

Functional Clones and Expressibility of Partition Functions*

Andrei Bulatov[†] Leslie Ann Goldberg[‡] Mark Jerrum[§] David Richerby[‡]
Stanislav Živný[¶]

19th April, 2017

Abstract

We study functional clones, which are sets of non-negative pseudo-Boolean functions (functions $\{0, 1\}^k \rightarrow \mathbb{R}_{\geq 0}$) closed under (essentially) multiplication, summation and limits. Functional clones naturally form a lattice under set inclusion and are closely related to counting Constraint Satisfaction Problems (CSPs). We identify a sublattice of interesting functional clones and investigate the relationships and properties of the functional clones in this sublattice.

1 Introduction

There is a considerable literature on the topic of relational clones, also called co-clones. These are sets of relations on a finite domain D that are closed under certain operations, the most interesting being conjunction of two relations and existential quantification over a variable. (Other closure operations, such as introduction of “fictitious arguments”, are technically but not conceptually important.) In this paper we focus on the Boolean domain, and presently we will assume that $D = \{0, 1\}$. It is well known that in the Boolean case, the set of relational clones is countably infinite and forms a lattice under set inclusion. The lattice has been explicitly described by Post [18].

It seems natural to widen this study to other algebraic structures. Functional clones were introduced formally by Bulatov, Dyer, Goldberg, Jerrum and McQuillan [2], with the motivation of studying the computational complexity of counting constraint satisfaction problems. A functional clone is a set of multivariate functions from a finite domain D to a semiring R that is closed under multiplication, summing over a variable and (optionally) taking a limit of a sequence of functions. (Other operations are needed for technical completeness. Formal

*Part of this work was done while the authors were visiting the Simons Institute for the Theory of Computing.

[†]Simon Fraser University, Burnaby, Canada. Supported by an NSERC Discovery Grant.

[‡]University of Oxford, UK. The research leading to these results has received funding from the European Research Council under the European Union’s Seventh Framework Programme (FP7/2007–2013) ERC grant agreement no. 334828 (Goldberg and Richerby; this grant also supported visits by Bulatov and Jerrum) and Horizon 2020 research and innovation programme (grant agreement no. 714532, Živný). The paper reflects only the authors’ views and not the views of the ERC or the European Commission. The European Union is not liable for any use that may be made of the information contained therein.

[§]Queen Mary, University of London, UK Supported by EPSRC grant EP/N004221/1, an NSERC Discovery Grant and the SFU Distinguished Lecturers program.

[¶]Supported by a Royal Society University Research Fellowship.

definitions are given in the following subsection.) In this paper, we focus attention on the case $D = \{0, 1\}$ and $S = (\mathbb{R}_{\geq 0}, \times, +)$. We reconsider functional clones as objects of interest in their own right, though the results we prove may yield insights in other areas.

There are at least three motivations for the current investigation.

The first, as indicated above, is intrinsic interest. Post’s lattice of relational clones has a fascinating structure. There is a Galois connection between sets of relations on D and sets of operations on D which establishes a beautiful duality between relational clones and clones of operations. Remarkably, the closure operator defined by the Galois connection exactly agrees with the one described earlier in terms of conjunction of relations and existential quantification over variables [9].

The situation with functional clones is not quite so clean. There is apparently no Galois connection between sets of functions on D and sets of (somehow appropriately generalised) operations on D that captures the closure under multiplication and summation described above. Moreover, the lattice of functional clones has the cardinality of the continuum (or even larger, depending on precise definitions) and there seems to be no hope of providing a complete description of it. Still, it is interesting to map out some of the main features of the lattice, to identify maximal functional clones, to identify sublattices of functional clones satisfying additional properties, to find alternative characterisations of certain functional clones in terms of generating sets or Fourier coefficients, etc. As a contribution in this direction we identify (Figure 1) a sublattice of what seem to us to be interesting functional clones.

The second motivation, hinted at earlier, is the desire to understand the computational complexity of certain counting problems. A classical (decision) Constraint Satisfaction Problem (CSP) is a generalised satisfiability problem. Instead of restricting clauses to being disjunctions of literals, as in standard satisfiability problems, we allow arbitrary relations between variables chosen from a specified set or “language” of relations Γ . We are interested in how the computational complexity of a CSP varies as a function of Γ . Clearly, extending the language Γ may increase the complexity of the corresponding CSP. It transpires that the complexity of a CSP depends not on the fine structure of Γ , but only on the relational clone generated by Γ . This observation makes feasible the detailed exploration of the complexity of classical CSPs.

A counting Constraint Satisfaction Problem ($\#$ CSP) asks for the number of satisfying assignments to a CSP. In their weighted form, counting CSPs are general enough to express many partition functions occurring in statistical physics. Just as with classical decision CSPs, the complexity of a counting CSP is determined by the functional clone generated by the constraint language, which now consists of functions taking, say, non-negative real values. Functional clones were introduced in [2] precisely as a tool for studying the complexity of $\#$ CSPs. Referring to Figure 1, the equality at the bottom of the lattice expresses the equivalence between (on the left) the partition function of the ferromagnetic Ising model and (on the right) the so-called high-temperature expansion in terms of even subgraphs. Counting CSPs at this level of the lattice can be approximated in polynomial time by an algorithm that exploits this equivalence [15]. Moving up the lattice, perhaps the most intriguing functional clone from the complexity point of view is \mathcal{M} which includes the counting CSPs that would become feasible to approximate if we were to discover a polynomial time approximation algorithm for counting matchings in a general (non-bipartite) graph.

A third motivation for our study is provided by the connection between functional clones and topics in statistical physics and machine learning. Many models in statistical physics are “spin models” defined by a graph or more generally a hypergraph on n vertices. To each

vertex is associated a variable taking on values from a set of “spins” which, in our case, is finite. A configuration of the system is an assignment of spins to the n variables. The edges of the graph or hypergraph specify local interactions between spins. These local interactions define a probability distribution on the set of all configurations. Take for example the Ising model, which is characterised by having just two spins. An instance of the Ising model is specified by an undirected graph; in other words, there are just pairwise interactions between spins. (Refer to Section 2 for details.) One question we may ask is: which k -way interactions may be induced in such a model? More precisely, what are the possible marginal distributions that may be observed on some k -subset of the vertex variables? This question is (modulo the normalising factor for the probability distribution in question) precisely a question about functional clones. The possible marginal distributions are the k -ary functions in the clone generated by the local pairwise interactions.

In the case of the antiferromagnetic Ising model, where the pairwise interactions favour unlike spins, the answer is given by Theorem 48: the possible marginal distributions are precisely those that are “self-dual”, i.e., invariant under exchange of 0 and 1. (It is clear that invariance under exchange of 0 and 1 is necessary; the point is that it is sufficient.) This result has an implication for the expressive power of Boltzmann machines in machine learning [1]. Specifically, if the bias parameters of the units are all zero, then the distributions realisable at the visible units are precisely those that are self-dual. Note that this is an expressibility result, in the spirit of Le Roux and Bengio [16], and says nothing about the feasibility of learning the distributions in question from examples.

The analogous question in the ferromagnetic case is seemingly harder. The three-variable marginals of a ferromagnetic Ising model can be described: they are the (normalised) functions of arity 3 in the functional clone associated with the Ising model, and are given in Theorem 64. Already at arity 4 the elements of the clone become hard to describe. Indeed, it is consistent with our current knowledge that membership in this clone is undecidable, even for functions of some fixed arity greater than three.

Finally, there is a connection between functional clones and the idea of “universal models” in statistical physics proposed by De las Cuevas and Cubitt [7]. In a sense, functional clones formalise De las Cuevas and Cubitt’s notion of “closure”. A spin model is “universal” in their sense if (very roughly) the functional clone generated by the model is the one at the top of the clone lattice, namely \mathcal{B} , that contains all functions. They identify the planar antiferromagnetic Ising model with external fields as an example of a universal model.

As we already noted, the antiferromagnetic Ising model generates the clone \mathcal{SD} of self-dual functions. Adding an external field takes us outside of \mathcal{SD} . Now, according to Lemma 49, the clone \mathcal{SD} is “maximal”, from which we deduce that the antiferromagnetic Ising model with fields generates \mathcal{B} , i.e., is universal in our sense. Note, however, that our framework does not incorporate the notion of planarity, and in any case our closures do not exactly correspond to those of De las Cuevas and Cubitt. However, the clone lattice gives a more nuanced account of the expressive power of various spin models than simple universality. For more on the expressive power of spin systems and their computational complexity, see Goldberg and Jerrum [13] and Chen, Dyer, Goldberg, Jerrum, Lu, McQuillan and Richerby [5].

1.1 Functional Clones

For every non-negative integer k , let \mathcal{B}_k be the set of all arity- k non-negative pseudo-Boolean functions (i.e., the set of all functions $\{0, 1\}^k \rightarrow \mathbb{R}_{\geq 0}$). Let \mathcal{B} be the set of all non-negative

pseudo-Boolean functions (of all arities), given by $\mathcal{B} = \mathcal{B}_0 \cup \mathcal{B}_1 \cup \mathcal{B}_2 \cup \dots$. Given a function $f \in \mathcal{B}_k$ and a permutation π of $\{1, \dots, k\}$, we write f^π for the function that maps $(x_1, \dots, x_k) \in \{0, 1\}^k$ to $f(x_{\pi(1)}, \dots, x_{\pi(k)})$. Functional clones are subsets of \mathcal{B} that are closed under certain operations. We start by defining the operations. Consider a set $\mathcal{F} \subseteq \mathcal{B}$.

- \mathcal{F} is *closed under the introduction of fictitious arguments* if, for every $k \geq 0$ and every k -ary function $f \in \mathcal{F}$, the $(k+1)$ -ary function g defined by $g(x_1, \dots, x_{k+1}) = f(x_1, \dots, x_k)$ is also in \mathcal{F} .
- \mathcal{F} is *closed under permuting arguments* if, for every $k \geq 1$, every k -ary function $f \in \mathcal{F}$ and every permutation π of $\{1, \dots, k\}$, the function f^π is also in \mathcal{F} .
- \mathcal{F} is *closed under product* if, for every $k \geq 0$, every k -ary function $f \in \mathcal{F}$ and every k -ary function $g \in \mathcal{F}$, the function h defined by $h(x_1, \dots, x_k) = f(x_1, \dots, x_k)g(x_1, \dots, x_k)$ is also in \mathcal{F} .
- \mathcal{F} is *closed under summation* if, for every $k \geq 1$ and every k -ary function $f \in \mathcal{F}$, the $(k-1)$ -ary function g defined by $g(x_1, \dots, x_{k-1}) = \sum_{x_k \in \{0,1\}} f(x_1, \dots, x_k)$ is also in \mathcal{F} .

Functional clones are defined in [2, Section 2]. The definition that we give here is equivalent to the one in [2], but is more suited to the setting of this paper. Let EQ be the binary equality function, which is the function in \mathcal{B}_2 defined by $\text{EQ}(0,0) = \text{EQ}(1,1) = 1$, and $\text{EQ}(0,1) = \text{EQ}(1,0) = 0$. Suppose that $\mathcal{F} \subseteq \mathcal{B}$ is a set of functions. The *functional clone* $\langle \mathcal{F} \rangle$ is defined to be the closure of $\mathcal{F} \cup \{\text{EQ}\}$ under the introduction of fictitious arguments, permuting arguments, product and summation.

Bulatov et al. [2, Proof of Lemma 2.1] show¹ that the set $\langle \mathcal{F} \rangle$ is unchanged if the order of closure is restricted in the following way. Let $\mathcal{A}(\mathcal{F})$ be the closure of $\mathcal{F} \cup \{\text{EQ}\}$ under the introduction of fictitious arguments and permuting arguments. Let $\prod(\mathcal{F})$ be the closure of $\mathcal{A}(\mathcal{F})$ under product. Then $\langle \mathcal{F} \rangle$ is the closure of $\prod(\mathcal{F})$ under summation. In the paper, we will use the fact that the order of closure can be restricted in this way. In particular, the definition of $\mathcal{A}(\mathcal{F})$ will be used.

The reason for defining functional clones is that they are closely connected to counting Constraint Satisfaction Problems (CSPs). Every function in $\langle \mathcal{F} \rangle$ can be represented by a *pps-formula* (“primitive product summation formula”), which is a summation of a product of atomic formulas representing functions in $\mathcal{A}(\mathcal{F})$.² The pps-formula can be viewed as the input to a counting CSP whose output is the value of the function. For example, consider the function XOR $\in \mathcal{B}_2$ defined by $\text{XOR}(0,0) = \text{XOR}(1,1) = 0$ and $\text{XOR}(0,1) = \text{XOR}(1,0) = 1$.

¹ Technically, the proof of Lemma 2.1 of [2] just shows that the closure of $\mathcal{A}(\mathcal{F})$ under product and summation is the same as the closure of $\prod(\mathcal{F})$ under summation. That is, to produce the closure of $\mathcal{A}(\mathcal{F})$ under product and summation it suffices to first close $\mathcal{A}(\mathcal{F})$ under product and then close the resulting set under summation. However, it is easy to show that $\prod(\mathcal{F})$ is closed under the introduction of fictitious arguments and permuting arguments, and so is the closure of $\prod(\mathcal{F})$ under summation, so without loss of generality, the three closures can be done in order: first close $\mathcal{F} \cup \{\text{EQ}\}$ under the introduction of fictitious arguments and permuting arguments, then close under product, then close under summation.

²There is one difference between pps-formulas as defined here, and pps-formulas as defined in [2], but it is not important. Consider an arity- k function f . Clearly, the arity- $(k-1)$ function defined by $g(x_1, \dots, x_{k-1}) = f(x_1, \dots, x_{k-1}, x_{k-1})$ is in $\langle \{f\} \rangle$ since $g(x_1, \dots, x_{k-1}) = \sum_{x_k \in \{0,1\}} f(x_1, \dots, x_k) \text{EQ}(x_{k-1}, x_k)$. The function g is not in the set $\mathcal{A}(\mathcal{F})$. Nevertheless, Bulatov et al. [2] view the formula ϕ_g that represents g as an “atomic formula” since they allow repeated arguments. For us, the formula ϕ_g is not atomic, but this makes no difference, since EQ is in all functional clones, so our functional clones are exactly the same as those of [2].

Let h be the function in \mathcal{B}_3 defined by $h(1, 1, 0) = h(0, 0, 1) = 1$ and $h(x_1, x_2, x_3) = 0$ for any $(x_1, x_2, x_3) \notin \{(1, 1, 0), (0, 0, 1)\}$. Let \mathbf{x} denote the tuple (x_1, x_2, x_3, x_4) . It is easy to see that h is in $\langle \{\text{XOR}\} \rangle$ since the functions $f_{i,j}(\mathbf{x}) = \text{XOR}(x_i, x_j)$ are in $\mathcal{A}(\mathcal{F})$ for any distinct i and j in $\{1, 2, 3, 4\}$ and the function $g(\mathbf{x}) = f_{1,4}(\mathbf{x})f_{2,4}(\mathbf{x})f_{1,3}(\mathbf{x})$ is in $\prod(\mathcal{F})$. Finally, $h(x_1, x_2, x_3) = \sum_{x_4 \in \{0,1\}} g(\mathbf{x})$. Now, for distinct i and j in $\{1, 2, 3, 4\}$, let $\phi_{i,j}(v_1, v_2, v_3, v_4)$ be an atomic formula representing the function $f_{i,j}$. The function g can be represented by the formula

$$\phi_g(v_1, v_2, v_3, v_4) = \phi_{1,4}(v_1, v_2, v_3, v_4)\phi_{2,4}(v_1, v_2, v_3, v_4)\phi_{1,3}(v_1, v_2, v_3, v_4).$$

This formula can be viewed as a CSP with variables $\{v_1, v_2, v_3, v_4\}$ and three XOR constraints. Finally, the function h can be represented by the pps-formula $\phi_h(v_1, v_2, v_3) = \sum_{v_4} \phi_g(v_1, v_2, v_3, v_4)$.

In order to study approximate counting CSPs it is necessary to go beyond functional clones by also allowing closure under limits. Given functions f and f' in \mathcal{B}_k , we write $\|f - f'\|_\infty$ for the L -infinity distance between f and f' , which is given by $\|f - f'\|_\infty = \max_{\mathbf{x} \in \{0,1\}^k} |f(\mathbf{x}) - f'(\mathbf{x})|$. We say that a k -ary function f is a *limit* of a set $\mathcal{F} \subseteq \mathcal{B}$ if there is some finite $S_f \subseteq \mathcal{F}$ such that, for every $\varepsilon > 0$, there is a k -ary function $f_\varepsilon \in \langle S_f \rangle$ such that $\|f - f_\varepsilon\|_\infty < \varepsilon$. We say that \mathcal{F} is *closed under limits* if, for every function f that is a limit of \mathcal{F} , $f \in \mathcal{F}$. The ω -clone $\langle \mathcal{F} \rangle_\omega$ is defined to be the closure of $\mathcal{F} \cup \{\text{EQ}\}$ under the introduction of fictitious arguments, permuting arguments, product, summation, and limits. In [2], the set $\langle \mathcal{F} \rangle_\omega$ is referred to as the “pps- ω -definable functional clone generated by \mathcal{F} ”. Bulatov et al. [2, Lemma 2.2] show that this set is unchanged if the order of closure is restricted so $\langle \mathcal{F} \rangle_\omega$ is the closure of $\langle \mathcal{F} \rangle$ under limits.³

The following lemma is straightforward, given that $\langle \mathcal{F} \rangle$ and $\langle \mathcal{F} \rangle_\omega$ are defined by closing a set (the set $\mathcal{F} \cup \{\text{EQ}\}$) using various operations. Nevertheless, we state the lemma here for future use. The lemma combines Lemmas 2.1 and 2.2 of [2]. (In that paper, the lemma was non-trivial, since the order of the closure operators was restricted.)

Lemma 1. *Suppose $\mathcal{F} \subseteq \mathcal{B}$. If $g \in \langle \mathcal{F} \rangle$ then $\langle \mathcal{F} \cup \{g\} \rangle = \langle \mathcal{F} \rangle$. If g is a limit of $\langle \mathcal{F} \rangle$ and h is a limit of $\langle \mathcal{F} \cup \{g\} \rangle$ then h is a limit of $\langle \mathcal{F} \rangle$. Equivalently, if $g \in \langle \mathcal{F} \rangle_\omega$ then $\langle \mathcal{F} \cup \{g\} \rangle_\omega = \langle \mathcal{F} \rangle_\omega$.*

1.2 Lattices

A *lattice* is a set L equipped with two commutative, associative binary operations \vee (*join*) and \wedge (*meet*) with the absorption property: $a \vee (a \wedge b) = a$ and $a \wedge (a \vee b) = a$ for all $a, b \in L$. The lattice operations \vee and \wedge induce a partial order on L as follows: for $a, b \in L$, $a \leq b$ if and only if $b = a \vee b$ (or, equivalently, $a = a \wedge b$). It is easy to see that, for any $a, b \in L$, the elements $a \vee b$ and $a \wedge b$ are the least upper bound and greatest lower bound of a and b , with respect to the order \leq . In other words, for any c such that $a \leq c$ and $b \leq c$ it holds that $a \vee b \leq c$, and for any d such that $d \leq a$ and $d \leq b$ it holds that $d \leq a \wedge b$. Conversely, if a set L has a partial order \leq such that any pair of elements has a least upper bound and a greatest lower bound, then it can be converted into a lattice by defining the operations of join and meet as the least upper bound and the greatest lower bound respectively. A subset $L' \subseteq L$ is called a *sublattice* if for all $a, b \in L'$, $a \vee b$ and $a \wedge b$ belong to L' . Note that \vee and \wedge here are the operations of L .

³Technically, the proof of Lemma 2.2 of [2] just shows that the closure of $\mathcal{A}(\mathcal{F})$ under product, summation and limits is the same as the closure of $\langle \mathcal{F} \rangle$ under limits. However, it is easy to see that the closure of $\langle \mathcal{F} \rangle$ under limits is closed under the introduction of fictitious arguments and permuting arguments.

1.3 Lattices of functional clones

Let \mathcal{L}_f and \mathcal{L}_ω denote the set of all functional clones and all ω -clones, respectively, ordered with respect to set inclusion. Then, for any two functional clones (or ω -clones) \mathcal{F} and \mathcal{G} , the least upper bound and the greatest lower bound are given by $\langle \mathcal{F} \cup \mathcal{G} \rangle$ (resp., $\langle \mathcal{F} \cup \mathcal{G} \rangle_\omega$) and $\mathcal{F} \cap \mathcal{G}$ (in both cases). Therefore \mathcal{L}_f and \mathcal{L}_ω can be viewed as lattices with operations of join and meet

$$\begin{aligned} \mathcal{F} \vee_f \mathcal{G} &= \langle \mathcal{F} \cup \mathcal{G} \rangle, & \mathcal{F} \wedge_f \mathcal{G} &= \mathcal{F} \cap \mathcal{G} & \text{for } \mathcal{L}_f, \\ \mathcal{F} \vee_\omega \mathcal{G} &= \langle \mathcal{F} \cup \mathcal{G} \rangle_\omega, & \mathcal{F} \wedge_\omega \mathcal{G} &= \mathcal{F} \cap \mathcal{G} & \text{for } \mathcal{L}_\omega. \end{aligned}$$

Since we are mostly concerned with ω -clones, we will omit the subscripts of \vee_ω and \wedge_ω .

As we will show in Theorem 13, the lattices \mathcal{L}_f and \mathcal{L}_ω are quite large, having cardinality $\beth_2 = 2^{2^{\aleph_0}}$. Therefore we will focus on the most interesting and important ω -clones.

Definition 2. An ω -clone \mathcal{F} is *maximal* in an ω -clone \mathcal{G} if $\mathcal{F} \subseteq \mathcal{G}$ and there is no ω -clone \mathcal{C} such that $\mathcal{F} \subset \mathcal{C} \subset \mathcal{G}$.

It is easily seen that \mathcal{F} is maximal in \mathcal{G} if and only if, for any function $g \in \mathcal{G} \setminus \mathcal{F}$, $\langle \mathcal{F} \cup \{g\} \rangle_\omega = \mathcal{G}$.

2 Notation and the clones that we study

We denote tuples in $\{0, 1\}^k$ by boldface letters. We use the notation $|\mathbf{x}|$ to denote the Hamming weight of \mathbf{x} . The symbols $\mathbf{0}$ and $\mathbf{1}$ are used to denote the all-zeroes and all-ones tuple of arity appropriate to the context. $\bar{\mathbf{x}}$ is the bitwise complement of \mathbf{x} . We define $[k] = \{1, \dots, k\}$.

