

Data verification and reconciliation with generalized error-control codes*

Mark G. Karpovsky

Lev B. Levitin

Ari Trachtenberg

Reliable Computing Lab
Boston University

Abstract

We consider the problem of data reconciliation, which we model as two separate multi-sets of data that must be reconciled with minimum communication. Under this model, we show that the problem of reconciliation is equivalent to a variant of the graph coloring problem and provide consequent upper and lower bounds on the communication complexity of reconciliation. More interestingly, we show by means of an explicit construction that the problem of reconciliation is, under certain general conditions, equivalent to the problem of finding good error-correcting codes. We show analogous results for the problem of multi-set verification, in which we wish to determine whether two multi-sets are equal using minimum communication. As a result, a wide body of literature in coding theory may be applied to the problems of reconciliation and verification.

Keywords: data reconciliation, verification, general errors, error-correcting codes

*thanksThis work was supported in part by the NSF under grants MIP9630096 and CCR-0133521.

1 Introduction

The problem of reconciling data is inherent to applications that require consistency among distributed information, including diverse examples such as gossip protocols for distributing networked data [1], resource discovery [2], mobile data [3, 4], and sequences of symbols from a given alphabet, such as nucleotide sequences in DNA or amino acids sequences in proteins [5]. In each of these examples, the system needs to determine and, thereafter, reconcile the differences between data stored in physically separate locations.

From the perspectives of scalability and performance, it is important that reconciliations occur with minimum communication, measured both by the number of transmitted bits and by the number of rounds of communication. When data are represented by sets, as can be reasonably modeled for the examples cited above, this problem is known as the *set reconciliation* problem [6, 7]. The *data reconciliation* problem is a natural generalization in which data is represented by multi-sets rather than sets; the case where only a single message of communication is permitted is termed *one-way* reconciliation.

This paper examines the one-way data reconciliation problem within a generalized framework in which differences between multi-sets correspond to evaluations of arbitrary “error” functions. We show that this problem of reconciliation is equivalent to a variation of the problem of graph coloring: second-order coloring or distance-2 coloring. A second-order coloring of a graph assigns colors to vertices in such a way that any two nodes separated by a path of length at most two are colored differently. Applying well known results from graph coloring, we then provide lower and upper bounds on the amount of information that must be sent between two hosts for this type of general reconciliation.

In many practical cases, it is not necessary to reconcile two multi-sets, but merely to determine whether they are in fact the same. This may be the case when testing of a remote device is performed by verification of its signature [8, 9]. Such a determination can often be made with substantially less communication than a full-scale reconciliation. In this context, we consider the problem of *one-way data verification*: verifying that two multi-sets are the same, subject to a given range of possible differences, by communication a single message. Again we show that such verification is equivalent to graph coloring and error detection. We also provide both lower and upper bounds on the amount of information that must be exchanged for data verification.

The main contribution of this work is a constructive connection between generalized error-correcting codes and one-way data reconciliation on the one hand, and between generalized error-detecting codes and one-way data verification on the other hand. More precisely, protocols for reconciling multi-sets that differ by some transformation in a set \mathcal{E} can be converted into codes capable of correcting errors represented by \mathcal{E} . Similarly, under certain general conditions, error-correcting codes for \mathcal{E} can be converted into a corresponding one-way data reconciliation protocol. The communication complexity of such a reconciliation protocol is linked to the size of the corresponding error-correcting code, so that good data reconciliation schemes result from good codes, and *vice versa*. In particular, perfect codes result in optimum one-way data reconciliation protocols. Similar results are provided for verification and error-detection.

1.1 Related work

We can broadly characterize the different techniques based on their model of the differences between two reconciling hosts.

One model involves synchronizing two discrete random variables with some known joint probability distribution using a minimum communication complexity. Witsenhausen [10] followed

by Alon and Orlitsky [11] show a connection between such random variable reconciliation and graph coloring, giving results analogous to those of Section 3.1 and 4.1. In addition, Orlitsky [12] showed how to use linear error-correcting codes for a specific instance of data reconciliation.

Another model involves two hosts reconciling files (or strings) that differ by a bounded number of insertions, deletions, or modifications (collectively: “edits”). The problem of efficient reconciliation under these circumstances, also known as the edit-distance problem [13], has been extensively studied [14, 15] because of its connections to important fields such as file synchronization and pattern recognition. Levenštejn [16] pioneered work in this area by developing error-correcting codes capable of correcting precisely these types of errors. Recently [17] also examined the problem of reconstructing a sequence from several copies distorted with these types of errors.

In our model, data on two reconciling hosts is represented by *multi-sets* that differ by a very general class of differences. Being stored as multi-sets, the data on the two hosts is inherently unindexed, meaning that only the content of the individual data items, and not their relative position, matters. Unlike other models, we also assume that reconciliation is agnostic to the roles of participating hosts, as explained in Section 2. Our results extend the bridge between coding theory and data reconciliation originally started in [12], providing statements of conditioned equivalence and corresponding bounds.

1.2 Organization

We begin in Section 2 with a brief formal introduction of the three problems connected in this paper. Section 3 addresses the general problem of data verification, proving connections between graph coloring, data verification, and error-detection, and describing consequent communication bounds. Section 4 provides analogous results for data reconciliation and error-correction. Finally, in Section 5 we describe several applications.

2 Background

2.1 Graph coloring

Definition 1. A *proper coloring* of a graph G with set of vertices V and edges E is an assignment of colors to each vertex in such a way that the vertices of any edge $e \in E$ are colored differently.

