Vintage Bit Cryptography

Bruce Christianson and Alex Shafarenko

University of Hertfordshire

We propose to use a Random High-Rate Binary (RHRB) stream for the purpose of key distribution. The idea is as follows. Assume availability of a high-rate (terabits per second) broadcaster sending random content. Members of the key group (*e.g.* {Alice, Bob}) share a weak secret (at least 60 bits) and use it to make a selection of bits from the RHRB stream at an extremely low rate (1 bit out of 10^{16} to 10^{18}). By the time that a strong key of reasonable size has been collected (1,000 bits), an enormous amount of data has been broadcast (10^{19}–10^{21} bits). This is 10^6 to 10^8 times current hard drive capacity, which makes it infeasible for the interceptor (Eve) to store the stream for subsequent cryptanalysis, which is what the interceptor would have to do in the absence of the shared secret. Alternatively Eve could record the selection of bits that correspond to every value of the weak shared secret, which under the above assumptions requires the same or greater amount of storage *i.e.* $2^{60} \times 10^3$. The members of the key group have no need to capture the whole stream, but store only the tiny part of it that is the key. Effectively this allows a pseudo-random sequence generated from a weak key to be leveraged up into a strong genuinely random key.

The stream observation time given a 10Tbit/sec broadcast rate is only 10^6 to 10^8 seconds, or a week to a few months. Over this time the shared secret is not used for any kind of communication and so the only possible threat is insufficient key storage security, which is present in any cryptographic scheme. It is interesting that in our approach the passage of time strengthens the resulting key: the longer we wait before the key is used, the less chance there is that any relevant part of the stream is present in a storage facility anywhere in the world, due to the sheer mass of data. This is, in a way, opposite to the standard assumption of cryptographic strength, that keys becomes weaker with time. Accordingly, we call this system Vintage Bit Cryptography.

It is interesting to note that vintage bits are not a hostage to future technology development: the ability to record more data per unit cost in future has no influence over the present time: vintage bits not recorded now will not become available later. Nor does leaking the weak secret compromise vintage bits obtained earlier, provided the time difference is sufficient to overwhelm the capacity of attacker's stream storage. In particular, schemes such as EKE [2,4] can be used to leverage the initial weak secret into a strong pseudo-random seed without fear that subsequent development of quantum computers (allowing the easy solution of discrete logarithm puzzles) will expose previously obtained vintage bit keys.

Beacon systems have been proposed before [9,12,10], particularly in connection with satellites [13]. A traditional beacon implementation based upon a

B. Christianson et al. (Eds.): Security Protocols 2006, LNCS 5087, pp. 261–265, 2009.

geostationary satellite would make the key distribution system available over a wide area at a very small cost to a consumer. But at present digital broadcast satellites lag far behind optical fibre in terms of bandwidth, transmitting only on the order of 10Gbits/sec, although this rate will increase with the use of higher microwave bands.

A satellite solution which could prove more interesting is a swarm of micro-satellites in a Low-Earth-Orbit (LEO). Such satellites could be equipped with an array of tuned silicon lasers that transmit on a number of wavelengths, and with physical random bit generators that control the lasers. Importantly, no radiation protection is required in this case. Indeed, the spacecraft need not have any processing power since all it broadcasts is random digital noise. LEO satellites could be tiny: less than a cubic decimeter undeployed size, with a small production and deployment cost: space scree (rather than dust).

Anyone with a few tens of thousands of dollars to spare can already have micro-satellites launched using a non-governmental space operator. These satellites can keep orbit for years without thrusters and can maintain their orientation by purely passive means. The overhead passage for one of these craft would last 20-30 min, so a continuous RHRB stream at terabit rates would require a hundred spacecraft or so. Using a polar orbit one can ensure that the continuous stream is available anywhere on the planet, and that the area of consistent observation (where all ground observers can see the same satellites at the same time), is of the order of 1000km across, which makes it quite suitable for European applications in particular. The XORing of streams produced from several satellites launched by mutually distrusting parties eliminates the need to trust any individual craft.

However optical fibres are an attractive alternative to satellites, and our primary interest in this paper is with very high bandwidth fibre-optic beacon systems. The first implementation issue to consider is feasibility.