Recall the function f^π from Section 1.1. We say that an arity- k function f is *symmetric* if, for all permutations π of $[k]$, $f = f^\pi$. We often write symmetric k -ary functions as $f = [f_0, \dots, f_k]$, where f_i is the value of f on arguments of Hamming weight i . Using this notation, the function EQ can be written as $\text{EQ} = [1, 0, 1]$. We make use of the following unary functions: $\delta_0 = [1, 0]$ and $\delta_1 = [0, 1]$.

Definition 3. The *Fourier transform* of a function $f: \{0, 1\}^k \rightarrow \mathbb{R}_{\geq 0}$ is the function $\hat{f}: \{0, 1\}^k \rightarrow \mathbb{R}$ defined by

$$\hat{f}(\mathbf{x}) = \frac{1}{2^k} \sum_{\mathbf{w} \in \{0, 1\}^k} (-1)^{|\mathbf{w} \wedge \mathbf{x}|} f(\mathbf{w}).$$

Note that, although we only consider functions whose range is the nonnegative reals, the Fourier transform of such a function may have negative numbers in its range. Readers who are familiar with Valiant's holant framework [20] will recognise that, if we represent k -ary functions as column vectors of length 2^k , the Fourier transform is equivalently defined as $\hat{f} = H^{\otimes k} f$ where $H = \frac{1}{2} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix} = \frac{1}{2} H^{-1}$. We will use this fact in the proof of Theorem 54.

Definition 4. For a real number $\lambda \geq 0$ and integer $k \geq 0$, the *k -ary hypergraph Ising function* is given by

$$I_k^\lambda(\mathbf{x}) = \begin{cases} 1 & \text{if } \mathbf{x} \in \{\mathbf{0}, \mathbf{1}\} \\ \lambda & \text{otherwise.} \end{cases}$$

The case $\lambda \leq 1$ is known as *ferromagnetic* and $\lambda \geq 1$ is *antiferromagnetic*.

Definition 5. An arity- k *match-circuit* is given by an undirected weighted graph G with vertex set $\{u_1, \dots, u_k\} \cup \{v_1, \dots, v_n\}$ for some $n \geq k$. Vertices u_1, \dots, u_k have degree 1 and are called “external vertices”. The edges adjacent to them (called “terminals”) are labelled y_1, \dots, y_k . Vertices v_1, \dots, v_n are called “internal vertices”. Each terminal edge has weight 1 and each non-terminal edge e is equipped with a positive weight w_e . Configurations assign spins 0 and 1 to edges. A configuration is a *perfect matching* if every internal vertex has exactly one spin-1 edge adjacent to it. The match-circuit implements the function f , where $f(y_1, \dots, y_k)$ is the sum, over perfect matchings, of the product of the weights of edges with spin 1, where the empty product has weight 1.

Note that, if f is implemented by a match-circuit then so are all functions $c \cdot f$ where c is a positive real number: just add an isolated edge of weight c to the match-circuit implementing f . Also, some authors require the underlying graphs of match-circuits to be planar, and some authors allow the edge weights to be negative.

Definition 6. An arity- k *even-circuit* is given by an undirected weighted graph G with vertex set $\{u_1, \dots, u_k\} \cup \{v_1, \dots, v_n\}$ for some $n \geq k$. Vertices u_1, \dots, u_k have degree 1 and are called “external vertices”. The edges adjacent to them (called “terminals”) are labelled y_1, \dots, y_k . Vertices v_1, \dots, v_n are called “internal vertices”. Each terminal edge has weight 1 and each non-terminal edge e is equipped with a weight $w_e \in (0, 1]$. Configurations assign spins 0 and 1 to edges. A configuration is an even subgraph if every internal vertex has an even number of spin-1 edge adjacent to it. The even-circuit implements the function f , where $f(y_1, \dots, y_k)$ is the sum, over even subgraphs, of the product of the weights of the edge with spin 1, where the empty product has weight 1.

Note that, for even-circuits, we require all weights to be in $(0, 1]$ whereas, for match-circuits, we only require that weights be positive. In fact, match-circuits implement the same class of functions when restricted to weights in $(0, 1]$ as they do with arbitrary positive weights, but we use the less restricted definition for convenience.

For convenience when discussing match-circuits and even-circuits, we associate an assignment σ of spins to the edges of a graph G with the spanning subgraph $H = (V(G), \{e \in E(G) \mid \sigma(e) = 1\})$.

Definition 7. Given a weighted graph H , we write $w(H) = \prod_{e \in E(H)} w_e$ for the weight of H .

Definition 8. We define the following subsets of \mathcal{B} .

- SD : all self-dual functions f , i.e., functions such that $f(\mathbf{x}) = f(\bar{\mathbf{x}})$ for all \mathbf{x} .
- \mathcal{P} : all functions f such that $\hat{f}(\mathbf{x}) \geq 0$ for all \mathbf{x} .
- \mathcal{PN} : all functions f such that $\hat{f}(\mathbf{x}) \geq 0$ when $|\mathbf{x}|$ is even and $\hat{f}(\mathbf{x}) \leq 0$ when $|\mathbf{x}|$ is odd.
- $SDP = SD \cap \mathcal{P} \cap \mathcal{PN}$.
- \mathcal{E} : all functions $c \cdot f$, where c is a non-negative real number and \hat{f} is implemented by an even-circuit.
- \mathcal{M} : all functions f such that \hat{f} is implemented by a match-circuit.

- $\text{Ferrolsing} = \{I_2^\lambda \mid 0 \leq \lambda \leq 1\}$. The functions in the set Ferrolsing model edge interactions in the ferromagnetic Ising model. See Cipra [6] for an introduction to the Ising model.
- $\text{AntiFerrolsing} = \{I_2^\lambda \mid \lambda \geq 1\}$. The functions in the set AntiFerrolsing model edge interactions in the anti-ferromagnetic Ising model.
- $\text{FerroHyperlsing} = \{I_k^\lambda \mid k \geq 2, \lambda \leq 1\}$. The functions in the set FerroHyperlsing model “many-body interactions” in a generalisation of the ferromagnetic Ising model which applies to hypergraphs — see [14] and [11, Section 2].

We emphasise that the sets we have defined are subsets of \mathcal{B} , the class of non-negative pseudo-Boolean functions. There are, for example, functions outside \mathcal{B} whose Fourier transforms are in \mathcal{B} , such as the symmetric, ternary function $f = [7, -1, -1, 7]$, which has Fourier transform $\hat{f} = [1, 0, 2, 0]$. Even though \hat{f} is nonnegative, f is not in \mathcal{P} because it is not in \mathcal{B} . Likewise, $f \notin \mathcal{M}$, even though \hat{f} is implemented by a match-circuit (as shown in the proof of Theorem 64).

Instead of \mathcal{M} , it may seem more natural to consider the set \mathcal{M}' of functions f that are implemented by match-circuits. However, \mathcal{M}' is not a functional clone: for example, it is not closed under the introduction of fictitious arguments. By a parity argument, any function f that is implemented by a match-circuit must have $f(\mathbf{x}) = 0$ for all \mathbf{x} with even Hamming weight, or $f(\mathbf{x}) = 0$ for all \mathbf{x} with odd Hamming weight. However, any function that is not everywhere zero and has a fictitious argument must be non-zero for inputs with both odd and even Hamming weights, so cannot be implemented by a match-circuit.

As we have remarked, the Fourier transform corresponds in the holant framework to a holographic transformation by (the appropriate tensor power of) the Hadamard matrix $\frac{1}{2} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}$. This corresponds, in a certain sense, to transforming the computation from using basis vectors $\begin{pmatrix} 1 \\ 0 \end{pmatrix}$ and $\begin{pmatrix} 0 \\ 1 \end{pmatrix}$ to using $\begin{pmatrix} 1 \\ 1 \end{pmatrix}$ and $\begin{pmatrix} 1 \\ -1 \end{pmatrix}$. It has been shown that the latter is the unique basis in which the equality function can be expressed using matchgates [4] and, thus, our use of the Fourier transform here is essential. Cai, Lu and Xia [4] have used the Fourier transform as a holographic transformation from counting CSPs to counting weighted perfect matchings, as the key tool to obtain polynomial-time algorithms for a wide range of weighted planar counting CSPs.

Note that, in the definition of \mathcal{E} , we allow scaling by a constant. We do this to allow the implementation of functions that have $f(\mathbf{0}) < 1$. This would be impossible without scaling, since the empty graph is an even subgraph of every even-circuit. It has weight 1 and the weight of the empty graph is one of the terms of the sum defining $f(\mathbf{0})$. In contrast, match-circuits can already implement functions with $f(\mathbf{0}) < 1$ without the need for scaling, and adding scaling to the definition of \mathcal{M} would not, in fact, change the class of implementable functions.

To avoid issues with scaling of Ising and hypergraph Ising functions, we work with the following clones rather than with $\langle \text{Ferrolsing} \rangle_\omega$, etc.

Definition 9.

$$\begin{aligned} \mathcal{I}_{\text{ferro}} &= \langle \text{Ferrolsing} \cup \mathcal{B}_0 \rangle_\omega \\ \mathcal{I}_{\text{anti}} &= \langle \text{AntiFerrolsing} \cup \mathcal{B}_0 \rangle_\omega \\ \mathcal{H}_{\text{ferro}} &= \langle \text{FerroHyperlsing} \cup \mathcal{B}_0 \rangle_\omega. \end{aligned}$$




Figure 1: The lattice \mathcal{L}' .

3 Main theorems

Let $\mathcal{L}' = \{\mathcal{B}, \mathcal{SD}, \mathcal{P}, \mathcal{PN}, \mathcal{SDP}, \langle\langle\mathcal{M}\rangle_\omega \cup \mathcal{H}_{\text{ferro}}\rangle_\omega, \langle\mathcal{M}\rangle_\omega, \mathcal{H}_{\text{ferro}}, \langle\mathcal{M}\rangle_\omega \cap \mathcal{H}_{\text{ferro}}, \mathcal{I}_{\text{ferro}}\}$.

Theorem 10. *The lattice \mathcal{L}' shown in Figure 1 is a sublattice of \mathcal{L}_ω . That is, all elements of \mathcal{L}' are distinct ω -clones, with the possible exceptions of \mathcal{SDP} and $\langle\langle\mathcal{M}\rangle_\omega \cup \mathcal{H}_{\text{ferro}}\rangle_\omega$, and $\langle\mathcal{M}\rangle_\omega \cap \mathcal{H}_{\text{ferro}}$ and $\mathcal{I}_{\text{ferro}}$, which might be equal. (This is indicated by the dotted lines in Figure 1.) Furthermore, the meets and joins of elements of \mathcal{L}' are as depicted in Figure 1 and*

- (i) $\mathcal{SD} = \mathcal{I}_{\text{anti}}$;
- (ii) $\mathcal{I}_{\text{ferro}} = \langle\mathcal{E}\rangle_\omega$;
- (iii) \mathcal{SD} , \mathcal{P} and \mathcal{PN} are maximal in \mathcal{B} ;
- (iv) \mathcal{SDP} is maximal in \mathcal{SD} .

Theorem 10 is proven in Section 9.

Theorem 11. *For any $\lambda > 1$, $\langle I_2^\lambda \cup \mathcal{B}_0 \rangle_\omega = \mathcal{I}_{\text{anti}}$. For any $\lambda \in (0, 1)$, $\langle I_2^\lambda \cup \mathcal{B}_0 \rangle_\omega = \mathcal{I}_{\text{ferro}}$.*

Proof. The two parts are Corollaries 22 and 24, respectively, from Section 5. \square

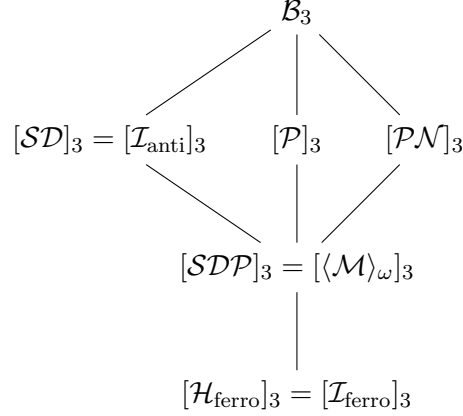


Figure 2: Ternary parts of the clones in \mathcal{L}' .

Theorem 12. \mathcal{E} and \mathcal{M} are functional clones, i.e., $\langle \mathcal{E} \rangle = \mathcal{E}$ and $\langle \mathcal{M} \rangle = \mathcal{M}$.

Proof. $\langle \mathcal{E} \rangle = \mathcal{E}$ is Theorem 53 and $\langle \mathcal{M} \rangle = \mathcal{M}$ is Theorem 51. □

Theorem 13. $|\mathcal{L}_f| = |\mathcal{L}_\omega| = \beth_2$.

Theorem 13 is proven in Section 11.

Theorem 62, proved in Section 10, shows that the set of monotone functions is an ω -clone and gives examples of ω -clones that generalise this clone.

3.1 Ternary functions

Given $n \geq 0$ and a set of functions $\mathcal{F} \subseteq \mathcal{B}$, we write $[\mathcal{F}]_n = \mathcal{F} \cap \mathcal{B}_n$. Note that $[\mathcal{B}]_n = \mathcal{B}_n$. Although $[\mathcal{F}]_n$ is a set of n -ary functions, it essentially includes all functions of smaller arity. In particular, if \mathcal{F} is a clone then it is closed under the introduction of fictitious arguments, as discussed in Section 1.1, and this allows functions of smaller arity to be “padded” to arity n . However, $[\mathcal{F}]_n$ is not, itself, a clone.

We now focus on the ternary parts of the clones from \mathcal{L}' , in which case certain distinctions, which were present in \mathcal{L}' , disappear.

Let $\mathcal{S}_3 = \{\mathcal{B}_3, [\mathcal{SD}]_3, [\mathcal{P}]_3, [\mathcal{PN}]_3, [\mathcal{SDP}]_3, [\langle \mathcal{M} \rangle_\omega]_3, [\mathcal{H}_{\text{ferro}}]_3, [\mathcal{I}_{\text{ferro}}]_3\}$.

Theorem 14. $[\mathcal{SDP}]_3 = [\langle \mathcal{M} \rangle_\omega]_3$, $[\mathcal{H}_{\text{ferro}}]_3 = [\mathcal{I}_{\text{ferro}}]_3$, and any other two elements of \mathcal{S}_3 are distinct.

Theorem 14 is proved in Section 12 and illustrated in Figure 2, where solid lines indicate *strict* set inclusions. The (non-strict) inclusions indicated in Figure 2 follow trivially from Theorem 10. The point of Theorem 14, in addition to the two collapses, is that all inclusions are strict. We note that, however, unlike in Figure 1, Figure 2 does not indicate any lattice order of \mathcal{S}_3 with respect to \wedge and \vee .

4 Finite generation

When we defined ω -clones in Section 1.1, we defined the limit of a set $\mathcal{F} \subseteq \mathcal{B}$ to be a function f which is approximated by a sequence of functions f_ε that are all in the functional clone of some finite subset S_f of \mathcal{F} . The finiteness restriction was present in the definitions of [2] and it is retained in this paper because it strengthens our results. Nevertheless, it causes slight technical problems, and to avoid these problems, we start the paper by defining a finite subset \mathcal{B}'_0 of \mathcal{B}_0 and showing that $\mathcal{B}_0 \subseteq \langle \mathcal{B}'_0 \rangle_\omega$. In the following definition, “ e ” is the base of the natural logarithm. The actual definition of \mathcal{B}'_0 is not very constrained, in the sense that we could have made other choices, but it is important to include an irrational number, and to include a number that is smaller than 1 and one that is larger than 1. We use a set of size four to simplify the argument.

Definition 15. $\mathcal{B}'_0 = \{1/e, 1/2, 2, e\}$

Lemma 16. $\mathcal{B}_0 \subseteq \langle \mathcal{B}'_0 \rangle_\omega$.

Proof. We will show that every nullary function in \mathcal{B}_0 is a limit of the closure of \mathcal{B}'_0 under product. Let $\alpha = \ln 2$. For any integers a and b , the quantity $e^{a+b\alpha}$ (viewed as a nullary function) is in $\langle \mathcal{B}'_0 \rangle$. So it suffices to show that, for every real number z (where e^z is viewed as a nullary function in \mathcal{B}_0) and any $\varepsilon > 0$, there are integers a and b such that $|e^{a+b\alpha} - e^z| < \varepsilon$. Given the universal quantification on ε , we can work instead with additive approximation — it suffices to show that for every real number z and every $\delta > 0$, there are integers a and b such that $|a + b\alpha - z| < \delta$. (To see this, suppose that we are given some z and ε . Let $\varepsilon' = \min(\varepsilon, 2e^z)$ and let $\delta = \varepsilon' e^{-z}/2$. Then since $\delta \leq 1$, we have $e^\delta - 1 \leq 2\delta$ so $e^{z+\delta} - e^z = e^z(e^\delta - 1) \leq 2\delta e^z = \varepsilon' \leq \varepsilon$. Similarly, $e^z - e^{z-\delta} = e^z(1 - e^{-\delta}) \leq 2\delta e^z \leq \varepsilon$.)

Now consider a real number $\delta > 0$. By Dirichlet’s approximation theorem, there are integers p and q such that $1 \leq q$ and $|p - q\alpha| < \delta$. Since α is positive and $q \geq 1$, it is clear that p is also positive if $\delta < \alpha$. Also, since α is irrational, $p - q\alpha$ is non-zero.

Consider any real number z . Let n be the integer such that $n \times |p - q\alpha| \leq z < (n + 1) \times |p - q\alpha|$. Then $|z - n \times |p - q\alpha|| < |p - q\alpha| < \delta$. If $p > q\alpha$ then $a = pn$ and $b = -qn$ suffices. Otherwise, $a = -pn$ and $b = qn$ suffices. \square

The proof of Lemma 16 is useful for one more technical finite generation result, so we state that here. For this, we need to define a class of parity functions.

Definition 17. For each $k \in \mathbb{N}$ and $\lambda \in \mathbb{R}_{\geq 0}$, we define the k -ary function

$$\text{Par}_k^\lambda(\mathbf{x}) = \begin{cases} 1 & \text{if } |\mathbf{x}| \text{ is even} \\ \lambda & \text{otherwise.} \end{cases}$$

By analogy to \mathcal{B}_0 , we also define a finite version.

Definition 18. $\text{Par}'_k = \{\text{Par}_k^{1/e}, \text{Par}_k^{1/2}, \text{Par}_k^2, \text{Par}_k^e\}$.

Lemma 19. For any even positive integer k and any $\lambda \in \mathbb{R}_{\geq 0}$, $\text{Par}_k^\lambda \in \langle \text{Par}'_k \rangle_\omega$.

Proof. Consider the k -ary function consisting of the product of a copies of Par_k^e and b copies of Par_k^2 . If the input has even parity, then the output is 1. Otherwise, the output is e^{a+ba} . Combinations of other functions in Par'_k are similar. So the proof is essentially the same as the proof of Lemma 16. \square

5 The Ising model

Recall the definition of I_2^λ from Definition 4, the definition of Ferrolsing and AntiFerrolsing from Definition 8 and the definition of $\mathcal{I}_{\text{anti}}$ from Definition 9. The following lemma is well known. We include it (with its standard proof) for completeness.

Lemma 20. $\text{Ferrolsing} \subseteq \langle \text{AntiFerrolsing}, \mathcal{B}_0 \rangle$.

Proof. We must show that $I_2^\lambda \in \langle \text{AntiFerrolsing}, \mathcal{B}_0 \rangle$ for all $\lambda \in [0, 1]$. For $\lambda = 0$, $I_2^\lambda = \text{EQ}$, which is in every functional clone by definition. For $\lambda = 1$, $I_2^\lambda \in \text{AntiFerrolsing}$ by definition. Any other function in Ferrolsing is of the form I_2^λ for some $\lambda \in (0, 1)$. Let $\lambda' = \lambda/(1 - \sqrt{1 - \lambda^2})$. Note that λ' is decreasing as λ increases, and that $\lambda' > 1$ so $I_2^{\lambda'} \in \text{AntiFerrolsing}$. Then note that $I_2^\lambda(x, y) = \frac{1}{1 + \lambda'^2} \sum_w I_2^{\lambda'}(x, w) I_2^{\lambda'}(w, y)$ since the weight is 1 if $x = y$ and $(2\lambda')/(1 + \lambda'^2) = \lambda$, otherwise. \square

The construction in the proof of the following lemma is based on one from the proof of [10, Lemma 3.3]. There are more efficient constructions, for example [12, Lemma 3.26] but we don't need them here.

Lemma 21. Consider I_2^λ and $I_2^{\lambda'}$ in AntiFerrolsing with $\lambda > 1$. Then $I_2^{\lambda'} \in \langle \{I_2^\lambda\} \cup \mathcal{B}'_0 \rangle_\omega$.

Proof. By the definition of AntiFerrolsing, $\lambda' \geq 1$. If $\lambda' = 1$ then $I_2^{\lambda'}$ is the arity-2 constant function (with output 1). This can be obtained from the constant 1 by introducing two fictitious arguments, so it is in $\langle \{I_2^\lambda\} \cup \mathcal{B}'_0 \rangle_\omega$.

So suppose $\lambda' > 1$. Let $y = 1/\lambda$ and let f be the symmetric arity-2 function $[y^{1/2}, y^{-1/2}, y^{1/2}]$, using the symmetric function notation from Section 2. For every positive integer t , let $F_{1,t}(x_1, x_2) = f(x_1, x_2)^t$. For every integer $\ell > 1$, let X_ℓ be the tuple of variables in $\{x_{i,j} \mid 1 \leq i \leq t, 1 \leq j \leq \ell - 1\}$ and let

$$F_{\ell,t}(x_1, x_2) = \sum_{X_\ell} \prod_{i=1}^t \left(f(x_1, x_{i,1}) \left(\prod_{j=1}^{\ell-2} f(x_{i,j}, x_{i,j+1}) \right) f(x_{i,\ell-1}, x_2) \right).$$

Note that the quantity $y^{1/2}$ can be viewed as a nullary function, so by Lemma 16, $y^{1/2}$ is a limit of $\langle \mathcal{B}'_0 \rangle$. Since $f = y^{1/2} I_2^\lambda$, Lemma 1 shows that f is a limit of $\langle \{I_2^\lambda\} \cup \mathcal{B}'_0 \rangle$. Finally, since $F_{\ell,t}$ is formed by summing products of functions in $\mathcal{A}(\{f\})$, Lemma 1 shows that $F_{\ell,t}$ is also a limit of $\langle \{I_2^\lambda\} \cup \mathcal{B}'_0 \rangle$.