A proper coloring using at most k colors will be called a *k -coloring* of the graph. The *chromatic number* of a graph, denoted $\gamma(G)$, is the minimum integer k for which there exists an k -coloring of G .

Definition 2. A *second-order coloring* of a graph is a proper coloring of a graph with the extra property that no two neighbors of any vertex have the same color.

A second-order coloring of G is also a proper coloring of the square of the graph, which is the graph G^2 obtained from G by additionally connecting with an edge each pair of vertices that are of distance two apart. The minimum number of colors needed to second-order color a graph is the *second order chromatic number* of the graph, denoted $\gamma_2(G)$.

2.2 Error detection and correction.

Consider the module \mathbb{Z}_q^n consisting of all n -dimensional vectors over the ring \mathbb{Z}_q . A q -ary code of length n is simply a subset of the elements of this module.

Definition 3. An *error set* for \mathbb{Z}_q^n is a set $\mathcal{E} = \{e_0, e_1, e_2, \dots, e_{|\mathcal{E}|}\}$ whose elements are functions $e_i : \mathbb{Z}_q^n \rightarrow \mathbb{Z}_q^n$, one of which is the identity function $e_0(x) = x$. If the functions $e_i \in \mathcal{E}$ are all bijections and their inverses e_i^{-1} are also in \mathcal{E} , then we shall call this set *bijective*. If the functions commute with each other, so that $e_i(e_j(x)) = e_j(e_i(x)) \forall x \in \mathbb{Z}_q^n, \forall e_i, e_j \in \mathcal{E}$, we shall call this set *commutative*.

We also generalize the concept of an error-ball around a vector with the following definition.

Definition 4. Given an error set \mathcal{E} and a vector $x \in \mathbb{Z}_q^n$, the \mathcal{E} -*image* of x is defined to be $\mathcal{E}[x] = \{e(x) \mid e \in \mathcal{E}\}$. The \mathcal{E} -*vicinity* of x is a closure of the \mathcal{E} -image defined as $\mathcal{E}(x) = \mathcal{E}[x] \cup \{z \in \mathbb{Z}_q^n \mid e(z) = x, e \in \mathcal{E}\}$. More generally, the \mathcal{E}^k -*vicinity* is defined to be

$$\mathcal{E}^k(x) = \bigcup_{y \in \mathcal{E}^{k-1}(x)} \mathcal{E}(y),$$

where $\mathcal{E}^1(x) = \mathcal{E}(x)$. The definition of \mathcal{E}^k - *image* is defined likewise with brackets replacing parentheses.

Note that if \mathcal{E} is bijective, then $\mathcal{E}^k[x] = \mathcal{E}^k(x)$ for all integers $k \geq 1$.

Definition 5. A code $C \in \mathbb{Z}_q^n$ *detects* the error set \mathcal{E} if $c_i \notin \mathcal{E}[c_j]$ for all $c_i \neq c_j \in C$ and *corrects* \mathcal{E} if

$$\mathcal{E}[c_i] \cap \mathcal{E}[c_j] = \emptyset \text{ for all } c_i \neq c_j \in C.$$

2.3 Set and multi-set reconciliation and verification.

The traditional formalization of the set reconciliation problem is as follows [6, 7]: given a pair of hosts A and B , each with a set (S_A and S_B respectively) of length- b bit-strings and no *a priori* knowledge of the other host's set, how can each host determine the mutual difference of the two sets with a minimal amount of communication.

In general, we may consider data represented as multi-sets whose elements are chosen from a finite, universal set U . A multi-set is defined to be a set whose element multiplicities (but not order) are significant. Thus, every multi-set M whose elements are taken from U may be associated uniquely with a *characteristic vector* $v(M)$ of length $n = |U|$ whose i -th component is j if and only if the i -th element of U occurs j times in M , for some canonical ordering of the elements of U .

We shall generally assume in this paper that the multiplicity of any element is bounded by $q - 1$ so that $v(M) \in \mathbb{Z}_q^{|U|}$. We further limit ourselves to the case where only one of the two hosts needs to determine the multi-set held by the other host, based on information transmitted in one message.

From the following definition, we see that one-way data reconciliation functions are precisely those functions that are injective over any given \mathcal{E} -vicinity.

Definition 6. The function $\sigma : \mathbb{Z}_q^n \rightarrow \Sigma$ is a *one-way data reconciliation function* for an error set \mathcal{E} if there exists a recovery function $R : (\Sigma \times \mathbb{Z}_q^n) \rightarrow \mathbb{Z}_q^n$ reconciling multi-sets that differ

by one of the functions in \mathcal{E} . More precisely, the recovery function must have the property that

$$\forall v_A, v_B \in \mathbb{Z}_q^n, \quad e \in \mathcal{E}, \quad v_A \in \mathcal{E}(v_B) \implies R(\sigma(v_A), v_B) = v_A.$$

The *transmission size* of such a reconciliation function is the the number of signals $|\Sigma|$ that need to be transmitted for reconciliation.

To perform a reconciliation with such a function, host A , which has the multi-set M_A with characteristic vector $v_A = v(M_A)$, would send $\sigma(v_A)$ to host B . By computing $R(\sigma(v_A), v_B)$, host B would then determine the characteristic vector v_A and, consequently, the multi-set M_A . Host B can then trivially complete the reconciliation by sending $\sigma(v_B)$ back to A .

We are also interested in the problem of data verification, due to its connections to set reconciliation and a variety of independent applications such as off-line testing [8] and signature analysis [9]. In these cases, two hosts seek to confirm that they have the same multi-set, subject to a known list \mathcal{E} of possible differences..

