# Cooling Codes: Thermal-Management Coding for High-Performance Interconnects

**Yeow Meng Chee**[x], **Tuvi Etzion**[*], **Han Mao Kiah**[x], **Alexander Vardy**[+]
[x]School of Physical and Mathematical Sciences, Nanyang Technological University, Singapore
[*]Computer Science Department, Technion, Israel Institute of Technology, Haifa 32000, Israel
[+]Department of Electrical and Computer Engineering, University of California, San Diego, CA 92093, USA
*ymchee@ntu.edu.sg*, *etzion@cs.technion.ac.il*, *hmkiah@ntu.edu.sg*, *avardy@ucsd.edu*

***Abstract*— High temperatures have dramatic negative effects on interconnect performance. Numerous techniques have been proposed to reduce the power dissipation of on-chip buses but they fall short of fully addressing the thermal challenges posed by high-performance interconnects. We introduce new efficient coding schemes that directly control the *peak temperature* of a bus by effectively cooling its hottest wires. This is achieved by avoiding state transitions on the hottest wires for as long as necessary until their temperature drops off. At the same time, we reduce the *average power consumption* by ensuring that the total number of state transitions on all the wires is bounded.**

**Our solutions call for redundancy: we use $n > k$ wires to encode a given $k$-bit bus. Therefore, it is important to determine the minimum possible number of wires $n$ needed to encode $k$ bits while satisfying the desired properties. We provide full analysis in each case, and show that the number of additional wires required to cool the $t$ hottest wires is negligible when $k$ is large. Moreover, the resulting encoders and decoders are fully practical. They do not require significant computational overhead and can be implemented without sacrificing a large circuit area.**

## 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. With the migration to process geometries of 65 nm and below, power dissipation has become as important an issue as timing and signal integrity. Aggressive technology scaling results in smaller feature size, greater packing density, increasing microarchitectural complexity, and higher clock frequencies. This is pushing chip level power consumption to the edge. It is not uncommon for on-chip hot spots to have temperatures exceeding 100°C, while inter-chip temperature differentials often exceed 20°C.

Power-aware design alone is insufficient to address this thermal challenge, 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.

High temperatures have dramatic negative effects on circuit behavior, with interconnects being among the most impacted circuit components. This is due, in part, to the ever decreasing interconnect pitch and the introduction of low-$\kappa$ dielectric insulation which has low thermal conductivity. For example, the Elmore delay [7] of an interconnect increases 5% to 6% for every 10°C increase in temperature, whereas the leakage current grows exponentially with temperature [1]. Therefore, minimizing the temperature of interconnects is of paramount importance for thermally-aware design.

### A. Related Work

Numerous encoding techniques have been proposed in the literature to reduce the overall power consumption of both on-chip and off-chip buses [2,12,13,15,16,20]. It is well established [4,11,14,20] that bus power is directly proportional to the product of line capacitance and the average number of state transitions on the bus wires. Thus the general idea is to encode the data transmitted over the bus so as to reduce the average number of transitions. For example, the "bus-invert" code [16] potentially complements the data on all the wires, according to the Hamming distance between consecutive transmissions, to ensure that the total number of state transitions on $n$ bus wires never exceeds $n/2$. Unfortunately, encoding techniques designed to minimize power consumption, do not directly address peak temperature minimization. To reduce the temperature of a wire, it is not sufficient to minimize its average switching activity. Rather, it is necessary to control the *temporal distribution of the state transitions* on the wire. To reduce the peak temperature of an interconnect, it is necessary to exercise such control for *all* of its constituent wires.

Wang et al. [20] proposed an efficient *thermal spreading* encoding scheme that evenly spreads the switching activity among all the wires, using a simple architecture consisting of a shift-register and a crossbar logic.

### B. Our Contributions

In this paper, we introduce new efficient coding schemes that simultaneously control both the *peak temperature* and the *average power consumption* of interconnects. The proposed coding schemes are distinguished from existing state-of-the-art by having some or all of the following features:

**A.** *We directly control the peak temperature of a bus by effectively cooling its hottest wires. This is achieved by avoiding state transitions on the hottest wires for as long as necessary until their temperature decreases.*

**B.** *We reduce the overall power dissipation by guaranteeing that the total number of transitions on the bus wires is below a specified threshold in every transmission.*

**C.** *We combine properties* **A** *and/or* **B** *with coding for improved reliability (e.g., for low-swing signaling), using existing error-correcting codes.*

To achieve these features, we propose to insert at the interface of the bus specialized circuits implementing encoding/decoding functions. The *resulting encoders/decoders are efficient:* they do not require significant computational and area overhead. The encoder requires knowledge of the hottest wires at every transmission. This can be obtained by using an analytical model to estimate current temperatures of wires [17], or to have actual temperature sensors for each wire [5].