We wish to show that $I_2^{\lambda'}$ is a limit of $\langle \{I_2^\lambda\} \cup \mathcal{B}'_0 \rangle$. To do this, we will show that, for every $0 < \varepsilon < 1$, there are positive integers t and ℓ and a non-negative constant c (viewed as a limit of $\langle \mathcal{B}'_0 \rangle$) such that

$$\max_{(x_1, x_2) \in \{0,1\}^2} |I_2^{\lambda'}(x_1, x_2) - c F_{\ell,t}(x_1, x_2)| < \varepsilon.$$

To see this, consider the following mutual recurrences.

$$m_\ell = \begin{cases} y^{1/2}, & \text{if } \ell = 1, \\ y^{1/2} m_{\ell-1} + y^{-1/2} b_{\ell-1}, & \text{if } \ell > 1. \end{cases}$$

$$b_\ell = \begin{cases} y^{-1/2}, & \text{if } \ell = 1, \\ y^{-1/2} m_{\ell-1} + y^{1/2} b_{\ell-1}, & \text{if } \ell > 1. \end{cases}$$

First, consider $t = 1$. Renaming the variables $\{x_{1,1}, \dots, x_{1,\ell-1}\}$ to $\{x_3, \dots, x_{\ell+1}\}$, the definition of $F_{\ell,t}$ (for $\ell > 1$) can be written as

$$F_{\ell,1}(x_1, x_2) = \sum_{(x_3, \dots, x_{\ell+1})} f(x_1, x_3) \left(\prod_{j=3}^{\ell} f(x_j, x_{j+1}) \right) f(x_{\ell+1}, x_2).$$

From the recurrences, it is easy to see that $F_{\ell,1}(0,0) = F_{\ell,1}(1,1) = m_\ell$ (“ m ” stands for “monochromatic”) and $F_{\ell,1}(0,1) = F_{\ell,1}(1,0) = b_\ell$ (“ b ” stands for “bichromatic”). Thus, for general t , $F_{\ell,t}(0,0) = F_{\ell,t}(1,1) = m_\ell^t$ and $F_{\ell,t}(0,1) = F_{\ell,t}(1,0) = b_\ell^t$.

Now the solution to the recurrences is

$$\begin{aligned} m_\ell &= y^{-\ell/2}((y+1)^\ell + (y-1)^\ell)/2 \\ b_\ell &= y^{-\ell/2}((y+1)^\ell - (y-1)^\ell)/2. \end{aligned}$$

Thus, since $0 < y < 1$, for odd ℓ we have

$$\frac{b_\ell}{m_\ell} = 1 + \frac{2}{\left(\frac{1+y}{1-y}\right)^\ell - 1}.$$

So finally, given $0 < \varepsilon < 1$, let ℓ be the smallest odd integer so that

$$\left(\frac{1+y}{1-y}\right)^\ell > 1 + \frac{2\lambda'}{\varepsilon}.$$

Let t be the smallest integer so that

$$\left(1 + \frac{2}{\left(\frac{1+y}{1-y}\right)^\ell - 1}\right)^t > \lambda'.$$

Let $c = m_\ell^{-t}$. Then $cF_{\ell,t}(0,0) = cF_{\ell,t}(1,1) = 1$. Also $cF_{\ell,t}(0,1) = cF_{\ell,t}(1,0) = (b_\ell/m_\ell)^t$ so

$$\begin{aligned} \lambda' &< cF_{\ell,t}(0,1) = cF_{\ell,t}(1,0) = \left(1 + \frac{2}{\left(\frac{1+y}{1-y}\right)^\ell - 1}\right)^t \\ &= \left(1 + \frac{2}{\left(\frac{1+y}{1-y}\right)^\ell - 1}\right)^{t-1} \left(1 + \frac{2}{\left(\frac{1+y}{1-y}\right)^\ell - 1}\right) \\ &\leq \lambda' \left(1 + \frac{\varepsilon}{\lambda'}\right) < \lambda' + \varepsilon, \end{aligned}$$

as required. \square

Corollary 22. For any $\lambda > 1$, $\langle I_2^\lambda \cup \mathcal{B}'_0 \rangle_\omega = \mathcal{I}_{\text{anti}}$.

Proof. Recall from definition 9 that $\mathcal{I}_{\text{anti}} = \langle \text{AntiFerrolsing} \cup \mathcal{B}_0 \rangle_\omega$ and from Definition 8 that for any $\lambda > 1$, $I_2^\lambda \in \text{AntiFerrolsing}$. This shows $\langle I_2^\lambda \cup \mathcal{B}'_0 \rangle_\omega \subseteq \mathcal{I}_{\text{anti}}$. To see that $\mathcal{I}_{\text{anti}} \subseteq \langle I_2^\lambda \cup \mathcal{B}'_0 \rangle_\omega$ we only need to show that for any $I_2^\lambda \in \mathcal{I}_{\text{anti}}$, $I_2^\lambda \in \langle I_2^\lambda \cup \mathcal{B}'_0 \rangle_\omega$, and this is Lemma 21. \square

Lemma 23. Consider I_2^λ and $I_2^{\lambda'}$ in Ferrolsing with $0 < \lambda < 1$. Then $I_2^{\lambda'} \in \langle \{I_2^\lambda\} \cup \mathcal{B}'_0 \rangle_\omega$.

Proof. As in the proof of Lemma 21, the proof is straightforward if $\lambda' \in \{0, 1\}$, so assume $0 < \lambda' < 1$. Define y , f and $F_{\ell,t}$ as in the proof of Lemma 21. Note that $y > 1$, so

$$\frac{m_\ell}{b_\ell} = 1 + \frac{2}{\left(\frac{y+1}{y-1}\right)^\ell - 1}.$$

Given $0 < \varepsilon < 1$, let ℓ be the smallest positive integer so that

$$\left(\frac{y+1}{y-1}\right)^\ell > 1 + \frac{2}{\varepsilon}.$$

Let t be the largest integer so that

$$\left(1 + \frac{2}{\left(\frac{y+1}{y-1}\right)^\ell - 1}\right)^{t-1} \leq \frac{1}{\lambda'}.$$

Let $c = \lambda' b_\ell^{-t}$. Then $cF_{\ell,t}(0, 1) = cF_{\ell,t}(1, 0) = \lambda'$. Also $cF_{\ell,t}(0, 0) = cF_{\ell,t}(1, 1) = \lambda' m_\ell^t / b_\ell^t$ and $m_\ell^t / b_\ell^t > 1/\lambda'$, so

$$\begin{aligned} 1 < cF_{\ell,t}(0, 0) &= cF_{\ell,t}(1, 1) = \lambda' \left(1 + \frac{2}{\left(\frac{y+1}{y-1}\right)^\ell - 1}\right)^t \\ &= \lambda' \left(1 + \frac{2}{\left(\frac{y+1}{y-1}\right)^\ell - 1}\right)^{t-1} \left(1 + \frac{2}{\left(\frac{y+1}{y-1}\right)^\ell - 1}\right) \\ &< 1 + \varepsilon, \end{aligned}$$

as required. □

The proof of the following corollary is straightforward and is essentially identical to the proof of Corollary 22.

Corollary 24. For any $\lambda \in (0, 1)$, $\langle I_2^\lambda \cup \mathcal{B}'_0 \rangle_\omega = \mathcal{I}_{\text{ferro}}$.

6 ω -clones defined by Fourier coefficients

6.1 Properties of Fourier coefficients

The proofs of the following three lemmas are routine calculations and we defer them to Appendix A.

Lemma 25. Let f and g be functions in \mathcal{B}_k .

(i) For any permutation π of $[k]$, $\widehat{f^\pi}(\mathbf{x}) = \widehat{f}(\pi(\mathbf{x}))$.

(ii) If $h(\mathbf{xz}) = f(\mathbf{x})$, then $\widehat{h}(\mathbf{x0}) = \widehat{f}(\mathbf{x})$ and $\widehat{h}(\mathbf{x1}) = 0$.

(iii) If $h(\mathbf{x}) = f(\mathbf{x0}) + f(\mathbf{x1})$, then $\widehat{h}(\mathbf{x}) = 2\widehat{f}(\mathbf{x0})$.

(iv) If $h(\mathbf{x}) = f(\overline{\mathbf{x}})$, then $\widehat{h}(\mathbf{x}) = (-1)^{|\mathbf{x}|}\widehat{f}(\mathbf{x})$.

(v) If $\|g - f\|_\infty < \varepsilon$, then $\|\widehat{g} - \widehat{f}\|_\infty < \varepsilon$.

(vi) If $k = 0$ then $\widehat{f} = f$.

It is also well-known (see, e.g., [8, 17]) that, if $f, g \in \mathcal{B}_k$, and h is defined by $h(\mathbf{x}) = f(\mathbf{x})g(\mathbf{x})$, then \widehat{h} is given by the convolution

$$\widehat{h}(\mathbf{x}) = \sum_{\mathbf{w} \in \{0,1\}^k} \widehat{f}(\mathbf{w})\widehat{g}(\mathbf{w} \oplus \mathbf{x}). \quad (1)$$

We will later need to know the Fourier coefficients of hypergraph Ising functions and of the parity functions defined in Definition 17.

Lemma 26. *For any k and λ ,*

$$\widehat{I}_k^\lambda(\mathbf{x}) = \begin{cases} \lambda + (1 - \lambda)/2^{k-1} & \text{if } \mathbf{x} = \mathbf{0} \\ (1 - \lambda)/2^{k-1} & \text{if } |\mathbf{x}| \text{ is even and positive} \\ 0 & \text{if } |\mathbf{x}| \text{ is odd.} \end{cases}$$

Lemma 27. *For any k and λ , $\widehat{\text{Par}}_k^\lambda(\mathbf{0}) = \frac{1}{2}(1 + \lambda)$, $\widehat{\text{Par}}_k^\lambda(\mathbf{1}) = \frac{1}{2}(1 - \lambda)$ and $\widehat{\text{Par}}_k^\lambda(\mathbf{x}) = 0$ for any $\mathbf{x} \notin \{\mathbf{0}, \mathbf{1}\}$.*

6.2 \mathcal{P} and \mathcal{PN}

Recall from Definition 8 that \mathcal{P} is the class of functions f such that $\widehat{f}(\mathbf{x}) \geq 0$ for all \mathbf{x} , and that \mathcal{PN} is the class of functions f such that $\widehat{f}(\mathbf{x}) \geq 0$ if $|\mathbf{x}|$ is even and $\widehat{f}(\mathbf{x}) \leq 0$ if $|\mathbf{x}|$ is odd. We first show that \mathcal{P} and \mathcal{PN} are ω -clones, and that they contain \mathcal{B}_0 .

Theorem 28. $\langle \mathcal{P} \rangle_\omega = \mathcal{P}$ and $\mathcal{B}_0 \subseteq \mathcal{P}$.

Proof. By the definition of ω -clones, the fact that \mathcal{P} is an ω -clone follows from the fact that it contains EQ and that it is closed under the various operations.

- It is easily verified (for example, apply Lemma 26 with $\lambda = 0$) that $\widehat{\text{EQ}} = \frac{1}{2}\text{EQ}$, which is a non-negative function. Therefore, $\text{EQ} \in \mathcal{P}$.
- For closure under permuting arguments, suppose that $f \in \mathcal{P}$ and let $h = f^\pi$ for some permutation π . By Lemma 25(i), \widehat{f} and \widehat{h} have the same range, so \widehat{h} is a nonnegative function, so $h \in \mathcal{P}$.
- For closure under introducing fictitious arguments, let $f \in \mathcal{P}$ and define $h(\mathbf{xy}) = f(\mathbf{x})$. Then $h \in \mathcal{P}$ because, by Lemma 25(ii), every Fourier coefficient of h is either zero or a Fourier coefficient of f .
- For closure under summation, let $f \in \mathcal{P}$ and define $h(\mathbf{x}) = f(\mathbf{x0}) + f(\mathbf{x1})$. By Lemma 25(iii), $\widehat{h}(\mathbf{x}) = 2\widehat{f}(\mathbf{x0}) \geq 0$ for any \mathbf{x} , so $h \in \mathcal{P}$.

- For closure under products, let $f, g \in \mathcal{P}$. $\widehat{fg}(\mathbf{x}) = \sum_{\mathbf{w}} \widehat{f}(\mathbf{w}) \widehat{g}(\mathbf{w} \oplus \mathbf{x}) \geq 0$, since every term of the sum is nonnegative, so $fg \in \mathcal{P}$.
- For closure under limits, let f be a function and suppose that, for every $\varepsilon > 0$, there is some $f_\varepsilon \in \mathcal{P}$ with $\|f_\varepsilon - f\|_\infty < \varepsilon$ (this is a weaker condition than requiring all such f_ε to be in $\langle \mathcal{G} \rangle$ for some finite $\mathcal{G} \subseteq \mathcal{P}$). Then, by Lemma 25(v), $\|\widehat{f}_\varepsilon - \widehat{f}\|_\infty < \varepsilon$. In particular, $\widehat{f}_\varepsilon(\mathbf{x}) \geq 0$ for all \mathbf{x} so, for all \mathbf{x} and all $\varepsilon > 0$, $\widehat{f}(\mathbf{x}) > -\varepsilon$. Therefore, $\widehat{f}(\mathbf{x}) \geq 0$ and $f \in \mathcal{P}$.

We now show that $\mathcal{B}_0 \subseteq \mathcal{P}$. Consider any $c \in \mathbb{R}_{\geq 0}$ and let f_c be the nullary function in \mathcal{B}_0 with range $\{c\}$. Let g_c be the unary function defined by $g_c(0) = g_c(1) = c/2$. $\widehat{g}_c(0) = c$ and $\widehat{g}_c(1) = 0$, so $g_c \in \mathcal{P}$. But $f_c(x_1) = \sum_{x_1} g_c(x_1)$ and ω -clones are closed under summation, so $f_c \in \mathcal{P}$. \square

Definition 29. Consider a function $f \in \mathcal{B}_k$. We define the *complement* \bar{f} of f by $\bar{f}(\mathbf{x}) = f(\bar{\mathbf{x}})$.

Theorem 30. $\langle \mathcal{PN} \rangle_\omega = \mathcal{PN}$ and $\mathcal{B}_0 \subseteq \mathcal{PN}$.

Proof. By Lemma 25(iv), $\mathcal{PN} = \{f \mid \bar{f} \in \mathcal{P}\}$. We first show that \mathcal{PN} is an ω -clone.

- Since $\overline{\text{EQ}} = \text{EQ}$ and \mathcal{P} contains EQ, \mathcal{PN} also contains EQ.
- For closure under permuting arguments, let f be a k -ary function in \mathcal{PN} and let π be a permutation of $[k]$. By Lemma 25(i), $\widehat{f^\pi}(\mathbf{x}) = \widehat{f}(\pi(\mathbf{x}))$ and, since $|\mathbf{x}| = |\pi(\mathbf{x})|$, we have $f^\pi \in \mathcal{PN}$.
- For closure under introducing fictitious arguments, let $f \in \mathcal{PN}$ and define $h(\mathbf{xy}) = f(\mathbf{x})$. Then $\bar{h}(\mathbf{xy}) = h(\overline{\mathbf{xy}}) = f(\bar{\mathbf{x}}) = \bar{f}(\mathbf{x}) \in \mathcal{P}$, so $h \in \mathcal{PN}$.
- For closure under summation, let $f \in \mathcal{PN}$, so $\bar{f} \in \mathcal{P}$. Define $h(\mathbf{x}) = f(\mathbf{x}0) + f(\mathbf{x}1)$. Then $\bar{h}(\mathbf{x}) = h(\bar{\mathbf{x}}) = f(\bar{\mathbf{x}}0) + f(\bar{\mathbf{x}}1) = \bar{f}(\mathbf{x}1) + \bar{f}(\mathbf{x}0) \in \mathcal{P}$, so $h \in \mathcal{PN}$.
- For closure under products, suppose $f, g \in \mathcal{PN}$ and let $h(\mathbf{x}) = f(\mathbf{x})g(\mathbf{x})$. Then $\bar{h}(\mathbf{x}) = h(\bar{\mathbf{x}}) = f(\bar{\mathbf{x}})g(\bar{\mathbf{x}}) = \bar{f}(\mathbf{x})\bar{g}(\mathbf{x}) \in \mathcal{P}$, so $h \in \mathcal{PN}$.
- For closure under limits, let f be a function and suppose that, for all $\varepsilon > 0$, there is some $f_\varepsilon \in \mathcal{PN}$ such that $\|f_\varepsilon - f\|_\infty < \varepsilon$. We must show that $f \in \mathcal{PN}$.
By Lemma 25(v), $\|\widehat{f}_\varepsilon - \widehat{f}\|_\infty < \varepsilon$. In particular, $\widehat{f}_\varepsilon(\mathbf{x}) \geq 0$ for all even-weight \mathbf{x} , and $\widehat{f}_\varepsilon(\mathbf{x}) \leq 0$ for all odd-weight \mathbf{x} . Therefore, for all even-weight \mathbf{x} , and all $\varepsilon > 0$, $\widehat{f}(\mathbf{x}) > -\varepsilon$, so $\widehat{f}(\mathbf{x}) \geq 0$. Similarly, $\widehat{f}(\mathbf{x}) \leq 0$ for all odd-weight \mathbf{x} , so $f \in \mathcal{PN}$.

The proof that $\mathcal{B}_0 \subseteq \mathcal{PN}$ is the same as the proof that $\mathcal{B}_0 \subseteq \mathcal{P}$ (see the proof of Theorem 28). \square

We now investigate the position of \mathcal{P} and \mathcal{PN} in the lattice \mathcal{L}' from Theorem 10. To do this, we use two technical lemmas, which we will also use in Section 7.2.

Lemma 31. *Let $f \in \mathcal{B}_n$. If $\widehat{f}(\mathbf{a}) < 0$ for some $\mathbf{a} \in \{0, 1\}^n$, then there is a function $g \in \langle \{f\} \rangle$ such that $\widehat{g}(\mathbf{1}) < 0$.*

Proof. Since $\widehat{f}(\mathbf{a}) \neq 0$, f cannot be the constant zero function. Therefore, $f(\mathbf{x}) > 0$ for some $\mathbf{x} \in \{0, 1\}^n$, which means that $\widehat{f}(\mathbf{0}) > 0$, so $\mathbf{a} \neq \mathbf{0}$. Since functional clones are closed under permuting arguments, and (by Lemma 25(i)), permuting arguments just permutes Fourier coefficients, we may assume that, for some $k \in [n]$, $a_1 = \dots = a_k = 1$ and $a_{k+1} = \dots = a_n = 0$. Let

$$g(x_1, \dots, x_k) = \sum_{x_{k+1}, \dots, x_n} f(x_1, \dots, x_n).$$

By Lemma 25(iii), $\widehat{g}(\mathbf{1}) = 2^{n-k} \widehat{f}(\mathbf{a}) < 0$. □

Definition 32. A function $f: \{0, 1\}^k \rightarrow \mathbb{R}_{\geq 0}$ is *permissive* if its range is $\mathbb{R}_{>0}$.

Lemma 33. Let $f \in \mathcal{B}_n$ with $\widehat{f}(\mathbf{1}) < 0$. Then, for every $k > 0$, there is a k -ary permissive function $h \in \langle \{f, \text{Par}_{k+n}^{1/2}\} \rangle$ such that $\widehat{h}(\mathbf{1}) < 0$.

Proof. Let $\mathbf{x} = (x_1, \dots, x_k)$ and $\mathbf{y} = (y_1, \dots, y_n)$. Define the functions

$$\begin{aligned} f'(\mathbf{x}, \mathbf{y}) &= f(\mathbf{y}) \\ g(\mathbf{x}, \mathbf{y}) &= \text{Par}_{k+n}^{1/2}(\mathbf{x}, \mathbf{y}) f'(\mathbf{x}, \mathbf{y}) \\ h(\mathbf{x}) &= \sum_{\mathbf{y} \in \{0, 1\}^n} g(\mathbf{x}, \mathbf{y}). \end{aligned}$$

Thus, $h \in \langle \{f, \text{Par}_{k+n}^{1/2}\} \rangle$. Further, since $\text{Par}_{k+n}^{1/2}$ is permissive and f is not the constant zero function (because $\widehat{f}(\mathbf{1}) \neq 0$), h is also permissive.

For the claim that $\widehat{h}(\mathbf{1}) < 0$, we have the following. The first equality is by Lemma 25(iii) and the second by Equation (1) from the beginning of Section 6. The third equality is because the first k arguments of f' are fictitious so, by Lemma 25(ii), $\widehat{f}'(\bar{\mathbf{v}}, \mathbf{w})$ is $\widehat{f}(\mathbf{w})$ when $\mathbf{v} = \mathbf{1}$ and is zero, otherwise. The final equality is because, by Lemma 27, $\widehat{\text{Par}}_{k+n}^{1/2}(\mathbf{1}, \mathbf{w})$ is $\frac{1}{4}$ when $\mathbf{w} = \mathbf{1}$ and is zero, otherwise.

$$\begin{aligned} \widehat{h}(\mathbf{1}) &= 2^n \widehat{g}(\mathbf{1}, \mathbf{0}) \\ &= 2^n \sum_{\mathbf{v}\mathbf{w} \in \{0, 1\}^{k+n}} \widehat{\text{Par}}_{k+n}^{1/2}(\mathbf{v}, \mathbf{w}) \widehat{f}'(\bar{\mathbf{v}}, \mathbf{w}) \\ &= 2^n \sum_{\mathbf{w} \in \{0, 1\}^n} \widehat{\text{Par}}_{k+n}^{1/2}(\mathbf{1}, \mathbf{w}) \widehat{f}(\mathbf{w}) \\ &= 2^{n-2} \widehat{f}(\mathbf{1}) < 0. \end{aligned} \quad \square$$

Lemma 34. \mathcal{P} is maximal in \mathcal{B} .

Proof. Consider any n -ary $f \in \mathcal{B} \setminus \mathcal{P}$. By definition, $\widehat{f}(\mathbf{a}) < 0$ for some $\mathbf{a} \in \{0, 1\}^n$ and, by Lemma 31, we may assume that $\widehat{f}(\mathbf{1}) < 0$. By Lemma 33, there is a permissive unary function $h \in \langle \{f, \text{Par}_{n+1}^{1/2}\} \rangle$ such that $\widehat{h}(\mathbf{1}) < 0$. By Lemma 27, $\text{Par}_{n+1}^{1/2} \in \mathcal{P}$ for every n , so we have $h \in \langle \mathcal{P} \cup \{f\} \rangle$.