Definition 7. A function $\alpha : \mathbb{Z}_q^n \rightarrow \Sigma$ is a *one-way data verification function* for an error set \mathcal{E} if there exists a decision function $\mathcal{D} : (\Sigma \times \mathbb{Z}_q^n) \rightarrow \{0, 1\}$ with the property that

$$\forall v_A, v_B \in \mathbb{Z}_q^n, \quad e \in \mathcal{E}, \quad v_A \in \mathcal{E}(v_B) \text{ and } \mathcal{D}(\alpha(v_A), v_B) = 1 \iff (v_A = v_B).$$

The *transmission size* of such a verification function is the number of signals $|\Sigma|$ that need to be transmitted for verification.

2.4 Differences between reconciliation and error-correction

There is a subtle, but important, difference between the definition of error-correction and that of data reconciliation, especially in the face of a non-bijective set \mathcal{E} . Error-correction presumes side-information about communicating parties: one host has an uncorrupted codeword of the code, and the other host has a corrupted transmission; however, both hosts know who has the uncorrupted codeword and who has the corrupted transmission. In the reconciliation case, neither host is restricted to having a codeword, and, more importantly, no side-information is available about the direction of the differences between the hosts.

To highlight this issue, consider the case of disseminating data in a peer-to-peer fashion along a branch of a multicast tree, where packet losses may occur in any communication. It is clear that any two hosts along this branch will have a subset relationship, meaning that one host will have a subset of the data held by the other. However, without global knowledge of the multicast tree, two arbitrary hosts would not initially know which has a subset of the other. In other words, reconciliation is assumed to be agnostic to the role (*e.g.*, subset or superset) played by the hosts. A similar situation occurs when reconciling output from two circuits, where one has incurred a stuck-at-fault causing several 0's to become 1's or vice-versa (depending on the location and nature of the error). One cannot know, *a priori*, which circuit has failed or the direction of its errors.

3 Data verification, coloring, and error-detection

The data verification problem can be reformulated formally as follows. Consider two hosts A and B with multi-sets M_A and M_B respectively. The goal of verification is to determine whether $M_A = M_B$, subject to the sole *a priori* assumption that M_B is in the \mathcal{E} -vicinity of M_A . The data verification problem is thus to determine the minimum amount of information A should send to B so that B can decide whether or not $M_A = M_B$.

3.1 Graph coloring

Consider a natural graph structure [10] corresponding to a given error set \mathcal{E} .

Definition 8. The *characteristic graph* of an error set \mathcal{E} is the undirected graph $G_{\mathcal{E}} = (V, E)$ whose vertices are characteristic vectors of multi-sets $M \subseteq \mathbb{Z}_q^n$. Any two vertices $v_1, v_2 \in V$ are connected by an edge in this graph iff there exists a non-identity error $e \in \mathcal{E}$ such that $e(v_1) = v_2$ or $v_1 = e(v_2)$.

Theorem 1. *Any proper coloring of $G_{\mathcal{E}}$ generates a one-way data verification function $\alpha(\cdot)$ for the error set \mathcal{E} . Conversely, any verification function $\alpha(\cdot)$ yields a proper coloring of $G_{\mathcal{E}}$. The minimum transmission size required for any such verification is precisely the chromatic number $\gamma(G_{\mathcal{E}})$.*

Theorem 1 follows from Definitions 1, 5, and 7 and the fact that each monochromatic set in a proper coloring may be selected as a *level set* (i.e., a set on which the function takes a constant value) for a corresponding one-way data verification function.

Corollary 1. *The minimum transmission size $T_V(\mathcal{E})$ for a one-way data verification function over an error set \mathcal{E} satisfies the inequalities*

$$T_V(\mathcal{E}) = \gamma(G_{\mathcal{E}}) \leq \max_{v \in \mathbb{Z}_q^n} |\mathcal{E}(v)| \leq 2|\mathcal{E}|. \quad (1)$$

For bijective errors, $T_V(\mathcal{E}) \leq |\mathcal{E}|$.

Proof. The left inequality follows directly from an application of Brook's Theorem [18] to Theorem 1, since the degree of a vertex v in the characteristic graph of \mathcal{E} is upper-bounded by $|\mathcal{E}(v)| - 1$. The right inequality follows from the definition of $\mathcal{E}(v)$. ■

Example 1: Consider the characteristic graph of an error set consisting of all odd-weight translations errors in \mathbb{Z}_2^n . Clearly this graph can be two-colored, indicating that set verification can be done with the transmission of one bit. However, if we simply change the error set to consist of all even-weight translations, then $n - 1$ bits of transmission are needed.

Though a non-optimal proper coloring satisfying the upper bound in (1) can be generated in linear time, practical use of such techniques is severely limited by the fact that the size of the characteristic graph grows exponentially in the size of underlying multi-sets being verified. For certain classes of errors, a more practical approach is based on error-detecting codes, described in the next section.

3.2 Error detection

The following theorem shows that monochromatic vertices in a proper coloring of $G_{\mathcal{E}}$ and level sets of a verification function each produce error-detecting codes for an error set \mathcal{E} .

