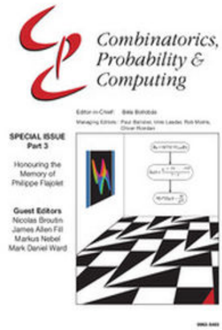

Local Maxima of Quadratic Boolean Functions

HUNTER SPINK

Combinatorics, Probability and Computing / Volume 25 / Issue 04 / July 2016, pp 633 - 640
DOI: 10.1017/S0963548315000322, Published online: 21 December 2015

Combinatorics, Probability and Computing (2016) 25, 633–640.
doi:10.1017/S0963548315000322

Local Maxima of Quadratic Boolean Functions

HUNTER SPINK

Trinity College, Cambridge University, United Kingdom
(e-mail: hunterspink@gmail.com)

Received 6 October 2013; revised 21 October 2014; first published online 21 December 2015

How many strict local maxima can a real quadratic function on $\{0,1\}^n$ have? Holzman conjectured a maximum of $\binom{n}{\lfloor n/2 \rfloor}$. The aim of this paper is to prove this conjecture. Our approach is via a generalization of Sperner's theorem that may be of independent interest.

2010 *Mathematics subject classification*: Primary 06E30
Secondary 05C65

1. Introduction

Let Θ be a real quadratic function (polynomial of total degree $\leqslant 2$) in n variables $x_1, \ldots, x_n$. A *strict local maximum* (or just *local maximum*) of Θ on the discrete cube $Q_n = \{0,1\}^n$ is a point whose value is strictly larger than all of its neighbours. As we are only concerned with the value of Θ on Q_n, we may assume when convenient that all terms are of degree 2, as the constant term is irrelevant, and we can replace x_i with x_i^2 if necessary. In this paper, we prove the following conjecture attributed to Ron Holzman (see [5, p. 3]; to the best of our knowledge it appeared first in the late 1980s, and was never formally published).

Theorem 1.1. *Let Θ be a quadratic function on Q_n. Then Θ has at most $\binom{n}{\lfloor n/2 \rfloor}$ local maxima.*

This bound is attained for example when $\Theta = -(x_1 + \cdots + x_n - \lfloor n/2 \rfloor)^2$. We will discuss all extremal (and close to extremal) examples at the end. From this, we can deduce some corollaries.

Corollary 1.2. *A (possibly degenerate) parallelepiped in $\mathbb{R}^n$ can have at most $\binom{n}{\lfloor n/2 \rfloor}$ vertices that are strictly closer to the origin than all of their neighbours.*

Proof. Let $\Theta'(\mathbf{x}) = -\Sigma x_i^2$, so we are counting strict local maxima of Θ' on the parallelepiped P. Perturb the vertices so that the parallelepiped is non-degenerate. Then there is a natural affine transformation τ mapping P to Q_n; letting $\Theta = \Theta' \circ \tau^{-1}$, we can apply Theorem 1.1 to Θ on Q_n. □

As a special case, we have the following.

Corollary 1.2′ (Littlewood–Offord theorem of Kleitman [4]). *The number of vertices of a (possibly degenerate) parallelepiped in* $\mathbb{R}^n$ *with all side lengths at least* 2 *that can land in the interior of a disc of radius* 1 *is at most* $\binom{n}{\lfloor n/2 \rfloor}$.

Proof. A vertex landing inside the disc must have all neighbours outside the disc, hence farther from the disc's centre. Thus we are done by Corollary 1.2. □

Let $[n] = \{1, 2, \ldots, n\}$, and let Δ denote symmetric set difference. We will find that a key ingredient in the proof of Theorem 1.1 is the following notion. Say that a family $\mathcal{F}$ is *weakly Sperner* if, for each $i \in [n]$, there exists a set $S_i \subseteq [n]$ such that the following holds. For every $A, B \in \mathcal{F}$ with $i \in A \,\Delta\, S_i$ and $i \notin B \,\Delta\, S_i$,

$$A \,\Delta\, S_i \not\supseteq B \,\Delta\, S_i.$$

An antichain is weakly Sperner, since then we can take $S_i = \emptyset$ for all $i \in [n]$.

