Turning Quantum Cryptography against itself: How to avoid indirect eavesdropping in quantum networks by passive and active adversaries

Stefan Rass
Alpen-Adria Universität Klagenfurt, Department of Applied Informatics
Universitätsstrasse 65-67
9020 Klagenfurt, Austria
stefan.rass@uni-klu.ac.at

Sandra König
Alpen-Adria Universität Klagenfurt
Universitätsstrasse 65-67
9020 Klagenfurt, Austria
sakoenig@edu.uni-klu.ac.at

***Abstract*—Quantum networks are communication networks in which adjacent nodes enjoy perfectly secure channels thanks to quantum key distribution (QKD). While QKD is renowned for perfect point-to-point security and its eavesdropping detection capabilities, end-to-end security is nontrivial to achieve. More importantly, the eavesdropping detection can indeed be turned against the system itself. It is known that perfect end-to-end security can be created from point-to-point security by means of multipath transmission (in fact, there is no other way to do this, assuming no pre-shared secrets and avoiding public-key cryptography). However, multipath transmission requires node-disjoint paths, which in turn are to be assured by the underlying routing protocol. At this point, an active or passive adversary may intentionally eavesdrop on the QKD protocol to temporarily cut a channel and to cause key-buffers running empty and enforcing local rerouting of packets towards nodes under his control. Consequently, the multipath transmission channels might no longer be non-intersecting, thus defeating the overall security by exploiting QKD's eavesdropping detection facilities. Alternatively, an active adversary may as well insert bogus traffic to cause local congestion, thus even sparing the effort of eavesdropping on a QKD link. In this work, we use Markov chains to model a multipath transmission, and we discuss the extent to which secure multipath transmission is resilient against local congestions caused by an adversary. We argue that a protection against an active adversary who uses bogus traffic to fiddle with the routing, calls for additional security measures, perhaps even beyond the capabilities of QKD or multipath transmission. It turns out that robustness against passive and active adversaries can be retained as long as no bogus traffic is observed.**

***Keywords*-Quantum Cryptography, Markov-Chain, Secure Routing, Information-Theoretic Security**

I. Introduction

Quantum key distribution (QKD) is known to provide perfect point-to-point security by virtue of its capability to detect passive eavesdropping. Despite considerable progress and ingenious concepts and results, QKD remains yet mostly limited to secure point-to-point connections. Although the theory of quantum repeaters is available in rich detail [2], these devices have not yet evolved beyond laboratory demonstrator status. On the classical road, perfect end-to-end security is achievable by means of multipath transmission. Remarkably, multiple paths have been proven to be both, a necessary and sufficient condition for perfect secrecy along a multihop connection (w.r.t. *not* assuming quantum repeater based transmission). The idea and security of such protocols (e.g., the one proposed by [3]) hinges on the chosen transmission paths to be pairwise non-intersecting. However, re-routing due to local congestions or intentionally caused blockages by the adversary can cause the network to temporarily allow intersections of paths and thus give an adversary an advantage when eavesdropping on relay nodes. More specifically, if the transmission uses t paths that are supposed to be disjoint (except for their respective end-points) then security against an adversary having up to k nodes under his control is not endangered as long as $t > k$ and the paths remain disjoint. More specifically, multipath transmission pursues the following general construction: to transmit a message m, the sender first puts it through a threshold secret sharing scheme, e.g., Shamir's (t, n)-scheme or plain (n, n)-sharing via the XOR of a sequence of random values, i.e., $m = s_1 \oplus s_2 \oplus \cdots \oplus s_n$, where $\oplus$ is the bitwise exclusive or. Each share s_i then travels over his own distinct path to the receiver, who reconstructs the message according to the chosen sharing scheme. In Shamir's case, this requires at least $t+1$ shares and in case of an XOR-sharing, all n shares are needed to recover m. In either case, the adversary needs to catch at least $t+1$ shares, respectively n shares, in order to learn anything. The simplest way to enforce a maximal number of corrupted nodes for that matter is having the paths pairwise non-intersecting, i.e., node-disjoint. If congestions cause local redirections such that multiple paths intersect in the same node, then the security of the transmission is doomed to fail, since the adversary may learn the required number of shares while perhaps having a much smaller number of nodes under his control. We introduce an attack in which the adversary exploits the eavesdropping detection facility of QKD without attempting to learn any of the secret key material. Instead, his only goal is to make the link run dry of key-material, so as to enforce the local neighborhood nodes to search for alternative paths over nodes that he controls. We call this an *indirect eavesdropping attack*.

The goal of this work is to investigate the resilience – in terms of end-to-end security – of quantum networks to such kind of incidents. We consider both, a passive and active adversary, being computationally unbounded and only constrained to have no more than k nodes in the network under his control. Moreover, we assume the routing to be under partial control of the sender, so that he can initiate a multipath transmission, but his chosen paths are potentially subject to temporary rerouting due to congestions. These congestions can be actively caused by the adversary, or coincidentally happen due to other reasons. In the latter case, we obtain simple sufficient criteria for perfectly secure communication remaining possible even if the routing is imperfect. The case of an active adversary causing traffic redirections is discussed based on these preliminary results.

Organization of the paper: We consider networks employing QKD for point-to-point- and multipath routing for end-to-end security, referred to as *quantum networks*. We briefly review the use of QKD with multipath transmission in Section IV. In Section V, we introduce a Markov-chain model for the path that a data packet takes from the sender to the receiver, with a particular focus to multipath transmission. Conditions under which an unreliable routing regime can yield perfect secrecy are derived in Section VI. Section VII is devoted to a discussion of active adversaries by extending the results from Section VI accordingly. Under suitable assumptions on his capabilities, we can retain security even against an active adversary. Dropping these assumptions, we demonstrate how an active adversary can indirectly influence the routing so as to direct the information flow towards his nodes without direct access to the routing. This form of indirect eavesdropping attack works even without using the eavesdropping facility of the underlying QKD protocol. An example supporting the practicability of our results is found in Section VIII. Final remarks are given in Section IX.

II. Related work

This work extends previous research described in [1]. Although eavesdropping detection in quantum key distribution [4] is quite well researched, only few authors deal with routing issues and even less consider problems arising from unreliable routing. Most closely related to ours is the work of [5], who provide a stochastic routing algorithm along with probabilistic measures of secrecy in a randomly compromised network. We improve on this by avoiding the assumption of some fixed routing algorithm. Instead, we formulate conditions under which a given routing protocol can provide perfect secrecy under random compromission. Consequently, the framework devised here is generic and requires simulations and empirical evaluation of the routing scheme at hand in order to be applied. Fortunately, simulation tools like OmNet++ [6] can rapidly provide such information. Practical QKD implementations are often subject to physical distance limitations (cf. [7], [8], [9] to name a few). Although unlimited distance QKD transmission is theoretically possible (see [10]), multipath transmission over disjoint channels remains up to now a practical necessity for perfect end-to-end security [11]. In particular, [3], [12], [13], [14] and references therein form the basis for our work, where our goal is to investigate a hidden assumption within these results: what happens if the routing is not fully reliable? Implementations of multipath transmission within the transmission control protocol (TCP) are currently under standardization, and many other protocols like stream control transmission protocol (SCTP [15]) as well facilitate concurrent transmission. Similarly as for a recently proposed extension of the secure socket layer (SSL) by QKD [16], [17], one could imagine QKD being integrated in such protocols. Load-balancing, local congestions and most importantly (adversarial) eavesdropping can all cause re-routing of packets and therefore make otherwise disjoint routes intersecting. Our work is an explicit account for security under such random distortions. To the best of our knowledge, such indirect eavesdropping attacks have not yet been considered elsewhere in the literature.