Theorem 2. *Any one-way data verification function $\alpha(\cdot)$ for an error set \mathcal{E} with transmission size τ generates a code in \mathbb{Z}_q^n which detects \mathcal{E} and has at least $\frac{q^n}{\tau}$ codewords. Moreover, for any \mathcal{E} , the level set of a one-way data verification function for \mathcal{E} is a code $\mathbb{C} \subseteq \mathbb{Z}_q^n$ detecting \mathcal{E} .*

Proof: Let \mathbb{C} be the code corresponding to a level set of a given verification function $\alpha(\cdot)$. Then, from Definition 7, it cannot be that $c_i \in \mathcal{E}(c_j)$ for codewords $c_i \neq c_j$. In other words, \mathbb{C} must detect errors in \mathcal{E} . Since, by assumption, there are τ level sets for α , at least one of

them must correspond to an error-detecting code of size $\geq \frac{q^n}{\tau}$. The second claim in the theorem follows from Definition 5, which implies that coloring codewords monochromatically in $G_{\mathcal{E}}$ will not violate a proper coloring of the graph. \square

Putting together Theorem 2 and Corollary 1 with Brooks' theorem gives the following generalization of the Gilbert-Varshamov bound to our general class of errors; Corollary 4 in a subsequent section provides a similar generalization for error correcting codes.

Corollary 2. *For any error set \mathcal{E} with characteristic graph $G_{\mathcal{E}}$ there exists an error-detecting code \mathbb{C} with number of codewords*

$$|\mathbb{C}| \geq \frac{q^n}{\gamma(G_{\mathcal{E}})} \geq \frac{q^n}{\max_{v \in \mathbb{Z}_q^n} |\mathcal{E}(v)|} \geq \frac{q^n}{2|\mathcal{E}|}.$$

For bijective errors, $|\mathbb{C}| \geq \frac{q^n}{|\mathcal{E}|}$.

The connection between verification and error detection established by Theorem 2 is unidirectional: given a verification function one immediately obtains at least one error-detecting code. The proof of this fact, as presented in Theorem 3, relies upon the notion of vector *orbits*, which partition the space \mathbb{Z}_q^n into equivalence classes.

Definition 9. The *orbit* of an element $v \in \mathbb{Z}_q^n$ under \mathcal{E} is the union of all \mathcal{E}^k -vicinities of v :

$$O(v) = \bigcup_{k=1}^{\infty} \mathcal{E}^k(v). \quad (2)$$

The following important property of orbits follows straightforwardly.

Lemma 1. *If $e_1(v) = e_2(v)$ for $e_1, e_2 \in \mathcal{E}$, for any bijective, commutative error set \mathcal{E} , then the functions $e_1(\cdot)$ and $e_2(\cdot)$ are identical on the entire orbit $O(v)$, that is $e_1(x) = e_2(x) \quad \forall x \in O(v)$.*

We now can state our main theorem for generating a verification function from an error-detecting code.

Theorem 3. *Any non-extendible¹ code $\mathbb{C} \subseteq \mathbb{Z}_q^n$ that detects a bijective and commutative error set \mathcal{E} also generates a one-way data verification function $\alpha(\cdot)$ with transmission size at most $\max_{c \in \mathbb{C}} |\mathcal{E}[c]| = \max_{c \in \mathbb{C}} |\mathcal{E}(c)|$.*

Proof. Because \mathbb{C} is non-extendible and \mathcal{E} is bijective, there exists a partitioning D of \mathbb{Z}_q^n into domains D_i such that $D_i \subseteq \mathcal{E}(c_i)$ for $c_i \in \mathbb{C}$. In addition, Lemma 1 implies that any given orbit O imposes a partitioning of \mathcal{E} into equivalence classes of functions acting on O . As such, we may designate a complete set of distinct error functions \mathcal{E}_i around the orbit of a codeword c_i with the defining property that if $e_1, e_2 \in \mathcal{E}_i$ then $e_1(v) \neq e_2(v) \quad \forall v \in O(c_i)$. It is assumed that both verifying hosts generate the same partitioning D and set of functions \mathcal{E}_i .

Computing the one-way data verification function $\alpha(v)$ then simply requires finding the domain D_i containing v and the error $e \in \mathcal{E}_i$ that maps the codeword of the domain, $c_i \in \mathbb{C}$, onto v (i.e., $e(c_i) = v$ implies $\alpha(v) = e$). Different vectors in the same domain necessarily have different values of $\alpha(\cdot)$. On the other hand, vectors in different domains either have different values of $\alpha(\cdot)$ or else need not be verified, as the following argument shows.

¹A code is extendible if codewords can be added to it without affecting its error-detecting/correcting capability.

Suppose $\alpha(v_i) = \alpha(v_j)$ for vectors $v_i \in D_i$ and $v_j \in D_j$ in different domains (*i.e.*, $i \neq j$). Then the commutativity and bijectivity of \mathcal{E} insures that v_i and v_j cannot differ by $\tilde{e} \in \mathcal{E}$ or else:

$$\begin{aligned} e(v_i) = c_i & \quad e(v_j) = c_j & \quad \tilde{e}(v_i) = v_j \\ e(\tilde{e}(v_i)) & = & \quad \tilde{e}(c_i) \\ e(v_j) & = & \quad \tilde{e}(c_i) \\ c_j & = & \quad \tilde{e}(c_i), \end{aligned}$$

with (3) contradicting the error-detecting capability of \mathbb{C} . Since v_i and v_j do not differ by an function in \mathcal{E} , they are beyond the verification requirements of $\alpha(\cdot)$.

