npj | Quantum Information

www.nature.com/npjqi

ARTICLE OPEN

Check for updates

Tight finite-key security for twin-field quantum key distribution

Guillermo Currás-Lorenzo[1]✉, Álvaro Navarrete[2], Koji Azuma[3,4], Go Kato[4,5], Marcos Curty[2] and Mohsen Razavi[1]

Quantum key distribution (QKD) offers a reliable solution to communication problems that require long-term data security. For its widespread use, however, the rate and reach of QKD systems must be improved. Twin-field (TF) QKD is a step forward toward this direction, with early demonstrations suggesting it can beat the current rate-versus-distance records. A recently introduced variant of TF-QKD is particularly suited for experimental implementation, and has been shown to offer a higher key rate than other variants in the asymptotic regime, where users exchange an infinite number of signals. Here, we extend the security of this protocol to the finite-key regime, showing that it can overcome the fundamental bounds on point-to-point QKD with ~10^{10} transmitted signals. In many practical regimes of interest, our analysis offers higher key rates than those of alternative variants. Moreover, some of the techniques we develop are applicable to the finite-key analysis of other QKD protocols.

npj Quantum Information (2021)7:22 ; https://doi.org/10.1038/s41534-020-00345-3

INTRODUCTION

Quantum key distribution (QKD) enables two remote parties, Alice and Bob, to generate a shared secret key in the presence of an eavesdropper, Eve, who may have unbounded computational power at her disposal[1–3]. While, ideally, the two parties can be at any distance, in practice, due to the loss and noise in the channel, point-to-point QKD is limited to a certain maximum distance at which secret-key bits can securely be exchanged. In fact, the longest distance achieved to date in a terrestrial QKD experiment is ~400 km[4,5]. The main limitation is the exponential decrease of the transmittance, η, with the channel length in optical fibres. Even with a high repetition rate of 10 GHz, it would take an average of ~2 min to send a single photon over a distance of 600 km of standard optical fibres, and ~300 years to send it over 1000 km[6]. Indeed, fundamental bounds[7–11] on the private capacity of repeaterless point-to-point QKD protocols show that their secret-key rate scales at best approximately linearly with η. A protocol that aims to overcome this linear scaling must then include at least one middle node. Interestingly, this is not a sufficient condition. A well-known counterexample is the so-called measurement-device-independent QKD (MDI-QKD)[12], which uses the middle node for an untrusted Bell-state measurement operation. There are, however, extensions of MDI-QKD that can improve its rate scaling from η to $\sqrt{\eta}$ by either using quantum memories[13,14] or quantum non-demolition measurements[15]. Such setups can, in fact, be considered to be the simplest examples of quantum repeaters[6,16], which are the ultimate solution to trust-free long-distance quantum communications[17]. However, even these simple versions may need more time to be efficiently implemented in practice[18,19].

Remarkably, the recently proposed twin-field QKD (TF-QKD)[20] can also overcome this linear scaling while using a relatively simple setup. TF-QKD is related to MDI-QKD, and it inherits its immunity to detector side channels. However, it relies on single-photon, rather than two-photon, interference for its entanglement swapping operation. The secret-key rate of this protocol was first conjectured[20] and then proven[21,22] to scale with $\sqrt{\eta}$ too, making this approach a strong candidate to beat the current QKD records[23–26] with today's technology. The main experimental challenge is that single-photon interference needs very precise phase stability, which makes it more demanding than two-photon interference. Also, some of its current security proofs[21,22] need Alice and Bob to randomly choose a global phase, and then post-select only those rounds in which their choices match, which causes a drop in the secret-key rate. Since the original proposal, several variants of TF-QKD have been developed[27–30], sharing the single-photon interference idea and its consequent $\sqrt{\eta}$ scaling, but differing in their experimental setups and security proofs. Moreover, some of these variants have been shown to be robust against phase reference mismatch[28–30], which simplifies their experimental implementation.

In this paper, we focus on the TF-QKD variant introduced in ref. [28], which has two key features: (i) it does not need phase post-selection, which results in a higher secret-key rate; and (ii) it is a convenient option for experimental implementation. Indeed, many of the current TF-QKD experiments use this variant[23,24,26]. One of its defining characteristics is its unconventional security proof; specifically, its estimation of the phase-error rate, a parameter needed to bound the amount of key information that may have leaked to an eavesdropper. In many QKD protocols, the phase-error rate of the single-photon emissions in one basis can be directly estimated by bounding the bit-error rate of the single-photon emissions in the other basis. In the above TF-QKD variant, however, the encoding bases are not mutually unbiased. To estimate the phase-error rate, the authors in ref. [28] use the complementarity[31] between the "phase" and the "photon number" of a bosonic mode. In this case, the security of a bit encoded in the relative phase of two coherent pulses can be related to the detection statistics of photon-number states. More specifically, in the asymptotic regime, the phase-error rate can be bounded by a non-linear function of infinitely many yield

[1]School of Electronic and Electrical Engineering, University of Leeds, Leeds, UK. [2]EI Telecomunicación, Dept. of Signal Theory and Communications, University of Vigo, Vigo, Spain. [3]NTT Basic Research Laboratories, NTT Corporation, Atsugi, Kanagawa, Japan. [4]NTT Research Center for Theoretical Quantum Physics, NTT Corporation, Atsugi, Kanagawa, Japan. [5]NTT Communication Science Laboratories, NTT Corporation, Atsugi, Kanagawa, Japan. ✉email: g.j.curraslorenzo@leeds.ac.uk

probabilities for even photon-number states[28], which can be estimated via the decoy-state method[32–34].

While, in the asymptotic regime, the protocol in ref. [28] can offer a higher key rate than its counterparts, it is not obvious if this advantage will still hold in a practical setting, where only a finite number of pulses is sent. In the finite-key regime, one should account for possible statistical fluctuations between the true phase-error rate and the measurement data used to estimate it. There are, however, two challenges in doing so. The first challenge is that the phase-error rate of the protocol is related to the measurement statistics of infinitely many combinations of photon-number states; in practice, one can only obtain bounds for a finite number of them, and dealing with the unbounded components is not as straightforward as in the asymptotic regime. The second challenge is that, unlike in many other QKD protocols, the encoding bases are not mutually unbiased. This opens the possibility that, under a coherent attack by Eve, the detection statistics of a particular round may depend on the basis choices made in previous rounds. Accounting for these correlations makes the analysis quite cumbersome.

In this work, we provide a rigorous security proof for the protocol in ref. [28] that accounts for these two issues in the finite-key setting. Our security proof provides a tight bound on the key rate against general coherent attacks. To overcome the two main challenges mentioned above, we borrow ideas from the finite-key analysis of MDI-QKD[35] and the loss-tolerant protocol[36,37], as well as introduce several methods of our own. To obtain a tighter result, we employ a recent technique to bound the deviation between a sum of correlated random variables and its expected value[38] which can be much tighter than the widely employed Azuma's inequality[39] when the success probability is low. Importantly, our numerical simulations show that the protocol can overcome the repeaterless bounds[8–10] for a block size of $\sim 10^{10}$ transmitted signals in nominal working conditions.

During the preparation of this manuscript, an alternative finite-key security analysis for an identical protocol setup has been reported in ref. [40], using an interesting, but different, approach. We would like to highlight that our analysis imposes fewer conditions on the setup parameters than that of ref. [40], and results in a higher key rate in most practical regimes. In the "Discussion" section, we compare both approaches. We also compare our results with those of the sending-or-not-sending TF-QKD protocol introduced in ref. [30], whose security has recently been extended to the finite-key regime[41]. We find that for reasonably large block sizes, and sufficiently low phase reference mismatch errors, the asymptotic key rate advantage of the scheme in ref. [28] is maintained in the finite-key regime, for most practical ranges of distance.

RESULTS

Protocol description

The setup of the TF-QKD protocol in ref. [28] is illustrated in Fig. 1 and its step-by-step description is given below. Alice and Bob generate quantum signals and send them to a middle node, Charlie, who would ideally couple them at a balanced 50:50 beamsplitter and perform a photodetection measurement. For simplicity, we assume the symmetric scenario in which the Alice–Charlie and Bob–Charlie quantum channels are identical. We note, however, that our analysis can be straightforwardly extended to the asymmetric scenario recently considered in refs. [42,43]. The emitted quantum signals belong to two bases, selected at random. In the X basis, Alice and Bob send phase-locked coherent states $|\pm\alpha\rangle$ with a random phase of either 0 or π with respect to a pre-agreed reference. In the Z basis, Alice and Bob generate phase-randomised coherent states (PRCSs), which are diagonal in the Fock basis. The X-basis states are used to generate the key, while the Z-basis data is used to estimate the detection statistics of Fock states, in combination with the decoy-state method. This is a crucial step in estimating the phase-error rate of the key, thus bounding the information that could have been leaked to a potential eavesdropper. The detailed steps of the protocol are:

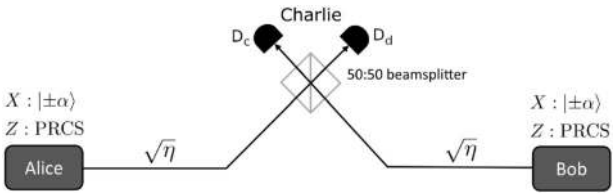

Fig. 1 Setup of the simple TF-QKD protocol[28] considered in this work. Alice and Bob generate their sifted key from the rounds in which they both select the X basis and Charlie declares that a single detector has clicked. The key bit is encoded in the phase of their coherent state. When the users select the same (a different) bit, the constructive (destructive) interference at Charlie's 50:50 beamsplitter should cause a click in detector D_c (D_d). The Z-basis PRCSs are only used to estimate the phase-error rate of the X-basis emissions.