A single optical fibre can already carry more than 1Tbit/sec with a bit-error rate (BER) better than 10^{-3} using an appropriate combination of Wavelength Division Multiplexing and Optical Time Division Multiplexing. Low BER is a key goal of conventional fibre optic communications, but this very tough restriction is not an issue for us. Transmission errors are easily mitigated against by using a simple protocol based on FEC and cryptographic hash functions:

$$A \longrightarrow B : P|Q$$

where $K = K_1|K_2|K_3$ are the vintage bits recorded by A: K_1 is the eventual shared secret with B, K_2 and K_3 are used as one-time pads;
h is a strong hash function, $P = K_2 \oplus h(K_1)$;
and F is a forward error correction function, $Q = K_3 \oplus F(K_1|K_2)$.

The protocol succeeds if B's calculated value for $h(K_1)$ based on the value of $K_1|K_2$ recovered from Q agrees exactly with the value for $h(K_1)$ recovered from P. Note that the message $P|Q$ can be sent over any open, moderately non-lossy channel: no endpoint authentication is required, and data integrity is an issue only if we are concerned with denial of service attacks. In particular, if the message is broadcast, the identity of Bob need not be revealed.

Because low BER is not a consideration for vintage bit cryptography, we are able to propose the use of cheap optical fibre technology which is not suited to the mainstream communication industry. This provides an attractive (cheap!) alternative to the optical fibre systems already being used for key distribution in industry, which use very low bit rates and quantum technology. These quantum-based systems make eavesdropping detectable, but come at a very high cost [6,7,8,14,15,16]. This form of quantum technology also depends crucially on the physical integrity of the optical cable: it eliminates passive eavesdropping but avoiding the man-in-the-middle attack requires at least a weak form of end-to-end authentication for the side-channel, which imposes constraints similar to the initial sharing of a weak secret in our proposal.

Ensuring the integrity of the communication path from a shared beacon is problematic with fibre-optic technology (in contrast with satellites). One simple possibility in the case of a point-to-point link is to co-locate the beacon with one of the participants (say Alice), as may be done in the quantum key agreement scenario. However a more interesting case is where we wish a single beacon on a fibre optic loop to be shared by all the loop nodes. In this case we would like to reduce the integrity requirement to reliance merely upon the integrity of the beacon itself, and not that of the fibre optic medium.

One possibility in this case is for clients to pre-share a weak secret with the beacon (or more accurately with a co-located trusted server). As they collect vintage bits to share with each other, Alice or Bob uses this weak secret to generate bits shared with the beacon service, over the same observation period and using the same protocol. The protocol between Alice and Bob now succeeds only if the vintage bits shared with the beacon have not been tampered with: if they are correct then the real beacon is the source of the bits shared between Alice and Bob. Otherwise the bits are corked and should not be used. The bits shared with the beacon can be discarded, or used to update the weak secret shared with the beacon. Optionally, the beacon service, since it is trusted anyway, can be used to share an initial secret between Alice and Bob in case they have not already been introduced.

However it may be a disadvantage for a beacon protocol to require per-client state to be kept at the server end, and individual communication between each node and the server along the side channel. An alternative is to use a variation of a Merkle-type protocol [3], combined with an additional lower-bandwidth authenticated broadcast by the server. In this case, whenever Alice and Bob collect vintage bits, at least one of them also takes a larger random sample of the beacon, at a rate of order 1 in 10^8–10^9. The beacon server also certifies (for example by public key signature or hash pre-image [1,11]) a random sample of the broadcast taken at a similar rate, which it publishes following sufficient delay to guard against the possibility of a replay attack. The beacon can sample blocks randomly, rather than individual bits. Alice or Bob can now guard against a false beacon by verifying (say, more than 80% match) sufficiently many of the bits which by chance occur in both server and client samples over the course of the collection period.

The number of shared bits increases linearly with the size of the sample being collected. Sampling at a rate of 1 in 10^8 for a base transmission rate of 10Tbps will thus require the beacon to certify about 1Gbyte per day. (If Alice also samples at the rate of 1 in 10^8 then over 80 Merkle bits will be shared per day.) There would be no technical difficulty for the beacon to send this amount of data down the optical medium given the terabit rate of the system. The beacon sample should be broadcast along with a sufficiently long hash, which is signed for authentication. However there is no real-time restriction on the broadcast of the signed hash, which may take place offline. The clients need to know that the beacon was authentic only before they commit to using the newly collected shared key, which as we indicated above takes a few weeks to a few months. This time scale also makes it feasible to employ authentication based on physical security (*e.g.* the delivery of physically authenticated records on tamper-evident media to the clients' sites) as an alternative.