In all, $\alpha(v)$ takes on a unique value for each $v \in \mathcal{E}_i$. However, Lemma 1 insures that $|\mathcal{E}_i| \leq |\mathcal{E}[c_i]|$, giving a maximum transmission size of $\max_{c \in \mathbb{C}} |\mathcal{E}[c]|$. ■

Theorem 3 shows an equivalence of data verification and error detection for bijective and commutative error sets \mathcal{E} ; the following simple example shows that the same result does not hold for non-bijective errors.

Counterexample 1: Consider the space \mathbb{Z}_3 under two error functions: the identity e_0 and the incrementor $e_1(x) = x + 1 \pmod{3}$. The code $\mathbb{C} = \{0\}$ is a non-extendible error-detecting code with $\max_{c \in \mathbb{C}} |\mathcal{E}[c]| = 2$. However, the characteristic graph of the error set cannot be colored with less than 3 colors, corresponding to a transmission size of $3 > \max_{c \in \mathbb{C}} |\mathcal{E}[c]|$ for the best verification function.

4 Data reconciliation, coloring, and error-correction

The problem of data reconciliation is to determine the minimum amount of information host A should send to host B so that B can determine the multi-set M_A held by A . As before, the only *a priori* assumption is that the set M_B held by B is a distortion of M_A via some error in $e \in \mathcal{E}$.

4.1 Graph coloring

Theorem 4. *Any second-order coloring of $G_{\mathcal{E}}$ generates a one-way data reconciliation function $\sigma(\cdot)$ for the error set \mathcal{E} . Conversely, any such reconciliation function $\sigma(\cdot)$ yields a second-order coloring of $G_{\mathcal{E}}$. The minimum transmission size required for reconciliation is precisely the second-order chromatic number $\gamma_2(G_{\mathcal{E}})$.*

The theorem follows from Definitions 2 and 6 and by an association of vertex colors to return values of σ . We may apply Brooks' theorem to the square of the characteristic graph of \mathcal{E} to get the following corollary.

Corollary 3. *The minimum transmission size $T_R(\mathcal{E})$ for a one-way data reconciliation function over an error set \mathcal{E} satisfies the inequalities*

$$\max_{v \in \mathbb{Z}_q^n} |\mathcal{E}(v)| \leq T_R(\mathcal{E}) \leq \max_{v \in \mathbb{Z}_q^n} |\mathcal{E}^2(v)|. \quad (3)$$

It is interesting to note that for certain error sets, such as those that form a group under composition, the transmission size for reconciliation and verification meet (*i.e.*, $T_V(\mathcal{E}) = \max_v |\mathcal{E}(v)| = T_R(\mathcal{E})$).

4.2 Error correction

The following theorem follows analogously to Theorem 2.

Theorem 5. *Any one-way data reconciliation function $\sigma(\cdot)$ for an error set \mathcal{E} with transmission size τ generates a code in \mathbb{Z}_q^n which corrects \mathcal{E} and has at least $\frac{q^n}{\tau}$ codewords. Moreover, each monochromatic set of vertices in a second-order coloring of $G_{\mathcal{E}}$ is a code that corrects \mathcal{E} .*

Applying Theorem 5 and Corollary 3 gives the following result, which reduces to the well-known Gilbert-Varshamov bound for Hamming errors.

Corollary 4. *For any error set \mathcal{E} with characteristic graph $G_{\mathcal{E}}$ there exists an error-correcting code \mathbb{C} with number of codewords*

$$\frac{q^n}{\max_{v \in \mathbb{Z}_q^n} |\mathcal{E}^2(v)|} \leq \frac{q^n}{\gamma_2(G_{\mathcal{E}})} \leq |\mathbb{C}| \leq \frac{q^n}{\max_{v \in \mathbb{Z}_q^n} |\mathcal{E}(v)|}.$$

By Theorem 2, any code that *detects* error set \mathcal{E} is a monochromatic set of vertices in a proper coloring of $G_{\mathcal{E}}$. In contrast with this, not every code that *corrects* \mathcal{E} is a monochromatic set in a second-order coloring of $G_{\mathcal{E}}$, as demonstrated by the following counterexample.

Counterexample 2: Consider \mathbb{Z}_2^3 under the error set $\mathcal{E} = \{e_0, e_1, e_2\}$, where e_0 is the identity function and e_1 and e_2 are given by the following table:

$x =$	000	001	010	011	100	101	110	111
$e_1(x) =$	000	010	100	001	000	110	100	011
$e_2(x) =$	000	000	011	010	101	001	101	110

It is easy to see that code $\mathbb{C} = \{000, 011, 110\}$ corrects the error set \mathcal{E} . However, \mathbb{C} cannot be a monochromatic set in any second-order coloring of the characteristic graph $G_{\mathcal{E}}$ because codewords 000 and 011 are of distance 2 from each other.

As with verification, an error-correcting code for a bijective, commutative error set generates a reconciliation function. In this case, the connection is to a *covering code* for \mathcal{E} , defined to be a code $\mathbb{C} \subseteq \mathbb{Z}_q^n$ such that $\bigcup_{c \in \mathbb{C}} \mathcal{E}(c) = \mathbb{Z}_q^n$.

Theorem 6. *Any covering code $\mathbb{C} \subseteq \mathbb{Z}_q^n$ that corrects a bijective and commutative error set \mathcal{E} generates a one-way data reconciliation function $\sigma(\cdot)$ with transmission size at most $\max_{c \in \mathbb{C}} |\mathcal{E}(c)|$.*

The proof follows similarly to that for Theorem 3. Note that any non-extendible code correcting \mathcal{E} is necessarily a covering code for \mathcal{E}^2 . We may thus formulate the following protocol for reconciliation using any code \mathbb{C} correcting \mathcal{E} .

