Low-Power Cooling Codes with Efficient Encoding and Decoding

Yeow Meng Chee*, Tuvi Etzion†, Han Mao Kiah*, Alexander Vardy‡, Hengjia Wei*

*School of Physical and Mathematical Sciences, Nanyang Technological University, Singapore

†Department of Computer Science, Technion, Israel

‡Dept. of Electrical and Computer Engineering and the Dept. of Computer Science and Engineering, University of California San Diego, USA

***Abstract*—A class of low-power cooling (LPC) codes, to control simultaneously both the peak temperature and the average power consumption of interconnects, were introduced recently. An (n, t, w)-LPC code is a coding scheme over n wires that (A) avoids state transitions on the t hottest wires (cooling), and (B) limit the number of transitions to w in each transmission (low-power).**

A few constructions for large LPC codes that have efficient encoding and decoding schemes, are given. In particular, when w is fixed, we construct LPC codes of size $(n/w)^{w-1}$ and show that these LPC codes can be modified to correct errors efficiently. We further present a construction for large LPC codes based on a mapping from cooling codes to LPC codes.

.

I. Introduction

Power and heat dissipation have emerged as first-order design constraints for chips, whether targeted for battery-powered devices or for high-end systems. High temperatures have dramatic negative effects on interconnect performance. Power-aware design alone is insufficient to address the thermal challenges, since it does not directly target the spatial and temporal behavior of the operating environment. For this reason, thermally-aware approaches have emerged as one of the most important domains of research in chip design today. Numerous techniques have been proposed to reduce the overall power consumption of on-chip buses (see [1] and the references therein). However, these techniques do not directly address peak temperature minimization.

Recently, Chee et al. [1] introduced several efficient coding schemes to directly control the peak temperature and the average power consumption. Among others, *low-power cooling (LPC) codes* are of particular interest as they control both the peak temperature and the average power consumption simultaneously. Specifically, an (n, t, w)-LPC code is a coding scheme for communication over a bus consisting of n wires, if the scheme has the features:

(A) every transmission does not cause state transitions on the t hottest wires;

(B) the number of state transitions on all the wires is at most w in every transmission.

Using *partial spreads*, Chee et al. [1] constructed LPC codes with efficient encoding and decoding schemes. When $t \leq 0.687n$ and $w \geq (n-t)/2$, these codes achieve optimal asymptotic rates. However, when w is small, the code rates are small and Chee et al. proposed another construction based on *decomposition of the complete hypergraph* into perfect matchings. While the construction results in LPC codes of large size, efficient encoding and decoding algorithms are not known.

In this work, we focus on this regime (w small) and construct LPC codes with efficient encoding and decoding schemes. Specifically, our contributions are as follows.

(I) We propose a method that takes a linear erasure code as input and constructs an LPC code. Using this method, we then construct a family of LPC codes of size $(n/w)^{w-1}$ that attains the asymptotic upper bound $O(n^{w-1})$ when w is fixed. We also use this method to construct a class of LPC codes of size $(n/w)^{w-e-1}$ that is able to correct e substitution errors.

(II) We propose efficient encoding / decoding schemes for the output LPC code. In particular, for the above family of LPC codes, we demonstrate encoding with $O(n)$ multiplications over $\mathbb{F}_q$ and decoding with $O(w^3)$ multiplications over $\mathbb{F}_q$, where $q = n/w$. Furthermore, the related class of LPC codes is able to correct e errors with complexity $O(n^3)$.

(III) A recursive construction for a class of (n, t, w)-LPC codes where $t \geq n/w$.

(IV) A construction for a class of (n, t, w)-LPC codes based on a mapping from (m, t)-cooling codes.

For lack of space, all the proofs of the results will be given in the full version of this paper.

II. Preliminary

Given a positive integer n, the set $\{1, 2, \ldots, n\}$ is abbreviated as $[n]$. The *Hamming weight* of a vector $\boldsymbol{x} \in \mathbb{F}_q^n$, denoted $\mathrm{wt}(\boldsymbol{x})$, is the number of nonzero positions in $\boldsymbol{x}$, while the *support* of $\boldsymbol{x}$ is defined as $\mathrm{supp}(\boldsymbol{x}) \triangleq \{i \in [n] : x_i \neq 0\}$.

A q-ary *code* $\mathcal{C}$ *of length* n is a subset of $\mathbb{F}_q^n$. If $\mathcal{C}$ is a subspace of $\mathbb{F}_q^n$, it is called *linear code*. An $[n, k, d]_q$-code is a linear code with dimension k and minimum distance d.