Since h is permissive and $\widehat{h}(\mathbf{1}) = \frac{1}{2}(h(0) - h(1))$, we have $h(1) > h(0) > 0$. We may further assume that $h(0) < \frac{1}{2}$ and $h(1) = 1$: if this is not the case, replace h with the function $h'(x) = (h(x)/h(1))^j$ for any sufficiently large integer j . The function h' is in $\langle \mathcal{P} \cup \{f\} \rangle$ since $h \in \langle \mathcal{P} \cup \{f\} \rangle$ and nullary functions such as $1/h(1)$ are in \mathcal{P} by Theorem 28.

Now, consider the symmetric binary function $g = [h(0)^{-2}, h(0)^{-1}, 0]$, which is in \mathcal{P} by the assumption on h . We have

$$\text{NAND}(x, y) = [1, 1, 0] = g(x, y) h(x) h(y) \in \langle \mathcal{P} \cup \{f\} \rangle.$$

By [2, Corollary 13.2(ii)], any ω -clone that contains NAND, a unary function h such that $h(1) > h(0) > 0$ and the nullary function $1/2$ also contains all of \mathcal{B}_1 . Therefore, $\mathcal{B}_1 \cup \{\text{NAND}\} \subseteq \langle \mathcal{P} \cup \{f\} \rangle_\omega$. By Lemmas 7.1 and 8.1 of [2], this implies that that $\langle \mathcal{P} \cup \{f\} \rangle_\omega = \mathcal{B}$, so \mathcal{P} is maximal in \mathcal{B} . \square

Corollary 35. \mathcal{PN} is maximal in \mathcal{B} .

Proof. Let $f \in \mathcal{B} \setminus \mathcal{PN}$. By Lemma 25(iv), we have $\bar{f} \notin \mathcal{P}$. We will now show that $\langle \mathcal{PN} \cup \{f\} \rangle = \{\bar{g} \mid g \in \langle \mathcal{P} \cup \{\bar{f}\} \rangle\}$. To see this, suppose that $g \in \langle \mathcal{PN} \cup \{f\} \rangle$. Then g is defined by a summation of a product of functions in $\mathcal{A}(\mathcal{PN} \cup \{f\})$. Complementing all of the functions in $\mathcal{A}(\mathcal{PN} \cup \{f\})$ exchanges the roles of 0's and 1's so the summing the product of the complements defines \bar{g} . Since the complements of the functions in $\mathcal{A}(\mathcal{PN} \cup \{f\})$ are in $\mathcal{A}(\mathcal{P} \cup \{\bar{f}\})$, this shows that \bar{g} is in $\langle \mathcal{P} \cup \{\bar{f}\} \rangle$. A similar argument gives the other direction.

Closing under limits, we get

$$\begin{aligned} \langle \mathcal{PN} \cup \{f\} \rangle_\omega &= \{\bar{g} \mid g \in \langle \mathcal{P} \cup \{\bar{f}\} \rangle_\omega\} \\ &= \{\bar{g} \mid g \in \mathcal{B}\} \\ &= \mathcal{B}. \end{aligned} \quad \square$$

Corollary 36. $\langle \mathcal{P} \cup \mathcal{PN} \rangle_\omega = \mathcal{B}$

Proof. Consider the symmetric function $f = [0, 1, 2]$. We have $\hat{f} = [1, -\frac{1}{2}, 0]$, so $f \in \mathcal{PN} \setminus \mathcal{P}$ and the result is immediate from maximality of \mathcal{P} in \mathcal{B} (Lemma 34). \square

7 Self-dual functions

Recall from Definition 8 that \mathcal{SD} is the class of self-dual functions, i.e., functions for which $f(\mathbf{x}) = f(\bar{\mathbf{x}})$ for all \mathbf{x} of appropriate arity.

Theorem 37. $\langle \mathcal{SD} \rangle_\omega = \mathcal{SD}$ and $\mathcal{B}_0 \subseteq \mathcal{SD}$.

Proof. By the definition of ω -clones, the fact that \mathcal{SD} is an ω -clone follows from the fact that it contains EQ and that it is closed under the various operations.

- The equality function is clearly self-dual so it is in \mathcal{SD} .
- For closure under permuting arguments, let $f \in \mathcal{SD}$ be a k -ary function and let π be a permutation of $[k]$. Then $f^\pi \in \mathcal{SD}$, since $f^\pi(\bar{\mathbf{x}}) = f(\pi(\bar{\mathbf{x}})) = f(\overline{\pi(\mathbf{x})}) = f(\pi(\mathbf{x})) = f^\pi(\mathbf{x})$.
- For closure under introducing fictitious arguments, let $f \in \mathcal{SD}$ and define $h(\mathbf{xy}) = f(\mathbf{x})$. Then $h(\bar{\mathbf{x}}\bar{y}) = f(\bar{\mathbf{x}}) = f(\mathbf{x}) = h(\mathbf{xy})$, so h is self-dual.
- For closure under summation, let $f \in \mathcal{SD}$ and define $h(\mathbf{x}) = f(\mathbf{x}0) + f(\mathbf{x}1)$. Then $h(\bar{\mathbf{x}}) = f(\bar{\mathbf{x}}0) + f(\bar{\mathbf{x}}1) = f(\mathbf{x}1) + f(\mathbf{x}0) = h(\mathbf{x})$, so $h \in \mathcal{SD}$.

- For closure under products, let $f, g \in \mathcal{SD}$ and consider $h(\mathbf{x}) = f(\mathbf{x})g(\mathbf{x})$. We have $h(\bar{\mathbf{x}}) = f(\bar{\mathbf{x}})g(\bar{\mathbf{x}}) = f(\mathbf{x})g(\mathbf{x}) = h(\mathbf{x})$, so $h \in \mathcal{SD}$.
- For closure under limits, let $f \in \mathcal{B}$ and suppose that, for all $\varepsilon > 0$, there is some $f_\varepsilon \in \mathcal{SD}$ such that $\|f - f_\varepsilon\|_\infty < \varepsilon$. We must show that $f \in \mathcal{SD}$.

For any \mathbf{x} , $|f(\mathbf{x}) - f_\varepsilon(\mathbf{x})| < \varepsilon$ and $|f(\bar{\mathbf{x}}) - f_\varepsilon(\bar{\mathbf{x}})| < \varepsilon$. But, since f_ε is self-dual, this gives $|f(\bar{\mathbf{x}}) - f_\varepsilon(\mathbf{x})| < \varepsilon$. It follows that $|f(\mathbf{x}) - f(\bar{\mathbf{x}})| < 2\varepsilon$ for all $\varepsilon > 0$, so $f(\mathbf{x}) = f(\bar{\mathbf{x}})$, so $f \in \mathcal{SD}$.

The proof that $\mathcal{B}_0 \subseteq \mathcal{SD}$ is the same as the proof that $\mathcal{B}_0 \subseteq \mathcal{P}$ in the proof of Theorem 28. \square

It turns out that the functions in \mathcal{SD} also have a natural characterisation in terms of their Fourier transforms. This allows us to study the relationship between \mathcal{SD} and the ω -clones from Section 6.

Lemma 38. *A k -ary function f is in \mathcal{SD} if and only if $\hat{f}(\mathbf{x}) = 0$ for all \mathbf{x} with odd Hamming weight.*

Proof. Suppose $f \in \mathcal{SD}$. We have $f(\mathbf{x}) = f(\bar{\mathbf{x}})$ so, by Lemma 25(iv), $\hat{f}(\mathbf{x}) = (-1)^{|\mathbf{x}|}\hat{f}(\mathbf{x})$. When $|\mathbf{x}|$ is odd, this implies that $\hat{f}(\mathbf{x}) = 0$.

Conversely, if $\hat{f}(\mathbf{x}) = 0$ for all \mathbf{x} with $|\mathbf{x}|$ odd, then

$$\begin{aligned} 2^k f(\mathbf{x}) &= \sum_{\mathbf{w} \in \{0,1\}^k} (-1)^{|\mathbf{w} \wedge \mathbf{x}|} \hat{f}(\mathbf{w}) \\ &= \sum_{\substack{\mathbf{w} \in \{0,1\}^k, \\ |\mathbf{w}| \text{ even}}} (-1)^{|\mathbf{w} \wedge \mathbf{x}|} \hat{f}(\mathbf{w}) \\ &= \sum_{\substack{\mathbf{w} \in \{0,1\}^k, \\ |\mathbf{w}| \text{ even}}} (-1)^{|\mathbf{w} \wedge \bar{\mathbf{x}}|} \hat{f}(\mathbf{w}) = 2^k f(\bar{\mathbf{x}}), \end{aligned}$$

so $f \in \mathcal{SD}$. \square

Theorem 39. $\mathcal{P} \cap \mathcal{SD} = \mathcal{PN} \cap \mathcal{SD} = \mathcal{P} \cap \mathcal{PN}$.

Proof. It is immediate from the definitions and Lemma 38 that each of these is the class of functions f such that $\hat{f}(\mathbf{x}) \geq 0$ if $|\mathbf{x}|$ is even and $\hat{f}(\mathbf{x}) = 0$ if $|\mathbf{x}|$ is odd. \square

Recall from Section 1.3 that the intersection of two ω -clones is an ω -clone. In the light of Theorem 39, we make the following definition.

Definition 40. Let \mathcal{SDP} be the ω -clone $\mathcal{SD} \cap \mathcal{P} \cap \mathcal{PN}$.

Theorem 39 makes it clear that $\mathcal{SDP} = \mathcal{P} \cap \mathcal{SD} = \mathcal{PN} \cap \mathcal{SD} = \mathcal{P} \cap \mathcal{PN}$.

Lemma 41. \mathcal{SD} , \mathcal{P} and \mathcal{PN} are pairwise incomparable under subset inclusion.

Proof. Consider the functions $f = [0, 1, 0]$ (the binary disequality function), $g = [2, 1, 0]$ and $h = [0, 1, 2]$. We have $\hat{f} = [\frac{1}{2}, 0, -\frac{1}{2}]$, $\hat{g} = [1, \frac{1}{2}, 0]$ and $\hat{h} = [1, -\frac{1}{2}, 0]$. Lemma 38 and the definitions of \mathcal{P} and \mathcal{PN} imply that, among the three ω -clones in the statement, f is only in \mathcal{SD} , g is only in \mathcal{P} and h is only in \mathcal{PN} . \square

For the relationship between \mathcal{SDP} and $\mathcal{H}_{\text{ferro}}$, we use the concept of *log-supermodular* functions. A function $f: \{0, 1\}^k \rightarrow \mathbb{R}_{\geq 0}$ is log-supermodular if $f(\mathbf{x} \vee \mathbf{y}) f(\mathbf{x} \wedge \mathbf{y}) \geq f(\mathbf{x}) f(\mathbf{y})$ for all $\mathbf{x}, \mathbf{y} \in \{0, 1\}^k$, where \vee and \wedge are applied bitwise.

Definition 42. Let \mathcal{LSM} be the set of all log-supermodular functions.

\mathcal{LSM} is an ω -clone [2, Lemma 4.2]. Note that all unary functions are trivially log-supermodular. Nullary functions are also log-supermodular, for example, using the proof that $\mathcal{B}_0 \subseteq \mathcal{P}$ (proof of Theorem 28).

The following characterisation of permissive log-supermodular functions of arity at least 2 is due essentially to Topkis [19] (see also [2, Lemma 5.1]). It provides a simple way to check that a permissive function is log-supermodular. A 2-pinning of a k -ary function f (with $k \geq 2$) is any binary function $g(x, y) = f(z_1, \dots, z_k)$ where each $z_i \in \{0, 1, x, y\}$ ($i \in [k]$), such that x and y each appear exactly once in the sequence z_1, \dots, z_k . It is immediate from the definition that every 2-pinning of a log-supermodular function f is also log-supermodular. The following lemma states that, for permissive functions, this condition is also sufficient.

Lemma 43 ([19]). *A permissive k -ary function is log-supermodular if, and only if, every 2-pinning of f is log-supermodular.*

Theorem 44. $\mathcal{H}_{\text{ferro}} \subset \mathcal{SDP}$.

Proof. Inspection of Lemma 26 shows that $\mathcal{H}_{\text{ferro}} \subseteq \mathcal{P} \cap \mathcal{PN}$ and, by Theorem 39, $\mathcal{P} \cap \mathcal{PN} = \mathcal{SDP}$. It remains to show that the inclusion is strict.

It is easy to check that every function in $\text{FerroHyperIsing} \cup \mathcal{B}_0$ is log-supermodular. It follows that $\mathcal{H}_{\text{ferro}}$ is a subset of the ω -clone of all log-supermodular functions so, in particular, every function in $\mathcal{H}_{\text{ferro}}$ is log-supermodular.

Consider the 4-ary function $f = [13, 4, 1, 4, 13]$. This function is not log-supermodular by Lemma 43, since the pinning $g(x, y) = f(x, y, 0, 0)$ has $g(1, 1)g(0, 0) = 13 < g(0, 1)g(1, 0) = 16$. Therefore, $f \notin \mathcal{H}_{\text{ferro}}$. However, $f \in \mathcal{SD}$ and we have $f = [4, 0, \frac{3}{2}, 0, 0]$ (the odd-weight coefficients are zero by Lemma 38) so $f \in \mathcal{P} \cap \mathcal{PN} \cap \mathcal{SD} = \mathcal{SDP}$. \square

7.1 Self-dual functions and Ising

In this section, we prove that $\mathcal{SD} = \mathcal{I}_{\text{anti}}$ (Theorem 48). To do this, we introduce a functional clone, $\mathcal{PAR}_{\text{ev}}$, of weighted, even-arity parity functions. Recall from Definition 17 that, for $k \in \mathbb{N}$ and $\lambda \in \mathbb{R}_{\geq 0}$, $\text{Par}_k^\lambda(\mathbf{x}) = 1$ if $|\mathbf{x}|$ is even, and $\text{Par}_k^\lambda(\mathbf{x}) = \lambda$, otherwise. Note that, when k is even, Par_k^λ is self-dual. Note also that $\text{Par}_2^\lambda = I_2^\lambda$. Our new clone is

$$\mathcal{PAR}_{\text{ev}} = \langle \{\text{Par}_k^\lambda \mid k \text{ is even}, \lambda \in \mathbb{R}_{\geq 0}\} \rangle.$$

Lemma 45. $\mathcal{SD} \subseteq \langle \mathcal{PAR}_{\text{ev}} \cup \mathcal{B}'_0 \rangle_\omega$.

Proof. Recall from Definition 18 that $\text{Par}'_k = \{\text{Par}_k^{1/e}, \text{Par}_k^{1/2}, \text{Par}_k^2, \text{Par}_k^e\}$. Let $\text{Par}'^{\text{ev}}_{\leq k} = \bigcup_{1 \leq j \leq \lfloor k/2 \rfloor} \text{Par}'_{2j}$. Recall from Definition 32 that a function $F: \{0, 1\}^k \rightarrow \mathbb{R}_{\geq 0}$ is *permissive* if its range is $\mathbb{R}_{>0}$. The proof splits into two parts. First, we show that every k -ary permissive function in \mathcal{SD} is a limit of $\langle \mathcal{B}'_0 \cup \text{Par}'^{\text{ev}}_{\leq k} \rangle$. Then we show the same for every other function in \mathcal{SD} .

Part One: Consider any permissive k -ary function $F \in \mathcal{SD}$. Let $f(\mathbf{y}) = \log F(\mathbf{y})$. (Note that f is not necessarily in \mathcal{B} , since its range may include negative numbers; this is not a problem.) By the definition of the Fourier transform,

$$f(\mathbf{y}) = 2^{-k} \sum_{\mathbf{w} \in \{0,1\}^k} (-1)^{|\mathbf{w} \wedge \mathbf{y}|} \widehat{f}(\mathbf{w}).$$

Exponentiating gives

$$F(\mathbf{y}) = \prod_{\mathbf{w} \in \{0,1\}^k} \exp(2^{-k} (-1)^{|\mathbf{w} \wedge \mathbf{y}|} \widehat{f}(\mathbf{w})).$$

Since f is self-dual, we may restrict the product to \mathbf{w} with even Hamming weight, since the odd-weight terms vanish by Lemma 38. Let

$$G_{\mathbf{w}}(\mathbf{y}) = \exp(2^{-k} (-1)^{|\mathbf{w} \wedge \mathbf{y}|} \widehat{f}(\mathbf{w})).$$

Then $F(\mathbf{y}) = \prod_{\mathbf{w} \in \{0,1\}^k: |\mathbf{w}| \text{ is even}} G_{\mathbf{w}}(\mathbf{y})$, so to finish (using Lemma 1) we just have to show that, for any $\mathbf{w} \in \{0,1\}^k$ with even Hamming weight, $G_{\mathbf{w}}$ is a limit of $\langle \mathcal{B}'_0 \cup \text{Par}'^{\text{ev}}_{\leq k} \rangle$.

Consider any such \mathbf{w} . Let $j = |\mathbf{w}|/2$. Given any \mathbf{y} , let \mathbf{z} be the arity- $2j$ tuple obtained from \mathbf{y} by deleting all positions that are 0 in \mathbf{w} . Let

$$G'_{\mathbf{w}}(\mathbf{z}) = \exp(2^{-k} (-1)^{|\mathbf{z}|} \widehat{f}(\mathbf{w})).$$

Note that the arity- k function $G_{\mathbf{w}}$ is constructed from the arity- $2j$ function $G'_{\mathbf{w}}$ by adding fictitious arguments. We will show that every arity- $2j$ function $G'_{\mathbf{w}}$ is a limit of a function in $\langle \mathcal{B}'_0 \cup \text{Par}'_{2j} \rangle$. This is all that we need since, by the closure of ω -clones under the addition of fictitious arguments, it also implies that $G_{\mathbf{w}}$ is a limit of $\langle \mathcal{B}'_0 \cup \text{Par}'^{\text{ev}}_{\leq k} \rangle$.

Now $G'_{\mathbf{w}}(\mathbf{z})$ is $\exp(2^{-k} \widehat{f}(\mathbf{w}))$ if $|\mathbf{z}|$ is even and it is $\exp(-2^{-k} \widehat{f}(\mathbf{w}))$ otherwise. Therefore, $G'_{\mathbf{w}}$ is $\lambda \times \text{Par}'_{2j}^{1/\lambda^2}$, where $\lambda = \exp(2^{-k} \widehat{f}(\mathbf{w}))$. To finish note that λ is a limit of $\langle \mathcal{B}'_0 \rangle$ (by Lemma 16) and $\text{Par}'_{2j}^{1/\lambda^2}$ is a limit of $\langle \text{Par}'_{2j} \rangle$ (by Lemma 19). Finally, their product is a limit of $\langle \mathcal{B}'_0 \cup \text{Par}'_{2j} \rangle$ by Lemma 1.

Part Two: Consider any non-permissive k -ary function $F \in \mathcal{SD}$. Let G be the $(k+2)$ -ary function

$$G(\mathbf{x}yz) = \begin{cases} F(\mathbf{x}) & \text{if } y \neq z \\ 1 & \text{if } y = z. \end{cases}$$

For every positive integer j , let H_j be the k -ary function

$$H_j(\mathbf{x}) = 2^{-(j+1)} \sum_{y,z \in \{0,1\}} G(\mathbf{x}yz) (I_2^2(yz))^j.$$

Note that H_j is a good approximation for F in the sense that $H_j(\mathbf{x}) = F(\mathbf{x}) + 2^{-j}$. Since F is self-dual, so is H_j ; since $H_j(\mathbf{x}) > F(\mathbf{x}) \geq 0$, H_j is permissive. By Part One, H_j is a limit of $\langle \mathcal{B}'_0 \cup \text{Par}'^{\text{ev}}_{\leq k} \rangle$.

For any $\varepsilon > 0$, there is j such that $2^{-j} < \varepsilon$ and thus $\|F - H_j\|_{\infty} < \varepsilon$. By transitivity of limits (Lemma 1), F itself is a limit of $\langle \mathcal{B}'_0 \cup \text{Par}'^{\text{ev}}_{\leq k} \rangle$. \square

We define the family of k -ary functions

$$\oplus_k^{\text{odd}}(\mathbf{x}) = \begin{cases} 1 & \text{if } |\mathbf{x}| \text{ is odd} \\ 0 & \text{otherwise.} \end{cases}$$

The function \oplus_4^{odd} turns out to be particularly useful.

Lemma 46. $\mathcal{PAR}_{\text{ev}} \subseteq \langle \text{AntiFerrolsing} \cup \mathcal{B}'_0 \cup \{\oplus_4^{\text{odd}}\} \rangle_\omega$.

Proof. For any $k > 1$, we have

$$\text{Par}_k^\lambda(x_1, \dots, x_k) = \lambda \sum_{y \in \{0,1\}} \oplus_k^{\text{odd}}(x_1, \dots, x_{k-1}, y) I_2^{1/\lambda}(y, x_k).$$

If $\lambda \leq 1$ then $I_2^{1/\lambda}$ is in AntiFerrolsing . Otherwise, by Lemma 20, there is a $c \in \mathcal{B}_0$ and a $I_2^{\lambda'} \in \text{AntiFerrolsing}$ such that $I_2^{1/\lambda} \in \langle \{c, I_2^{\lambda'}\} \rangle$. The nullary functions λ and c are limits of $\langle \mathcal{B}'_0 \rangle_\omega$ by Lemma 16.

