Continuous Distribution of Entanglement in Quantum Networks with Regular Topologies

Lars Talsma

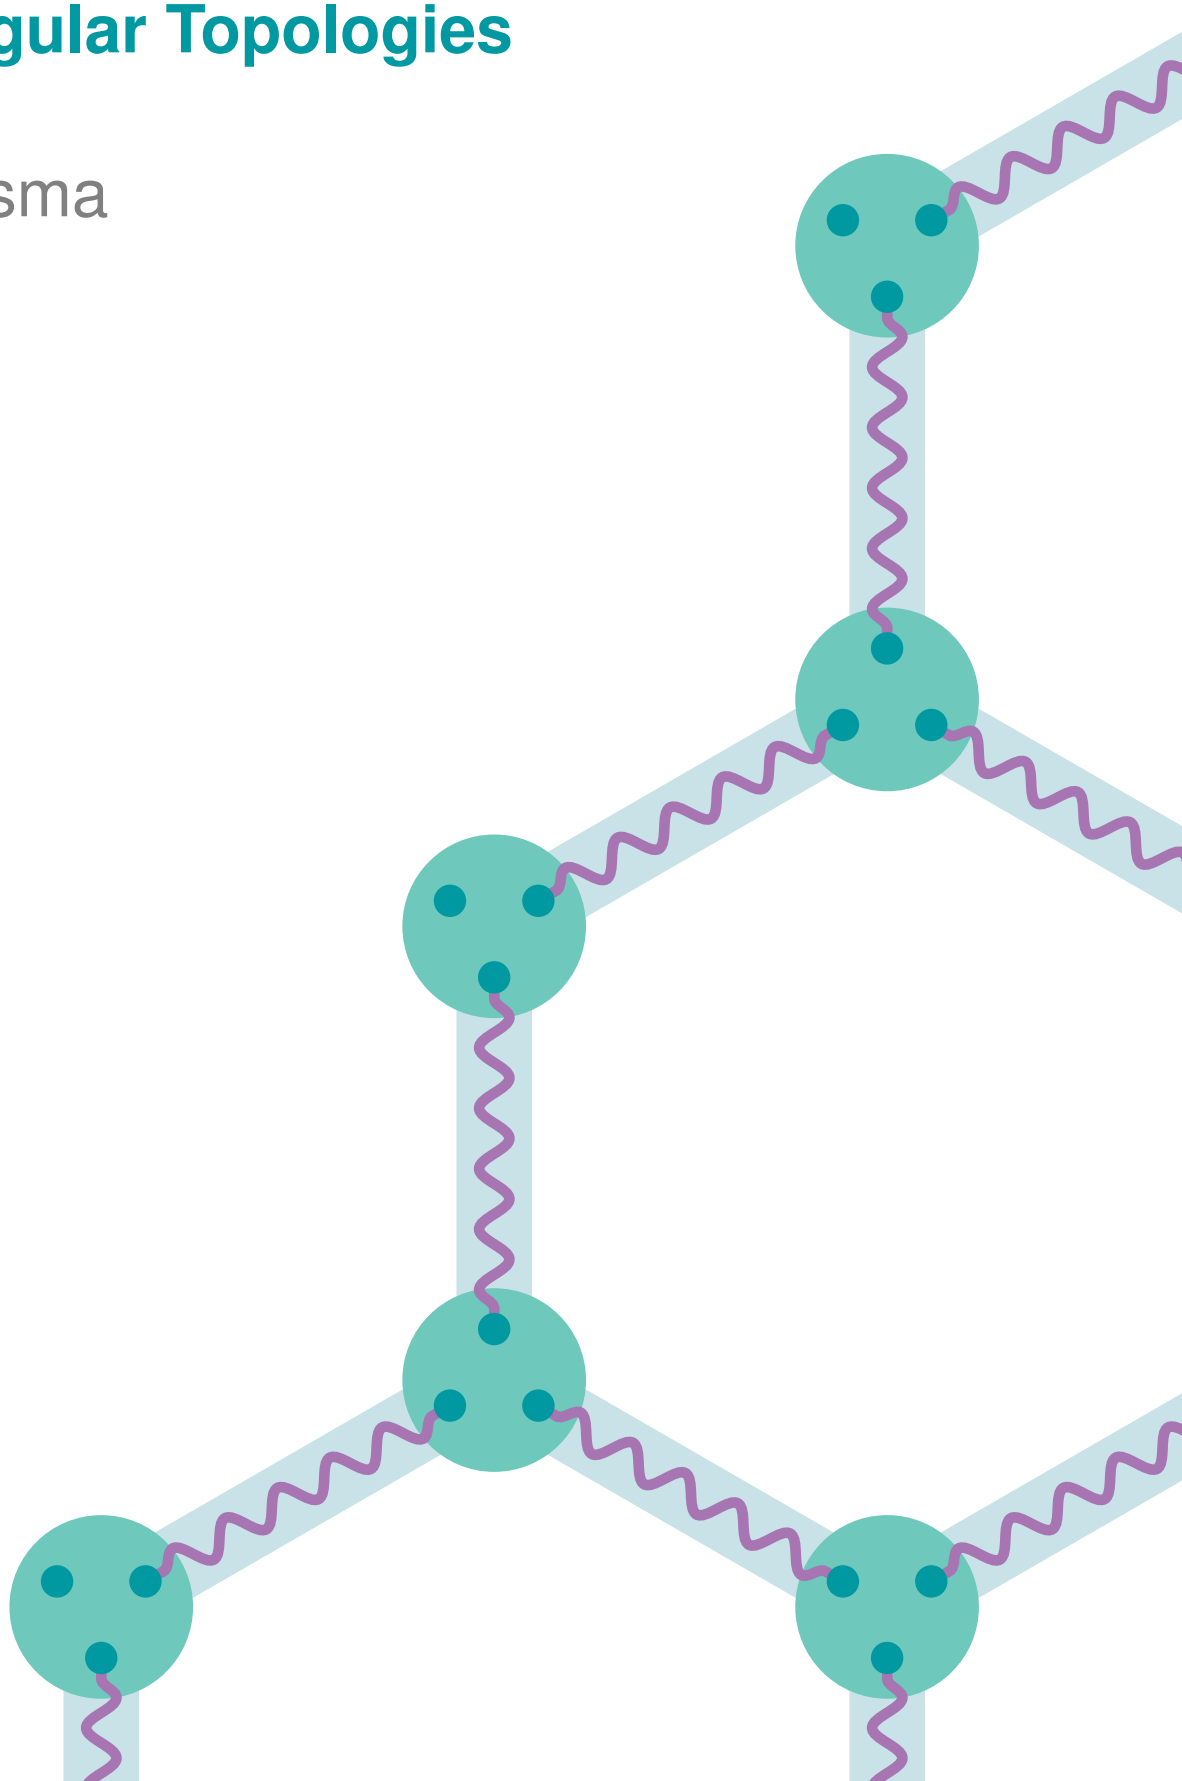

Continuous Distribution of Entanglement in Quantum Networks with Regular Topologies

by

Lars Talsma

to obtain the degree of Master of Science

at the Delft University of Technology,

to be defended publicly on Thursday January 12, 2023 at 9:00 AM.

Student number:	4477766	
Project duration:	February 28, 2022 – January 12, 2023	
Thesis committee:	Prof. Dr. S. D. C. Wehner,	TU Delft, supervisor
	Dr. M. Blaauboer,	TU Delft
	Dr. J. Borregaard,	TU Delft
Daily supervisor:	PhD candidate Á. G. Iñesta,	TU Delft

If you need to know the best counter to the queen's gambit,
you ask a chess grandmaster.

If you need to figure out the average number of steps it would take
before a randomly moving knight returns to its starting square,
you ask a mathematician.

— Kelsey Houston-Edwards

Abstract

Quantum computers can solve certain problems faster than classical computers, but they require many qubits to solve valuable problems. *Distributed quantum computing* provides a scalable approach to increasing the number of qubits by interconnecting small-capacity quantum devices, or *quantum nodes*. Entangled states shared between nodes, so-called *entangled links*, can serve as a resource for implementing nonlocal operations. A better understanding of distributing links in a network of quantum nodes can guide the design of hardware and protocols for distributed quantum computing systems. We used two metrics to measure the performance of such entanglement distribution protocols considering the network's objectives. Specifically, the *virtual node degree* reflects the requirement for many links to perform nonlocal operations, while the *virtual neighbourhood size* reflects the need for links between remote nodes to increase the number of qubits available for computation. Contrary to most prior research, these metrics explicitly consider the time-dependent fidelity of entangled links. We used discrete-time simulations to investigate the performance of a protocol that continuously distributes entanglement in a quantum network with a regular topology. The number of entangled links in the network evolves as quantum nodes create new links through entanglement generation and entanglement swaps, and remove low-fidelity links. The nodes probabilistically attempt swaps and can maximise the performance metrics by varying this probability.

We found that the performance metrics exhibit qualitatively similar behaviour for various network parameters, such as coherence time and entanglement generation fidelity. However, the network parameters shift the swap attempt probability that maximises the virtual neighbourhood size differently. The effect of network boundaries on performance metrics depends on the network topology.

Acknowledgements

It is my pleasure to thank the wonderful people I have collaborated with and have been surrounded by during the last nine months working on this thesis.

Most notably, I would like to thank my supervisor, Prof. Dr. Stephanie Wehner, and my daily supervisor, PhD candidate Álvaro Gómez Iñesta, for allowing me the opportunity to conduct my Master's thesis project in their group and setting up the amazingly supportive and productive research environment. Álvaro, thank you for your excellent guidance throughout my project! Our great weekly meetings always provided me with an excess of motivation and objectives. Your expert knowledge and incredible help have propelled this project to fruition!

Furthermore, I would like to thank the people in the Wehner group for being so welcoming. I have thoroughly enjoyed my time in the group, especially during the various activities and the (much-needed) coffee breaks. A special thanks to Helena and Judith for taking care of the administrative aspects of the project.

Moreover, QuTech has proven to be a great community with countless highly motivated people driving the second quantum revolution. Specifically, a shout-out to all the QuTech Master's students in the student's room sharing joy and sorrow; I'll fondly remember our coffee breaks and the vrijmibo's down in the TPKV. A special mention goes to Rick and Otmar, without whom I would not have enjoyed quantum *that* much!

Perhaps most importantly, I would like to thank Dr. Blaauboer and Dr. Borregaard for kindly accepting the invitation to join my thesis committee.

Additionally, I would like to thank all my friends for providing me with some well-deserved distractions. In particular, thanks to my roommates for always being down for dinner and drinks, and allowing me to vent about why Python *won't listen to me*. Also, thanks to the *intelligen't gentlemen* of the Caldera for the delightful borrels and holidays we have enjoyed together. Thanks to my tennis teammates for pushing me to maintain my physical (and mental) fitness. A shout-out to the fantastic students at Forze Hydrogen Racing for providing me (sometimes a little too much) distraction in the form of *how can we connect these two cooling tubes?* – I enjoy that greatly. Even more enjoyable are the Formula One sessions and associated drinks with my former Forze teammates. For my *friends from back in the day*, I am grateful we still have such a great connection.

Lastly, the people I have to thank the most are my family: Heit, Mem, thanks for your everlasting love and support – I wouldn't be where I am today without you. I hope we will enjoy the beautiful Côte d'Azur and the magnificent snow-covered Alps many more times together! Gitte, thank you for showing me there is more than *just* thinking.

Lars
Delft, December 2022

Contents

1

Introduction

Quantum entanglement between qubits is a uniquely quantum-mechanical phenomenon resulting in strong correlations between them – stronger than classically possible [1]. Quantum computers can leverage entangled qubits to facilitate quantum algorithms that offer an advantage over classical algorithms. For example, Shor's quantum algorithm for integer factorisation [2] enables an exponential speedup compared to classical algorithms. Such algorithms require many qubits – for example, factoring a 2048-bit semiprime requires an estimated 20 million qubits [3]. However, the number of qubits on a single processing node has increased from 2 qubits with the first experimental demonstrations of quantum algorithms in 1998 [4, 5] to more than 50 qubits in recent experiments [6–8]. Hence, the number of qubits has to increase significantly to solve valuable tasks.

Distributed quantum computing provides a scalable path for growing the number of qubits by employing a quantum network of smaller-capacity quantum devices interconnected over physical channels. The quantum devices cooperate by spreading the computations of a complex quantum problem among themselves to lower the computational requirements of single devices [9]. The remote quantum devices can use shared entangled states, which we refer to as *entangled links*, to implement nonlocal operations [10–12]. Such operations allow cooperating quantum devices to realise Shor's algorithm [13]. However, imperfect physical channels result in the generation of lower-quality entangled links. Fortunately, the quality of the entangled links, which we characterise by their *fidelity* [14], can be improved. For example, entanglement distillation may turn many low-fidelity entangled links into fewer higher-fidelity states [15, 16], and quantum error correction methods can utilise entangled links to make the network fault-tolerant [17–19]. Evidently, quantum networks for distributed quantum computing require many entangled links.

To produce many entangled links, quantum networks can use a *continuous distribution* (CD) of entanglement protocol [20]. This protocol contrasts with *on-demand delivery* of entanglement, where quantum devices request entangled links [21–24]. Generally, such an approach is more efficient regarding quantum resources. However, it requires a scheduling policy to tell the devices what to do based on the demands elsewhere in the network, which can become complex for large networks. Therefore, a distributed quantum computing network will likely use a CD protocol to distribute entangled links.

In this thesis, we are specifically interested in quantum networks for distributed quantum computing operating a CD protocol. We adopt the network model of Iñesta and Wehner [20]: a *quantum network* consists of *quantum nodes* for generating, processing and storing quantum information; and *physical channels* to distribute quantum information between the nodes. The network's objective is to distribute bipartite entangled links between nodes. The *virtual network* consists of a set of quantum nodes and their shared entangled links. Furthermore, nodes sharing entangled links are *virtual neighbours*.

We consider quantum networks with a *regular topology* and a *physical node degree* k_p. In such regular networks, each quantum node shares physical channels with k_p nodes. The nodes form a chain for $k_\mathrm{p} = 2$; a honeycomb lattice for $k_\mathrm{p} = 3$; a square lattice for $k_\mathrm{p} = 4$; and a triangular lattice for $k_\mathrm{p} = 6$. The two-dimensional topologies tile the plane regularly. We assume that all quantum nodes and physical channels are identical. The network model refrains from implementing specific realisations of qubits but abstracts away the details of physical platforms in model parameters. Furthermore, we assume that each node has an *infinite* number of quantum memories.

We simulate the evolution of virtual networks by dividing time into discrete slots and operating the CD protocol within each time slot, following Iñesta and Wehner [20]. Virtual networks evolve as quantum nodes (1) create entangled links through entanglement generation and (2) entanglement swaps, and (3) remove entangled links when their fidelity has decreased too much due to decoherence and swaps.

Motivated by the need for *many* entangled links between *remote* quantum nodes to implement nonlocal operations, we use the *virtual node degree* and *virtual neighbourhood size* [20] to characterise the performance of the CD protocol. The virtual node degree denotes the number of entangled links stored by a node at a particular time. The virtual neighbourhood size denotes the number of virtual neighbours a node has at a specific time. In contrast to most previous research, these figures of merit explicitly consider the time dependence of the virtual network.

Overall, this thesis aims to characterise the performance of quantum networks with a regular topology operating a CD protocol – performing entanglement generation, entanglement swaps and removing low-fidelity entangled links – in terms of the virtual node degree and the virtual neighbourhood size.

This thesis is organised as follows:

- In Chapter 2, we motivate our research in more detail by discussing the concepts of distributed quantum computing. Subsequently, we introduce the quantum network model and performance metrics we use to characterise quantum networks running a CD protocol. We then relate our research to state-of-the-art hardware for quantum networks and discuss how to analyse, organise and manage such networks.
- In Chapter 3, we introduce our simulation framework and present how we extract the performance metrics from the simulations.
- In Chapter 4, we present and discuss the results of our simulations and extract heuristics for designing quantum networks with regular topologies running a CD protocol.
- In Chapter 5, we summarise and reflect on our results and look beyond this work.

2

Background

We start this chapter by motivating our research in Section 2.1, introducing quantum networks and their application in distributed quantum computing. Then, in Section 2.2, we present the quantum network model that we employ in our investigation. In Section 2.3, we introduce the metrics to quantify the performance of quantum networks and compare them to other figures of merit. We finalise this chapter by putting our research in perspective in Section 2.4, discussing state of the art in quantum network hardware and how to analyse, organise and manage such quantum networks.

2.1 Quantum networks: towards distributed quantum computing

A *Quantum network* is a system of interconnected quantum devices that can leverage quantum-mechanical effects such as superposition and entanglement [25–27] to facilitate technologies such as

- *quantum communication*, enabling provably secure communication utilising quantum key distribution [28, 29],
- *quantum sensing*, allowing interferometric telescopes with longer baselines [30], and
- *distributed quantum computing*, where smaller capacity quantum computers work together to achieve results usually reserved for large-scale quantum computers [9].

We continue this section by zooming in on the concepts of distributed quantum computing, sometimes referred to as networked or modular quantum computing.

Distributed quantum computing

Interconnected quantum computers solve complex quantum computational tasks by distributing the computations between them [9, 31, 32]. By growing the number of qubits in the system with the number of quantum devices in the network, distributed quantum computing allows ensembles of quantum computers with few qubits to cooperate and perform tasks that normally require quantum computers with many qubits [9].

The small-capacity quantum computers can use short-range, high-precision interactions for local control and entanglement for nonlocal coupling [16]. Entanglement shared between quantum devices provides a physical resource to implement nonlocal operations, allowing for universal quantum computation [10–12]. We refer to an entangled state shared between two quantum nodes as an *entangled link*.

The quality of the shared entangled links, characterised by their fidelity [14], decreases over time. At some point, they might be unsuitable for implementing nonlocal operations. Fortunately, there exist techniques to supply high-fidelity entangled links. For example, entanglement distillation can turn many low-fidelity entangled links into fewer higher-fidelity states using local operations and classical communication [15, 16]. Furthermore, quantum information can be encoded in multi-qubit states to correct errors resulting in fault-tolerant quantum computing [17, 18]. For instance, nonlocal joint measurements of multi-qubit Greenberger-Horne-Zeilinger (GHZ) states can correct errors [17–19]. The GHZ states can be created by *fusing* entangled links. However, their fidelity decreases with the number of entangled links resulting in the need for entanglement distillation [19].

Quantum networks for distributed quantum computing require many entangled links, motivating us to investigate the distribution of entanglement in such networks. We now continue with presenting our approach to this investigation.

2.2 Quantum network model

In Section 2.2.1, we start by defining what constitutes a quantum network. Then, in Section 2.2.2, we present the network dynamics we consider: entanglement generation, entanglement swapping and discarding entangled links. In Section 2.2.3, we specify the network topologies that we investigate. Finally, in Section 2.2.4, we introduce the entanglement distribution protocol we use to measure the network's performance.

2.2.1 Definitions

We define the elements of a quantum network (illustrated in Figure 2.1) in order to define the network itself, adopting the definitions of the model of Iñesta and Wehner [20]:

Definition 1. A *quantum node* is a station that can generate, process and store quantum information, and exchange quantum states with other nodes over physical channels.

Definition 2. A *physical channel* can transport quantum states between two quantum nodes. Two nodes are *physical neighbours* if they are connected via a physical channel.

Definition 3. A *quantum network* is a set of quantum nodes, some of which are connected via physical channels.

Quantum nodes can, for example, use nitrogen-vacancy centres in diamond [33, 34] or trapped ions [35, 36] as qubits and be connected over optical fibres [37]. We further discuss potential physical implementations in Section 2.4.1.

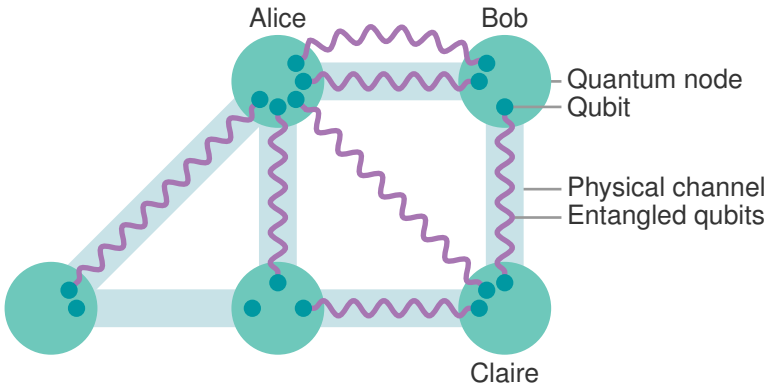

Figure 2.1: The elements of a quantum network. Quantum nodes can share any number of entangled links, and use entanglement swapping to create links between nodes that are not physical neighbours.

2.2.2 Dynamics

Our quantum networks aim to distribute entangled links among remote network nodes. The quantum nodes can use these entangled links to implement nonlocal operations between them for distributed quantum computing [12, 13]. We now elaborate on the network operations to produce entangled links suitable for nonlocal operations: entanglement generation, entanglement swapping, and discarding entangled links.

Entanglement generation

Physical neighbours can attempt to generate a shared entangled link in a *heralded* fashion. Detecting a photon, for example from the resonance fluorescence of a nitrogen-vacancy defect centre in diamond [33], can herald the success of an entangling attempt [38–40].

The entanglement generation attempt is successful with probability p, while the operation fails with probability $1-p$. No entangled link is generated in case of failure.

We model the entanglement generation process subject to *quantum noise* as a perfect entanglement generation process producing a Bell state $|\phi^+\rangle = (|00\rangle + |11\rangle)/\sqrt{2}$ followed by the application of a *depolarising* channel. Depolarising noise is a worst-case noise model where the initial $|\phi^+\rangle$ state is unaffected with probability x. However, the initial state depolarises with probability $1-x$; that is, the initial state is replaced by the completely mixed state $\mathbb{I}_4/4$ [41], where $\mathbb{I}_4$ is the 4-dimensional identity. The resulting Werner state [42] after the application of the depolarising channel on $|\phi^+\rangle$ is

$$\rho = \frac{4F-1}{3}|\phi^+\rangle\langle\phi^+| + \frac{1-F}{3}\mathbb{I}_4, \tag{2.1}$$

where $F = F(\rho, |\phi^+\rangle) \equiv \langle\phi^+|\rho|\phi^+\rangle = \frac{3}{4}x + \frac{1}{4}$ is the fidelity [14] of the newly generated Werner state relative to the Bell state $|\phi^+\rangle$. We assume that all nodes in the network generate entangled Werner states with the same fidelity F_{gen}.

Entanglement swapping

When two nodes do not share a physical channel, they can create a shared entangled link via an intermediary node using *entanglement swapping* [43]. For example, suppose Bob is a physical neighbour of Alice and Claire, while Alice and Claire do not share a physical channel. When Bob has heralded entangled links with Alice and Claire, Bob can perform a Bell-state measurement on his entangled qubits. Using local operations and classical communication, the three can implement the swap, consuming the initial links to generate a link shared between Alice and Claire, as demonstrated in Figure 2.2.

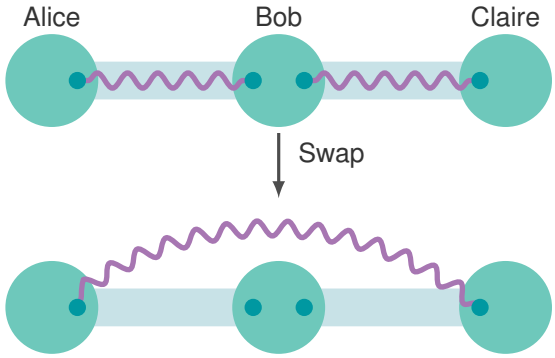

Figure 2.2: The principle of entanglement swapping. Bob can directly generate entangled links with his physical neighbours, Alice and Claire. Even though they do not share a physical channel, Alice and Claire can produce an entangled link using entanglement swapping via Bob.