Definition 1. For n and t, an (n, t)*-cooling code* $\mathbb{C}$ of size M is defined as a collection $\{\mathcal{C}_1, \mathcal{C}_2, \ldots, \mathcal{C}_M\}$, where $\mathcal{C}_1, \mathcal{C}_2, \ldots, \mathcal{C}_M$ are disjoint subsets of $\{0,1\}^n$ satisfying the following property: for any set $S \subseteq [n]$ of size $|S| = t$ and for $i \in [M]$, there exists a vector $\boldsymbol{u} \in \mathcal{C}_i$ with $\mathrm{supp}(\boldsymbol{u}) \cap S = \varnothing$. We refer to $\mathcal{C}_1, \mathcal{C}_2, \ldots, \mathcal{C}_M$ as *codesets* and the vectors in them as *codewords*.

Definition 2. For n, t and w with $t + w \leq n$, an (n, t, w)*-low-power cooling (LPC) code* $\mathbb{C}$ of size M is defined as a collection of codesets $\{\mathcal{C}_1, \mathcal{C}_2, \ldots, \mathcal{C}_M\}$, where $\mathcal{C}_1, \mathcal{C}_2, \ldots, \mathcal{C}_M$ are disjoint subsets of $\{\boldsymbol{u} \in \{0,1\}^n : \mathrm{wt}(\boldsymbol{u}) \leq w\}$ satisfying the following property: for any set $S \subseteq [n]$ of size $|S| = t$ and for $i \in [M]$, there exists a vector $\boldsymbol{u} \in \mathcal{C}_i$ with $\mathrm{supp}(\boldsymbol{u}) \cap S = \varnothing$.

978-1-5386-4780-6/18/$31.00©2018 IEEE

In this paper, we focus on a class of (n,t,w)-LPC codes where every transmission results in exactly w state transitions. We call such codes (n,t,w)-*constant-power cooling (CPC) codes*. In particular, let $J(n,w) \triangleq \{\boldsymbol{u} \in \{0,1\}^n : \mathrm{wt}(\boldsymbol{u}) = w\}$, and $J^+(n,w) \triangleq \{\boldsymbol{u} \in \{0,1\}^n : \mathrm{wt}(\boldsymbol{u}) \leq w\}$. Then an (n,t,w)-CPC code is an (n,t,w)-LPC code such that $\mathcal{C}_i \subseteq J(n,w)$ for $i \in [M]$.

A. Set Systems

For a finite set X of size n, 2^X denotes the collection of all subsets of X, i.e., $2^X \triangleq \{A : A \subseteq X\}$. A *set system* of order n is a pair $(X,\mathcal{B})$, where X is a finite set of n *points* and $\mathcal{B} \subseteq 2^X$. The elements of $\mathcal{B}$ are called *blocks*. A set system $(X,2^X)$ is a *complete set system*. Two set systems $(X,\mathcal{B}_1)$ and $(X,\mathcal{B}_2)$ with the same point set are called *disjoint* if $\mathcal{B}_1 \cap \mathcal{B}_2 = \varnothing$.

A *partial parallel class* of a set system $(X,\mathcal{B})$ is a collection of pairwise disjoint blocks. If a partial parallel class partitions the point set X, it is called *parallel class*. A set system $(X,\mathcal{B})$ is called *resolvable* if the block set $\mathcal{B}$ can be partitioned into parallel classes.

There is a canonical one-to-one correspondence between the set of all codes of length n and the set of all set systems of order n: the coordinates of vectors in $\{0,1\}^n$ correspond to the points in $[n]$, and each vector $\boldsymbol{u} \in \{0,1\}^n$ corresponds to the block $\mathrm{supp}(\boldsymbol{u})$. Thus we may speak of the set system of a code or the code of a set system.

B. Upper Bounds

Given a t-subset S and a vector $\boldsymbol{u} \in \{0,1\}^n$, we shall say that $\boldsymbol{u}$ *avoids* S if $\mathrm{supp}(\boldsymbol{u}) \cap S = \varnothing$. Observing that a vector of weight w avoids exactly $\binom{n-w}{t}$ different t-subsets, we have the following bounds on LPC codes and CPC codes.

Theorem 3. *Let $\mathbb{C}$ be an (n,t,w)-LPC code of size M, then*

$$M \leq \frac{1}{\binom{n}{t}}\left[\sum_{i=0}^{w}\binom{n}{i}\binom{n-i}{t}\right] = \sum_{i=0}^{w}\binom{n-t}{i}.$$

Furthermore, if $\mathbb{C}$ is an (n,t,w)-CPC code, then

$$M \leq \binom{n}{w}\binom{n-w}{t}\Big/\binom{n}{t} = \binom{n-t}{w}.$$

Hence both LPC codes and CPC codes share the same asymptotic upper bound $O(n^w)$.

Next, we improve this upper bound in some parameters. To do so, we need results on Turán systems.

