A Appendix

A.1 Experiments

The number of variables and clauses in each model counting problem are given in Table 2.

Table 2. Number of variables and clauses in each model counting problem.

INSTANCE	VARIABLES	CLAUSES
LANG12	576	13584
LANG15	1024	32320
LANG16	1024	32320
LANG19	1444	54226
LANG20	1600	63280
LANG23	2116	96370
LANG24	2304	109536
LANG27	2916	156114
LANG28	3136	174160
LS8	301	1603
LS9	456	2864
LS10	657	4761
LS11	910	7480
LS12	1221	11231
LS13	1596	16248
LS14	2041	22789
LS15	2562	31136
20RDR45	190	20349
23RDR45	253	42504
2BITMAX6	252	766
9SYMML	2604	36994
APEX75	1983	15358
FCLQ18	603	23312
FCLQ20	730	33662
VDAGRRCSW9	6498	130997

A.2 Proofs

Proof of Theorem 1. We closely follow the line of reasoning used in the proof of Theorem 3 of Ermon et al. (2014), which was for the specific case short XORs being the hash family. That line of reasoning can be distilled down to three main steps starting with a set $S \subseteq \{0,1\}^n$:

1. Fix an arbitrary $x \in S$ and, for any $w \in [n]$, let y_w be an arbitrary element of S at Hamming distance w from x. Then define a quantity of interest based on the clash probability:

$$\frac{1}{q-1}\sum_{w=1}^{n} h(w|x)\big(\Pr[h(x)=h(y_w)]\big)^m \quad (4)$$

where $h(w|x)$ denotes the number of elements of the set S at distance w from x.

2. Note that $h(w|x)$ is difficult to analyze. To circumvent this issue, observe that for the case of short XORs, the clash probability $\Pr[h(x) = h(y_w)]$ is a non-increasing function of w.

3. Using the above observation, upper bound the quantity in Eq. (4) by "squeezing" as many y's at small distances w as possible, until the total number is $|S|$:

$$\frac{1}{q-1}\sum_{w=1}^{w^*}\binom{n}{w}\big(\Pr[h(x)=h(y_w)]\big)^m \quad (5)$$

We describe how these three steps can be modified to work with functions f other than short XORs.

For step 1, instead of computing the clash probability for a particular hash family, we simply use Lemma 1 and replace $\Pr[h(x) = h(y_w)]$ with $1 - \mathrm{NS}'_w(h)$, which, by Proposition 2, is precisely $1 - \mathrm{NS}'_w(f)$.

Next, for step 2, rather than relying on monotonicity of the clash probability,[1] we define an ordering τ over $[n]$ under which $\mathrm{NS}'_{\tau(w)}(f)$ is non-increasing. Such an ordering must clearly exist. It is, in fact, often very close to the identity permutation.

Finally, for step 3, we use the τ ordering (rather than the identity ordering) and "squeeze" as many ys at distances w but *under this ordering* (rather than smallest-distance-first) as possible, until the total number is $|S|$. This yields:

$$\frac{1}{q-1}\sum_{w=1}^{w^*}\binom{n}{\tau(w)}\Big(1-\mathrm{NS}'_{\tau(w)}(f)\Big)^m \quad (6)$$

The rest of the argument follows precisely that in the proof of Theorem 3 of Ermon et al. (2014). □

Proof of Proposition 3. For $j, k \in [n]$, it will help to define $p(k, n, j)$ as:

$$p(k,n,j) = \frac{\sum_{\substack{\ell=1 \\ \ell:odd}}^{\min(k,j)}\binom{j}{\ell}\binom{n-j}{k-\ell}}{\binom{n}{k}} \quad (7)$$

This is the probability that if one chooses k balls uniformly out of n balls, then an odd number of balls are among some subset of j "special" balls.