(1) Preparation
Alice (Bob) chooses the key-generation basis X with probability p_X or the parameter estimation basis Z with probability $p_Z = 1 - p_X$, and

(1.1) If she (he) chooses the X basis, she (he) generates a random bit b_A (b_B), prepares an optical pulse in the coherent state $|(-1)^{b_A}\alpha\rangle$ ($|(-1)^{b_B}\alpha\rangle$), and sends it to Charlie.

(1.2) If she (he) chooses the Z basis, she (he) sends an optical pulse in a PRCS of intensity μ, selected from the set $\boldsymbol{\mu} = \{\mu_0, \mu_1, \ldots, \mu_{d-1}\}$ with probability p_μ, where d is the number of decoy intensities used.

They repeat step (1) for N rounds.

(2) Detection
An honest Charlie measures each round separately by interfering Alice and Bob's signals at a 50:50 beamsplitter, followed by threshold detectors D_c and D_d placed at the output ports corresponding to constructive and destructive interference, respectively. After the measurement, Charlie reports the pair (k_c, k_d), where $k_c = 1$ ($k_d = 1$) if detector D_c (D_d) clicks and $k_c = 0$ ($k_d = 0$) otherwise. If he is dishonest, Charlie can measure all rounds coherently using an arbitrary quantum measurement, and report N pairs (k_c, k_d) depending on the result. A round is considered successful (unsuccessful) if $k_c \neq k_d$ ($k_c = k_d$).

(3) Sifting
For all successful rounds, Alice and Bob disclose their basis choices, keeping only those in which they have used the same basis. Let $\mathcal{M}_X$ ($\mathcal{M}_Z$) be the set of successful rounds in which both users employed the X (Z) basis, and let $M_X = |\mathcal{M}_X|$ ($M_Z = |\mathcal{M}_Z|$) be the size of this set. Alice and Bob disclose their intensity choices for the rounds in $\mathcal{M}_Z$ and learn the number of rounds $M^{\mu\nu}$ in $\mathcal{M}_Z$ in which they selected intensities $\mu \in \boldsymbol{\mu}$ and $\nu \in \boldsymbol{\mu}$, respectively. Also, they generate their sifted keys from the values of b_A and b_B corresponding to the rounds in $\mathcal{M}_X$. For those rounds in which $k_c = 0$ and $k_d = 1$, Bob flips his sifted key bit.

(4) Parameter estimation
Alice and Bob apply the decoy-state method to $M^{\mu\nu}$, for $\mu, \nu \in \boldsymbol{\mu}$, obtaining upper bounds M^{U}_{nm} on the number of rounds M_{nm} in $\mathcal{M}_Z$ in which they sent n and m photons, respectively. They do this for all $n, m \geq 0$ such that $n + m$ is

even and $n+m \le S_{\text{cut}}$ for a prefixed parameter S_{cut}. Then, they use this data to obtain an upper bound N_{ph}^{U} on the number of phase errors, N_{ph}, in their sifted keys.

(5) Postprocessing

(5.1) Error correction: Alice sends Bob a prefixed amount λ_{EC} of syndrome information bits through an authenticated public channel, which Bob uses to correct errors in his sifted key.

(5.2) Error verification: Alice and Bob compute a hash of their error-corrected keys using a random universal hash function, and check whether they are equal. If so, they continue to the next step; otherwise, they abort the protocol.

(5.3) Privacy amplification: Alice and Bob extract a secret-key pair (S_A, S_B) of length $|S_A| = |S_B| = \ell$ from their error-corrected keys using a random two-universal hash function.

Parameter estimation and secret-key rate analysis

The main contribution of this work—see "Methods" section for the details—is a procedure to obtain a tight upper bound N_{ph}^{U} on the total number of phase errors N_{ph} in the finite-key regime for the protocol described above. Namely, we find that, except for an arbitrarily small failure probability ε, it holds that

$$N_{\text{ph}} \le N_{\text{ph}}^{\text{U}} := \frac{p_x^2}{p_z^2}\sum_{j=0}^{1}\left[\sum_{\substack{n,m\in\mathbb{N}_j\\ n+m\le S_{\text{cut}}}}\sqrt{\frac{p_{nm|X}}{p_{nm|Z}}}\sqrt{M_{nm}^{\text{U}}+\Delta_{nm}}+\sqrt{M_Z+\Delta}\sum_{\substack{n,m\in\mathbb{N}_j\\ n+m>S_{\text{cut}}}}\sqrt{\frac{p_{nm|X}}{p_{nm|Z}}}\right]^2+\Delta, \tag{1}$$

where $p_{nm|X}$ ($p_{nm|Z}$) is the probability that Alice and Bob's joint X (Z) basis pulses contain n and m photons, respectively, given by

$$p_{nm|X} = |\langle a|n\rangle|^2|\langle a|m\rangle|^2, \tag{2}$$

$$p_{nm|Z} = \sum_{\mu,\nu\in\underline{\boldsymbol{\mu}}} p_\mu p_\nu p_{n|\mu} p_{m|\nu}, \tag{3}$$

with $p_{n|\mu} = \mu^n \exp(-\mu)/n!$ being the Poisson probability that a PRCS pulse of intensity μ will contain n photons; $\mathbb{N}_0$ ($\mathbb{N}_1$) is the set of non-negative even (odd) integers; and Δ and Δ_{nm} are statistical fluctuation terms defined in step (4) of subsection "Instructions for experimentalists", where we provide a step-by-step instruction list to apply our results to the measurement data obtained in an experimental setup. The rest of the parameters have been introduced in the protocol description.

When it comes to finite-key analysis, there is one key difference between the protocol considered in this work and several other protocols, such as, for example, decoy-state BB84[44], decoy-state MDI-QKD[35], and sending-or-not-sending TF-QKD[41]. In all the latter setups, when there are no state preparation flaws, the single-photon components of the two encoding bases are mutually unbiased; in other words, they look identical to Eve once averaged by the bit selection probabilities. This implies that such states could have been generated from a maximally entangled bipartite state, where one of its components is measured in one of the two orthogonal bases, and the other half represents an encoded key bit. In fact, the user(s) could even wait until they learn which rounds have been successfully detected to decide their measurement basis, effectively delaying their choice of encoding basis. This possibility allows the application of a random sampling argument: since the choice of the encoding basis is independent of Eve's attack, the bit-error rate of the successful X-basis emissions provides a random sample of the phase-error rate of the successful Z-basis emissions, and vice versa. Then, one can apply tight statistical results, such as the Serfling inequality[45], to bound the phase-error rate in one basis, using the measured bit-error rate in the other basis. This approach, however, is not directly applicable to the protocol considered here, in which the secret key is extracted from all successfully detected X-basis signals, not just from their single-photon components. Moreover, the encoding bases are not mutually unbiased: the Z-basis states are diagonal in the Fock basis, while the X-basis states are not. This will require a different, perhaps more cumbersome, analysis as we highlight below.

To estimate the X-basis phase-error rate from the Z-basis measurement data, we construct a virtual protocol in which the users learn their basis choice by measuring a quantum coin after Charlie/Eve reveals which rounds were successful. Note that, because of the biased basis feature of the protocol, the statistics of the quantum coins associated to the successful rounds could depend on Eve's attack. This means that the users cannot delay their choice of basis, which prevents us from applying the random sampling argument. Still, it turns out that the quantum coin technique now allows us to upper bound the average number of successful rounds in which the users had selected the X basis and obtained a phase error. This bound is a non-linear function of the average number of successful rounds in which they had selected the Z basis and respectively sent n and m photons, with $n+m$ even. More details can be found in the "Methods" section; see Eq. (19).

The main tool we use to relate each of the above average terms to their actual occurrences, N_{ph} and M_{nm}, is Azuma's inequality[39], which is widely used in security analyses of QKD to bound sums of observables over a set of rounds of the protocol (in our case, the set of successful rounds after sifting) when the independence between the observables corresponding to different rounds cannot be guaranteed. When using Azuma's inequality, the deviation term Δ scales with the square root of the number of terms in the sum. In our case, Δ scales with $\sqrt{M_s}$, where M_s is the number of successful rounds after sifting. For parameters of comparable magnitude to M_s, this provides us with a reasonably tight bound. Whenever the parameter of interest is small, however, the provided bound could instead be loose. This is the case for the crucial term M_{00}^{U} in Eq. (1), as vacuum states are unlikely to result in successful detection events, and thus the bound obtained with Azuma's inequality can be loose. This is important because, in Eq. (1), the coefficient associated to the vacuum term is typically the largest. To obtain a better bound for this term, we employ a remarkable recent technique to bound the deviation between a sum of dependent random variables and its expected value[38]. This technique provides a much tighter bound than Azuma's inequality when the value of the sum is much lower than the number of terms in the sum. In particular, it provides a tight upper bound for the vacuum component M_{00}. In "Methods" section, we provide a statement of the result and we explain how we apply it to our protocol.