We assume that nodes successfully swap entanglement with probability p_{s}. If the swap succeeds, both initial entangled links are consumed to create the new, longer link. Swaps fail with a probability $1-p_{\text{s}}$ to consume both initial states and yields no resulting state.

When quantum nodes swap two Werner states of the form (2.1) with fidelities F_1 and F_2, the resulting state is also a Werner state with fidelity [24, 44]

$$F' = F_1F_2 + \frac{(1-F_1)(1-F_2)}{3}. \tag{2.2}$$

Generally, the swap operation reduces the fidelity, $F' \leq F_1, F_2$.

Discarding entangled links

Quantum nodes can use entanglement generation and entanglement swaps to produce entangled links between many remote nodes. They can use these entangled links for nonlocal operations in distributed quantum computing. However, the entangled links should be of sufficient quality for those operations. We now discuss two processes that result in low-fidelity states and how nodes ensure that all links are of sufficient quality.

Firstly, entangled links can interact with their environment and *decohere* – their fidelity decreases over time. The initial fidelity $F(t)$ of a Werner state at time t evolves as

$$F(t+\Delta t) = \frac{1}{4} + \left(F(t) - \frac{1}{4}\right) e^{-\Delta t/T_2}, \tag{2.3}$$

under the successive application of a depolarising channel during a time interval Δt [24]. Here, T_2 is an abstract coherence time expressing the exponential decay rate of fidelity [24].

The quantum nodes discard entangled links that exist as long as some *cutoff time* t_{cut} to guarantee that all entangled links have a fidelity larger than some threshold fidelity F_{min} [24, 45–47]. The quantum nodes keep track of the *age* of entangled links – the time that has passed since the creation of the entangled links. When the age of the entangled links is equal to the cutoff time, the quantum nodes discard the entangled links.

Secondly, we assume that an entangled link created in an entanglement swap adopts the maximum age of the two initial entangled links, following Reference [24]. As the fidelity of a newly generated state after swapping is generally smaller than the fidelities of the initial entangled links (by Equation (2.2)), the entangled links can be involved in a maximum number of M swaps before their fidelity is below the threshold fidelity F_{min}. We refer to M as the *maximum swap distance*.

When quantum nodes generate entangled links with fidelity F_{gen}, the relation

$$t_{\text{cut}} \leq -T_2 \ln\left(\frac{3}{4F_{\text{gen}} - 1}\left(\frac{4F_{\text{min}} - 1}{3}\right)^{1/(M+1)}\right) \tag{2.4}$$

must be satisfied to ensure that the fidelities of the entangled links between the nodes in the network exceed the threshold fidelity F_{min} [24].

We conclude this section on network dynamics by noting that quantum nodes can use entanglement distillation to turn n low-fidelity entangled pairs into $m \leq n$ pairs of higher fidelity using local operations and classical communication. For example, when nodes share entangled links with a fidelity too low for implementing nonlocal operations, they can distil entangled links until the fidelities of the new entangled links meet their requirements. Many distillation schemes exist – for example, bipartite distillation protocols [48, 49] may distil two initial entangled links of fidelity $F > \frac{1}{2}$ to a state of fidelity $F' > F$. Hence we set $F_{\text{min}} = \frac{1}{2}$ as higher-fidelity entangled links can then be generated from the lower-fidelity initial states. However, implementing distillation protocols is outside the scope of this thesis.

2.2.3 Topology

We assume that we are free to choose the topology of quantum networks for distributed quantum computing. To modularly scale the number of nodes in the network, it makes sense to investigate network topologies that grow in a regular pattern.

Definition 4. In a quantum network with a *regular topology* and *physical node degree* k_p, each quantum node shares physical channels with k_p quantum nodes.

The quantum nodes form a chain for $k_\mathrm{p} = 2$; a honeycomb lattice for $k_\mathrm{p} = 3$; a square grid for $k_\mathrm{p} = 4$; and a triangular lattice for $k_\mathrm{p} = 6$ (illustrated in Figure 2.3). The honeycomb (or hexagonal), square and triangular tilings are the three regular tilings of the plane.

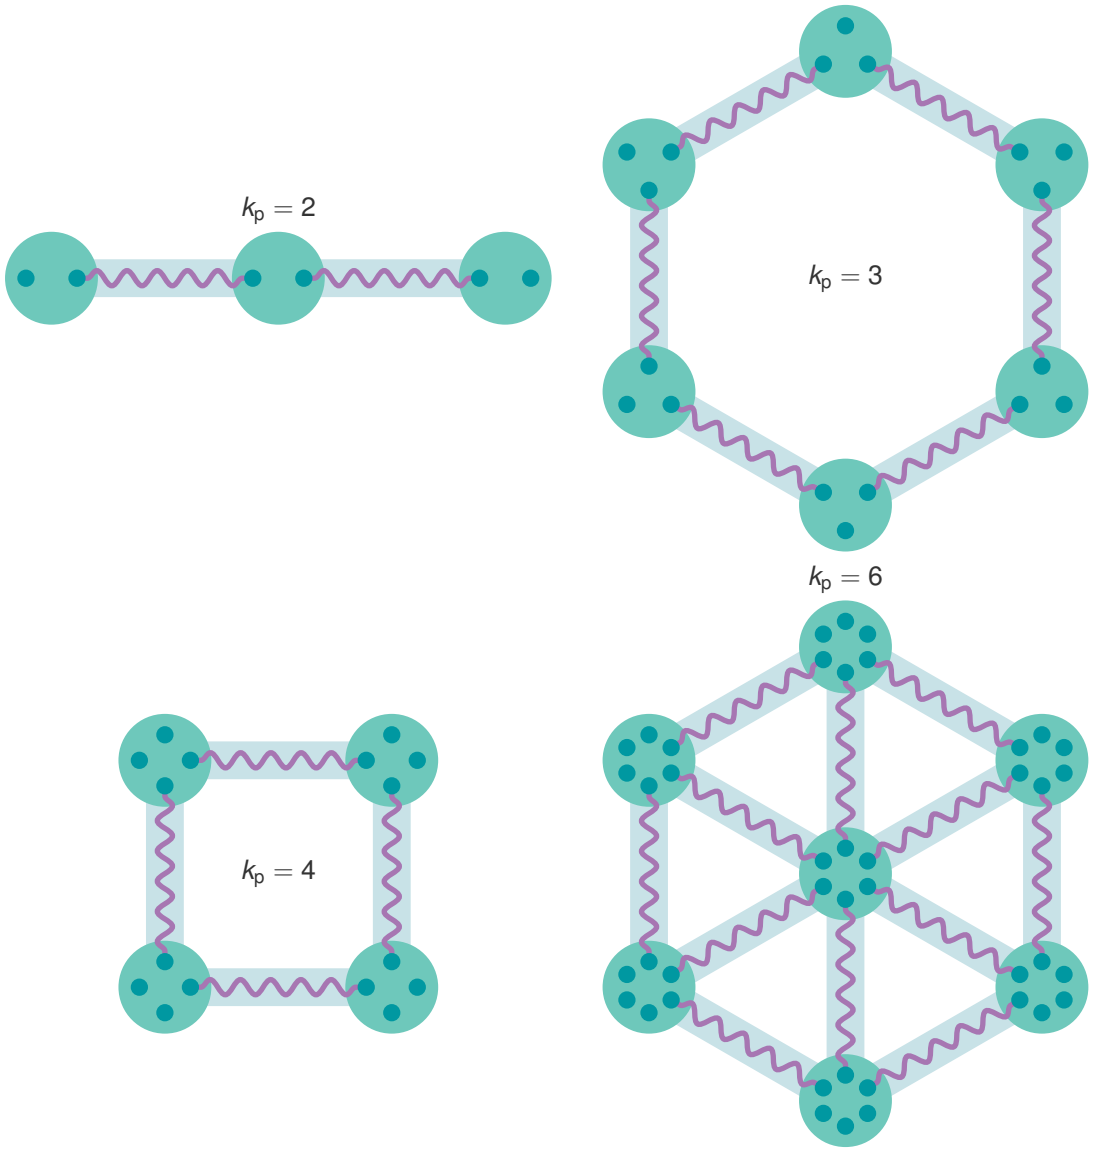

Figure 2.3: The structures emerging for the various physical node degrees: a chain of nodes, a honeycomb lattice, a square grid, and a triangular lattice. The qubits in the quantum nodes are oriented in the directions of the physical channels.

Chains of quantum nodes have been widely researched in the form of quantum repeater chains where distant nodes can generate a shared entangled link by swapping entangled links between intermediary nodes over shorter segments [44, 50]. The goal of quantum repeaters is to create end-to-end entanglement, which contrasts with the goal of our networks: distribute entangled links among all pairs of nodes in the network.

We assume that quantum nodes have an infinite number of memories – there is always a free memory to store a new entangled link. Furthermore, when a node generates an entangled link over a physical link, the entangled qubit is assigned the direction of the physical channel. We refer to this direction as the *orientation* of the qubit. For example, when a quantum node in a chain generates entanglement with a physical neighbour to its left, the entanglement is stored in a left-oriented qubit. A qubit retains its orientation over its lifetime – for example, the qubit orientation remains the same after being involved in an entanglement swap.

Lastly, we assume that the quantum networks are *homogeneous* to conclude that:

- all nodes generate entangled links with the same fidelity F_{gen} and with the same probability of success p,
- all nodes successfully attempt entanglement swaps with the same probability p_{s},
- all nodes have an infinite number of memories,
- all entangled links decohere according to the same coherence time T_2, and
- all nodes discard entangled links with an age equal to the cutoff time t_{cut} or have been involved in more than M swaps.

2.2.4 Continuous distribution of entanglement protocol

The quantum networks aim to distribute many entangled links between many quantum nodes such that the nodes can implement nonlocal operations in distributed quantum computing. The quantum networks employ a *continuous distribution* (CD) of entanglement protocol to continuously establish entanglement in the network, as proposed by Iñesta and Wehner [20]. The network discretises time, and all quantum nodes implement the CD protocol simultaneously during each time step, as prescribed by:

Continuous distribution of entanglement protocol

Cutoff time Discard entangled links with ages equal to the cutoff time t_{cut}. Nodes first apply cutoffs to ensure they do not use old links later in the protocol.

Entanglement generation Attempt to generate shared Werner states (entangled links) with physical neighbours and succeed with a probability p.

Entanglement swapping Attempt to swap two entangled links with a probability q and succeed with a probability p_{s}. A quantum node randomly chooses the first entangled link from its memory. The node chooses the second link randomly from the set of links stored in a differently-oriented qubit. If the swap succeeds, the two initial links transform into a new link; if it fails, the nodes discard the two initial links. Nodes do not know what swaps the other nodes implement.

Maximum swap distance Discard entangled links that have been involved in more than M entanglement swaps. The quantum nodes communicate their results to conclude which entangled links they have swapped.

The swap attempt probability q is a protocol parameter that nodes can vary to optimise entanglement distribution. We assume that all quantum nodes attempt swaps simultaneously. Furthermore, nodes only swap entangled links with different qubit orientations. For the chain, this means that nodes swap entangled links extending to the left and right of the nodes, preventing swaps between the same two initial quantum nodes. There is no time to communicate and coordinate desired entanglement swaps as nodes attempt swaps simultaneously. We recall that qubit orientation is preserved during a swap. We conclude this section by summarising the quantum network parameters in Table 2.1.

Table 2.1: Summary of quantum network parameters.

Parameter	Explanation
k_{p}	Physical node degree
p	Probability of successfully generating entanglement
q	Probability of attempting an entanglement swap
p_{s}	Probability of successfully implementing an entanglement swap
T_2	Abstract coherence time
t_{cut}	Cutoff time
F_{gen}	Fidelity of generated entangled links
F_{min}	Minimum required entangled link fidelity
M	Maximum swap distance

2.3 Quantum network performance metrics

Now that we have a model to measure the performance of distributing entanglement, we introduce two figures of merit to quantify the performance in Section 2.3.1. Then, in Section 2.3.2, we discuss other figures of merit used to characterise quantum networks and compare them to ours.

2.3.1 Virtual node degree and virtual neighbourhood size

It is convenient to define an auxiliary network for analysing the performance of quantum networks: the *virtual network*. We assume that the quantum networks are static in time. However, the number of entangled links in the network evolves as they can be created and destroyed via entanglement generation, entanglement swaps and discarding low-fidelity states. We analyse the performance of our networks in the framework of a virtual network, defined by its constituents, as introduced by Iñesta and Wehner [20]:

Definition 5. An *entangled link* between two quantum nodes is a bipartite Werner state between them. Two nodes sharing an entangled link are referred to as *virtual neighbours*.

Definition 6. A *virtual network* consists of a set of quantum nodes and the time-dependent set of entangled links shared among them.

Subsequently, we define the performance metrics, which we often refer to as the virtual metrics, to analyse the time-dependent connectivity of quantum networks, as introduced by Iñesta and Wehner [20]. We can divide the goal of distributing entanglement between the nodes into two parts. Firstly, quantum nodes want to share *many* entangled links with other nodes such that they can implement many nonlocal operations. This motivates us to define the first figure of merit:

Definition 7. The *virtual node degree* of node i at time t, $k_i(t)$, denotes the number of entangled links stored by node i at time t.

Secondly, quantum nodes want to share entangled links with *many remote* nodes. If a quantum node shares entangled links with many remote nodes, the virtual network they constitute contains many qubits for quantum computations. This motivates us to define the second figure of merit:

Definition 8. The *virtual neighbourhood size* of node i at time t, $v_i(t)$, denotes the number of virtual neighbours of node i at time t.

A large virtual node degree $k_i(t)$ indicates that node i shares many entangled links with other nodes at a time t. A large virtual neighbourhood size $v_i(t)$ indicates that node i has many virtual neighbours at a time t. For example, in Figure 2.1, Alice stores five entangled links, so her virtual node degree $k_{\text{Alice}} = 5$. Furthermore, she has four virtual neighbours, so her virtual neighbourhood size $v_{\text{Alice}} = 4$. Claire stores three entangled links shared with three virtual neighbours; therefore $k_{\text{Claire}} = 3$ and $v_{\text{Claire}} = 3$.

We can bound the virtual node degree and virtual neighbourhood size (Table 2.2) for *infinite* quantum networks where all quantum nodes have a physical node degree k_{p} and an infinite number of quantum memories (see Appendix B for derivations).

Table 2.2: Bounds on virtual node degree and virtual neighbourhood size.

	$k_i(t)$	$v_i(t)$
$k_\mathrm{p} = 2$	$2t_\mathrm{cut}$	$2\min(t_\mathrm{cut}, M+1)$
$k_\mathrm{p} = 3$	$3t_\mathrm{cut}$	$3\min\left(t_\mathrm{cut}, \frac{1}{2}(M+1)(M+2)\right)$
$k_\mathrm{p} = 4$	$4t_\mathrm{cut}$	$4\min\left(t_\mathrm{cut}, \frac{1}{2}(M+1)(M+2)\right)$
$k_\mathrm{p} = 6$	$6t_\mathrm{cut}$	$6\min\left(t_\mathrm{cut}, \frac{1}{2}(M+1)(M+2)\right)$

Finally, we note that optimising either the virtual node degree $k_i(t)$ or the virtual neighbourhood size $v_i(t)$ can result in contradicting strategies. For example, if nodes optimise over the probability of implementing swaps q in the CD protocol, they maximise $k_i(t)$ when they do not attempt swaps. Then, the nodes only lose entangled links when they discard them for reaching the cutoff time, and not because they consumed them in swaps or discarded them for being involved in too many swaps. Conversely, the number of virtual neighbours can exceed the number of physical neighbours when nodes apply swaps, resulting in higher $v_i(t)$. If nodes attempt too many swaps, they lose many entangled links because they consume them and because they are involved in too many swaps, decreasing $v_i(t)$. Hence, to maximise $v_i(t)$, quantum nodes must find a balance between attempting too few and too many swaps, in contradiction with optimising $k_i(t)$. In Chapter 4, we investigate the influence of the swap attempt probability on the optimal virtual metrics in detail. We first discuss why the virtual node degree and the virtual neighbourhood size are suitable figures of merit by comparing them to metrics used in different investigations.

2.3.2 Other metrics

Classical approaches to analysing classical networks have been adapted to quantum networks. For example, Reference [51] used connectivity metrics such as the node degree distribution to analyse quantum networks. Also, classical percolation theory revealed, for example, entanglement phase transitions in regular quantum networks [52]. However, our quantum networks aim to distribute many entangled links between many (remote) nodes. The classical approaches do not take dynamic phenomena like decoherence – which decreases the number of entangled links – into account and are unsuitable for analysing quantum networks.

Furthermore, one can study the time it takes to distribute end-to-end entanglement over a quantum network – often referred to as the waiting time – or the end-to-end entanglement generation rate [53]. For example, waiting times in quantum repeater chains [24, 47, 54, 55], as well as the capacity of quantum repeaters [22] and quantum network switches [23, 56, 57] have been studied. These metrics do not describe how many entangled links nodes share with remote nodes, so they are unsuitable for analysing quantum networks for distributed quantum computing.

2.4 State of the art on quantum networks

We put our research in perspective by relating it to recent developments in the area of quantum networks. In Section 2.4.1, we discuss state of the art on experimental quantum networks and couple them to implementations for distributed quantum computing. Then, in Section 2.4.2, we look beyond the hardware and discuss how to analyse, organise and manage such quantum networks.

2.4.1 Hardware for quantum networks

Linking quantum nodes over photonic channels is a promising approach for implementing quantum networks. These channels can be realised in optical fibres [37] and free space via a ground- [58] or satellite-based connection [59]. Various physical platforms have realised photon-mediated remote entanglement generation – for example, nitrogen-vacancy (NV) defect centres in diamond [33, 34], trapped ions [35, 36], neutral atoms [60, 61], quantum dots [62, 63] and rare-earth-ion doped crystals [64]. Hybrid setups, where the optically-active communication qubits cooperate with memory qubits, have been implemented to provide robust storage of quantum information, for example, by addressing the carbon-13 nuclear spins near an NV centre in diamond [65–67] or by trapping multiple species of ions [68, 69]. This combination of a photonic interface for heralded entanglement generation and a robust quantum state storage platform in a single node make NV centres in diamond and trapped ions well-suited for use in distributed quantum computing networks [17, 18, 70]. Such networks have demonstrated key protocols. For instance, experiments with NV centres have demonstrated the distillation of two low-fidelity entangled states into a higher-fidelity state [71]. Also, a deterministic remote entanglement generation protocol has been built [72] upon the intrinsically probabilistic entangling procedures (for instance, $p \approx 4 \cdot 10^{-5}$ in Reference [73]). A three-node quantum network has demonstrated entanglement swapping through an intermediary node, generating GHZ states between the three nodes and quantum teleportation of qubit states between non-neighbouring remote nodes [73, 74]. Experimentally-achieved robust memories (coherence time T_2 on the order of a second) could enable fundamental network primitives, such as the creation of entangled four-qubit GHZ states and deterministic nonlocal two-qubit gates [67]. Nonlocal gates have been implemented in trapped-ion systems [75]. Lastly, both platforms can implement deterministic entanglement swaps as complete Bell-state measurements (that is, succeeding with 100% probability) have been demonstrated using NV centres in diamond [73, 74, 76, 77] and trapped ions [78, 79].

Linear optics quantum computing offers a different but related approach to scalable quantum computation [80]. For example, *measurement-based quantum computation* implements quantum computations via projective single-qubit measurements on entangled cluster states on a (regular) lattice [81]. Small photonic cluster states can be *fused* into larger ones using Bell-state measurements [80, 82]. On a similar notion, *fusion-based quantum computing* uses resource-state generators to produce (constant-sized) few-photon entangled states and fusion devices to implement quantum computations via multi-qubit projective entanglement measurements such as (two-qubit) Bell-state measurements [83]. Combining identical modules containing resource-state generators, fusion devices, and fibres yields large-scale fault-tolerant quantum computers [84]. Although these concepts differ from the matter-based systems we discussed, they also use entanglement generation between lattice neighbours and Bell-state measurements to *grow* the computational space.

Bell-state measurements based on linear optics can generally identify two out of four Bell states: the entanglement swap succeeds with a probability $p_s = \frac{1}{2}$ [39, 85]. Adding ancillary photons to the linear optics setup can increase the success probability [86, 87]. Recent experimental work has achieved a success probability of around 58% using a pair of ancillary photons [88], and they foresee achieving a success rate of 75% with an additional ancillary photon pair or using a Bell state as an auxiliary state.

2.4.2 Quantum networks: beyond hardware

Next to advancements in hardware performance, we need insight into the performance of quantum network protocols, requirements on the underlying quantum hardware and reliable control of the quantum devices to scale up complex physics experiments into full-fledged quantum networks. We first present some analyses of quantum networks using analytical and numerical methods, and then we discuss steps towards reducing the complexity of quantum network design.

Quantum network performance

Analytical methods can be used to investigate the quantum network protocols and hardware requirements for network architectures. For example, using Markov chains, one can find the expected delivery time of end-to-end entanglement in a repeater chain [54] or investigate the capacity and the performance of distribution policies in a star-shaped quantum switch [56, 57]. The optimal entanglement distribution policy can be found using Markov decision processes for quantum repeater chains with [24] and without cutoffs [55]. Furthermore, protocols that minimise the number of Bell states necessary to produce high-fidelity GHZ states for distributed quantum computing have been investigated [19].

When the behaviour of a quantum network becomes too complex for mathematical analysis, numerical simulations can be used [53]. Simulations can be purpose-built to investigate a specific quantum network architecture or protocol – like in this thesis – or can be used to simulate a variety of quantum network architectures. For example, the software tool NetSquid – the NETwork Simulator for QUantum Information – can simulate quantum networking and modular computing systems, and model time-dependent physical non-idealities [89]. Using the NetSquid simulator, Reference [89] recovered the capacity of the previously mentioned star-shaped quantum network switch [56], extended the model to a larger range of parameters and subjected it to a more sophisticated noise model.

Quantum network stack

To reduce the complexity of quantum network design, a quantum network stack divides the functionalities of a quantum network into vertical layers of abstraction [90–92]. With these layers of abstraction, high-level layers – for instance, applications – do not have to deal with the details of low-level protocols – such as generating entanglement [91]. For example, a stack can contain [92]

- a transport layer to transmit qubits on top of
- a network layer to establish end-to-end entanglement on top of
- a link layer providing a robust entanglement generation service on top of
- a physical layer where the hardware attempts entanglement generation.

Research can be realised separately on the layers of the stack – for instance, developing a quantum data plane protocol that enables end-to-end quantum communication to fulfil the role of a quantum network layer [93] or designing and experimentally implementing a link layer protocol [92, 94].

Extending this idea to distributed quantum computing, a virtual quantum processor, assembled on top of lower-level abstractions and unaware of the physical implementation, implements the quantum algorithms using entangled qubits [9]. Exemplifying the abstraction, a physical architecture containing a one-dimensional array of resource-state generators can serve as a two-dimensional configuration of fusion measurements in fusion-based quantum computing [83]. Similarly, the physical implementations of the quantum networks can be different from the regular topologies we investigate. Furthermore, we can relate our CD protocol to the network layer – the protocol distributes entanglement over the network, but it is not involved in the specifics of generating entanglement or the details of hardware implementations.

3

Methods

In this chapter, we present the methods we use to investigate entanglement distribution in quantum networks with regular topologies running a CD protocol. First, in Section 3.1, we introduce the simulation framework. Then, in Section 3.2, we discuss how we extract and analyse the virtual node degree and virtual neighbourhood size from the simulation, noting that the virtual metrics are random variables.

3.1 Simulation

We simulate the evolution of virtual networks by discretising time into slots and operating the CD protocol in each time slot[1], following Iñesta and Wehner [20]. At the end of each time slot, we record the virtual node degree and virtual neighbourhood size, and increment the ages of the entangled links. We treat time as a dimensionless parameter such that a round of the CD protocol takes one unit of time. Furthermore, we relate the cutoff time such that a cutoff time $t_{\text{cut}} = 1$ corresponds to discarding an entangled link after one round of the CD protocol.