Let $n \geq k \geq r$. Let $|X| = n$ and define $\binom{X}{r}$ be the collection of r-subsets of X. A *Turán (n,k,r)-system* is a set system $(X,\mathcal{B})$ where $|X| = n$ and $\mathcal{B} \subseteq \binom{X}{r}$ such that every k-subset of X contains at least one of the blocks. The *Turán number* $T(n,k,r)$ is the minimum number of blocks in such system. This number is determined only for $r = 2$ and some sporadic examples. De Caen [3] proved the general lower bound.

$$T(n,k,r) \geq \frac{n-k+1}{n-r+1}\cdot\frac{k/r}{\binom{k}{r}}\binom{n}{r}. \tag{1}$$

The following is immediate from the definition of Turán systems.

Proposition 4. *A family of codesets $\{\mathcal{C}_1,\mathcal{C}_2,\ldots,\mathcal{C}_M\}$ is an (n,t,w)-CPC code if and only if the set system of each $\mathcal{C}_i$ is a Turán $(n,n-t,w)$-system and these M set systems are pairwise disjoint.*

Combining the bound (1) and Proposition 4, we have the following upper bound on the size of CPC codes.

Theorem 5. *Let $\mathbb{C}$ be an (n,t,w)-CPC code of size M. Then*

$$M \leq \frac{n-w+1}{t+1}\binom{n-t-1}{w-1}.$$

When $t = \Theta(n)$, we have $(n-w+1)/(t+1) = O(1)$, and so the upper bound for (n,t,w)-CPC codes is improved from $O(n^w)$ to $O(n^{w-1})$.

For an (n,t,w)-LPC code, the number of codesets containing at least one codeword of weight $< w$ is at most $\sum_{i=0}^{w-1}\binom{n}{i}$. Thus, the size of a LPC code is at most $\sum_{i=0}^{w-1}\binom{n}{i} + \frac{n-w+1}{t+1}\binom{n-t-1}{w-1}$, which is $O(n^{w-1})$ when t is proportional to n.

C. Previous Constructions

Chee et. al [1] provided the following construction of LPC / CPC codes.

Proposition 6 (Decomposition of Complete Hypergraphs)**.** *Let $n = (t+1)w$. There exists an (n,t,w)-CPC code of size $\binom{n-1}{w-1}$.*

When w is fixed, we have t proportional to n and the above construction attains the asymptotic upper bound $O(n^{w-1})$. Unfortunately, no efficient encoding and decoding methods are known for this construction and the only known encoding method involves listing all $\binom{n-1}{w-1}$ codesets.

In the same paper, Chee et. al then proposed the following constructions of LPC codes that have efficient coding schemes.

Proposition 7 (Concatenation)**.** *Suppose that $q \leq \sum_{i=0}^{w'}\binom{s}{i}$ and q is a prime power and $t \leq s$.*

(i) If $t+1 \leq m/2$, then there exisits an (ms,t,mw')-LPC code of size q^{m-t-1}.

(ii) If $t+1 \leq m \leq q+1$, then there exisits an (ms,t,mw')-LPC code of size q^{m-t}.

Proposition 8 (Sunflower Construction)**.** *Let $r+t \leq (n+s)/2$. If a linear $[n,s,w+1]_2$-code exists and a linear $[n-t,r,w+1]_2$-code does not exist, then there exists an (n,t,w)-LPC code of size 2^{n-t-r}.*

III. Constant-Power Cooling Codes

In this section, we present our main construction of CPC codes. Our construction is based on the following generalization of Proposition 6.

Proposition 9. *Let $(X,\mathcal{B})$ be a set system of order n with $\mathcal{B}$ partitioned into M paritial parallel classes $\mathcal{P}_1,\mathcal{P}_2,\ldots,\mathcal{P}_M$. If $\mathcal{B} \subseteq \binom{X}{w}$ and each $\mathcal{P}_i$ has at least $t+1$ blocks, then the codesets $\mathcal{P}_1,\mathcal{P}_2,\ldots,\mathcal{P}_M$ constitute an (n,t,w)-CPC code.*

Since a decomposition of the complete k-uniform hypergraph with vertex set X is a resolvable set system $(X,\binom{X}{k})$, we recover Proposition 6.

A. CPC Codes Based on Linear Codes

Let $\mathcal{C}$ be an $[N, K, D]_q$-code. Using the codewords of $\mathcal{C}$, we construct a set system with q^{K-1} partial parallel classes and then Proposition 9 yields an CPC code $\mathbb{D}$. To equip $\mathbb{D}$ with efficient coding, we then utilize the erasure-correcting algorithms of the linear code $\mathcal{C}$. We discuss this in detail in Section III-B.

For a coordinate set I and a vector $\boldsymbol{\sigma} \in \mathbb{F}_q^{|I|}$, we say $\boldsymbol{\sigma}$ *appears* λ *times in* $\mathcal{C}$ *at* I if there are λ codewords in $\mathcal{C}$ whose restrictions on I are $\boldsymbol{\sigma}$. Since any two codewords of $\mathcal{C}$ agree in at most $N - D$ positions, we have the following observations.