Following ODonnell (2003), Proposition 2.3.1, we can write:

$$\mathrm{NS}'_w(f) = \frac{1}{2} - \frac{1}{2}\sum_{T\subseteq[n]}\mathbf{E}[x_T y_T]\hat{f}(T)^2$$

[1] Fixed distance clash probabilities behave very closely to the provably monotone model where each bit is flipped independently. However, they aren't always strictly monotone in the fixed distance model.

where $\hat{f}(T)$ are Fourier coefficients of f, $x_T y_T = \prod_{i \in T} x_i y_i$, and the expectation is taken over x chosen uniformly from $\{-1,1\}^n$ and y uniformly from $N_w(x)$. The only difference from ODonnell (2003) is in the distribution of y, which is now take from the fixed-distance model rather than flipping each bit of x independently. Observing that $x_T y_T$ is really only a function of the Hamming distance between x and y within the index set defined by T, we can set $x = 1^n$ without loss of generality. The resulting term, $\mathbf{E}[y_T]$, is determined by how often y has an *odd* number of -1 values in the bits belonging to T, which is precisely $p(w,n,|T|)$.

Specifically, $\mathbf{E}[x_T y_T] = 1 - 2\,p(w,n,|T|)$. Further, when $T \setminus \sup(f)$ is non-empty (here $\sup(f)$ refers to the support of the function f), i.e., T contains a bit that is not in the support of f, then $\hat{f}(T) = 0$. Thus we have:

$$\begin{aligned}\mathrm{NS}'_w(f) &= \frac{1}{2} - \frac{1}{2} \sum_{T \subseteq \sup(f)} \Big(1 - 2\,p(w,n,|T|)\Big) \hat{f}(T)^2 \\ &= \sum_{T \subseteq \sup(f)} p(w,n,|T|) \hat{f}(T)^2\end{aligned}$$

where the last equality follows from Parseval's identity which states that for any boolean function f, $\sum_{T \subseteq \sup(f)} \hat{f}(T)^2 = 1$. This proves the first part of the claim.

As further insight, we consider the case where f is XOR_j. Then the only non-zero Fourier coefficient is the one with $T = \sup(f)$, and this coefficient is 1, yielding $\mathrm{NS}'_w(\mathrm{XOR}_j) = p(w,n,j)$.

For the second part of the claim, we begin with some shorthand: $q(w) \equiv \binom{n}{w-1} \mathrm{NS}'_{w-1}(h)$, where q is defined over $w = 1, \ldots, n+1$. This offset is designed to simplify the application of the Vandermonde inverse. Note that our goal is to compute the $\mathrm{NS}'_w(h)$'s for $w = 0, \ldots, n$, which we can easily recover from the $q(w)$'s.

Using this notation, we can recover $n+1$ evaluations of a polynomial in ϵ whose coefficients are the $q(w)$'s as follows:

$$\begin{aligned}\mathrm{NS}_\epsilon(h) &= \sum_{w=0}^{n} \binom{n}{w} \epsilon^w (1-\epsilon)^{n-w} \mathrm{NS}'_w(h) \\ &= (1-\epsilon)^n \sum_{w=1}^{n+1} (\epsilon/(1-\epsilon))^{w-1} q(w)\end{aligned}$$

One can evaluate this polynomial at $\epsilon_1, \ldots, \epsilon_{n+1}$ evenly spaced between 0 and $1/2$. Defining $\overrightarrow{\epsilon} = [\epsilon_1, \ldots, \epsilon_{n+1}]$, and $\alpha_i = \epsilon_i/(1-\epsilon_i)$, rewriting the above in matrix form we have

$$\overrightarrow{\mathrm{NS}_\epsilon} = \mathrm{diag}((\overrightarrow{1} - \overrightarrow{\epsilon})^n) V_{\overrightarrow{\alpha}} \overrightarrow{q(w)} \tag{8}$$

where $V_{\overrightarrow{\alpha}}$ is the classical form of the Vandermonde matrix with parameters $\overrightarrow{\alpha}$, whose inverse B can be computed in closed-form. Using the closed-form Vandermonde inverse $[b_{ij}]$ we have

$$b_{ij} = (-1)^{j-1} \frac{\prod_{\substack{S \subseteq [n+1] \\ |S| = n+1-j \\ i \notin S}} \prod_{s \in S} \alpha_s}{\alpha_i \prod_{\substack{m=1 \\ m \neq i}}^{n+1} (\alpha_m - \alpha_i)}$$

and use that to solve for $\overrightarrow{q(w)}$ in (8), to calculate

$$q(w) = \frac{\epsilon_w}{(1-\epsilon_w)^n} \sum_{i=1}^{n+1} b_{iw} \mathrm{NS}_{\epsilon_i}(h)$$

as desired. Finally, $\mathrm{NS}'_w(f) = q(w+1)/\binom{n}{w}$. □