Having obtained the upper bound $e_{\text{ph}}^{\text{U}} := N_{\text{ph}}^{\text{U}}/M_X$ on the phase-error rate, we show in Supplementary Note A that, if the length of the secret key obtained after the privacy amplification step satisfies

$$\ell \le M_X\left[1-h(e_{\text{ph}}^{\text{U}})\right] - \lambda_{\text{EC}} - \log_2\frac{2}{\epsilon_{\text{c}}} - \log_2\frac{1}{4\epsilon_{\text{PA}}^2}, \tag{4}$$

the protocol is guaranteed to be ϵ_{c}-correct and ϵ_{s}-secret, with $\epsilon_{\text{s}} = \sqrt{\varepsilon} + \epsilon_{\text{PA}}$; where ε is the failure probability associated to the estimation of the phase-error rate, $h(x) = -x\log_2 x - (1-x)\log_2(1-x)$ is the Shannon binary entropy function, and λ_{EC} is number of bits that are spent in the error correction procedure. Here, our security analysis follows the universal composable security framework[46,47], according to which a protocol is ϵ_{sec}-secure if it is both ϵ_{c}-correct and ϵ_{s}-secret, with $\epsilon_{\text{sec}} \le \epsilon_{\text{c}} + \epsilon_{\text{s}}$.

Instructions for experimentalists

Here, we provide a step-by-step instruction list to apply our security analysis to a real-life experiment:

(1) Set the security parameters ϵ_c and ϵ_{PA}, as well as the failure probabilities ε_c and ε_a for the inverse multiplicative Chernoff bound and the concentration bound for sums of dependent random variables, respectively. Set S_{cut}. Calculate the overall failure probability ε of the parameter estimation process, which depends on the number of times that the previous two inequalities are applied. In general, $\varepsilon = d^2\varepsilon_c + \left(\left\lfloor \frac{S_{cut}}{2} \right\rfloor + 1\right)^2 \varepsilon_a + \varepsilon_a$, where d is the number of decoy intensities employed by each user. For $S_{cut} = 4$ and three decoy intensities, we have that $\varepsilon = 9\varepsilon_c + 10\varepsilon_a$.

(2) Use prior information about the channel to obtain a prediction $\tilde{M}^{U}_{00}$ on M^{U}_{00}, the upper bound on the number of Z-basis vacuum events that will be obtained after applying the decoy-state method.

(3) Run steps (1)–(3) of the protocol, obtaining a sifted key of length M_X, and Z-basis measurement counts $M^{\mu\nu}$ for $\mu, \nu \in \underline{\boldsymbol{\mu}}$. Let $M_s = M_X + M_Z$ be the number of successful rounds after sifting.

(4) Use the analytical decoy-state method included in the Supplementary Note B and the measured values of $M^{\mu\nu}$ to obtain upper bounds M^{U}_{nm}, for all n, m such that $n + m$ is even and $n + m \le S_{cut}$. Alternatively, use the numerical estimation method introduced in the Supplementary Notes of ref. [35].

(5) Set $\Delta = \sqrt{\frac{1}{2} M_s \ln \varepsilon_a^{-1}}$ and $\Delta_{nm} = \Delta$ for all n, m except for $m = n = 0$. Substitute $\tilde{\Lambda}_n \to \tilde{M}^{U}_{00}$ in Eq. (32) to find parameters a and b. Set

$$\Delta_{00} = \left[b + a\left(\frac{2M^{U}_{00}}{M_s} - 1 \right) \right] \sqrt{M_s}, \tag{5}$$

(6) Use Eq. (1) to find N^{U}_{ph} and set $e^{U}_{ph} = N^{U}_{ph}/M_X$.

(7) Use Eq. (4) to specify the required amount of privacy amplification and to find the corresponding length of the secret key that can be extracted. The key obtained is ϵ_{sec}-secure, with $\epsilon_{sec} = \epsilon_c + \epsilon_s$ and $\epsilon_s = \sqrt{\varepsilon} + \epsilon_{PA}$.

DISCUSSION

In this section, we analyse the behaviour of the secret-key rate as a function of the total loss. We simulate the nominal scenario in which there is no Eve and Charlie is honest. In this case, the total Alice–Bob loss includes the loss in the quantum channels, as well as the inefficiency of Charlie's detectors. We compare the key rate for the protocol in Fig. 1, using the finite-key security analysis introduced in the previous section, with that of the sending-or-not-sending TF-QKD protocol[30,41], as well as with the finite-key analysis presented in ref. [40]. We also include the asymptotic secret-key capacity for repeaterless QKD systems over lossy channels, known as the PLOB bound[9], for comparison. It is given by $-\log_2(1-\eta)$, where η is the transmittance of the Alice–Bob quantum channel, which includes the efficiency of Charlie's detectors. While specific bounds for the finite-key setting have recently been studied[10,48], in the practical regimes of interest to this work, they numerically offer a negligible difference to the PLOB bound. The latter has then been used in all relevant graphs for consistency. To simulate the data that would be obtained in all protocols, we use the simple channel model described in Supplementary Note C, which accounts for phase reference mismatch and polarisation misalignment. Also, we assume that both users employ three decoy-state intensities $\mu_0 > \mu_1 > \mu_2$. Since the optimal value $\mu_2 = 0$ is typically difficult to achieve in practice, we set $\mu_2 = 10^{-4}$ and optimise the secret-key rate over the value of μ_0 and μ_1. We also optimise it over the selection probabilities, as well as over p_X and a.

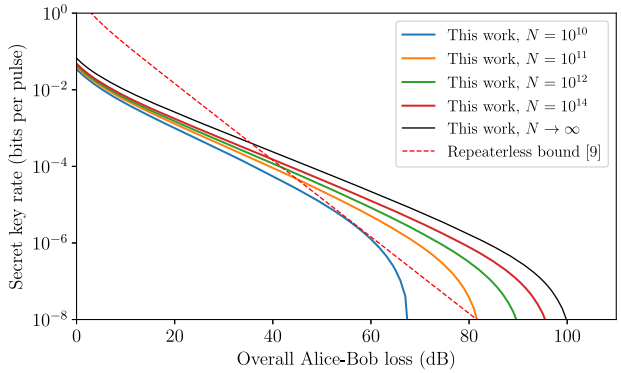

Fig. 2 Secret-key rate obtainable as a function of the channel loss. We consider different values of the block size N, which represents the total number of rounds in the protocol. The overall Alice–Bob loss includes the loss in both quantum channels and in Charlie's detectors. The simulation parameters are stated in the main text.

In our simulations, we model the phase reference mismatch between Alice and Bob's pulses by shifting Bob's signals by an angle $\phi = \delta_{ph}\pi$, where $\delta_{ph} = 9.1\%$. This corresponds to a QBER of ~2% for most attenuations, matching the experimental results in ref. [23]. For brevity, we do not consider the effect of polarisation misalignment in our numerical results, but one can use the provided analytical model to study different scenarios of interest. In principle, even if the mechanism used for polarisation stability is not perfect, one can use polarisation filters to ensure that the same polarisation modes are being coupled at the 50:50 beamsplitter, at the cost of introducing additional loss. We assume a per-pulse dark count probability $p_d = 10^{-8}$ for each detector. We assume an error correction leakage of $\lambda_{EC} = fM_X h(e_X)$, where e_X is the bit-error rate of the sifted key, and f is the error correction inefficiency, which we assume to be $f = 1.16$. For the security bounds, we set $\epsilon_c = \epsilon_s = 10^{-10}$, and for simplicity we set $\varepsilon = \epsilon_{PA} = \epsilon_s/3$.

In Fig. 2, we display the secret-key rate per pulse achievable for different values of the block size, N, of transmitted signals. It can be seen that the protocol could outperform the repeaterless bound for a block size of ~10^{10} transmitted signals per user, at an approximate total loss of 50 dB. For standard optical fibres, this corresponds to a total distance of 250 km, if we neglect the loss in the photodetectors. At a 1 GHz clock rate, it takes only ~10 s to collect the required data. For a block size of 10^{11} transmitted signals, the protocol can already outperform the repeaterless bound for a total loss ranging from 45 to over 80 dB. By increasing N, we approach the asymptotic performance of the protocol. We note that our choice of dark count probability, $p_d = 10^{-8}$, may be conservative, since a dark count rate of 1 c.p.s., corresponding to $p_d = 10^{-9}$ with a repetition rate of 1 GHz, may be achievable with state-of-the-art SSPD[49]. In Supplementary Note D, we show an additional graph for $p_d = 10^{-9}$. We find that, for sufficiently large block sizes, the maximum distance increases when the dark count probability decreases. Interestingly, however, this is not the case for $N = 10^{10}$, for which the two curves are almost identical.