Lemma 10. *Let $\mathcal{C}$ be an $[N, K, D]_q$-code.*

(i) For any $(N - D + 1)$-subset I and any $\boldsymbol{\sigma} \in \mathbb{F}_q^{N-D+1}$, $\boldsymbol{\sigma}$ appears at most once in $\mathcal{C}$ at I.

(ii) For any $(N - D)$-subset J and any $\boldsymbol{\tau} \in \mathbb{F}_q^{N-D}$, $\boldsymbol{\tau}$ appears at most q times in $\mathcal{C}$ at J.

Lemma 11. *Let $\boldsymbol{G}$ be a generator matrix of an $[N, K, D]_q$-code. Then every $K \times (N - D)$ submatrix of $\boldsymbol{G}$ has rank K or $K - 1$. Furthermore, there is a $K \times (N - D)$ submatrix of rank $K - 1$.*

We are ready to present our main construction.

Construction 1. Let $\boldsymbol{G}$ be a generator matrix of an $[N, K, D]_q$-code $\mathcal{C}$. From Lemma 11, we assume that the last $N - D$ columns of $\boldsymbol{G}$ form a submatrix of rank $K - 1$.

- Partition $\mathcal{C}$ into disjoint codesets $\mathcal{C}_1, \mathcal{C}_2, \ldots, \mathcal{C}_M$ such that $\boldsymbol{u}$ and $\boldsymbol{v}$ are in the same codeset if and only if they agree on the last $N - D$ symbols.
- For $i \in [M]$, we truncate the codewords in $\mathcal{C}_i$ to length w by removing their last $N - w$ symbols. In other words, set $\mathcal{C}'_i = \{\boldsymbol{u}|_{[w]} : \boldsymbol{u} \in \mathcal{C}_i\}$ for $i \in [M]$.
- Let $X = \mathbb{F}_q \times [w]$. We construct the partial parallel classes $\mathbb{D} = \{\mathcal{D}_1, \mathcal{D}_2, \ldots, \mathcal{D}_M\}$, where

$$\mathcal{D}_i = \{\{(x_j, j) : j \in [w]\} \ : \ \boldsymbol{x} = x_1 x_2 \cdots x_w \in \mathcal{C}'_i\}.$$

If $N - D + 1 \leq w \leq D$, then $\mathbb{D}$ is an (n, t, w)-CPC code of size $M = q^{K-1}$ where $n = qw$ and $t \leq q - 1$.

Observe that for Construction 1, in addition to the input code $\mathcal{C}$, we need to find a minimum weight word in $\mathcal{C}$ in order to determine the $K \times (N - D)$-submatrix. As it is difficult to determine the minimum weight word for general linear codes, we focus on certain families of codes where this is well known.

Recall that a linear $[N, K, D]_q$-code is *maximum distance separable (MDS)* if $D = N - K + 1$. If the input code $\mathcal{C}$ is MDS, every K columns of G are linearly independent and so any $K \times (N - D)$ submatrix of G has rank $K - 1$ as $N - D = K - 1$. Therefore, we use any $N - D$ coordinate positions to partition $\mathcal{C}$. It is well known that MDS codes exist for the following all parameters.

Theorem 12 (see [2, Ch.11])**.** *Let q be a prime. If $D \geq 3$, then there exists an $[N, K, D]_q$-MDS code if $N \leq q + 1$ for all q and $2 \leq K \leq q - 1$, except when q is even and $K \in \{3, q - 1\}$, in which case $N \leq q + 2$.*

Setting $N = q + 1$, $K = w$, $D = q - w + 2$ and using an $[N, K, D]_q$-MDS code as the input code, Construction 1 yields the following corollary.

Corollary 13. *If $q = n/w$ is a prime power and $q \geq \max\{2w - 2, t + 1\}$, then there exists an (n, t, w)-CPC code of size $(n/w)^{w-1}$.*

In Corollary 13, when w is fixed, we choose t to be propotional to n. In this case, the codes constructed asymptotically attain the upper bound $O(n^{w-1})$. We also note that in some parameters, these CPC codes are much larger than the LPC codes provided by Propositions 7 and 8.

Example 14. Setting $n = 96$, $w = 6$ and $t = 15$ in Corollary 13 yields a $(96, 15, 6)$-CPC code of size $16^5 = 2^{20}$.

In contrast, suppose we use Proposition 7 to construct an $(96, t, 6)$-LPC code with $t \leq 15$. The largest size $16^5 = 2^{20}$ is obtained by setting $m = 6$, $t = 1$, $s = 16$, $w' = 1$, and $q = 16$. While the resulting $(96, 1, 6)$-LPC code has the same size as the CPC from Corollary 13, the cooling capability of the former is clearly much weaker. Proposition 8, on the other hand, yields an $(96, 15, 6)$-LPC code of size 2^{16} by setting $s = 81$ and $r = 65$.

As before, in some parameters, these CPC codes are much larger than the LPC codes provided by Propositions 7 and 8.

