repeaters,* Physical Review Letters **109**, 070503 (2012).

[7] P. Kómár, E. M. Kessler, M. Bishof, L. Jiang, A. S. Sørensen, J. Ye, and M. D. Lukin, *A quantum network of clocks,* Nature Physics **10**, 582 (2014).

[8] D. L. Moehring, P. Maunz, S. Olmschenk, K. C. Younge, D. N. Matsukevich, L.-M. Duan, and C. Monroe, *Entanglement of single-atom quantum bits at a distance,* Nature **449**, 68 (2007).

[9] L. Stephenson, D. Nadlinger, B. Nichol, S. An, P. Drmota, T. Ballance, K. Thirumalai, J. Goodwin, D. Lucas, and C. Ballance, *High-Rate, High-Fidelity Entanglement of Qubits Across an Elementary Quantum Network,* Physical Review Letters **124**, 110501 (2020).

[10] S. Ritter, C. Nölleke, C. Hahn, A. Reiserer, A. Neuzner, M. Uphoff, M. Mücke, E. Figueroa, J. Bochmann, and G. Rempe, *An elementary quantum network of single atoms in optical cavities,* Nature **484**, 195 (2012).

[11] J. Hofmann, M. Krug, N. Ortegel, L. Gérard, M. Weber, W. Rosenfeld, and H. Weinfurter, *Heralded Entanglement Between Widely Separated Atoms,* Science **337**, 72 (2012).

[12] H. Bernien, B. Hensen, W. Pfaff, G. Koolstra, M. S. Blok, L. Robledo, T. H. Taminiau, M. Markham, D. J. Twitchen, L. Childress, and R. Hanson, *Heralded entanglement between solid-state qubits separated by three metres,* Nature **497**, 86 (2013).

[13] N. Kalb, A. A. Reiserer, P. C. Humphreys, J. J. W. Bakermans, S. J. Kamerling, N. H. Nickerson, S. C. Benjamin, D. J. Twitchen, M. Markham, and R. Hanson, *Entanglement distillation between solid-state quantum network nodes,* Science **356**, 928 (2017).

[14] P. C. Humphreys, N. Kalb, J. P. J. Morits, R. N. Schouten, R. F. L. Vermeulen, D. J. Twitchen, M. Markham, and R. Hanson, *Deterministic delivery of remote entanglement on a quantum network,* Nature **558**, 268 (2018).

5

[15] M. Pompili, S. L. N. Hermans, S. Baier, H. K. C. Beukers, P. C. Humphreys, R. N. Schouten, R. F. L. Vermeulen, M. J. Tiggelman, L. d. S. Martins, B. Dirkse, S. Wehner, and R. Hanson, *Realization of a multinode quantum network of remote solid-state qubits,* Science **372**, 259 (2021).

[16] A. Delteil, Z. Sun, W.-b. Gao, E. Togan, S. Faelt, and A. Imamoğlu, *Generation of heralded entanglement between distant hole spins,* Nature Physics **12**, 218 (2016).

[17] R. Stockill, M. Stanley, L. Huthmacher, E. Clarke, M. Hugues, A. Miller, C. Matthiesen, C. Le Gall, and M. Atatüre, *Phase-Tuned Entangled State Generation between Distant Spin Qubits,* Physical Review Letters **119**, 010503 (2017).

[18] L. Aparicio, R. Van Meter, and H. Esaki, *Protocol design for quantum repeater networks,* in *Proceedings of the 7th Asian Internet Engineering Conference on - AINTEC '11* (ACM Press, Bangkok, Thailand, 2011) pp. 73–80.

[19] A. Dahlberg, M. Skrzypczyk, T. Coopmans, L. Wubben, F. Rozpędek, M. Pompili, A. Stolk, P. Pawełczak, R. Knegjens, J. de Oliveira Filho, R. Hanson, and S. Wehner, *A link layer protocol for quantum networks,* in *Proceedings of the ACM Special Interest Group on Data Communication*, SIGCOMM '19 (Association for Computing Machinery, New York, NY, USA, 2019) pp. 159–173.

[20] A. Pirker and W. Dür, *A quantum network stack and protocols for reliable entanglement-based networks,* New Journal of Physics **21**, 033003 (2019).

[21] W. Kozlowski, A. Dahlberg, and S. Wehner, *Designing a quantum network protocol,* in *Proceedings of the 16th International Conference on emerging Networking EXperiments and Technologies* (ACM, Barcelona Spain, 2020) pp. 1–16.

[22] A. Aguado, V. Lopez, J. P. Brito, A. Pastor, D. R. Lopez, and V. Martin, *Enabling Quantum Key Distribution Networks via Software-Defined Networking,* in *2020 International Conference on Optical Network Design and Modeling (ONDM)* (IEEE, Castelldefels, Barcelona, Spain, 2020) pp. 1–5.

[23] W. Kozlowski, F. Kuipers, and S. Wehner, *A P4 Data Plane for the Quantum Internet,* in *Proceedings of the 3rd P4 Workshop in Europe* (ACM, Barcelona Spain, 2020) pp. 49–51.

[24] M. Alshowkan, B. P. Williams, P. G. Evans, N. S. Rao, E. M. Simmerman, H.-H. Lu, N. B. Lingaraju, A. M. Weiner, C. E. Marvinney, Y.-Y. Pai, B. J. Lawrie, N. A. Peters, and J. M. Lukens, *Reconfigurable Quantum Local Area Network Over Deployed Fiber,* PRX Quantum **2**, 040304 (2021).

[25] R. Van Meter, T. Satoh, T. D. Ladd, W. J. Munro, and K. Nemoto, *Path selection for quantum repeater networks,* Networking Science **3**, 82 (2013).

[26] M. Caleffi, *Optimal Routing for Quantum Networks,* IEEE Access **5**, 22299 (2017).

[27] L. Gyongyosi and S. Imre, *Decentralized base-graph routing for the quantum internet,* Physical Review A **98**, 022310 (2018).

5

[28] M. Pant, H. Krovi, D. Towsley, L. Tassiulas, L. Jiang, P. Basu, D. Englund, and S. Guha, *Routing entanglement in the quantum internet,* npj Quantum Information **5**, 1 (2019).

[29] K. Chakraborty, F. Rozpedek, A. Dahlberg, and S. Wehner, *Distributed Routing in a Quantum Internet,* arXiv:1907.11630 [quant-ph] (2019).

[30] S. Shi and C. Qian, *Concurrent Entanglement Routing for Quantum Networks: Model and Designs,* in *Proceedings of the Annual conference of the ACM Special Interest Group on Data Communication on the applications, technologies, architectures, and protocols for computer communication* (ACM, Virtual Event USA, 2020) pp. 62–75.

[31] K. Chakraborty, D. Elkouss, B. Rijsman, and S. Wehner, *Entanglement Distribution in a Quantum Network: A Multicommodity Flow-Based Approach,* IEEE Transactions on Quantum Engineering **1**, 1 (2020).

[32] M. Skrzypczyk and S. Wehner, *Dynamic time-division multiple access in noisy intermediate-scale quantum networks,* in preparation (2021).

[33] *NetQASM SDK,* https://github.com/QuTech-Delft/netqasm.

[34] A. Dahlberg, B. van der Vecht, C. D. Donne, M. Skrzypczyk, I. t. Raa, W. Kozlowski, and S. Wehner, *NetQASM – A low-level instruction set architecture for hybrid quantum-classical programs in a quantum internet,* arXiv:2111.09823 [quant-ph] (2021).

[35] *FreeRTOS real-time operating system for microcontrollers,* https://www.freertos.org/.