We simulate both finite and *infinite* quantum networks. Infinite networks with a regular topology and a physical node degree k_{p} are replicated by finite networks with periodic boundary conditions such that all nodes share a physical channel with k_{p} physical neighbours. For example, in the infinite chain, the nodes on the chain's boundaries share a physical channel. The virtual metrics of nodes in a finite network with periodic boundary conditions should converge to those in an infinite network as the number of nodes in the finite system goes to infinity. However, we can not simulate an infinite number of nodes. We, therefore, choose the number of quantum nodes in the network that minimises the influence of the finite number of nodes. For example, in an infinite chain where quantum nodes discard entangled links involved in more than M swaps, the simulated chain contains $2(M+1)+1$ nodes. We see this by noting that quantum nodes store links involved in, at most, M swaps. Then, in a chain, nodes sharing a link involved in M swaps are separated by $M+1$ physical links. Furthermore, nodes can store links in two directions, meaning that each node has $2(M+1)$ potential virtual neighbours.

[1]The Python simulations for the regular-topology quantum networks running a CD protocol investigated in this thesis can be found at `https://github.com/lars-talsma/CD-regular-networks`.

Thus, including the node itself, the simulation requires at least $2(M+1)+1$ to minimise the influence of the finite number of nodes in the network. If the simulation uses fewer quantum nodes, a node in the finite chain with periodic boundary conditions can *grow* entangled links in both directions that end up being stored in the same node, resulting in a miscount of the number of virtual neighbours $v_i(t)$. Using more quantum nodes in the simulation does not change results significantly but increases the computational requirements. In conclusion, to minimise the influence on the virtual metrics when a finite number of quantum nodes replicates an *infinite* network and the computational requirements, the simulation contains $L = 2(M+1)+1$ nodes for $k_\mathrm{p} = 2$; L^2 for $k_\mathrm{p} = 4, 6$; and $2L^2$ for $k_\mathrm{p} = 3$ (the extra factor 2 stems from implementation considerations).

Infinite networks provide a convenient platform for analysing the behaviour of virtual metrics in regular networks as each node behaves the same (see Section 3.2). Experimentally, infinite networks may be accomplished by connecting nearby boundary nodes over physical channels. For example, Reference [84] suggests appropriately connecting modules to achieve periodic boundary conditions in fusion-based quantum computing. Furthermore, for a regular finite network with a physical node degree k_p, the interior quantum nodes have k_p physical neighbours (the nodes on the network's boundary share fewer than k_p physical links with neighbours). Then, quantum nodes far away from the boundaries in finite networks perform similarly to nodes in infinite quantum networks (see Section 4.3), allowing the performance of large-scale finite networks to be approximated by infinite networks, resulting in lower computational requirements.

From the network parameters related by Equation (2.4), the simulation directly implements the cutoff time t_cut and the maximum swap distance M. We vary the coherence time T_2 and entangled link generation fidelity F_gen by associating values of t_cut and M satisfying Equation (2.4). Specifically, motivated by the behaviour of Equation (2.4), we associate the coherence time with the cutoff time and the entanglement generation fidelity with the maximum swap distance. For example, if the entanglement generation fidelity F_gen is too low, then consuming two entangled links produces a new link of fidelity $F_\mathrm{gen}^2 + (1-F_\mathrm{gen})^2/3 < F_\mathrm{min}$ (see Equation (2.2)). That is, the nodes should not implement swaps if F_gen is too low. Furthermore, when the coherence time T_2 is too short, quantum nodes must discard entangled links very early (short cutoff time t_cut) regardless of F_gen.

Lastly, we assume that each quantum node has an *infinite* number of memories. This is equivalent to the quantum nodes having at least $k_\mathrm{p} t_\mathrm{cut}$ memories. According to our CD protocol, quantum nodes in a regular network with physical node degree k_p can generate at most k_p entangled links per time slot. Quantum nodes discard entangled links when their ages are equal to the cutoff time, so they will never store more than $k_\mathrm{p} t_\mathrm{cut}$ entangled links. Hence, in our simulations, quantum nodes have $k_\mathrm{p} t_\mathrm{cut}$ memories to limit the computational requirements. We note that if quantum nodes had fewer than $k_\mathrm{p} t_\mathrm{cut}$ memories, they would require a protocol to decide which entangled links to keep in case they wanted to generate new links. Such a protocol is outside the scope of this thesis, so we assume that nodes have an infinite number of memories.

3.2 Virtual metrics

We can extract the virtual metrics from our simulation, but we note that the virtual node degree $k_i(t)$ and virtual neighbourhood size $v_i(t)$ are time-dependent random variables. However, Iñesta and Wehner [20] have shown that a unique steady-state value exists for the expected virtual degree node degree and the expected virtual neighbourhood size in quantum networks operating a CD protocol with probabilistic entanglement generation. The existence of the steady-state values motivates us to analyse the virtual metrics by averaging many samples.

Generally, we should run many simulations extracting samples of the virtual metrics $k_i(t), v_i(t)$ from a specific node i at a particular time t. However, to lessen the computational requirements, we assume that virtual metrics converge to the same steady-state value when averaged over equivalent nodes and over time (after the virtual metrics have attained their steady-state values). For infinite networks, the virtual metrics of all nodes behave the same due to the network's translational (and some rotational) symmetries. Hence, we average over all nodes. For finite networks, the virtual metrics depend on the nodes' location in the network. For example, nodes close to a network boundary can have lower virtual metrics because they have fewer physical neighbours and fewer potential virtual neighbours. However, due to symmetries, some nodes in finite networks will have the same virtual metrics. For example, the finite chain is symmetric around the centre node, so we average over nodes with the same distance to the centre node. When we analyse the performance of finite networks, we specify which nodes we average over.

For a quantum node in a quantum network with N equivalent quantum nodes simulated over T time slots, we define the average values

$$\langle k_i(t)\rangle = \langle k_i(t)\rangle_{i,t} \equiv \frac{1}{N}\frac{1}{T-t_{\mathrm{SS}}}\sum_{i=1}^{N}\sum_{t=t_{\mathrm{SS}}}^{T} k_i(t), \tag{3.1}$$

$$\langle v_i(t)\rangle = \langle v_i(t)\rangle_{i,t} \equiv \frac{1}{N}\frac{1}{T-t_{\mathrm{SS}}}\sum_{i=1}^{N}\sum_{t=t_{\mathrm{SS}}}^{T} v_i(t). \tag{3.2}$$

We only average over time after the virtual metrics reach their steady state, which we assume happens after the steady-state time t_{SS}. The virtual metrics start at zero and grow to their steady-state values during the initial time slots of the simulation. The steady-state time depends on the network parameters, but for simplicity, we set a relatively large value $t_{\mathrm{SS}} = 10t_{\mathrm{cut}}$ to encompass all scenarios. For example, it generally takes at least t_{cut} time slots to reach the steady-state regime (see Figure A.2). However, if the swap attempt probability q is large, it can take more time slots.

As discussed for infinite quantum networks, we relate the maximum swap distance to the number of quantum nodes: the simulation contains $L = 2(M+1)+1$ nodes for $k_{\mathrm{p}} = 2$; L^2 for $k_{\mathrm{p}} = 4, 6$; and $2L^2$ for $k_{\mathrm{p}} = 3$. Generally, we average the performance metrics of the infinite chain over 100 000 time slots and those of the two-dimensional lattices over 5000 time slots. With this combination of parameters, we ensure that we average the virtual metrics over (approximately) the same number of samples (the product of the number of nodes and time slots). In finite networks, we average over an increased number of time slots (specified with the results) to account for the decrease in equivalent nodes.

Finally, we note that the proof for the unique steady-state value for the expected number of virtual neighbours and expected virtual degree of any quantum node in a quantum network running the CD protocol provided by Iñesta and Wehner [20] is under the assumption that entanglement generation is probabilistic ($p < 1$). However, they expect a unique steady state for deterministic generation ($p = 1$), although they do not provide formal proof. Furthermore, we have compared simulations using $p = 0.999$ and $p = 1$ and only found marginal differences in the average virtual metrics. Although we do not conclusively show that there are steady state values for the expected virtual node degree and expected virtual neighbourhood size when $p = 1$, we feel sufficiently motivated to assume so. Deterministic entanglement generation proves to be a convenient assumption for interpreting and explaining our results.

4

Results

In this chapter, we present and discuss the results of our simulations. In Section 4.1, we start by presenting the baseline network parameters we use in our simulations. Then, in Section 4.2, we analyse the performance of infinite regular networks as a function of the swap attempt probability q and the other network parameters in terms of the virtual node degree and virtual neighbourhood size. In Section 4.3, we compare the performance in infinite networks to that in finite regular networks, focussing on the influence of the network's boundaries. Lastly, in Section 4.4, we summarise our findings as heuristics for designing quantum networks with regular networks.

4.1 Baseline network parameters

We investigate how quantum nodes in regular networks running a CD protocol can maximise their virtual node degree and virtual neighbourhood size for different combinations of network parameters. In the CD protocol, the nodes vary the swap attempt probability q to optimise the virtual metrics. Before presenting the results, we discuss how we trivialise the network parameters when we are not varying them.

We assume that entanglement generation is deterministic ($p = 1$). Individual entanglement generation attempts are generally probabilistic. However, devised upon these intrinsically probabilistic generation attempts, a deterministic entanglement distribution protocol can guarantee the delivery of entangled states at specified intervals [72]. Such a robust deterministic entanglement distribution service can be part of the link layer of a quantum network stack. Recent experiments require many entanglement generation attempts to successfully create an entangled link, corresponding to an entanglement generation success probability $p \ll 1$ (for instance, $p \approx 4 \cdot 10^{-5}$ in Reference [73]). Low entanglement generation success probabilities require many simulation steps per successful generation attempt, resulting in demanding simulation requirements compared to deterministic generation. Furthermore, deterministic entanglement generation offers a convenient platform for exploring and interpreting the simulations, as all nodes in the network generate the same number of entangled links.

We assume that entanglement swaps succeed deterministically ($p_\mathrm{s} = 1$). Platforms such as NV centres in diamond and trapped ions can realise complete Bell-state measurements to implement deterministic entanglement swaps [76, 78] (see Section 2.4.1). As we do not need to consider failed swaps, deterministic swaps are convenient for analysing networks.

We assume that the quantum nodes have an *infinite* number of memories. This means that nodes can store all entangled links until they discard them when they age to the cutoff time when they do not attempt swaps. Hence, quantum nodes never store more than $k_\mathrm{p} t_\mathrm{cut}$ entangled links (see Section 3.1 for more details). The number of required memories remains limited with short cutoff times t_cut, relatively close to experimentally achieved numbers (for example, Reference [66] achieves a ten-qubit register in diamond).

We combine a coherence time $T_2 = 100$, a cutoff time $t_\mathrm{cut} = 7$ (treating time as a dimensionless parameter; see Section 3.1), an entanglement generation fidelity $F_\mathrm{gen} = 0.9$, a maximum swap distance $M = 4$ and a minimum required fidelity $F_\mathrm{min} = \frac{1}{2}$ (quantum nodes can distil entangled links with $F > \frac{1}{2}$; see Section 2.2.2). This combination of network parameters satisfies Equation (2.4). The parameters are large enough for the behaviour of the virtual metrics to be nontrivial but manageable in terms of computational demands. For longer cutoff times, the number of occupied qubits per node grows, resulting in increased computational requirements for storing and processing the increased number of links. For the two-dimensional networks, the number of simulated nodes, and hence the simulation time, grows quadratically with the maximum swap distance (see Section 3.2). Although out of reach for recent experiments, Humphreys *et al.* [72] expect to achieve $F_\mathrm{gen} = 0.9$ with near-term improvements in the ratio of entangling and decoherence rates. Furthermore, the Bell-state measurement implementation used by Pompili *et al.* [73] takes about 1 ms while Bradley *et al.* [67] reach coherence times on the order of a second, further motivating that these network parameters are obtainable in the future.

4.2 Infinite regular networks

Infinite networks offer a practical platform for analysis as all quantum nodes behave the same way due to the system's translational (and some rotational) symmetries. This allows us to average performance metrics over the quantum nodes in the network and the number of simulation time slots (see Section 3.2), resulting in favourable simulation times. We analyse the performance of the CD protocol as a function of the swap probability q and the other network parameters using the average virtual node degree and average virtual neighbourhood size of the quantum nodes in the network. In Section 4.2.1, we first discuss the qualitatively similar behaviour of the average virtual metrics as a function of the swap attempt probability q. Then, in Sections 4.2.2 and 4.2.3, we present the results for varying coherence time T_2 and entanglement generation fidelity F_{gen}. In Section 4.2.4, we combine these results and investigate the maximum average virtual neighbourhood size as a function of T_2 and F_{gen}. Finally, in Sections 4.2.5 and 4.2.6, we investigate what happens when swaps and entanglement generation are probabilistic instead of deterministic. Let us start with investigating the general behaviour of the virtual metrics in infinite regular networks.

4.2.1 Virtual metrics as a function of swap probability

The quantum nodes can vary the probability of attempting a swap q to trade between having many entangled links between physical neighbours (not attempting swaps) or trying to share entangled links with remote quantum nodes in the network (attempting many swaps). Preferably, the quantum nodes have both a sizeable virtual node degree and an extensive virtual neighbourhood size to enable many nonlocal operations in a network containing many qubits.

We note that the behaviour of the virtual metrics is qualitatively similar for all network parameter combinations we explored (see, for example, Figure 4.1). Firstly, the average virtual node degree $\langle k_i(t)\rangle$ is largest for $q = 0$, decreases for $q > 0$ and (almost) converges to zero for $q = 1$. Secondly, the average virtual neighbourhood size $\langle v_i(t)\rangle = k_{\text{p}}$ for $q = 0$ and increases and peaks for $q > 0$ before it decreases again to (almost) zero for $q = 1$.

When nodes do not attempt swaps ($q = 0$), all entangled links are generated only to be discarded when the entangled links reach the cutoff time; that is, all quantum nodes store $k_i(t) = k_{\text{p}}t_{\text{cut}}$ entangled links. The virtual node degree decreases for $q > 0$ as nodes lose entangled links by consuming them during swaps and by discarding them if they are involved in more than M swaps (we recall that swaps succeed deterministically). Then, for $q = 1$, (almost) all entangled links are involved in more than M swaps. Quantum nodes in networks with physical node degree $k_{\text{p}} = 2, 4, 6$ generate an even number of entangled links (deterministic entanglement generation). When the nodes always attempt swaps ($q = 1$), all the initial links are consumed to create new, longer links. As all nodes implement swaps, these links are involved in more than M swaps and subsequently discarded, resulting in $k_i(t) = 0$ for all nodes. In contrast, for $k_{\text{p}} = 3$, quantum nodes generate an odd number of entangled links during entanglement generation. As entanglement swaps consume an even number of entangled links, some links might not be used in more than M swaps when $q = 1$. Hence, for $k_{\text{p}} = 3$, always swapping results in nonzero (but still almost zero) average virtual node degree $\langle k_i(t)\rangle$.

The quantum nodes only share entangled links with their physical neighbours when they do not attempt swaps ($q = 0$), resulting in $v_i(t) = k_\mathrm{p}$. The virtual neighbourhood size grows for $q > 0$ as the nodes can generate entangled links with nodes they do not share physical channels with. Increasing the number of swaps keeps increasing the virtual neighbourhood size initially. However, attempting more swaps also means that nodes lose more entangled links, resulting in fewer virtual neighbours. There is a threshold swap probability where losing virtual neighbours is balanced with generating new virtual neighbours. The average virtual neighbourhood size attains its maximum value at this swap probability – we say that this is the optimal swap probability that provides the maximum $\langle v_i(t)\rangle$. For larger swap probabilities, $\langle v_i(t)\rangle$ decreases. The average virtual neighbourhood size converges to (almost) zero for $q = 1$, as nodes discard all entangled because they are involved in too many swaps (again, not all links are removed for $k_\mathrm{p} = 3$).

Lastly, we recall that the virtual node degree and virtual neighbourhood size are random variables. We refer to Appendix C for a discussion of their probability distributions. The protocol optimisation depends on the choice of network parameters. We now present the average virtual metrics as a function of the swap probability q and network parameters.

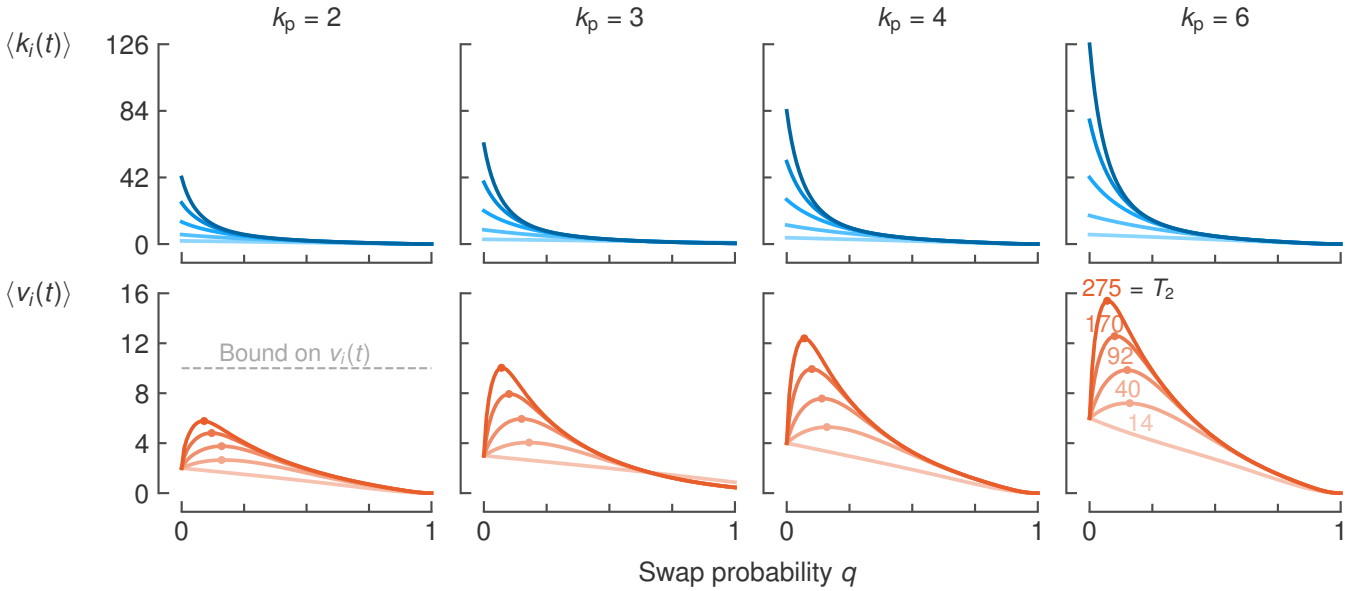

Figure 4.1: Longer coherence times increase the virtual metrics and decrease the optimal swap probability. The average virtual node degree $\langle k_i(t)\rangle$ increases for longer T_2, as nodes hold links longer before cutting them off. The maximum average virtual neighbourhood size $\langle v_i(t)\rangle$ increases for longer T_2, as nodes can have more virtual neighbours when storing more links. The optimal swap attempt probability q decreases with T_2. We vary $T_2 = 14, 40, 92, 170, 275$ (and associated cutoff times $t_\mathrm{cut} = 1, 3, 7, 13, 21$; coloured from light to dark) and use the baseline parameters $p = 1$, $p_\mathrm{s} = 1$, $F_\mathrm{gen} = 0.9$, $M = 4$, and $F_\mathrm{min} = \frac{1}{2}$. The distributions of the virtual metrics are shown in Figure C.1.

4.2.2 Coherence time

For a growing coherence time T_2 (and associated cutoff time t_cut, see Section 3.1), the maximum average virtual neighbourhood size $\langle v_i(t)\rangle$ increases (Figure 4.1; the distribution of the virtual metrics is shown in Figure C.1 for clarity). The optimal swap probability q decreases for longer T_2. As T_2 increases, entangled links live longer, meaning that nodes can store more links and have more virtual neighbours, resulting in an increased $\langle v_i(t)\rangle$.

The virtual neighbourhood size attains the maximum $\langle v_i(t)\rangle$ for decreasing optimal q as quantum nodes attempt swaps conservatively because there is more time to grow the entangled links before cutting them off. Then, attempting fewer swaps means that nodes lose fewer entangled links due to consuming links in swaps and discarding them for being involved in too many swaps (we recall that swaps always succeed, $p_\mathrm{s} = 1$), resulting in higher $\langle v_i(t)\rangle$. We note that the maximum values of $\langle v_i(t)\rangle$ scale relatively linearly with increasing T_2, but the returns slowly diminish. We note that (a function of) the maximum swap distance M bounds the virtual neighbourhood size when T_2 (and associated t_cut) is long relative to M (bounds in Table 2.2). For $k_\mathrm{p} = 2$, the virtual neighbourhood size seems to approach the limit $v_i(t) \leq 2(M+1) = 10$ as T_2 increases. However, further increasing T_2 is left to future investigations due to computational constraints.

As the physical node degree k_p increases, the maximum $\langle v_i(t)\rangle$ also increases. Specifically, for $T_2 = 14$ ($t_\mathrm{cut} = 1$), $\langle v_i(t)\rangle$ increases by a factor equal to the ratio of physical node degrees: $\langle v_i(t)\rangle$ increases due to an increase in physical neighbours but not due to swaps (nodes immediately remove swapped links in the next time slot for being too old). Then, for longer coherence times (and corresponding cutoff times), the entangled links can be involved in swaps without nodes discarding them immediately. Comparing $k_\mathrm{p} = 2$ to $k_\mathrm{p} = 3$ when $T_2 = 275$ ($t_\mathrm{cut} = 21$), $\langle v_i(t)\rangle$ increases by a factor of approximately 1.7, which is more than the ratio of the physical node degrees. We explain this by noting that the number of potential virtual nodes grows quicker for nodes in a $k_\mathrm{p} = 3$ network than one with $k_\mathrm{p} = 2$ (see also Appendix B.2). For example, when entangled links are involved in one or fewer swaps, the set of potential virtual neighbours is 4 for $k_\mathrm{p} = 2$, but 9 for $k_\mathrm{p} = 3$. Hence, with only entangled links involved in one swap, $k_\mathrm{p} = 3$ nodes can already have a sizeable $\langle v_i(t)\rangle$ compared to $k_\mathrm{p} = 2$ nodes. This means that $k_\mathrm{p} = 3$ nodes can swap less frequently to grow their $\langle v_i(t)\rangle$ compared to $k_\mathrm{p} = 2$ nodes. The optimal q associated with the maximum $\langle v_i(t)\rangle$ reflects this: $q \approx 0.09$ for $k_\mathrm{p} = 2$, but $q \approx 0.07$ for $k_\mathrm{p} = 3$. That nodes attempt fewer swaps means that they lose fewer entangled links due to consuming and discarding them for being involved in too many swaps. This increases $\langle k_i(t)\rangle$ by a factor of approximately 1.7 at the optimal q comparing $k_\mathrm{p} = 2$ to $k_\mathrm{p} = 3$.