III. Preliminaries and Notation

Let $M \in \{0,1\}^*$ denote a binary string of arbitrary length. Let $|M|$ be its length (in bits), and let $H(M)$ denote the Shannon-entropy of a random message source M. A quantum network is an undirected graph $G = (V, E)$ in which each pair of adjacent nodes shares a channel that is secured by means of quantum key distribution. The sets of nodes and edges in G are denoted by $V(G)$ and $E(G)$, respectively. An $s-r$-*path* in a graph is an ordered sequence of adjacent nodes starting with $s \in V$ and ending in $r \in V$. We will denote a (general) path by ρ, and its set of nodes will be written as $V(\rho)$. Two $s-r$-paths ρ_1, ρ_2 are said to be *node-disjoint*, if $V(\rho_1) \cap V(\rho_2) = \{s, r\}$, i.e., the paths do not intersect elsewhere than in their start- and end-nodes. For any node $v \in V(G)$, we denote the collection of its immediate neighbors as $\mathrm{nb}(v) := \{u \in V | (v, u) \in E\}$.

Markov chains: As our routing model will be based on Markov chains, we briefly review the respective basics for convenience of the reader. We will straightforwardly focus on graph models for our introduction: once Alice has handed over her encrypted payload to the network for delivery to Bob, the actual journey of the packet can be considered as a random walk through the network until it reaches its final destination. Though the routing itself is essentially deterministic, randomness comes into play due to local congestions and subsequent re-routing. Consequently, we can consider the packet as describing a trajectory of a *stochastic process*, or more specifically a *Markov chain*, whose state space is the set $V(G)$, i.e., the set of all relay nodes that the packet can possibly visit. For any two nodes

$u, v \in V(G)$, assume that the packet travels from u to v with probability $p_{uv} = \Pr[u \to v]$. Since $V(G)$ is finite, we can fix any enumeration $V(G) = \{1, 2, \ldots, n\}$ and write p_{ij} for the chance of the packet traveling from i to j within one hop. To model this hop-by-hop forwarding, let us introduce the random variable $X(\tau) \in V(G)$ for $\tau = 1, 2, 3, \ldots$ telling us the node that hosts the data packet at time-step $\tau \in \mathbb{N}$. A *trajectory* is the sequence $(X(0), X(1), X(2), \ldots)$ describing the packet's trace, starting off at the sender $X(0)$ until it reaches its final destination (the receiver) at some later point in time. In terms of conditional probability, we have $p_{ij} = \Pr[X(\tau + 1) = j | X(\tau) = i]$ describing the one-step transition probability. The (one-step) *transition matrix* is defined as the $(n \times n)$-matrix $P = (p_{ij})_{i,j=1}^{n}$.

As we are dealing with multipath transmission in the following, consider t independent copies of a trajectory, named $1, 2, \ldots, t$. The particular state of the i-th trajectory at time τ is written as $X_i(\tau)$. Let the function $\pi_i(\tau, v) : \mathbb{N} \times V \to [0, 1]$ describe the chance that the i-th trajectory ($i = 1, 2, \ldots, t$) is within node v at time $\tau \in \mathbb{N}$, i.e., $\pi_i(\tau, v) = \Pr[X_i(\tau) = v]$. The whole distribution (supported on the set of nodes $V(G)$) is denoted as $\pi_i(\tau)$, and the whole ensemble of t trajectories is denoted as $\pi(\tau) = (\pi_1(\tau), \ldots, \pi_t(\tau))$.

Adversary Model: Our attacker will be a computationally unbounded active threshold adversary named Eve. That is, given a network $G = (V, E)$, with a sender s and receiver r (both in V), the adversary can compromise up to $k \leq |V \setminus \{s, r\}|$ *nodes* in G (thanks to QKD, an activity on any of the links would be detected anyway). Moreover, Eve knows all relevant protocol specification and the network topology, and is not bound to follow the protocol. A weaker notion is assuming her to stick passively to the protocol in order to extract secret information. We call this behavior *passive*, as opposed to an *active* adversary, as described previously and refined later in Section VII. Throughout the remainder of this work, the adversary's threshold will be denoted as k.

Security Model: Our notion of security is based on the concepts used in [11]. We need some notation: a general *protocol* Π is an interactive process between a sender and a receiver. In the course of Π, Alice exchanges a set $C = \{C_1, \ldots, C_n\}$ of messages with Bob in order to secretly transmit a message $M \in \{0, 1\}^*$ of entropy $H(M)$. The full set C is called the protocol's *transcript*. A subset $\mathrm{adv}(M) \subseteq \{C_1, \ldots, C_n\}$ of the transcript obtained by eavesdropping of the adversary is called his *view* in the protocol Π (a closely related equivalent notion is found used in [13]).

Definition III.1. *Let $\varepsilon > 0$, and let Π be a message transmission protocol. We call a protocol* ε-secure*, if the following two conditions are satisfied:*

1) *$H(M|adv(M)) \in \{0, H(M)\}$ and*
2) *$Pr[H(M|adv(M)) = 0] \leq \varepsilon$,*

i.e., the adversary can disclose M with a chance of at most ε. We call the protocol Π efficient*, if the size of the transcript, i.e., $\sum_{i=1}^{n} |C_i|$, is polynomial in the size of the message M, the size of underlying network (in terms of nodes), and $\log \frac{1}{\varepsilon}$. A protocol that is ε-secure for any $\varepsilon > 0$ is said to enjoy* perfect secrecy.

It is easy to see that if a protocol is ε-secure with $\varepsilon < 2^{-|M|}$, then simply guessing the message is more likely than breaking the protocol itself.

IV. QKD-Based Multipath Transmission

Multipath transmission pursues a simple idea: having t paths from s to r that are node-disjoint, the sender can transmit a message m by first putting it through a (t', t)-secret sharing (Shamir's for instance), giving the shares $s_1, \ldots, s_t$ and sending each share over its own (distinct) path to r. The adversary is successful if and only if he catches at least t' shares. Obviously, the scheme is unconditionally secure if $t' > k$ (where k is the adversary's threshold), but in addition, we require full knowledge of the topology, and assured delivery over the chosen disjoint paths. The general interplay between network connectivity and unconditional security has been studied extensively (cf. [14], [13], [3]). However, common to all these results is the implicit assumption of secure and reliable routing. That is, most existing multipath transmission regimes prescribe a fixed set of chosen node-disjoint paths. These paths are assumed stable and unchanged over the duration of a transmission; the adversary might intercept the paths but cannot redirect them. Hence, our goal in the next section is to find out whether or not unconditional security can be retained if the paths are not reliably under the sender's control. In other words, what happens if the adversary indirectly fiddles with the routing?

V. A Markov-Chain Routing Model

To simplify technicalities, let us assume a *synchronous* forwarding regime, i.e., the nodes simultaneously forward their packets at fixed times. This permits us to use a discrete time variable $\tau \in \mathbb{N}$. This assumption is not too restrictive, since even an asynchronous forwarding regime can be reasonably approximated by choosing a small unit of time and letting some nodes remain occasionally inactive in some steps.

Consider an arbitrary but fixed trajectory i among an ensemble of t independent trajectories in the following. It is well known from the theory of Markov chains that the state of the i-th chain at time $\tau \in \mathbb{N}$ is governed by $\pi_i(\tau) = P^{\tau} \cdot p_i(0)$, where P is the transition matrix. Our chain has only a *single absorbing state*, which is the receiver's state r (the receiver will surely not pass on his message any further). Furthermore, it can be assumed irreducible, because if it were not, then there would be at least two nodes u, v in the network whose chance of getting a packet from u to v is zero, so they could never communicate.