We now show that, for any even $k > 0$, $\oplus_k^{\text{odd}} \in \langle \oplus_4^{\text{odd}} \rangle$. For $k = 2$, we have

$$\oplus_2^{\text{odd}}(x, y) = \sum_{w,z} \oplus_4^{\text{odd}}(x, y, w, z) \text{EQ}(x, w) \text{EQ}(x, z),$$

and the case $k = 4$ is trivial. For $k \geq 6$,

$$\oplus_k^{\text{odd}}(x_1, \dots, x_k) = \sum_{y,z \in \{0,1\}} \oplus_4^{\text{odd}}(x_1, x_2, x_3, y) \oplus_2^{\text{odd}}(y, z) \oplus_{k-2}^{\text{odd}}(z, x_4, \dots, x_k),$$

and the claim follows by induction on even k . The lemma follows by Lemma 1. \square

Lemma 47. *There exist $\lambda, \lambda' > 1$ such that $\oplus_4^{\text{odd}} \in \langle \{I_2^\lambda, I_2^{\lambda'}\} \cup \mathcal{B}'_0 \rangle_\omega$.*

Proof. Let $\lambda' = 2$. (Any value that is greater than one would do, but we take $\lambda' = 2$ for concreteness.) Let $\lambda = \sqrt{\lambda'^4 + \sqrt{\lambda'^8 - 1}}$. Note that $\lambda > 1$. Consider the (self-dual, symmetric) function

$$f(x_1, x_2, x_3, x_4) = \sum_{y \in \{0,1\}} \prod_{i \in [4]} I_2^\lambda(x_i, y) \prod_{i \in [4], j \neq i \in [4]} I_2^{\lambda'}(x_i, x_j).$$

Note that $f(0, 0, 0, 0) = (\lambda^4 + 1) = 2\lambda^2\lambda'^4 = f(0, 0, 1, 1) = 1023$. Also, $f(0, 0, 0, 1) = (\lambda + \lambda^3)\lambda'^3 \approx 1491$. Now, define the function $g(\mathbf{x}) = f(\mathbf{x})/f(0, 0, 0, 1)$. Note that $f(\mathbf{x}) \in \langle \{I_2^\lambda, I_2^{\lambda'}\} \rangle$ and $1/f(0, 0, 0, 1) \in \langle \mathcal{B}'_0 \rangle_\omega$ by Lemma 16. So, by Lemma 1, for every positive integer j , $g^j \in \langle \{I_2^\lambda, I_2^{\lambda'}\} \cup \mathcal{B}'_0 \rangle_\omega$.

We have $g(\mathbf{x}) = 1$ if $|\mathbf{x}|$ is odd, and $g(\mathbf{x}) < 1$, otherwise. This gives $\|g - \oplus_4^{\text{odd}}\|_\infty = g(0, 0, 0, 0) < 1$ so, for any $\varepsilon > 0$, we can choose an integer j such that $\|g^j - \oplus_4^{\text{odd}}\|_\infty < \varepsilon$. \square

Theorem 48. $\mathcal{SD} = \mathcal{I}_{\text{anti}}$.

Proof. Since every function in $\text{AntiFerrolsing} \cup \mathcal{B}_0$ is self-dual by definition and \mathcal{SD} is an ω -clone by Theorem 37, we have $\mathcal{I}_{\text{anti}} \subseteq \mathcal{SD}$.

We now show $\mathcal{SD} \subseteq \mathcal{I}_{\text{anti}}$. Lemmas 45 and 46 (together with Lemma 1) imply $\mathcal{SD} \subseteq \langle \text{AntiFerrolsing} \cup \mathcal{B}'_0 \cup \{\oplus_4^{\text{odd}}\} \rangle_\omega$. From this and Lemma 47 (together with Lemma 1) we have $\mathcal{SD} \subseteq \langle \text{AntiFerrolsing} \cup \mathcal{B}'_0 \rangle_\omega$. \square

7.2 Maximality

Lemma 49. \mathcal{SD} is a maximal ω -clone in \mathcal{B} ; i.e., for any function $f \notin \mathcal{SD}$, $\langle \mathcal{SD} \cup \{f\} \rangle_\omega = \mathcal{B}$.

Proof. Let k be the arity of f . First, we show that $\langle \mathcal{SD} \cup \{f\} \rangle_\omega$ contains $\delta_0 = [1, 0]$ or $\delta_1 = [0, 1]$. If $f(\mathbf{1}) > f(\mathbf{0})$, we have $f(\mathbf{1}) > 0$ so the nullary function $f_1 = 1/f(\mathbf{1})$ is well-defined, and it is in $\langle \mathcal{SD} \cup \{f\} \rangle_\omega$ since $\mathcal{B}_0 \subseteq \mathcal{SD}$ by Theorem 37. In this case,

$$\delta_1(x) = \lim_{n \rightarrow \infty} \sum_{x_2, \dots, x_k} f(x_1, x_2, \dots, x_k)^n \left(\prod_{i=1}^{k-1} \text{EQ}(x_i, x_{i+1}) \right) \left(\frac{1}{f(\mathbf{1})} \right)^n,$$

so $\delta_1(x) \in \langle \mathcal{SD} \cup \{f\} \rangle_\omega$. If $f(\mathbf{1}) < f(\mathbf{0})$, we similarly show $\delta_0 \in \langle \mathcal{SD} \cup \{f\} \rangle_\omega$.

If $f(\mathbf{0}) = f(\mathbf{1})$ there is some $\mathbf{a} \in \{0, 1\}^k$ such that $f(\mathbf{a}) \neq f(\bar{\mathbf{a}})$ so $k \geq 2$. Because ω -clones are closed under permuting arguments, we may assume without loss of generality that $a_1 = \dots = a_\ell = 0$ and $a_{\ell+1} = \dots = a_k = 1$. Let

$$g(x_1, x_k) = \sum_{x_2, \dots, x_{k-1}} f(x_1, \dots, x_k) \left(\prod_{i=1}^{\ell-1} \text{EQ}(x_i, x_{i+1}) \right) \left(\prod_{i=\ell+1}^{k-1} \text{EQ}(x_i, x_{i+1}) \right).$$

Clearly, $g \in \langle f \rangle$. The function g satisfies $g(0, 0) = g(1, 1)$ and $g(0, 1) \neq g(1, 0)$. Set $h(x) = g(x, 0) + g(x, 1)$. We have $h(0) = g(0, 0) + g(0, 1) \neq g(1, 0) + g(1, 1) = h(1)$ and $h \in \langle f \rangle$. An argument similar to the one in the first paragraph of this proof shows that δ_0 or δ_1 is in $\langle \mathcal{SD} \cup \{h\} \rangle_\omega$. By Lemma 1, δ_0 or δ_1 is in $\langle \mathcal{SD} \cup \{f\} \rangle_\omega$.

For the rest of the proof, suppose that $\delta_0 \in \langle \mathcal{SD} \cup \{f\} \rangle_\omega$; the case where $\delta_1 \in \langle \mathcal{SD} \cup \{f\} \rangle_\omega$ is very similar. Take any $g \in \mathcal{B}$, say of arity k . Let h be the $(k+1)$ -ary function defined as follows: for any $\mathbf{x} \in \{0, 1\}^k$, $h(\mathbf{x}0) = h(\bar{\mathbf{x}}1) = g(\mathbf{x})$. As is easily seen, $h \in \mathcal{SD}$. It is also easy to see that $g(\mathbf{x}) = \sum_y h(\mathbf{x}y) \delta_0(y)$. Thus, by Lemma 1, $g \in \langle \mathcal{SD} \cup \{f\} \rangle_\omega$. \square

Next, we prove that \mathcal{SDP} is a maximal ω -clone in \mathcal{SD} . Recall the functions Par_k^λ from Definition 17.

Theorem 50. \mathcal{SDP} is a maximal ω -clone in \mathcal{SD} ; i.e., for any n -ary self-dual function $f \in \mathcal{SD} \setminus \mathcal{SDP}$, $\langle \mathcal{SDP} \cup \{f\} \rangle_\omega = \mathcal{SD}$.

Proof. Since $\mathcal{B}_0 \subseteq \mathcal{SDP}$ by Theorems 37 and 28, it suffices by Theorem 48, Corollary 22 and Lemma 1, to show that $I_2^\lambda \in \langle \mathcal{SDP} \cup \{f\} \rangle_\omega$ for some $\lambda > 1$. Since $f \in \mathcal{SD} \setminus \mathcal{SDP}$, there is some $\mathbf{a} \in \{0, 1\}^n$ such that $\widehat{f}(\mathbf{a}) < 0$. By Lemma 31, we may assume that $\mathbf{a} = \mathbf{1}$. Then, by Lemma 33, there is a permissive binary function $h \in \langle \{f, \text{Par}_{n+2}^{1/2}\} \rangle$ such that $\widehat{h}(1, 1) < 0$.

Because the n -ary function f is self-dual and $\widehat{f}(\mathbf{1}) \neq 0$, n must be even by Lemma 38. For all even n , $\text{Par}_{n+2}^{1/2}$ is self-dual. Also, $\text{Par}_{n+2}^{1/2} \in \mathcal{P}$ by Lemma 27, so it is in \mathcal{SDP} and, by Lemma 1, $h \in \mathcal{SDP}$.

Since $\widehat{h}(1, 1) < 0$ and h is permissive, there are constants $c > b > 0$ such that $h(0, 0) = h(1, 1) = b$ and $h(0, 1) = h(1, 0) = c$. Therefore, the function $(1/b)h$ is $I_2^{c/b}$, with $c/b > 1$, and this function belongs to $\langle \mathcal{SDP} \cup \{f\} \rangle_\omega$. \square




Figure 3: Match-circuits used in the proof of Theorem 51. Every edge has weight 1 unless otherwise indicated. a) The equality function. b) The product of functions implemented by the match-circuits F and G .

8 Match-circuits and even-circuits

8.1 Match-circuits

We first show that \mathcal{M} and \mathcal{E} are functional clones. It is still open whether they are ω -clones. As far as we know, there may be a function in $\langle \mathcal{M} \rangle_\omega$ that is the limit of a sequence of functions f_ε where each f_ε is implemented by a match-circuit with its own underlying graph. It is not clear in this case whether f itself can be implemented by a match-circuit. A similar comment applies to \mathcal{E} .

Theorem 51. $\langle \mathcal{M} \rangle = \mathcal{M}$ and $\mathcal{B}_0 \subseteq \mathcal{M}$.

Proof. We show that \mathcal{M} contains the equality function and has all the closure properties required by the definition of functional clone.

- For the equality function, we have $\widehat{\text{EQ}} = \frac{1}{2}\text{EQ}$. This function is implemented by the graph shown in Figure 3(a).
- Permuting arguments corresponds directly to renaming the terminals of the circuit, so it is clear that \mathcal{M} is closed under this operation.
- For closure under the introduction of fictitious arguments, let $g(\mathbf{x}z) = f(\mathbf{x})$ for some k -ary $f \in \mathcal{M}$. By Lemma 25(ii), $\widehat{g}(\mathbf{x}0) = \widehat{f}(\mathbf{x})$ and $\widehat{g}(\mathbf{x}1) = 0$. The match-circuit for \widehat{g} is the disjoint union of the match-circuit F for \widehat{f} and a weight-1 path on new vertices u_{k+1} , v and v' (in that order). If y_{k+1} is assigned 0, then any perfect matching is the union of the edge (v, v') and a perfect matching of F , so has weight $\widehat{f}(y_1, \dots, y_k)$; if y_{k+1} is assigned 1, there is no perfect matching, so the assignment has weight 0, as required.
- For closure under summation, let $g(\mathbf{x}) = \sum_z f(\mathbf{x}z)$ for some $(k+1)$ -ary function $f \in \mathcal{M}$. By Lemma 25(iii), $\widehat{g}(\mathbf{x}) = 2\widehat{f}(\mathbf{x}0)$, so we obtain a match-circuit for \widehat{g} from the circuit F for \widehat{f} by deleting the vertex u_{k+1} (which is equivalent to forcing its adjacent edge in F

to be spin-0) and adding a new weight-2 edge between two new vertices (which doubles the weight of any perfect matching).

- For closure under products, let $h(\mathbf{x}) = f(\mathbf{x})g(\mathbf{x})$ for k -ary functions $f, g \in \mathcal{M}$. Let \widehat{f} and \widehat{g} be implemented by match-circuits F and G , with terminal vertices $u_{f,1}, \dots, u_{f,k}$ and $u_{g,1}, \dots, u_{g,k}$, respectively. For each $i \in [k]$, let $y_{f,i}$ be the unique edge adjacent to $u_{f,i}$ in F and define $y_{g,i}$ similarly in G . Recall that $\widehat{h}(\mathbf{y}) = \sum_{\mathbf{w} \in \{0,1\}^k} \widehat{f}(\mathbf{w}) \widehat{g}(\mathbf{w} \oplus \mathbf{y})$.

Let σ be any assignment of spins 0 and 1 to the edges of the match-circuit H shown in Figure 3(b). We claim that, if σ is a perfect matching then, for all $i \in [k]$, $\sigma(y_{f,i}) = \sigma(y_{g,i}) \oplus \sigma(y_i)$.

If $\sigma(y_i) = 0$ then we must have $\sigma(v_i, v_{k+i}) = 1$. We may have $\sigma(y_{f,i}) = 0$ or $\sigma(y_{f,i}) = 1$ but, in either case, $\sigma(y_{f,i}) = \sigma(y_{g,i}) = \sigma(y_{g,i}) \oplus 0$.

Otherwise, $\sigma(y_i) = 1$ and we must have $\sigma(v_i, v_{k+i}) = 0$. Now there are two cases. If $\sigma(v_{k+i}, v_{f,i}) = 1$, then $\sigma(y_{f,i}) = 0 = \sigma(y_{g,i}) \oplus 1$; if $\sigma(v_{k+i}, v_{g,i}) = 1$, then $\sigma(y_{f,i}) = 1 = \sigma(y_{g,i}) \oplus 1$. This completes the proof of the claim.


For any choice of spins x_1, \dots, x_k for the edges y_1, \dots, y_k we can choose any spins w_1, \dots, w_k for the edges $y_{f,1}, \dots, y_{f,k}$. Doing so forces us to assign the spin $w_i \oplus x_i$ to $y_{g,i}$. Therefore, the value computed by the match-circuit is $\sum_{\mathbf{w}} \widehat{f}(\mathbf{w}) \widehat{g}(\mathbf{w} \oplus \mathbf{x})$, as required.

The fact that $\mathcal{B}_0 \subseteq \mathcal{M}$ comes from the definition of match-circuit. Any positive $c \in \mathcal{B}_0$ can be implemented by a match-circuit with no terminals containing one edge with weight c . The constant 0 is implemented by a match-circuit with no terminals whose three edges form a 3-cycle. \square

Theorem 52. $\mathcal{I}_{\text{ferro}} \subseteq \langle \mathcal{M} \rangle_{\omega} \cap \mathcal{H}_{\text{ferro}}$.

Proof. It is trivial that $\mathcal{I}_{\text{ferro}} \subseteq \mathcal{H}_{\text{ferro}}$, so it remains to prove that $\mathcal{I}_{\text{ferro}} \subseteq \langle \mathcal{M} \rangle_{\omega}$.

By Corollary 24, it suffices to show that $f = I_2^{1/2} \in \mathcal{M}$. By Lemma 26, we have $\widehat{f} = [\frac{3}{4}, 0, \frac{1}{4}]$. This is implemented by the following match-circuit, with terminals (u_1, v_1) and (u_2, v_2) .



\square

We do not know whether the inclusion in the statement of Theorem 52 is strict. This corresponds to the dotted line in Figure 1.

8.2 Even-circuits

The proof that \mathcal{E} is a functional clone is similar to Theorem 51.

Theorem 53. $\langle \mathcal{E} \rangle = \mathcal{E}$ and $\mathcal{B}_0 \subseteq \mathcal{E}$.

Proof. We show that \mathcal{E} contains the equality function and has all the closure properties required by the definition of functional clone.

- For the equality function, we have $\widehat{\text{EQ}} = \frac{1}{2}\text{EQ}$. The definition of \mathcal{E} accounts for the multiplication by $1/2$, so we need only show that EQ is implemented by an even-circuit. Indeed, it is implemented by a three-edge path between two terminals (where all edges have weight 1).
- Permuting arguments corresponds directly to renaming the terminals of the circuit, so it is clear that \mathcal{E} is closed under this operation.
- For closure under the introduction of fictitious arguments, let $g(\mathbf{x}z) = f(\mathbf{x})$ for some k -ary $f \in \mathcal{E}$. By Lemma 25(ii), $\widehat{g}(\mathbf{x}0) = \widehat{f}(\mathbf{x})$ and $\widehat{g}(\mathbf{x}1) = 0$. The even-circuit for \widehat{g} is the disjoint union of the even-circuit F for \widehat{f} and a weight-1 edge on new vertices u_{k+1} and v_{k+1} . Even subgraphs with $y_{k+1} = 0$ correspond to even subgraphs of F . There are no even subgraphs with $y_{k+1} = 1$.
- For closure under summation, let $g(\mathbf{x}) = \sum_z f(\mathbf{x}z)$ for some $(k+1)$ -ary function $f \in \mathcal{E}$. By Lemma 25(iii), $\widehat{g}(\mathbf{x}) = 2\widehat{f}(\mathbf{x}0)$, so we obtain an even-circuit for $\widehat{g}/2$ from the circuit F for \widehat{f} by deleting the vertex u_{k+1} (which is equivalent to forcing its adjacent edge in F to be spin-0).
- For closure under products, let $h(\mathbf{x}) = f(\mathbf{x})g(\mathbf{x})$ for k -ary functions $f, g \in \mathcal{E}$. Let \widehat{f} and \widehat{g} be implemented by even-circuits F and G , with terminal vertices $u_{f,1}, \dots, u_{f,k}$ and $u_{g,1}, \dots, u_{g,k}$, respectively. For each $i \in [k]$, let $y_{f,i}$ be the unique edge adjacent to $u_{f,i}$ in F and define $y_{g,i}$ similarly in G . Recall that, by (1), $\widehat{h}(\mathbf{y}) = \sum_{\mathbf{w} \in \{0,1\}^k} \widehat{f}(\mathbf{w})\widehat{g}(\mathbf{w} \oplus \mathbf{y})$. Let H be the even-circuit that is the same as the one shown in Figure 3(b) except that the edges $(u_{f,j}, u_{g,j})$ are deleted. Let σ be any assignment of spins 0 and 1 to the edges of H . We claim that, if σ is an even subgraph then, for all $i \in [k]$, $\sigma(y_{f,i}) = \sigma(y_{g,i}) \oplus \sigma(y_i)$. If $\sigma(y_i) = 1$ then we must have $\sigma(v_i, v_{k+i}) = 1$. Thus, exactly one of the edges $(v_{k+i}, u_{f,i})$ and $(v_{k+i}, u_{g,i})$ has spin one. So $\sigma(y_{f,i})$ and $\sigma(y_{g,i})$ differ. Otherwise, $\sigma(y_i) = 0$ so $\sigma(v_i, v_{i+i}) = 0$ so $\sigma(y_{f,i})$ and $\sigma(y_{g,i})$ agree. This completes the proof of the claim.
For any choice of spins x_1, \dots, x_k for the edges y_1, \dots, y_k we can choose any spins w_1, \dots, w_k for the edges $y_{f,1}, \dots, y_{f,k}$. Doing so forces us to assign the spin $w_i \oplus x_i$ to $y_{g,i}$. Therefore, the value computed by the even-circuit is $\sum_{\mathbf{w}} \widehat{f}(\mathbf{w})\widehat{g}(\mathbf{w} \oplus \mathbf{x})$, as required.

The fact that $\mathcal{B}_0 \subseteq \mathcal{E}$ comes from the definition of even-circuit. The nullary zero function $f = 0$ is in \mathcal{E} , since $f = 0 \cdot g$ for any function g implemented by an even-circuit. Any non-zero nullary function $f = c$ is in \mathcal{E} , since $f = c \cdot 1$, where 1 is the constant one function implemented by the empty graph, whose unique even subgraph is the empty graph, which has weight 1. \square

The next theorem shows that the functional clone \mathcal{E} is the same as the clone generated by nullary functions and ferromagnetic Ising model interactions. Something very close to this equivalence is seen in the “high-temperature expansion” of the Ising model, first elucidated by Van der Waerden [21]. In our proof, we employ the framework of holants and holographic transformations. See Cai, Lu and Xia [4] for the wider context, particularly the introduction to that paper and Theorem IV.1.

Theorem 54. $\langle \text{Ferrolsing} \cup \mathcal{B}_0 \rangle = \mathcal{E}$.

Proof. Theorem 53 shows that $\mathcal{B}_0 \subseteq \mathcal{E}$. It is also true that $I_2^1 \in \mathcal{E}$ since it can be constructed from \mathcal{B}_0 by introducing fictitious arguments. To see that $\text{Ferrolsing} \subseteq \mathcal{E}$ consider any function I_2^λ with $0 \leq \lambda < 1$. Let $f(x_1, x_2) = (2/(1 + \lambda))I_2^\lambda$. By Lemma 26, $\widehat{f}(0, 0) = 1$, $\widehat{f}(0, 1) = \widehat{f}(1, 0) = 0$ and $\widehat{f}(1, 1) = (1 - \lambda)/(1 + \lambda)$. But \widehat{f} can be implemented by an even-circuit consisting of a three-edge path between two terminals in which the middle edge has weight $(1 - \lambda)/(1 + \lambda)$. Thus, $I_2^\lambda \in \langle \mathcal{E} \rangle$. By Lemma 1, $\langle \text{Ferrolsing} \cup \mathcal{B}_0 \rangle \subseteq \langle \mathcal{E} \rangle$, so by Theorem 53, $\langle \text{Ferrolsing} \cup \mathcal{B}_0 \rangle \subseteq \mathcal{E}$.

We now show that $\mathcal{E} \subseteq \langle \text{Ferrolsing} \cup \mathcal{B}_0 \rangle$. It is obvious that any nullary function is in $\langle \text{Ferrolsing} \cup \mathcal{B}_0 \rangle$ so consider any function $g \in \mathcal{E}$ with arity $k \geq 1$. From the definition of \mathcal{E} , $g = c \cdot f$ for some non-negative real number c and \widehat{f} is implemented by an even-circuit G .

We will show that f is in $\langle \text{Ferrolsing} \cup \mathcal{B}_0 \rangle$, which implies that g is also in $\langle \text{Ferrolsing} \cup \mathcal{B}_0 \rangle$. To do this, we start by viewing the even-circuit G as an instance of a holant problem. A holant problem [4] consists of a graph in which, for all d , every degree- d vertex v is equipped with a function $f_v \in \mathcal{B}_d$. A configuration assigns spins 0 and 1 to the edges, and the weight of a configuration is the product, over all vertices v , of $f_v(\mathbf{x})$, where \mathbf{x} is the string of spins of edges around v (in some appropriate order). The partition function is the sum of the weights of the configurations.

To represent the relevant holant problem cleanly, we first construct G' from G by two-stretching the internal edges of G (turning them all into two-edge paths). That is, if G contains an edge $e = (v_i, v_j)$, we add a new vertex v_e to G' . We view G' has a holant problem, so configurations assign spins 0 and 1 to the edges of G' . At each vertex v_i of G' we add a function f_{v_i} which is 1 if an even number of its arguments have spin-1 and is 0 otherwise. At each new vertex v_e of G' we add a function f_{v_e} which is the symmetric arity-2 function $[1, 0, w_e]$. This ensures that, in configurations with non-zero weight, the two edges adjacent to the new vertex v_e get the same spin (so non-zero configurations of G' correspond to even subgraphs of G). It also ensures also that the weight w_e of the edge e of the even-circuit G is accounted for. It is easy to see that the partition function of the holant problem G' is the same as the function implemented by the even-circuit G , which is \widehat{f} .