978-1-5090-4096-4/17/$31.00 ©2017 IEEE

We consider both adaptive and nonadaptive (memoryless) coding schemes. The advantage of nonadaptive schemes is that they are easier to implement and do not require memory. The disadvantage is that it is not possible to implement Property **A** with nonadaptive encoding. For this reason, most of the coding schemes developed in this paper are adaptive, based on the idea of *differential encoding*. Notably, however, all of our schemes require the encoder and decoder circuits to keep track of *only one* previous (the most recent) transmission.

Unlike the thermal spreading methods of [20], the solutions we propose introduce redundancy: we require $n > k$ wires to encode a given $k$-bit bus. A key consideration is the *area overhead due to the additional* $n-k$ *wires*. Therefore, it is important to determine the theoretically minimum possible number of wires $n$ needed to encode $k$ bits while satisfying the desired properties. We show that the number of additional wires required to satisfy Property **A** becomes negligible when $k$ is large.

### C. Organization

The rest of this paper is organized as follows. The next section gives formulations of the coding problems that result from the thermal-management features we propose to implement. In Section III, we present a nonadaptive coding scheme that combines Property **B** with the thermal spreading approach of [20]. Our constructions in Section III are based on the notions of *anticodes* and *quorum systems*, and use key results from the theory of combinatorial designs. Section IV is devoted to Property **A**: we show how state transitions on the $t$ hottest wires can be avoided by using only $t+1$ additional bus lines. This optimal construction is based on combining *differential coding* with the notion of *spreads* in projective geometry. In Section V, we show how Properties **A**, **B**, and **C** can all be achieved *at the same time*. We design coding schemes that simultaneously control peak temperature and average power consumption in every transmission, while also correcting transmission errors on the bus wires. Owing to lack of space, we omit most of the proofs, the analysis, examples, as well as some more constructions and the asymptotic behaviour of our codes. They appear in an arxiv version [3].

## II. Problem Formulation and Preliminaries

We elaborate upon Properties **A**, **B**, **C** introduced in the previous section. For each of these properties, we characterize the performance of the corresponding coding scheme by a *single integer parameter*. All of our coding schemes use $n > k$ wires to encode a $k$-bit bus. We assume that communication across the bus is synchronous, occurring in clocked cycles called *transmissions*. This leads to the following definition.

**Definition 1.** *Consider a coding scheme for a bus of $n$ wires. Let $t, w, e$ be positive integers less than $n$. We say that the coding scheme has*

- ***Property*** $\mathbf{A}(t)$**:** *if every transmission does not cause state transitions on the $t$ hottest wires;*
- ***Property*** $\mathbf{B}(w)$**:** *if the total number of state transitions on all the wires is at most $w$, in every transmission;*
- ***Property*** $\mathbf{C}(e)$**:** *if up to $e$ transmission errors on the $n$ wires can be corrected.*

*We presume that, at the time of transmission, it is known which $t$ wires are the hottest; Property $\mathbf{A}(t)$ is required to hold assuming that **any** $t$ wires can be designated as the hottest.*

The values of $t$, $w$, $e$ are *design parameters*, to be determined by the specific thermal requirements of specific interconnects. The proposed coding schemes work for any $t$, $w$, $e$.

We view the collective state of the $n$ wires during each transmission as a binary vector $\boldsymbol{x} = (x_1, x_2, \ldots, x_n)$. The set of all such binary vectors is $\mathbb{F}_2^n$ and also the *Hamming $n$-space* $\mathcal{H}(n) = \{0,1\}^n$. Given any $\boldsymbol{x}, \boldsymbol{y} \in \mathcal{H}(n)$, the *Hamming distance* $d(\boldsymbol{x}, \boldsymbol{y})$ is the number of positions where $\boldsymbol{x}$ and $\boldsymbol{y}$ differ. The *Hamming weight* of a vector $\boldsymbol{x} \in \mathcal{H}(n)$, denoted $\mathrm{wt}(\boldsymbol{x})$, is the number of nonzero positions in $\boldsymbol{x}$. Thus $\mathrm{wt}(\boldsymbol{x}) = d(\boldsymbol{x}, \mathbf{0})$.

Conventionally, a binary *code* $\mathbb{C}$ *of length $n$* is simply a subset of $\mathcal{H}(n)$. Given a code $\mathbb{C}$, its *minimum distance* $d(\mathbb{C})$ and *diameter* $\mathrm{diam}(\mathbb{C})$ are defined as follows:

$$d(\mathbb{C}) = \min_{\boldsymbol{x},\boldsymbol{y} \in \mathbb{C}} d(\boldsymbol{x}, \boldsymbol{y}) \quad \text{and} \quad \mathrm{diam}(\mathbb{C}) = \max_{\boldsymbol{x},\boldsymbol{y} \in \mathbb{C}} d(\boldsymbol{x}, \boldsymbol{y})$$

## III. Nonadaptive Low-Power Codes

The encoding schemes considered in this section belong to the *nonadaptive* kind, in that the choice which codeword to transmit across the bus in the current transmission does not depend on codewords that have been transmitted earlier. Such coding schemes are also known as *memoryless*. The advantage of nonadaptive schemes is that they are simpler to implement: they do not need a continuously changing data model, and they do not require memory to track the history of previous transmissions.