The dependence of the secret-key rate on the block size N has been shown in Fig. 3, at a fixed total loss of 50 dB and for several values of the phase reference mismatch δ_{ph}. In all cases, there is a minimum required block size to obtain a positive key rate. This minimum block size can be even <10^9 in the ideal case of no phase reference mismatch, and it goes up to ~10^{10} at $\delta_{ph} = 20\%$. There is a sharp increase in the secret-key rate once one goes over

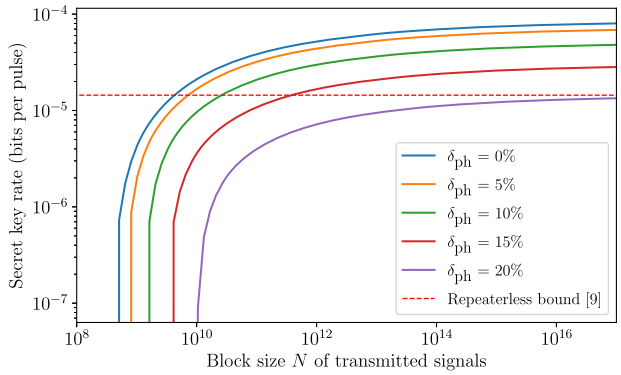

Fig. 3 Secret-key rate obtainable as a function of the block size *N*. We assume a total loss of 50 dB and consider several values of the phase reference mismatch δ_{ph}. All other simulation parameters are stated in the main text.

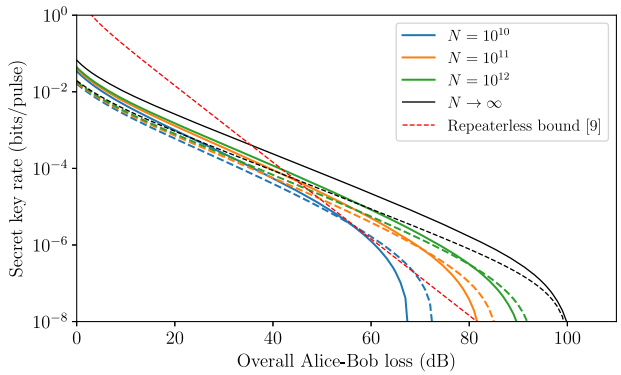

Fig. 4 Comparison between this work (solid) and sending-or-not-sending TF-QKD[30,41] (dashed). We consider different values for the block size *N* of transmitted signals. All other simulation parameters are stated in the main text.

this minimum required block size, after which one slowly approaches the key rate in the asymptotic limit. The latter behaviour is likely due to the use of Azuma's inequality. One can, nevertheless, overcome the repeaterless bound at a reasonable block size in a practical regime where $\delta_{ph} \leq 15\%$. At higher values of total loss this crossover happens at even larger values of δ_{ph}.

In Fig. 4, we compare the performance of our protocol with that of the sending-or-not-sending TF-QKD protocol presented in refs. [30,41]. To compute the results of the sending-or-not-sending protocol, we have used the analysis in ref. [41], after correcting a mistake present in Appendix A of that work. Namely, according to Eqs. (S14)–(S19) of ref. [50], if the failure probability of the phase-error rate estimation is $\bar{\varepsilon}$, then the smooth max entropy term in the left-hand side of Eq. (A5) should be $H_{\max}^{\sqrt{\bar{\varepsilon}}}$ instead of $H_{\max}^{\bar{\varepsilon}}$. In the asymptotic regime, the protocol considered in this work outperforms the sending-or-not-sending protocol at all values of total loss. For a block size of 10^{12} transmitted signals, this is still the case up to 80 dB of total loss, after which the key rate is already $<10^{-6}$ bits per pulse for both protocols. For a block size of 10^{10} transmitted signals, however, the curves for the two protocols cross at ~55 dB, after which the sending-or-not-sending protocol offers a better performance. This behaviour is due to the different statistical fluctuation analyses applied to the two protocols. As explained in the "Results" section, the single-photon components in the sending-or-not-sending protocol are mutually unbiased, allowing for a simpler and tighter estimation of the phase-error rate. This is not the case for our TF-QKD protocol, for which this estimation involves the application of somewhat looser bounds for several terms in Eq. (1). We conclude that for sufficiently large

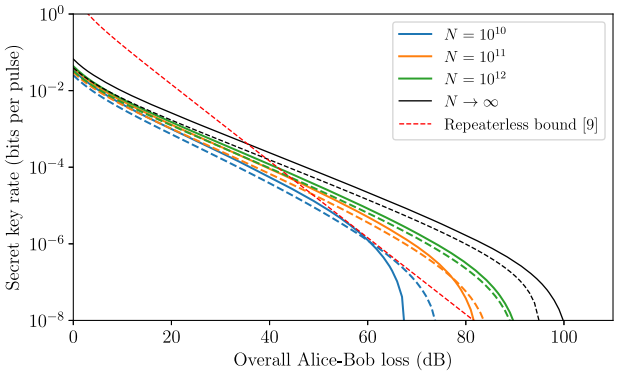

Fig. 5 Comparison between this work (solid) and the alternative analysis in ref. [40] (dashed). We consider different values for the block size *N* of transmitted signals. All other simulation parameters are stated in the main text.

block sizes, and a sufficiently low phase reference mismatch, the protocol considered in this work maintains its better key rate performance over the sending-or-not-sending variant. We note that for smaller block sizes and higher values of phase reference mismatch, this comparative advantage is reduced, or even inverted in some regimes. For completeness, in Supplementary Note D, we provide additional simulation results for a broader range of parameter values.

Finally, in Fig. 5, we compare our results with those of the alternative analysis in ref. [40]. To compute the secret-key rate of the latter, we use the code provided by the authors, except for the adjustments needed to match it to the channel model described in Supplementary Note C. It can be seen that, in most regimes, the analysis introduced in this paper provides a higher key rate than that of ref. [40]. Moreover, we remark that the security proof presented in ref. [40], in its current form, is only applicable when the state generated by the weakest decoy intensity μ_2 is a perfect vacuum state of intensity $\mu_2 = 0$. The security analysis presented in this work, however, can be applied to any experimental value of μ_2, and we assume a value of $\mu_2 = 10^{-4}$, which may be easier to achieve in practice. That said, the security proof in ref. [40] adopts an interesting approach that results in a somehow simpler statistical analysis. In particular, unlike in the analysis presented in this paper, the authors in ref. [40] do not estimate the detection statistics of photon-number states as an intermediate step to bounding the phase-error rate. Instead, they show that the operator corresponding to a phase error can be bounded by a linear combination of the *Z*-basis decoy states. While this linear bound is asymptotically looser than the non-linear formula in Eq. (1), it allows the application of a simpler statistical analysis based on a double use of Bernoulli sampling. Given that the finite-key analysis of a protocol could be part of the software package of a product, we believe that the additional key rate achievable by our analysis in many regimes justifies its slightly more complex approach.

In conclusion, we have proven the security of the protocol proposed in ref. [28], in the finite-key regime and against coherent attacks. Our results show that, under nominal working conditions experimentally achievable by today's technology, this scheme could outperform the repeaterless secret-key rate bound in a key exchange run of ~10 s, assuming a 1 GHz clock rate. In terms of key rate, it would also outperform other TF-QKD variants, as well as alternative security proofs, in many practical regimes of interest.

METHODS

In this section, we introduce the procedure that we use to prove the security of the protocol, referring to the Supplementary Notes when appropriate. For notation clarity, we assume the symmetric scenario in

which Alice and Bob employ the same X-basis amplitude a and the same set of Z-basis intensities $\boldsymbol{\mu}$, which is optimal when the Alice–Charlie and Bob–Charlie channels are identical. However, the analysis can be applied as well to the asymmetric scenario[42,43] by appropriately redefining the parameters $p_{nm|X}$ and $p_{nm|Z}$.

Virtual protocol

To bound the information leakage to Eve, we construct an entanglement-based virtual protocol that is equivalent to the actual protocol. In this virtual protocol, Alice and Bob measure their local ancilla systems in a basis that is conjugate to that used to generate the key. We refer to the error rate of the virtual protocol as the phase-error rate e_{ph}. The objective of the security analysis is to find an upper bound $e_{\mathrm{ph}}^{\mathrm{U}}$ such that $\Pr(e_{\mathrm{ph}}>e_{\mathrm{ph}}^{\mathrm{U}}) \leq \varepsilon$. In Supplementary Note A, we show how this can be used to prove the security of the key obtained in the actual protocol.