Now we apply a standard trick from the holant literature. Let $H = \frac{1}{2}[1, 1, -1]$ be the symmetric arity-2 Hadamard/FFT function. Construct a new holant instance G'' from G' by three-stretching every edge of G' and equipping every new vertex with the function H . Since $H = \frac{1}{2}H^{-1}$, the partition function of G'' is $2^{-|E(G')|}$ times the partition function of G' , so it is, upto a constant factor, \widehat{f} .

Now construct a new holant problem G''' from G'' by considering all of the original vertices of G' .

- For any arity- d vertex v_i (which is an original internal vertex of G), replace the subgraph consisting of v_i (with its “arity- d even parity” function) and all of its neighbours (which have H functions) with an equivalent arity- d vertex equipped with the arity- d equality function. This leaves the partition function unchanged.
- For any vertex v_e (one of the new nodes with $[1, 0, w_e]$ functions added in the construction of G') let $\lambda_e = (1 - w_e)/(1 + w_e)$ and replace v_e together with its two neighbours (which have H functions) with the equivalent degree-2 vertex whose function is $\frac{1}{4}(1 + w_e)[1, \lambda_e, 1]$. Again, this does not change the partition function.
- The only remaining vertices with H functions are adjacent to the external vertices of G . Replacing these functions with arity-2 equality, we obtain a holant problem G''' whose

partition function is the Fourier transform of that of G'' . Thus, its partition function is f .

We now have a holant problem G''' implementing f , up to a constant factor. All of the functions at the vertices of G''' are equality (of any arity) or $\frac{1}{4}(1 + w_e)[1, \lambda_e, 1]$ for some $0 \leq \lambda_e < 1$ so they are all in $\langle \text{Ferrolsing} \cup \mathcal{B}_0 \rangle$. Moreover, the equality constraints correspond to the original internal vertices v_i of G and the new Ising constraints correspond to edges between internal vertices of G . Thus, G''' implements a sum (over the spins of the internal vertices of G) of a product (over the spins of the internal edges of G) of constraints in $\langle \text{Ferrolsing} \cup \mathcal{B}_0 \rangle$. This shows that f is in the closure of $\langle \text{Ferrolsing} \cup \mathcal{B}_0 \rangle$ under product and summation, so f itself is in $\langle \text{Ferrolsing} \cup \mathcal{B}_0 \rangle$. \square

Recall that $\mathcal{I}_{\text{ferro}} = \langle \text{Ferrolsing} \cup \mathcal{B}_0 \rangle_\omega$. The following corollary follows immediately from Theorem 54 and Theorem 53 using the definition of an ω -clone.

Corollary 55. $\mathcal{I}_{\text{ferro}} = \langle \mathcal{E} \rangle_\omega$.

8.3 Relationship of $\langle \mathcal{M} \rangle_\omega$ with other clones

In this section, we give, in Lemma 56, a necessary condition for a 4-ary function to be in \mathcal{M} . Moreover, we show, in Lemma 57, that for symmetric functions this condition is also sufficient. We then use these results to study the relationship between $\langle \mathcal{M} \rangle_\omega$ and the clones around it in the lattice \mathcal{L}' .

Lemma 56. *For every 4-ary function $f \in \mathcal{M}$,*

$$\widehat{f}(0011) \widehat{f}(1100) + \widehat{f}(0101) \widehat{f}(1010) + \widehat{f}(0110) \widehat{f}(1001) \geq \widehat{f}(0000) \widehat{f}(1111). \quad (2)$$

Proof. Consider an arity-4 match-circuit G that implements \widehat{f} as described in Definition 5. Let $S = \{u_1, u_2, u_3, u_4\}$ be the set of external vertices of G . For $A \subseteq S$, let M_A denote the set of perfect matchings which include terminals adjacent to A (by assigning them spin 1) and exclude terminals adjacent to $S \setminus A$ (by assigning them spin 0). We exhibit an injective map

$$\nu: M_\emptyset \times M_S \rightarrow M_{\{u_1, u_2\}} \times M_{\{u_3, u_4\}} \cup M_{\{u_1, u_3\}} \times M_{\{u_2, u_4\}} \cup M_{\{u_1, u_4\}} \times M_{\{u_2, u_3\}}$$

which is weight-preserving in the sense that, for matchings m_1, \dots, m_4 with $\nu(m_1, m_2) = (m_3, m_4)$, we have $w(m_1)w(m_2) = w(m_3)w(m_4)$. The existence of ν implies (2).

Given $(m_1, m_2) \in M_\emptyset \times M_S$, consider $m_1 \oplus m_2$ and note that this is a collection of cycles together with two paths π and π' . Let π_1 be the path connecting vertex u_1 to one of the other external vertices; the other path connects the remaining external vertices. If π joins u_1 to u_2 , then $m_3 := m_1 \oplus \pi \in M_{\{u_1, u_2\}}$ and $m_4 := m_2 \oplus \pi \in M_{\{u_3, u_4\}}$, with similar claims for π joining u_1 to u_3 or u_4 . The construction is invertible, since $m_3 \oplus m_4 = m_1 \oplus m_2$, from which we can recover π and, hence, m_1 and m_2 . Therefore, $\nu: (m_1, m_2) \mapsto (m_3, m_4)$ is an injection as claimed.

To see that ν is weight-preserving, observe that the edges of π each appear in exactly one of m_1 and m_2 and in exactly one of m_3 and m_4 and that, for $i \in \{1, 2\}$, $m_i \setminus \pi = m_{i+2} \setminus \pi$.

$$w(m_1)w(m_2) = \prod_{e \in m_1 \setminus \pi} w_e \prod_{e \in m_2 \setminus \pi} w_e \prod_{e \in \pi} w_e = \prod_{e \in m_3 \setminus \pi} w_e \prod_{e \in m_4 \setminus \pi} w_e \prod_{e \in \pi} w_e = w(m_3)w(m_4). \quad \square$$




Figure 4: Match-circuits used in the proof of Lemma 57. Every edge has weight 1 unless otherwise indicated. a) The case $C_0 > 0$, $C_2 = C_4 = 0$. b) The case $C_0, C_2 > 0$, $C_4 \geq 0$.

We give the converse of Lemma 56 for symmetric functions.

Lemma 57. *If f is a symmetric, arity-4, self-dual function such that*

$$3\widehat{f}(0011)^2 \geq \widehat{f}(0000) \widehat{f}(1111), \quad (3)$$

then $f \in \mathcal{M}$.

Note that (3) is just (2) specialised to symmetric functions.

Proof. For ease of notation, let $C_0 = \widehat{f}(0,0,0,0)$, $C_2 = \widehat{f}(0,0,1,1)$ and $C_4 = \widehat{f}(1,1,1,1)$. Since $C_0 = \frac{1}{16} \sum_{\mathbf{z} \in \{0,1\}^4} f(\mathbf{z})$ is a sum of nonnegative terms, if $C_0 = 0$, then f is the constant zero function, which is in \mathcal{M} by Theorem 51. For the rest of the proof, we assume that $C_0 > 0$.

If $C_2 = C_4 = 0$, then \widehat{f} is implemented by the match-circuit shown in Figure 4(a). If at most one of C_2 and C_4 is zero, then (3) implies that $C_2 > 0$.

We will construct a match-circuit G for \widehat{f} (see Figure 4(b)). In addition to the terminal edges $y_i = (u_i, v_i)$ for $i \in [4]$, G will have edges (v_i, w_i) and (w_i, x_i) . It will also contain a clique on the six vertices $x_1, x_2, x_3, x_4, x_5, x_6$. The edge (x_5, x_6) has weight λ and the edges (w_i, x_i) have weight μ . All other edges have weight 1.

Let $S = \{u_1, u_2, u_3, u_4\}$. Following the proof of Lemma 56, for $A \subseteq S$, M_A denotes the set of perfect matchings which include terminals adjacent to A (by assigning them spin 1) and exclude terminals adjacent to $S \setminus A$ (by assigning them spin 0). Let Z_A denote the sum of the weights of the perfect matchings in M_A .

The perfect matchings in M_\emptyset contain all of the edges (v_i, w_i) and none of the edges (u_i, v_i) or (w_i, x_i) . There are three perfect matchings of the clique that include the weight- λ edge (x_5, x_6) and twelve perfect matchings of the clique that do not include this edge, so $Z_\emptyset = 3\lambda + 12$. Similarly, the single perfect matching in M_S contains all of the edges (u_i, v_i) and (w_i, x_i) (which have weight μ) and none of the edges (v_i, w_i) . The edge (x_5, x_6) is present, so $Z_S = \lambda\mu^4$. Finally, the perfect matchings in $M_{\{u_1, u_2\}}$ contain the two weight- μ edges (w_1, x_1) and (w_2, x_2) but not the two weight- μ edges (w_3, x_3) and (w_4, x_4) . There are three matchings of the 4-clique containing x_3, x_4, x_5, x_6 , one of which has weight λ , so $Z_{\{u_1, u_2\}} = (\lambda + 2)\mu^2$

and the same is true for Z_A for any other size-two set $A \subseteq S$. Now let

$$z(\lambda) = \frac{Z_\emptyset Z_S}{Z_{\{u_1, u_2\}} Z_{\{u_1, u_2\}}}.$$

Note that $z(\lambda) = \lambda(3\lambda + 12)/(\lambda + 2)^2$, and that the range of $z(\lambda)$ includes the interval $[0, 3)$. There are now three cases.

If $3C_2^2 > C_0C_4 > 0$, we can choose λ so that $z(\lambda) = C_0C_4/C_2^2$. Now choose μ to obtain $(Z_\emptyset, Z_{\{u_1, u_2\}}, Z_S) \propto (C_0, C_2, C_4)$. In order to get the constant multiple correct, G can be supplemented with an additional edge.

If $3C_2^2 > C_0C_4 = 0$ (so $C_4 = 0$, since C_0 and C_2 are positive), we can simulate $\lambda = 0$ by deleting the edge (x_5, x_6) . We have $Z_\emptyset = 12$ and $Z_{\{u_1, u_2\}} = 2\mu^2$ and, as before, we can choose μ so that $(Z_\emptyset, Z_{\{u_1, u_2\}}) \propto (C_0, C_2)$ and add an edge to G for the required constant multiple.

Finally, if $3C_2^2 = C_0C_4 > 0$, must achieve $z(\lambda) = 3$. This can be done by effectively setting $\lambda = \infty$ by removing the vertices x_5 and x_6 and their incident edges. \square

Theorem 58. $\langle \mathcal{M} \rangle_\omega \subset \mathcal{SDP}$.

Proof. Let G be a match-circuit with terminals y_1, \dots, y_k , where $y_i = (u_i, v_i)$ for each $i \in [k]$, which implements the function $\widehat{f}(\mathbf{y})$. For any assignment \mathbf{a} to \mathbf{y} , $\widehat{f}(\mathbf{a})$ is the total weight of the perfect matchings of the graph $G - \{v_i \mid a_i = 1\}$. Since a graph with an odd number of vertices has no perfect matchings, $\widehat{f}(\mathbf{a}) = 0$ whenever $|\mathbf{a}|$ is odd or $\widehat{f}(\mathbf{a}) = 0$ whenever $|\mathbf{a}|$ is even. (This is the so-called ‘‘parity condition’’ of match-circuits; see, e.g., [3].)

We first show that $\mathcal{M} \subseteq \mathcal{SDP}$. Suppose that $f \in \mathcal{M}$. We will show that $f \in \mathcal{SDP}$. Since Theorems 37 and 28 guarantee that $\mathcal{B}_0 \in \mathcal{SDP}$, we can assume without loss of generality that the arity, k , of f is positive. Since \widehat{f} is implemented by a match-circuit, $\widehat{f}(\mathbf{x}) \geq 0$ for all \mathbf{x} , so $f \in \mathcal{P}$. If f is the constant arity- k zero function, then $f \in \mathcal{SDP}$ trivially, so assume that $f(\mathbf{a}) > 0$ for at least one $\mathbf{a} \in \{0, 1\}^k$. This implies that $\widehat{f}(\mathbf{0}) = 2^{-k} \sum_{\mathbf{x}} f(\mathbf{x}) > 0$ so, by the parity condition, $\widehat{f}(\mathbf{x}) = 0$ whenever $|\mathbf{x}|$ is odd. By Lemma 38, $f \in \mathcal{SD}$, and we have established that $f \in \mathcal{SD} \cap \mathcal{P}$. But $\mathcal{SDP} = \mathcal{SD} \cap \mathcal{P}$ by Theorem 39, so $f \in \mathcal{SDP}$.

To show that $\mathcal{M} \subset \mathcal{SDP}$, consider the function $g = I_4^{1/4} = [1, \frac{1}{4}, \frac{1}{4}, \frac{1}{4}, 1]$, which is in \mathcal{SD} . By Lemma 26, $\widehat{g}(0000) = \frac{11}{32}$, $\widehat{g}(\mathbf{x}) = \frac{3}{32}$ when $|\mathbf{x}| \in \{2, 4\}$ and $\widehat{g}(\mathbf{x}) = 0$, otherwise, so $g \in \mathcal{P} \cap \mathcal{PN} = \mathcal{SDP}$. However, $g \notin \mathcal{M}$ by Lemma 56.

It remains to ‘‘lift’’ the result to ω -clones. We have $\langle \mathcal{M} \rangle_\omega \subseteq \langle \mathcal{SDP} \rangle_\omega = \mathcal{SDP}$, so it is enough to show that $g = I_4^{1/4} \notin \langle \mathcal{M} \rangle_\omega$. Suppose, towards a contradiction, that $g \in \langle \mathcal{M} \rangle_\omega$. By the definition of ω -clones, for every $\varepsilon > 0$, there is a function $f \in \langle \mathcal{M} \rangle$ such that $\|f - g\|_\infty < \varepsilon/32$. Since \mathcal{M} is a functional clone by Theorem 51, $f \in \mathcal{M}$. Then Lemma 25(v) implies that $\|\widehat{f} - \widehat{g}\|_\infty < \varepsilon/32$, also. Thus, for all $\mathbf{x} \in \{0, 1\}^4$,

$$\widehat{g}(\mathbf{x}) - \varepsilon/32 < \widehat{f}(\mathbf{x}) < \widehat{g}(\mathbf{x}) + \varepsilon/32. \quad (4)$$

Since $f \in \mathcal{M}$, (2) must hold. Plugging (4) and the values of \widehat{g} into (2) gives

$$3\left(\frac{3}{32} + \frac{\varepsilon}{32}\right)^2 > \left(\frac{11}{32} - \frac{\varepsilon}{32}\right)\left(\frac{3}{32} - \frac{\varepsilon}{32}\right),$$

But this only holds for sufficiently large positive values of ε , contradicting the assumption that $g \in \langle \mathcal{M} \rangle_\omega$. \square

Lemmas 56 and 57 also allow us to separate $\mathcal{H}_{\text{ferro}}$ from $\langle \mathcal{M} \rangle_\omega$.

Lemma 59. $\langle \mathcal{M} \rangle_\omega$ and $\mathcal{H}_{\text{ferro}}$ are incomparable under \subseteq .

Proof. Let $f = [13, 4, 1, 4, 13]$. We saw in the proof of Theorem 44 that $f \notin \mathcal{H}_{\text{ferro}}$ and that $\widehat{f} = [4, 0, \frac{3}{2}, 0, 0]$. However, $f \in \langle \mathcal{M} \rangle_\omega$ by Lemma 57.

Now, let $g = I_4^{1/2} \in \mathcal{H}_{\text{ferro}}$. By Lemma 26, $\widehat{g} = [\frac{9}{16}, 0, \frac{1}{16}, 0, \frac{1}{16}]$ and $g \notin \langle \mathcal{M} \rangle_\omega$ by Lemma 56, since $3(\frac{1}{16})^2 < \frac{9}{16} \cdot \frac{1}{16}$. \square

9 The lattice \mathcal{L}'

In this section, we prove Theorem 10.

Theorem 10. *The lattice \mathcal{L}' shown in Figure 1 is a sublattice of \mathcal{L}_ω . That is, all elements of \mathcal{L}' are distinct ω -clones, with the possible exceptions of \mathcal{SDP} and $\langle \langle \mathcal{M} \rangle_\omega \cup \mathcal{H}_{\text{ferro}} \rangle_\omega$, and $\langle \mathcal{M} \rangle_\omega \cap \mathcal{H}_{\text{ferro}}$ and $\mathcal{I}_{\text{ferro}}$, which might be equal. (This is indicated by the dotted lines in Figure 1.) Furthermore, the meets and joins of elements of \mathcal{L}' are as depicted in Figure 1 and*

(i) $\mathcal{SD} = \mathcal{I}_{\text{anti}}$;

(ii) $\mathcal{I}_{\text{ferro}} = \langle \mathcal{E} \rangle_\omega$;

(iii) \mathcal{SD} , \mathcal{P} and \mathcal{PN} are maximal in \mathcal{B} ;

(iv) \mathcal{SDP} is maximal in \mathcal{SD} .

Proof. First, we check that the vertices of \mathcal{L}' are, indeed, ω -clones. \mathcal{B} is trivially an ω -clone. $\langle \mathcal{M} \rangle_\omega$, $\mathcal{H}_{\text{ferro}}$, $\langle \langle \mathcal{M} \rangle_\omega \cup \mathcal{H}_{\text{ferro}} \rangle_\omega$ and $\mathcal{I}_{\text{ferro}}$, are ω -clones by definition. Hence, so is the intersection $\langle \mathcal{M} \rangle_\omega \cap \mathcal{H}_{\text{ferro}}$. \mathcal{P} , \mathcal{PN} and \mathcal{SD} are ω -clones by Theorems 28, 30 and 37, so their intersection \mathcal{SDP} is also an ω -clone.

Next, we check the lattice structure.

We start with the strict inclusions (indicated by the solid lines in Figure 1). It is easy to see, using the definitions, that \mathcal{SD} , \mathcal{P} and \mathcal{PN} are strict subsets of \mathcal{B} . \mathcal{SD} , \mathcal{P} and \mathcal{PN} are pairwise-incomparable under \subseteq by Lemma 41. \mathcal{SDP} is a subset of \mathcal{SD} , \mathcal{P} and \mathcal{PN} by definition; it is a strict subset because \mathcal{SD} , \mathcal{P} and \mathcal{PN} are distinct. $\langle \mathcal{M} \rangle_\omega$ and $\mathcal{H}_{\text{ferro}}$ are \subseteq -incomparable by Lemma 59. Consequently, $\langle \mathcal{M} \rangle_\omega \cap \mathcal{H}_{\text{ferro}}$ is a strict subset of both $\langle \mathcal{M} \rangle_\omega$ and $\mathcal{H}_{\text{ferro}}$ and $\langle \langle \mathcal{M} \rangle_\omega \cup \mathcal{H}_{\text{ferro}} \rangle_\omega$ is a strict superset of both $\langle \mathcal{M} \rangle_\omega$ and $\mathcal{H}_{\text{ferro}}$. For the remaining two inclusions (which we do not know to be strict, as indicated by the dotted lines in Figure 1), $\mathcal{I}_{\text{ferro}} \subseteq \langle \mathcal{M} \rangle_\omega \cap \mathcal{H}_{\text{ferro}}$ by Theorem 52. Also, $\langle \mathcal{M} \rangle_\omega \subset \mathcal{SDP}$ by Theorem 58 and $\mathcal{H}_{\text{ferro}} \subset \mathcal{SDP}$ by Theorem 44. Hence $\langle \langle \mathcal{M} \rangle_\omega \cup \mathcal{SDP} \rangle_\omega \subseteq \mathcal{SDP}$ since \mathcal{SDP} is an ω -clone.

Note that since the meet of any two clones is defined as their intersection the meets of any two ω -clones from \mathcal{L}' are indeed as shown in Figure 1 (this uses Theorem 39). We now show that also the joins of any two ω -clones from \mathcal{L}' are as shown in Figure 1. Lemma 49 implies $\mathcal{SD} \vee \mathcal{P} = \mathcal{SD} \vee \mathcal{PN} = \mathcal{B}$. By Corollary 36, $\mathcal{P} \vee \mathcal{PN} = \mathcal{B}$. $\langle \mathcal{M} \rangle_\omega \vee \mathcal{H}_{\text{ferro}} = \langle \langle \mathcal{M} \rangle_\omega \cup \mathcal{H}_{\text{ferro}} \rangle_\omega$ by definition.

$\mathcal{SD} = \mathcal{I}_{\text{anti}}$ by Theorem 48. $\mathcal{I}_{\text{ferro}} = \langle \mathcal{E} \rangle_\omega$ by Corollary 55.

Maximality of \mathcal{SD} , \mathcal{P} and \mathcal{PN} in \mathcal{B} is by Lemma 49, Lemma 34 and Corollary 35, respectively. Maximality of \mathcal{SDP} in \mathcal{SD} is by Theorem 50. \square

10 Clones of monotone functions

For $\mathbf{x}, \mathbf{y} \in \{0, 1\}^n$, $\mathbf{x} \leq \mathbf{y}$ denotes the fact that $x_i \leq y_i$ for all $i \in [n]$. For any function $f(x_1, \dots, x_n) \in \mathcal{B}_n$, let \sim_f denote the equivalence relation on $[n]$ given by $i \sim_f j$ if and only if for every $\mathbf{x} \in \{0, 1\}^n$, $f(\mathbf{x}) = 0$ whenever $x_i \neq x_j$. Let V_1, \dots, V_ℓ be the equivalence classes of \sim_f and let \tilde{f} be the function in \mathcal{B}_ℓ defined as follows. For any $\mathbf{x} \in \{0, 1\}^\ell$, construct $\mathbf{y} \in \{0, 1\}^n$ as follows. For all $i \in [n]$, if $i \in V_j$, then set $y_i = x_j$. Then $\tilde{f}(\mathbf{x}) = f(\mathbf{y})$.

Lemma 60. *For any $f \in \mathcal{B}$, $\tilde{f} \in \langle f \rangle$ and $f \in \langle \tilde{f} \rangle$.*

Proof. The function f is constructed from \tilde{f} by introducing fictitious arguments and adding EQ factors. The function \tilde{f} is constructed from f by summing out variables. \square

Definition 61. (Definition of monotone and block monotone functions.)