[36] M. H. Abobeih, J. Cramer, M. A. Bakker, N. Kalb, M. Markham, D. J. Twitchen, and T. H. Taminiau, *One-second coherence for a single electron spin coupled to a multi-qubit nuclear-spin environment,* Nature Communications **9**, 2552 (2018).

[37] C. Bradley, J. Randall, M. Abobeih, R. Berrevoets, M. Degen, M. Bakker, M. Markham, D. Twitchen, and T. Taminiau, *A Ten-Qubit Solid-State Spin Register with Quantum Memory up to One Minute,* Physical Review X **9**, 031045 (2019).

[38] J. Serrano, M. Lipinski, T. Wlostowski, E. Gousiou, E. van der Bij, M. Cattin, and G. Daniluk, *The white rabbit project,* (2013).

[39] C. Granade, C. Ferrie, I. Hincks, S. Casagrande, T. Alexander, J. Gross, M. Kononenko, and Y. Sanders, *QInfer: Statistical inference software for quantum applications,* Quantum **1**, 5 (2017).

[40] C. Granade, J. Combes, and D. G. Cory, *Practical Bayesian tomography,* New Journal of Physics **18**, 033024 (2016).

[41] B. Nachman, M. Urbanek, W. A. de Jong, and C. W. Bauer, *Unfolding quantum computer readout noise,* npj Quantum Information **6**, 1 (2020).

[42] G. Vardoyan, S. Guha, P. Nain, and D. Towsley, *On the Stochastic Analysis of a Quantum Entanglement Distribution Switch,* IEEE Transactions on Quantum Engineering **2**, 1 (2021).

5

[43] G. Vardoyan, S. Guha, P. Nain, and D. Towsley, *On the Capacity Region of Bipartite and Tripartite Entanglement Switching,* ACM SIGMETRICS Performance Evaluation Review **48**, 45 (2021).

[44] B. C. Rose, D. Huang, Z.-H. Zhang, P. Stevenson, A. M. Tyryshkin, S. Sangtawesin, S. Srinivasan, L. Loudin, M. L. Markham, A. M. Edmonds, D. J. Twitchen, S. A. Lyon, and N. P. d. Leon, *Observation of an environmentally insensitive solid-state spin defect in diamond,* Science **361**, 60 (2018).

[45] C. Nguyen, D. Sukachev, M. Bhaskar, B. Machielse, D. Levonian, E. Knall, P. Stroganov, R. Riedinger, H. Park, M. Lončar, and M. Lukin, *Quantum Network Nodes Based on Diamond Qubits with an Efficient Nanophotonic Interface,* Physical Review Letters **123**, 183602 (2019).

[46] M. E. Trusheim, B. Pingault, N. H. Wan, M. Gündoğan, L. De Santis, R. Debroux, D. Gangloff, C. Purser, K. C. Chen, M. Walsh, J. J. Rose, J. N. Becker, B. Lienhard, E. Bersin, I. Paradeisanos, G. Wang, D. Lyzwa, A. R.-P. Montblanch, G. Malladi, H. Bakhru, A. C. Ferrari, I. A. Walmsley, M. Atatüre, and D. Englund, *Transform-Limited Photons From a Coherent Tin-Vacancy Spin in Diamond,* Physical Review Letters **124**, 023602 (2020).

[47] N. T. Son, C. P. Anderson, A. Bourassa, K. C. Miao, C. Babin, M. Widmann, M. Niethammer, J. Ul Hassan, N. Morioka, I. G. Ivanov, F. Kaiser, J. Wrachtrup, and D. D. Awschalom, *Developing silicon carbide for quantum spintronics,* Applied Physics Letters **116**, 190501 (2020).

[48] W. Kozlowski and S. Wehner, *Towards Large-Scale Quantum Networks,* in *Proceedings of the Sixth Annual ACM International Conference on Nanoscale Computing and Communication*, NANOCOM '19 (Association for Computing Machinery, New York, NY, USA, 2019) pp. 1–7.

[49] *Data and software supporting "Experimental demonstration of entanglement delivery using a quantum network stack",* Available after peer review is complete (2021), 10.4121/16912522.v1.

6

Conclusion

6.1 Summary of results

The focus of this thesis has been advancing the state of the art of entanglement-based quantum networks. Here we give a brief summary of the results of each chapter:

- In Chapter 3, we present the realization of a multi-node entanglement-based quantum network. We develop a new phase stabilization scheme for quantum network nodes, that allows single-photon entanglement generation over multiple nodes. This enables, together with a newly built third node, the distribution of Greenberger-Horne-Zeilinger states over the network without the use of post-selection. Additionally, we demonstrate the entanglement swapping protocol, which enables the distribution of heralded entangled states between nodes that do not share a direct physical link.

- In Chapter 4, we improve the capabilities of the multi-node network, by increasing the fidelity of the elementary links, introducing a new repetitive readout scheme, and adding control over a fifth qubit in the network. These advances allow for the demonstration of unconditional qubit teleportation beyond nearest-neighbors. Additionally, we show that memory qubit lifetimes under network node activity can be improved by using decoupling sequences

- In Chapter 5, we develop, implement, and demonstrate experimentally, the first two layers of a Quantum Internet stack—the physical layer and the link layer. The use of a stack abstracts the generation of entangled states into a platform-independent service, which can be used by network applications running at higher layers. To benchmark our system, we execute three platform-independent applications: quantum tomography of the delivered states, fidelity and latency estimation, and remote state preparation.

6.2 Near-future research directions

6.2.1 Automated tuning of network nodes

One of the main challenges we faced when performing the experiments of Chapters 3 and 4, was the amount of human intervention required. Combined over the two experiments, approximately one month of pure data collection and calibration was necessary—due to the low experimental data rate. On this topic, you can read Sophie Hermans' post in the QuTech Blog [1].

While the measurement and calibrations are executed automatically, a person is required to verify that all the calibrations are looking acceptable (and if not, figure out what is wrong). Additionally, specifically for the NV center, one needs to *babysit* the setups in case a significant jump in the resonance condition occurs during the experiment (which would put the measurement on hold).

The first challenge, automatic calibration, has been already addressed in literature [2, 3] for superconducting qubits using dependency graphs. Implementing a similar technique for NV centers would allow the user to specify which features of the setup are required, and let the node decide what should be calibrated and what not. For example, a qubit frequency calibration for the NV center could get triggered only if a significant change in magnetic field has been detected, or if the quality of the gates looks non-optimal from other measurements.

The second challenge, maintaining the NV center on resonance over extended periods of time, has been the topic of experimentation in our group over the last decade, but never resulted in a complete solution that *just works* [4]. Each NV center, and each resonance tuning point, is unique in terms of stability and dynamics: we find that some NV centers might be subject to discrete jumps of the optical transition frequencies—on the order of 100 MHz—every couple of minutes, while others might remain relatively unchanged for hours at a time. We also observe that intense optical activity (strong optical pulses during entanglement generation) increases the rate of said jumps. The hearth of the complexity—in keeping the device on resonance—lies in the two charge states of the NV center. Obtaining few counts during the CR check procedure can be a symptom of several things: correct charge state but incorrect frequency of the measurement and/or reset lasers, incorrect charge state (NV^0) and incorrect frequency of the charge reset laser, or a combination of them (the most challenging scenario). One needs to explore the parameter space to understand what is the most likely situation, and then inch towards the optimal solution. Additionally, a new jump might occur while trying to get the NV center on resonance, making previously acquired information outdated. Empirically, we see that a person can learn how to autonomously find back the resonance of an NV center after a jump in a few weeks of training.

Such high-dimensional *playgrounds* have been effectively dealt with in literature using artificial intelligence tools [5]. A first step in this direction—in the context of NV centers— has been taken in Ref. [6], but further developments are necessary to reach a reliable tool usable in practice.

