Preface

This book studies two closely related notions, namely locally decodable codes (LDCs) and private information retrieval (PIR) schemes.

Locally decodable codes are a class of error-correcting codes. Error-correcting codes help ensure reliable transmission of information over noisy channels, as well as reliable storage of information on a medium that may become partially corrupted over time or whose reading device is subject to errors. Such codes allow one to add redundancy, or bit strings, to messages, encoding them into longer bit strings, called codewords, in such a way that the message can still be recovered even if a certain fraction of the codeword bits are corrupted. In typical applications of error-correcting codes the message is first partitioned into small blocks, each of which is then encoded separately. This encoding strategy allows efficient random-access retrieval of the information, since one needs to decode only the portion of data in which one is interested. Unfortunately, this strategy yields poor noise resilience, since when even a single block (out of possibly tens of thousands) is completely corrupted, some information is lost. In view of this limitation, it would seem preferable to encode the whole message into a single codeword of an error-correcting code. Such a solution improves the robustness to noise but is hardly satisfactory, since one needs to look at the whole codeword in order to recover any particular bit of the message (at least when using classical error-correcting codes). Such decoding complexity is prohibitive for today’s massive data sets.

Locally decodable codes are codes that simultaneously provide efficient random-access retrieval and high noise resilience by allowing reliable reconstruction of an arbitrary bit of the message from looking at only a small number of randomly chosen codeword bits. Local decodability comes at the price of a certain loss in terms of code efficiency. Specifically, locally decodable codes require longer codeword lengths than their classical counterparts.

Private information retrieval schemes are cryptographic protocols designed to safeguard the privacy of database users. They allow clients to retrieve records from public databases while completely hiding the identity of the retrieved records from the database owners. The possibility of retrieving database records without revealing their identities to the owner of the database may seem beyond hope. Note, however, that a trivial solution is available: when users want a single record, they can ask for a copy of the whole database. This solution involves enormous communication overhead, however, and is likely to be unacceptable. It turns out that for users who want to keep their privacy fully protected (in the "information-theoretic" sense), this trivial solution is optimal.

Fortunately, this negative result applies only to databases stored on a single server, and not to databases replicated across several servers. In 1995, Chor et al. came up with PIR schemes that enable private retrieval of records from replicated databases, with a nontrivially small amount of communication. In such protocols, users query each server holding the database. The protocol ensures that each individual server (by observing only the query it receives) gets no information about the identity of the items of interest to the user.

In this book, we provide a fresh algebraic look at the theory of locally decodable codes and private information retrieval schemes. We obtain new families of LDCs and PIR schemes that have much better parameters than those of previously known constructions. We also prove some limitations on two server PIR schemes in a restricted setting that covers all currently known schemes.

This version. This book is essentially the same as a dissertation filed with the Massachusetts Institute of Technology in 2007. A few sections and proofs have been expanded to provide more intuition and perspective. In addition, this book has an addendum at the end of every chapter to bring the reader up to date with the various developments in the subjects covered in the book during the period from mid 2007 to mid 2010.

Springer

http://www.springer.com/978-3-642-14357-1

Locally Decodable Codes and Private Information
Retrieval Schemes
Yekhanin, S.
2010, XII, 82 p., Hardcover
ISBN: 978-3-642-14357-1