We write H_{jA} for the time (measured in *hops*) that it takes a trajectory to get from node j to any of the target nodes in the set $A \subseteq V$,

$$H_{jA} = \min\{\tau \geq 0 : X(\tau) \in A | X(0) = j\}.$$

The probability h_{jA} of the chain ever reaching A from j is therefore $h_{jA} = \Pr[H_{jA} < \infty]$, and the family $(h_{jA}; j \in V)$ is the smallest non-negative solution of the equation system

$$h_{jA} = \sum_{i \in V} p_{ji} h_{iA}, \tag{1}$$

where $h_{jA} = 1$ for all $j \in A$ and p_{ji} is the probability of passing from node j onwards to node i (see [18, p.123] for details). Writing down this system for, say 5 equations with $A = \{1, 3\}$, we get (after some minor algebra),

$$\begin{array}{rcccc} -p_{21} - p_{23} & = & (p_{22} - 1)h_{2A} & + & p_{24}h_{4A} \\ -p_{41} - p_{43} & = & p_{42}h_{2A} & + & (p_{44} - 1)h_{4A}, \end{array}$$

where we additionally substituted $h_{rA} = 0$, as r is the only absorbing state of our chains. Let us write (in a slight abuse of notation) $P_{-R,-C}$ to denote the matrix P with all rows in R and all columns in C *deleted*. Similarly, we use the notation $P_{R,C}$ to denote the matrix P only with the rows in R and columns in C retained. To ease notation, let us put $Q := P_{-r,-r}$, i.e., Q is P without the r-th row and column. If I is the identity matrix, and $\mathbf{1}$ is the vector of all 1's, then the above equation system takes the compact form

$$-Q_{-A,A} \cdot \mathbf{1} = (Q_{-A,-A} - I)h_A, \tag{2}$$

where h_A is the family $(h_{1A}, h_{2A}, \ldots, h_{rA})$, excluding $h_{rA} = 0$ and $h_{jA} = 1$ for all $j \in A$. In order to have the values h_j for $j \neq r$ and $j \notin A$ well-defined, we ought to show that $(Q_{-A,-A} - I)$ is invertible. This is our first

Lemma V.1. *Let P be a stochastic matrix of an irreducible Markov-chain with the state space V and exactly one absorbing state $r \in V$. Select any set of states $A \subset V$ with $r \in A$, and let $Q = P_{-A,-A}$ be the submatrix of P that describes transitions between states outside of A. Then $Q - I$ is invertible.*

Proof: Partition the state set V into $V_1 = A$ and $V_2 = V \setminus A$, then $r \in V_1$ and Q describes transitions *within* V_2. For each $v \in V_2$, write $\pi_{V_2}(\tau, v)$ for the chance of the chain being in state v after τ steps. From the theory of Markov-chains, we know that the vector $\pi_{V_2}(\tau) = (\pi_{V_2}(\tau, v))_{v \in V_2}$ is given by $\pi_{V_2}(\tau) = Q^\tau \pi_{V_2}(0)$. As the chain is irreducible, we will eventually reach r from *any* state in V_2, and since r is absorbing, this means that $Q^\tau \to 0$ as $\tau \to \infty$. Now, put $(Q - I)x = 0$. Then $Qx = x$ and on iterating $Q^\tau x = x$. As $\tau \to \infty$, $Q^\tau x = x \to 0$, so $Q - I$ is invertible. ∎

Lemma V.1 helps constructing a formula giving us the chance that exactly l trajectories pass through a given area $A \subseteq V$ that is under the adversary's control. We can solve the system (2) for any given set A and see whether it is passed with certainty. Similarly as for the binomial distribution, we can ask for the probability of a subset of l trajectories hitting A within finite time, with the remaining ones never reaching A. The probability we are after is the sum over all subsets of size l. Formally, we have

Proposition V.2. *Let a graph $G = (V, E)$ be given, and assume a random walk of t trajectories starting at nodes $1, 2, \ldots, t$. For a given $A \subseteq V$, the chance of l trajectories passing through A is given by*

$$p(A, l) = \sum_{\substack{M \subseteq [1:t] \\ |M| = l}} \left[\prod_{i \in M} h_{iA} \prod_{i \in ([1:t] \setminus M)} (1 - h_{iA}) \right],$$

where the vector $(h_{iA})_{i \in V}$ is calculated by putting $h_{rA} = 0, h_{jA} = 1$ for all $j \in A$, and calculating the remaining probabilities by solving (2)*. Here, $[1:t]$ is a shorthand notation for the set $\{1, 2, \ldots, t\}$.*

VI. Secrecy Against Passive Adversaries

According to Proposition V.2, the adversary will not learn anything unless he conquers some set A that is passed by sufficiently many, say l, trajectories. Consequently, his best strategy is attacking the set with maximum likelihood of seeing sufficiently many trajectories. It follows that the most vulnerable subset of nodes in the network is

$$A^* = \operatorname*{argmax}_{A \subseteq V} \Pr[l \text{ trajectories traverse } A] = \operatorname*{argmax}_{A \subseteq V} p(A, l). \tag{3}$$

The following result is an immediate consequence of the above discussion:

Theorem VI.1. *A network with a routing regime described by a transition matrix P can provide perfect secrecy without pre-shared end-to-end secrets, if and only if for some integer $l \geq 1$, we have $p(A, l) < 1$ for all $A \subseteq V$ that the adversary can compromise.*

Proof: Assume that $p(A, l) < 1$ for any set A and choose $\varepsilon > 0$ arbitrarily small. Put the message through a (n, n) secret sharing scheme, giving the shares $s_1, s_2, \ldots, s_n$. Send each s_i over l paths to the receiver. The adversary is successful if and only if he catches all shares, but the chance for this to happen decays exponentially fast as $p(A, l)^n \to 0$ as $n \to \infty$. It remains to choose n sufficiently large so as to have $p(A, l)^n < \varepsilon$.

Conversely, if $p(A, l) = 1$ for some set A, then there is no way to avoid the adversary when transmitting something over the network. Hence, secret communication is impossible. ∎

Despite this maximum likelihood optimization problem being sound, it is yet infeasible to evaluate as the number of subsets to test is exponential (in the adversary's threshold). We shall therefore set out to find sufficient criteria that are easier to test.

For a 1-passive adversary, we have the following test:

Theorem VI.2. *Let $t = |nb(s)| \geq 1$ count the sender s's neighbors. If, for each $v \in V$, we have $\sum_{i=1}^{t} h_{iv} < t$, then the network provides perfect secrecy against a 1-passive adversary.*

Proof: Put the secret message through a (t,t)-secret sharing and let each share take its own individual path through the network (i.e., do a random walk according to the transition matrix P). With the random indicator variable

$$\mathbb{I}_{i,j} := \begin{cases} 1, & \text{if } h_{ij} > 0 \\ 0, & \text{otherwise,} \end{cases}$$

the number of trajectories passing through a node $v \in V$ is given by $N_v := \sum_{i=1}^{t} \mathbb{I}_{i,v}$, and its expected value is $\mathsf{E}(N_v) = \mathsf{E}(\sum_{i=1}^{t} \mathbb{I}_{i,v}) = \sum_{i=1}^{t} h_{iv}$. The assertion now directly follows from Markov's inequality, since

$$\Pr[N_v \geq t] \leq \frac{\mathsf{E}(N_v)}{t} < \frac{t}{t} = 1,$$

which holds for all $v \in V$. The network thus provides perfect secrecy by Theorem VI.1. ■

Theorem VI.3. *Let $G = (V, E)$ be a graph, and let the sender and receiver be $s, r \in V$. Let the adversary be k-passive, i.e., up to k nodes in G can be compromised. For perfect secrecy, it is necessary that $|nb(s)| > k$. In that case, with $V^* := V \setminus \{s, r\}$, if*

$$\forall i \in nb(s) : h_{ij} \leq \frac{1}{ek} \;\; \forall j \in V^* \setminus \{i\}\,, \tag{4}$$

then the network provides perfect secrecy.

Proof: Without loss of generality, assume s's neighbors to be the nodes $\{1, 2, \ldots, t\}$, and put the secret message m through a (t,t)-secret-sharing scheme, transmitting the i-th share over the i-th neighbor of s (the remaining path of each is individual and determined by the network's transition matrix P). Observe that the adversary will not learn anything unless he gathers all t shares.

If $t \leq k$, then the adversary can "cut off" s from the rest of the network, thus reading all information conveyed by s, and perfect secrecy is impossible by Theorem VI.1.