The automatic calibration of the quantum network nodes, combined with a tool that can reliably maintain the devices on resonance, would boost significantly our platform, enabling more complex demonstrations and continuous operation of, for example, the physical layer of a quantum network.

6.2.2 Advanced decoupling sequences for memory qubits

Storing entangled states in memory qubit is a fundamental requirement for quantum network end-nodes. In Chapter 3 we have presented evidence suggesting that the entanglement generation procedure is not a limiting factor for the lifetime of memory qubits. Indeed, in Chapter 4, by applying a single decoupling pulse to the memory qubit, we increased the lifetime by a factor of five. However, the effective gain on the protocol rate was only a factor two with respect to Chapter 3. This is because of our implementation of the decoupling pulse: the limited capabilities of the AWG used at the time, forced us to first play the entanglement attempts, then—in case of successful entanglement—apply the decoupling pulse. To rephase the memory qubit, we had to wait several milliseconds before being able to use the memory qubit again.

A near-term improvement would be to maintain a fixed decoupling routine for the memory qubit, and play entanglement attempts both before and after the decoupling pulse. In case of successful entanglement generation, the unnecessary remaining entanglement attempts could be replaced by idle time until the memory qubit is rephased. This should immediately result in a factor of two improvement in memory qubit lifetime with respect to Chapter 4, and open the door to applying several decoupling pulses, interleaved by blocks of entanglement attempts. Additionally, the decoupling pulses could be played at strategic phases—such as XY8—to further shield the memory qubit from quasi-static noise and gate errors [7].

These improvements might push the lifetime of ^{13}C memory qubits to the regime where long-distance entanglement can be generated deterministically, while a memory qubit is holding onto an arbitrary quantum state, bringing a protocol such as deterministic teleportation within reach.

6.2.3 Larger networks

Solving the challenges outlined in the previous two sections—automated calibration and control, and increasing memory qubit lifetime—could enable protocols involving more than three nodes. Establishing a multipartite entangled state across four physically separated nodes should be achievable with current device technology. Considering four nodes in line configuration (e.g. a node Daisy connected only to Charlie), one could establish entanglement between Daisy and Charlie, while Alice establishes entanglement with Bob—independently and simultaneously. The average waiting time before both links would be established is only marginally higher than that required by a single link (50 % more, in case the links have the same generation rate[1]). The following steps—swapping the states to the two memory qubits and generating the entangled state between Bob and Charlie—would take the same time as they do in a three-node experiment. Therefore, given one could control and calibrate the four nodes reliably at the same time, a four-node GHZ state would be within reach with current NV center devices. A GHZ state shared between four nodes (with each node holding exactly one qubit of the state) would be the fundamental ingredient to demonstrate quantum anonymous transmission [8], which allows transmission of

---

[1]Given X_1 and X_2 the exponentially distributed random variables (with parameters λ_1 and λ_2) describing the number of entanglement attempts necessary to establish the two links, $\max(X_1, X_2)$ is the number of entanglement attempts necessary before both links are ready. $\mathbb{E}[\max(X_1, X_2)] = 1/\lambda_1 + 1/\lambda_2 - 1/(\lambda_1 + \lambda_2)$. For $\lambda_1 = \lambda_2 = \lambda$, $\mathbb{E}[\max(X_1, X_2)] = 3/2\lambda$.

classical and quantum bits, while hiding both sender and receiver, even when all physical transmissions can be monitored. It is worth noting that the full protocol—transmission of a qubit from a secret sender to a secret receiver—requires multiple rounds of four-node GHZ state generation, and would therefore be significantly slower than the experiments presented in this thesis.

Moving past four nodes will in all likelihood require advances in the rate of entanglement generation. With bulk NV center devices, like the ones used throughout this thesis, only marginal improvements are expected to be possible [9]. Embedding point defects in optical cavities, both to increase photon extraction efficiency and emission into the ZPL, would provide an avenue towards order-of-magnitude improvements in entanglement generation rate, see Ruf *et al.* [9] for a recent review. In this regard, progress has recently been reported [10] in coupling a single NV center to a fiber-based micro-cavity in a closed-cycle cryostat. Near-future improvements to the experimental setup are expected to deliver one-to-two orders of magnitude improvements in the ZPL collection efficiency, and therefore in remote entanglement generation rate (using a single-photon protocol).

An alternative option to fiber-based micro-cavities is to embed the defect into a photonic crystal cavity (usually nanometer-sized). Due to their sensitivity to external electric fields, NV centers are not good candidates for this kind of devices [9]. Group-IV color centers, such as the SiV [11], GeV [12], SnV [13] and PbV [14] in diamond, are more suited to this implementation because of their crystallographic symmetry, which makes them insensitive, to first order, to external electric fields and strain [9]. This insensitivity, however, comes at the cost of a more challenging frequency tuning of devices—which is required for emission-based entanglement generation protocols. Indeed, at the time of writing, entanglement between physically separated group-IV color centers has not been reported.

6.2.4 Network layer

A natural step forward from the results of Chapter 5 would be to implement the *network layer*. The task of the network layer would be to deliver entangled states between nodes that are not nearest-neighbors [15], such as Alice and Charlie in our quantum network. This would allow applications to request entangled states between any two nodes, without having to deal with the network topology: if two nodes are nearest-neighbors, the network layer will request the entangled state directly to the link layer; if they are not, the network layer will need to generate a number of entangled states—via the link layer—and perform entanglement swapping operations. The strategy that the network layer will use to *route* the entanglement between the end nodes is an active topic of research [15, 16]. In general, it will depend on the capabilities of the network (i.e. the gate fidelities, the qubits lifetimes, the entanglement generation rates, etc.), the network topology (there might be multiple paths to establish a given state) and other concurring operations on the network (an application might be using a link for a higher priority task).

While demonstrating the operation of a network layer in a three-node network, such as the one used in Chapters 3 and 4, will not involve solving all the above-mentioned challenges, it will require the real-time control of memory qubits via the network stack—which is an open challenge—as well as the scheduling of entanglement over two links.

References

[1] S. Hermans, *Ready, set, measure!* (2021), `https://blog.qutech.nl/2021/02/25/ready-set-measure/`.

[2] J. Kelly, P. O'Malley, M. Neeley, H. Neven, and J. M. Martinis, *Physical qubit calibration on a directed acyclic graph,* arXiv:1803.03226 [quant-ph] (2018).

[3] M. Rol, *Control for Programmable Superconducting Quantum Systems,* Ph.D. thesis, Delft University of Technology (2020).

[4] N. Kalb, *Progression of technology for quantum control or: jumping and drifting of NV centres,* (2016), `https://blog.qutech.nl/2016/12/14/progression-of-technology-for-quantum-control-or-jumping-and-drifting-of-nv-centres/`.

[5] R. Durrer, B. Kratochwil, J. Koski, A. Landig, C. Reichl, W. Wegscheider, T. Ihn, and E. Greplova, *Automated Tuning of Double Quantum Dots into Specific Charge States Using Neural Networks,* Physical Review Applied **13**, 054019 (2020).

[6] A. Marif, *Automated Tuning of Nitrogen Vacancy Center in Diamond,* Master's thesis, Delft University of Technology (2021).

[7] C. Bradley, J. Randall, M. Abobeih, R. Berrevoets, M. Degen, M. Bakker, M. Markham, D. Twitchen, and T. Taminiau, *A Ten-Qubit Solid-State Spin Register with Quantum Memory up to One Minute,* Physical Review X **9**, 031045 (2019).