In the virtual protocol, Alice replaces her X-basis emissions by the preparation of the state

$$|\psi_X\rangle_{Aa} = \frac{1}{\sqrt{2}}(|+\rangle_A|\alpha\rangle_a + |-\rangle_A|-\alpha\rangle_a), \tag{6}$$

where A is an ancilla system at Alice's lab, a is the photonic system sent to Eve, and $|\pm\rangle = \frac{1}{\sqrt{2}}(|0\rangle \pm |1\rangle)$; while Bob replaces his X-basis emissions by a similarly defined $|\psi_X\rangle_{Bb}$. After Eve's attack, Alice and Bob measure systems A and B in the Z-basis $\{|0\rangle, |1\rangle\}$, which is conjugate to the X basis $\{|+\rangle, |-\rangle\}$ that they would use to generate the key. It is useful to write the state in Eq. (6) as

$$|\psi_X\rangle_{Aa} = |0\rangle_A|C_0\rangle_a + |1\rangle_A|C_1\rangle_a, \tag{7}$$

where $|C_0\rangle$ and $|C_1\rangle$ are the (unnormalised) cat states

$$|C_0\rangle = \frac{1}{2}(|\alpha\rangle + |-\alpha\rangle), \qquad |C_1\rangle = \frac{1}{2}(|\alpha\rangle - |-\alpha\rangle). \tag{8}$$

Alice's Z-basis emissions are diagonal in the Fock basis, and the virtual protocol replaces them by their purification

$$|\psi_Z\rangle_{Aa} = \sum_{n=0}^{\infty} \sqrt{p_{n|Z}}|n\rangle_A|n\rangle_a, \tag{9}$$

where $p_{n|Z} = \sum_{\mu\in\boldsymbol{\mu}} p_\mu p_{n|\mu}$ is the probability that Alice's Z-basis pulse contains n photons, averaged over the selection of μ. Unlike in the actual protocol, in the virtual protocol Alice and Bob learn the photon number of their signals by measuring systems A and B after Eve's attack.

Lastly, Alice's emission of $|\psi_X\rangle_{Aa}$ with probability p_X and $|\psi_Z\rangle_{Aa}$ with probability p_Z is replaced by the generation of the state

$$|\psi\rangle_{A_cAa} = \sqrt{p_X}|0\rangle_{A_c}|\psi_X\rangle_{Aa} + \sqrt{p_Z}|1\rangle_{A_c}|\psi_Z\rangle_{Aa}, \tag{10}$$

where A_c is a quantum coin ancilla at Alice's lab; while Bob's is replaced by an equally defined $|\psi\rangle_{B_cBb}$. Alice and Bob measure systems A_c and B_c after Eve's attack, delaying the reveal of their basis choice. The full description of the virtual protocol is the following:

(1) Preparation
Alice and Bob prepare N copies of the state $|\phi\rangle = |\psi\rangle_{A_cAa} \otimes |\psi\rangle_{B_cBb}$ and send all systems a and b to Eve over the quantum channel.

(2) Detection
Eve performs an arbitrary general measurement on all the subsystems a and b of $|\phi\rangle^{\otimes N}$ and publicly announces N bit pairs (k_c, k_d). Without loss of generality, we assume that there is a one-to-one correspondence between her measurement outcome and her set of announcements. A round is considered successful (unsuccessful) if $k_c \neq k_d$ ($k_c = k_d$). Let $\mathcal{M}$ ($\overline{\mathcal{M}}$) represent the set of successful (unsuccessful) rounds.

(3) Virtual sifting
For all rounds, Alice and Bob jointly measure the systems A_c and B_c, learning whether they used the same or different bases, but not the specific basis they used. Let $\mathcal{M}_s$ ($\mathcal{M}_d$) denote the set of successful rounds in which they used the same (different) bases.

(4) Ancilla measurement

(4.1) For all rounds in $\mathcal{M}_s$, Alice (Bob) first measures the system A_c (B_c) in $\{|0\rangle, |1\rangle\}$, learning her (his) choice of basis. If the result is $|0\rangle_{A_c}$ ($|0\rangle_{B_c}$), she (he) measures system A (B) in $\{|0\rangle, |1\rangle\}$; if the result is $|1\rangle_{A_c}$ ($|1\rangle_{B_c}$), she (he) measures system A (B) in the Fock basis.

(4.2) For all rounds in $\mathcal{M}_d$, Alice (Bob) measures the systems A_c (B_c) and A (B), using the same strategy as in step (4.1).

(5) Intensity assignment
For all rounds in $\mathcal{M}$ in which Alice (Bob) obtained $|1\rangle_{A_c}$ ($|1\rangle_{B_c}$), she (he) assigns each n-photon state to intensity μ with probability $p_{\mu|n}$.

(6) Classical communication
For all rounds in $\mathcal{M}$, Alice and Bob announce their basis and intensity choices over an authenticated public channel.

(7) Estimation of the number of phase errors
Alice and Bob calculate an upper bound on N_{ph} using their Z-basis measurement data.

Two points from the virtual protocol above require further explanation. The first is that, in the real protocol, Bob flips his key bit when Eve reports $k_c = 0$ and $k_d = 1$. This step is omitted from the virtual protocol, since the X-basis bit flip gate σ_z has no effect on Bob's Z-basis measurement result. The second point concerns step (5), which may appear to serve no purpose, but is needed to ensure that the classical information exchanged between Alice and Bob is equivalent to that of the real protocol. The term $p_{\mu|n}$ is the probability that Alice's (Bob's) Z-basis n-photon pulse originated from intensity μ, and it is given by

$$p_{\mu|n} = \frac{p_\mu p_{n|\mu}}{\sum_{\mu\in\boldsymbol{\mu}} p_\mu p_{n|\mu}}. \tag{11}$$

Phase-error rate estimation

We now turn our attention to Alice and Bob's measurements in step (4.1) of the virtual protocol. Let $u \in \{1, 2, ..., M_s\}$ index the rounds in $\mathcal{M}_s$, and let ξ_u denote the measurement outcome of the uth round. The possible outcomes are $\xi_u = X_{ij}$, corresponding to $|00\rangle_{A_cB_c}|ij\rangle_{AB}$, where $i, j \in \{0, 1\}$; and $\xi_u = Z_{nm}$, corresponding to $|11\rangle_{A_cB_c}|n, m\rangle_{AB}$, where n and m are any non-negative integers. Note that the outcomes $|10\rangle_{A_cB_c}$ and $|01\rangle_{A_cB_c}$ are not possible due to the previous virtual sifting step. A phase error occurs when $\xi_u \in \{X_{00}, X_{11}\}$. In Supplementary Note E, we prove that the probability to obtain a phase error in the uth round, conditioned on all previous measurement outcomes in the protocol, is upper bounded by

$$\Pr(\xi_u \in \{X_{00}, X_{11}\}|\mathcal{F}_{u-1}) \leq \frac{p_X^2}{p_Z^2}\sum_{j=0}^{1}\left[\sum_{n,m\in\mathbb{N}_j}\sqrt{\frac{p_{nm|X}}{p_{nm|Z}}\Pr(\xi_u = Z_{nm}|\mathcal{F}_{u-1})}\right]^2, \tag{12}$$

where $\mathcal{F}_{u-1}$ is the σ-algebra generated by the random variables $\xi_1, ..., \xi_{u-1}$, $\mathbb{N}_0$ ($\mathbb{N}_1$) is the set of non-negative even (odd) numbers, and the probability terms $p_{nm|X}$ and $p_{nm|Z}$ have been defined in Eqs. (2) and (3). In Eq. (12), for notation clarity, we have omitted the dependence of all probability terms on the outcomes of the measurements performed in steps (2) and (3) of the virtual protocol.

Applying the concentration bound in Eq. (30), we have that, except with probability ε_a,

$$N_{\mathrm{ph}} \leq \sum_{u=1}^{M_s}\Pr(\xi_u \in \{X_{00}, X_{11}\}|\mathcal{F}_{u-1}) + \Delta, \tag{13}$$

where N_{ph} is the number of events of the form $\xi_u \in \{X_{00}, X_{11}\}$ in $\mathcal{M}_s$, and $\Delta = \sqrt{\frac{1}{2}M_s \ln \varepsilon_a^{-1}}$ is a deviation term. Similarly, from Eq. (30), we have that, except with probability ε_a,

$$\sum_{u=1}^{M_s}\Pr(\xi_u = Z_{nm}|\mathcal{F}_{u-1}) \leq M_{nm} + \Delta, \tag{14}$$

where M_{nm} is the number of events of the form $\xi_u = Z_{nm}$ in $\mathcal{M}_s$. As we will explain later, this bound is not tight when applied to the vacuum counts M_{00}. For this term, we use the alternative bound in Eq. (33), according to which, except with probability ε_a,

$$\sum_{u=1}^{M_s}\Pr(\xi_u = Z_{00}|\mathcal{F}_{u-1}) \leq M_{00} + \Delta_{00}. \tag{15}$$

In this case, the deviation term is given by

$$\Delta_{00} = \left[b + a\left(\frac{2M_{00}}{M_s} - 1\right)\right]\sqrt{M_s}, \tag{16}$$

where a and b can be found by substituting $\tilde{\Lambda}_n$ by $\tilde{M}_{00}^{\mathrm{U}}$ in Eq. (31).