Assume $t > k$ henceforth, so there exists at least one honest neighbor of s in every attack scenario. Let $A \subseteq V$ with $A = \{j_1, \ldots, j_k\}$ be a set of compromised nodes. The (mutually dependent) events $T_l^{j_i}$ for $i = 1, 2, \ldots, k$ occur when the trajectory starting off the node l reaches node j_i. For each (starting node) $l = 1, 2, \ldots, t$, we have

$$\Pr\left[T_l^{j_i}\right] = h_{lj_i} \leq \max\{h_{lv} | v \in V \setminus \{l, s, r\}\} \leq \frac{1}{ek}, \tag{5}$$

where the last inequality follows from our hypothesis. Since $\Pr\left[T_l^{j_i}\right] \leq \frac{1}{ek}$, then Lovász local lemma (symmetric version) implies

$$\Pr\left[\bigcap_{\nu=1}^{k} \overline{T_i^{j_\nu}}\right] > 0. \tag{6}$$

Protocol skeleton for secret and eff cient delivery of a message over an untrusted network.
Input: Message m, round number n and number t of shares per round.
Protocol steps for the sender:

1) Put m through a (n,n)-secret sharing, giving the shares $s_1, \ldots, s_n$.
2) For $i = 1, 2, \ldots, n$ do the following: put the i-th share s_i through a (t,t)-secret sharing, where $t = |\mathrm{nb}(s)|$, and transmit the j-th share of s_i over the j-th neighbor of s (cf. Theorem VI.3).

Figure 1. Multi-round multi-path transmission

In other words, the l-th trajectory has a positive chance of *evading* the set $\{j_1, \ldots, j_k\}$. Since inequality (5) holds independently of the particular j_i's, (6) is true for all these sets. If condition (5) holds for all $l = 1, 2, \ldots, t$, then in every attack scenario there is at least one trajectory with a positive chance of *not* passing through the compromised area in the graph. So, for every $A \subset V$ with $|A| \leq k$, it holds that $p(A, t) < 1$ and the network can provide perfect security by Theorem VI.1. ■

Eff ciency

Regarding the bandwidth demand, we require the overall network traff c (bit complexity) and round complexity to be polynomial in $\log \frac{1}{\varepsilon}$ for any chosen $\varepsilon > 0$. Assume the network satisf es the condition for perfect secrecy in Theorem VI.1.

Fix some $\varepsilon > 0$. We will prove the following transmission regime to enjoy eff cient bit- and round-complexity, i.e., polynomial efforts in $\log \frac{1}{\varepsilon}$. Let the secret message m be transmitted from s to r by virtue of the framework protocol shown in Figure 1. For a passive adversary with a threshold k, the number of shares t must be larger than k. The number n of rounds will be determined now.

Obviously, the attacker will not learn anything unless he gets all the information f owing over the network (due to the (n,n)- and (t,t)-sharings). Our task is proving n to be polynomial in $\log \frac{1}{\varepsilon}$ and the size of the network. For the proof, def ne an indicator variable for each round $i = 1, 2, \ldots, n$ via

$$\mathbb{I}_i = \begin{cases} 1, & \text{if the share } s_i \text{ was disclosed;} \\ 0, & \text{otherwise,} \end{cases}$$

so that $\mathbb{I}_i$ measures the adversary's success (in a binary scale) in the i-th round. By our hypothesis, Theorem VI.1 implies $\Pr[\mathbb{I}_i = 1] < 1$ for all rounds i and all sets of nodes that the adversary could have conquered (recall that the adversary is k-passive). Put $\rho := \max_{i=1,2,\ldots,n} \Pr[\mathbb{I}_i = 1]$, then $\rho < 1$. Since $0 \leq \mathbb{I}_i \leq 1$ for all i, the f rst moment $\mathsf{E}(\mathbb{I}_i)$ exists

and $\mathbb{I}_i$'s deviation from its mean is bounded by $-1 \leq \mathbb{I}_i - \mathsf{E}(\mathbb{I}_i) \leq 1$ for all i. Define $S := \sum_{i=1}^{n} \mathbb{I}_i$, then since $\mathsf{E}(\mathbb{I}_i) \leq \rho$, we get $\mathsf{E}(S) = \sum_{i=1}^{n} \mathsf{E}(\mathbb{I}_i) \leq n\rho$. Moreover, $S - \mathsf{E}(S) \geq S - n\rho \geq \tau$ for some τ to be fixed later. Application of a variant of Hoeffding's inequality (with relaxed independence constraints; see [19]) gives

$$\Pr[S - n\rho \geq \tau] \leq \Pr[S - \mathsf{E}(S) \geq \tau] \leq \exp\left(-\frac{\tau^2}{2n}\right)$$

Since $\frac{1}{n}S \geq \min_i \mathbb{I}_i$, we can choose τ to satisfy $\frac{\tau}{n} \leq \min_i \mathbb{I}_i - \rho \leq \frac{1}{n}S - \rho$. So we can continue the chain of inequalities on the left-side as

$$\Pr\left[\min_i \mathbb{I}_i - \rho \geq \frac{\tau}{n}\right] \leq \Pr[S - n\rho \geq \tau] \leq \exp\left(-\frac{\tau^2}{2n}\right),$$

and by taking $\delta := \frac{\tau}{n}$ we conclude that

$$p := \Pr\left[\min_i \mathbb{I}_i \geq \rho + \delta\right] \leq \exp\left(-\frac{n\delta^2}{2}\right)$$

for all $\delta \geq 0$. By construction, the adversary is successful if and only if $\mathbb{I}_i = 1$ for all rounds $i = 1, 2, \ldots, n$, or equivalently, $\min_i \mathbb{I}_i = 1$. Choosing $\delta := 1 - \rho > 0$, the number n of rounds until $\Pr[\min_i \mathbb{I}_i \geq \rho + \delta = 1] < \varepsilon$ is achieved comes to $n \in \mathcal{O}\left(\log \frac{1}{\varepsilon}\right)$. The bit-complexity is $n \cdot t \cdot |m|$, where $|m|$ is the length of the message, and as such in $\mathcal{O}\left(|m| \cdot |\mathrm{nb}(s)| \cdot \log \frac{1}{\varepsilon}\right)$, i.e., polynomial in the network size and $\log \frac{1}{\varepsilon}$. Summarizing the discussion, we have proved

Theorem VI.4. *If a given network provides perfect secrecy according to Theorems VI.1, VI.2 or VI.3, then there is an efficient protocol achieving this.*

VII. Secrecy Against Active Adversaries

It is easy to see that the results of Section VI no longer hold when the adversary becomes active. Picking up our line of arguments that led to Theorem VI.4, the adversary can destroy the message simply by fiddling with one of its shares. Equally obvious is a quick-fix by attaching a checksum to the message, which lets the receiver *detect* (not necessarily correct) this kind of manipulation upon combining the incoming shares. For later reference, we state this as remark:

Remark VII.1. *One can prove (see [20]) that if error detection is required reliably with a probability of at least $1 - \varepsilon$ for $\varepsilon > 0$, then the size of the share grows by at least $\log \frac{1}{\varepsilon}$ additional bits. So, attaching an appropriate checksum to the secret before sharing it is close to optimal in terms of additional overhead.*

To ease technicalities in the following, let us distinguish two different forms of activity for the adversary:

1) he participates only in the protocol, but is allowed to actively deviate from it as he wishes,
2) he participates in the protocol and additionally runs parallel sessions over the network.

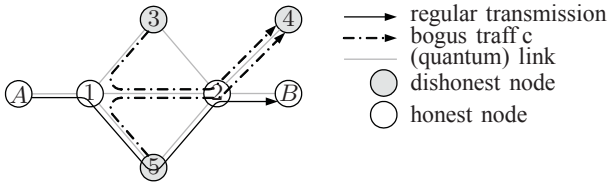

Figure 2. Path alteration via bogus traffic

The first kind of active adversary is easer to deal with, since his activity is basically focused on active modifications to the messages that he gets to pass his nodes. Modifying the routing information in order to redirect these messages differently than intended by Alice will not help him learn anything (simply because the packet is in his possession already). On the other hand, he cannot redirect packets that he does not get to see in order to acquire them. Theorem VII.2 is concerned with security against such an attacker.