- Given a function $f \in \mathcal{B}_n$ and an index $i \in [n]$, the argument x_i is said to be *fictitious* in f if, for all \mathbf{x} and \mathbf{x}' that differ only at position i , $f(\mathbf{x}) = f(\mathbf{x}')$.
- For any non-negative integer n and any $\alpha \geq 0$, the function $f(x_1, \dots, x_n) \in \mathcal{B}_n$ is said to be α -*monotone* if, for every argument x_i that is not fictitious in f , and every \mathbf{x} with $x_i = 0$, $\alpha f(\mathbf{x}) \leq f(x_1, \dots, x_{i-1}, \bar{x}_i, x_{i+1}, \dots, x_n)$.
- A function $f \in \mathcal{B}$ is said to be *block- α -monotone* if \tilde{f} is α -monotone.
- 1-monotone functions are called *monotone* functions. Block-1-monotone functions are called *block-monotone* functions.
- The set of all block-monotone functions is denoted by \mathcal{MON} , and the set of all block- α -monotone functions is denoted by \mathcal{MON}_α .

Note that if $\alpha \geq \beta$ then $\alpha f(\mathbf{x}) \leq f(\mathbf{y})$ implies $\beta f(\mathbf{x}) \leq f(\mathbf{y})$. Thus, every α -monotone function is β -monotone and $\mathcal{MON}_\alpha \subseteq \mathcal{MON}_\beta$. We next show that, for any $\alpha \geq 1$, \mathcal{MON}_α is an ω -clone. Since $\mathcal{MON}_1 = \mathcal{MON}$, this implies that \mathcal{MON} is an ω -clone.

Theorem 62. *For any $\alpha \geq 1$, $\langle \mathcal{MON}_\alpha \rangle_\omega = \mathcal{MON}_\alpha$ and $\mathcal{B}_0 \subseteq \mathcal{MON}_\alpha$.*

Proof. Fix any $\alpha \geq 1$. We will show that $\text{EQ} \in \mathcal{MON}_\alpha$ and that it is closed under the usual operations.

- $\widetilde{\text{EQ}}$ is the unary constant function $\widetilde{\text{EQ}}(x) = 1$. Its only argument is fictitious. Thus, it is in \mathcal{MON}_α .
- For closure under permuting arguments, suppose that $f \in \mathcal{MON}_\alpha$ and that g is formed from f by permuting arguments. Then \tilde{g} is formed from \tilde{f} by permuting arguments. Since \tilde{f} is α -monotone, so is \tilde{g} , so $g \in \mathcal{MON}_\alpha$.
- For closure under introducing fictitious arguments, let $f \in \mathcal{MON}_\alpha$ and define $h(\mathbf{xy}) = f(\mathbf{x})$. Then the argument y is in its own equivalence class in \sim_h so $\tilde{h}(\mathbf{xy}) = \tilde{f}(x)$. Since \tilde{f} is α -monotone, and y is fictitious, \tilde{h} is α -monotone.
- For closure under summation, let $f \in \mathcal{MON}_\alpha$ and define $h(\mathbf{x}) = f(\mathbf{x}0) + f(\mathbf{x}1)$. Let $n + 1$ be the arity of f . There are two possibilities.

- First, suppose that x_{n+1} is equivalent to some other argument under \sim_f (for convenience, assume that it is equivalent to x_n). Then $f(x_1, \dots, x_n, x_n) = h(x_1, \dots, x_n)$ so $\tilde{f} = \tilde{h}$, and \tilde{h} is α -monotone because f is.
- Otherwise, let $V_1, \dots, V_{\ell+1}$ be the equivalence classes of \sim_f , where $V_{\ell+1}$ contains only x_{n+1} . Then V_1, \dots, V_ℓ are the equivalence classes of \sim_h . We claim that if the argument corresponding to V_i is not fictitious in \tilde{h} then it is not fictitious in \tilde{f} . Then, since \tilde{f} is α -monotone, changing the value of this argument increases the value of the function by a factor of α , both when $x_{n+1} = 0$ and when $x_{n+1} = 1$. Thus, changing the value of the argument also increases the value of \tilde{h} by a factor of α , and \tilde{h} is α -monotone.
- For closure under products, let $f, g \in \mathcal{MON}_\alpha$ and consider $h(\mathbf{x}) = f(\mathbf{x})g(\mathbf{x})$. Let n be the arity of h, f and g . Suppose that V_1, \dots, V_ℓ are the equivalence classes of \sim_f . Consider an equivalence class V_i and a string $\mathbf{x} \in \{0, 1\}^n$ which sets all arguments in V_i to 0. Let \mathbf{x}' be the string constructed from \mathbf{x} by changing the value of the arguments in V_i to 1. Since \tilde{f} is α -monotone, there are two possibilities:
 1. If the argument of \tilde{f} corresponding to V_i is fictitious then $f(\mathbf{x}') = f(\mathbf{x})$.
 2. If the argument of \tilde{f} corresponding to V_i is not fictitious then $f(\mathbf{x}') \geq \alpha f(\mathbf{x})$.

A similar comment applies to g . Now let the equivalence classes of \sim_h be $V'_1, \dots, V'_{\ell'}$. Consider some equivalence class V'_i — this is a union of \sim_f classes and a union of \sim_g classes. Suppose that the argument corresponding to V'_i is not fictitious in \tilde{h} . We want to argue that there is at least one of the \sim_f and \sim_g classes corresponding to V'_i that is not fictitious. To see this, suppose for contradiction that they are all fictitious. Start with a string \mathbf{x} in which all of the arguments in V'_i are the same. First consider changing, one-by-one all of the values of arguments in the \sim_f classes corresponding to V'_i . Since they are fictitious, this does not change the value of f . Similarly, changing-one-by-one all of the values in the \sim_g classes corresponding to V'_i does not change the value of g . So if \mathbf{x}' is the string derived from \mathbf{x} by changing the spin of V'_i , then $f(\mathbf{x}') = f(\mathbf{x})$ and $g(\mathbf{x}') = g(\mathbf{x})$. But this is a contradiction, since the argument corresponding to V'_i is not fictitious in \tilde{h} . Now suppose that \mathbf{x} takes value 0 on V'_i . Changing the value of some \sim_f or \sim_g class inside V'_i increases the value of the function by a factor of α . Changing each other \sim_f or \sim_g class inside V'_i either leaves the value alone, or increases it by another factor of α . Hence, \tilde{h} is α -monotone.

- For closure under limits, let $f \in \mathcal{B}_n$ and suppose that, for all integers $i > 0$, there is some $g_i \in \mathcal{MON}_\alpha$ such that $\|f - g_i\|_\infty < 2^{-i}$. We must show that $f \in \mathcal{MON}_\alpha$.

There must be some equivalence relation \sim_g on $[n]$ such that $\sim_{g_i} = \sim_g$ for infinitely many i . In fact, we may assume that $\sim_{g_i} = \sim_g$ for all i : if not, let g'_1, g'_2, \dots be the subsequence of functions whose equivalence relation is \sim_g , note that $\|f - g'_i\|_\infty < 2^{-i}$ for all i and use the sequence g'_1, g'_2, \dots in place of g_1, g_2, \dots .

Now, every equivalence class of \sim_f is a union of equivalence classes of \sim_g . To see this suppose that $r \sim_g s$. For all i and all $\mathbf{x} \in \{0, 1\}^n$ with $x_r \neq x_s$, we have $g_i(\mathbf{x}) = 0$. Therefore, $|f(\mathbf{x})| < 2^{-i}$ for all i , so $f(\mathbf{x}) = 0$ and $r \sim_f s$.⁴

⁴Note that we do not necessarily have $\sim_f = \sim_g$. For example, the 2-monotone symmetric binary function

Now, consider some argument x_j that is not fictitious in f . Let $\mathbf{x} \in \{0,1\}^n$ be a tuple such that $x_j = 0$ and $x_r = x_s$ whenever $r \sim_f s$. Let \mathbf{y} be the tuple with $y_s = x_s$ for all $s \not\sim_f j$ and $y_s = 1$ for $s \sim_f j$. We must show that $f(\mathbf{y}) \geq \alpha f(\mathbf{x})$. This is trivial when $f(\mathbf{x}) = 0$ so we consider the case that $f(\mathbf{x}) = \lambda > 0$. Then, for all large enough i , $g_i(\mathbf{x}) > \lambda - 2^{-i} > 0$ so, by block- α -monotonicity of g_i , $g_i(\mathbf{y}) > \alpha(\lambda - 2^{-i})$. So, for all large enough i , $g(\mathbf{y}) > \alpha(\lambda - 2^{-i}) - 2^{-i}$, so $g(\mathbf{y}) \geq \alpha\lambda$, as required.

The proof that $\mathcal{B}_0 \subseteq \mathcal{MON}_\alpha$ is straightforward since a function $f \in \mathcal{B}_0$ has no arguments, fictitious or otherwise. \square

11 Cardinality of the set of clones

In this section we determine the cardinality of the lattices of functional and ω -clones, proving Theorem 13.

Theorem 13. $|\mathcal{L}_f| = |\mathcal{L}_\omega| = \beth_2$.

Since $|\mathcal{B}| = \beth_1$, we have $|\mathcal{L}_f|, |\mathcal{L}_\omega| \leq \beth_2$. Therefore, we focus on proving the inverse inequality. As every ω -clone is also a functional clone, it suffices to prove that $|\mathcal{L}_\omega| \geq \beth_2$. We construct a set of functions, $\mathcal{F} \subseteq \mathcal{B}_2$ with $|\mathcal{F}| = \beth_1$ that has the following property: For any $\mathcal{G} \subseteq \mathcal{F}$, $\langle \mathcal{G} \rangle_\omega \cap \mathcal{F} = \mathcal{G}$. This immediately implies that, for any $\mathcal{G}_1, \mathcal{G}_2 \subseteq \mathcal{F}$ with $\mathcal{G}_1 \neq \mathcal{G}_2$ we have $\langle \mathcal{G}_1 \rangle_\omega \neq \langle \mathcal{G}_2 \rangle_\omega$. Therefore, $|\mathcal{L}_\omega| \geq 2^{|\mathcal{F}|} = \beth_2$.

For any real $\alpha > 2$, let f_α denote the binary function given by $f_\alpha(0,0) = 1$, $f_\alpha(0,1) = f_\alpha(1,0) = 2$ and $f_\alpha(1,1) = 2\alpha$. Let \mathcal{F} denote the set $\{f_\alpha \mid \alpha > 3\}$. Note that f_α is 2-monotone for any $\alpha > 2$. Therefore $\langle \mathcal{F} \rangle_\omega \subseteq \mathcal{MON}_2$.

Lemma 63. *Let $\mathcal{G} \subseteq \mathcal{F}$ be a finite set and let $\beta = \min\{\alpha \mid f_\alpha \in \mathcal{G}\}$. If $f \in \langle \mathcal{G} \rangle$ is a binary function such that $\tilde{f} = f$, then either $\gamma f \in \mathcal{G}$ for some constant γ or $f(0,1) \geq \min\left\{4, \frac{2+\beta}{2}\right\} f(0,0)$, or $f(1,0) \geq \min\left\{4, \frac{2+\beta}{2}\right\} f(0,0)$.*

Proof. Suppose f is a binary function in $\langle \mathcal{G} \rangle$. As noted in the introduction to the paper, f can be expressed as

$$f(x, y) = \sum_{x_1, \dots, x_k} \prod_{j=1}^t g'_j(x, y, x_1, \dots, x_k), \quad (5)$$

where t and k are non-negative integers and each function g'_j is a $(k+2)$ -ary function in $\mathcal{A}(\mathcal{G})$. Recall that $\mathcal{A}(\mathcal{G})$ is the closure of $\mathcal{G} \cup \{\text{EQ}\}$ under the introduction of fictitious arguments and permuting arguments. Thus, every function g'_j in (5) is constructed from a binary function $h \in \mathcal{G} \cup \{\text{EQ}\}$ by introducing fictitious arguments and permuting arguments. If $h = \text{EQ}$ then g'_j can be removed from the expression on the right-hand-side of (5) without changing the function $f(x, y)$ by instead allowing re-use of variables. (If h forces x_i and x_j to be equal, then we can just remove all instances of x_j and replace them with x_i and we can also remove x_j from the sum.) Also, we can remove the fictitious arguments in the g'_j functions, replacing each g'_j with the corresponding binary function $g_j \in \mathcal{G}$. Suppose that $f = \tilde{f}$ (so the variables x and y are in different \sim_f classes and removing the $h = \text{EQ}$ functions from (5) does not remove

$f(x, y) = [0, 0, 1]$ has just one equivalence class, but it is the limit of the 2-monotone functions $f_i(x, y) = [0, 2^{-i}, 1]$ for $i \geq 1$, and each of these functions has two equivalence classes.

either x or y). Then, by these transformations, (5) shows that there are non-negative integers s and m so that

$$f(x, y) = \sum_{u_1, \dots, u_m} \prod_{j=1}^s g_j(x_{j,1}, x_{j,2}), \quad (6)$$

where, for all $j \in [s]$, $g_j \in \mathcal{G}$ and $x_{j,1}$ and $x_{j,2}$ are in $\{x, y, u_1, \dots, u_m\}$ (though $x_{j,1}$ and $x_{j,2}$ may not necessarily be distinct).

Without loss of generality we assume that there are $0 \leq p, q, r \leq s$ such that functions g_j for $0 < j \leq p$ involve both x and y ; functions g_j for $p < j \leq p + q$ involve x but not y ; functions g_j for $p + q < j \leq p + q + r$ involve y but not x ; and the remaining functions do not involve x or y . Note that since none of x, y is fictitious, $p + q > 0$ and $p + r > 0$. For $x, y \in \{0, 1\}$ and $\mathbf{u} \in \{0, 1\}^m$, let also

$$\begin{aligned} T_{xy}(x, y) &= \prod_{j=1}^p g_j(x_{j,1}, x_{j,2}), & T_x(x, \mathbf{u}) &= \prod_{j=p+1}^{p+q} g_j(x_{j,1}, x_{j,2}), \\ T_y(y, \mathbf{u}) &= \prod_{j=p+q+1}^{p+q+r} g_j(x_{j,1}, x_{j,2}), & T_0(\mathbf{u}) &= \prod_{j=p+q+r+1}^s g_j(x_{j,1}, x_{j,2}). \end{aligned}$$

Thus,

$$f(x, y) = \sum_{\mathbf{u} \in \{0, 1\}^m} T_{xy}(x, y) T_x(x, \mathbf{u}) T_y(y, \mathbf{u}) T_0(\mathbf{u}).$$

If $p = 1$ and $q = r = 0$ then $\gamma f \in \mathcal{G}$ for $1/\gamma = \sum_{\mathbf{u} \in \{0, 1\}^m} T_0(\mathbf{u})$. If $p + q > 1$ we have for any $y \in \{0, 1\}$ and $\mathbf{u} \in \{0, 1\}^m$

$$4T_{xy}(0, y) T_x(0, \mathbf{u}) T_y(y, \mathbf{u}) T_0(\mathbf{u}) \leq T_{xy}(1, y) T_x(1, \mathbf{u}) T_y(y, \mathbf{u}) T_0(\mathbf{u}),$$

because every $g_j \in \mathcal{G}$ is 2-monotone. This implies that $4f(0, y) \leq f(1, y)$ for all y so, in particular, $4f(0, 0) \leq f(1, 0)$. Similarly, $4f(0, 0) \leq f(0, 1)$ if $p + r > 1$. Therefore, the only remaining case is $p = 0$ and $q = r = 1$.

Define α so that g_1 is the function $f_\alpha(x, u_1)$, where $f_\alpha \in \mathcal{G}$. Then

$$\begin{aligned} f(1, 0) &= \sum_{\mathbf{u} \in \{0, 1\}^m} f_\alpha(1, u_1) T_y(0, \mathbf{u}) T_0(\mathbf{u}) \\ &= \sum_{\mathbf{u}' \in \{0, 1\}^{m-1}} f_\alpha(1, 0) T_y(0, 0, \mathbf{u}') T_0(0, \mathbf{u}') + \sum_{\mathbf{u}' \in \{0, 1\}^{m-1}} f_\alpha(1, 1) T_y(0, 1, \mathbf{u}') T_0(1, \mathbf{u}') \\ &= 2 \sum_{\mathbf{u}' \in \{0, 1\}^{m-1}} f_\alpha(0, 0) T_y(0, 0, \mathbf{u}') T_0(0, \mathbf{u}') + \alpha \sum_{\mathbf{u}' \in \{0, 1\}^{m-1}} f_\alpha(0, 1) T_y(0, 1, \mathbf{u}') T_0(1, \mathbf{u}') \\ &= 2f(0, 0) + (\alpha - 2) \sum_{\mathbf{u}' \in \{0, 1\}^{m-1}} f_\alpha(0, 1) T_y(0, 1, \mathbf{u}') T_0(1, \mathbf{u}') \\ &\geq 2f(0, 0) + (\alpha - 2) \sum_{\mathbf{u}' \in \{0, 1\}^{m-1}} \frac{1}{2} (f_\alpha(0, 0) T_y(0, 0, \mathbf{u}') T_0(0, \mathbf{u}') + f_\alpha(0, 1) T_y(0, 1, \mathbf{u}') T_0(1, \mathbf{u}')) \\ &= \frac{2 + \alpha}{2} f(0, 0). \end{aligned}$$

The inequality here holds because all the functions involved are monotone and therefore

$$f_\alpha(0, 0) T_y(0, 0, \mathbf{u}') T_0(0, \mathbf{u}') \leq f_\alpha(0, 1) T_y(0, 1, \mathbf{u}') T_0(1, \mathbf{u}').$$

The result follows by the choice of β . \square

We can now prove Theorem 13.

Proof. As we observed at the beginning of the section, to prove that $|\mathcal{L}_\omega| \geq \beth_2$, it suffices to show that for any $\mathcal{G}_1, \mathcal{G}_2 \subseteq \mathcal{F}$ with $\mathcal{G}_1 \neq \mathcal{G}_2$, $\langle \mathcal{G}_1 \rangle_\omega \neq \langle \mathcal{G}_2 \rangle_\omega$. Suppose $f_\alpha \in \mathcal{G}_1 \setminus \mathcal{G}_2$. We show that $f_\alpha \notin \langle \mathcal{G}_2 \rangle_\omega$. The function f_α is symmetric and binary and it satisfies $f_\alpha = \widetilde{f_\alpha}$ so take any binary symmetric function $f \in \langle \mathcal{G}_2 \rangle_\omega$ with $\widetilde{f} = f$. We will show that $f \neq f_\alpha$.

By the definition of ω -clone, there is a finite set $\mathcal{G}'' \subseteq \langle \mathcal{G}_2 \rangle$ such that $f = \lim_{n \rightarrow \infty} h_n$ where each h_n is a binary function in $\langle \mathcal{G}'' \rangle$. Each of the finitely many functions in \mathcal{G}'' can be written as a finite sum of a product of (finitely many) functions in $\mathcal{A}(\mathcal{G}_2)$. Let \mathcal{G}' be the finite set of functions in \mathcal{G}_2 which correspond to the relevant functions in $\mathcal{A}(\mathcal{G}_2)$. Then clearly each $h_n \in \langle \mathcal{G}' \rangle$.

Since $\widetilde{f} = f$ there are at most finitely many n such that $\widetilde{h_n} \neq h_n$ — so we will remove these from the sequence of functions $\{h_n\}$ and of course it is still true that $f = \lim_{n \rightarrow \infty} h_n$.

Let $\beta = \min\{\mu \mid f_\mu \in \mathcal{G}'\}$. By Lemma 63, either there is a γ_n such that $\gamma_n h_n \in \mathcal{G}'$ or $h_n(0, 1) \geq \min\left\{4, \frac{2+\beta}{2}\right\} h_n(0, 0)$, or $h_n(1, 0) \geq \min\left\{4, \frac{2+\beta}{2}\right\} h_n(0, 0)$.

If there are infinitely many h_n such that $\gamma_n h_n \in \mathcal{G}'$ for some constants γ_n then, since \mathcal{G}' is finite, infinitely many of the $\gamma_n h_n$ functions are equal. Since $f(0, 0) = 1$, $\lim_{n \rightarrow \infty} \gamma_n = 1$ for such functions. Therefore, each of $\gamma_n h_n$ is equal to f . Thus $f \in \mathcal{G}'$ in this case. Clearly, $f \neq f_\alpha$ since $f_\alpha \notin \mathcal{G}_2$ so $f_\alpha \notin \mathcal{G}'$.

If there are finitely many h_n with $\gamma_n h_n \in \mathcal{G}'$ then we can remove these from the sequence of functions $\{h_n\}$ and as before it is still true that $f = \lim_{n \rightarrow \infty} h_n$. From now on, we therefore assume that none of the functions h_n belong to \mathcal{G}' . Suppose that $h_n(1, 0) \geq \min\left\{4, \frac{2+\beta}{2}\right\} h_n(0, 0)$ for infinitely many h_n . Then $f(1, 0) \geq \min\left\{4, \frac{2+\beta}{2}\right\} f(0, 0)$, and $f \notin \mathcal{F}$ so clearly $f \neq f_\alpha$. The case when there are infinitely many h_n with $h_n(1, 0) \geq \min\left\{4, \frac{2+\beta}{2}\right\} h_n(0, 0)$ is similar. \square

12 Ternary functions

In this section, we prove Theorem 14, which we restate here for convenience.

Recall that $\mathcal{S}_3 = \{\mathcal{B}_3, [\mathcal{SD}]_3, [\mathcal{P}]_3, [\mathcal{PN}]_3, [\mathcal{SDP}]_3, [(\mathcal{M})_\omega]_3, [\mathcal{H}_{\text{ferro}}]_3, [\mathcal{I}_{\text{ferro}}]_3\}$.

Theorem 14. $[\mathcal{SDP}]_3 = [(\mathcal{M})_\omega]_3$, $[\mathcal{H}_{\text{ferro}}]_3 = [\mathcal{I}_{\text{ferro}}]_3$, and any other two elements of \mathcal{S}_3 are distinct.

Proof. The two collapses are proved in Theorem 64 and Theorem 65, respectively.