Now, we will transform Eq. (12) to apply Eqs. (13)–(15). Let us denote the right-hand side of Eq. (12) as $f(\vec{p}_u)$, where $\vec{p}_u$ is a vector of probabilities composed of $\Pr(\xi_u = Z_{nm}|\mathcal{F}_{u-1})\ \forall\, n, m$. If we expand the square in $f(\vec{p}_u)$, we can see that all addends are positive and proportional to $\sqrt{p_1 p_2}$, where p_1 and p_2 are elements of $\vec{p}_u$, implying that $f(\vec{p}_u)$ is a concave function. Thus, by Jensen's inequality[51], we have

$$\frac{1}{M_s}\sum_{u=1}^{M_s} f(\vec{p}_u) \leq f\left(\frac{1}{M_s}\sum_{u=1}^{M_s}\vec{p}_u\right). \tag{17}$$

After taking the average over all rounds M_s on both sides of Eq. (12), applying Eq. (17) on the right-hand side, and cancelling out the term $1/M_s$ on both sides of the inequality, we have that

$$\sum_{u=1}^{M_s}\Pr(\xi_u \in \{X_{00}, X_{11}\}|\mathcal{F}_{u-1}) \leq \frac{p_X^2}{p_Z^2}\sum_{j=0}^{1}\left[\sum_{n,m\in\mathbb{N}_j}\sqrt{\frac{p_{nm|X}}{p_{nm|Z}}\sum_{u=1}^{M_s}\Pr(\xi_u = Z_{nm}|\mathcal{F}_{u-1})}\right]^2. \tag{18}$$

We are now ready to apply Eqs. (13)–(15) to substitute the sums of probabilities in Eq. (18) by N_{ph} and M_{nm}. However, note that, in their application of the decoy-state method, Alice and Bob only estimate the value of M_{nm} for terms of the form $n+m \leq S_{\mathrm{cut}}$, so it is only useful to substitute Eq. (14) for these terms. With this in mind, we obtain

$$N_{\mathrm{ph}} - \Delta \leq \frac{p_X^2}{p_Z^2}\sum_{j=0}^{1}\left[\sum_{\substack{n,m\in\mathbb{N}_j\\ n+m\leq S_{\mathrm{cut}}}}\sqrt{\frac{p_{nm|X}}{p_{nm|Z}}}\sqrt{M_{nm}+\Delta_{nm}} + \sum_{\substack{n,m\in\mathbb{N}_j\\ n+m>S_{\mathrm{cut}}}}\sqrt{\frac{p_{nm|X}}{p_{nm|Z}}\sum_{u=1}^{M_s}\Pr(\xi_u = Z_{nm}|\mathcal{F}_{u-1})}\right]^2, \tag{19}$$

where $\Delta_{nm} = \Delta$ except for Δ_{00}.

We still need to deal with the sum over the infinitely many remaining terms of the form $n+m > S_{\mathrm{cut}}$. For them, we apply the following upper bound

$$\sum_{u=1}^{M_s}\Pr(\xi_u = Z_{nm}|\mathcal{F}_{u-1}) \leq \sum_{u=1}^{M_s}\Pr(\xi_u = Z|\mathcal{F}_{u-1}) \leq M_Z + \Delta, \tag{20}$$

where $\xi_u = Z$ denotes that Alice and Bob learn that they have used the Z basis in the uth round in $\mathcal{M}_s$; and M_Z is the number of events of the form $\xi_u = Z$ obtained by Alice and Bob. In the last step, we have used Eq. (30), using an identical argument as in Eq. (13). When we apply Eq. (20) to Eq. (19), we end up with the term

$$\sum_{\substack{n,m\in\mathbb{N}_j\\ n+m>S_{\mathrm{cut}}}}\sqrt{\frac{p_{nm|X}}{p_{nm|Z}}}\sqrt{M_Z+\Delta} = \sqrt{M_Z+\Delta}\sum_{\substack{n,m\in\mathbb{N}_j\\ n+m>S_{\mathrm{cut}}}}\sqrt{\frac{p_{nm|X}}{p_{nm|Z}}}. \tag{21}$$

It can be shown that the infinite sum in Eq. (21) converges to a finite value if

$$\max\{\boldsymbol{\mu}\} > a^2. \tag{22}$$

Substituting Eq. (20) into Eq. (19), and isolating N_{ph}, we obtain

$$N_{\mathrm{ph}} \leq \frac{p_X^2}{p_Z^2}\sum_{j=0}^{1}\left[\sum_{\substack{n,m\in\mathbb{N}_j\\ n+m\leq S_{\mathrm{cut}}}}\sqrt{\frac{p_{nm|X}}{p_{nm|Z}}}\sqrt{M_{nm}+\Delta_{nm}} + \sqrt{M_Z+\Delta}\sum_{\substack{n,m\in\mathbb{N}_j\\ n+m>S_{\mathrm{cut}}}}\sqrt{\frac{p_{nm|X}}{p_{nm|Z}}}\right]^2 + \Delta. \tag{23}$$

Note that the right-hand side of Eq. (23) is a function of the measurement counts M_{nm}, which cannot be directly observed. They must be substituted by the upper bounds M_{nm}^{U} obtained via the decoy-state analysis, as explained below. After doing so, we obtain Eq. (1). The failure probability ε associated to the estimation of N_{ph} is upper bounded by summing the failure probabilities of all concentration inequalities used. That includes each application of Eqs. (30) and (33), which fail with probability ε_{a}; and each application of the multiplicative Chernoff bound in the decoy-state analysis, which fails with probability ε_{c}. In the case of three decoy intensities and $S_{\mathrm{cut}} = 4$, we have $\varepsilon = 9\varepsilon_{\mathrm{c}} + 10\varepsilon_{\mathrm{a}}$. In our simulations, we set $\varepsilon_{\mathrm{c}} = \varepsilon_{\mathrm{a}}$ for simplicity.

Decoy-state analysis

Since Alice and Bob's Z-basis emissions are a mixture of Fock states, the measurement counts M_{nm} have a fixed value, which is nevertheless unknown to them. Instead, the users have access to the measurement counts $M^{\mu\nu}$, the number of rounds in $\mathcal{M}_Z$ in which they selected intensities μ and ν, respectively. To bound M_{nm}, we use the decoy-state method[32–34]. This technique exploits the fact that Alice and Bob could have run an equivalent virtual scenario in which they directly send Fock states $|n,m\rangle$ with probability $p_{nm|Z}$, and then randomly assign each of them to intensities μ and ν with probability

$$p_{\mu\nu|nm} = \frac{p_{\mu\nu}p_{nm|\mu\nu}}{p_{nm|Z}}, \tag{24}$$

where $p_{\mu\nu} = p_\mu p_\nu$ and $p_{nm|\mu\nu} = p_{n|\mu}p_{m|\nu}$. In particular, each of the instances in which Alice and Bob chose the Z basis, sent n and m photons, and Eve announced a detection is assigned to intensities μ and ν with a fixed probability $p_{\mu\nu|nm}$, even if Eve employs a coherent attack. This implies that these assignments can be regarded as an independent Bernoulli trial, and $M^{\mu\nu}$ can be regarded as a sum of independent Bernoulli trials. The average value of $M^{\mu\nu}$ is

$$\mathbb{E}[M^{\mu\nu}] = \sum_{n,m=0}^{\infty} p_{\mu\nu|nm}M_{nm}. \tag{25}$$

In the actual protocol, Alice and Bob know the realisations $M^{\mu\nu}$ of these random variables. By using the inverse multiplicative Chernoff bound[52,53], stated in Supplementary Note F, they can compute lower and upper bounds $\mathbb{E}^{\mathrm{L}}[M^{\mu\nu}]$ and $\mathbb{E}^{\mathrm{U}}[M^{\mu\nu}]$ for $\mathbb{E}[M^{\mu\nu}]$. These will set constraints on the possible value of the terms M_{nm}. We are interested in the indices (i, j) such that $i+j \leq S_{\mathrm{cut}}$ and $i+j$ is even, and an upper bound on each M_{ij} can be found by solving the following linear optimisation problem

$$\begin{aligned} &\max M_{ij}\\ &\text{s.t. } \forall \mu,\nu\ \mathbb{E}^{\mathrm{U}}[M^{\mu\nu}] \geq \sum_{n,m=0}^{\infty} p_{\mu\nu|nm}M_{nm},\\ &\mathbb{E}^{\mathrm{L}}[M^{\mu\nu}] \leq \sum_{n,m=0}^{\infty} p_{\mu\nu|nm}M_{nm}. \end{aligned} \tag{26}$$

This problem can be solved numerically using linear programming techniques, as described in the Supplementary Note 2 of ref. [35]. While accurate, this method can be computationally demanding. For this reason, we have instead adapted the asymptotic analytical bounds of refs. [42,54] to the finite-key scenario and used them in our simulations. The results obtained using these analytical bounds are very close to those achieved by numerically solving Eq. (26). This analytical method is described in Supplementary Note B.