This is the major difference to the second kind of adversary, who can attempt to redirect packets by intentionally congesting links that he does not directly control. To illustrate the problem, consider the simple topology displayed in Figure 2. In this scenario, Alice wishes to transmit a message to Bob, which would be possible over the path over the nodes 1 and 2. However, even though the adversary does not control this path, he can nevertheless congest the link from 1 to 2 with bogus traffic so as to enforce re-routing over node 5 (or node 3), which is under his control.

Testing whether this kind of attack is possible is highly nontrivial, because we now face an adversary who can manipulate the graph topology, while only being constrained by the link capacities. For instance, the adversary could look for a path cover of the graph G that respects the existing bandwidth limits. Indeed, even without the bandwidth restriction, the problem of finding a minimal path cover of this kind on a general graph is NP-complete, but becomes solvable in linear time for certain classes of graphs (see e.g., [21]). On the contrary, the adversary could as well compute a maximal multi-source multi-sink flow between his nodes in order to maximally congest the network. Abusing the Ford-Fulkerson approach, he could choose the flow-augmenting paths in a way so as to use as many links between honest nodes as possible. However, up to now, this is a mere heuristic and not yet a provably optimal attack strategy. Even worse, from the perspective of the honest parties, one would have to compute such a flow for all scenarios of attacking, which again boosts the computational efforts for analysis far out into infeasibility. The most trivial way of fixing this is to abandon all kinds of rerouting due to congestions and designing the relay nodes as mere queues, where messages are temporarily stored.

It appears that guarding against such kind of attack is

more a matter of congestion control. Consequently, going into more detail is thus beyond the scope of this work, as we did not presume any particular congestion control or routing scheme here.

However, if an active adversary of the first kind is assumed (i.e., bound to only manipulating, inserting or blocking of messages), we can reformulate our previous results accordingly to remain valid. The basic trick is to use the following property of secret-sharing and Reed-Solomon codes. It is well-known that Shamir's (t, n) secret sharing corresponds to a Reed-Solomon code of length n with t information words (cf. [22]). Consequently, we can recover from up to $\lfloor (n-t)/2 \rfloor$ modified shares by virtue of the Welch-Berlekamp algorithm [23] (in fact, this technique is standard in multipath transmission; cf. [3] for instance). From the error correction capacity of the code and the condition that the adversary should have less than t shares in his possession, we easily deduce the (also well known) fact that secret-sharing is robust against an active adversary with a threshold less than $n/3$. Hence, up to a third of the shares (i.e., paths) can be compromised and packets along them can be modified and the message remains concealed and intact upon reconstruction. This is the basic fact that yields to straightforward generalizations of the results in Section VI stated in the following.

Formally, a $(t, n)-$secret-sharing scheme is secure against a k-active adversary as long as its threshold k satisfies $k < \frac{n}{3} < t$ In analogy to Theorem VI.2 we get the following criterion for a 1-active adversary:

Theorem VII.1. *Let $t = |nb(s)|$ count the sender s's neighbors. If, for each $v \in V$, we have $\sum_{i=1}^{t} h_{iv} < \frac{t}{3}$, then the network provides perfect secrecy against a* 1*-active adversary.*

Proof: Put the secret message through a (t, t)-secret sharing and let each share take its own individual path through the network. With the random indicator variable

$$\mathbb{I}_{i,j} := \begin{cases} 1, & \text{if } h_{ij} > 0 \\ 0, & \text{otherwise,} \end{cases}$$

the number of trajectories passing through a node $v \in V$ is given by $N_v := \sum_{i=1}^{t} \mathbb{I}_{i,v}$, and its expected value is $\mathsf{E}(N_v) = \sum_{i=1}^{t} h_{iv}$. An active modification is possible if at least $t/3$ shares get compromised, so we can use Markov's inequality to conclude

$$\Pr[N_v \geq t/3] \leq \frac{\mathsf{E}(N_v)}{t/3} < \frac{t/3}{t/3} = 1,$$

which holds for all $v \in V$. The network thus provides perfect secrecy since the adversary can not intercept enough shares. ■

Unfortunately, Theorem VI.3 no longer holds for active adversaries. Still, we can use it to guard a transmission against an active adversary as well, yet we need some additional requirements on the network.

In fact, multipath transmission protocols usually hinge on the sender's ability to choose the paths in a way that he likes. This assumption is rarely stated explicitly (as for instance, it is used in [3] or [13]), but nevertheless of crucial importance. By specification [24, p.19], the internet protocol (IP) provides the following feature: the sender of a message can prescribe the list and order of intermediate nodes over which the packet must be forwarded until it reaches the receiver. The *Session Initiation Protocol* (SIP), specified in [25], defines a functional strict source routing mechanism, meaning that the sender can choose his relay nodes and no other nodes must be visited during a transmission. For our purposes, a weaker notion is sufficient, namely the symmetric answer property, which is introduced here:

Definition VII.1 (Symmetric Answer Property (SAP))**.** *Let a message transmission be over the relay nodes $v_1, v_2, \ldots, v_n$. If each relay node keeps the so-defined channel open for a subsequent response (e.g., an acknowledge message), i.e., the receiver can respond over the path $v_n, v_{n-1}, \ldots, v_2, v_1$, then the network is said to satisfy the* symmetric answer property.

In fact, it is this particular feature that is implicitly used in recent work on multipath transmission such as [3] or [13], although it is not explicitly stated there (usually, it is implicitly assumed in a sloppy form as saying that "the sender responds over the same channel over which he received the information"). Here, we will explicitly use this to construct a communication protocol that enjoys robustness against an active adversary. In the light of the previous discussion, this appears to be a mild and reasonable assumption, as it is included and supported by the common technological standards for data transmission, as referenced above.

Theorem VII.2. *Let $G = (V, E)$ be a graph, and let the sender and receiver be $s, r \in V$. Let the adversary be k-active, i.e., up to k nodes in G can be compromised. For perfect secrecy, it is necessary that $t = |nb(s)| > 3k$. If the network satisfies condition* (4) *and the symmetric answer property (SAP), then it permits perfect secrecy and resilience against an active adversary of the first kind.*

Notice that only the necessary condition has changed, but the sufficient condition was only augmented by assuming the SAP, since the line of arguments in the proof of Theorem VI.3 can no longer be used to prove that the adversary gets to see at most a third of the trajectories (as would be required). Nevertheless, we can use Theorem VI.3 to construct a protocol that guards us against active adversaries too.

The proof of Theorem VII.2 will partially rely on the robustness of secret sharing against modification of shares. The required result along these lines is summarized as

follows:

Lemma VII.3. *Let a general (u, v)-secret-sharing be given, and assume that the adversary has modified up to k shares. Then,*

- *if $0 \leq k < v/3 < u$ then there is no harm; all errors can be corrected.*
- *if $v/3 \leq k < u$, then the message cannot be disclosed by the attacker, but he can still thwart a correct reconstruction.*
- *if $u \leq k$, then the attacker can disclose the message without notice.*

This fact is quite well-known (cf. [26]), yet proofs can be found in [27].

Proof of Theorem VII.2: Without loss of generality, assume s's neighbors to be the nodes $\{1, 2, \ldots, t\}$, and put the message m through a (t, t)-secret-sharing scheme, transmitting the i-th share over the i-th neighbor of s.

If $t \leq 3k$, then the adversary can gain enough information to know and perhaps modify (replace) the message already after one transmission, hence $t > 3k$ is necessary for perfect security.

In the following, suppose that the sender transmits a message m along with a checksum $H(m)$ over node-disjoint channels to the receiver. The checksum (e.g., a cryptographic CRC; cf. [28]) will provide an additional mean of detecting manipulations once the error correction (and detection) capabilities of the encoding failed (cf. Remark VII.1).