[8] M. Christandl and S. Wehner, *Quantum Anonymous Transmissions,* in *Advances in Cryptology - ASIACRYPT 2005,* Lecture Notes in Computer Science, edited by B. Roy (Springer, Berlin, Heidelberg, 2005) pp. 217–235.

[9] M. Ruf, N. H. Wan, H. Choi, D. Englund, and R. Hanson, *Quantum networks based on color centers in diamond,* Journal of Applied Physics **130**, 070901 (2021).

[10] M. Ruf, M. Weaver, S. van Dam, and R. Hanson, *Resonant Excitation and Purcell Enhancement of Coherent Nitrogen-Vacancy Centers Coupled to a Fabry-Perot Microcavity,* Physical Review Applied **15**, 024049 (2021).

[11] M. K. Bhaskar, R. Riedinger, B. Machielse, D. S. Levonian, C. T. Nguyen, E. N. Knall, H. Park, D. Englund, M. Lončar, D. D. Sukachev, and M. D. Lukin, *Experimental demonstration of memory-enhanced quantum communication,* Nature **580**, 60 (2020).

[12] R. Høy Jensen, E. Janitz, Y. Fontana, Y. He, O. Gobron, I. P. Radko, M. Bhaskar, R. Evans, C. D. Rodríguez Rosenblueth, L. Childress, A. Huck, and U. Lund Andersen, *Cavity-Enhanced Photon Emission from a Single Germanium-Vacancy Center in a Diamond Membrane,* Physical Review Applied **13**, 064016 (2020).

[13] A. E. Rugar, S. Aghaeimeibodi, D. Riedel, C. Dory, H. Lu, P. J. McQuade, Z.-X. Shen, N. A. Melosh, and J. Vučković, *Quantum Photonic Interface for Tin-Vacancy Centers in Diamond,* Physical Review X **11**, 031021 (2021).

6

[14] P. Wang, T. Taniguchi, Y. Miyamoto, M. Hatano,   and T. Iwasaki, *Low-Temperature Spectroscopic Investigation of Lead-Vacancy Centers in Diamond Fabricated by High-Pressure and High-Temperature Treatment,* arXiv:2106.03413 [cond-mat, physics:quant-ph] (2021).

[15] W. Kozlowski and S. Wehner, *Towards Large-Scale Quantum Networks,* in *Proceedings of the Sixth Annual ACM International Conference on Nanoscale Computing and Communication*, NANOCOM '19 (Association for Computing Machinery, New York, NY, USA, 2019) pp. 1–7.

[16] R. Van Meter, *Quantum networking*, Networks and telecommunications series (John Wiley & Sons, 2014).

6

A

Correcting quantum tomography for known measurement errors

The process of measuring physical qubits is inevitably affected by errors. If the result of a measurement is required as input for a subsequent application (such as the feed-forward gates in Chapter 3 and Chapter 4), errors will propagate throughout the computation and the correction methods outlined here do not apply. However, when measuring a quantum state, for example to obtain its fidelity with an ideal target state, it is possible to estimate what the state would have been in the case of a measurement not affected by known errors. This is useful for a number of reasons: when calibrating gates, such as π or $\pi/2$ rotations, it is important to know the population without the readout error, in order to choose the correct amplitudes and duration for the MW pulses; when performing quantum state tomography of multi-qubit states, for example remote Bell states, to correctly assign sources of infidelity; additionally, having correct uncertainty bars on the corrected measured values is crucial to identify patterns in the data, or to avoid drawing conclusions when the uncertainty is dominating. Some parts of this appendix are taken from Section 3.5.4 and rearranged here for completeness.

The measurement error on a (multi-)qubit system can in general be described by the following relation:

$$\mathbf{m} = \hat{R}\,\mathbf{p}, \tag{A.1}$$

where $\mathbf{p} = (p_0, \ldots, p_s)^T$ is the (column) vector of expected populations, $\mathbf{m} = (m_0, \ldots, m_s)^T$ is the (column) vector of measured populations, $R_{ij} = \text{Prob}(\text{measured state is } i \mid \text{true state is } j)$ and s is the number of quantum states possible, which usually is $s = 2^q$ with q the number of qubits.

In the case of a single qubit ($q = 1, s = 2$):

$$\hat{R} = \begin{pmatrix} R_{00} & R_{01} \\ R_{10} & R_{11} \end{pmatrix} = \begin{pmatrix} F_0 & 1 - F_1 \\ 1 - F_0 & F_1 \end{pmatrix}$$

where $F_{0/1}$ are the measurement (or readout) fidelities of the states $|0\rangle$ and $|1\rangle$.

As an example,

$$m_0 = F_0 p_0 + (1 - F_1) p_1,$$

i.e. the measured population in $|0\rangle$ is given by the correctly assigned population in $|0\rangle$ plus the incorrectly assigned population in $|1\rangle$.

The process of measurement error correction is that of estimating $\mathbf{p}$ from $\mathbf{m}$ and $\hat{R}$, which are measured experimentally and are therefore subject to statistical uncertainty. In the following, experimental estimates are marked with a star ($\star$), $\mathbf{p}^\star, \mathbf{m}^\star$ and $R^\star$.

A.1 Matrix inversion

The first—and most intuitive—solution to the problem is that of matrix inversion. Indeed, it is possible to rewrite equation (A.1) as

$$\mathbf{p} = \hat{R}^{-1}\mathbf{m}, \tag{A.2}$$

with $\hat{R}^{-1}$ the inverse matrix of $\hat{R}$ ($\hat{R}^{-1}\hat{R} = \mathbb{I}$). If $\mathbf{m}$ and $\hat{R}$ are known exactly, $\mathbf{p}$ can be calculated by inverting the $\hat{R}$ matrix.

In practice, however, both $\mathbf{m}$ and $\hat{R}$ need to be determined experimentally. To estimate $\mathbf{m}$, one performs N quantum measurements and produces a histogram $\mathbf{n} = (n_0, ..., n_s)$ of the measured states. $\mathbf{n}$ will be distributed according to a Multinomial distribution,

$$\begin{aligned}
\mathrm{E}[n_i] &= N m_i \\
\mathrm{Var}[n_i] &= N m_i (1 - m_i) \\
\mathrm{Cov}[n_i, n_k] &= -N m_i m_k \quad (i \neq k)
\end{aligned} \tag{A.3}$$

An experimental estimate of $\mathbf{m}$ is $\mathbf{m}^\star = \mathbf{n}/N$,

$$\begin{aligned}
m_i^\star &= n_i/N \\
\mathrm{Var}[m_i^\star] &= m_i^\star(1 - m_i^\star)/N \\
\mathrm{Cov}[m_i^\star, m_k^\star] &= -m_i^\star m_k^\star/N \quad (i \neq k)
\end{aligned} \tag{A.4}$$

Similarly, the matrix $\hat{R}$ can be estimated via `measure_R_histogram`, detailed in Sequence A.1. The system is prepared and measured for N times, and the measurement outcomes are accumulated in the resulting histogram matrix R^h. Dividing the histogram by N yields an estimation $R^\star$ of the error matrix $\hat{R}$.

$$\begin{aligned}
\mathrm{E}[R_{ij}^\star] &= R_{ij}^h/N \\
\mathrm{Var}[R_{ij}^\star] &= R_{ij}^\star(1 - R_{ij}^\star)/N \\
\mathrm{Cov}[R_{ij}^\star, R_{kj}^\star] &= -R_{ij}^\star R_{ik}^\star/N \quad (i \neq k)
\end{aligned} \tag{A.5}$$

Here, we assume that the states can be prepared without errors. This is of course not the case in practice, but on our platform the measurement errors are considerably larger than the initialization errors (several percents, compared to a fraction of a percent), so the approximation is fair.