For $k_\mathrm{p} = 4$, the maximum $\langle v_i(t)\rangle$ increases by a factor of approximately 2.1 compared to $k_\mathrm{p} = 2$ when $T_2 = 275$ ($t_\mathrm{cut} = 21$). The average virtual node degree $\langle v_i(t)\rangle$ increases as the set of potential virtual neighbours grows quicker than the physical node degree, but its influence diminishes. We explain this by the lack of direction given by the swapping routine. Specifically, entangled links are swapped randomly (as long as they are stored in qubits oriented differently), resulting in some virtual neighbours sharing multiple entangled states and some potential virtual neighbours sharing zero. The optimal swap probability q is approximately equal for $k_\mathrm{p} = 3$ and $k_\mathrm{p} = 4$. Again, attempting fewer swaps means that nodes lose fewer links increasing $\langle k_i(t)\rangle$ by a factor of about 2.3 for $k_\mathrm{p} = 4$ compared to $k_\mathrm{p} = 2$. The diminishing returns are clear for $k_\mathrm{p} = 6$, where the maximum $\langle v_i(t)\rangle$ increases by a factor of approximately 2.7 compared to $k_\mathrm{p} = 2$ when $T_2 = 275$ ($t_\mathrm{cut} = 21$). For $k_\mathrm{p} = 6$, the optimal swap probability q is approximately equal to that of $k_\mathrm{p} = 3, 4$ and $\langle k_i(t)\rangle$ increases by a factor of about 3.4 compared to $k_\mathrm{p} = 2$.

The average virtual node degree $\langle k_i(t)\rangle$ increases with T_2 as the quantum nodes can store more links due to their infinite number of quantum memories. However, they always converge qualitatively similarly to (almost) zero as $q \to 1$.

4.2.3 Entanglement generation fidelity

Increasing the fidelity of generated entangled links also increases the maximum $\langle v_i(t)\rangle$ (Figure 4.2). However, better entanglement generation fidelity increases the optimal swap probability q associated with the maximum $\langle v_i(t)\rangle$. As F_{gen} grows, links can be involved in more swaps before quantum nodes discard them for having a too-low fidelity. So if the nodes do not swap at a sufficient rate, they discard links for living to the cutoff time. However, increasing q too much means that nodes lose entangled links by attempting swaps and discarding them for being involved in too many swaps, but there can be more swaps before removal if F_{gen} is large. We see diminishing returns on the maximum values of $\langle v_i(t)\rangle$ by increasing F_{gen}. By the CD protocol, quantum nodes implement swaps by choosing two entangled links randomly; hence the set of virtual neighbours also grow randomly. If there is enough time (large T_2) and $v_i(t)$ is only constrained by M, there is a significant probability of connecting to each possible virtual neighbour. However, when there is little time, there is a significant probability that some virtual neighbours share multiple entangled links and others none. The maximum $\langle v_i(t)\rangle$ increases for increasing physical node degree with approximately the same ratios as for the results varying T_2.

Lastly, the performance of $\langle k_i(t)\rangle$ for various F_{gen} is similar: the same maximum at $q = 0$, decrease for $q > 0$ and convergence to 0 for $q = 1$. However, we note that the curves of $\langle k_i(t)\rangle$ are lower for decreasing F_{gen} because the quantum nodes discard more entangled links for having a too-low fidelity after more than M swaps.

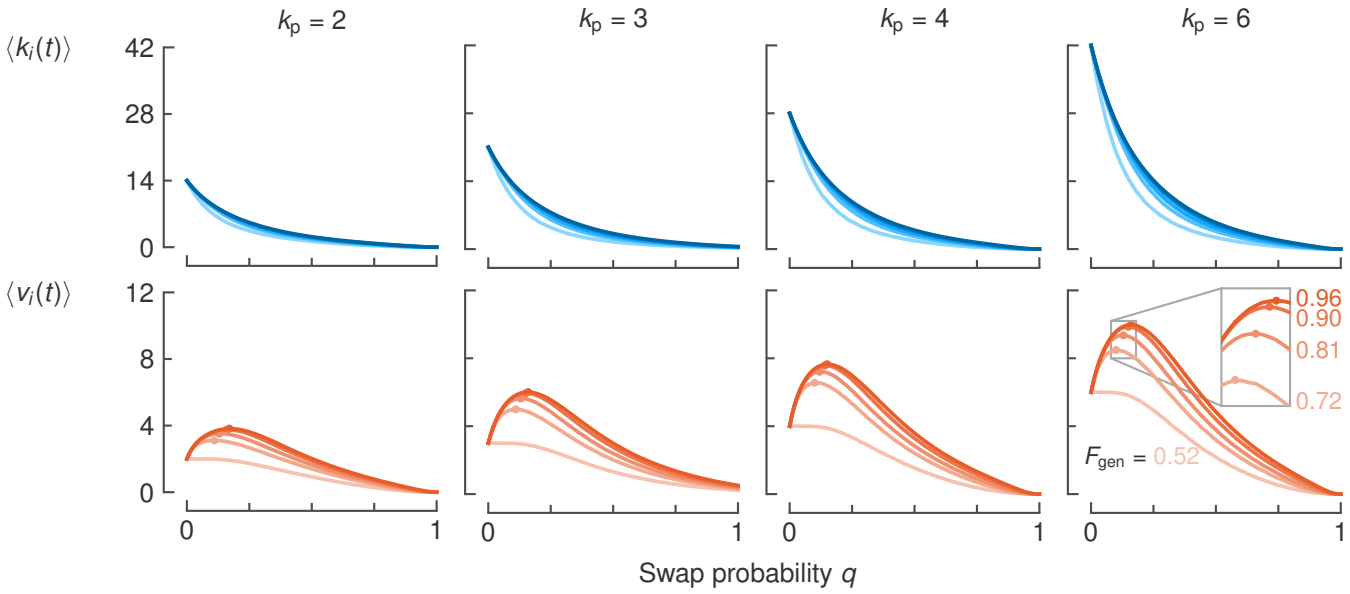

Figure 4.2: A better entanglement generation fidelity increases both the maximum virtual neighbourhood size and the optimal swap probability. The average virtual node degree $\langle k_i(t)\rangle$ is slightly higher for better entanglement generation fidelities F_{gen}, as links can be involved in more swaps before nodes discard them for being involved in too many swaps. The maximum average virtual neighbourhood size $\langle v_i(t)\rangle$ increases for higher F_{gen}, as links can be involved in more swaps before nodes discard them, meaning that nodes have a larger set of potential virtual neighbours. The optimal swap probability q increases for increasing F_{gen}. We vary $F_{\text{gen}} = 0.52, 0.72, 0.81, 0.90, 0.96$ (associated maximum swap distance $M = 0, 1, 2, 4, 7$; coloured from light to dark) while using the baseline parameters $p = 1$, $p_s = 1$, $T_2 = 100$, $t_{\text{cut}} = 7$, and $F_{\text{min}} = \frac{1}{2}$. The distributions of the virtual metrics are shown in Figure C.3.

4.2.4 Optimising over coherence time and entanglement generation fidelity

Having seen how the maximum value of $\langle v_i(t) \rangle$ depends on the coherence time T_2 and the fidelity of generated entangled links F_{gen}, we now optimise the CD protocol over varying combinations of T_2 and F_{gen} (Figure 4.3). Specifically, we set $M = 4$ and $F_{\text{min}} = \frac{1}{2}$ and vary both T_2 and F_{gen} by setting t_{cut} to the maximum value satisfying Equation (2.4). For example, $t_{\text{cut}} = 1$ satisfies Equation (2.4) for $M = 4$, $F_{\text{min}} = \frac{1}{2}$, $T_2 = 14$ and $F_{\text{gen}} = 0.9$, and $t_{\text{cut}} = 15$ for $M = 4$, $F_{\text{min}} = \frac{1}{2}$, $T_2 = 70$ and $F_{\text{gen}} = 1$ (see Figure C.2 for details). We then extract the maximum $\langle v_i(t) \rangle$ over the range of swap probabilities q. The maximum value of $\langle v_i(t) \rangle$ grows (relatively) linear with both T_2 and F_{gen} (more details in Figure C.4). Probabilistic swaps ($p_{\text{s}} = \frac{1}{2}$) do not meaningfully change the behaviour of the maximum $\langle v_i(t) \rangle$ compared to deterministic swaps, except for lower overall values.

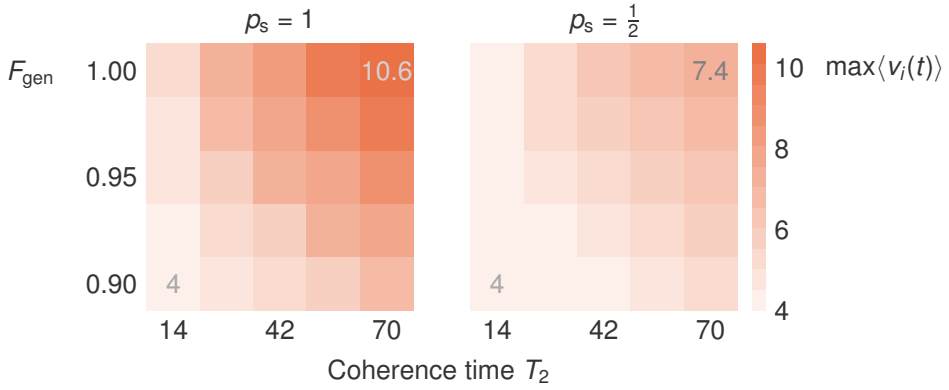

Figure 4.3: The maximum virtual neighbourhood size increases relatively linearly with coherence time and entanglement generation fidelity in an infinite chain. We vary $T_2 = 14, 28, 42, 56, 70$ and $F_{\text{gen}} = 0.90, 0.925, 0.95, 0.975, 1$ and associate the maximum cutoff times t_{cut} satisfying Equation (2.4) while using the baseline parameters $p = 1$, $p_{\text{s}} = 1$, $M = 4$ and $F_{\text{min}} = \frac{1}{2}$. We average over 2500 time slots. See Figure C.2 for details on t_{cut} and Figure C.4 for details on linearity and probability distributions.

4.2.5 Probabilistic swaps

Entanglement swaps that succeed probabilistically decrease the maximum value of the virtual neighbourhood size $\langle v_i(t)\rangle$ compared to deterministically succeeding swaps (Figure 4.4). When swaps fail, nodes consume the initial entangled links but yield no resulting state, meaning that nodes lose virtual neighbours (and entangled links). At the same time, the optimal swap probability q decreases for lower swap success probabilities p_s. If a swap can fail, it would be beneficial not to swap too often, as there is a risk of losing the initial links resulting in lower $\langle v_i(t)\rangle$. The maximum $\langle v_i(t)\rangle$ increases for increasing physical node degree with (approximately) the same ratios as for the results varying T_2. The virtual node degree $\langle k_i(t)\rangle$ behaves similarly for deterministic and probabilistic swaps, just worse.

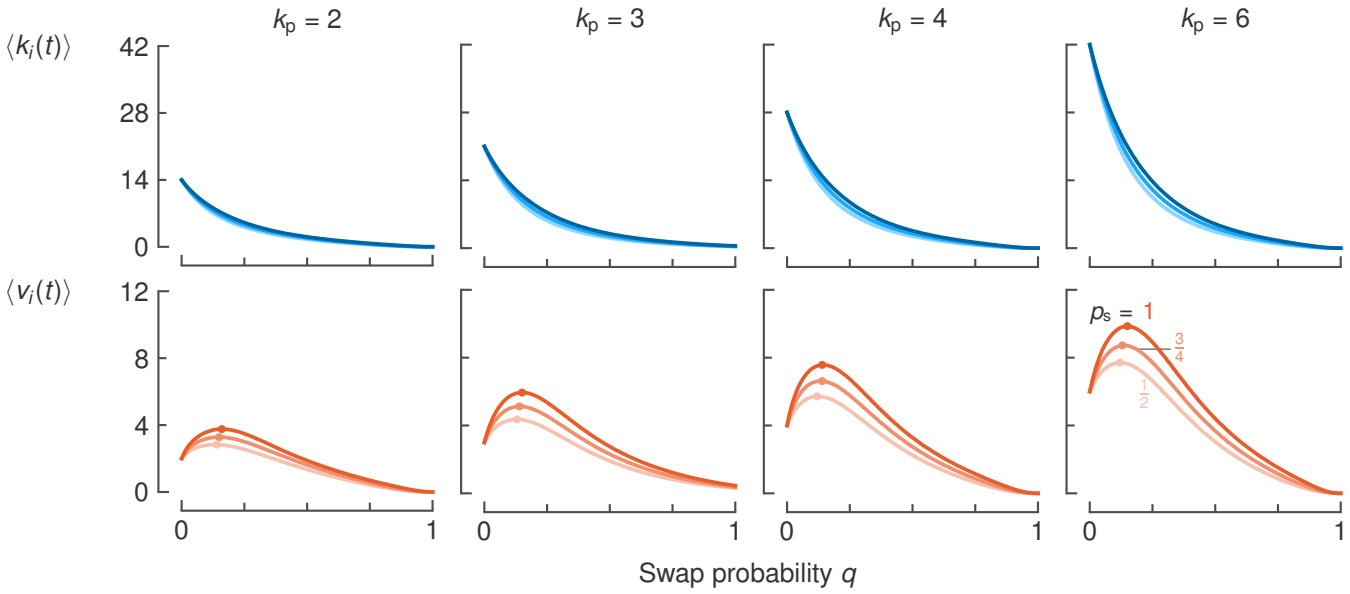

Figure 4.4: Probabilistic swaps decrease the maximum virtual neighbourhood size and the optimal swap probability. The average virtual node degree $\langle k_i(t)\rangle$ decreases slightly faster as a function of the swap attempt probability q for decreasing probabilities of successfully implementing a swap p_s. The maximum value of the average virtual neighbourhood size $\langle v_i(t)\rangle$ and the associated optimal q both decrease for lower p_s: it is beneficial to attempt fewer swaps when swaps can fail, destroying the initial links in case of failure. We vary $p_s = \frac{1}{2}, \frac{3}{4}, 1$ (coloured from light to dark) while using the baseline parameters $p = 1$, $T_2 = 100$, $t_{cut} = 7$, $F_{gen} = 0.9$, $M = 4$, and $F_{min} = \frac{1}{2}$. The probability distributions of the virtual metrics are shown in Figure C.5.

4.2.6 Probabilistic entanglement generation

Compared to deterministic entanglement generation, probabilistic generation with $p = \frac{1}{2}$ roughly halves the optimal swap probability q associated with the maximum average virtual neighbourhood size $\langle v_i(t)\rangle$ (Figure 4.5). When the entanglement generation probability halves, it takes, on average, twice as long to generate an entangled link. To replicate the behaviour for deterministic generation, we should halve the swap rate accordingly; that is, halve the swap attempt probability. If we assume that the protocols generating entanglement with probability $p = \frac{1}{2}, 1$ are devised on the same underlying (intrinsically probabilistic) entanglement generation rate, we correspondingly double the coherence time T_2 (and the cutoff time t_{cut}) when we decrease $p = 1$ to $p = \frac{1}{2}$.

Subsequently, probabilistic entanglement generation only decreases the maximum value of $\langle v_i(t)\rangle$ slightly in both infinite chains and infinite square grids. Specifically, for the infinite square grid, the optimal q decreases from approximately 0.14 ($\langle v_i(t)\rangle \approx 7.6$) for deterministic entanglement generation to 0.08 ($\langle v_i(t)\rangle \approx 7.3$) when entanglement is successfully generated with a probability $p = \frac{1}{2}$. Then, halving the entanglement generation probability again results in another halving of the optimal swap probability: $q \approx 0.04$ and $\langle v_i(t)\rangle \approx 7.1$ when $p = \frac{1}{4}$. The results are similar for the infinite chain where the optimal swap probability $q = 0.16$ ($\langle v_i(t)\rangle = 3.8$) for $p = 1$ decreases to $q = 0.1$ ($\langle v_i(t)\rangle = 3.6$) for $p = \frac{1}{2}$ and $q = 0.05$ ($\langle v_i(t)\rangle = 3.5$) for $p = \frac{1}{4}$. Lastly, generating fewer entangled links, on average, also results in a lower average virtual node degree $\langle k_i(t)\rangle$.

We note that $\langle k_i(t)\rangle, \langle v_i(t)\rangle$ do not converge to zero if nodes always swap ($q = 1$) when entanglement is generated probabilistically. The reasoning is similar to the nonzero virtual metrics for the honeycomb lattice ($k_\mathrm{p} = 3$): if a node has an odd number of entangled links, there is a probability that an entangled link will not be involved in too many swaps even if nodes always swap. With probabilistic generation, it can happen that nodes generate an odd number of entangled links, resulting in nonzero $k_i(t), v_i(t)$.

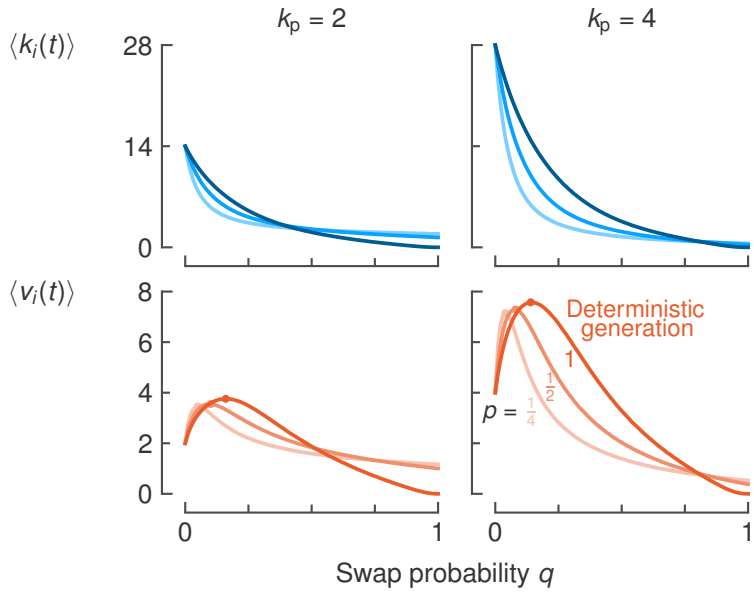

Figure 4.5: Halving the success probability of entanglement generation roughly halves the optimal swap probability. The average virtual node degree $\langle k_i(t)\rangle$ decreases faster as a function of the swap probability q for decreasing probabilities of successfully generating entanglement p. Furthermore, probabilistic generation results in non-zero virtual metrics as $q = 1$, as nodes do not always generate an even number of entangled links. The maximum average virtual neighbourhood size $\langle v_i(t)\rangle$ decreases slightly for lower p. Furthermore, the optimal q associated with the maximum $\langle v_i(t)\rangle$ decreases proportionally to the decrease in p: when p halves, it takes, on average, twice as long to generate an entangled link. We vary $p = \frac{1}{4}, \frac{1}{2}, 1$ (coloured from light to dark) while correspondingly increasing the $T_2 = 100, 200, 400$ ($t_\mathrm{cut} = 7, 14, 28$) and using the baseline parameters $p_\mathrm{s} = 1$, $F_\mathrm{gen} = 0.9$, $M = 4$, and $F_\mathrm{min} = \frac{1}{2}$. The probability distributions of the virtual metrics are shown in Figure C.6.

4.3 Finite regular networks

Infinite quantum networks offer a convenient platform to analyse protocol performance, but such networks might not always be feasible. We now analyse the influence of boundaries in finite networks on the performance of the CD protocol. As nodes closer to the network's boundary perform differently than nodes in the network's centre, we can not average the virtual metrics over all nodes. In Section 4.3.1, we discuss the virtual metrics of nodes in a finite chain and, in Section 4.3.2, those in a finite square grid.

4.3.1 Finite chain

In a finite chain, the virtual metrics reflect that the edge nodes only have one physical neighbour (Figure 4.6): they are half that of the other nodes when nodes do not attempt swaps. Furthermore, by the design of the CD protocol, the edge nodes can not implement swaps, so they do not consume links as $q > 0$. They only lose entangled links when they are involved in too many swaps or if they are too old. Therefore, the average virtual node degree $\langle k_i(t)\rangle$ only decreases slowly as q increases. Initially, the average number of virtual neighbours $\langle v_i(t)\rangle$ increases as other nodes implement swaps but then flattens for a wide range of q. Both $\langle k_i(t)\rangle$ and $\langle v_i(t)\rangle$ converge quickly to zero if nodes always attempt swaps, as the nodes discard all the low-fidelity links.

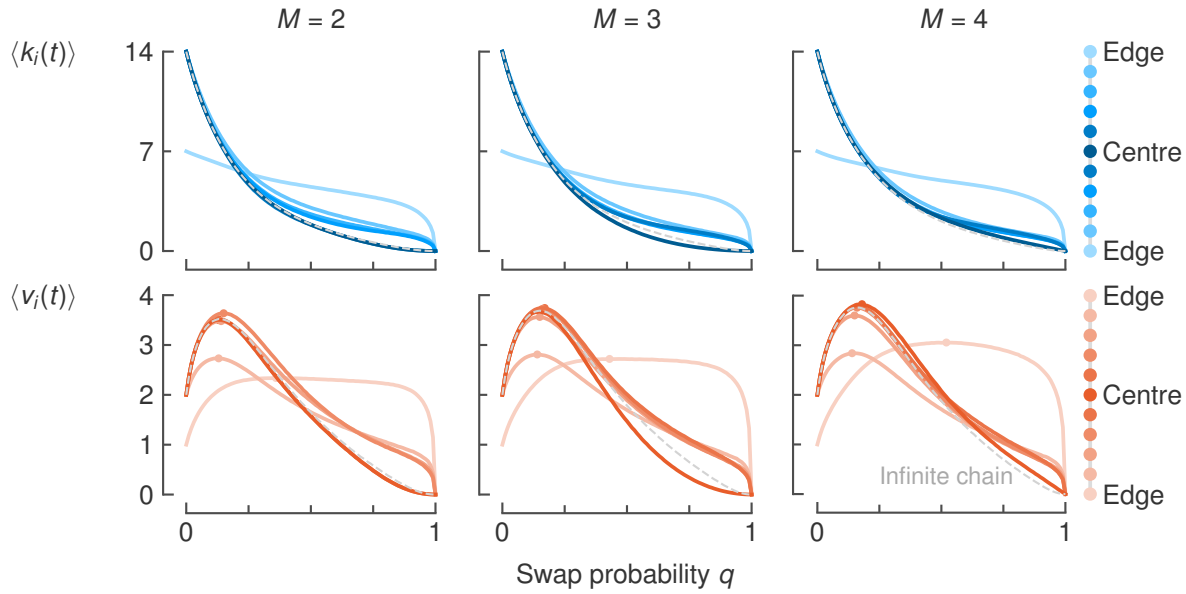

Figure 4.6: Virtual metrics depend strongly on a node's location in a finite chain. Central nodes have similar virtual metrics to nodes in an infinite chain. Edge nodes have qualitatively different virtual metrics because they have a single physical neighbour. By the swap routine implementation, they can not attempt swaps. Hence, the edge nodes only lose links when they cut off old links or discard them for being involved in too many swaps. The virtual metrics of some off-centre nodes converge slower to zero as $q \to 1$. As nodes attempt swaps between two links stored in differently-oriented qubits, off-centre nodes close to the edge will never store links that are swapped too many times in the direction of the edge (as the edge nodes can not attempts swaps). They are, therefore, more likely to store entangled links at the end of the CD protocol when nodes attempt many swaps. Additionally, nodes closer to the edge have lower virtual metrics because they have fewer potential virtual neighbours. We use the baseline parameters $p = 1, p_s = 1$, $T_2 = 100$, $t_{cut} = 7$, $F_{gen} = 0.9$, $M = 4$, and $F_{min} = \frac{1}{2}$. We average over 500 000 time slots and equivalent nodes (that is, nodes with the same distance to the centre). The probability distributions of the virtual metrics are shown in Figure C.7.