Once the transmission has started, the proof of Theorem VI.3 ultimately concludes that at least one trajectory will bypass the adversary on its way from the sender to the receiver. The active adversary can either modify or not modify the shares that he intercepts. Not modifying anything literally means a passive adversary, which has been covered in the course of Theorem VI.3, hence we consider an active attacker in the following.

The protocol described now establishes a shared end-to-end secret between a sender and receiver. First, we transmit a (random) message m along with a cryptographic checksum $H(m)$ via a (t, t)-secret-sharing and (hopefully) disjoint paths over the network, and act as if the adversary were passive. This transmission process is repeated for several rounds, each of which yields a partial key K_i (for the i-th round) that we can use (e.g., concatenate and hash) to distill the final key for communication (e.g., to be used as a one-time pad over a classical, perhaps insecure, channel).

We have *two* mechanisms of error detection: the inherent error correction that comes with the secret-sharing (via the Welch-Berlekamp-Algorithm in case of Shamir's polynomial secret sharing), and the cryptographic checksum after the reconstruction. Let us abbreviate the error-correction as EC and the checksum verification as CV hereafter. Each of these can (independently) yield a positive or negative outcome, giving us four cases to distinguish in the i-th round:

1) EC points out no errors and the CV confirms the checksum: in that case, the adversary (with high probability) has either learnt nothing or everything, since the only case in which no error is determined by the error correction algorithm occurs when the adversary managed to replace *all* shares. If that happens, it is easy to replace the hidden secret by something else along with a matching checksum (hence the CV can be expected to return positive).
   Anyway, since there is a positive chance that the adversary has indeed discovered the secret, the receiver will discard any results in this case.
2) EC points out no errors but the CV fails: in that case, the adversary managed to replace all the shares, but has used a secret that is inconsistent with the reconstructed checksum. This would technically point out a manipulation while the adversary would have been capable of avoiding this detection. So, there is no point in acting like this, and this case is to be treated equally as case 1.
3) EC points out errors, but the CV confirms the checksum: in this case, the adversary must have managed to replace sufficiently many shares (cf. Lemma VII.3) to trick the error correction into wrongly indicating correct shares as malicious. Yet at least one original share has not been intercepted, because the error correction pointed out at least one error. Since we do not know which share is the correct one, but know that there must be at least one, we take the protocols output as the bitwise exclusive-or of all shares $s_1, \ldots, s_t$, that is we create
   $$K_i := s_1 \oplus s_2 \oplus \cdots \oplus s_t,$$
   knowing that the partial key K_i is entirely unknown to the adversary since at least one share in it acts like a one-time pad encryption key.
4) EC points out errors and the CV fails: in this case and by Lemma VII.3, at least $t/3$ but less than t shares must have been manipulated, since the adversary was unable to replace the secret consistently. In that case, as before, we use the bitwise XOR of all shares to distill the partial key K_i as the output of round i.

This protocol is repeated for several rounds until a sufficient amount of key-material (partial keys $K_1, K_2, \ldots$) has been produced. Notice that the actual information m transmitted through the secret-sharing is of no real value, and merely serves to create a redundancy scheme that we can use to detect a manipulation.

More importantly, observe that if case 1 occurs, then the adversary can easily make the protocol output to look like any of the other cases occurred. If this happens, then he has gained the correct information c that the receiver will use. However, the proof of Theorem VI.3 implies that with

a positive probability, cases 2, 3 or 4 must occur, hence he cannot entirely intercept the communication.

It is crucial for the sender to get notif ed in which rounds the protocol output has been discarded by the receiver. He does this by telling the sender which of the four cases above occurred, and sends this information identically back over all paths over which he received the shares originally. If the adversary missed one of the shares, then this channel will safely deliver the notif cation to the sender, thanks to the symmetric answer property. Hence, upon any two mismatching notif cations, the sender will automatically be notif ed of the attack attempt. Even in case 1, if the adversary managed to intercept all channels, he can either replace the notif cations or remain passive. In the former case, he would indicate an attack while he could have convinced the sender that there was no attack at all, so there is no point in acting like this. However, if an attack like in case 1 was successful, then the receiver would discard key-material that the sender would use, making the two end up using different (and hence useless) keys.

This kind of person-in-the-middle situation can be detected by letting the sender and receiver sacrif ce some key-bits for public comparison on a possibly insecure channel. Suitable protocols for this are well-known from quantum cryptography and we will therefore not go into further details here. If the two keys turn out different, then both discard their key-material and rerun the protocol from scratch. ■

As far as eff ciency of the transmission in the presence of an active adversary is concerned, the transmission's eff ciency is basically determined by the chance of at least one trajectory avoiding the adversaries premises. While there is a positive chance that this will happen eventually (thanks to Theorem VI.3), the number of repetitions until this occurs suff ciently often, is diff cult to determine without knowledge of the precise likelihoods. These can be obtained from simulations, but in any case, the protocol is to be repeated until the f nal verif cation indicates a correct and useable key. Nevertheless, in the next section, we use an example to show how the number of repetitions can be computed at least partially.

VIII. Application to Quantum Networks

It is important to emphasize that Theorems VI.1, VI.2 and VI.3 *should not* directly be applied to the communication network at hand. Instead, we are interested in estimating the harm that any deviation from a prescribed routing strategy causes. Going back to multipath transmission, our goal is using the results from the previous section to classify a given network as (in)secure under the assumption of random detours that a packet takes upon local congestions or empty local quantum-key-buffers.

We illustrate the application of Theorem VI.3 by using a simple example, which shall demonstrate the general line of reasoning. Take the network shown in Figure 3, with

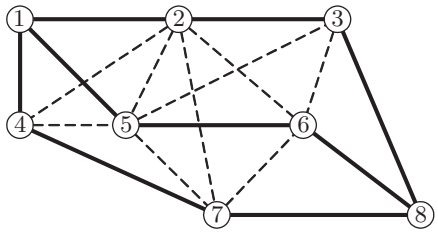

Figure 3. Example multipath transmission from 1 to 8

each link secured by means of QKD. Alice (node 1) performs a multipath communication over three disjoint channels $\rho_1 = (1 \to 2 \to 3 \to 8), \rho_2 = (1 \to 5 \to 6 \to 8), \rho_3 = (1 \to 4 \to 7 \to 8)$ (shown bold) to Bob's node 8. Assume that each node does the packet forwarding reliably, up to some chance of α for the packet to def ect from the prescribed route. Thus, assuming stochastic independence for the sake of simplicity, with probability $1 - \alpha^{\text{length}(\rho_i)-2}$, the packet will travel over ρ_i as desired. Notice that any path is accessible from any other, and that an adversary will surely not waste resources by attacking anywhere else than on the chosen paths. Hence, we can create an abstract model for such a multipath transmission by restricting the focus on whether the packets travel as desired (likelihood determined by the reliability of routing, i.e., the probability of the packet not deviating from its prescribed route), or whether they take detours (should happen with a small chance only) that could yield to intersecting paths and disclosure of the secret message.

For the analysis of a general network $G = (V, E)$ under a multipath transmission scenario, we therefore consider the auxiliary graph $G' = (V', E')$: let $\rho_1, \ldots, \rho_t$ be paths in G, then each of these becomes a node in G', which is connected to the sender and receiver, so put $V' := \{\rho_1, \ldots, \rho_t\} \cup \{s, r\}$. Attacking elsewhere than on the paths $\rho_1, \ldots, \rho_t$ is less paying for the adversary than compromising the paths themselves, so we may safely disregard any nodes in the network that are not on a chosen path. Also, assume that a packet can jump from any path to any other, so the nodes $\rho_1, \ldots, \rho_t$ form a clique. Finally, each path ρ_i is connected to the receiver r in a one-way manner, as the receiver is absorbing and will not pass anything further. Similarly, the sender is (one-way-)connected to all his chosen paths, though these transitions are of no further interest, since an accidental jump from a path back to the sender can trivially be corrected by the sender putting the packet back on its correct path. The set of edges therefore comes to $E' = \{\rho_1, \ldots, \rho_t\}^2 \cup \{(\rho_i, r), (s, \rho_i) | i = 1, 2, \ldots, t\}$. The resulting transition graph for the example is depicted in Figure 4, with arrows indicating possible state transitions.