Once $R^\star$ and $\mathbf{m}^\star$ are obtained, it is possible to calculate $\mathbf{p}^\star$ with equation (A.2). The propagation of the uncertainty from $R^\star$ and $\mathbf{m}^\star$ to $\mathbf{p}^\star$ needs to take into account the

Sequence A.1: ESTIMATION OF THE MEASUREMENT ERROR MATRIX

```
1  def measure_R_histogram(N, s):
2      '''Measure an histogram that, when normalized, estimates the measurement error matrix.
3      Args:
4          N: number of measurements to perform per initial state.
5          s: number of states.
6      '''
7      R_histogram[0:s][0:s] = 0  # Initialize an s×s matrix that holds the histograms.
8      for state in range(s):
9          for repetition in range(N):
10             initialize_in_state(state)
11             result = measure_system()
12             R[result][state] += 1
13
14     return R_histogram
```

covariance terms $\text{Cov}[m_i^\star, m_k^\star]$ (and it is therefore very tedious even for the two qubit case $q = 2$).

The $\hat{R}$ matrix of a multi-qubit system can be decomposed into the tensor product of several single-qubit 2×2 matrices if the measurement processes are independent of each other (for example, if the qubits are in two separate devices). In the case of NV-center memory qubits, since they are measured via the communication qubit, we use the $\hat{R}$ of the communication qubit also for the memory qubit; this is because measuring the communication qubit is known to induce phase noise (but no bit-flipping noise) on the memory qubits (see Chapter 3 and Chapter 4).

A.2 Iterative Bayesian unfolding

The matrix inversion method can produce non-physical populations, i.e. probabilities smaller than 0 or greater than 1. One could, in principle, clip the values generated by the matrix inversion to the $[0, 1]$ interval, but then the problem of renormalizing the non-clipped probabilities arises.

An alternative method, which solves the non-physical outcomes of the matrix inversion, is the iterative Bayesian unfolding, introduced for high-energy physics by D'Agostini [1] and revised to generate uncertainties via a Monte Carlo simulation in Ref. [2]. Only recently, the technique has been introduced to the quantum computation and information field [3].

Bayes' rule, in terms of measured states E (events) and true states C (causes) is the following:

$$P(C_j|E_i) = \frac{P(E_i|C_j)P(C_j)}{\sum_{k=0}^{s} P(E_i|C_k)P(C_k)}. \tag{A.6}$$

Following the notation used in the previous section, $P(E_i|C_j) = R_{ij}$ and $P(C_j) = p_j$, therefore:

$$U_{ji} \equiv P(\text{true state is } j \mid \text{measured state is } i) = \frac{R_{ij}p_j}{\sum_{k=0}^{s} R_{ik}p_k}, \tag{A.7}$$

where U_{ji} is the *unfolding matrix*[1], that, given a certain initial probability distribution of causes, gives the probability that a certain cause j is behind an observed event i. It is worth noting that $\hat{U}$ is not the algebraic inverse of $\hat{R}$ and cannot be computed from $\hat{R}$ alone.

It is possible to obtain the required probability via

$$
\begin{aligned}
p_j &= P(\text{true state is } j) \\
&= \sum_{i=0}^{s} P(\text{true state is } j \mid \text{measured state is } i)\, P(\text{measured state is } i) \\
&= \sum_{i=0}^{s} m_i U_{ji}
\end{aligned}
\tag{A.8}
$$

This recursive relation can be used to obtain improved estimates of p_j. In the iterative Bayesian unfolding, one starts with a flat, zero knowledge, $p_j^0 = 1/s$. The next iteration of the probability distribution is given by

$$
p_j^{t+1} = \sum_{i=0}^{s} m_i U_{ji}^t = \sum_{i=0}^{s} \frac{m_i R_{ij} p_j^t}{\sum_{k=0}^{s} R_{ik} p_k^t}.
\tag{A.9}
$$

At each iteration, the probability distribution will get closer to the *correct* one [1]. Indeed, if the p_j are the true probabilities, the m_i are the expected probabilities, and R_{ij} are the ideal error matrix elements, the iteration will leave the probability distribution untouched. If we instead use a guess $p_j^\star$, the measured $m_i^\star$ and the measured $R_{ij}^\star$, we will get an improved guess for p_j at each round.

In practice, after approximately 10 iterations, the probability distribution has stabilized. By measuring the χ^2 distance between two successive iterations, it is possible to terminate the unfolding procedure once the required stability has been reached.

A.3 Monte Carlo uncertainty estimation

Once the estimates $R^\star$, $\mathbf{m}^\star$ and $\mathbf{p}^\star$ are measured and computed, it is possible to calculate, using a Monte Carlo simulation, the uncertainty on $\mathbf{p}^\star$ and its covariance matrix (which is necessary when computing functions of the populations, such as the correlators).

The Monte Carlo simulation essentially produces random variations of $R^\star$ and $\mathbf{m}^\star$ using an appropriate probability distribution. With each random variation, it calculates one more step of the iterative Bayesian unfolding using $\mathbf{p}^\star$. The generated ensemble of resampled $\mathbf{p}^\star$ is used to calculate the required statistical quantities, such as $\mathrm{Var}[p_j^\star]$ and $\mathrm{Cov}[p_j^\star, p_k^\star]$. The advantage of this method, compared to analytical uncertainty propagation from the matrix inversion, is that there is no assumption of normality (i.e. being Gaussian-distributed) of the data.

A.3.1 Resampling a measured probability distribution
Both $\mathbf{m}^\star$ and the columns of $R^\star$ are measured approximations of probability distributions. Following Ref. [2], it is possible to resample them using the Dirichlet distribu-

---

[1] The first index of the matrix is always the row, and the second is always the column. Here the matrix element's indices are ji (instead of ij) to keep the i for measured states and the j for true states.

tion. The Dirichlet distribution $\text{Dir}(\boldsymbol{\alpha})$, with $\boldsymbol{\alpha} = (\alpha_0, \ldots, \alpha_s), \alpha_i > 0$, is a probability distribution of probability distributions. Its support (the values it produces) is $(X_0, \ldots, X_s)$ with $X_i \geq 0, \sum_{i=0}^{s} X_i = 1$. An important property of the Dirichlet distribution is that it is the conjugate prior of the Multinomial distribution (which is what produces the data we observe): given a certain prior $\text{Dir}(\boldsymbol{\alpha})$ for the probabilities behind a Multinomial distribution, and given an observed histogram $\mathbf{n} = (n_0, \ldots, n_s)$, the updated (posterior) distribution of probabilities is still a Dirichlet distribution with parameter $\boldsymbol{\alpha} + \mathbf{n}$. The vector of parameters is therefore usually called *pseudo-counts* of the Dirichlet distribution.

Starting from an uninformative flat prior $\mathbf{1} = (1, \ldots, 1)$, we obtain the probability distribution of the probabilities behind the observed Multinomial distribution, $\text{Dir}(\mathbf{n} + \mathbf{1})$. Sampling from this posterior Dirichlet distribution will generate random variations of the probability distribution that was originally measured. An immediate advantage of using the Dirichlet distribution is that observing zero events on a certain measured state, does not automatically assign zero probability to that population.

A.3.2 Calculating the resampled populations

Once N_{MC} Monte Carlo variations of $\mathbf{m}^\star$ and $R^\star$ have been generated, indicated here with a tilde $\tilde{m}^\star, \tilde{R}^\star$, we can use them to generate N_{MC} variations of the measured probability distribution $\mathbf{p}^\star$. As mentioned above, we are essentially doing one more step of the iterative Bayesian unfolding procedure.