Protocol 1. (Data reconciliation using a code) *All hosts interested in reconciling must agree upon fixed methods for generating:*

- a code \mathbb{C} correcting \mathcal{E} ,
- a partitioning D of \mathbb{Z}_q^n into domains $D_i \subseteq \mathcal{E}^2(c_i)$, and
- a complete and distinct set of errors $\mathcal{E}_i \subseteq \mathcal{E}$ acting upon each orbit $O(c_i)$.

Host B with characteristic vector v_B can then reconcile with host A with vector v_A as follows:

1. A sends to B the unique value $\sigma(v_A) = e_A$ such that $e_A(c_A) = v_A$ for $c_A \in \mathbb{C}$ and $e_A \in \mathcal{E}_A^2$

2. B finds the unique $e \in \mathcal{E}_B$ with the property that $e_A^{-1}(v_B) = e(c)$ for $c \in \mathbb{C}$. If such an e does not exist, then v_A and v_B are irreconcilable.
3. B determines A 's characteristic vector as $v_A = e^{-1}(v_B)$.

In the case of Hamming errors (i.e., \mathcal{E} consists of $e_i(x) = x + v$ for v of Hamming weight at most some fixed t), one can understand Theorem 6 in terms of the covering radius of a code [19], defined to be the minimum value ρ for which balls of radius ρ will completely cover \mathbb{Z}_q^n .

Corollary 5. *For every non-extendible code \mathbb{C} of length n , covering radius ρ , and minimum distance $2t + 1$, there corresponds a one-way data reconciliation and verification function for Hamming errors with transmission size*

$$\sum_{i=0}^{\rho} \binom{n}{i} (q-1)^i.$$

A consequence of Corollary 5 is that perfect codes produce optimal one-way data reconciliation schemes. For example, the length 23 binary Golay code [20] can be used to produce an optimal scheme for reconciling subsets of \mathbb{Z}_2^{23} .

5 Examples and Applications

5.1 Groups under composition

Our first examples are of error sets that form a group under composition, in which case several of the bounds in the paper meet.

Example 2: (Cyclic shift) Consider the set of errors corresponding to left or right cyclic shifts of up to $t \leq n$ positions of a vector,

$$\mathcal{E}_{\text{cal}} = \{e_i(\langle x_0 x_1 x_2 \dots x_{n-1} \rangle) = \langle x_i x_{1+i} x_{2+i} \dots x_{n-1+i} \rangle \mid -t \leq i \leq t\},$$

where indices of x are taken mod n . For the case where $t \geq \lfloor \frac{n}{2} \rfloor$, the set of shifts of a given vector forms a clique in the characteristic graph $G_{\mathcal{E}_{\text{cal}}}$ so that the minimum transmission size is

$$T_V(\mathcal{E}_{\text{cal}}) = \gamma(G_{\mathcal{E}_{\text{cal}}}) = n = T_R(\mathcal{E}_{\text{cal}}),$$

meaning that the upper bound of Corollary 1 and the lower bound of Corollary 3 is attained. In the alternate case where $t < \lfloor \frac{n}{2} \rfloor$, the transmission size is

$$T_V(\mathcal{E}_{\text{cal}}) = \gamma(G_{\mathcal{E}_{\text{cal}}}) = t + 1 < \max_x |\mathcal{E}_{\text{cal}}(x)| = 2t + 1 = T_R(\mathcal{E}_{\text{cal}}).$$

Example 3: (Permutations) Consider the error set of permutation errors, in which vectors may be perturbed according to any permutation of their bits. This error set, though non-commutative, affords a simple reconciliation scheme derived from a corresponding error-correcting code \mathbb{C}_{Π} . The codewords of \mathbb{C}_{Π} are vectors whose rightmost bits are 1 and remaining bits 0. Reconciling two vectors v_A and v_B merely involves transmitting the permutations π_A and π_B that transform these vectors to respective codewords $c_A, c_B \in \mathbb{C}_{\Pi}$. Having received, π_B , host A can simply compute $\pi_B^{-1}(c_A)$ to determine v_B if the two vectors differ by a permutation error.

Since the characteristic graph for permutation errors consists of disjoint cliques (one for each codeword c_i), the amount of communication required for both verification and reconciliation of permutations of length n vectors is equal to the size of the largest clique:

$$T_V(\mathcal{E}_{\text{perm}}) = T_R(\mathcal{E}_{\text{perm}}) = \binom{n}{\lfloor \frac{n}{2} \rfloor}.$$

5.2 Classical reconciliation

Consider the traditional formulation of the set reconciliation problem, as stated in Section 2.3. This form of set reconciliation has many applications in networking, including gossip protocols [1, 21, 22], resource discovery [2, 6] and synchronization [3, 4].

In the context of this work, we may view classical set reconciliation as a one-way data reconciliation problem under Hamming errors of weight at most t . Corollary 3 shows that the minimum number of bits needed for such reconciliation when $t = n$ is $\log_2(\max_x |\mathcal{E}(x)|) = 2^b$. In other words, the optimal way to surely reconcile two sets in the face of arbitrary additive errors is to transmit the entire set. On the other hand, if a t is known *a priori* then the size of the corresponding error-set is

$$\sum_{i=0}^t \binom{2^b}{i} \xrightarrow{\{\frac{t}{2^b} \rightarrow 0\}} \approx 2^{bt}. \quad (4)$$

The upper bound transmission size of bt bits is essentially achieved in [7] using polynomial interpolation, and can be alternatively achieved by using reconciliation schemes based on BCH codes.