The topology of the auxiliary graph G', excluding the transitions from s to each ρ_i (for obvious reasons) def nes the Markov-chain on which we can invoke the results from

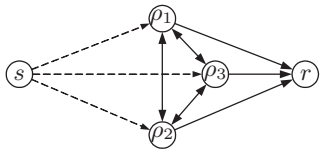

Figure 4. Auxiliary graph G' describing state transitions

Section VI. For the analysis, it remains to specify the following likelihoods:

- $\Pr[\rho_i \to r]$: with the parameter α as above, this is $\Pr[\rho_i \to r] = 1-\alpha^{\text{length}(\rho_i)-2}$. Notice that several events of node failure are not necessarily independent, and correlations among these must be considered in a more accurate (perhaps more realistic) model.
- $\Pr[\rho_i \to \rho_j]$: this quantity depends on the particular chances of jumping from a node on ρ_i to any node on ρ_j, and must be worked out individually for the network at hand. For the sake of simplicity and illustration, we assume an equal likelihood of jumping on any other path once ρ_i is left. For the example, we take $\Pr[\rho_i \to \rho_j] = \frac{1}{t-1}(1 - \Pr[\rho_i \to r])$.

Since the jumps from the sender to each of his chosen paths are uninteresting, we do not need to model the corresponding transition probabilities, nor must these appear in the transition matrix of the Markov-chain. These links are merely included to have G' consistent with our criteria, and are therefore shown dashed.

With $\alpha = 0.01$, we end up finding the transition matrix:

$$P = \begin{array}{c} \\ \rho_1 \\ \rho_2 \\ \rho_3 \\ r \end{array} \begin{array}{c} \begin{array}{cccc} \rho_1 & \rho_2 & \rho_3 & r \end{array} \\ \begin{pmatrix} 0 & 0.01 & 0.01 & 0.98 \\ 0.01 & 0 & 0.01 & 0.98 \\ 0.01 & 0.01 & 0 & 0.98 \\ 0 & 0 & 0 & 1 \end{pmatrix} \end{array}$$

Now, we can use Theorem VI.3 on this matrix to see that the network is indeed secure against a 2-passive adversary: with $V^* = \{1, 2, 3\}$ and by solving (2) for $A = \{1\}, \{2\}, \{3\}$, we find $h_{ij} = \frac{1}{99} < \frac{1}{2e} \approx 0.184$, for each $i, j \in V^*, j \neq j$. It follows that the network remains secure even under much less reliable routing. Indeed, we can tolerate up to $\alpha \approx 0.155$, i.e., a more than 15% chance of the packets becoming re-routed via indirect eavesdropping or congestion control. Finally, Theorem VI.4 tells that resilience against such incidents can be retained efficiently.

In order to illustrate Theorem VII.2, let us consider a network whose auxiliary graph has a similar topology as shown in Figure 4, but has 7 paths connecting Alice and Bob. The adversary is 2-active ($k = 2$), so that the necessary condition of more than $3k = 6$ neighbors is satisfied. Moreover, let the reliability of the network transmission be $\alpha = 80\%$, i.e., there is a chance of roughly 4% for the packet jumping from one path to another. Then, condition (4) is satisfied and the network provides perfect secrecy against a 2-active adversary by Theorem VII.2. With the concrete figures in hand, we can even compute the required number of protocol repetitions: it is the precise lower bound to the strictly positive probability (6) for a trajectory to bypass the adversary's servants. The sought bound is provided by the asymmetric version of the Lovasz local lemma from which the symmetric version of the Lovasz local lemma can be concluded. We spare the details for brevity, and draw the bound

$$\Pr\left[\bigcap_{\nu=1}^{k} \overline{T_i^{j_\nu}}\right] \geq \left(1 - \frac{1}{k+1}\right)^k$$

from the asymmetric (general) version of the lemma, where $T_i^{j_\nu}$ is the event of the l-th trajectory visiting the adversarial node ν starting from the sender's neighboring node i. In our case with $k = 2$, this bound evaluates to 0.44, so that there is quite a good chance for the adversary to miss at least one trajectory. This means that there is a $1 - 0.44 \approx 55.55\%$ chance for cases 2, 3 or 4 in the proof of Theorem VII.2 to occur. Since case 2 will never be observed for a reasonably acting adversary, we have a chance of $p = 0.55$ to distill key material in each round thanks to the remaining cases 3 and 4. So, the expected amount of key-material comes to $\approx 0.55n$ Bit for n rounds, and the required number of repetitions can be computed from the required amount of key-material. Still, this does not mean that case 1 is impossible and the adversary could have tricked the sender and receiver into thinking that cases 3 or 4 apply in some rounds. So, the final decision whether or not to use the key is up to the public comparison. The number of repetitions upon failure of this last step is geometrically distributed, yet the distribution parameter, namely the required success probability of a single Bernoulli trial (which is nothing else than a protocol execution), unfortunately cannot be computed from the given information.

IX. Conclusion

We have obtained simple criteria for protection against passive and active adversaries, if the activity is constrained to modifications and no bogus traffic. In case of coincidental redirection of packets along alternative routes, we have shown sufficient criteria for the transmission remaining secure in such cases. Based on these results, we have sketched how an active adversary can successfully be repelled by techniques of secret sharing, multipath transmission and error correction. Roughly speaking, our proposed protocols extend the purpose of QKD to create point-to-point secrets, to an application using QKD to establish end-to-end secrets.

Let us briefly review the results in chronological and condensed form. Our first main result is Theorem VI.1, which states that perfect secrecy is achievable if and only if the sender has a strictly positive chance to circumvent the adversary's corrupted nodes somehow. Theorem VI.2 and

Theorem VI.3 give sufficient conditions for this to happen, assuming a passive adversary listening. These conditions are derived from a Markov-chain model of the transmission. Basically, the analysis works by solving a linear equation system (1) for the vector of hitting probabilities h_{jA} (remember that h_{jA} is the chance for a packet starting off node j eventually reaching any node in the set A), where the hitting probabilities go directly into the criteria for secure communication. For solving (1), all we need are the likelihoods p_{ij} for a packet to travel to node j from node i. This is the description of the routing scheme as a Markov chain model. The model can of course be put to question, however, judging from the vast variety of routing strategies, routing table update procedures and possible flow control mechanisms, the Markov model appears to be sufficiently flexible to cover a large number of cases. If any of these sufficient criteria for perfect secrecy turns out satisfied, then Theorem VI.4 assures that the transmission is not only secure but also efficiently doable.

Regarding active adversaries, things are much more involved, and Figure 2 sketched a simple rerouting enforcement by inserting bogus traffic and exploiting load balancing and flow control. In alignment to our previous results, Theorem VII.1 transfers the known condition for 1-passive adversaries to its analogous form for 1-active adversaries. The transition from a 1-active to a k-active adversary calls for the additional hypothesis of symmetric answers, that is, the receiver must be able to reliably respond over the same channel over which he received a share in the first place. We call this the symmetric answer property, and Theorem VII.2 states that security against a k-active adversary can be achieved under roughly the same conditions as for a k-passive adversary, except for the additional assumption on symmetric answer channels. Unfortunately, all of these results refer to adversaries that do not run parallel sessions and particularly are not congesting links by bogus traffic. Defending the system against this kind of attack is beyond the capabilities of the given criteria and up to security systems linked to the flow and congestion control system within the quantum network.