Given the $\mathbf{p}^\star$ estimated in the unfolding procedure (using the measured data and the measured error matrix, not the resampled ones), we can calculate, for each $\tilde{R}^\star$, an unfolding matrix $\tilde{U}^\star$ as

$$\tilde{U}_{ji}^\star = \frac{\tilde{R}_{ij}^\star p_j^\star}{\sum_{k=0}^{s} \tilde{R}_{ik}^\star p_k^\star}. \tag{A.10}$$

Now, for each pair $(\tilde{m}^\star, \tilde{U}^\star)$ generated by the Monte Carlo simulation, we can calculate the resampled $\tilde{p}^\star$:

$$\tilde{p}_j^\star = \sum_{i=0}^{s} \tilde{m}_i^\star \tilde{U}_{ji}^\star = \sum_{i=0}^{s} \frac{\tilde{m}_i^\star \, \tilde{R}_{ij}^\star \, p_j^\star}{\sum_{k=0}^{s} \tilde{R}_{ik}^\star p_k^\star}. \tag{A.11}$$

The $\tilde{p}^\star$ will be an ensemble of N_{MC} resampled unfolded state populations, which we can use to calculate statistical quantities such as variances, covariances and uncertainties.

A.3.3 Calculating the uncertainties and the covariance matrix

With the ensemble of resampled probabilities calculated in the last section, we can proceed to calculate the statistical estimators that we are interested in—the uncertainties and the covariance matrix. In practice, we will extract the uncertainties from the variance of the populations, which is the diagonal of the covariance matrix.

The covariance between two random variables X and Y is:

$$\begin{aligned} \text{Cov}[X, Y] &= \text{E}\left[(X - \text{E}[X])(Y - \text{E}[Y])\right] \\ &= \text{E}[XY] - \text{E}[X]\text{E}[Y]. \end{aligned} \tag{A.12}$$

In our case we want the covariance between any two p_i and p_j, which gives us the covariance matrix $\mathbf{\Sigma}^p$ of the populations:

$$
\begin{aligned}
\Sigma^p_{ij} = \mathrm{Cov}[p_i, p_j] &= \mathrm{E}\left[(p_i - \mathrm{E}[p_i])(p_j - \mathrm{E}[p_j])\right] \\
&= \mathrm{E}[p_i p_j] - \mathrm{E}[p_i]\mathrm{E}[p_j].
\end{aligned}
\tag{A.13}
$$

The quantities in the last line can be computed by simple multiplication and averages of the Monte Carlo populations ensemble. Once the covariance matrix is available, the variance and the uncertainty are easily calculated from its diagonal:

$$
\mathrm{Var}[p_i] = \Sigma^p_{ii},
\tag{A.14}
$$

$$
\sigma[p_i] = \sqrt{\Sigma^p_{ii}}.
\tag{A.15}
$$

A.3.4 Correlators and other functions

It is often the case that one needs to combine the populations via some function to obtain interesting quantities. One such example, often recurring in quantum experiments, is estimating the correlator $\langle ZZ \rangle = p_{00} - p_{01} - p_{10} + p_{11}$. This is a linear combination with coefficients $\mathbf{A} = (1, -1, -1, 1)$. The variance of any linear combination of random variables is

$$
\mathrm{Var}[\mathbf{A}p] = \mathbf{A}\,\mathbf{\Sigma}^p\,\mathbf{A}^{\mathrm{T}}
\tag{A.16}
$$

In the case of uncorrelated variables—in which the covariance matrix has only elements in the diagonal—the equation above reduces to the simple known formula of sums of variances with squared coefficients. However, in our case, the random variables are elements of a discrete probability distribution and are therefore strongly correlated.

As a simple example, consider calculating the sum of all the populations. Their sum is 1, and repeating the experiment will never change that, so the uncertainty should be 0. Simply summing all the variances would lead to a non-zero uncertainty, which would not make sense. Taking into account the covariance, correctly estimates the uncertainty for linear combinations of populations.

In the case of functions that are not linear combinations, a good estimate of the uncertainty can be obtained by Taylor expansion of the function. Alternatively, the ensemble of populations calculated in the Monte Carlo simulation can be used directly to calculate the required function, and then estimate the uncertainty statistically on the ensemble.

References

[1] G. D'Agostini, *A multidimensional unfolding method based on Bayes' theorem,* Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **362**, 487 (1995).

[2] G. D'Agostini, *Improved iterative Bayesian unfolding,* arXiv:1010.0632 [physics] (2010).

[3] B. Nachman, M. Urbanek, W. A. de Jong, and C. W. Bauer, *Unfolding quantum computer readout noise,* npj Quantum Information **6**, 1 (2020).

Glossary

ADC Analog to digital converter.

AOM Acousto-optic modulator, Section 2.7.

APD Avalanche photodiode.

AWG Arbitrary waveform generator, Section 2.7.

BAR Basis-alternating repetitive (readout), Chapter 4.

BSM Bell state measurement, Chapters 3 and 4.

CPLD Complex programmable logic device.

CR check Charge and resonance check, Section 2.4.6.

DAC Digital to analog converter.

DIO Digital input/output.

DQP Distributed queue protocol.

EOM Electro-optic modulator, Section 2.7.

ESR Electronic spin resonance.

FC/APC Ferrule connector, angled physical contact, Section 2.7.

FC/PC Ferrule connector, physical contact, Section 2.7.

FPGA Field-programmable gate array, Section 2.7.

GUI Graphical user interface.

HAL Hardware Abstraction Layer, Chapter 5.

IQ In-phase and Quadrature, signal modulation, Section 2.7.

LAN Local area network.

MCU Microcontroller unit, ADwin Pro II in this thesis, Section 2.7.

MHP Midpoint Heralding Protocol, Chapter 5.

MW Microwave.

NV Nitrogen-vacancy, Chapter 2.

PC Personal computer.

PER Polarization extinction ratio, Section 2.7.

PM Polarization maintaining (fiber, optical component), Section 2.7.

PSB Phonon sideband, Chapter 2.

QEGP Quantum Entanglement Generation Protocol, Chapter 5.

RF Radio frequency.

SIL Solid immersion lens, Section 2.3.

SNSPD Superconducting nanowire single photon detector, Section 2.7.

SPI Serial peripheral interface.

SSRO Single-shot readout, Chapter 2.

TCP/IP Transmission Control Protocol / Internet Protocol. Also known as the Internet Protocol suite.

TDMA Time Division Multiple Access, Chapter 5.

ZPL Zero-phonon line, Chapter 2.

Acknowledgments

Finally, we made it to the most read part of any thesis. I have to thank many people for what I achieved in these four years of PhD, both from a research perspective, and in terms of personal growth.

Ronald, you are a great scientist and an inspiring team leader. Thank you for giving me the opportunity of doing a PhD in your group. Your guidance and support made me, and the rest of the team, face the ~~problems~~ challenges we encountered one by one, without losing track of the long-term goal. Thank you also for creating an amazing work environment, in your group, in Team Diamond as a whole, and at QuTech in general. I have deeply enjoyed covering your office in post-its for your Spinoza prize, I hope they will never completely go away. Good luck on your upcoming projects—I am sure they will be, as always, great successes.

Stephanie, thank you for being my copromotor. I always enjoyed our discussions on the future Quantum Internet, and how design choices we make now will have an impact on its future. I admired your supervision of the Quantum Internet Alliance, it was a huge effort, and it resulted in great science and collaborations. I am really happy we got to demonstrate the operation of the link layer before the end of my PhD. Onto the network layer, and beyond!