The performance of nodes close to the centre is quite similar to the central node (and the performance of nodes in an infinite chain), although the edges have an influence. Specifically, for q larger than the optimal swap probability, some off-centre nodes converge to $\langle k_i(t)\rangle, \langle v_i(t)\rangle = 0$ slower than the central node. We explain this with an example: suppose we look at the node to the left of the edge node. We recall that nodes attempt swaps between two links stored in qubits with different orientations and that the orientations are preserved during a swap. When this node does not attempt to swap, it will store a link in a left- and right-oriented qubit (recall that generation is deterministic). By the swap implementation, these entangled links can only grow in the direction of their qubit's orientation. The right-oriented link is already connected to the edge node and can not grow longer, so the node will not discard it for being involved in too many swaps. The left-oriented link, however, can grow towards and beyond the centre and be involved in more than M swaps. The node will then discard the link, leaving the node with a single right-oriented link; that is, nonzero $\langle k_i(t)\rangle, \langle v_i(t)\rangle$. Such a scenario, where most nodes attempt swaps but one does (or some do) not, is likely to happen for a large but not deterministic swap attempt probability. This is reflected by the slower convergence of the virtual metrics for off-centre nodes as $q \to 1$. When $q = 1$, all nodes attempt swaps, meaning that all links will be involved in too many swaps resulting in $\langle k_i(t)\rangle, \langle v_i(t)\rangle = 0$.

Comparing different off-centre nodes, we see that the $\langle v_i(t)\rangle$ behaviour of nodes close to the edge is qualitatively similar to those closer to the centre but quantitatively worse. Nodes close to the edge have fewer potential virtual neighbours in the direction of the edge, and they, therefore, have fewer virtual neighbours; that is, a smaller $\langle v_i(t)\rangle$. Conversely, $\langle k_i(t)\rangle$ of nodes close to the edge is quantitatively similar to nodes close to the centre. Although the number of potential neighbours is smaller for nodes closer to the edge, all nodes lose a similar number of entangled links due to swaps and discarding them for being involved in too many swaps. In fact, the average virtual node degree $\langle k_i(t)\rangle$ of nodes close to the edge can perform slightly better than that of nodes close to the centre, as entangled links can only be involved in too many swaps in one direction when the node is close to the edge.

In an 11-node chain with a maximum swap distance $M = 4$, the central node will never store entangled links that are involved in more than M swaps. Therefore, it will not discard links for being involved in too many swaps, so the virtual metrics are (slightly) better than those of nodes in the infinite chain. Contrary to the off-centre nodes, the central node is likely to have an equal number of entangled links oriented to either side (recall that generation and swaps are deterministic). For a significant swap probability, the central node is more likely to swap all the links in its memory than off-centre nodes. This results in $\langle k_i(t)\rangle, \langle v_i(t)\rangle$ of the central node converging to zero faster as $q \to 1$ than off-centre nodes (and similar to the virtual metrics of nodes in the infinite chain).

For $M = 2, 3$ in an 11-node chain, the maximum $\langle v_i(t)\rangle$ of the central node is approximately the same as that of nodes in an infinite chain. The two nodes next to the central node perform the same as the central node when $M = 2$ (overlapped with the central node in Figure 4.6). These nodes are unaware of the edges; they can store entangled links involved in too many swaps (before they discard them). Hence they perform similarly to the nodes in the infinite chain. Some off-centre nodes have a more sizeable $\langle v_i(t)\rangle$ than nodes in an infinite chain. This happens because of a combination of circumstances.

First, these nodes can only grow (too) long links in one direction as they are relatively close to the edge, meaning that they discard fewer links for being involved in too many swaps. Secondly, the set of potential virtual neighbours is not too limited if they are not too close to the edge.

Finally, all nodes except the edge nodes have relatively similar optimal swap probabilities. In particular, the optimal swap probability is between approximately $q = 0.13$ and $q = 0.15$ for the 11-node chain with $M = 2$, between $q = 0.14$ and $q = 0.17$ when $M = 3$, and between $q = 0.14$ and 0.18 when $M = 4$. The optimal swap probabilities are higher when M is longer because the nodes can implement more swaps before they discard entangled links for being involved in too many swaps, resulting in more virtual neighbours. Generally, a larger maximum $\langle v_i(t)\rangle$ corresponds to a higher optimal q as more attempted swaps result in more virtual neighbours.

4.3.2 Finite square grid

For $k_\mathrm{p} = 4$, the virtual metrics of all nodes are qualitatively similar, and those of nodes with equal physical node degrees are quantitatively similar (Figure 4.7). For instance, in a finite square grid, quantum nodes close to the centre perform similarly as they all have four physical neighbours. Precisely, the average virtual node degree $\langle k_i(t)\rangle$ and virtual neighbourhood size $\langle v_i(t)\rangle$ overlap with those of nodes in an infinite grid. Closer to the edge, nodes with four physical neighbours have the same $\langle k_i(t)\rangle$ but slightly smaller $\langle v_i(t)\rangle$, as they have fewer potential virtual neighbours.

When nodes do not attempt swaps, the corner nodes have $\langle k_i(t)\rangle, \langle v_i(t)\rangle$ half that of the central nodes as they only have two physical neighbours. Converse to the edge nodes in the finite chain that could not attempt swaps because they had only one physical neighbour, the corner nodes can attempt entanglement swaps. Their performance is, therefore, qualitatively similar to the other nodes for $q > 0$. Furthermore, the maximum $\langle v_i(t)\rangle$ for nodes with two physical neighbours is roughly half the value of the maximum $\langle v_i(t)\rangle$ for nodes with four physical neighbours.

The side nodes (those on the boundary but not in the corner) have three physical neighbours and therefore have higher $\langle k_i(t)\rangle, \langle v_i(t)\rangle$ than the corner nodes. When nodes do not attempt swaps, their $\langle k_i(t)\rangle, \langle v_i(t)\rangle$ is $\frac{3}{4}$ that of the central node. Similarly, the maximum $\langle v_i(t)\rangle$ is also approximately $\frac{3}{4}$ that of the central node. We note that when nodes always swap, $\langle k_i(t)\rangle, \langle v_i(t)\rangle$ of these nodes converge to a nonzero value. Similar to the explanation for nodes in an infinite honeycomb lattice ($k_\mathrm{p} = 3$), there is a probability for nodes to be left with an entangled link after the CD protocol when always attempting swaps. This is due to the swap implementation with an odd number of generated links (generation is deterministic). The average virtual node degree $\langle k_i(t)\rangle$ of the different side nodes is approximately equal, but there is a (slight) difference in the average physical neighbourhood sizes $\langle v_i(t)\rangle$. Side nodes closer to a corner, e.g., the left-bottom corner, only have potential virtual neighbours in one quadrant, the top-right quadrant. A central side node has potential virtual neighbours in two quadrants. Hence they have more potential neighbours resulting in a larger $\langle v_i(t)\rangle$ than that of a node closer to the corner. The virtual node degree $\langle k_i(t)\rangle$ is not dependent on the number of potential virtual neighbours; it decreases when swaps are involved in (too many) swaps.

Lastly, we note that all nodes have similar optimal swap probabilities regardless of how many physical neighbours they have, ranging from approximately $q = 0.13$ to $q = 0.15$.

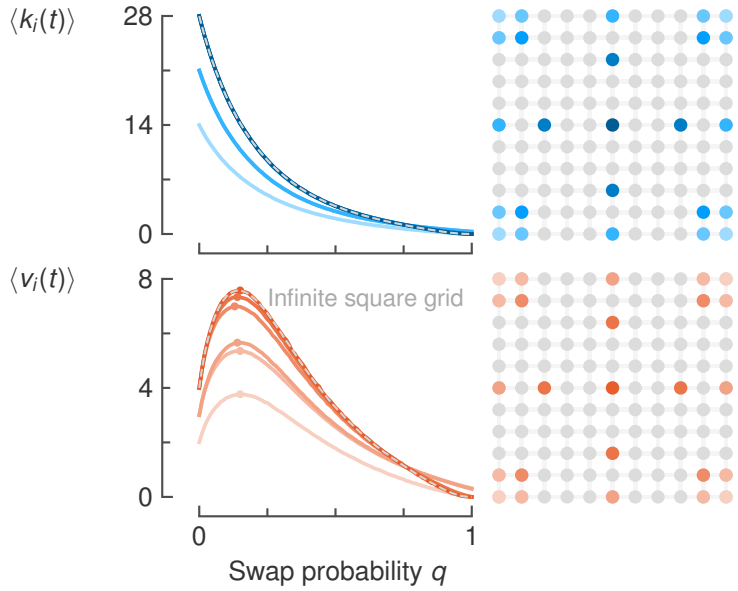

Figure 4.7: Virtual metrics of nodes in a finite square grid are qualitatively similar, with the physical node degrees driving quantitative differences. Quantum nodes with four physical neighbours have the same average virtual node degree $\langle k_i(t)\rangle$, and almost the same average virtual neighbourhood size $\langle v_i(t)\rangle$ – nodes close to the boundaries have fewer potential virtual neighbours resulting in smaller $\langle v_i(t)\rangle$. The virtual metrics of side nodes (nodes on the boundary but not in the corner) have similar behaviour but are smaller by roughly a factor $\frac{3}{4}$ as they only have three physical neighbours. As the corner nodes have two physical neighbours, their virtual metrics behave similarly but lower by roughly a factor $\frac{2}{4}$. We use the baseline parameters $p = 1, p_s = 1$, $T_2 = 100$, $t_{cut} = 7$, $F_{gen} = 0.9$, $M = 4$ and $F_{min} = \frac{1}{2}$. We average over 25 000 time slots and equivalent nodes (as indicated in the legend). The probability distributions of the virtual metrics are shown in Figure C.8.

4.4 Heuristics for designing regular-topology quantum networks

Having investigated various quantum networks with regular topologies, we now summarise our findings with heuristics. For instance, if a quantum network engineer wants to design a quantum network with regular topologies running a CD protocol, what are some general guidelines for designing such networks?

Firstly, increasing the values of network parameters, such as the coherence time and entanglement generation fidelity, shifts the optimal swap attempt probability q providing the maximum average virtual neighbourhood size $\langle v_i(t) \rangle$. The direction of the shift depends on the network parameter. In particular, increasing the coherence time T_2 decreases the optimal q. Conversely, increasing the entanglement generation probability F_{gen} increases the optimal q.

Furthermore, in the regime we investigated, the maximum value of $\langle v_i(t) \rangle$ increases relatively linearly with an increase in the coherence time T_2 and the entanglement generation probability F_{gen}.

Also, the maximum value of $\langle v_i(t) \rangle$ increases with the physical node degree of the quantum network. Increasing the network from $k_{\text{p}} = 2$ to $k_{\text{p}} = 3$ increases the maximum value more than the ratio of the physical node degrees. However, the returns on increasing the physical node are diminishing. Specifically, increasing the physical node degree from $k_{\text{p}} = 2$ to $k_{\text{p}} = 4$ roughly doubles the maximum $\langle v_i(t) \rangle$, while increasing $k_{\text{p}} = 2$ to $k_{\text{p}} = 6$ results in less than tripling $\langle v_i(t) \rangle$.

Compared to deterministic implementations, probabilistic entanglement generation and probabilistic entanglement swaps influence the average virtual neighbourhood size differently. In particular, probabilistic entanglement generation primarily influences the optimal swap probability q: halving the entanglement generation probability p roughly halves the optimal q. The associated maximum $\langle v_i(t) \rangle$ only decreases marginally. In contrast, probabilistic swaps decrease the maximum $\langle v_i(t) \rangle$ considerably while only slightly decreasing the optimal q.

Lastly, the performance of quantum nodes in finite quantum networks depends on the network topology. For example, the virtual metrics of quantum nodes in a finite chain strongly depend on the combination of the node's distance to the chain's edge and the maximum swap distance. Conversely, all nodes have qualitatively similar virtual metrics in a finite square grid, with the physical node degree of a node primarily driving quantitative differences.

5

Conclusions

In this chapter, we summarise our investigation, reflect on its results and offer potential steps for research beyond this thesis.

To summarise, we have investigated the performance of a protocol that continuously distributes entanglement in quantum networks with regular topologies. Motivated by the need for entangled links between nodes to implement nonlocal gates in distributed quantum computing, we have used two figures of merit that encapsulate the goals of such networks and the time-dependent nature of the fidelity of entangled links. Firstly, the virtual node degree manifests the need for many entangled links to implement many nonlocal operations. Secondly, the virtual neighbourhood size represents the need to share entangled links with many remote quantum nodes to grow the number of qubits that are available for computations.

Employing a protocol that continuously delivers entanglement, the distribution of entangled links in the quantum network evolves due to the creation of links by entanglement generation and entanglement swaps, and the removal of low-fidelity links. The quantum nodes can optimise the virtual node degree and virtual neighbourhood size as a function of the probability of attempting swaps. We have found that the network parameters, such as the coherence time and entanglement generation fidelity, influence the optimal swap attempt probability associated with the maximum average virtual neighbourhood size differently. Furthermore, the network topology, characterised by the number of physical neighbours per node in regular networks, and the presence or absence of boundaries in the network strongly affect the network performance. For example, increasing the number of physical neighbours per node from two to three increases the maximum average virtual neighbourhood size by more than the ratio of physical neighbours. However, the returns diminish when further increasing the number of physical neighbours.

Further research could implement a process to consume entangled links (for example, as implemented by Iñesta and Wehner [20]). In quantum networks for distributed quantum computing, quantum nodes consume entangled links to implement nonlocal operations in quantum computations. A process consuming entangled links then emulates implementing nonlocal operations. It would be interesting to see the effect of consuming entangled links on the virtual metrics and the corresponding optimal swap attempt probability.

Furthermore, nodes could implement a more elaborate entanglement swap routine to grow the virtual neighbourhood size more efficiently than the current routine that randomly pairs two entangled links. For example, a routine could guide specific combinations of entangled links resulting in more virtual neighbours.

By increasing computational resources, we can investigate more extensive quantum networks or nodes that generate entangled links with better fidelities and longer coherence times. For example, it would be interesting to see how the relatively linear relationship between coherence time, entanglement generation fidelity and the maximum virtual neighbourhood size develops with longer coherence times. Moreover, for more accurate analysis, the simulation could track the fidelities of the entangled links instead of setting threshold values such as the cutoff time and maximum swap distance. Then, an elaborate swap routine could match entangled links with certain fidelities in a way that optimises the virtual metrics.

Lastly, analytical results – more precise than the bounds presented here – would be a valuable addition to the heuristics derived from the numerical results in this work and can allow for a deeper understanding of the dynamics of entanglement distribution in quantum networks with regular topologies.

Bibliography

[1] John S Bell. "On the Einstein Podolsky Rosen paradox". In: *Physics Physique Fizika* 1.3 (1964), p. 195.

[2] Peter W Shor. "Algorithms for quantum computation: discrete logarithms and factoring". In: *Proceedings 35th annual symposium on foundations of computer science*. Ieee. 1994, pp. 124–134.

[3] Craig Gidney and Martin Ekerå. "How to factor 2048 bit RSA integers in 8 hours using 20 million noisy qubits". In: *Quantum* 5 (2021), p. 433.

[4] Isaac L Chuang, Neil Gershenfeld, and Mark Kubinec. "Experimental implementation of fast quantum searching". In: *Physical review letters* 80.15 (1998), p. 3408.

[5] Jonathan A Jones and Michele Mosca. "Implementation of a quantum algorithm on a nuclear magnetic resonance quantum computer". In: *The Journal of chemical physics* 109.5 (1998), pp. 1648–1653.

[6] Frank Arute, Kunal Arya, Ryan Babbush, Dave Bacon, Joseph C Bardin, Rami Barends, Rupak Biswas, Sergio Boixo, Fernando GSL Brandao, David A Buell, et al. "Quantum supremacy using a programmable superconducting processor". In: *Nature* 574.7779 (2019), pp. 505–510.

[7] Han-Sen Zhong, Hui Wang, Yu-Hao Deng, Ming-Cheng Chen, Li-Chao Peng, Yi-Han Luo, Jian Qin, Dian Wu, Xing Ding, Yi Hu, et al. "Quantum computational advantage using photons". In: *Science* 370.6523 (2020), pp. 1460–1463.

[8] Yulin Wu, Wan-Su Bao, Sirui Cao, Fusheng Chen, Ming-Cheng Chen, Xiawei Chen, Tung-Hsun Chung, Hui Deng, Yajie Du, Daojin Fan, et al. "Strong quantum computational advantage using a superconducting quantum processor". In: *Physical review letters* 127.18 (2021), p. 180501.

[9] Daniele Cuomo, Marcello Caleffi, and Angela Sara Cacciapuoti. "Towards a distributed quantum computing ecosystem". In: *IET Quantum Communication* 1.1 (2020), pp. 3–8.

[10] Daniel Gottesman and Isaac L Chuang. "Demonstrating the viability of universal quantum computation using teleportation and single-qubit operations". In: *Nature* 402.6760 (1999), pp. 390–393.

[11] Xinlan Zhou, Debbie W Leung, and Isaac L Chuang. "Methodology for quantum logic gate construction". In: *Physical Review A* 62.5 (2000), p. 052316.

[12] Jens Eisert, Kurt Jacobs, Polykarpos Papadopoulos, and Martin B Plenio. "Optimal local implementation of nonlocal quantum gates". In: *Physical Review A* 62.5 (2000), p. 052317.

[13] Anocha Yimsiriwattana and Samuel J Lomonaco Jr. "Distributed quantum computing: A distributed Shor algorithm". In: *Quantum Information and Computation II*. Vol. 5436. SPIE. 2004, pp. 360–372.

[14] Richard Jozsa. "Fidelity for mixed quantum states". In: *Journal of modern optics* 41.12 (1994), pp. 2315–2323.

[15] Wolfgang Dür and H-J Briegel. "Entanglement purification for quantum computation". In: *Physical review letters* 90.6 (2003), p. 067901.

[16] Liang Jiang, Jacob M Taylor, Anders S Sørensen, and Mikhail D Lukin. "Distributed quantum computation based on small quantum registers". In: *Physical Review A* 76.6 (2007), p. 062323.

[17] Naomi H Nickerson, Ying Li, and Simon C Benjamin. "Topological quantum computing with a very noisy network and local error rates approaching one percent". In: *Nature communications* 4.1 (2013), pp. 1–5.

[18] Naomi H Nickerson, Joseph F Fitzsimons, and Simon C Benjamin. "Freely scalable quantum technologies using cells of 5-to-50 qubits with very lossy and noisy photonic links". In: *Physical Review X* 4.4 (2014), p. 041041.

[19] Sébastian de Bone, Runsheng Ouyang, Kenneth Goodenough, and David Elkouss. "Protocols for creating and distilling multipartite GHZ states with Bell pairs". In: *IEEE Transactions on Quantum Engineering* 1 (2020), pp. 1–10.

[20] Álvaro G Iñesta and Stephanie Wehner. "Performance metrics for the continuous distribution of entanglement in multi-user quantum networks". In: *arXiv preprint arXiv:2307.01406* (2023).

[21] Kaushik Chakraborty, Filip Rozpedek, Axel Dahlberg, and Stephanie Wehner. "Distributed routing in a quantum internet". In: *arXiv preprint arXiv:1907.11630* (2019).

[22] Stefano Pirandola. "End-to-end capacities of a quantum communication network". In: *Communications Physics* 2.1 (2019), pp. 1–10.

[23] Gayane Vardoyan, Saikat Guha, Philippe Nain, and Don Towsley. "On the stochastic analysis of a quantum entanglement distribution switch". In: *IEEE Transactions on Quantum Engineering* 2 (2021), pp. 1–16.

[24] Álvaro G Iñesta, Gayane Vardoyan, Lara Scavuzzo, and Stephanie Wehner. "Optimal entanglement distribution policies in homogeneous repeater chains with cutoffs". In: *arXiv preprint arXiv:2207.06533* (2022).

[25] H Jeff Kimble. "The quantum internet". In: *Nature* 453.7198 (2008), pp. 1023–1030.

[26] Stephanie Wehner, David Elkouss, and Ronald Hanson. "Quantum internet: A vision for the road ahead". In: *Science* 362.6412 (2018), eaam9288.

[27] Wojciech Kozlowski, Stephanie Wehner, Rodney Van Meter, Bruno Rijsman, Angela Sara Cacciapuoti, Marcello Caleffi, and Shota Nagayama. *Architectural Principles for a Quantum Internet*. Internet-Draft draft-irtf-qirg-principles-11. Work in Progress. Internet Engineering Task Force, Aug. 2022. 46 pp. URL: `https://datatracker.ietf.org/doc/draft-irtf-qirg-principles/11/`.

[28] Charles H. Bennett and Gilles Brassard. "Quantum cryptography: Public key distribution and coin tossing". In: *Theoretical Computer Science* 560 (2014), pp. 7–11.

[29] Artur K. Ekert. "Quantum cryptography based on Bell's theorem". In: *Phys. Rev. Lett.* 67 (6 1991), pp. 661–663.

[30] Daniel Gottesman, Thomas Jennewein, and Sarah Croke. "Longer-baseline telescopes using quantum repeaters". In: *Physical review letters* 109.7 (2012), p. 070503.

[31] Richard Cleve and Harry Buhrman. "Substituting quantum entanglement for communication". In: *Physical Review A* 56.2 (1997), p. 1201.

[32] Lov K Grover. "Quantum telecomputation". In: *arXiv preprint quant-ph/9704012* (1997).

[33] Hannes Bernien, Bas Hensen, Wolfgang Pfaff, Gerwin Koolstra, Machiel S Blok, Lucio Robledo, Tim H Taminiau, Matthew Markham, Daniel J Twitchen, Lilian Childress, et al. "Heralded entanglement between solid-state qubits separated by three metres". In: *Nature* 497.7447 (2013), pp. 86–90.

[34] Bas Hensen, Hannes Bernien, Anaïs E Dréau, Andreas Reiserer, Norbert Kalb, Machiel S Blok, Just Ruitenberg, Raymond FL Vermeulen, Raymond N Schouten, Carlos Abellán, et al. "Loophole-free Bell inequality violation using electron spins separated by 1.3 kilometres". In: *Nature* 526.7575 (2015), pp. 682–686.

[35] David L Moehring, Peter Maunz, Steve Olmschenk, Kelly C Younge, Dzmitry N Matsukevich, L-M Duan, and Christopher Monroe. "Entanglement of single-atom quantum bits at a distance". In: *Nature* 449.7158 (2007), pp. 68–71.

[36] LJ Stephenson, DP Nadlinger, BC Nichol, S An, P Drmota, TG Ballance, K Thirumalai, JF Goodwin, DM Lucas, and CJ Ballance. "High-rate, high-fidelity entanglement of qubits across an elementary quantum network". In: *Physical review letters* 124.11 (2020), p. 110501.

[37] Wolfgang Tittel, Jürgen Brendel, Hugo Zbinden, and Nicolas Gisin. "Violation of Bell inequalities by photons more than 10 km apart". In: *Physical review letters* 81.17 (1998), p. 3563.

[38] Carlos Cabrillo, J Ignacio Cirac, Pablo Garcia-Fernandez, and Peter Zoller. "Creation of entangled states of distant atoms by interference". In: *Physical Review A* 59.2 (1999), p. 1025.

[39] L-M Duan, Mikhail D Lukin, J Ignacio Cirac, and Peter Zoller. "Long-distance quantum communication with atomic ensembles and linear optics". In: *Nature* 414.6862 (2001), pp. 413–418.

[40] Sean D Barrett and Pieter Kok. "Efficient high-fidelity quantum computation using matter qubits and linear optics". In: *Physical Review A* 71.6 (2005), p. 060310.

[41] Michael A Nielsen and Isaac Chuang. *Quantum computation and quantum information*. 2002.

[42] Reinhard F Werner. "Quantum states with Einstein-Podolsky-Rosen correlations admitting a hidden-variable model". In: *Physical Review A* 40.8 (1989), p. 4277.

[43] Marek Zukowski, Anton Zeilinger, Michael A Horne, and Aarthur K Ekert. """ Event-ready-detectors" Bell experiment via entanglement swapping." In: *Physical Review Letters* 71.26 (1993).

[44] William J Munro, Koji Azuma, Kiyoshi Tamaki, and Kae Nemoto. "Inside quantum repeaters". In: *IEEE Journal of Selected Topics in Quantum Electronics* 21.3 (2015), pp. 78–90.