In the nonadaptive case, an *$n$-bit coding scheme* for a source $\mathcal{S} \subseteq \mathcal{H}(k)$ is a triple $\mathscr{E} = \langle \mathbb{C}, \mathcal{E}, \mathcal{D} \rangle$, where

1) $\mathbb{C}$ is a binary code of length $n$,
2) $\mathcal{E}\colon \mathcal{S} \to \mathbb{C}$ is a bijective map called an *encoding function*,
3) $\mathcal{D}\colon \mathbb{C} \to \mathcal{S}$ is a bijective map called a *decoding function*, such that $\mathcal{D}(\mathcal{E}(\boldsymbol{u})) = \boldsymbol{u}$ for all $\boldsymbol{u} \in \mathcal{S}$.

Suppose $\boldsymbol{u}, \boldsymbol{v} \in \mathcal{S}$ are two words that are to be communicated across an $n$-bit bus during consecutive transmissions. Then the total switching activity of the bus is $d(\mathcal{E}(\boldsymbol{u}), \mathcal{E}(\boldsymbol{v})) \leqslant \mathrm{diam}(\mathbb{C})$. It follows that the coding scheme *satisfies Property* $\mathbf{B}(w)$ *if and only if* $\mathrm{diam}(\mathbb{C}) \leqslant w$. We call such a code $\mathbb{C}$ an $(n,w)$***-low-power code***.

In this section, we are interested in $(n,w)$-low-power codes that also achieve low peak temperatures by *spreading the switching activities* among the wires as uniformly as possible. In doing so, we follow the analysis and the resulting *thermal spreading* approach of [20]. To quantify the thermal spreading achieved by a given coding scheme $\mathscr{E} = \langle \mathbb{C}, \mathcal{E}, \mathcal{D} \rangle$, we treat the source $\mathcal{S}$ as a random variable taking on values in $\mathcal{H}(k)$, and assume that $\mathcal{S}$ is uniformly distributed. This is a common assumption in bus analysis [14]. If the expected switching activity of the bus wires is uniformly distributed then we say that the code $\mathbb{C}$ is ***thermal-optimal***. This leads to the following problem:

$$\text{Given } n \text{ and } w\text{, determine the maximum size of a thermal-optimal } (n,w)\text{-low-power code.} \tag{1}$$

The size of $\mathbb{C}$ is important because we wish to minimize the *area overhead* introduced by our coding scheme, This overhead is largely determined by the number $n-k$ of additional wires that we need to encode.

In a thermal-optimal code $\mathbb{C}$, the number of codewords $(x_1, x_2, \ldots, x_n) \in \mathbb{C}$ having $x_i = 1$ is the same for all $i$. Such codes are said to be ***equireplicate*** in the combinatorics literature. To construct such codes, we need tools from the theory of set systems.

### A. Set Systems

Given a positive integer $n$, the set $\{1, 2, \ldots, n\}$ is abbreviated as $[n]$. For a finite set $X$ and $k \leqslant |X|$, we let

$$2^X = \{A : A \subseteq X\} \quad \text{and} \quad \binom{X}{k} = \{A \in 2^X : |A| = k\}.$$

A *set system of order n* is a pair $(X, \mathcal{A})$, where $X$ is a finite set of *n points* and $\mathcal{A} \subseteq 2^X$. The elements of $\mathcal{A}$ are called *blocks*. The *replication number* of $x \in X$ is the number of blocks containing $x$. A set system is *equireplicate* if its replication numbers are all equal.

There is a natural one-to-one correspondence between the Hamming $n$-space $\mathcal{H}(n)$ and the set system $([n], 2^{[n]})$ of order $n$. For a vector $\boldsymbol{x} = (x_1, x_2, \ldots, x_n) \in \mathcal{H}(n)$, the *support of* $\boldsymbol{x}$ is defined as $\operatorname{supp}(\boldsymbol{x}) = \{i \in [n] : x_i \neq 0\}$. With this, the positions of vectors in $\mathcal{H}(n)$ correspond to points in $[n]$, each vector $\boldsymbol{x} \in \mathcal{H}(n)$ corresponds to the block $\operatorname{supp}(\boldsymbol{x})$, and $d(\boldsymbol{x}, \boldsymbol{y}) = |\operatorname{supp}(\boldsymbol{x}) \triangle \operatorname{supp}(\boldsymbol{y})|$, where $\triangle$ stands for the symmetric difference. Thus, we may speak of the *set system of a code* or the *code of a set system*.

### B. Thermal-Optimal Low-Power Codes

The set system $([n], \mathcal{A})$ of a thermal-optimal $(n, w)$-low-power code is defined by the following properties:

1) $|A_1 \triangle A_2| \leqslant w$ for all $A_1, A_2 \in \mathcal{A}$, and
2) $([n], \mathcal{A})$ is equireplicate.