5.3 Page errors

The model of page errors assumes that errors occur only in the same region. Consider, for sake of example, that two hosts each have subsets of \mathbb{Z}_{16} , with page regions defined every four elements; thus the first page contains elements $\{1, 2, 3, 4\}$, the next page contains $\{5, 6, 7, 8\}$, *etc.* The error set $\mathcal{E}_{\text{page}}$ for this model contains all functions that corrupt a single page. For example, a corruption of the first page by a toggling the existence of set elements 1 and 3 is given by

$$e_3(x) = x \oplus 0000\ 0000\ 0000\ 0101$$

where x is the characteristic vector of the corrupted set. Since $|\mathcal{E}(x)| = 61$ for any x , Corollary 1 implies that one-way set verification requires at most 6 bits of communication. One-way set reconciliation, on the other hand, requires $6 \leq T_R(\mathcal{E}_{\text{page}}) \leq \log(|\mathcal{E}^2(X)|) = \log(1411) < 11$ bits of communication. Using extended Reed-Solomon codes of length 4, we can attain $T_V(\mathcal{E}_{\text{page}}) \leq 4$ bits and $T_R(\mathcal{E}_{\text{page}}) \leq 8$ bits of communication.

In general, if sets are chosen as subsets of $\mathbb{Z}_2^{2^b}$ and each page contains p elements any t of which might be corrupted, then the amount of communication needed for reconciliation when $\frac{t}{2^b} \rightarrow 0$ and $\frac{p}{2^b} \rightarrow 0$ is given by

$$\log_2 T_R(\mathcal{E}_{\text{page}}) \leq \log_2 \left[\sum_{i=0}^{2t} (2^p - 1)^i \binom{\frac{2^b}{p}}{i} \right] \approx 2bt \quad (5)$$

bits, whereas a classical function correcting any pt errors would require roughly ptb bits of communication. The bound in 5 corresponds to the bound in [23] based on using Reed-Solomon codes for reconciliation of these types of errors.

5.4 Client-server reconciliation

Consider the case of a client host maintaining a subset of data on a serving host. The error set for verification and reconciliation is the set of unidirectional errors \mathcal{E}_U , that is $e(x) = y, e \in \mathcal{E}_U$ if and only if the set $v^{-1}(x) \subseteq v^{-1}(y)$, where v^{-1} is the mapping transforming characteristic vectors to sets. This error set is a commutative, but not bijective.

Consider the characteristic graph of this error set over binary vectors of length n and for up to t unidirectional errors. The vector of all zeroes is clearly contained in a clique of size t consisting of vectors starting with up to t ones. Applying Corollary 1 we see that verification requires a transmission size of at least $T_V(\mathcal{E}_U) \geq t + 1$ signals. In fact, this lower bound is easily achievable since the server can send to the client the number of entries it has, mod t , from which the client can verify equality.

For reconciliation, we note that the vectors in the ball of errors around the all zero vector must each be colored differently, giving $T_R(\mathcal{E}_U) \geq \sum_{i=0}^t \binom{n}{i}$. Note that this is identical to the lower bound (4) for reconciling under classical errors. Thus, though it is much easier to verify unidirectional errors than classical errors, it is just as hard to reconcile under either error set.

6 Conclusions

In this work we have studied the problems of verification and reconciliation of remote data with a minimum amount of communication. We have demonstrated connections between one-way data reconciliation/verification, error-control codes, and graph coloring over a general error set. In particular, we have described in Section 4 how to transform an arbitrary code that corrects a general class of commutative and bijective errors into an algorithm for data reconciliation, and *vice versa*; similarly in Section 3 we have shown how such an error-detecting code can be used to perform data verification. The quality of the derived error correcting/detecting codes is dependent on the quality of the chosen data reconciliation/verification schemes, and vice versa, with the particular example that perfect codes generate optimal schemes. Thus, we have the following constructive relations over \mathbb{Z}_q^n :

$$\begin{array}{ll} \text{length } n, \text{ size } \frac{q^n}{\tau} \text{ code detecting } \mathcal{E} & \longrightarrow \text{length } n, \text{ transmission size } \geq \tau \text{ verification for } \mathcal{E}; \\ \text{length } n, \text{ size } \geq \frac{q^n}{\tau} \text{ code detecting } \mathcal{E} & \longleftarrow \text{length } n, \text{ transmission size } \tau \text{ verification for } \mathcal{E}. \end{array}$$

$$\begin{array}{ll} \text{length } n, \text{ size } \frac{q^n}{\tau} \text{ covering code for } \mathcal{E} & \longrightarrow \text{length } n, \text{ transmission size } \geq \tau \text{ reconciliation for } \mathcal{E}; \\ \text{length } n, \text{ size } \geq \frac{q^n}{\tau} \text{ code correcting } \mathcal{E} & \longleftarrow \text{length } n, \text{ transmission size } \tau \text{ reconciliation for } \mathcal{E}. \end{array}$$

Nevertheless, there is also a fundamental difference between data communication schemes and error-control for non-bijective or non-commutative errors, in that the former needs to be agnostic to the role of communicating parties whereas the latter cannot be.

