

# ON THE NECESSITY OF ANAMORPHIC ENCRYPTION FOR COVERT COMMUNICATION

TREY LI

The University of Manchester, Manchester, UK  
 trey.li@manchester.ac.uk

We formulate the problem that motivates anamorphic encryption [1] as *covert communication* under a *supervisor* who demands access to every secret key whose existence is evident to the supervisor.<sup>1,2,3,4</sup> An observation is that any bidirectional channel permitting the transmission of user-chosen independent uniform strings with sufficient aggregate capacity in the respective directions enables public-key covert communication.<sup>5,6</sup> Consider a public-key encryption scheme such that for every fixed message, a public key and a ciphertext encrypting that message under the public key are jointly computationally indistinguishable from two independent uniform strings of the corresponding lengths.<sup>7</sup> Over the permitted channel, Alice sends a public key and Bob returns a ciphertext under that public key. No public information indicates that the secret key exists or that the encryption scheme is being used.<sup>8</sup> Indeed, a single such encryption scheme suffices for all such channels.<sup>9,10</sup> An example is Regev encryption over  $\mathbb{Z}_{2^r}$  with a public key consisting of  $m$  LWE samples satisfying  $m \geq (n+1)r + 2\lambda$  for which decisional LWE and the strong leftover hash lemma imply that the public key and the ciphertext of every fixed message bit are jointly computationally indistinguishable from two independent uniform strings of lengths  $m(n+1)r$  and  $(n+1)r$  respectively with distinguishing advantage at most the decisional LWE advantage plus  $2^{-\lambda-1}$ .

## REFERENCE

- [1] Giuseppe Persiano, Duong Hieu Phan, and Moti Yung. Anamorphic encryption: Private communication against a dictator. In Orr Dunkelman and Stefan Dziembowski, editors, *Advances in Cryptology—EUROCRYPT 2022, Part II*, volume 13276 of Lecture Notes in Computer Science, pages 34–63. Springer, 2022. doi:10.1007/978-3-031-07085-3\_2.

## NOTES

<sup>1</sup> The formulation of anamorphic encryption (AE) imposes requirements beyond those necessary for covert communication (CC) (see Note 2), and may be viewed as CC under additional constraints: its overt public key provides public evidence that the corresponding ordinary secret key exists and may therefore be demanded by the supervisor, whereas the anamorphic double key is intended to leave no such evidence and hence need not be revealed. Our observation is that CC itself admits much simpler solutions once the additional constraints are removed. This raises the question of what additional value AE provides beyond such simple solutions for CC.

<sup>2</sup> Examples of additional constraints imposed by AE beyond CC include scheme-specificity and the focus on pre-existing permitted schemes (see also Notes 9 and 10). The former typically requires finding an anamorphic mode for each underlying scheme, while the latter excludes introducing new schemes that support covert communication. These constraints substantially restrict the design space for achieving covert communication.

<sup>3</sup> Two typical forms of evidence are publicly registered public keys and recognizable ciphertexts sent over public channels.

<sup>4</sup> We impose no restriction on when the supervisor may demand access to a secret key. The only requirement is evidence that the key exists. Once such evidence exists, the supervisor may demand access to the key at any time.

<sup>5</sup> A complete formulation must address further details. For example, side-channel information such as data sizes and transmission frequency should follow normal protocol patterns. If one round does not provide enough capacity, we may require the channel to allow multiple rounds with sufficient aggregate capacity. Here we focus on the possibility of CC, leaving bandwidth optimization open. A relevant clarification is that randomness need not be transmitted in plain form as long as it can be efficiently recovered by the receiver.

<sup>6</sup> More generally, beyond uniformly random strings, for every public-key encryption scheme, any channel that can transmit recoverable information with the same joint distribution as a public key and a ciphertext encrypting a fixed message can serve as a covert channel, provided that doing so does not reveal the existence of the corresponding secret key. How general this class of channels is and whether it approaches the most general possible class remains an open problem. A natural theoretical limit, however, is clear. Any channel supporting CC, including AE, must leave some freedom of choice among the outputs permitted by the supervisor. If the permitted output is uniquely determined by the information available to the supervisor, as in deterministic encryption and deterministic or unique signatures, there is no room to encode covert information without a side channel.

<sup>7</sup> The requirement here is different from pseudorandom public keys and pseudorandom ciphertexts. First, we must consider the joint distribution rather than the two distributions separately. Second, requiring “pseudorandom ciphertexts” alone is not enough. For example, ElGamal ciphertexts may be pseudorandom as tuples of group elements, but need not be uniform over all bit strings of the same length due to the underlying group structure. Nonetheless, ElGamal can still be used with channels that carry recoverable random group elements, such as hash inputs in an ElGamal-based Fujisaki–Okamoto key encapsulation mechanism.

<sup>8</sup> Here we do not give the supervisor access to users’ private local states or other internal private information. Such access would also reveal the double keys used in AE.

<sup>9</sup> In contrast, AE has usually been scheme-specific: an anamorphic mode is typically developed separately for each underlying scheme.

<sup>10</sup> This also raises a question about the assumption in AE [1] that one must stick to existing schemes rather than create new ones out of concern that the adversary would forbid them. In fact, new mechanisms, including both overt channels and covert encryption schemes, may not be easy to forbid. For the former, if the property required for covert communication is shared by a broad class of ordinary channels, as in our case where many channels can carry user-chosen randomness (up to differences in bandwidth), it is difficult to forbid all of them, although an active supervisor may sanitise such channels, which may also be costly. For the latter, if there is no evidence that an encryption scheme is being used, a prohibition against it cannot be effectively enforced, although an active supervisor may sanitise channels that might carry ciphertexts of such a scheme; yet this does not amount to forbidding the encryption scheme itself. The same principle also underlies AE solutions, which assume that the anamorphic encryption algorithm corresponding to the double key cannot be effectively forbidden if its use leaves no evidence.