It follows that our problem in (1) can be recast as an equivalent problem in extremal set systems, as follows:

*Given $n$ and $w$, determine $T(n, w)$, the maximum size of an equireplicate set system $(X, \mathcal{A})$ of order $n$ such that $|A_1 \triangle A_2| \leqslant w$ for all $A_1, A_2 \in \mathcal{A}$.* (2)

If the equireplication condition is removed, the set system is known as an *anticode of length n and diameter w*. Hence, thermal-optimal low-power codes are *equireplicate anticodes*.

The maximum size of an anticode (not necessarily equireplicate) has been completely determined by Kleitman [9] and Katona [10]. When $w$ is even, the extremal set system is equireplicate, which implies a solution for (1) and (2).

**Corollary 1.**

$$T(n, w) = \sum_{i=0}^{w/2} \binom{n}{i} \qquad \text{when } w \equiv 0 \pmod 2.$$

The situation when $w$ is odd is much more difficult, but we can derive that for all odd $w$, we have:

$$\sum_{i=0}^{\frac{w-1}{2}} \binom{n}{i} \leqslant T(n, w) \leqslant \binom{n-1}{\frac{w-1}{2}} + \sum_{i=0}^{\frac{w-1}{2}} \binom{n}{i}.$$

Next, some exact values of $T(n, w)$ for odd $w$ are established.

**Proposition 2.**

$$T(n, 1) = 1 \quad \text{for } n \geqslant 2,$$
$$T(n, 3) = n + 1 \quad \text{for } n \geqslant 5,$$
$$T(n, n-1) = 2^{n-1} \quad \text{for } n \geqslant 3.$$

For other odd values of $w$, we start with the extremal anticode $\mathcal{A}$ of diameter $w - 1$ and add blocks to $\mathcal{A}$ while maintaining the equireplication requirement. All such blocks must contain exactly $(w+1)/2$ points and any two of them must intersect in at least one point. Interestingly, these properties precisely define a *regular uniform quorum system* of rank $(w+1)/2$. Quorum systems have been studied extensively in the literature — see [19] for a recent survey.

## IV. Cooling Codes

Unfortunately, it is not possible to satisfy Property $\mathbf{A}(t)$ with nonadaptive coding schemes, even for $t = 1$. In this section, we shall see that if the encoder and decoder keep track of just *one previous transmission* then Property $\mathbf{A}(t)$ can be satisfied for any $t$ with only $t + 1$ additional wires if $2(t + 1) \leqslant n$, by using spreads, a notion from projective geometry. If $t + 1 > n/2$ we propose a construction which generalizes the one with spreads. We also consider error-correction in our schemes.

### A. Differential Encoding and Decoding

The main idea of our differential encoding method is to encode the data to be communicated across the bus in the *difference* between the current transmission and the previous one.

To explain differential encoding in more detail, let us first consider a nonadaptive coding scheme $\mathscr{E} = \langle \mathbb{C}, \mathcal{E}, \mathcal{D} \rangle$ such as the one described in the previous section. Suppose the encoder $\mathcal{E}$ puts out a sequence of codewords $\boldsymbol{x}_1, \boldsymbol{x}_2, \ldots$ in successive transmissions. We propose to pipe the output of $\mathcal{E}$ into a *differential encoder* $\mathcal{E}_\Delta$ that, in response, puts out the sequence of codewords $\boldsymbol{0}, \boldsymbol{y}_1, \boldsymbol{y}_2, \ldots$ defined as follows:

$$\boldsymbol{y}_i = \boldsymbol{y}_{i-1} + \boldsymbol{x}_i \qquad \text{for } i = 1, 2, \ldots$$

where the addition is in $\mathbb{F}_2^n$. At the receiving end, we prepend the decoder $\mathcal{D}$ with the *differential decoder* $\mathcal{D}_\Delta$ which reconstructs the original codeword sequence $\boldsymbol{x}_1, \boldsymbol{x}_2, \ldots$ from the sequence $\boldsymbol{y}_1, \boldsymbol{y}_2, \ldots$ as follows:

$$\boldsymbol{x}_i = \boldsymbol{y}_i + \boldsymbol{y}_{i-1} \qquad \text{for } i = 1, 2, \ldots$$

Henceforth, we *always* assume that a coding scheme $\langle \mathbb{C}, \mathcal{E}, \mathcal{D} \rangle$ is augmented with the differential encoder $\mathcal{E}_\Delta$ and decoder $\mathcal{D}_\Delta$.

Differential coding is useful in our context, since in transmitting a codeword $\boldsymbol{x} = (x_1, x_2, \ldots, x_n)$, there is a state transition on wire $i$ if and only if $x_i = 1$, and the total number of transitions is precisely $\operatorname{wt}(\boldsymbol{x})$. This makes it possible to reduce the area overhead significantly beyond the best overhead achievable with nonadaptive schemes. For example, under differential encoding, a code $\mathbb{C}$ satisfies Property $\mathbf{B}(w)$ — and so is an $(n, w)$-low-power code — if and only if $\operatorname{wt}(\boldsymbol{x}) \leqslant w$ for all $\boldsymbol{x} \in \mathbb{C}$. It follows that the thermal-optimal $(n, w)$-low-power code of maximum size is given by

$$J^+(n, w) \stackrel{\text{def}}{=} \{\, \boldsymbol{x} \in \{0,1\}^n : \operatorname{wt}(\boldsymbol{x}) \leqslant w \,\}.$$

This set, is clearly equireplicate and its size is much larger than the size of the largest anticode of diameter $w$.