[45] OA Collins, SD Jenkins, A Kuzmich, and TAB Kennedy. "Multiplexed memory-insensitive quantum repeaters". In: *Physical review letters* 98.6 (2007), p. 060502.

[46] Filip Rozpedek, Kenneth Goodenough, Jeremy Ribeiro, Norbert Kalb, V Caprara Vivoli, Andreas Reiserer, Ronald Hanson, Stephanie Wehner, and David Elkouss. "Parameter regimes for a single sequential quantum repeater". In: *Quantum Science and Technology* 3.3 (2018), p. 034002.

[47] Sumeet Khatri, Corey T Matyas, Aliza U Siddiqui, and Jonathan P Dowling. "Practical figures of merit and thresholds for entanglement distribution in quantum networks". In: *Physical Review Research* 1.2 (2019), p. 023032.

[48] Charles H Bennett, Gilles Brassard, Sandu Popescu, Benjamin Schumacher, John A Smolin, and William K Wootters. "Purification of noisy entanglement and faithful teleportation via noisy channels". In: *Physical review letters* 76.5 (1996), p. 722.

[49] David Deutsch, Artur Ekert, Richard Jozsa, Chiara Macchiavello, Sandu Popescu, and Anna Sanpera. "Quantum privacy amplification and the security of quantum cryptography over noisy channels". In: *Physical review letters* 77.13 (1996), p. 2818.

[50] H-J Briegel, Wolfgang Dür, Juan I Cirac, and Peter Zoller. "Quantum repeaters: the role of imperfect local operations in quantum communication". In: *Physical Review Letters* 81.26 (1998), p. 5932.

[51] Samuraí Brito, Askery Canabarro, Rafael Chaves, and Daniel Cavalcanti. "Statistical properties of the quantum internet". In: *Physical Review Letters* 124.21 (2020), p. 210501.

[52] Antonio Acín, J Ignacio Cirac, and Maciej Lewenstein. "Entanglement percolation in quantum networks". In: *Nature Physics* 3.4 (2007), pp. 256–259.

[53] Koji Azuma, Stefan Bäuml, Tim Coopmans, David Elkouss, and Boxi Li. "Tools for quantum network design". In: *AVS Quantum Science* 3.1 (2021), p. 014101.

[54] Evgeny Shchukin, Ferdinand Schmidt, and Peter van Loock. "Waiting time in quantum repeaters with probabilistic entanglement swapping". In: *Physical Review A* 100.3 (2019), p. 032322.

[55] Evgeny Shchukin and Peter van Loock. "Optimal entanglement swapping in quantum repeaters". In: *Physical Review Letters* 128.15 (2022), p. 150502.

[56] Gayane Vardoyan, Saikat Guha, Philippe Nain, and Don Towsley. "On the stochastic analysis of a quantum entanglement switch". In: *ACM SIGMETRICS Performance Evaluation Review* 47.2 (2019), pp. 27–29.

[57] Gayane Vardoyan, Saikat Guha, Philippe Nain, and Don Towsley. "On the capacity region of bipartite and tripartite entanglement switching". In: *ACM SIGMETRICS Performance Evaluation Review* 48.3 (2021), pp. 45–50.

[58] Richard J Hughes, Jane E Nordholt, Derek Derkacs, and Charles G Peterson. "Practical free-space quantum key distribution over 10 km in daylight and at night". In: *New journal of physics* 4.1 (2002), p. 43.

[59] Juan Yin, Yuan Cao, Yu-Huai Li, Sheng-Kai Liao, Liang Zhang, Ji-Gang Ren, Wen-Qi Cai, Wei-Yue Liu, Bo Li, Hui Dai, et al. "Satellite-based entanglement distribution over 1200 kilometers". In: *Science* 356.6343 (2017), pp. 1140–1144.

[60] Stephan Ritter, Christian Nölleke, Carolin Hahn, Andreas Reiserer, Andreas Neuzner, Manuel Uphoff, Martin Mücke, Eden Figueroa, Joerg Bochmann, and Gerhard Rempe. "An elementary quantum network of single atoms in optical cavities". In: *Nature* 484.7393 (2012), pp. 195–200.

[61] Julian Hofmann, Michael Krug, Norbert Ortegel, Lea Gérard, Markus Weber, Wenjamin Rosenfeld, and Harald Weinfurter. "Heralded entanglement between widely separated atoms". In: *Science* 337.6090 (2012), pp. 72–75.

[62] Aymeric Delteil, Zhe Sun, Wei-bo Gao, Emre Togan, Stefan Faelt, and Ataç Imamoğlu. "Generation of heralded entanglement between distant hole spins". In: *Nature Physics* 12.3 (2016), pp. 218–223.

[63] Robert Stockill, MJ Stanley, Lukas Huthmacher, Edmund Clarke, Maxim Hugues, AJ Miller, Clemens Matthiesen, Claire Le Gall, and Mete Atatüre. "Phase-tuned entangled state generation between distant spin qubits". In: *Physical review letters* 119.1 (2017), p. 010503.

[64] Dario Lago-Rivera, Samuele Grandi, Jelena V Rakonjac, Alessandro Seri, and Hugues de Riedmatten. "Telecom-heralded entanglement between multimode solid-state quantum memories". In: *Nature* 594.7861 (2021), pp. 37–40.

[65] Andreas Reiserer, Norbert Kalb, Machiel S Blok, Koen JM van Bemmelen, Tim H Taminiau, Ronald Hanson, Daniel J Twitchen, and Matthew Markham. "Robust quantum-network memory using decoherence-protected subspaces of nuclear spins". In: *Physical Review X* 6.2 (2016), p. 021040.

[66] Conor E Bradley, Joe Randall, Mohamed H Abobeih, RC Berrevoets, MJ Degen, Michiel A Bakker, Matthew Markham, DJ Twitchen, and Tim H Taminiau. "A ten-qubit solid-state spin register with quantum memory up to one minute". In: *Physical Review X* 9.3 (2019), p. 031045.

[67] CE Bradley, SW de Bone, PFW Möller, S Baier, MJ Degen, SJH Loenen, HP Bartling, M Markham, DJ Twitchen, R Hanson, et al. "Robust quantum-network memory based on

spin qubits in isotopically engineered diamond". In: *npj Quantum Information* 8.1 (2022), pp. 1–9.

[68] Ismail Volkan Inlek, Clayton Crocker, Martin Lichtman, Ksenia Sosnova, and Christopher Monroe. "Multispecies trapped-ion node for quantum networking". In: *Physical review letters* 118.25 (2017), p. 250502.

[69] P Drmota, D Main, DP Nadlinger, BC Nichol, MA Weber, EM Ainley, A Agrawal, R Srinivas, G Araneda, CJ Ballance, et al. "Robust Quantum Memory in a Trapped-Ion Quantum Network Node". In: *arXiv preprint arXiv:2210.11447* (2022).

[70] C Monroe, R Raussendorf, A Ruthven, KR Brown, P Maunz, L-M Duan, and J Kim. "Large-scale modular quantum-computer architecture with atomic memory and photonic interconnects". In: *Physical Review A* 89.2 (2014), p. 022317.

[71] Norbert Kalb, Andreas A Reiserer, Peter C Humphreys, Jacob JW Bakermans, Sten J Kamerling, Naomi H Nickerson, Simon C Benjamin, Daniel J Twitchen, Matthew Markham, and Ronald Hanson. "Entanglement distillation between solid-state quantum network nodes". In: *Science* 356.6341 (2017), pp. 928–932.

[72] Peter C Humphreys, Norbert Kalb, Jaco PJ Morits, Raymond N Schouten, Raymond FL Vermeulen, Daniel J Twitchen, Matthew Markham, and Ronald Hanson. "Deterministic delivery of remote entanglement on a quantum network". In: *Nature* 558.7709 (2018), pp. 268–273.

[73] Matteo Pompili, Sophie LN Hermans, Simon Baier, Hans KC Beukers, Peter C Humphreys, Raymond N Schouten, Raymond FL Vermeulen, Marijn J Tiggelman, Laura dos Santos Martins, Bas Dirkse, et al. "Realization of a multinode quantum network of remote solid-state qubits". In: *Science* 372.6539 (2021), pp. 259–264.

[74] SLN Hermans, M Pompili, HKC Beukers, S Baier, J Borregaard, and R Hanson. "Qubit teleportation between non-neighbouring nodes in a quantum network". In: *Nature* 605.7911 (2022), pp. 663–668.

[75] Severin Daiss, Stefan Langenfeld, Stephan Welte, Emanuele Distante, Philip Thomas, Lukas Hartung, Olivier Morin, and Gerhard Rempe. "A quantum-logic gate between distant quantum-network modules". In: *Science* 371.6529 (2021), pp. 614–617.

[76] Wolfgang Pfaff, Bas J Hensen, Hannes Bernien, Suzanne B van Dam, Machiel S Blok, Tim H Taminiau, Marijn J Tiggelman, Raymond N Schouten, Matthew Markham, Daniel J Twitchen, et al. "Unconditional quantum teleportation between distant solid-state quantum bits". In: *Science* 345.6196 (2014), pp. 532–535.

[77] Raustin Reyes, Takaya Nakazato, Nobuaki Imaike, Kazuyasu Matsuda, Kazuya Tsurumoto, Yuhei Sekiguchi, and Hideo Kosaka. "Complete Bell state measurement of diamond nuclear spins under a complete spatial symmetry at zero magnetic field". In: *Applied Physics Letters* 120.19 (2022), p. 194002.

[78] MD Barrett, J Chiaverini, T Schaetz, J Britton, WM Itano, JD Jost, E Knill, C Langer, D Leibfried, R Ozeri, et al. "Deterministic quantum teleportation of atomic qubits". In: *Nature* 429.6993 (2004), pp. 737–739.

[79] Mark Riebe, H Häffner, CF Roos, W Hänsel, J Benhelm, GPT Lancaster, TW Körber, C Becher, F Schmidt-Kaler, DFV James, et al. "Deterministic quantum teleportation with atoms". In: *Nature* 429.6993 (2004), pp. 734–737.

[80] Mihir Pant, Don Towsley, Dirk Englund, and Saikat Guha. "Percolation thresholds for photonic quantum computing". In: *Nature communications* 10.1 (2019), pp. 1–11.

[81] Robert Raussendorf and Hans J Briegel. "A one-way quantum computer". In: *Physical review letters* 86.22 (2001), p. 5188.

[82] Daniel E Browne and Terry Rudolph. "Resource-efficient linear optical quantum computation". In: *Physical Review Letters* 95.1 (2005), p. 010501.

[83] Sara Bartolucci, Patrick Birchall, Hector Bombin, Hugo Cable, Chris Dawson, Mercedes Gimeno-Segovia, Eric Johnston, Konrad Kieling, Naomi Nickerson, Mihir Pant, et al. "Fusion-based quantum computation". In: *arXiv preprint arXiv:2101.09310* (2021).

[84] Hector Bombin, Isaac H Kim, Daniel Litinski, Naomi Nickerson, Mihir Pant, Fernando Pastawski, Sam Roberts, and Terry Rudolph. "Interleaving: Modular architectures for fault-tolerant photonic quantum computing". In: *arXiv preprint arXiv:2103.08612* (2021).

[85] John Calsamiglia and Norbert Lütkenhaus. "Maximum efficiency of a linear-optical Bell-state analyzer". In: *Applied Physics B* 72.1 (2001), pp. 67–71.

[86] Warren P Grice. "Arbitrarily complete Bell-state measurement using only linear optical elements". In: *Physical Review A* 84.4 (2011), p. 042331.

[87] Fabian Ewert and Peter van Loock. "3/4-efficient Bell measurement with passive linear optics and unentangled ancillae". In: *Physical review letters* 113.14 (2014), p. 140403.

[88] Matthias J Bayerbach, Simone E D'Aurelio, Peter van Loock, and Stefanie Barz. "Bell-state measurement exceeding 50% success probability with linear optics". In: *arXiv preprint arXiv:2208.02271* (2022).

[89] Tim Coopmans, Robert Knegjens, Axel Dahlberg, David Maier, Loek Nijsten, Julio de Oliveira Filho, Martijn Papendrecht, Julian Rabbie, Filip Rozpedek, Matthew Skrzypczyk, et al. "Netsquid, a network simulator for quantum information using discrete events". In: *Communications Physics* 4.1 (2021), pp. 1–15.

[90] Rodney Van Meter, Thaddeus D Ladd, William J Munro, and Kae Nemoto. "System design for a long-line quantum repeater". In: *IEEE/ACM Transactions On Networking* 17.3 (2008), pp. 1002–1013.

[91] Alexander Pirker and Wolfgang Dür. "A quantum network stack and protocols for reliable entanglement-based networks". In: *New Journal of Physics* 21.3 (2019), p. 033003.

[92] Axel Dahlberg, Matthew Skrzypczyk, Tim Coopmans, Leon Wubben, Filip Rozpedek, Matteo Pompili, Arian Stolk, Przemysław Pawełczak, Robert Knegjens, Julio de Oliveira Filho, et al. "A link layer protocol for quantum networks". In: *Proceedings of the ACM Special Interest Group on Data Communication*. 2019, pp. 159–173.

[93] Wojciech Kozlowski, Axel Dahlberg, and Stephanie Wehner. "Designing a quantum network protocol". In: *Proceedings of the 16th International Conference on emerging Networking EXperiments and Technologies*. 2020, pp. 1–16.

[94] Matteo Pompili, Carlo Delle Donne, Ingmar te Raa, Bart van der Vecht, Matthew Skrzypczyk, Guilherme Ferreira, Lisa de Kluijver, Arian J Stolk, Sophie LN Hermans, Przemysław Pawełczak, et al. "Experimental demonstration of entanglement delivery using a quantum network stack". In: *npj Quantum Information* 8.1 (2022), p. 121.

A

Virtual metrics convergence

In this appendix, we motivate the assumption that the virtual node degree $k_i(t)$ and virtual neighbourhood size $v_i(t)$ converge to a steady-state value when averaged over samples, time or network nodes.

First, $k_i^s(t)$ and $v_i^s(t)$ are the virtual metrics associated with a simulation sample s. Then, the sample-averaged ($S = 5000$) virtual metrics of node any node i in an infinite chain,

$$\langle k_i(t) \rangle = \sum_{s=1}^{S} \frac{k_i^s(t)}{S}, \qquad \langle v_i(t) \rangle = \sum_{s=1}^{S} \frac{v_i^s(t)}{S}, \tag{A.1}$$

converge to a steady state value (Figure A.1 shows the metrics for node $i = 0$).

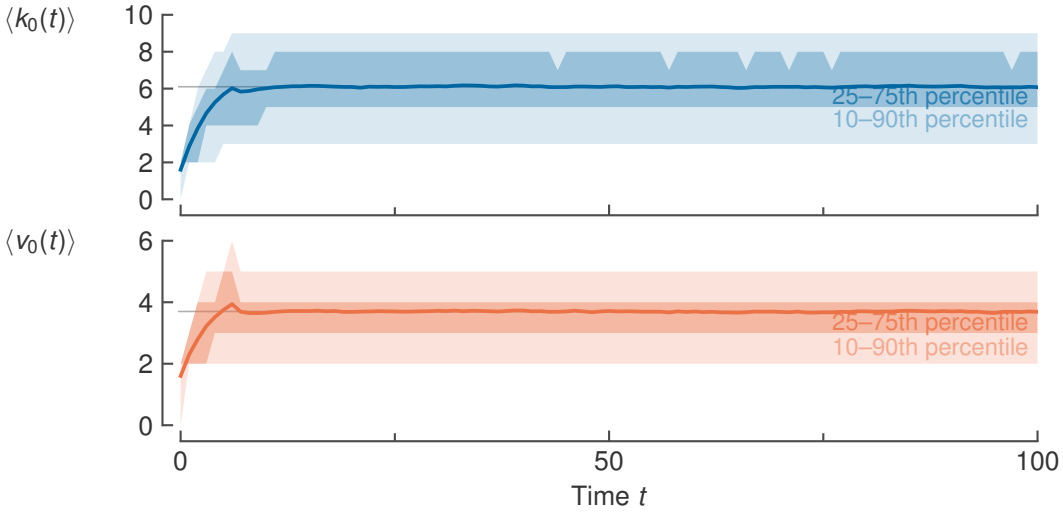

Figure A.1: The sample-averaged virtual metrics converge to a steady state. Time evolution showing the virtual node degree $k_i(t)$ and virtual neighbourhood size $v_i(t)$ of a specific node $i = 0$ in an infinite chain averaged over $S = 5000$ samples. The step-wise behaviour of the percentile bands is due to the integer nature of the virtual metrics. We use a swap probability $q = 0.2$ and the baseline parameters $p = 1, p_s = 1$, $T_2 = 100$, $t_{cut} = 7$, $F_{gen} = 0.9$, $M = 4$, $F_{min} = \frac{1}{2}$.

Similarly, the node-averaged ($N = 5000$) virtual metrics in an infinite chain,

$$\langle k_i(t)\rangle_i = \sum_{i=1}^{N} \frac{k_i(t)}{N}, \qquad \langle v_i(t)\rangle_i = \sum_{i=1}^{N} \frac{v_i(t)}{N}, \tag{A.2}$$

converge to the same steady state value in an infinite network (Figure A.2).

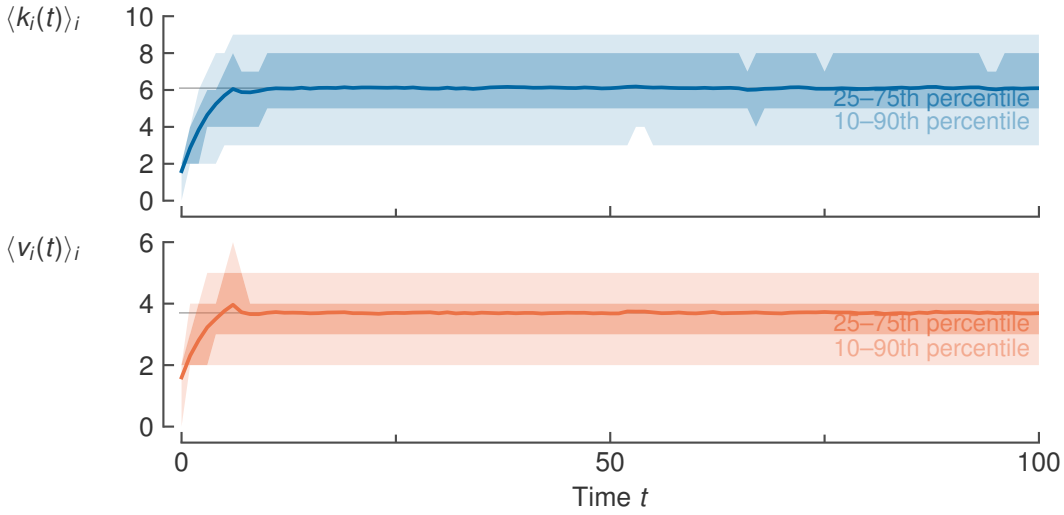

Figure A.2: The node-averaged virtual metrics converge to a steady state. Time evolution showing the virtual node degree $k_i(t)$ and virtual neighbourhood size $v_i(t)$ averaged over a chain of 5000 nodes with periodic boundary conditions. We use a swap probability $q = 0.2$ and the baseline parameters $p = 1, p_s = 1$, $T_2 = 100$, $t_{cut} = 7$, $F_{gen} = 0.9$, $M = 4$, $F_{min} = \frac{1}{2}$.

Lastly, the probability distributions of the virtual node degree $\Pr[k_i(t) = k]$ and the virtual neighbourhood size $\Pr[v_i(t) = v]$ are (almost) the same for all nodes after a simulation containing 10^7 time slots for an infinite chain (Figure A.2).

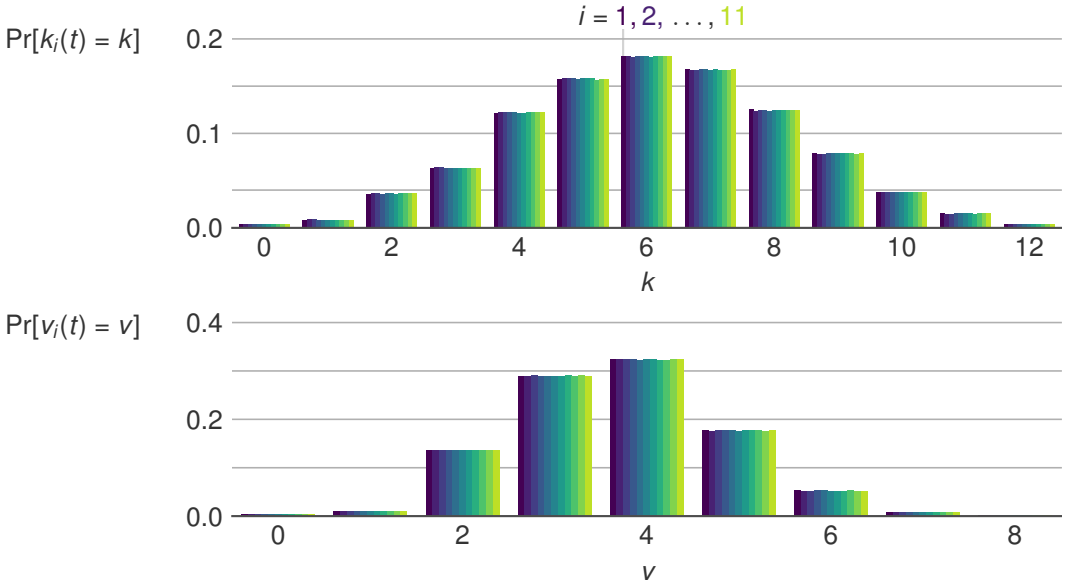

Figure A.3: All nodes in the *infinite* chain exhibit the same virtual metrics probability distributions. Histogram with probability distributions of the virtual node degree $k_i(t)$ and virtual neighbourhood size $v_i(t)$ in an 11-node chain with periodic boundary conditions (10^7 simulation steps). We use $q = 0.2$ and the baseline parameters $p = 1, p_s = 1$, $T_2 = 100$, $t_{cut} = 7$, $F_{gen} = 0.9$, $M = 4$, $F_{min} = \frac{1}{2}$.

B

Bounds on the virtual metrics

In this appendix, we derive the bounds on the virtual node degree and the virtual neighbourhood size for infinite networks with a regular topology (Table 2.2).

B.1 Bounds on the virtual node degree

A quantum node with an infinite number of quantum memories in a regular-topology network with physical node degree k_p and cutoff time t_cut can store at most $k_\mathrm{p}t_\mathrm{cut}$ entangled links. We note that quantum nodes generate at most k_p entangled links per time slot according to our CD protocol. Also, quantum nodes store most links when they do not attempt swaps, as swaps consume the initial links. Then, nodes only discard entangled links when their fidelity is too low due to living too long. At that point, the quantum node stores $k_\mathrm{p}t_\mathrm{cut}$ entangled links.

Thus, by definition of the virtual node degree – the number of entangled links stored by a node at a particular time – we conclude that $k_i(t) \leq k_\mathrm{p}t_\mathrm{cut}$.

B.2 Bounds on the virtual neighbourhood size

Generally, a node can not have more virtual neighbours than it stores entangled links, so the virtual node degree bounds the virtual neighbourhood size, $v_i(t) \leq k_i(t)$. Furthermore, if nodes do not attempt swaps, they only share entangled links with their physical neighbours. However, the number of virtual neighbours can increase when nodes attempt swaps, as they can create entangled links with remote, non-neighbouring nodes. We show that the number of virtual neighbours $v_i(t)$ increases in a similar pattern for increasing swap distance M.

The proof is similar for different topologies: initially, for $M = 0$, a quantum node only shares entanglement with its k_p physical neighbours. Then, for an increase from arbitrary $M \to M + 1$, the virtual neighbourhood size $v_i(t)$ increases in a regular pattern as a function of M (Figure B.1), concluding our proof by induction. Specifically:

- For $k_\mathrm{p} = 2$, the maximum virtual neighbourhood size $v_i(t)$ grows by 2 as $M \to M + 1$: the central node can have an additional virtual neighbour both to the left and right. Hence $v_i(t) \leq \sum_{k=0}^{M} 2 = 2(M + 1)$.