The trust assumptions in our fibre-optic approach are very limited, and are nearly the same as those of the competing quantum approach: the beacon has to be trusted to be authentically random, and a man-in-the middle attack must be detected by end-to-end use of a weak secret. However we make no assumptions about the physical integrity of the fibre-optic link.

While the idea of cryptographic use of a beacon is not in itself new, previous work has tended to focus upon satellite implementations. The threat model for the fibre optic context introduced here is rather different to that for the satellite, and the ramifications of this should lead to interesting new developments.

References

1. Anderson, R., Bergadano, F., Crispo, B., Lee, J.-H., Manifavas, C., Needham, R.: A new family of authentication protocols. Operating Systems Review 32(4), 9–20 (1998)
2. Bellovin, S.M., Merritt, M.: Encrypted Key Exchange: Password-Based Protocols Secure Against Dictionary Attacks. In: Proceedings of the I.E.E.E. Symposium on Research in Security and Privacy, Oakland (May 1992)
3. Christianson, B., Wheeler, D.: Merkle Puzzles Revisited – Finding Matching Elements between Lists. In: Christianson, B., Crispo, B., Malcolm, J.A., Roe, M. (eds.) Security Protocols 2001. LNCS, vol. 2467, pp. 87–90. Springer, Heidelberg (2002)
4. Christianson, B., Roe, M., Wheeler, D.: Secure Sessions from Weak Secrets. In: Christianson, B., Crispo, B., Malcolm, J.A., Roe, M. (eds.) Security Protocols 2003. LNCS, vol. 3364, pp. 190–205. Springer, Heidelberg (2005)
5. Ding, X., Mazzocchi, D., Tsudik, G.: Experimenting with Server-Aided Signatures. In: Proceedings of Network and Distributed System Security Symposium, NDSS 2002 (2002)
6. Gisin, N., Ribordy, G., Tittel, W., Zbinden, H.: Quantum cryptography. Rev. Mod. Phys. 74, 145–195 (2002)
7. Gobby, S.C., Yuan, Z.L., Shields, A.J.: Quantum key distribution over 122km standard telecom fiber. Appl. Phys. Lett. 84, 3762–3764 (2004)
8. Hughes, R.J., Morgan, G.L., Peterson, C.G.: Quantum key distribution over a 48 km optical fibre network. J. Mod. Phys. 47, 533–547 (2000)

9. Maurer, U.: Conditionally-perfect secrecy and a provably-secure randomized cipher. Journal of Cryptology 5, 53–66 (1992)
10. Cachin, C., Maurer, U.M.: Unconditional Secrecy against Memory-Bounded Adversaries. In: Kaliski Jr., B.S. (ed.) CRYPTO 1997. LNCS, vol. 1294, pp. 292–306. Springer, Heidelberg (1997)
11. Merkle, R.C.: A digital signature based on a conventional encryption function. In: Pomerance, C. (ed.) CRYPTO 1987. LNCS, vol. 293, pp. 369–378. Springer, Heidelberg (1988)
12. Mitchell, C.J.: A storage complexity based analogue of Maurer key establishment using public channels. In: Boyd, C. (ed.) Cryptography and Coding 1995. LNCS, vol. 1025, pp. 84–93. Springer, Heidelberg (1995)
13. Rabin, M., Ding, Y.Z.: Hyper-Encryption and Everlasting Security. In: Alt, H., Ferreira, A. (eds.) STACS 2002. LNCS, vol. 2285, p. 1. Springer, Heidelberg (2002)
14. Wu, B.B., Narimanov, E.E.: A method for secure communications over a public fiber-optical network. Opt. Express 14, 3738–3751 (2006)
15. Yoshizawa, A., Kaji, R., Tsuchida, H.: 10.5 km fiber-optic quantum key distribution at 1550 nm with a key rate of 45 kHz. Japanese J. Appl. Phys. 43, L735–L737 (2004)
16. Yuan, Z., Shields, A.: Continuous operation of a one-way quantum key distribution system over installed telecom fibre. Opt. Express 13, 660–665 (2005)