### B. Definition of Cooling Codes

Even under differential encoding, it is still not possible to satisfy Property $\mathbf{A}(t)$ with conventional binary codes. Consequently, we henceforth modify our notion of a code $\mathbb{C}$ as follows. The elements of $\mathbb{C}$ are *sets of binary vectors* of length $n$, say $C_1, C_2, \ldots, C_M$. We refer to $C_1, C_2, \ldots, C_M$ as *codesets*. We do not require these codesets to be of the same size, but we do require them to be disjoint: $C_i \cap C_j = \varnothing$ for all $i \neq j$. The elements of each codeset $C_i$ are called *codewords*. The goal is to guarantee that no matter which codeset $C_i$ is chosen, for each possible designation of $t$ wires as the hottest, there is at least one codeword in $C_i$ with *zeroes* on the corresponding $t$ positions. This leads to the following definition.

**Definition 2.** *For positive integers $n$ and $t < n$, an $(n,t)$-**cooling code** $\mathbb{C}$ of size $M$ is defined as a set $\{C_1, C_2, \ldots, C_M\}$, where $C_1, C_2, \ldots, C_M$ are disjoint subsets of $\mathcal{H}(n)$ satisfying the following property: for any set $\mathcal{S} \subset [n]$ of size $|\mathcal{S}| = t$ and for all $i \in [M]$, there exists a codeword $\boldsymbol{x} \in C_i$ with $\operatorname{supp}(\boldsymbol{x}) \cap \mathcal{S} = \varnothing$.*

Given the foregoing definition of cooling codes, we also need to modify our notions of an encoding function and a decoding function, introduced in Section III. As before, we assume that the data to be communicated across the bus is represented by a source $\mathcal{S}$ taking on some $M \leqslant 2^k$ values in $\mathcal{H}(k)$. The input to the encoding function $\mathcal{E}$ now comprises, in addition to a word $\boldsymbol{u} \in \mathbb{S}$, also a set $\mathcal{S} \subset [n]$ of size $t$ representing the positions of the $t$ hottest wires. We let

$$\mathcal{C} = C_1 \cup C_2 \cup \cdots \cup C_M.$$

Then the output of the encoding function $\mathcal{E}$ is a vector $\boldsymbol{x} \in \mathcal{C}$ such that $\operatorname{supp}(\boldsymbol{x}) \cap \mathcal{S} = \varnothing$. For every possible $\mathcal{S}$, the function $\mathcal{E}(\cdot, \mathcal{S})$ is a bijective map from $\mathbb{S}$ to $\mathbb{C}$. Since the codesets $C_1, C_2, \ldots, C_M$ are disjoint, this allows the decoding function $\mathcal{D}$ to recover $\boldsymbol{u} \in \mathbb{S}$ from the encoder output $\boldsymbol{x} \in \mathcal{C}$. We summarize the foregoing discussion in the next definition.

**Definition 3.** *For integers $n$ and $t < n$, an $(n,t)$-**cooling coding scheme** for a source $\mathbb{S} \subseteq \mathcal{H}(k)$ is a triple $\mathcal{E} = \langle \mathbb{C}, \mathcal{E}, \mathcal{D} \rangle$, where*

1) *The code $\mathbb{C}$ is an $(n,t)$-cooling code;*
2) *The encoding function $\mathcal{E}$: $\mathbb{S} \times \binom{[n]}{t} \to \mathcal{C}$ is such that for all $\mathcal{S} \subset [n]$ of size $t$ and all $\boldsymbol{u} \in \mathbb{S}$, we have*
$$\operatorname{supp}\big(\mathcal{E}(\boldsymbol{u}, \mathcal{S})\big) \cap \mathcal{S} = \varnothing\,;$$
3) *The decoding function $\mathcal{D}$: $\mathcal{C} \to \mathbb{S}$ is such that for all $\boldsymbol{u} \in \mathbb{S}$, we have $\mathcal{D}\big(\mathcal{E}(\boldsymbol{u}, \mathcal{S})\big) = \boldsymbol{u}$ regardless of the value of $\mathcal{S}$.*

### C. Bounds on the Size of Cooling Codes

In this subsection, we show that realizing an $(n,t)$-cooling coding scheme requires at least $t+1$ additional wires. In the next subsection, we present a construction that achieves this bound. Herein, let us begin with the following lemma.