Example 15. There exists an $[17, 8, 9]_9$-code (see [4]). Setting $w = 9$ and $t = 8$ in Construction 1 yields a $(81, 8, 9)$-CPC code of size $9^7 \approx 2^{22.189}$.

In contrast, the largest $(81, 8, 9)$-LPC resulting from Proposition 7 is of size $9 \approx 2^{3.17}$ by setting $m = s = q = 9$, $w' = 1$. Proposition 8, on the other hand, yields an $(81, 8, 9)$-LPC of size 2^{21} by setting $s = 54$ and $r = 52$.

B. Encoding and Decoding Schemes

Now, we discuss the encoding and decoding schemes for the code $\mathbb{D}$ in Construction 1. Let $\boldsymbol{G}$ be a generator matrix of the input $[N, K, D]_q$-linear code $\mathcal{C}$ such that the last $N - D$ columns of $\boldsymbol{G}$ form a matrix $\hat{\boldsymbol{G}}$ of rank $K - 1$. Furthermore, we assume that $\hat{\boldsymbol{G}}$ has the following form.

$$\hat{\boldsymbol{G}} = \begin{pmatrix} \boldsymbol{A} & \boldsymbol{I}_{k-1} \\ 0 \cdots 0 & 0 \cdots 0 \end{pmatrix},$$

where $\boldsymbol{I}_{k-1}$ is the identity matrix of size $K - 1$.

Since the number of codesets is q^{K-1}, we index the codesets in $\mathbb{D}$ using vectors in $\mathbb{F}_q^{K-1}$. For $\boldsymbol{\sigma} \in \mathbb{F}_q^{K-1}$, let $\mathcal{C}_{\boldsymbol{\sigma}}$ be the set of q codewords whose $(K - 1)$-suffix is $\boldsymbol{\sigma}$. Then $\mathcal{C}'_{\boldsymbol{\sigma}}$ and $\mathcal{D}_{\boldsymbol{\sigma}}$ are derived as in Construction 1.

Given a t-subset S of $\mathbb{F}_q \times [w]$ and $\boldsymbol{\sigma} = (\sigma_1, \sigma_2, \ldots, \sigma_{K-1}) \in \mathbb{F}_q^{K-1}$, our objective for encoding is to find $D \in \mathcal{D}_{\boldsymbol{\sigma}}$ such that $D \cap S = \varnothing$. Let $\boldsymbol{\beta}_i$ be the i-th row of $\boldsymbol{G}$ and $\boldsymbol{\beta}_i|_{[w]}$ be the w-prefix of $\boldsymbol{\beta}_i$. First compute

$$\boldsymbol{r} = \sum_{i=1}^{K-1} \sigma_i \boldsymbol{\beta}_i|_{[w]}.$$

Then the codeset $\mathcal{C}'_{\boldsymbol{\sigma}}$ can be determined as follows.

$$\mathcal{C}'_{\boldsymbol{\sigma}} = \{\boldsymbol{r} + \lambda \boldsymbol{\beta}_k|_{[w]} : \lambda \in \mathbb{F}_q\}.$$

Consequently, the codeset $\mathcal{D}_{\boldsymbol{\sigma}}$ can be derived as in Construction 1, and so we can compare all the blocks in $\mathcal{D}_{\boldsymbol{\sigma}}$ with S to find the block D such that $D \cap S = \varnothing$.

Hence, for our encoding, since $n = qw$, we need to perform $O(n)$ multiplications over $\mathbb{F}_q$ to find $\mathcal{D}_{\boldsymbol{\sigma}}$ and $O(tn)$ comparisons to find D.

For the decoding, suppose that we have a codeword $\{(x_1,1),(x_2,2),\ldots,(x_w,w)\}$. Since $w \geq N-D+1$, we can run the erasure correcting algorithm of $\mathcal{C}$ on $(x_1,x_2,\ldots,x_w,?,?,\ldots,?)$ to recover the last $K-1$ symbols, $x_{N-K+2},x_{N-K+3},\ldots,x_N$. Then the message can be decoded as $(x_{N-K+2},x_{N-K+3},\ldots,x_N)$. In particular, if the input code $\mathcal{C}$ is a Reed-Solomon code, by using Lagrange interpolation, we perform $O(w^3)$ multiplications to decode.

IV. Error-Correcting CPC Codes

In this section we consider LPC codes that can correct transmission errors ('0' received as '1', or '1' received as '0'). An (n,w,t)-CPC that is able to correct up to e errors is denoted as an (n,t,w,e)-CPECC code.

We adapt Construction 1 to produce CPECC codes.

Theorem 16. *If the input code $\mathcal{C}$ is an $[N,K,D]_q$-code, then the output code $\mathbb{D}$ of Construction 1 is an (n,t,w,e)-CPECC code of size $M = q^{K-1}$, where $n = qw$, $t \leq q$, $e = w-(N-D)-1$.*

For error-correcting, we focus on a special case, where $\mathcal{C}$ is a Reed-Solomon code and $K = N-D+1 = w-e$.