Trivially, $[\mathcal{SD}]_3$, $[\mathcal{P}]_3$, and $[\mathcal{PN}]_3$ are strict subsets of $[\mathcal{B}]_3$. We now show that these three sets are distinct. Recall the binary functions f, g, h from the proof of Lemma 41 with $f \in \mathcal{SD} \setminus (\mathcal{P} \cup \mathcal{PN})$, $g \in \mathcal{P} \setminus (\mathcal{SD} \cup \mathcal{PN})$ and $h \in \mathcal{PN} \setminus (\mathcal{SD} \cup \mathcal{P})$. Define $f'(x, y, z) = f(x, y)$, $g'(x, y, z) = g(x, y)$, and $h'(x, y, z) = h(x, y)$. From the proof Lemma 41, $\widehat{f} = [\frac{1}{2}, 0, -\frac{1}{2}]$, $\widehat{g} = [1, \frac{1}{2}, 0]$ and $\widehat{h} = [1, -\frac{1}{2}, 0]$. By Lemma 25(ii), $\widehat{f}'(x, y, 0) = \widehat{f}(x, y)$ and $\widehat{f}'(x, y, 1) = 0$, and similarly for \widehat{g}' and \widehat{h}' . It is now easy to verify that $f' \in [\mathcal{SD}]_3 \setminus ([\mathcal{P}]_3 \cup [\mathcal{PN}]_3)$, $g' \in [\mathcal{P}]_3 \setminus ([\mathcal{SD}]_3 \cup [\mathcal{PN}]_3)$ and $h' \in [\mathcal{PN}]_3 \setminus ([\mathcal{SD}]_3 \cup [\mathcal{P}]_3)$.

Since $[\mathcal{SD}]_3$, $[\mathcal{P}]_3$, $[\mathcal{PN}]_3$ are distinct, we have $[\mathcal{SDP}]_3 \subset [\mathcal{SD}]_3, [\mathcal{P}]_3, [\mathcal{PN}]_3$.




Figure 5: The match-circuit used in the proof of Theorem 64. Every edge has weight 1 unless otherwise indicated.

Finally, we show that $[\mathcal{H}_{\text{ferro}}]_3 \subseteq [\mathcal{SDP}]_3$. The non-strict inclusion $[\mathcal{H}_{\text{ferro}}]_3 \subseteq [\mathcal{SDP}]_3$ comes from the fact that $\mathcal{H}_{\text{ferro}} \subseteq \mathcal{SD}$ (from Theorem 10). To get the strict inclusion, we will exhibit a ternary function that is in $[\mathcal{SDP}]_3$, but is not in $[\mathcal{LSM}]_3$ (so is not in $[\mathcal{H}_{\text{ferro}}]_3$ since $\mathcal{H}_{\text{ferro}} \subseteq \mathcal{LSM}$, as discussed in the proof of Theorem 44, and hence $[\mathcal{H}_{\text{ferro}}]_3 \subseteq [\mathcal{LSM}]_3$).

Consider the 3-ary self-dual function f defined by $f(0,0,0) = 6$, $f(0,0,1) = 4$, and $f(0,1,0) = f(1,0,0) = 5$. It can be verified that $f \in \mathcal{P}$ and thus $f \in [\mathcal{SDP}]_3$. We use Lemma 43 to show that $f \notin \mathcal{LSM}$: for 2-pinning $g(x,y) = f(x,y,1)$ we have $g(0,0)g(1,1) = 4 \cdot 6 = 24 < 25 = 5 \cdot 5 = g(0,1)g(1,0)$. \square

Theorem 64. $[\mathcal{SDP}]_3 = [\langle \mathcal{M} \rangle_\omega]_3$.

Proof. By Theorem 10, $\langle \mathcal{M} \rangle_\omega \subseteq \mathcal{SDP}$ and thus $[\langle \mathcal{M} \rangle_\omega]_3 \subseteq [\mathcal{SDP}]_3$. It remains to show the other inclusion, $[\mathcal{SDP}]_3 \subseteq [\langle \mathcal{M} \rangle_\omega]_3$. Consider a 3-ary function $f(x,y,z) \in \mathcal{SDP}$. If f is the constant zero function then $f \in [\langle \mathcal{M} \rangle_\omega]_3$ since \hat{f} is also the constant zero function, and it can be implemented by a match-circuit with three terminals and a disjoint triangle.

If f is not the constant zero function, then $\hat{f}(0,0,0) > 0$. There are values $a > 0$ and $b, c, d \geq 0$ such that $\hat{f}(0,0,0) = a$, $\hat{f}(0,1,1) = b$, $\hat{f}(1,0,1) = c$, $\hat{f}(1,1,0) = d$ and (by Lemma 38) $\hat{f}(\mathbf{x}) = 0$ when $|\mathbf{x}|$ is odd. It is easily verified that, for $b, c, d > 0$, \hat{f} is implemented by the match-circuit shown in Figure 5. Definition 5 does not allow zero-weight edges but we can implement \hat{f} in cases where some of b, c and d are zero by deleting the corresponding edge or edges from the match-circuit in Figure 5. Hence, $f \in [\mathcal{M}]_3$, so it is in $[\langle \mathcal{M} \rangle_\omega]_3$. \square

Theorem 65. $[\mathcal{H}_{\text{ferro}}]_3 = [\mathcal{I}_{\text{ferro}}]_3$.

Proof. By Theorem 10, $\mathcal{I}_{\text{ferro}} \subseteq \mathcal{H}_{\text{ferro}}$ and thus $[\mathcal{I}_{\text{ferro}}]_3 \subseteq [\mathcal{H}_{\text{ferro}}]_3$. It remains to show the other inclusion, $[\mathcal{H}_{\text{ferro}}]_3 \subseteq [\mathcal{I}_{\text{ferro}}]_3$. In fact, we prove something stronger, namely that $[\mathcal{SDP} \cap \mathcal{LSM}]_3 \subseteq [\mathcal{I}_{\text{ferro}}]_3$. By the proof of Theorem 44, $\mathcal{H}_{\text{ferro}} \subseteq \mathcal{SDP} \cap \mathcal{LSM}$ so the required inclusion follows.

An arbitrary 3-ary function $f \in \mathcal{SDP}$ is given by $f(0,0,0) = f(1,1,1) = \lambda$, $f(0,0,1) = f(1,1,0) = a$, $f(0,1,0) = f(1,0,1) = b$, and $f(1,0,0) = f(0,1,1) = c$.

For f to be in \mathcal{LSM} , by Lemma 43 (and in particular the fact that the necessary condition of Lemma 43 holds even for non-permissive functions, as discussed in Section 7), the following functions must also be in \mathcal{LSM} : $f_1, \dots, f_6 \in \mathcal{LSM}$ where $f_1(x,y) = f(0,x,y)$, $f_2(x,y) = f(x,0,y)$, $f_3(x,y) = f(x,y,0)$, $f_4(x,y) = f(1,x,y)$, $f_5(x,y) = f(x,1,y)$ and $f_6(x,y) = f(x,y,1)$. This gives

$$\lambda c \geq ab, \quad \lambda b \geq ac \quad \text{and} \quad \lambda a \geq bc. \quad (7)$$

Assume that f is permissive. Without loss of generality (by scaling since $\mathcal{B}_0 \subseteq \mathcal{I}_{\text{ferro}}$), let $\lambda = 1$. Let $g(x, y, z) = I_2^{\lambda_1}(x, y) I_2^{\lambda_2}(x, z) I_2^{\lambda_3}(y, z)$, where $\lambda_1 = \sqrt{bc/a}$, $\lambda_2 = \sqrt{ac/b}$ and $\lambda_3 = \sqrt{ab/c}$. By (7), $\lambda_1, \lambda_2, \lambda_3 \leq 1$, and hence $g \in [\mathcal{I}_{\text{ferro}}]_3$. We now verify that $f = g$. By the definition of g , $g(0, 0, 0) = g(1, 1, 1) = 1$, $g(0, 0, 1) = g(1, 1, 0) = 1 \cdot \lambda_2 \cdot \lambda_3 = a$, $g(0, 1, 0) = g(1, 0, 1) = \lambda_1 \cdot 1 \cdot \lambda_3 = b$, and $g(1, 0, 0) = g(0, 1, 1) = \lambda_1 \cdot \lambda_2 \cdot 1 = c$.

It remains to deal with non-permissive f . If $\lambda = 0$ then (7) implies that at most one of a , b , and c is non-zero. If all three are zero then f is the constant zero function and thus trivially in $[\mathcal{I}_{\text{ferro}}]_3$ since $\mathcal{B}_0 \subseteq \mathcal{I}_{\text{ferro}}$. Otherwise, let $a > 0$ and $b = c = 0$; the other two cases are symmetric. Since $\hat{f}(1, 1, 0) = -a/4 < 0$, $f \notin \mathcal{SDP}$, a contradiction. If $\lambda > 0$ then, by scaling, let $\lambda = 1$. The inequalities (7) imply that at most one of a , b , and c is non-zero. If all three are zero then $f(x, y, z) = \text{EQ}(x, y) \text{EQ}(y, z)$. Otherwise, let $a > 0$ and $b = c = 0$; the other two cases are symmetric. In this case $f(0, 0, 0) = f(1, 1, 1) = 1$, $f(0, 0, 1) = f(1, 1, 0) = a$ and $f(x, y, z) = 0$ otherwise. Thus $f(x, y, z) = \text{EQ}(x, y) I_2^a(y, z)$. Now, because $f \in \mathcal{LSM}$, we must have $f(0, 0, 0) f(1, 1, 1) = 1 \geq f(0, 0, 1) f(1, 1, 0) = a^2$, so $a \leq 1$ and $f \in [\mathcal{I}_{\text{ferro}}]_3$. \square

Acknowledgments

We thank the anonymous referee for their helpful comments.

References

- [1] David H. Ackley, Geoffrey E. Hinton, and Terrence J. Sejnowski. A learning algorithm for Boltzmann machines. *Cognitive Science*, 9(1):147–169, 1985.
- [2] Andrei A. Bulatov, Martin E. Dyer, Leslie Ann Goldberg, Mark Jerrum, and Colin McQuillan. The expressibility of functions on the Boolean domain, with applications to counting CSPs. *J. ACM*, 60(5):32, 2013.
- [3] Jin-Yi Cai and Aaron Gorenstein. Matchgates revisited. *Theory of Computing*, 10(7):167–197, 2014.
- [4] Jin-Yi Cai, Pinyan Lu, and Mingji Xia. Holographic algorithms with matchgates capture precisely tractable planar #CSP. In *2010 IEEE 51st Annual Symposium on Foundations of Computer Science FOCS 2010*, pages 427–436. IEEE Computer Soc., Los Alamitos, CA, 2010.
- [5] Xi Chen, Martin Dyer, Leslie Ann Goldberg, Mark Jerrum, Pinyan Lu, Colin McQuillan, and David Richerby. The complexity of approximating conservative counting CSPs. *J. Comput. System Sci.*, 81(1):311–329, 2015.
- [6] Barry A. Cipra. An introduction to the Ising model. *Amer. Math. Monthly*, 94(10):937–959, 1987.
- [7] Gemma De las Cuevas and Toby S. Cubitt. Simple universal models capture all classical spin physics. *Science*, 351(6278):1180–1183, 2016.
- [8] Ronald de Wolf. A brief introduction to Fourier analysis on the Boolean cube. *Theory of Computing, Graduate Surveys*, 1:1–20, 2008.

- [9] David Geiger. Closed systems of functions and predicates. *Pacific J. Math.*, 27:95–100, 1968.
- [10] Leslie Ann Goldberg and Mark Jerrum. The complexity of ferromagnetic Ising with local fields. *Combinatorics, Probability & Computing*, 16(1):43–61, 2007.
- [11] Leslie Ann Goldberg and Mark Jerrum. Approximating the partition function of the ferromagnetic Potts model. *J. ACM*, 59(5):25, 2012.
- [12] Leslie Ann Goldberg and Mark Jerrum. Inapproximability of the Tutte polynomial of a planar graph. *Computational Complexity*, 21(4):605–642, 2012.
- [13] Leslie Ann Goldberg and Mark Jerrum. A complexity classification of spin systems with an external field. *Proc. Natl. Acad. Sci. USA*, 112(43):13161–13166, 2015.
- [14] Geoffrey Grimmett. Potts models and random-cluster processes with many-body interactions. *J. Statist. Phys.*, 75(1-2):67–121, 1994.
- [15] Mark Jerrum and Alistair Sinclair. Polynomial-time approximation algorithms for the Ising model. *SIAM Journal on Computing*, 22(5):1087–1116, 1993.
- [16] Nicolas Le Roux and Yoshua Bengio. Representational power of restricted Boltzmann machines and deep belief networks. *Neural Comput.*, 20(6):1631–1649, 2008.
- [17] Ryan O’Donnell. *Analysis of Boolean Functions*. Cambridge University Press, 2014.
- [18] Emil L. Post. *The Two-Valued Iterative Systems of Mathematical Logic*. Annals of Mathematics Studies, no. 5. Princeton University Press, Princeton, N. J., 1941.
- [19] Donald M. Topkis. Minimizing a submodular function on a lattice. *Operations Research*, 26(2):305–321, 1978.
- [20] Leslie G. Valiant. Holographic Algorithms. *SIAM Journal on Computing*, 37(5):1565–1594, 2008.
- [21] B. L. van der Waerden. Die lange Reichweite der regelmäßigen Atomanordnung in Mischkristallen. *Zeitschrift für Physik*, 118:473–488, July 1941.

A Fourier transforms

In this appendix, we prove Lemmas 25–27.

Lemma 25. *Let f and g be functions in \mathcal{B}_k .*

- (i) *For any permutation π of $[k]$, $\widehat{f}^\pi(\mathbf{x}) = \widehat{f}(\pi(\mathbf{x}))$.*
- (ii) *If $h(\mathbf{xz}) = f(\mathbf{x})$, then $\widehat{h}(\mathbf{x0}) = \widehat{f}(\mathbf{x})$ and $\widehat{h}(\mathbf{x1}) = 0$.*
- (iii) *If $h(\mathbf{x}) = f(\mathbf{x0}) + f(\mathbf{x1})$, then $\widehat{h}(\mathbf{x}) = 2\widehat{f}(\mathbf{x0})$.*
- (iv) *If $h(\mathbf{x}) = f(\overline{\mathbf{x}})$, then $\widehat{h}(\mathbf{x}) = (-1)^{|\mathbf{x}|}\widehat{f}(\mathbf{x})$.*
- (v) *If $\|g - f\|_\infty < \varepsilon$, then $\|\widehat{g} - \widehat{f}\|_\infty < \varepsilon$.*

(vi) If $k = 0$ then $\widehat{f} = f$.

Proof. (i) In the following, the third equality is because permuting a Boolean vector doesn't change its Hamming weight and the fourth equality is reordering the terms of the sum.

$$\begin{aligned}
\widehat{f^\pi}(\mathbf{x}) &= \frac{1}{2^k} \sum_{\mathbf{w} \in \{0,1\}^k} (-1)^{|\mathbf{x} \wedge \mathbf{w}|} f^\pi(\mathbf{w}) \\
&= \frac{1}{2^k} \sum_{\mathbf{w} \in \{0,1\}^k} (-1)^{|\mathbf{x} \wedge \mathbf{w}|} f(\pi(\mathbf{w})) \\
&= \frac{1}{2^k} \sum_{\mathbf{w} \in \{0,1\}^k} (-1)^{|\pi(\mathbf{x}) \wedge \pi(\mathbf{w})|} f(\pi(\mathbf{w})) \\
&= \frac{1}{2^k} \sum_{\mathbf{w} \in \{0,1\}^k} (-1)^{|\pi(\mathbf{x}) \wedge \mathbf{w}|} f(\mathbf{w}) \\
&= \widehat{f}(\pi(\mathbf{x})).
\end{aligned}$$

(ii) For any $z \in \{0,1\}$,

$$\begin{aligned}
\widehat{h}(\mathbf{x}z) &= \frac{1}{2^{k+1}} \sum_{\mathbf{w} \in \{0,1\}^k} (-1)^{|\mathbf{x}z \wedge \mathbf{w}0|} f(\mathbf{w}) + \frac{1}{2^{k+1}} \sum_{\mathbf{w} \in \{0,1\}^k} (-1)^{|\mathbf{x}z \wedge \mathbf{w}1|} f(\mathbf{w}) \\
&= \frac{1}{2^{k+1}} \sum_{\mathbf{w} \in \{0,1\}^k} (-1)^{|\mathbf{x} \wedge \mathbf{w}|} f(\mathbf{w}) ((-1)^{|z \wedge 0|} + (-1)^{|z \wedge 1|}) \\
&= \frac{1}{2} \widehat{f}(\mathbf{x}) (1 + (-1)^{|z|}) \\
&= \begin{cases} \widehat{f}(\mathbf{x}) & \text{if } z = 0 \\ 0 & \text{if } z = 1. \end{cases}
\end{aligned}$$

(iii) For any \mathbf{x} ,

$$\begin{aligned}
\widehat{h}(\mathbf{x}) &= \frac{1}{2^{k-1}} \sum_{\mathbf{w} \in \{0,1\}^{k-1}} (-1)^{|\mathbf{x} \wedge \mathbf{w}|} (f(\mathbf{w}0) + f(\mathbf{w}1)) \\
&= \frac{1}{2^{k-1}} \sum_{\mathbf{w} \in \{0,1\}^{k-1}} (-1)^{|\mathbf{x}0 \wedge \mathbf{w}0|} f(\mathbf{w}0) + \frac{1}{2^{k-1}} \sum_{\mathbf{w} \in \{0,1\}^{k-1}} (-1)^{|\mathbf{x}0 \wedge \mathbf{w}1|} f(\mathbf{w}1) \\
&= \frac{1}{2^{k-1}} \sum_{\mathbf{w} \in \{0,1\}^k} (-1)^{|\mathbf{x}0 \wedge \mathbf{w}|} f(\mathbf{w}) \\
&= 2\widehat{f}(\mathbf{x}0).
\end{aligned}$$

(iv) Noting that $(-1)^{a-b} = (-1)^{a+b}$, we have

$$\begin{aligned}
\widehat{h}(\mathbf{x}) &= \frac{1}{2^k} \sum_{\mathbf{w} \in \{0,1\}^k} (-1)^{|\mathbf{x} \wedge \mathbf{w}|} f(\overline{\mathbf{w}}) \\
&= \frac{1}{2^k} \sum_{\mathbf{w} \in \{0,1\}^k} (-1)^{|\mathbf{x} \wedge \overline{\mathbf{w}}|} f(\mathbf{w}) \\
&= \frac{1}{2^k} \sum_{\mathbf{w} \in \{0,1\}^k} (-1)^{|\mathbf{x}| - |\mathbf{x} \wedge \mathbf{w}|} f(\mathbf{w}) \\
&= \frac{1}{2^k} (-1)^{|\mathbf{x}|} \sum_{\mathbf{w} \in \{0,1\}^k} (-1)^{|\mathbf{x} \wedge \mathbf{w}|} f(\mathbf{w}) \\
&= (-1)^{|\mathbf{x}|} \widehat{f}(\mathbf{x}).
\end{aligned}$$

(v)

$$\|\widehat{g} - \widehat{f}\|_\infty = \left\| \frac{1}{2^k} \sum_{\mathbf{w} \in \{0,1\}^k} (-1)^{|\mathbf{x} \wedge \mathbf{w}|} (g(\mathbf{w}) - f(\mathbf{w})) \right\|_\infty < \frac{1}{2^k} 2^k \varepsilon = \varepsilon. \quad \square$$

(vi) Suppose $k = 0$ so $f() = c$. Consider the unary function h defined by $h(0) = h(1) = c/2$. Then $f() = h(0) + h(1)$ so by item iii, $\widehat{f}() = 2\widehat{h}() = c$.

Lemma 26. For any k and λ ,

$$\widehat{I}_k^\lambda(\mathbf{x}) = \begin{cases} \lambda + (1 - \lambda)/2^{k-1} & \text{if } \mathbf{x} = \mathbf{0} \\ (1 - \lambda)/2^{k-1} & \text{if } |\mathbf{x}| \text{ is even and positive} \\ 0 & \text{if } |\mathbf{x}| \text{ is odd.} \end{cases}$$

Proof. We have

$$\begin{aligned}
2^k \widehat{I}_k^\lambda(\mathbf{x}) &= \sum_{\mathbf{w} \in \{0,1\}^k} (-1)^{|\mathbf{x} \wedge \mathbf{w}|} I_k^\lambda(\mathbf{w}) \\
&= (1 - \lambda)(1 + (-1)^{|\mathbf{x}|}) + \lambda \sum_{\mathbf{w} \in \{0,1\}^k} (-1)^{|\mathbf{x} \wedge \mathbf{w}|} \\
&= (1 - \lambda)(1 + (-1)^{|\mathbf{x}|}) + \lambda 2^{k-|\mathbf{x}|} \sum_{\mathbf{u} \in \{0,1\}^{|\mathbf{x}|}} (-1)^{|\mathbf{u}|},
\end{aligned}$$

where we adopt the convention that $\{0,1\}^0$ contains exactly one tuple, which has Hamming weight 0. This means the sum evaluates to 1 if $\mathbf{x} = \mathbf{0}$ and to zero, otherwise. \square

Lemma 27. For any k and λ , $\widehat{\text{Par}}_k^\lambda(\mathbf{0}) = \frac{1}{2}(1 + \lambda)$, $\widehat{\text{Par}}_k^\lambda(\mathbf{1}) = \frac{1}{2}(1 - \lambda)$ and $\widehat{\text{Par}}_k^\lambda(\mathbf{x}) = 0$ for any $\mathbf{x} \notin \{\mathbf{0}, \mathbf{1}\}$.

Proof. The first two equalities are straightforward from the definition. Let $x \in \{0,1\}^n \setminus \{\mathbf{0}, \mathbf{1}\}$.

$$\begin{aligned}
\widehat{\text{Par}}_k^\lambda(\mathbf{x}) &= \frac{1}{2^n} \sum_{\mathbf{w} \in \{0,1\}^n} (-1)^{|\mathbf{w} \wedge \mathbf{x}|} \text{Par}_k^\lambda(\mathbf{x}) \\
&= \frac{1}{2^n} \sum_{|\mathbf{w}| \text{ even}} (-1)^{|\mathbf{w} \wedge \mathbf{x}|} + \frac{\lambda}{2^n} \sum_{|\mathbf{w}| \text{ odd}} (-1)^{|\mathbf{w} \wedge \mathbf{x}|}.
\end{aligned}$$

Suppose without loss of generality that $x_1 = 0$ and $x_2 = 1$. Then, for every $\mathbf{w} \in \{0, 1\}^n$, the tuples $\mathbf{w} = (w_1, \dots, w_n)$ and $\mathbf{w}' = (\overline{w_1}, \overline{w_2}, w_3, \dots, w_n)$ have the same parity, but $|\mathbf{w} \wedge \mathbf{x}| \neq |\mathbf{w}' \wedge \mathbf{x}|$. Therefore

$$\sum_{|\mathbf{w}| \text{ even}} (-1)^{|\mathbf{w} \wedge \mathbf{x}|} = \sum_{|\mathbf{w}| \text{ odd}} (-1)^{|\mathbf{w} \wedge \mathbf{x}|} = 0. \quad \square$$