**Lemma 3.** *Let $\mathbb{C}$ be an $(n,t)$-cooling code of size $|\mathbb{C}| = M$. Then*

$$M \leqslant \frac{t!(n-t)!}{n!} \sum_{w=0}^{n-t} \binom{n}{w} \binom{n-w}{t} = 2^{n-t}.$$

An $(n,t)$-cooling code $\mathbb{C}$ of size $|\mathbb{C}| = 2^{n-t}$ would be *perfect*. But, it can be proved that such cooling codes do not exist, unless $t = 1$ or $t = n-1$ Therefore,

**Theorem 4.** *If $1 < t < n-1$, then the size of any $(n,t)$-cooling code $\mathbb{C}$ is bounded by $|\mathbb{C}| \leqslant 2^{n-t} - 1$. Consequently, such a code cannot support the transmission of $n-t$ or more bits over an $n$-wire bus.*

### D. Construction of Optimal Cooling Codes

In this subsection, we construct optimal $(n,t)$-cooling codes that support the transmission of up to $n-t-1$ bits, where $2(t+1) \leqslant n$, over an $n$-wire bus. Our construction is based on the notion of *spreads* in projective geometry.

Loosely speaking, a $\tau$-spread of the vector space $\mathbb{F}_2^n$ is a partition of $\mathbb{F}_2^n$ into disjoint subspaces of dimension $\tau$. Formally, a collection $V_1, V_2, \ldots, V_M$ of $\tau$-dimensional subspaces of $\mathbb{F}_2^n$ is said to be a *$\tau$-spread of* $\mathbb{F}_2^n$ if

$$V_i \cap V_j = \{\mathbf{0}\} \quad \text{for all } i \neq j, \text{ and} \qquad (3)$$
$$\mathbb{F}_2^n = V_1 \cup V_2 \cup \cdots \cup V_M.$$

It is well known that such $\tau$-spreads exist if and only if $\tau$ divides $n$, in which case $M = (2^n - 1)/(2^\tau - 1) > 2^{n-\tau}$. For the case where $\tau$ does not divide $n$, *partial spreads* with $M \geqslant 2^{n-\tau}$ have been constructed in [8, Theorem 11]. In what follows, we take $\tau = t+1$. To simplify the terminology, we assume w.l.o.g. that $t+1$ divides $n$ (which is justified by [8]).

**Theorem 5.** *Let $V_1, V_2, \ldots, V_M$ be a $(t+1)$-spread of $\mathbb{F}_2^n$, and define the code $\mathbb{C} = \{V_1^*, V_2^*, \ldots, V_M^*\}$, where $V_i^* = V_i \setminus \{\mathbf{0}\}$ for all $i$. Then $\mathbb{C}$ is an $(n,t)$-cooling code of size $M \geqslant 2^{n-t-1}$.*

*Proof.* It is obvious from (3) that the $M$ codesets of $\mathbb{C}$ are disjoint subsets of $\mathcal{H}(n)$. It remains to verify that for any set $\mathcal{S} \subset [n]$ of size $t$, each of $V_1^*, V_2^*, \ldots, V_M^*$ contains at least one vector whose support is disjoint from $\mathcal{S}$. To this end, consider an arbitrary $(t+1)$-dimensional subspace $V$ of $\mathbb{F}_2^n$, and suppose $\{\boldsymbol{v}_1, \boldsymbol{v}_2, \ldots, \boldsymbol{v}_{t+1}\}$ is a basis for $V$. Let $\boldsymbol{v}_1', \boldsymbol{v}_2', \ldots, \boldsymbol{v}_{t+1}'$ denote the projections of the basis vectors on the $t$ positions in $\mathcal{S}$. These $t+1$ vectors lie in a $t$-dimensional vector space — the projection of $\mathbb{F}_2^n$ on $\mathcal{S}$. Hence, these vectors must be linearly dependent, and there exist binary coefficients $a_1, a_2, \ldots, a_{t+1}$, not all zero, with $a_1\boldsymbol{v}_1' + a_2\boldsymbol{v}_2' + \cdots + a_{t+1}\boldsymbol{v}_{t+1}' = \mathbf{0}$. But then $\boldsymbol{x} = a_1\boldsymbol{v}_1 + a_2\boldsymbol{v}_2 + \cdots + a_{t+1}\boldsymbol{v}_{t+1}$ is a nonzero vector in $V$ whose support does not include any of the positions in $\mathcal{S}$. As this holds for an arbitrary $(t+1)$-dimensional subspace, it must hold for each of the subspaces $V_1, V_2, \ldots, V_M$ in the spread. ∎

Whether we start with a spread or a partial spread, the size of our code $\mathbb{C}$ is usually strictly larger than $2^{n-t-1}$. In order to use such a code to communicate $k = n-t-1$ bits, one can choose a subset of $\mathbb{C}$ arbitrarily. Finally, we note that there are several efficient algorithms for coding into spreads are known. In particular a simple and powerful method was developed by Dumer [6] in the context of coding for memories with defects, and can be used for our codes too.