Construction 2. Let w and e be positive integers and q be a prime power such that $q \geq 2w-e-1$. Let $a_1,a_2,\ldots,a_w,b_1,b_2,\ldots,b_{w-e-1}$ be $2w-e-1$ pairwise distinct elements of $\mathbb{F}_q$.

- For each polynomial $f(X) \in \mathbb{F}_q[X]$, define
$$C_f = \{(f(a_j),j) : j \in [w]\}.$$
- For each $\boldsymbol{\sigma} = (\sigma_1,\sigma_2,\ldots,\sigma_{w-e-1}) \in \mathbb{F}_q^{w-e-1}$, let
$$\mathcal{E}_{\boldsymbol{\sigma}} = \{C_f : f \in \mathbb{F}_q[X], \deg(f) \leq w-e-1, \\ f(b_\ell) = \sigma_\ell \text{ for } \ell \in [w-e-1]\}.$$

Denote the collection of $\mathcal{E}_{\boldsymbol{\sigma}}$ as $\mathbb{E}$. Then $\mathbb{E}$ is an (n,t,w,e)-CPECC code of size q^{w-e-1} where $n = qw$ and $t < q$.

The encoding scheme in Section III-B can be easily adapted for the encoding of the CPECC code $\mathbb{E}$. Now we use Algorithm 1 to illustrate the decoding scheme.

Theorem 17. *Algorithm 1 is correct. In other words, suppose that $\boldsymbol{c} \in \mathbb{E}$ and $\boldsymbol{u}$ is the received word from $\boldsymbol{c}$ with at most e errors. Then Algorithm 1 returns $\boldsymbol{\sigma} \in \mathbb{F}_q^{w-e-1}$ such that $\boldsymbol{c} \in \mathcal{E}_{\boldsymbol{\sigma}}$.*

The Berlekamp-Welch algorithm corrects errors in time $O(q^3)$ [5], and hence, Algorithm 1 has complexity $O(n^3)$.

V. Recursive Construction

Notice that an (n,t,w)-CPC code resulting from Proposition 6 and Construction 1 has $t < n/w$. In this section, we present a recursive construction that yields (n,t,w)-CPC codes with $t \geq n/w$.

To do so, we revisit Proposition 9. Let $(X,\mathcal{B})$ be a set system of order n with $\mathcal{B} \subseteq \binom{X}{w}$ and $\mathcal{B}$ partitioned into M paritial parallel classes $\mathcal{P}_1,\mathcal{P}_2,\ldots,\mathcal{P}_M$. Suppose further than each $\mathcal{P}_i$ has exactly q blocks.

Let S be a t-subset of X and fix a partial parallel class $\mathcal{P}_i$. If $t \geq q$, it is not always possible to choose a block/codeword in $\mathcal{P}_i$ that avoids S. However, by pigeonhole principle, we can find a block/codeword that contains at most $\lfloor t/q \rfloor$ elements in S. Suppose we use a $(w,\lfloor t/q \rfloor,w')$-LPC code $\mathbb{H}$ to break up this codeword into codewords of weight at most w'. Then we are able to find a block/codeword of weight w' that avoids S.

We formalise this idea in the following recursive construction. To simplify our construction, we specialise our recursive construction akin to Construction 2.

Algorithm 1 Error-Correction for the CPECC codes in Construction 2

Input: a binary word $\boldsymbol{u}$ indexed by $\mathbb{F}_q \times [w]$
Output: a message $\boldsymbol{\sigma} \in \mathbb{F}_q^{w-e-1}$

$C \leftarrow \text{supp}(\boldsymbol{u})$. Note that $C \subseteq \mathbb{F}_q \times [w]$
for each $i \in [w]$ **do**
  if $|\{(y,i) : (y,i) \in C\}| = 1$ **then**
    $y_i \leftarrow y$, where (y,i) is the unique pair in $\{(y,i) : (y,i) \in C\}$;
  else
    $y_i \leftarrow$ '#';
$\boldsymbol{y} \leftarrow (y_1,y_2,\ldots,y_w)$;
$\hat{\boldsymbol{y}} \leftarrow$ the vector obtained by deleting all the '#' from $\boldsymbol{y}$
run the decoding algorithm for Reed-Solomon codes on $\hat{\boldsymbol{y}}$, returning a polynomial $L(x)$ of degree $w-e-1$;
$\boldsymbol{\sigma} \leftarrow (L(b_1),L(b_2),\ldots,L(b_{w-e-1}))$;
return $\boldsymbol{\sigma}$;

Construction 3. Let $\mathbb{H}$ be an (n,t,w)-CPC code of size m. Let $q \geq n+w-1$ be a prime power and choose $a_1,a_2,\ldots,a_n,b_1,b_2,\ldots,b_{w-1}$ be $n+w-1$ distinct elements of $\mathbb{F}_q$.