- For $k_\mathrm{p} = 3$, the maximum $v_i(t)$ grows by $3(M+1)$ as $M \to M+1$. For $M \geq 1$, the number of virtual neighbours visually grows in layers: alternately, a layer with six equal sides of length $M+1$ (a hexagon), and a layer with three short sides of length M and three long sides of length $M+1$ are added. Hence $v_i(t) \leq \sum_{k=0}^{M} 3(k+1) = \frac{3}{2}(M+1)(M+2)$.
- For $k_\mathrm{p} = 4$, the maximum $v_i(t)$ grows by $4(M+1)$ as $M \to M+1$. Similar to $k_\mathrm{p} = 3$, the number of virtual neighbours visually grows in layers: for each step $M \to M+1$, a square layer (angled at $\pi/4$) with side length $M+1$ is added. Hence $v_i(t) \leq \sum_{k=0}^{M} 4(k+1) = 2(M+1)(M+2)$.
- For $k_\mathrm{p} = 6$, the maximum $v_i(t)$ grows by $6(M+1)$ as $M \to M+1$. Again, the number of virtual neighbours grows in layers: for each step $M \to M+1$, a hexagonal layer with a base length of $M+1$ is added. Hence $v_i(t) \leq \sum_{k=0}^{M} 6(k+1) = 3(M+1)(M+2)$.

Combining these bounds with $v_i(t) \leq k_i(t)$, we have derived the bounds from Table 2.2.

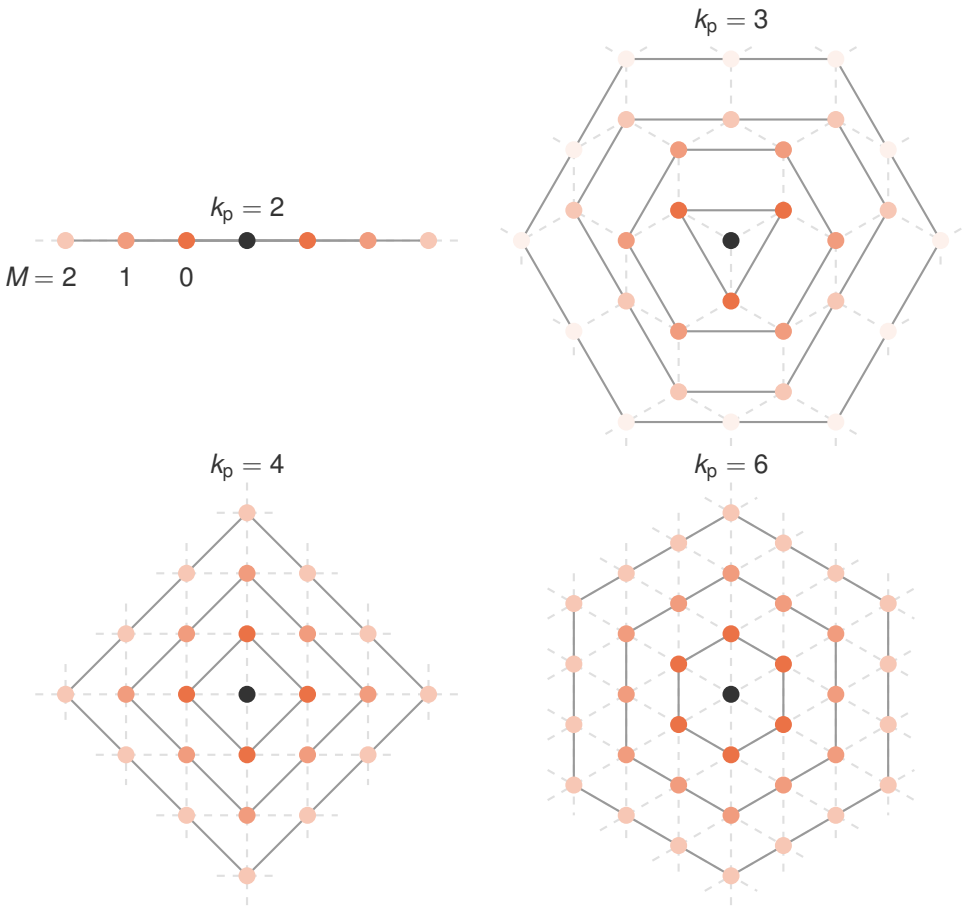

Figure B.1: The maximum virtual neighbourhood size grows similarly for increasing *M*. The maximum number of virtual neighbours, $v_i(t)$, of a node (black) grows similarly for each increase $M \to M+1$ (emerging structure illustrated by solid line): $v_i(t)$ grows with 2 for $k_\mathrm{p} = 2$; with $3(\mathrm{M}+1)$ for $k_\mathrm{p} = 3$; with $4(\mathrm{M}+1)$ for $k_\mathrm{p} = 4$; and with $6(M+1)$ for $k_\mathrm{p} = 6$. Note the underlying physical structures: chains, honeycombs, squares and triangles (dashed lines).

C

Virtual metrics distribution

In this appendix, we present the mean virtual metrics as a function of the swap probability and varying simulation parameters including their distribution in the form of percentiles.

The virtual node degree $k_i(t)$ and virtual neighbourhood size $v_i(t)$ are time-dependent random variables. We analyse their performance by averaging many samples to obtain the average virtual node degree $\langle k_i(t) \rangle$ and average virtual neighbourhood size $\langle v_i(t) \rangle$ and their corresponding probability distributions (see Section 3.2 for more details). For clarity, we have not shown the probability distributions in the main text, but we do so in this appendix. Specifically, we analyse the 25–75th and 10–90th percentiles associated with the samples that average to $\langle k_i(t) \rangle$, $\langle v_i(t) \rangle$. The percentiles offer a suitable approach to analysing the probability distributions of $k_i(t)$ and $v_i(t)$ as their distributions are often not symmetric due to upper and lower bounds on the virtual metrics. Histograms exemplify the probability distributions of $k_i(t)$ and $v_i(t)$ more explicitly for a specific combination of network parameters and swap probability (Figure A.3). Lastly, we note the step-like behaviour of the percentiles due the integer nature of the virtual metrics.

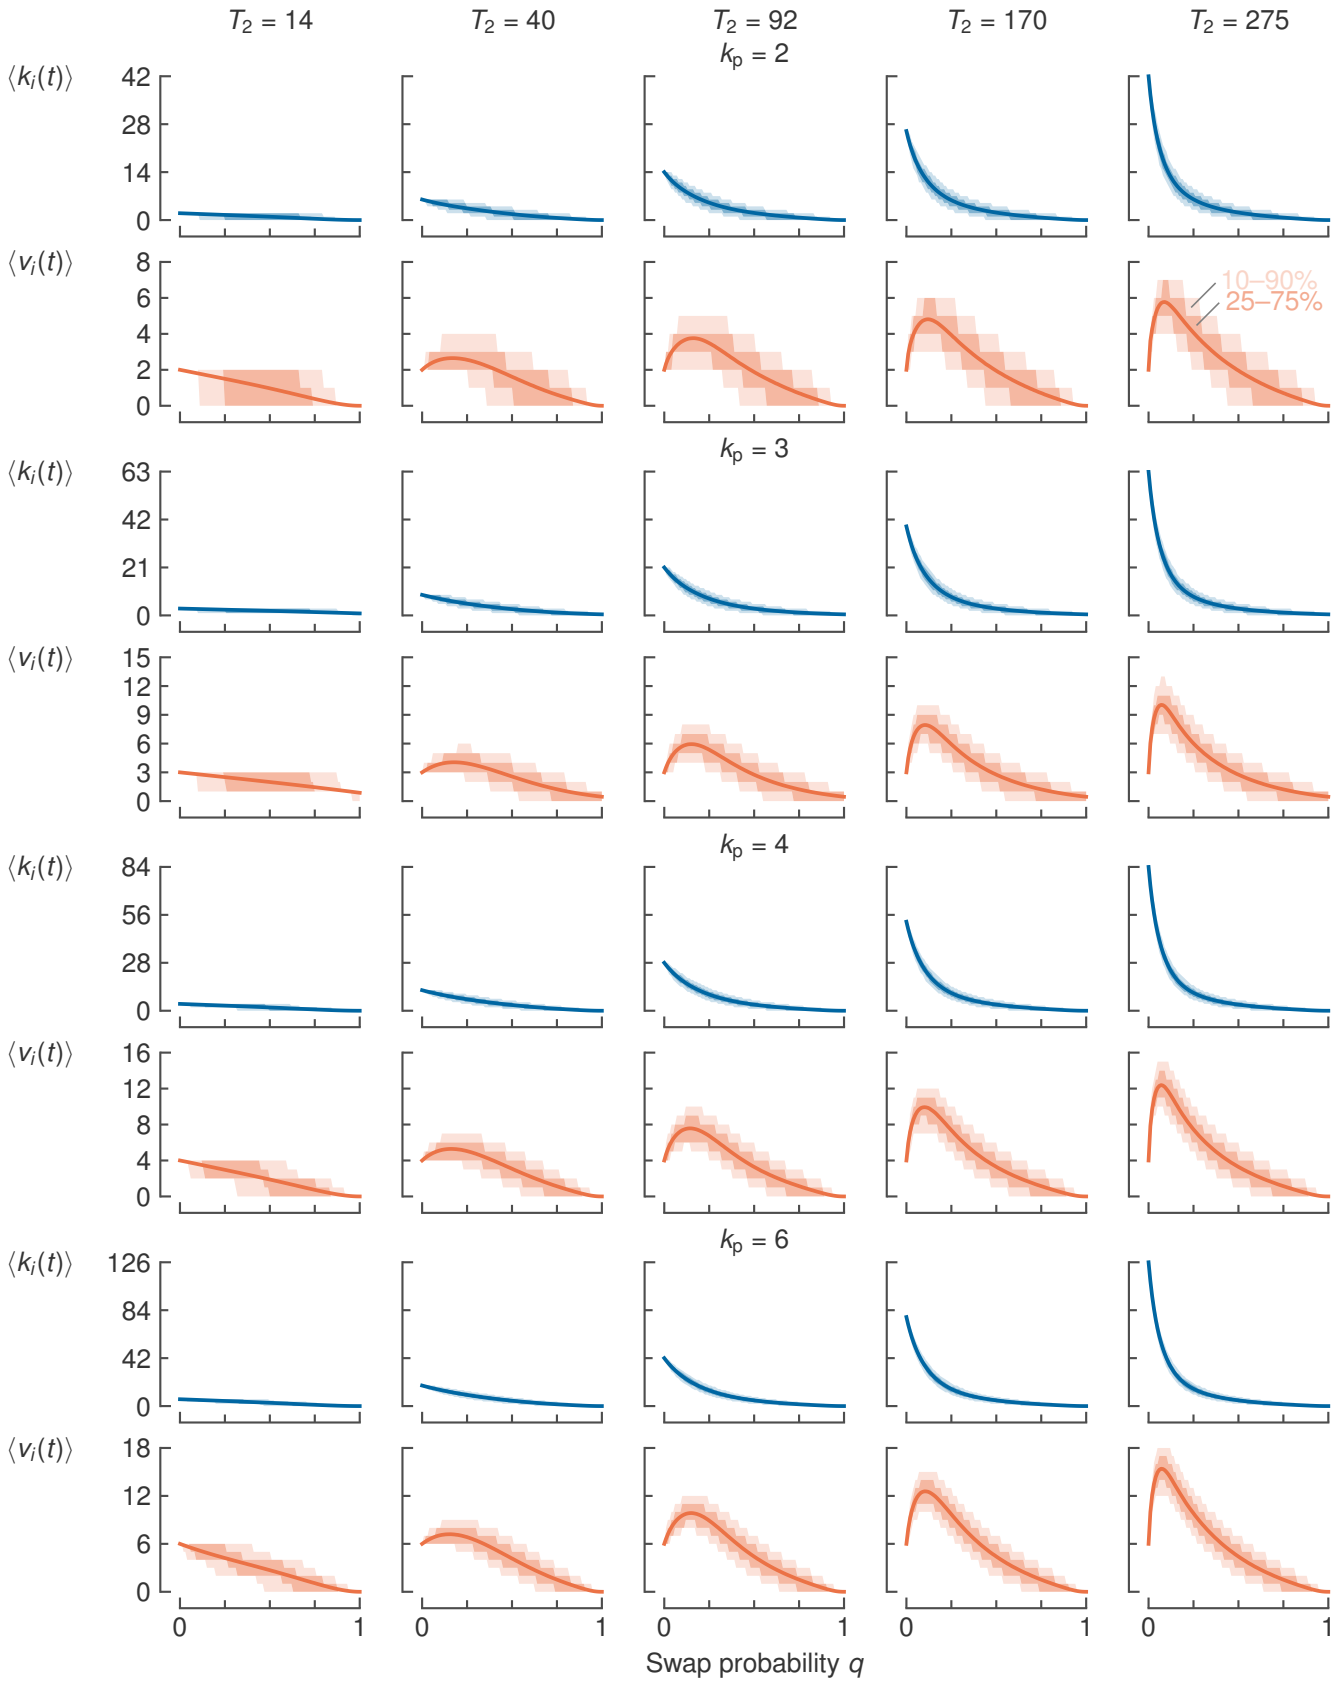

Figure C.1: Average virtual metrics and their underlying probability distribution as a function of the swap attempt probability and coherence time of infinite regular networks. We vary the coherence time $T_2 = 14, 40, 92, 170, 275$ (and associated cutoff times $t_{\text{cut}} = 1, 3, 7, 13, 21$) while using the baseline parameters: entanglement generation success probability $p = 1$, swap success probability $p_s = 1$, entanglement generation fidelity $F_{\text{gen}} = 0.9$, maximum swap distance $M = 4$, and minimum entangled link fidelity $F_{\text{min}} = \frac{1}{2}$. The probability distributions are expressed as 25–75th and 10–90th percentile bands.

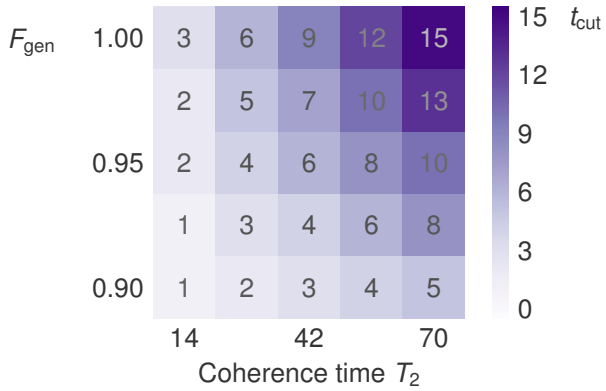

Figure C.2: Cutoff time as a function of the entanglement generation fidelity and coherence time. We vary the coherence time $T_2 = 14, 28, 42, 56, 70$ and entanglement generation fidelity $F_{\text{gen}} = 0.90, 0.925, 0.95, 0.975, 1$ and associate the maximum cutoff times t_{cut} satisfying Equation (2.4) while using the baseline parameters: entanglement generation success probability $p = 1$, swap success probability $p_{\text{s}} = 1$, maximum swap distance $M = 4$, and minimum entangled link fidelity $F_{\text{min}} = \frac{1}{2}$.

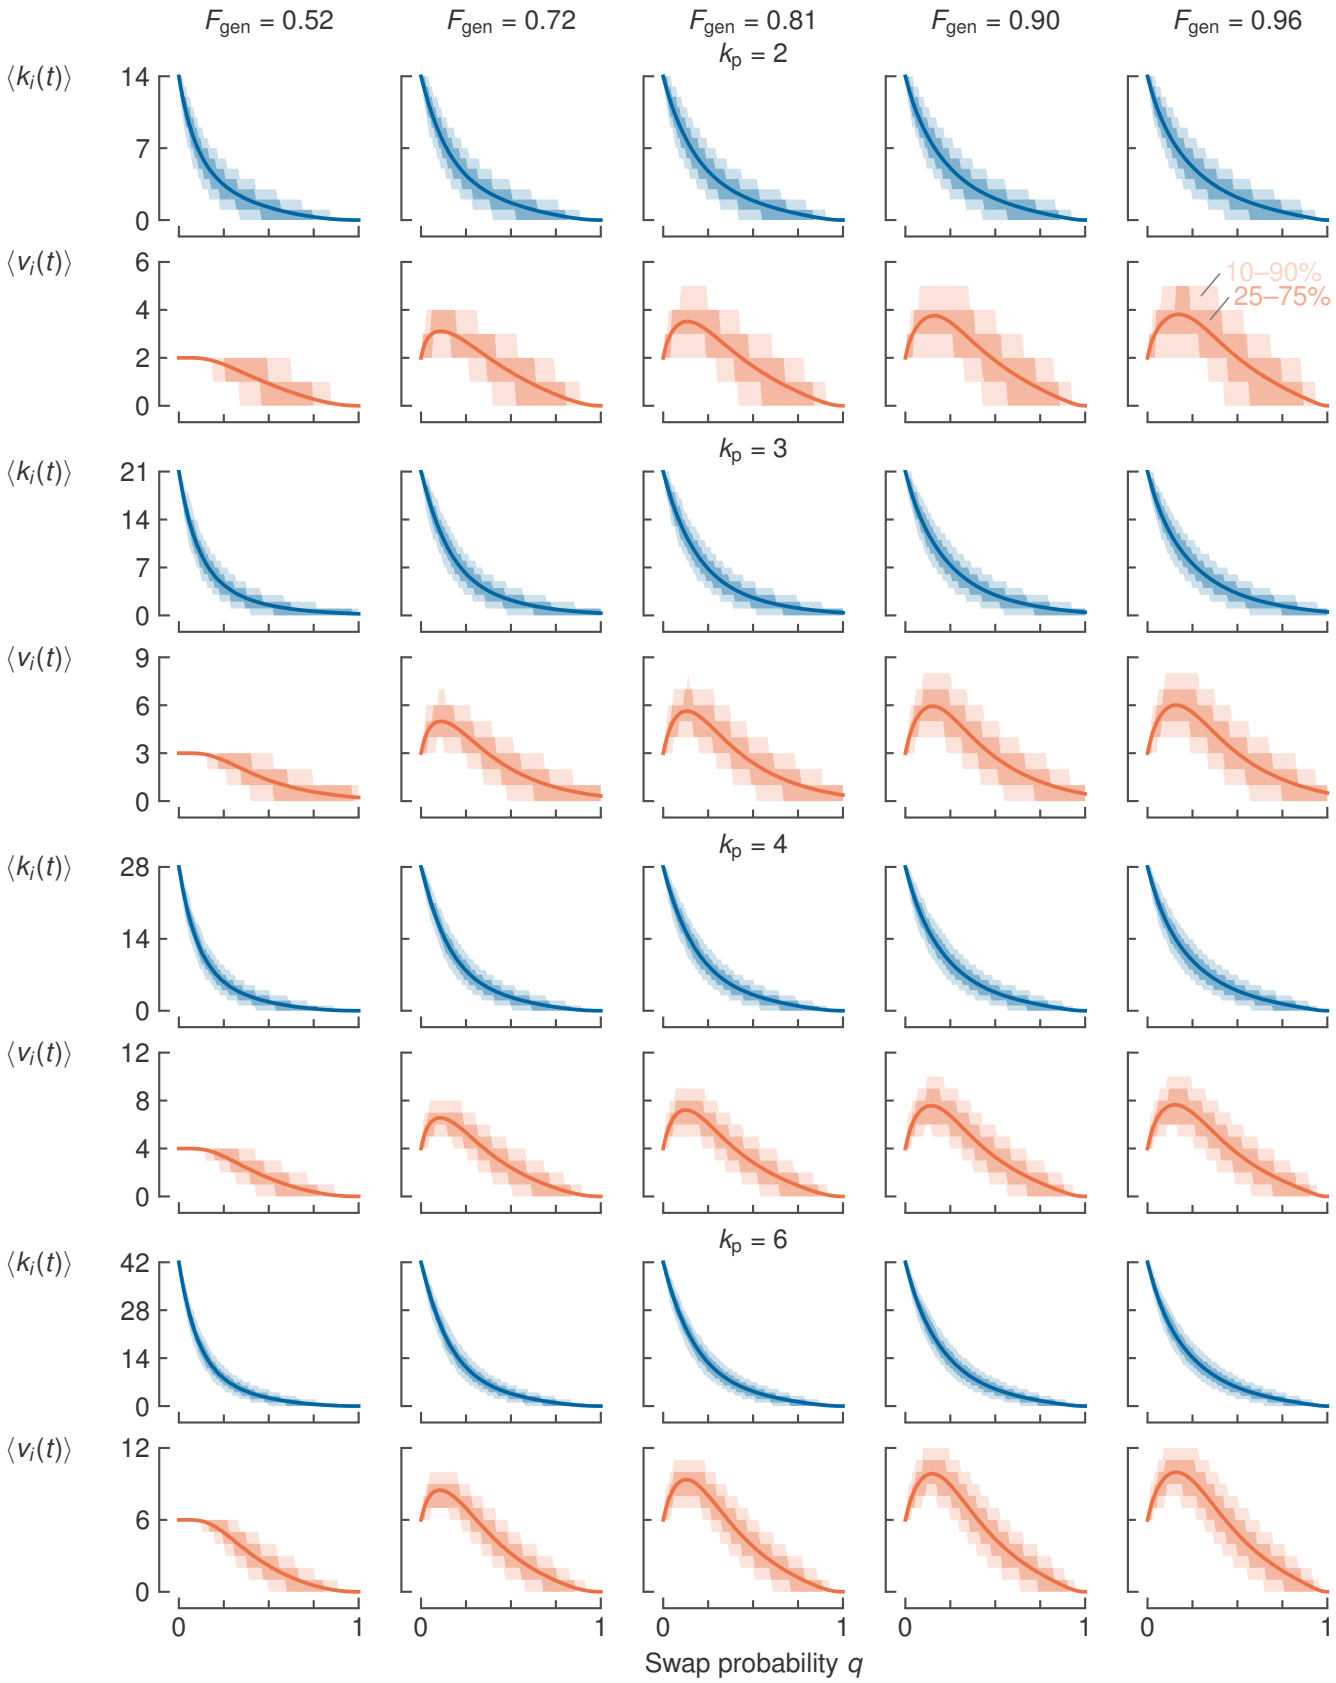

Figure C.3: Average virtual metrics and their underlying probability distribution as a function of the swap attempt probability and entanglement generation fidelity of infinite regular networks. We vary the entanglement generation fidelity $F_{\text{gen}} = 0.52, 0.72, 0.81, 0.90, 0.96$ (associated maximum swap distance $M = 0, 1, 2, 4, 7$) while using the baseline parameters: entanglement generation success probability $p = 1$, swap success probability $p_{\text{s}} = 1$, coherence time $T_2 = 100$, cutoff time $t_{\text{cut}} = 7$, and minimum entangled link fidelity $F_{\text{min}} = \frac{1}{2}$. The probability distributions are expressed as 25–75th and 10–90th percentile bands.