### E. Cooling Codes for Large t

The cooling codes constructed based on spreads or partial spreads imply that $t+1 \leqslant n/2$. When $t+1 > n/2$ we can use another construction which generates codes of large size. The construction is based on a sunflower whose heart of seeds is a linear code $C$ and his flowers are codes obtain by the sum of $C$ and elements of a spread.

**Theorem 6.** *If $n, t, \ell, r, d$ are integers, such that $r + t \leqslant (n + \ell)/2$, and the following two requirements are satisfied:*

- *there exists an $[n, \ell, d]$ code (a linear code of length $n$, dimension $\ell$, and minimum Hamming distance $d$);*
- *an $[n-t, r, d]$ linear code does not exist,*

*then there exists an $(n,t)$-cooling code of size $M > 2^{n-t-r}$.*

Theorem 6 can be applied in various ways. For example, we can derive the following result.

**Corollary 7.** *If $n = 2^m$, $m > 1$, and $t = (n + m - 1)/2$, then there exists an $(n,t)$-cooling code of size $M > 2^{n-t-2}$.*

Finally, Theorem 6 is generalized with weaker conditions [3].

### F. Error-Correcting Cooling Codes

In this subsection, we construct $(n,t)$-cooling codes that satisfy simultaneously Properties $\mathbf{A}(t)$ and $\mathbf{C}(e)$. The idea is to take your favourite linear error-correcting code $\mathbb{C}$, of length $n$ and dimension $\kappa \geqslant 2(t+1)$, which corrects $e$ errors. In addition we take a $(t+1)$-spread (or partial spread) of $\mathbb{F}_2^\kappa$. Since $\mathbb{C}$ has dimension $\kappa$, there exists at least one set of $\kappa$ coordinates whose projection spans $\mathbb{F}_2^\kappa$. We form the $(t+1)$-spread on these coordinates, and partition the codewords of $\mathbb{C}$ into codesets related to the spread. Given any $t$ coordinates, each codeset has at least one codeword with *zeroes* in these $t$ coordinates since the spread has dimension $t+1$. Moreover, each codeset can correct at least $e$ errors as the code $\mathbb{C}$. Thus, we have an $(n,t)$-cooling code which can correct at least $e$ errors.

**Theorem 8.** *If there exists a binary $[n, \kappa, 2e+1]$ code and $\kappa \geqslant 2(t+1)$, then there exists an $(n,t)$-cooling code which correct at least $e$ errors, whose size $M$ is greater than $2^{\kappa-t-1}$.*

## V. Low-Power Cooling Codes

In this section, we present coding schemes that satisfy Properties $\mathbf{A}(t), \mathbf{B}(w)$ (with or without $\mathbf{C}(e)$) simultaneously in every transmission. We call the corresponding codes $(n,t,w)$***-low-power cooling codes*** (or $(n,t,w)$-LPC codes for short). We suggest two types of constructions. The first is based on resolutions in block designs. The second, which will be presented in the full paper, is based on cooling codes over $\mathbb{F}_q$, dual codes of $[n, \kappa, t+1]$ codes, MDS codes, spreads, and concatenation codes. The results of this construction are summarized in the following two theorems.

**Theorem 9.** *If $q \leqslant \sum_{i=0}^{w'} \binom{s}{i}$ and $t+1 \leqslant m/2$, then there exists an $(ms, t, mw')$-LPC code of size $M > q^{m-t-1}$.*

**Theorem 10.** *If $q \leqslant \sum_{i=0}^{w'} \binom{s}{i}$ and $m \leqslant q+1$, then there exists an $(ms, t, mw')$-LPC code of size $q^{m-t}$.*

As before, we assume that the coding schemes constructed in what follows are augmented by differential encoding. Since the codes are also $(n,t)$-cooling codes, they conform to Definition 2. Thus a code $\mathbb{C}$ is a collection of codesets $C_1, C_2, \ldots, C_M$, which are disjoint subsets of $\mathcal{H}(n)$. In order to satisfy Property $\mathbf{B}(w)$, we now make sure that

$$C_1, C_2, \ldots, C_M \subset J^+(n,w) .$$

As shown in Section IV-A, this guarantees that the total number of state transitions on the $n$ bus wires is at most $w$.

The first construction applies the well known Baranyai's theorem [18, p. 536]. The theorem known in the context of resolutions in block design or the decomposition of complete hypergraphs asserts that the set of $\binom{w(t+1)}{w}$ binary words of length $w(t+1)$ and weight $w$ can be partitioned into pairwise disjoint sets of size $t+1$, where two words in the same set do not intersect. It is easy to verify that these sets can be taken as the codesets of a $(w(t+1), t, w)$-LPC code.

To correct $e$ errors and detect $e+1$ errors, in this construction, we should use a constant weight code of length $w(t+1)$, weight $w$, and minimum Hamming distance $2e+2$ (instead of all the $\binom{w(t+1)}{w}$) words of length $w(t+1)$). Related partitions are well known for $e = 1$ and $w = 3, 4$, in the context of resolvable Steiner systems [18, p. 353].