I would like to thank the members of the committee, **Leo DiCarlo**, **Simon Gröblacher**, **Tracy Northup**, **Anders Sørensen** and **Rodney Van Meter**—thank you for reading my thesis, providing feedback, and taking part in the defense ceremony. **Leo**, thank you also for all the fun times in the band, and for the tips on mixer calibrations. **Simon**, even though we only spoke a few times, it was great having you as my PhD mentor. I knew that, should I have needed to, you would have been there. **Tracy**, thank you for welcoming me in Innsbruck, it was fun controlling ions surrounded by the mountains.

Sophie, this thesis is also very much your thesis. We have worked together from the very beginning. I don't think I could have wished for a better teammate. Taking turns at aligning optics to not go crazy, disentangling years of messy and botched code, debugging microseconds of AWG sequences for weeks. Not a single chapter of this thesis would have been half as good—or even possible—without you. Thank you. You are a great researcher and a great person, I wish you and Daan a happy future.

Thanks to my mate, **Will**. You were the first non-Diamond person I connected with here in Delft, and you have been part of my family for the last four years. When things were not working in lab, I could always have a chat with you, come to your place and play some tunes while having a biertje. We shared the struggles of long-distance relationships during the pandemic—it really helped me to get through it, thank you. And thanks to both you and **Luca** for getting me to go to the gym—at least for a while.

Simon, speed-coupling optical fibers with the Hulk-mode playlist blasting throughout the lab will always be one of my most fond PhD memories. You bring positivity in any discussion: everything is fine—within error bars. Making remote-remote entanglement

was an achievement in and of its own; I already miss Florian randomly dashing through (his) room calling you for dinner. Good luck in starting your group, I cannot wait to visit *your* lab in Innsbruck!

Hans, you started your master's project essentially when I started my PhD, which in practice meant I had no idea on how to answer most of your questions. Together, we learned about NV centers, cryostats, $50\,\Omega$ terminators, and to never deploy sensitive fibers in a compressor room 🐢 Thank you also for guiding me through the Dutch culture, and for teaching me how to pronounce Groningen right. I am really happy you decided to come back to Team Diamond for your PhD, good luck on your SnV project, I think it's the key to larger diamond networks.

Lisa, opposite to Hans, you started when most things were already figured out, this meant you had to do most of the catching-up by yourself. I was surprised by your ability to swiftly pick-up ADwin programming, many people struggle after years! Thank you for helping us achieve the link layer demonstration, I had a lot of fun fixing all the bugs we created ourselves. Good luck on the science communication project, and on what you decide to do next!

Team Diamond really is an incredible group of people. **Tim**, thank you for all the discussions throughout the years, I appreciated your feedback on both the nitty-gritty details and the long-term vision. I always look forward to your questions on any topic. I remember you asking me, during my Go-NoGo defense, how would a link layer experiment be different from attaching an external trigger to the existing experiment—I hope the last chapter clears that out! **Max**, when you asked me to have a look at one of your thesis' chapters, I realized I had to hurry up on mine! We started a few months from each other, but I always looked up to you as a more mature and experienced researcher. Obtaining Purcell enhancement in a pounding cryostat is an incredible achievement, and your review on diamond quantum networks is excellent. Thanks for all the fun gatherings at your house, for your beers and the board games. **Conor**, you have been a constant source of knowledge and fun in these four years. I hope Sophie and I did not annoy you *too much* with our infinite flow of questions on nuclear-spin control. I am really looking forward to working with you in Chicago, and to exploring restaurants with you, Effy and Grace. **Arian**, I am honored to have introduced you to the marvels of pizza fritta. It's always fun to play, or to have you play, devil's advocate in our discussions. Congratulations to you, **Kian** and **Marie-Christine** on TPQI. You are a great team, and I am sure you will get those nodes entangled. I will miss our *Netowrks* meetings and our coconut-iced-lattes! **Matthew**, my very first day at QuTech was also your visit to apply as a postdoc, I will always remember our clueless Koehandel game that night. Even if we never worked on a project together, I always enjoyed our lunch discussion about what was going on in lab. Good luck at QPhox! **Dr. Joe**, I had so much fun playing music, chatting, traveling and messing around with you. Thank you and **Lucy** for that amazing cocktail party at your place, and I wish you both a happy future. **Mohamed**, congratulations on the Rubicon grant! You are an amazing physicist, and I am always looking forward to your point-of-view on research and non-research topics. I hope the visa nonsense will get resolved before long, see you soon in the US. **Hans**, supervising the home-version of the RP with you was fun, especially exchanging new interpretations of quantum mechanics the students came up with. Great work on the nuclear spin pairs, it's mind-blowing that

they can keep a coherent state for minutes. **Maarten**, I had a lot of fun exploring New York with you. Good luck on your next adventure. **Matteo**, I am sure you and the rest of the team will ace all the challenges that will come up in the tin-vacancy project. In bocca al lupo! **Sjoerd**, you are the new *Git King*, thank you for taking over. And good luck to you and **Guido** on M3. **Alejandro**, **Mariagrazia**, **Nicolas** and **Benjamin**, good luck on the multi-nodes! We only had a bit of overlap, but I can already see you will have success with Alice, Bob and Charlie. They can be annoying at times, but deep down they are good kids, take care of them! Also, Mariagrazia, in four years, Ronald will be looking for a new Mariagrazia ;) **Laura**, thank you for your support during the three-node experiment. I never thought a weird wiggle in a plot would result in a great master's project. Congratulations on your new PhD position in Paris. **Nina** and **Julia**, the Cleanroom Queens. It is always fun chatting with you at the coffee table and in meetings here and there. Good luck on your PhDs! **Yanik**, despite the delays on the cryostat, you managed to get a setup built from scratch in record time. It already looks next-generation, amazing work. **Christopher**, I am happy you decided to join the group. It was fun sharing the office for a bit. Good luck with your PhD! **Damian**, thanks for our discussions on quantum-to-classical transition—it's nice sometimes to take a step back from the lasers and think about what it all means. **Peter** and **Norbert**, much of what we achieved is based on your work. Thanks for laying the foundations of single-click entanglement, and for getting Sophie and I started. **Suzanne**, thanks for introducing me to the QuTech Blog. I had a lot of fun there. **Aletta**, thanks for taking over my position in the QuTech Blog. As I said, it is a lot of fun! And good luck on your PhD. **Lorenzo**, you visited in the middle of the pandemic, so you did not get to experience Team Diamond to its fullest. I am sure when you come back, you will have a lot of fun. Thanks for the occasional remote wavemeter fix! **Takashi**, I am always looking forward to the next colorful diamond slide from your setup. **Josh**, even if not officially Team Diamond, you are always most welcome to our gatherings. It is always so interesting discussing with you. Good luck on the next QKD project. **Julius**, welcome to Team Diamond! Thanks also to all the students that have been in Team Diamond over the years, and made it an even more enjoyable place to work in, **Airat**, **Annick**, **Asier**, **Casper**, **Chris**, **Christian**, **Floris**, **Guus**, **Kamiel**, **Lars**, **Laurens**, **Magriet**, **Marianne**, **Mark**, **Otmar**, **Remon**, **Romy**, **Santi**, **Sarel**, **Stein** and **Thomas**. **Sergey** and **Sjaak**, I really enjoyed our time in B106!

Wolfgang, **Hannes**, **Bas**, **Norbert**, I wore thin your theses. Thank you for the legacy you have left in Team Diamond. We stood on your shoulders. **Hannes**, thank you for offering me a position in Chicago, I am very much looking forward to working with you!

I have to thank all the people in Stephanie's and David's group, especially for the last year. I learned a lot about classical and quantum networks. Thank you, **Carlo**, **Wojciech**, **Ingmar**, **Bart**, **Guilherme**, **Tim**, **Bas**, **Filip**, **Axel** and **Matt**. **Przemek**, thank you for the feedback, discussions and supervision during the link layer experiment. It was always fun talking with you! The update to the new software framework would have not been possible without the work and support of the QuTech software team, **Ludo**, **Ingmar**, **Pieter**, **Henri**, **Önder**, **Sidney** and **Joris**. You guys are great, keep it up!

