Abstract

Given the breath-taking pace at which the amount of data being generated on a daily basis is growing, and the keen desire to extract information from this data, there is strong interest within the storage industry, at finding means to efficiently and reliably store this data. Given that individual storage units are prone to failure, data pertaining to a single _le is distributed across storage nodes. Such storage of `Big Data' across a spatially distributed network of nodes, calls for codes than can efficiently handle the issue of node repair. The need for node repair could arise on account of device failure, need for a maintenance reboot, or simply because the node is busy serving other demands. A new branch of coding theory has sprung up in response. Regenerating codes minimize data download during a repair operation, while codes with locality ensure that local operations suffice for node repair. The present thesis makes contributions to both regenerating codes as well as codes with locality.

While the reliable storage of data in disks with minimum storage overhead is a well-studied problem, the problem of node repair is relatively new. From the viewpoint of the storage industry, the cost to repair a failed node can be measured in two distinct ways: firstly, in terms of the number of surviving nodes accessed, and secondly, in terms of the amount of data transmitted over the network to ensure repair of the failed node.

In a regenerating code having parameter set (n; k; d; (_; _);B), and over a _finite _field Fq, a file consisting of B symbols from Fq is encoded into n_ symbols, and the coded symbols are stored in n distinct nodes, each node storing _ symbols. The parameter _ is referred to as sub-packetization level. The entire _le can be recovered from downloading k_ symbols from any set of k nodes. Under the exact-repair (ER) setting that is of interest here, the contents of a failed node are repaired exactly from a total of d_ symbols downloaded from any d helper nodes, each helper node transmitting _ symbols. In the alternative setting of functional repair (FR), the contents of the replacement node can be different from that of the failed node, however, following node repair, the new configuration is still required to satisfy the data-recovery and node repair properties of a regenerating code. Under the FR setting, there is a trade-off known as storage-repair-bandwidth (S-RB) trade-off between the values of _ and d_, that optimize the size of the _le being stored. The two extreme points of the trade-off, one corresponding to the minimum possible value of _ and the other, to the minimum possible value of d_, are known as the minimum-storage-regenerating (MSR) and minimum-bandwidth-regenerating (MBR) points respectively. It has been shown that the extreme points of the ER trade-off coincide with those of the FR trade-off. However, the characterization of trade-off in the case of ER remains an open and challenging problem.

High-rate MSR code constructions the _first problem reported in the thesis, is one of designing exact-repair MSR codes that possess certain additional desirable properties, namely, high rate (i.e., rate greater than half) and a small value of sub-packetization level _. In the present thesis, we construct a family of MSR codes for d = (n − 1) having a fixed rational rate R = (t − 1)=t and a sub-packetization that is polynomial in k for a fixed rate. In the construction, _ varies with k as _ = O(kt). This was the first high-rate MSR code construction to have such small value of sub-packetization.

Canonical codes: Regenerating codes for the interior points of the S-RB trade-off the next pair of results relate to the construction of interior-point regenerating codes. For the parameter set (n; k; d = k), we first construct as our second result, a class of codes referred to as canonical codes, having an auxiliary parameter w, with w in the range n−k < w < k, these codes operate in the region between the MSR point and the MBR point, and perform significantly better than operating points obtained through space-sharing. The canonical code constructed for the case (n; k = n − 1; d = n − 1),

turn out to achieve an interior point of the FR trade-off, thus characterizing for the first time, an interior point of the ER trade-off.

Generalized Canonical codes We next build on top of a canonical code having parameters (n; d; d) and construct codes corresponding to a more general parameter set of the form (n; k < d; d) using two different approaches. In the first approach, leading to what we refer to as non-canonical codes, an exponential expansion in field size is needed while the parameters _ and _ remain the same as in the case of the (n; d; d) code. In the second approach, the value of _ is increased from that of the (n; d; d) code, but no expansion in field size is needed. The resultant codes are referred to as improved layered codes. Both approaches lead to codes that perform much better than the space-sharing line.

Outer bounds on the storage-repair-bandwidth trade-off the next set of results presented here take on the form of two distinct outer bounds for the S-RB trade-off. The first is referred to here as the repair-matrix bound and the second as the improved Muhajir-Tandon bound. The repair-matrix bound lies strictly away from the FR trade-off for every parameter (n; k; d), and thus proves that the ER trade-off is bounded away from the FR trade-off, even when the _le size grows to infinity. This result solved an open problem in the literature. In conjunction with the improved layered codes constructed in the thesis and described above, these outer bounds characterize the ER trade-off for the case when (n; k = 3; d = n 1). This is the first family of parameters for which the ER trade-off is characterized. Furthermore, for the parameter set (n; k = 4; d = n 1), the operating point of the improved layered code coincides with the outer bound, characterizing yet another interior point of the S-RB trade-off.

Codes with hierarchical locality the final set of results relate to the second class of codes developed for efficient node repair, namely, codes with locality. The aim of this class of codes is to reduce the number of nodes accessed while repairing a failed node. In a code with locality having locality parameter r, every symbol can be repaired by accessing r other symbols, where r is typically, much smaller than the dimension k of the code. We introduce in the thesis the notion of hierarchical locality, where the local codes are of varying locality and possess a hierarchical structure. In the case of two-level hierarchical locality, for instance, every symbol is protected both by an inner, short-block-length local code, as well as a middle code of larger block length. The code symbols in the inner code are a subset of the code symbols in the middle code. We first present the case of codes with two-level hierarchy, derive an upper bound on the minimum distance and provide optimal code constructions for a wide parameter range. The minimum-distance bound, and code construction are then generalized for the case of a hierarchical structure having an arbitrary number of levels.