- Consider the point set $A = \mathbb{F}_q \times [n]$ and set $\mathcal{B} = \{\{(f(a_j),j) : j \in [n]\} : f \in \mathbb{F}_q[x], \deg(f) \leq w-1\}$. As before, using $b_1,b_2,\ldots,b_{w-1}$, we can partition $\mathcal{B}$ into q^{w-1} parallel classes. Label the blocks such that the parallel classes are $\mathcal{P}_i = \{B_{ij} : j \in [q]\}$ for $i \in [q^{w-1}]$.
- Since the block B_{ij} is of size n, we can use B_{ij} to relabel the point set of the (n,t,w)-CPC code $\mathbb{H}$ so that each codeword in $\mathbb{H}$ corresponds to a w-subset of B_{ij}. Then each codeword B_{ij} gives rise to a collection of codesets, denoted as $\mathcal{E}_{ij\ell}$ for $\ell \in [m]$.
- For $(i,\ell) \in [q^{w-l}] \times [m]$, the codeset $\mathcal{E}_{i\ell}$ is given by the union $\bigcup_{j=1}^{q} \mathcal{E}_{ij\ell}$.

Therefore, $\mathbb{E} = \{\mathcal{E}_{i\ell} : i \in [q^{w-1}], \ell \in [m]\}$ is an (nq,tq,w)-LPC code of size mq^{w-1}.

Corollary 18. *Let q be a prime power. If $t+w \leq n$ and $q \geq n+w-1$, then*

(i) there exists an (nq,tq,w)-CPC code of size q^{w-1};
(ii) there exists an (nq,tq,w)-LPC code of size $\sum_{i=0}^{w-1} q^i$.

Example 19. We compare certain CPC codes resulting from Construction 3 and Corollary 18 with the LPC codes resulting from Proposition 8.

(i) Consider the set of five disjoint 3-$(10,4,1)$-designs constructed by Etzion and Hartmann [6]. By taking the complements of the blocks, we obtain a $(10,3,6)$-CPC code of size five. Applying Construction 3 with

$q = 16$, we obtain a $(160, 48, 6)$-CPC code of size $5 \cdot 16^5 \approx 2^{22.322}$.
In contrast, Proposition 8 yields a $(160, 48, 6)$-LPC code of size 2^{17} by setting $s = 137$ and $r = 95$.

(ii) Setting $n = 9$, $t = 2$, $w = 7$, and $q = 16$ in Corollary 18 yields a $(144, 32, 7)$-LPC code of size $\sum_{i=0}^{6} 16^i \approx 2^{24.093}$.
In contrast, Proposition 8 yields a $(144, 32, 7)$-LPC code of size 2^{18} by setting $s = 121$ and $r = 94$.

In the regime where w is fixed and t is proportional to n, we show that the codes in this section are asymptotically larger than those resulting from Proposition 8. The CPC codes resulting from Construction 3 and Corollary 18 attain the asymptotic upper bound $O((nq)^{w-1})$ when w is fixed. In contrast, if we apply Proposition 8 with $s = nq - \lceil \log_2(\sum_{i=0}^{w-1} \binom{nq-1}{i}) \rceil$ (the GV lower bound) and $r = nq - tq - \lfloor \log_2(\sum_{i=0}^{2} \binom{nq-tq}{i}) \rfloor$ (the Hamming upper bound), we obtain an (nq, tq, w)-LPC code of size $O((nq)^{w/2}) = o((nq)^{w-1})$.

VI. LPC Codes from Cooling Codes

In this section we use a novel method to transform cooling codes into low-power cooling codes, while preserving the efficiency of the cooling codes. The construction is based on two mappings.

Given a surjection

$$g : [n] \longrightarrow [m] \ ,$$

called a *coordinates mapping* and an injection

$$\varphi : \mathbb{F}_2^m \longrightarrow J^+(n, w) \ ,$$

the mapping $\varphi(m, n, w)$ is called a *domination mapping* if the following condition holds:

For any binary word $(v_1, v_2, \ldots, v_m)$ of length m and its image $(u_1, u_2, \ldots, u_n)$, under the injection φ, $u_i = 1$ implies that $v_{g(i)} = 1$ for each $1 \leq i \leq n$.

Theorem 20. *If there exists an (m, t)-cooling code $\mathbb{C}$ and a domination mapping $\varphi(m, n, w)$ then the code*

$$\mathbb{C}' \triangleq \{\varphi(x) \ : \ x \in \mathbb{C}\},$$

in an (n, t, w)-LPC code.

Domination mappings are not difficult to find. There are many with efficient encoding and decoding. For example, we have such mappings for the following triples (m, n, w): (5,8,2), (9,15,3), (12,20,4).