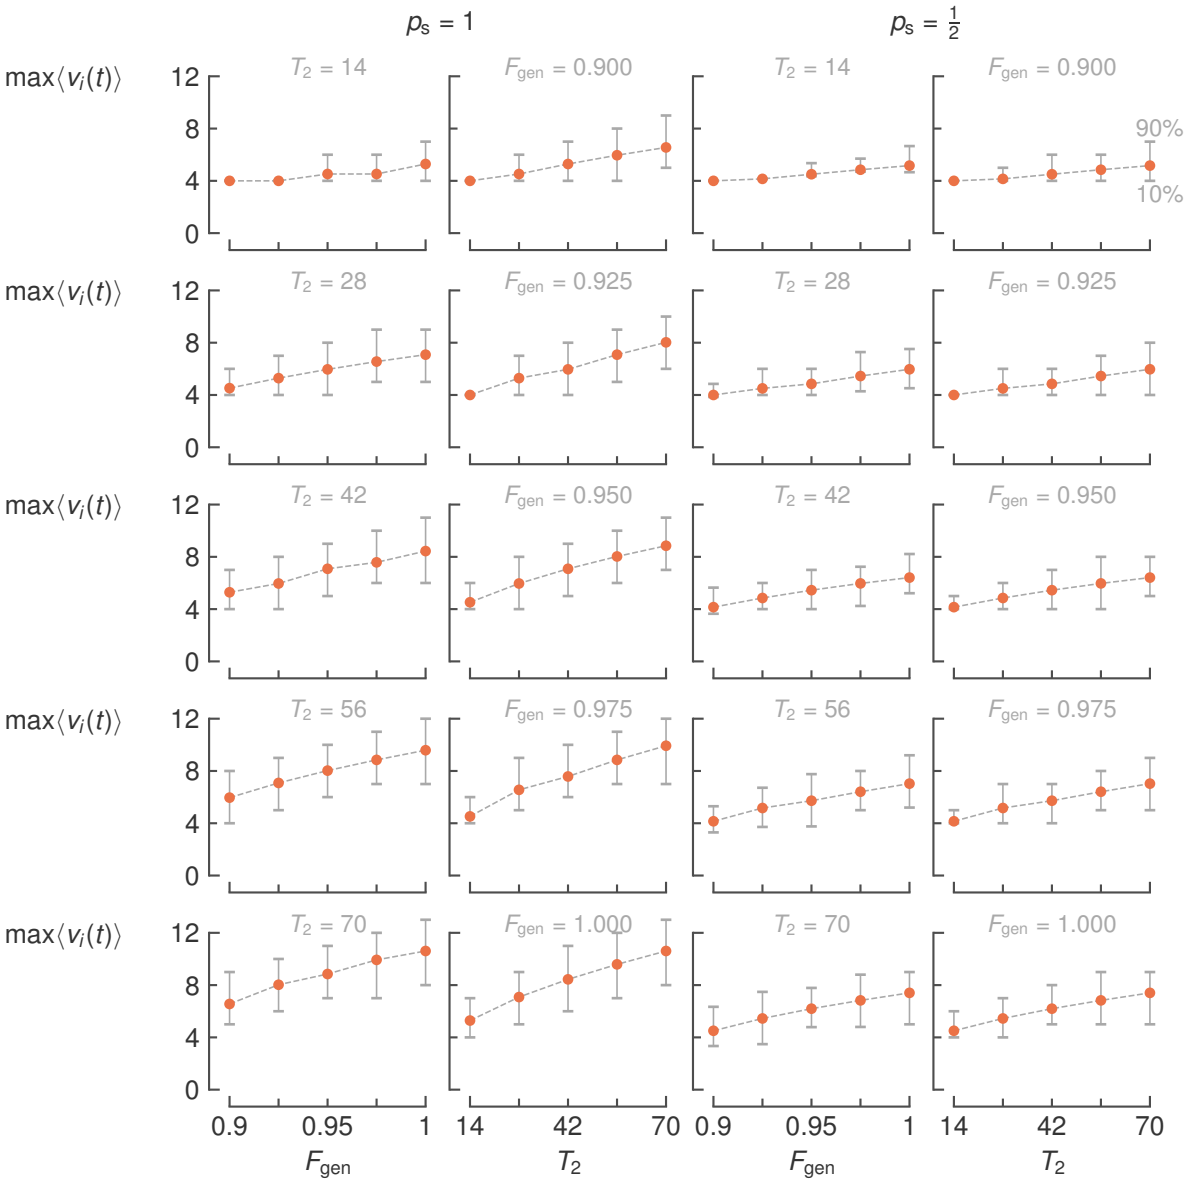

Figure C.4: Maximum average virtual neighbourhood size and the underlying probability distribution as a function of the coherence time and entanglement generation fidelity of infinite regular networks. We vary the coherence time $T_2 = 14, 28, 42, 56, 70$ and entanglement generation fidelity $F_{gen} = 0.90, 0.925, 0.95, 0.975, 1$ and associate the maximum cutoff times t_{cut} satisfying Equation (2.4) while using the baseline parameters: entanglement generation success probability $p = 1$, swap success probability $p_s = 1$, maximum swap distance $M = 4$, and minimum entangled link fidelity $F_{min} = \frac{1}{2}$. The probability distributions are expressed as 10–90th percentiles.

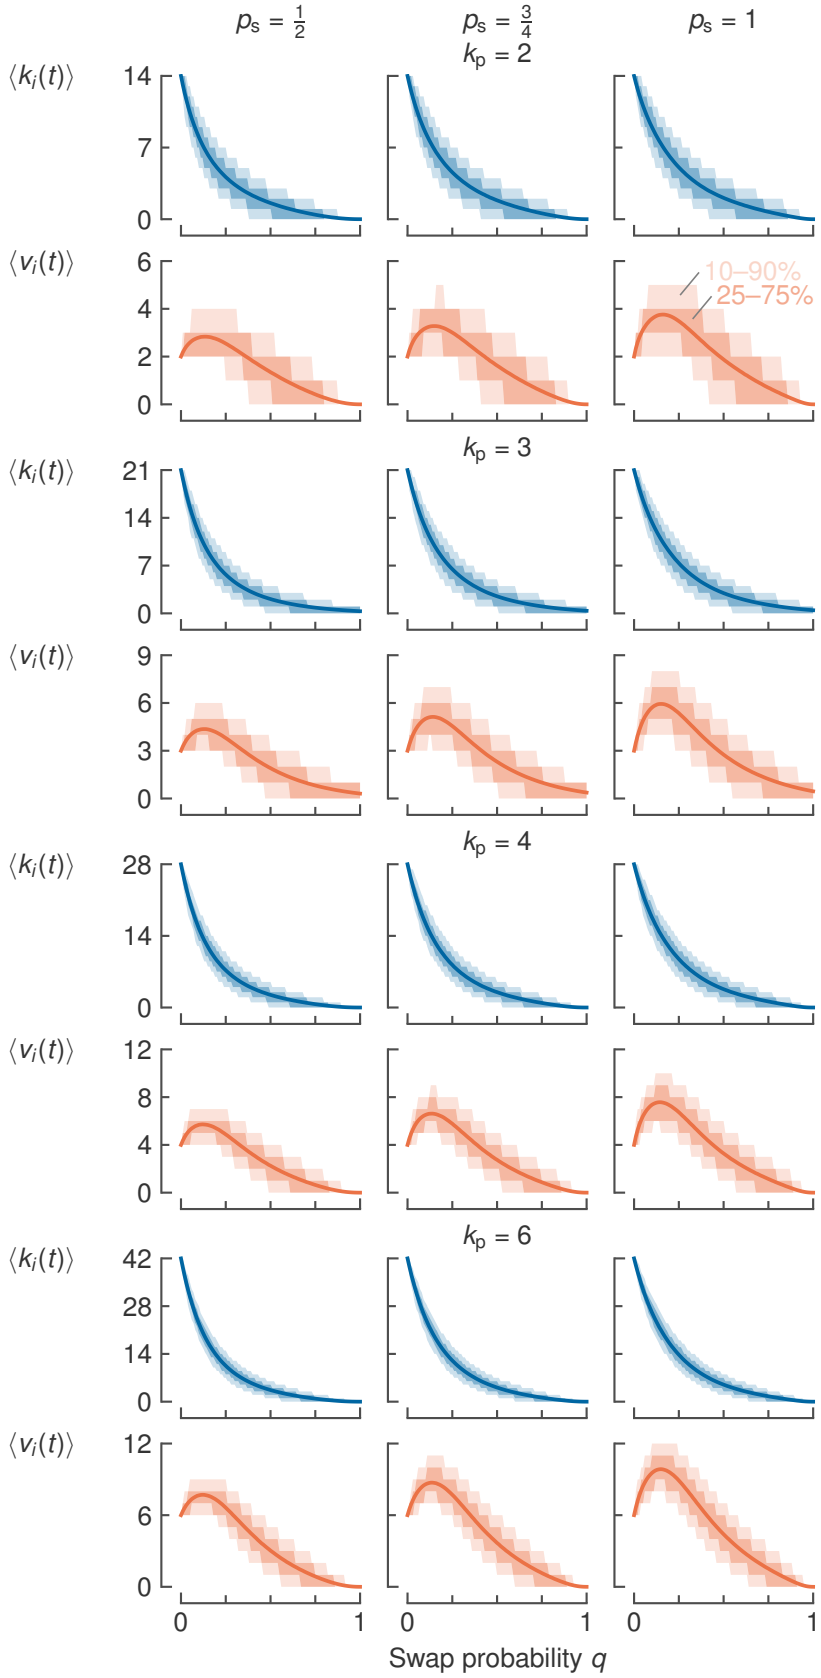

Figure C.5: Average virtual metrics and their underlying probability distribution as a function of the swap attempt probability and swap success probability of infinite regular networks. We vary the swap success probability $p_s = \frac{1}{2}, \frac{3}{4}, 1$ while using the baseline parameters: entanglement generation success probability $p = 1$, coherence time $T_2 = 100$, cutoff time $t_{cut} = 7$, entanglement generation fidelity $F_{gen} = 0.9$, maximum swap distance $M = 4$, and minimum entangled link fidelity $F_{min} = \frac{1}{2}$. The probability distributions are expressed as 25–75th and 10–90th percentile bands.

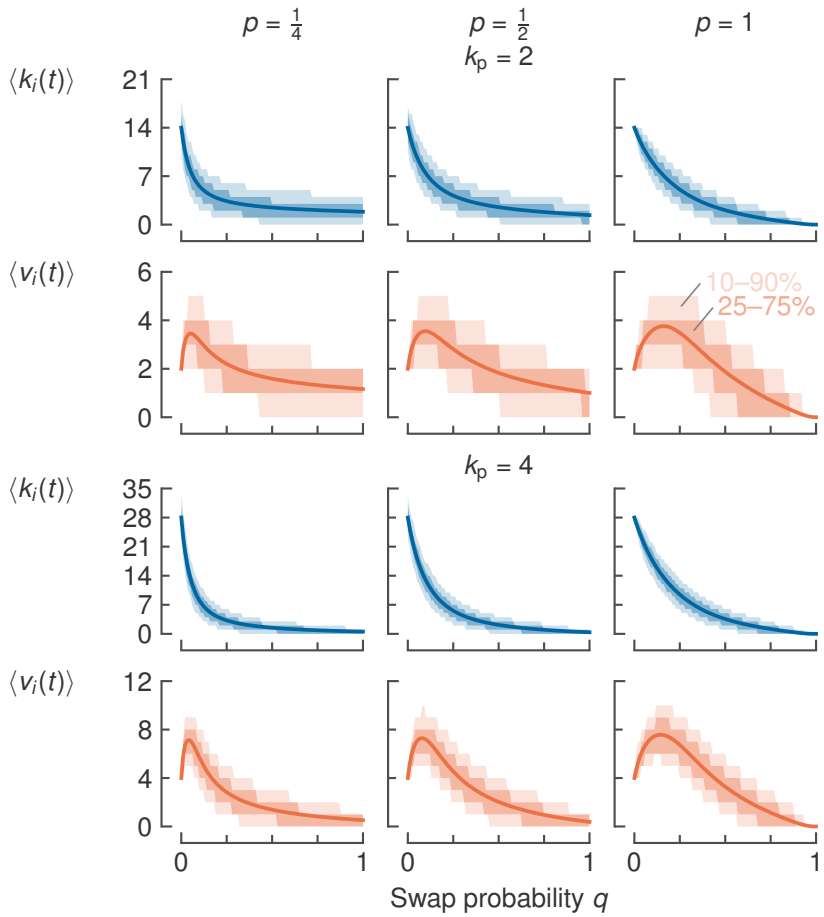

Figure C.6: Average virtual metrics and their underlying probability distribution as a function of the swap attempt probability and entanglement generation success probability of infinite regular networks. We vary $p = \frac{1}{4}, \frac{1}{2}, 1$ while correspondingly increasing the $T_2 = 100, 200, 400$ ($t_{\text{cut}} = 7, 14, 28$) and using the baseline parameters: swap success probability $p_s = 1$, entanglement generation fidelity $F_{\text{gen}} = 0.9$, maximum swap distance $M = 4$, and minimum entangled link fidelity $F_{\text{min}} = \frac{1}{2}$. The probability distributions are expressed as 25–75th and 10–90th percentile bands.

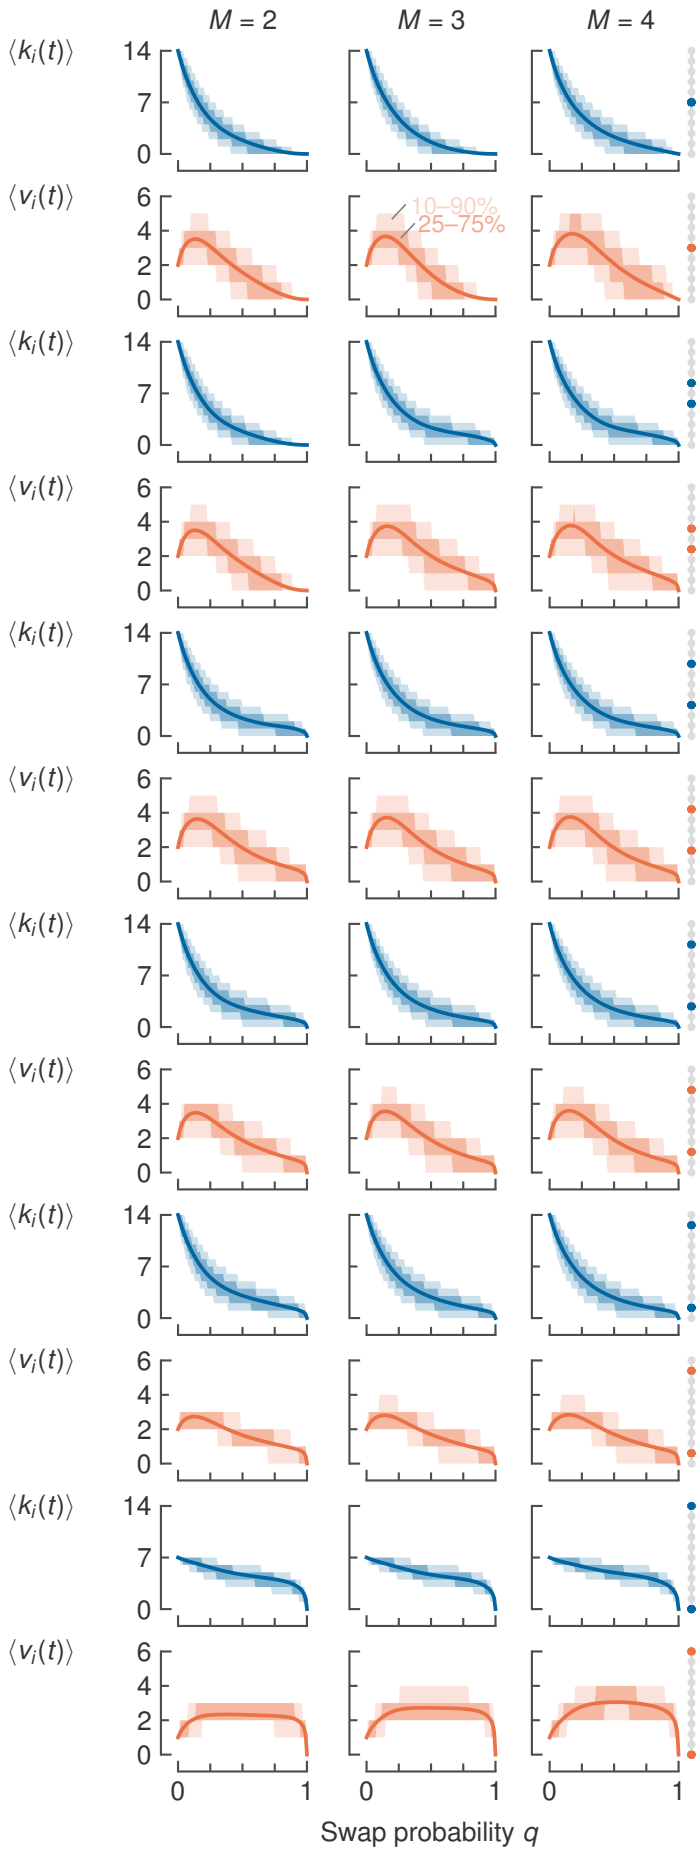

Figure C.7: Average virtual metrics and their underlying probability distribution as a function of the swap attempt probability, maximum swap distance and location in an 11-node chain. We vary the maximum swap distance $M = 2, 3, 4$ (associated entanglement generation fidelity $F_{\text{gen}} = 0.81, 0.87, 0.90$) while using the baseline parameters: $p = 1$, $p_{\text{s}} = 1$, $T_2 = 100$, $t_{\text{cut}} = 7$, and $F_{\text{min}} = \frac{1}{2}$. The probability distributions are expressed as 25–75th and 10–90th percentile bands.

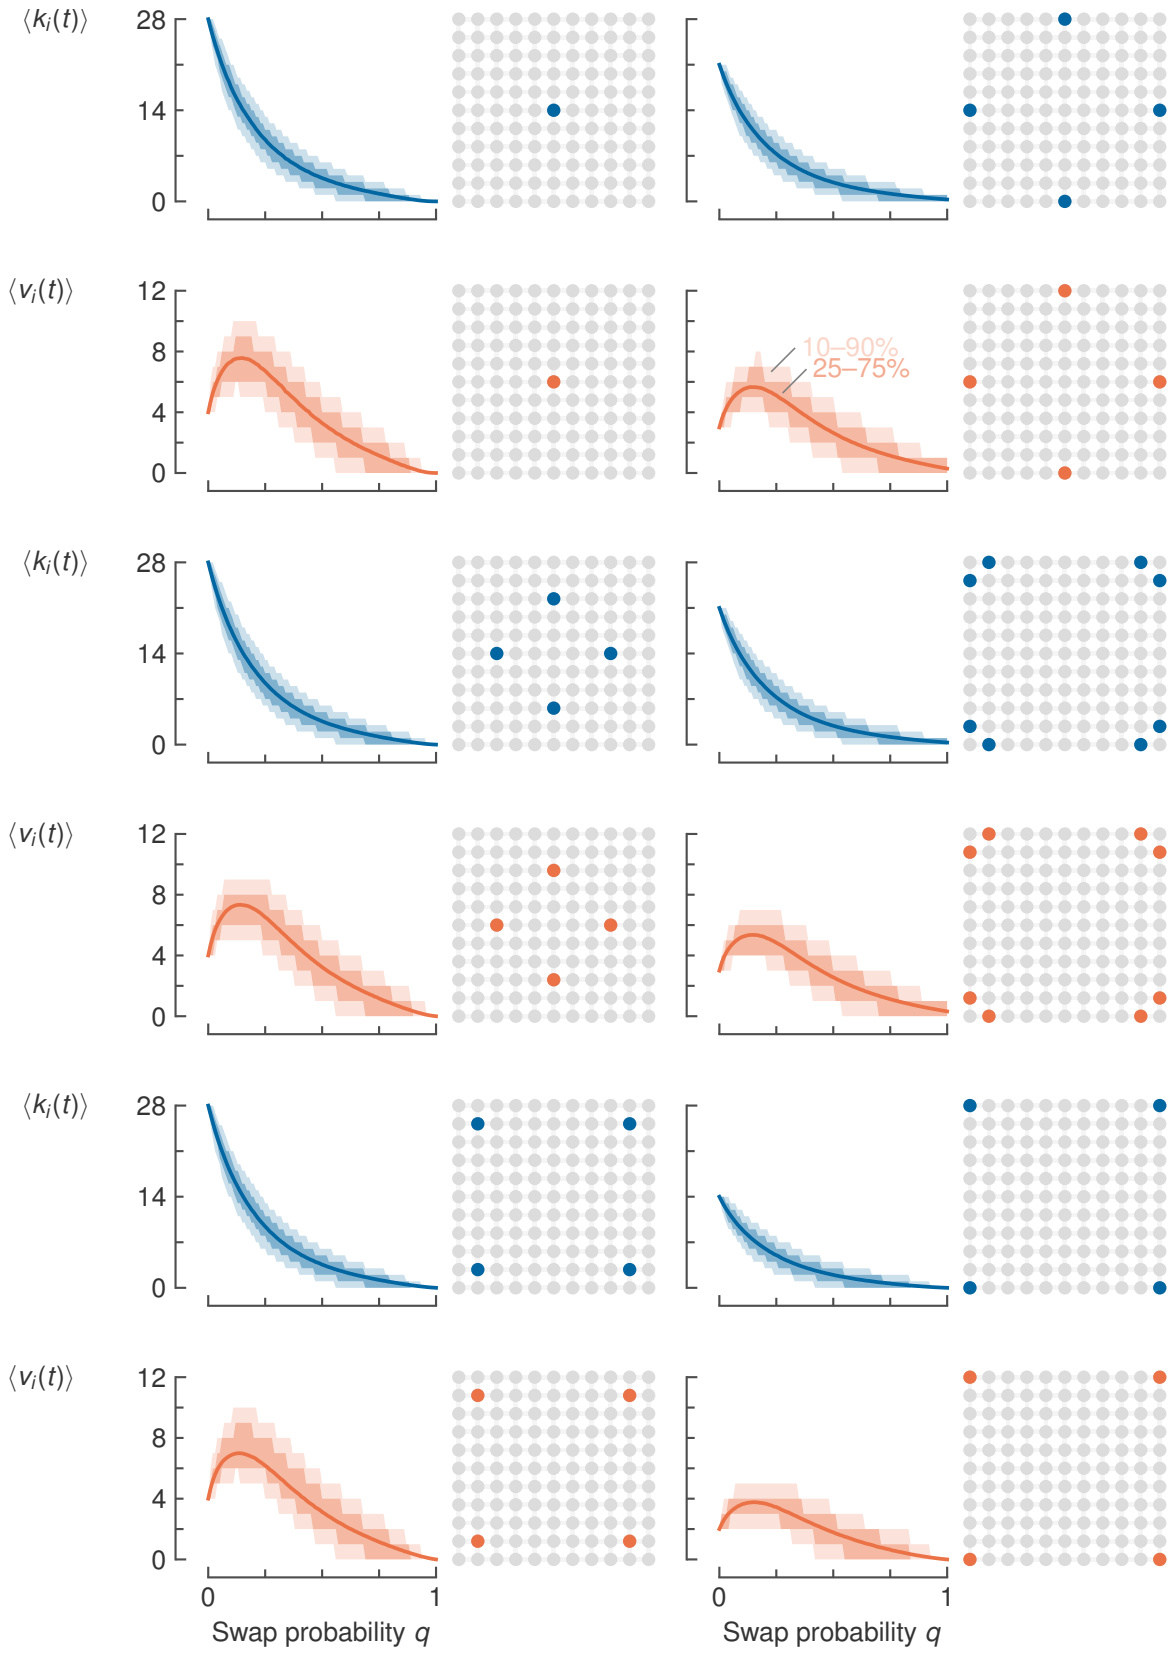

Figure C.8: Average virtual metrics and their underlying probability distribution as a function of the swap attempt probability and location in 11 by 11 square grid of nodes. We use the baseline parameters: entanglement generation success probability $p = 1$, swap success probability $p_s = 1$, coherence time $T_2 = 100$, cutoff time $t_{\text{cut}} = 7$, entanglement generation fidelity $F_{\text{gen}} = 0.9$, maximum swap distance $M = 4$, and minimum entangled link fidelity $F_{\text{min}} = \frac{1}{2}$. The probability distributions are expressed as 25–75th and 10–90th percentile bands.