Our results are only indirectly dependent on the quantum nature of the network, as the attack targets the multipath transmission regime only by *exploiting* general QKD properties. These are, moreover, independent of the particular QKD-implementation, and equally well apply to discrete or continuous quantum information encodings. In general, any successful denial-of-service attack, regardless of whether on a conventional or quantum line, can be used for indirect eavesdropping in the described form, as soon as secure multipath transmission is used.

This work is an explicit account for an adversary who turns the QKD eavesdropping detection against the network. If end-to-end security is set up by means of multipath transmission, then ”disconnecting” (by eavesdropping) otherwise adjacent nodes may enforce local re-routing of packets and in turn direct the information flow right into the adversary’s hands. We presented various sufficient criteria for a network to retain its security under indirect eavesdropping attacks by passive and certain active adversaries. Our results provide sufficient criteria to conclude that a network retains perfect secrecy under randomly compromised nodes and routes. Necessary criteria have not been given here, and are subject of future research.

Another interesting open problem is how to act against attacks involving bogus traffic in the quantum network. As has been demonstrated by a simple example scenario, an adversary can redirect traffic ”remotely” by cleverly overloading certain links and nodes (passive eavesdropping might as well yield such effects). Guarding a multipath transmission against this kind of attack is yet an open problem, and an interesting challenge of future research.

References

[1] S. Rass and S. König, “Indirect eavesdropping in quantum networks,” in *Proc. of the 5th Int. Conf. on Quantum-, Nano and Micro-technologies (ICQNM)*. Xpert Publishing Services, PO Box 7382. Wilmington, DE 19803, USA: Expert Publishing Services (XPS), 2011, pp. 83–88.

[2] Z.-S. Yuan, Y.-A. Chen, B. Zhao, S. Chen, J. Schmiedmayer, and J.-W. Pan, “Experimental demonstration of a BDCZ quantum repeater node,” *Nature*, vol. 454, pp. 1098–1101, 2008.

[3] M. Fitzi, M. K. Franklin, J. Garay, and S. H. Vardhan, “Towards optimal and efficient perfectly secure message transmission,” in *Proc. of 4th Theory of Cryptography Conf. (TCC)*, ser. LNCS 4392. Springer, 2007, pp. 311–322.

[4] C. Bennett and G. Brassard, “Public key distribution and coin tossing,” in *IEEE Int. Conf. on Computers, Systems, and Signal Processing*. IEEE Press, 1984, pp. 175–179.

[5] H. Wen, Z. Han, Y. Zhao, G. Guo, and P. Hong, “Multiple stochastic paths scheme on partially-trusted relay quantum key distribution network,” *Science in China Series F: Information Sciences*, vol. 52, no. 1, pp. 18–22, 2009.

[6] A. Varga and R. Hornig, “An overview of the OMNeT++ simulation environment,” in *Simutools ’08: Proc. of the 1st Int. Conf. on Simulation tools and techniques for communications, networks and systems & workshops*. ICST, Brussels, Belgium, Belgium: ICST (Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering), 2008, pp. 1–10. [Online]. Available: http://portal.acm.org/citation.cfm?id=1416222.1416290, last access: 06/19/2012.

[7] T. Schmitt-Manderbach, “Long distance free-space quantum key distribution,” Ph.D. dissertation, Ludwig–Maximilians–University Munich, Faculty of Physics, 2007.

[8] E. Diamanti, H. Takesue, C. Langrock, M. M. Fejer, and Y. Yamamoto, “Field trial of differential-phase-shift quantum key distribution using polarization independent frequency up-conversion detectors,” *Optics Express*, vol. 15, pp. 15 920–15 927, 2007.

[9] H. Xu, L. Ma, A. Mink, B. Hershman, and X. Tang, "1310-nm quantum key distribution system with up-conversion pump wavelength at 1550 nm," *Optics Express*, vol. 15, pp. 7247–7260, Jun. 2007.

[10] H.-K. Lo and H. F. Chau, "Unconditional security of quantum key distribution over arbitrarily long distances," *Science*, vol. 283, pp. 2050–2056, 1999.

[11] M. Franklin and R. Wright, "Secure communication in minimal connectivity models," *J. of Cryptology*, vol. 13, no. 1, pp. 9–30, 2000.

[12] M. Franklin and M. Yung, "Secure hypergraphs: privacy from partial broadcast," in *Proc. of the 27th annual ACM Symp. on Theory of computing*, ser. STOC '95. New York, NY, USA: ACM, 1995, pp. 36–44.

[13] Y. Wang and Y. Desmedt, "Perfectly secure message transmission revisited," *IEEE Trans. on Information Theory*, vol. 54, no. 6, pp. 2582–2595, 2008.

[14] M. Ashwin Kumar, P. R. Goundan, K. Srinathan, and C. Pandu Rangan, "On perfectly secure communication over arbitrary networks," in *PODC '02: Proc. of the 21st annual Symp. on Principles of distributed computing*. New York, NY, USA: ACM, 2002, pp. 193–202.

[15] R. Stewart, "RFC4960: Stream Control Transmission Protocol," http://tools.ietf.org/html/rfc4960, September 2007, last access: 05/17/2011.

[16] M. Pivk, C. Kollmitzer, and S. Rass, "SSL/TLS with quantum cryptography," in *Proc. of the 3rd Int. Conf. on Quantum, Nano and Micro Technologies*. IEEE Computer Society, February 2009, pp. 96–101.

[17] A. Mink, S. Frankel, and R. Perlner, "Quantum key distribution (qkd) and commodity security protocols: Introduction and integration," *Int. J. of Network Security & its Applications (IJNSA)*, vol. 1, no. 2, pp. 101–112, July 2009.

[18] D. Stirzaker, *Stochastic Processes & Models*. Oxford University Press, 2005.

[19] W. D. Smith, "Tail bound for sums of bounded random variables," URL: http://www.math.temple.edu/∼wds/homepage/works.html, April 2005, last access: 05/17/2011.

[20] M. Carpentieri, A. De Santis, and U. Vaccaro, "Size of shares and probability of cheating in threshold schemes," in *EUROCRYPT '93: Workshop on the theory and application of cryptographic techniques on Advances in cryptology*. Secaucus, NJ, USA: Springer-Verlag New York, Inc., 1994, pp. 118–125.

[21] R. Srikant, R. Sundaram, K. S. Singh, and C. P. Rangan, "Optimal path cover problem on block graphs and bipartite permutation graphs," *Theoretical Computer Science*, vol. 115, no. 2, pp. 351–357, July 19 1993.

[22] R. McElice and D. Sarwate, "On sharing secrets and Reed-Solomon codes," *Communications of the ACM*, vol. 24, no. 9, pp. 583–584, 1981.

[23] E. Berlekamp and L. Welch, "Error correction of algebraic block codes, US Patent Nr. 4,633,470,", 1986.

[24] J. Postel (ed.), "RFC 791 Internet Protocol - DARPA Internet Programm, Protocol Specif cation", Internet Engineering Task Force, September 1981, http://www.ietf.org/rfc/rfc791.txt, last access: 06/19/2012.

[25] M. Handley, H. Schulzrinne, E. Schooler, and J. Rosenberg, "SIP: Session Initiation Protocol," RFC 2543 (Proposed Standard), Internet Engineering Task Force, March 1999, obsoleted by RFCs 3261, 3262, 3263, 3264, 3265. [Online]. Available: http://www.ietf.org/rfc/rfc2543.txt, last access: 06/19/2012.

[26] T. Rabin and M. Ben-Or, "Verif able secret sharing and multiparty protocols with honest majority," in *STOC '89: Proc. of the twenty-f rst annual ACM Symp. on Theory of computing*. New York, NY, USA: ACM, 1989, pp. 73–85.

[27] S. Rass, "On information-theoretic security: Contemporary problems and solutions," Ph.D. dissertation, Klagenfurt University, Institute of Applied Informatics, June 2009.

[28] H. Krawczyk, "LFSR-based hashing and authentication," in *CRYPTO '94: Proc. of the 14th Annual Int. Cryptology Conf. on Advances in Cryptology*. London, UK: Springer, 1994, pp. 129–139.