Concentration inequality for sums of dependent random variables

A crucial step in our analysis is the substitution of the sums of probabilities in Eq. (18) by their corresponding observables in the protocol. Typically, this is done by applying the well-known Azuma's inequality[39]. Instead, we use the following recent result[38]:

Let $\xi_1, \ldots, \xi_n$ be a sequence of random variables satisfying $0 \leq \xi_l \leq 1$, and let $\Lambda_l = \sum_{u=1}^{l}\xi_u$. Let $\mathcal{F}_l$ be its natural filtration, i.e. the σ-algebra generated by $\{\xi_1, \ldots, \xi_l\}$. For any n, and any a, b such that $b \geq |a|$,

$$\Pr\left[\sum_{u=1}^{n}E(\xi_u|\mathcal{F}_{u-1}) - \Lambda_n \geq \left[b + a\left(\frac{2\Lambda_n}{n} - 1\right)\right]\sqrt{n}\right] \leq \exp\left[\frac{-2(b^2-a^2)}{\left(1+\frac{4a}{3\sqrt{n}}\right)^2}\right]. \tag{27}$$

By replacing $\xi_l \to 1-\xi_l$ and $a \to -a$, we also derive

$$\Pr\left[\Lambda_n - \sum_{u=1}^{n}E(\xi_u|\mathcal{F}_{u-1}) \geq \left[b + a\left(\frac{2\Lambda_n}{n} - 1\right)\right]\sqrt{n}\right] \leq \exp\left[\frac{-2(b^2-a^2)}{\left(1-\frac{4a}{3\sqrt{n}}\right)^2}\right]. \tag{28}$$

In our analysis, we apply Eqs. (27) and (28) to sequences $\xi_1, \ldots, \xi_n$ of Bernoulli random variables, for which $E(\xi_u|\mathcal{F}_{u-1}) = \Pr(\xi_u = 1|\mathcal{F}_{u-1})$.

Now, if we set $a = 0$ on Eqs. (27) and (28), we obtain

$$\begin{aligned} &\Pr\left[\Lambda_n - \sum_{u=1}^{n}\Pr(\xi_u = 1|\mathcal{F}_{u-1}) \geq b\sqrt{n}\right] \leq \exp\left[-2b^2\right],\\ &\Pr\left[\sum_{u=1}^{n}\Pr(\xi_u = 1|\mathcal{F}_{u-1}) - \Lambda_n \geq b\sqrt{n}\right] \leq \exp\left[-2b^2\right] \end{aligned}. \tag{29}$$

This is a slightly improved version of the original Azuma's inequality, whose right-hand side is $\exp\left[-\frac{1}{2}b^2\right]$. Equating the right-hand sides of

Eq. (29) to ε_a, and solving for b, we have that

$$\begin{aligned}&\sum_{u=1}^{n}\Pr(\xi_u=1|\xi_1,\ldots,\xi_{u-1})\leq\Lambda_n+\Delta,\\&\Lambda_n\leq\sum_{u=1}^{n}\Pr(\xi_u=1|\xi_1,\ldots,\xi_{u-1})+\Delta,\end{aligned}\tag{30}$$

with $\Delta=\sqrt{\frac{1}{2}n\ln\varepsilon_a^{-1}}$, and where each of the bounds in Eq. (30) fail with probability at most ε_a.

The bound in Eq. (30) scales with $\sqrt{n}$, and it is only tight when Λ_n is of comparable magnitude to n. When $\Lambda_n\ll n$, one can set a and b in Eq. (27) appropriately to obtain a much tighter bound. To do so, one can use previous knowledge about the channel to come up with a prediction $\tilde{\Lambda}_n$ of Λ_n before running the experiment. Then, one obtains the values of a and b that would minimise the deviation term if the realisation of Λ_n equalled $\tilde{\Lambda}_n$, by solving the optimisation problem

$$\begin{aligned}\min_{a,b}\quad&\left[b+a\left(\frac{2\tilde{\Lambda}_n}{n}-1\right)\right]\sqrt{n}\\\text{s.t.}\quad&\exp\left[\frac{-2(b^2-a^2)}{(1+\frac{4a}{3\sqrt{n}})^2}\right]=\varepsilon_a,\\&b\geq|a|.\end{aligned}\tag{31}$$

The solution to Eq. (31) is

$$\begin{aligned}a&=\frac{3\left(72\sqrt{n}\tilde{\Lambda}_n(n-\tilde{\Lambda}_n)\ln\varepsilon_a-16n^{3/2}\ln^2\varepsilon_a+9\sqrt{2}(n-2\tilde{\Lambda}_n)\sqrt{-n^2\ln\varepsilon_a(9\tilde{\Lambda}_n(n-\tilde{\Lambda}_n)-2n\ln\varepsilon_a)}\right)}{4(9n-8\ln\varepsilon_a)(9\tilde{\Lambda}_n(n-\tilde{\Lambda}_n)-2n\ln\varepsilon_a)},\\b&=\frac{\sqrt{18a^2n-\left(16a^2+24a\sqrt{n}+9n\right)\ln\varepsilon_a}}{3\sqrt{2n}}.\end{aligned}\tag{32}$$

After fixing a and b, we have that

$$\sum_{u=1}^{n}\Pr(\xi_u=1|\xi_1,\ldots,\xi_{u-1})\leq\Lambda_n+\Delta',\tag{33}$$

except with probability ε_a, where

$$\Delta'=\left[b+a\left(\frac{2\Lambda_n}{n}-1\right)\right]\sqrt{n}.\tag{34}$$

In our numerical simulations, we have found the simple bound in Eq. (30) to be sufficiently tight for all components except the vacuum contribution M_{00}. For this latter component, we use Eq. (33) instead. However, note that the users do not know the true value of M_{00}, even after running the experiment. Instead, they will obtain an upper bound M_{00}^{U} on M_{00} via the decoy-state method, and they will apply Eq. (33) to this upper bound. Therefore, to optimise the bound, the users should come up with a prediction $\tilde{M}_{00}^{U}$ on the value of M_{00}^{U} that they expect to obtain after running the experiment and performing the decoy-state analysis, and then substitute $\tilde{\Lambda}_n\to\tilde{M}_{00}^{U}$ in Eq. (31) to obtain the optimal values of a and b. To find $\tilde{M}_{00}^{U}$, one can simply use their previous knowledge of the channel to come up with predictions $\tilde{M}^{\mu\nu}{}^{U}$ of $M^{\mu\nu}$, and run the decoy-state analysis using these values to obtain $\tilde{M}_{00}^{U}$.

DATA AVAILABILITY

All data generated in this study can be reproduced using the equations and methodology introduced in this paper and its Supplementary Notes, and are available from the corresponding author upon reasonable request.

Received: 19 February 2020; Accepted: 12 December 2020;
Published online: 05 February 2021