The product construction:

Given ℓ domination mappings $\varphi_i(m_i, n_i, w_i)$, $1 \leq i \leq \ell$, and a binary word $(x_1, x_2, \ldots, x_\ell)$, where the length of x_i is m_i, for each $1 \leq i \leq \ell$. The mapping $\varphi(m, n, w)$, defined by

$$\varphi(x_1, x_2, \ldots, x_\ell) = (\varphi_1(x_1), \varphi_2(x_2), \ldots, \varphi_\ell(x_\ell)),$$

is also a domination mapping for $m = \sum_{i=1}^{\ell} m_i$, $n = \sum_{i=1}^{\ell} n_i$, and $w = \sum_{i=1}^{\ell} w_i$.

The main construction is given with specific parameters for simplicity and better understanding, but it can be generalized easily to other sets of parameters.

Construction 4. Assume we are given $w \geq 6$, $m = 3w = 9\alpha + 12\beta$, $n = 5w = 15\alpha + 20\beta$, t, and an (m, t)-cooling code $\mathbb{C}$ with 2^k codesets. We will construct an (n, t, w)-LPC code $\mathbb{C}'$. Let u be the information word of length k and let $\mathbb{C}_u$ be its related codeset in $\mathbb{C}$. The encoder partitions the set of m coordinates into $\alpha + \beta$ subsets, α subsets of size 9 and β subsets of size 12. Similarly, it partitions the set n coordinates of codewords from $\mathbb{C}'$ into $\alpha + \beta$ subsets, α subsets of size 15 and β subsets of size 20. Let $\varphi_1(9, 15, 3)$ and $\varphi_2(12, 20, 4)$ be two domination mappings and let $\varphi(m, n, w)$ be the domination mapping implied by the product construction on α copies of φ_1 and β copies of φ_2. Let g be the coordinates mapping implied by the product construction. Let $T = \{i_1, \ldots, i_t\}$ be the t hottest wires on the n-wire bus and let $T' = \{g(i_1), \ldots, g(j_t)\}$ be a t'-subset of $[m]$, where $t' \leq t$. The encoder finds the vector v in $\mathbb{C}_u$ related to the set T', i.e. v has *zeroes* in the coordinates of T', as required by the encoder of $\mathbb{C}$. Finally the encoder parse v into $v_1 v_2 \ldots v_\alpha v'_1 v'_2 \ldots v'_\beta$, where v_i is of length 9 and v'_i is of length 12. By using the encoding of the mappings $\varphi_1(9, 15, 3)$ and $\varphi_2(12, 20, 4)$, the encoder maps $v_1 v_2 \ldots v_\alpha v'_1 v'_2 \ldots v'_\beta$ to the word $\varphi(v_1 v_2 \ldots v_\alpha v'_1 v'_2 \ldots v'_\beta)$ of the (n, t, w)-LPC code, where each v_i is mapped by φ_1 to a word of length 15 and each v'_i is mapped by φ_2 to a word of length 20.

The decoder is applied in reverse order to generate the information word of length k from a word of length n of the code $\mathbb{C}'$, by first generating a word from $\mathcal{C}$ and after that using the decoder of $\mathbb{C}$.

Note that Construction 4 can be viewed as a modification of the Concatenation construction (See Proposition 7). Construction 4 has an advantage on the Concatenation of Proposition 7 and other constructions with larger size for the same weight w and where the number of hottest wires is t. A complete analysis and comparison between all the constructions will be given in the full version of this paper.

Acknowledgment

Y. M. Chee was supported in part by the Singapore Ministry of Education under grant MOE2017-T3-1-007. T. Etzion was supported in part by the BSF-NSF grant 2016692 Binational Science Foundation (BSF), Jerusalem, Israel, under Grant 2012016. The research of T. Etzion and A. Vardy was supported in part by the National Science Foundation under grant CCF-1719139. The research of Y. M. Chee, H. M. Kiah and H. Wei was supported in part by the Singapore Ministry of Education under grant MOE2015-T2-2-086.

References

[1] Y. M. Chee, T. Etzion, and H. M. Kiah, and A. Vardy, "Cooling codes: thermal-management coding for high-performance interconnects," *IEEE Trans. Information Theory*, to appear, Accepted Oct 2017.

[2] F. J. MacWilliams, N. J. Sloane NJ, *The theory of error-correcting codes*. North-Holland, 1977.

[3] D. De Caen, "Extension of a theorem of Moon and Moser on complete subgraphs," *Ars Combinatoria*, vol. 16, pp. 5–10, 1983.

[4] M. Grassl, "Bounds on the minimum distance of linear codes and quantum codes," online available at http://www.codetables.de.

[5] L. R. Welch and E. R Berlekamp. "Error correction for algebraic block codes," December 30 1986. US Patent 4,633,470.

[6] T. Etzion, and A. Hartman, "Towards a large set of Steiner quadruple systems," *SIAM J. Discrete Math.*, vol. 4, no. 2, pp. 182-195, 1991.