Using the connections discovered in this paper, we have also presented the following bounds on the number of signals that need to be transmitted for data verification ($T_V(\mathcal{E})$) and data reconciliation ($T_R(\mathcal{E})$) for an arbitrary set of corruptions \mathcal{E} that form a characteristic graph $G_{\mathcal{E}}$:

$$\gamma(G_{\mathcal{E}}) \leq T_V(\mathcal{E}) \leq \max_x |\mathcal{E}(x)| \tag{6}$$

$$\max_x |\mathcal{E}(x)| \leq \gamma_2(G_{\mathcal{E}}) = T_R(\mathcal{E}) \leq \max_x |\mathcal{E}^2(x)|. \tag{7}$$

Both upper bounds (6) and (7) are constructive in that we have described an explicit means of attaining them for a given error set. As mentioned in Section 1.1, the lower bounds have already appeared in various forms in the literature.

Finally, we have presented a number of examples throughout the work and in Section 5, thereby demonstrating the applicability of this work to such diverse areas as testing, file synchronization, and client-server network updates.

Acknowledgments

The authors wish to thank Alon Orlitsky, Krishnamurthy Viswanathan, and Junan Zhang for their interest and thoughtful discussions. One of the authors would also like to thank Yaron Minsky for stimulating discussions and Alexander Vardy for encouragement.

References

- [1] R. van Renesse, Y. Minsky, and M. Hayden, "A gossip-style failure detection service," in *Middleware '98: IFIP International Conference on Distributed Systems Platforms and Open Distributed Processing*, Nigel Davies, Kerry Raymond, and Jochen Seitz, Eds. 1998, pp. 55–70, Springer Verlag.
- [2] M. Harchol-Balter, T. Leighton, and D. Lewin, "Resource discovery in distributed networks," in *18th Annual ACM-SIGACT/SIGOPS Symposium on Principles of Distributed Computing*, Atlanta, GA, May 1999.
- [3] A. Trachtenberg and D. Starobinski, "Towards global synchronization," *Large Scale Networks workshop*, March 2001, <http://ana.lcs.mit.edu/sollins/LSN-Workshop/papers/>.
- [4] A. Trachtenberg, D. Starobinski, and S. Agarwal, "Fast PDA synchronization using characteristic polynomial interpolation," *Proc. INFOCOM*, June 2002.
- [5] R. Durbin, S. Eddy, A. Krogh, and G. Mitchéson, *Biological sequence analysis*, Cambridge university press, 1998.
- [6] Y. Minsky, A. Trachtenberg, and R. Zippel, "Set reconciliation with nearly optimal communication complexity," Tech. Rep. TR1999-1778, TR2000-1796, TR2000-1813, Cornell University, 2000.
- [7] Y. Minsky, A. Trachtenberg, and R. Zippel, "Set reconciliation with nearly optimal communication complexity," in *International Symposium on Information Theory*, June 2001, p. 232.
- [8] D.P. Siewiorek and R.S. Swarz, *Reliable Computer Systems: Design and Evaluation*, Digital Press, 1992.
- [9] M. G. Karpovsky and P. Nagvajara, "Design of self-diagnostic boards by signature analysis," *IEEE Trans. on Industrial Electronics*, pp. 241–246, May 1989.
- [10] H.S. Witsenhausen, "The zero-error side information problem and chromatic numbers," *IEEE Trans. on Info. Theory*, vol. 22, no. 5, September 1976.
- [11] N. Alon and A. Orlitsky, "Source coding and graphs entropies," *IEEE Transactions on Information Theory*, vol. 42, no. 5, pp. 1329–1339, September 1996.
- [12] Alon Orlitsky, "Interactive communication of balanced distributions and correlated files," *SIAM Journal on Discrete Mathematics*, vol. 6, no. 4, pp. 548–564, November 1993.
- [13] Alon Orlitsky, "Interactive communication: Balanced distributions, correlated files, and average-case complexity," in *Proceedings of the 32nd Annual Symposium on Foundations of Computer Science*, 1991, pp. 228–238.
- [14] G. Cormode, M. Paterson, S.C. Sahinhalp, and U. Vishkin, "Communication complexity of document exchange," *ACM-SIAM Symposium on Discrete Algorithms*, January 2000.

- [15] T. Schwarz, R.W. Bowdidge, and W.A. Burkhard, "Low cost comparisons of file copies," *Proceedings of the International Conference on Distributed Computing Systems*, pp. 196–202, 1990.
- [16] V.I. Levenšteĭn, "Binary codes capable of correcting spurious insertions and deletions of ones," *Problems of Information Transmission*, vol. 1, no. 1, pp. 8–17, 1965.
- [17] V.I. Levenšteĭn, "Efficient reconstruction of sequences," *IEEE Trans. on Info. Theory*, vol. 47, pp. 2–22, January 2001.
- [18] R.L. Brooks, "On colouring the nodes of a network," *Proc. Cambridge Phil. Soc.*, vol. 37, pp. 194–197, 1941.
- [19] G.Cohen, I.Honkala, S.Litsyn, and A.Lobstein, *Covering Codes*, Elsevier, 1997.
- [20] F.J. MacWilliams and N.J.A. Sloane, *The Theory of Error-Correcting Codes*, North-Holland Publishing Company, New York, 1977.
- [21] K. Guo, M. Hayden, R. van Renesse, W. Vogels, and K. P. Birman, "GSGC: An efficient gossip-style garbage collection scheme for scalable reliable multicast," Tech. Rep., Cornell University, December 1997.
- [22] M. Hayden and K. Birman, "Probabilistic broadcast," Tech. Rep., Cornell University, 1996.
- [23] K.A.S. Abdel-Ghaffar and A.E. Abbadi, "An optimal strategy for comparing file copies," *IEEE Transactions on Parallel and Distributed Systems*, vol. 5, no. 1, pp. 87–93, January 1994.