REFERENCES

1. Scarani, V. et al. The security of practical quantum key distribution. *Rev. Mod. Phys.* **81**, 1301–1350 (2009).
2. Lo, H.-K., Curty, M. & Tamaki, K. Secure quantum key distribution. *Nat. Photonics* **8**, 595–604 (2014).
3. Pirandola, S. et al. Advances in quantum cryptography. *Adv. Opt. Photonics* **12**, 1012–1236, https://doi.org/10.1364/AOP.361502 (2020).
4. Yin, H.-L. et al. Measurement-device-independent quantum key distribution over a 404 km optical fiber. *Phys. Rev. Lett.* **117**, 190501 (2016).
5. Boaron, A. et al. Secure quantum key distribution over 421 km of optical fiber. *Phys. Rev. Lett.* **121**, 190502 (2018).
6. Sangouard, N., Simon, C., de Riedmatten, H. & Gisin, N. Quantum repeaters based on atomic ensembles and linear optics. *Rev. Mod. Phys.* **83**, 33–80 (2011).
7. Pirandola, S., García-Patrón, R., Braunstein, S. L. & Lloyd, S. Direct and reverse secret-key capacities of a quantum channel. *Phys. Rev. Lett.* **102**, 050503 (2009).
8. Takeoka, M., Guha, S. & Wilde, M. M. Fundamental rate-loss tradeoff for optical quantum key distribution. *Nat. Commun.* **5**, 5235 (2014).
9. Pirandola, S., Laurenza, R., Ottaviani, C. & Banchi, L. Fundamental limits of repeaterless quantum communications. *Nat. Commun.* **8**, 15043 (2017).
10. Wilde, M. M., Tomamichel, M. & Berta, M. Converse bounds for private communication over quantum channels. *IEEE Trans. Inf. Theory* **63**, 1792–1817 (2017).
11. Pirandola, S. et al. Theory of channel simulation and bounds for private communication. *Quantum Sci. Technol.* **3**, 035009 (2018).
12. Lo, H.-K., Curty, M. & Qi, B. Measurement-device-independent quantum key distribution. *Phys. Rev. Lett.* **108**, 130503 (2012).
13. Panayi, C., Razavi, M., Ma, X. & Lütkenhaus, N. Memory-assisted measurement-device-independent quantum key distribution. *New J. Phys.* **16**, 043005 (2014).
14. Abruzzo, S., Kampermann, H. & Bruß, D. Measurement-device-independent quantum key distribution with quantum memories. *Phys. Rev. A* **89**, 012301 (2014).
15. Azuma, K., Tamaki, K. & Munro, W. J. All-photonic intercity quantum key distribution. *Nat. Commun.* **6**, 10171 (2015).
16. Duan, L.-M., Lukin, M., Cirac, J. & Zoller, P. Long-distance quantum communication with atomic ensembles and linear optics. *Nature* **414**, 413–418 (2001).
17. Piparo, N. L. & Razavi, M. Long-distance trust-free quantum key distribution. *IEEE J. Sel. Top. Quantum Electron.* **21**, 123–130 (2015).
18. Bhaskar, M. K. et al. Experimental demonstration of memory-enhanced quantum communication. *Nature* **580**, 60–64 (2020).
19. Trényi, R., Azuma, K. & Curty, M. Beating the repeaterless bound with adaptive measurement-device-independent quantum key distribution. *New J. Phys.* **21**, 113052 (2019).
20. Lucamarini, M., Yuan, Z., Dynes, J. & Shields, A. Overcoming the rate–distance limit of quantum key distribution without quantum repeaters. *Nature* **557**, 400–403 (2018).
21. Tamaki, K., Lo, H.-K., Wang, W. & Lucamarini, M. Information theoretic security of quantum key distribution overcoming the repeaterless secret key capacity bound. Preprint at https://arxiv.org/abs/1805.05511 (2018).
22. Ma, X., Zeng, P. & Zhou, H. Phase-matching quantum key distribution. *Phys. Rev. X* **8**, 031043 (2018).
23. Minder, M. et al. Experimental quantum key distribution beyond the repeaterless secret key capacity. *Nat. Photonics* **13**, 334–338 (2019).
24. Zhong, X., Hu, J., Curty, M., Qian, L. & Lo, H.-K. Proof-of-principle experimental demonstration of twin-field type quantum key distribution. *Phys. Rev. Lett.* **123**, 100506 (2019).
25. Liu, Y. et al. Experimental twin-field quantum key distribution through sending or not sending. *Phys. Rev. Lett.* **123**, 100505 (2019).
26. Wang, S. et al. Beating the fundamental rate-distance limit in a proof-of-principle quantum key distribution system. *Phys. Rev. X* **9**, 021046 (2019).
27. Lin, J. & Lütkenhaus, N. Simple security analysis of phase-matching measurement-device-independent quantum key distribution. *Phys. Rev. A* **98**, 042332 (2018).
28. Curty, M., Azuma, K. & Lo, H.-K. Simple security proof of twin-field type quantum key distribution protocol. *npj Quantum Inf.* **5**, 64 (2019).
29. Cui, C. et al. Twin-field quantum key distribution without phase postselection. *Phys. Rev. Appl.* **11**, 034053 (2019).
30. Wang, X.-B., Yu, Z.-W. & Hu, X.-L. Twin-field quantum key distribution with large misalignment error. *Phys. Rev. A* **98**, 062323 (2018).
31. Koashi, M. Simple security proof of quantum key distribution based on complementarity. *New J. Phys.* **11**, 045018 (2009).
32. Hwang, W.-Y. Quantum key distribution with high loss: toward global secure communication. *Phys. Rev. Lett.* **91**, 057901 (2003).
33. Lo, H.-K., Ma, X. & Chen, K. Decoy state quantum key distribution. *Phys. Rev. Lett.* **94**, 230504 (2005).
34. Wang, X.-B. Beating the photon-number-splitting attack in practical quantum cryptography. *Phys. Rev. Lett.* **94**, 230503 (2005).
35. Curty, M. et al. Finite-key analysis for measurement-device-independent quantum key distribution. *Nat. Commun.* **5**, 3732 (2014).
36. Tamaki, K., Curty, M., Kato, G., Lo, H.-K. & Azuma, K. Loss-tolerant quantum cryptography with imperfect sources. *Phys. Rev. A* **90**, 052314 (2014).
37. Mizutani, A., Curty, M., Lim, C. C. W., Imoto, N. & Tamaki, K. Finite-key security analysis of quantum key distribution with imperfect light sources. *New J. Phys.* **17**, 093011 (2015).
38. Kato, G. Concentration inequality using unconfirmed knowledge. Preprint at https://arxiv.org/abs/2002.04357 (2020).
39. Azuma, K. Weighted sums of certain dependent random variables. *Tohoku Math. J.* **19**, 357–367 (1967).

40. Maeda, K., Sasaki, T. & Koashi, M. Repeaterless quantum key distribution with efficient finite-key analysis overcoming the rate-distance limit. *Nat. Commun.* **10**, 3140 (2019).
41. Jiang, C., Yu, Z.-W., Hu, X.-L. & Wang, X.-B. Unconditional security of sending or not sending twin-field quantum key distribution with finite pulses. *Phys. Rev. Appl.* **12**, 024061 (2019).
42. Grasselli, F., Navarrete, Á. & Curty, M. Asymmetric twin-field quantum key distribution. *New J. Phys.* **21**, 113032 (2019).
43. Wang, W. & Lo, H.-K. Simple method for asymmetric twin-field quantum key distribution. *New J. Phys.* **22**, 013020 (2020).
44. Lim, C. C. W., Curty, M., Walenta, N., Xu, F. & Zbinden, H. Concise security bounds for practical decoy-state quantum key distribution. *Phys. Rev. A* **89**, 022307 (2014).
45. Serfling, R. Probability inequalities for the sum in sampling without replacement. *Ann. Stat.* **2**, 39–48 (1974).
46. Ben-Or, M., Horodecki, M., Leung, D. W., Mayers, D. & Oppenheim, J. in *Theory of Cryptography Conference*, Vol. 3378, 386–406 (Springer, Heidelberg, 2005).
47. Renner, R. & König, R. in *Theory of Cryptography Conference*, Vol. 3378, 407–425 (Springer, Heidelberg, 2005).
48. Laurenza, R. et al. Tight bounds for private communication over bosonic Gaussian channels based on teleportation simulation with optimal finite resources. *Phys. Rev. A* **100**, 042301 (2019).
49. Marsili, F. et al. Detecting single infrared photons with 93% system efficiency. *Nat. Photonics* **7**, 210–214 (2013).
50. Tomamichel, M., Lim, C. C. W., Gisin, N. & Renner, R. Tight finite-key analysis for quantum cryptography. *Nat. Commun.* **3**, 634 (2012).
51. Jensen, J. Sur les fonctions convexes et les inégalités entre les valeurs moyennes. *Acta Math.* **30**, 175–193 (1906).
52. Zhang, Z., Zhao, Q., Razavi, M. & Ma, X. Improved key-rate bounds for practical decoy-state quantum-key-distribution systems. *Phys. Rev. A* **95**, 012333 (2017).
53. Bahrani, S., Elmabrok, O., Currás Lorenzo, G. & Razavi, M. Wavelength assignment in quantum access networks with hybrid wireless-fiber links. *J. Opt. Soc. Am. B* **36**, B99 (2019).
54. Grasselli, F. & Curty, M. Practical decoy-state method for twin-field quantum key distribution. *New J. Phys.* **21**, 073001 (2019).

ACKNOWLEDGEMENTS

We thank Margarida Pereira, Kiyoshi Tamaki, and Mirko Pittaluga for valuable discussions. We thank Kento Maeda, Toshihiko Sasaki, and Masato Koashi for the computer code used to generate Fig. 5, as well as for insightful discussions. This work was supported by the European Union's Horizon 2020 research and innovation programme under the Marie Sklodowska-Curie grant agreement number 675662 (QCALL). M.C. also acknowledges support from the Spanish Ministry of Economy and Competitiveness (MINECO), and the Fondo Europeo de Desarrollo Regional (FEDER) through the grant TEC2017-88243-R. K.A. thanks support, in part, from PRESTO, JST JPMJPR1861. A.N. acknowledges support from a FPU scholarship from the Spanish Ministry of Education. M.R. acknowledges the support of UK EPSRC grant EP/M013472/1. G.K. acknowledges financial support by the JSPS Kakenhi (C) No. 17K05591.

AUTHOR CONTRIBUTIONS

G.C.-L. performed the analytical calculations and the numerical simulations. A.N. constructed the analytical decoy-state estimation method. G.K. derived the security bounds in Appendix A. All the authors contributed to discussing the main ideas of the security proof, checking the validity of the results, and writing the paper.

COMPETING INTERESTS

The authors declare no competing interests.

ADDITIONAL INFORMATION

Supplementary information Supplementary information The online version contains supplementary material available at https://doi.org/10.1038/s41534-020-00345-3.

Correspondence and requests for materials should be addressed to G.C.-L.

Reprints and permission information is available at http://www.nature.com/reprints

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.