## Acknowledgment

Y. M. Chee was supported in part by the Singapore Ministry of Education under Research Grant MOE2015-T2-2-086. T. Etzion and A. Vardy were supported in part by the United States — Israel Binational Science Foundation (BSF), Jerusalem, Israel, under Grant 2012016. H. M. Kiah was supported in part by the Singapore Ministry of Education under Research Grants MOE2015-T2-2-086 and MOE2016-T1-001-156.

## References

[1] A. H. Ajami, K. Banerjee, M. Pedram, and L.P.P.P. van Ginneken, "Analysis of non-uniform temperature-dependent interconnect performance in high performance ICs," in *DAC 2001: Proceedings of the 38-th Annual Design Automation Conference*. ACM Press, 2001, pp. 567–572.

[2] Y.M. Chee, C.J. Colbourn and A.C.H. Ling, "Optimal memoryless encoding for low power off-chip data buses," in *ICCAD 2006: Proceedings of the International Conference on Computer-Aided Design*. IEEE Press, 2006, pp. 369–374.

[3] Y.M. Chee, T. Etzion, H.M. Kiah, and A. Vardy, "Cooling codes: thermal-management coding for high-performance interconnects," *arxiv.org/abs/1701.07872*, January 2017.

[4] T.Y. Chiang, K. Banerjee, and K.C. Saraswat, "Compact modeling and SPICE-based simulation for electrothermal analysis of multilevel ULSI interconnects," in *ICCAD 2001: Proceedings of the International Conference on Computer-Aided Design*. IEEE Press, 2001, pp. 165–172.

[5] B. Datta and W. Burleson, "Collaborative sensing of on-chip wire temperatures using interconnect based ring oscillators," in *GLSVLSI 2008: Proceedings of the 18th Great Lakes Symposium on VLSI.* Orlando 2008, pp. 41–46.

[6] I. Dumer, "Asymptotically optimal codes correcting memory defects of fixed multiplicity," *Probl. Pered. Inform.*, vol. 25, no. 4, pp. 3-10, 1989.

[7] W. Elmore, "The transient response of damped linear network with particular regard to wide-band amplifier," *Journ. of Applied Physics*, vol. 10, pp. 55-63, 1948.

[8] T. Etzion and A. Vardy, "Error-correcting codes in projective space," *IEEE Trans. Inform. Theory*, vol. 57, no. 2, pp. 1165–1173, February 2011.

[9] D.J. Kleitman, "On a combinatorial conjecture of Erdős," *J. Combin. Theory*, vol. 1, pp. 209–214, September 1966.

[10] G. Katona, "Intersection theorems for systems of finite sets," *Acta Math. Acad. Sci. Hungar.*, vol. 15, pp. 329–337, 1964.

[11] E. Macii, M. Pedram, and F. Somenzi, "High-level power modeling, estimation, and optimization," *IEEE Trans. Computer-Aided Design Integr. Circuits and Syst.*, vol. 17, no. 11, pp. 1061–1079, November 1998.

[12] S. Ramprasad, N.R. Shanbhag, and I.N. Hajj, "Information-theoretic bounds on the average signal transition activity," *IEEE Trans. Very Large Scale Integr. Syst.*, vol. 7, no. 3, pp. 359–368, September 1999.

[13] P.P. Sotiriadis and A.P. Chandrakasan, "Reducing bus delay in submicron technology using coding," in *IEEE Asia and South Pacific Design Automation Conference,* 2001, pp. 109–114.

[14] P.P. Sotiriadis and A.P. Chandrakasan, "A bus energy model for deep submicron technology," *IEEE Trans. Very Large Scale Integr. Syst.*, vol. 10, no. 3, pp. 341–350, March 2002.

[15] P.P. Sotiriadis and A.P. Chandrakasan, "Bus energy reduction by transition pattern coding using a detailed deep submicrometer bus model," *IEEE Trans. on Circ. and Syst. – Fund. Theory and Appl.*, vol. 50, no. 10, pp. 1280–1295, October 2003.

[16] M.R. Stan and W.P. Burleson, "Bus-invert coding for low power I/O," *IEEE Trans. Very Large Scale Integr. Syst.*, vol. 3, no. 1, pp. 49–58, 1995.

[17] K. Sundaresan and N.R. Mahapatra, "Accurate energy dissipation and thermal modeling for nanometer-scale buses," in *HPCA 2005: Proceedings of the 11th International Symposium on High-Performance Computer Architecture*. IEEE Computer Society, 2005, pp. 51–60.

[18] J.H. van Lint and R.M. Wilson, *A course in Combinatorics*. Cambridge University Press, 1992.

[19] M. Vukolić, *Quorum Systems: With Applications to Storage and Consensus*. San Rafael, CA.: Morgan & Claypool Publishers, February 2012.

[20] F. Wang, M. De Bole, X. Wu, Y. Xie, N. Vijaykrishnan, and M.J. Irwin, "On-chip bus thermal analysis and optimization," *IET Computers & Digital Techniques*, vol. 1, no. 5., pp. 590–599, September 2007.