Eliška, **Achmed** and **Anton**, it was great thinking about how to use AI to control NV centers. I am sure someday measurements will babysit themselves!

I had the pleasure to play in the QuTech Q2 band. We had a lot of fun, I played songs

I never thought I would, and in general rehearsing was a great way of unplugging from work. Thank you all, **Anne-Marije**, **Christian**, **Christopher**, **Conor**, **Floor**, **Gustavo**, **Hans**, **Jelle**, **Joe**, **Julia**, **Leo**, **Marina**, **Nicolas**, **Remco**, **Romy**, **Ruud**, **Will**.

Editing the QuTech Blog was great. Thanks to **Adriaan**, **Aletta**, **Anne-Marije**, **Guan**, **Jonas**, **Sébastian**, **Siddharth**, **Suzanne** and **Tim** for our lunches spent together and for all the fun hours spent doing anything *but* editing. **Adriaan**, thank you also for getting me started with Quantify, I am sure the coming year will be full of exciting developments.

Spin-NANO allowed me to meet great people across Europe, **Alessandro**, **Luca**, **Alejandro**, **Najwa**, **Matt**, **Samarth**, **Marco**, **Riccardo**, **Sigurd**, **Yanick**, **Stephan** and **Fabio**. Thanks for the fun times in Sheffield, Konstanz, Cambridge, Brussels and Toulouse.

Raymond, **Raymond**, **Marijn**, **Roy** thank you for your electronics magic. You saved us countless hours of debugging and always came up with top-quality devices ready to use. **Jason**, **Jelle**, **Mariya**, **Mark**, **Matt**, **Nico**, **Olaf**, **Remco**, **Roy**, **Siebe**, QuTech would not work as it does without you. Thank you for all the help and tips throughout the years. Thanks to the people from TNO, **Erwin**, **Jaco**, **Anna**, **Dirk**, **Martin**, **Ronald**. It is always good to hear an engineer's point-of-view of the mess we make in the lab, and you were always there to help. **Chantal**, thank you for always having an answer to my questions! You, **Marjia**, **Nico**, **Ingrid** and the rest of the management team and project managers keep the QuTech machine running. Thanks to a lot of people for just being around, it was great having occasional chats with all of you, **Alberto**, **Anta**, **Brian**, **Fenglei**, **Leonie**, **Liubov**, **Jake**, **Johannes**, **Mario**, **Mohammad**, **Mohsen**, **Ravi**, **Sara**, **Slava**.

Krissy, **Rick**, **Ann**, **Peter** and all the rest of the Seattle family, thank you for the weekly Zoom calls, they really helped me during the worst times of the pandemic. I am looking forward to seeing you all in person soon.

Ayla, **Federica**, **Francesco**, **Giulia**, **Giuls**, **Lorenzo**, **Luigi**, **Mattia**, **Tommaso**, siete da anni (decenni!) i pilastri della mia salute mentale. Grazie per tutte le chiamate, le vacanze passate insieme, le risate e le incursioni in Olanda. Non vedo l'ora di fare i cretini per Chicago insieme a voi. Vi voglio bene.

Tata, **Stefania**, **Nonna**, **Rosa**, **Daniele**, **Marcello**, **Cristina**, **Massimo**, **Alessandro** e tutto il resto del **Parentame**, vi voglio bene! **Fede**, ti ho visto crescere dalla Renault 5 con i fiori adesivi del liceo, alla macchinona da direttrice aziendale per una multinazionale. Sei grande, ti voglio bene. **Marta**, grazie per tutti i messaggi vocali che mi ricordano di non scomparire, ti voglio bene. **Valerio** e **Flavio**, siete fantastici e non vedo l'ora di passare pomeriggi interi su Minecraft con voi.

Mamma e **Papà**, grazie per avermi dato l'opportunità di seguire le mie passioni e i miei interessi, anche—soprattutto—quando si traducevano inevitabilmente in qualche Filippica. **Keira** e **Raksha**, grazie di essere sempre contente di vedermi, che sia dopo 10 mesi che non vi vedo, o dopo 10 minuti.

Finally, I think that—above all else—I need to thank my time in Delft for giving me the incredibly fortuitous chance of meeting you, **Grace**. Thank you for supporting me over the last years, I could not have made it without you. You make me truly happy, and I am so grateful for your love. I cannot wait to be with you in Chicago, and for what comes next. I love you 🌽

Matteo
Delft, November 2021

Curriculum Vitæ

Matteo Pompili

2017 – 2021	Ph.D. in Experimental Physics, *Technische Universiteit Delft* Thesis: Multi-Node Quantum Networks with Diamond Qubits Group: Hanson Lab, QuTech and Kavli Institute of Nanoscience Promotor: Prof. dr. ir. Ronald Hanson Copromotor: Prof. dr. Stephanie Wehner Delft, The Netherlands
2015 – 2017	M.Sc. in Physics *cum Laude*, *La Sapienza, Università di Roma* Graduation research in the Quantum Information Group Advisor: Prof. dr. Fabio Sciarrino Rome, Italy
2012 – 2015	B.Sc. in Physics *cum Laude*, *La Sapienza, Università di Roma* Graduation research in the Quantum Optics Group Advisor: Prof. dr. Paolo Mataloni Rome, Italy
2007 – 2012	Secondary School, *L. S. S. G. Peano* Rome, Italy
14/06/1993	Born in Rome, Italy

List of Publications

6. **M. Pompili***, C. Delle Donne*, I. te Raa, B. van der Vecht, M. Skrzypczyk, G. Ferreira, L. de Kluijver, A. J. Stolk, S. L. N. Hermans, P. Pawełczak, W. Kozlowski, & R. Hanson, S. Wehner. *Experimental demonstration of entanglement delivery using a quantum network stack.* arXiv:2111.11332 (2021).

5. S. L. N. Hermans*, **M. Pompili***, H. K. C. Beukers, S. Baier, J. Borregaard & R. Hanson. *Qubit teleportation between non-neighboring nodes in a quantum network.* arXiv:2110.11373 (2021).

4. **M. Pompili***, S. L. N. Hermans*, S. Baier*, H. K. C. Beukers, P. C. Humphreys, R. N. Schouten, R. F. L. Vermeulen, M. J. Tiggelman, L. dos Santos Martins, B. Dirkse, S. Wehner & R. Hanson, *Realization of a multinode quantum network of remote solid-state qubits.* Science, Vol. **372**, Issue 6539, pp. 259-264 (2021).

3. B. Dirkse, **M. Pompili**, R. Hanson, M. Walter & S. Wehner, *Witnessing entanglement in experiments with correlated noise.* 2020 Quantum Sci. Technol. **5** 035007.

2. A. Dahlberg*, M. Skrzypczyk*, T. Coopmans, L. Wubben, F. Rozpędek, **M. Pompili**, A. Stolk, P. Pawełczak, R. Knegjens, J. de Oliveira Filho, R. Hanson & S. Wehner, *A link layer protocol for quantum networks.* In Proceedings of the ACM Special Interest Group on Data Communication (2019) pp. 159–173.

1. T. Giordani, F. Flamini, **M. Pompili**, N. Viggianiello, N. Spagnolo, A. Crespi, R. Osellame, N. Wiebe, M. Walschaers, A. Buchleitner & F. Sciarrino, *Experimental statistical signature of many-body quantum interference.* Nature Photon **12**, 173–178 (2018).

Included in this thesis.

* Equally contributing authors.