We will show without difficulty that the set of local maxima of a quadratic boolean function is in fact a weakly Sperner family, and hence reduce the problem to the following purely combinatorial result.

Theorem 1.3. *Let* $\mathcal{F}$ *be a weakly Sperner family of subsets of* $[n]$. *Then* $|\mathcal{F}| \leqslant \binom{n}{\lfloor n/2 \rfloor}$.

Theorem 1.3 is perhaps interesting in its own right, as it gives a generalization of Sperner's theorem [6] (which is the case when all $S_i = \emptyset$). Moreover, Theorems 1.1 and 1.3 seem closely related to the influence of boolean functions, specifically the $k = 2$ case of the Gotsman–Linial conjecture [3] on the maximal total influence of a polynomial threshold function $\mathrm{sgn}(p(\mathbf{x}))$ on Q_n, where p is a real degree k polynomial.

We note that Erdős [2] originally resolved the one-dimensional case of the Littlewood–Offord problem using Sperner's theorem, and that Kleitman [4] developed his symmetric chain decomposition method to solve the n-dimensional case. It is therefore quite fitting that we use Kleitman's methods to prove a generalization of Sperner's theorem.

At the end of the paper, we analyse the quadratic forms that attain the maximum number of local maxima. It turns out that our method allows us not only to deduce the structure of the quadratic function when we attain equality, but also when we are within $\frac{1}{n}\binom{n}{\lfloor n/2 \rfloor}$ of the optimal solution $\binom{n}{\lfloor n/2 \rfloor}$. We do not know how close this bound of $\frac{1}{n}\binom{n}{\lfloor n/2 \rfloor}$ is to being optimal.

This paper is self-contained. See [1] for general background on set systems and Sperner's theorem.

2. Proof

We frequently view the discrete cube as the family of subsets of $[n]$ in the obvious way (using the x_i as indicator functions). Assuming all terms of Θ are of degree 2 (so that it is a quadratic form), we have its associated symmetric matrix (q_{ij}), where $\Theta(\mathbf{x}) = \sum_{i,j} q_{ij}x_ix_j$.

2.1. Proof that Theorem 1.3 implies Theorem 1.1

The fact that Theorem 1.3 implies Theorem 1.1 is an immediate consequence of the following.

Claim 2.1. *Let $\Theta : Q_n \to \mathbb{R}$ be a quadratic function, and set*

$$\mathcal{F} = \{A \subseteq [n] : A \text{ is a local maximum of } \Theta\}.$$

Then $\mathcal{F}$ is weakly Sperner with respect to

$$S_i = \{j \in [n]\backslash\{i\} : q_{ij} > 0\}.$$

Proof. Fix $i \in [n]$. The difference between the value of Θ on the $x_i = 1$ plane and the $x_i = 0$ plane as a function of the remaining coordinates is given by the linear function $\theta_i(\mathbf{x}) = q_{ii} + 2\sum_{j\neq i} q_{ij}x_j$. Since $i \notin S_i$, we need to show that we cannot have $A, B \in \mathcal{F}$ with $i \in A$ and $i \notin B$ with $A \,\Delta\, S_i \supseteq B \,\Delta\, S_i$. Indeed, as A is a local maximum, we must have $q_{ii} + 2\sum_{i\neq j\in A} q_{ij} > 0$ and similarly $q_{ii} + 2\sum_{i\neq j\in B} q_{ij} < 0$, so taking the difference we get

$$2\sum_{i\neq j\in A} q_{ij} - 2\sum_{i\neq j\in B} q_{ij} = 2\Bigg(\sum_{i\neq j\in S_i^c\cap(A\backslash B)} q_{ij} - \sum_{i\neq j\in S_i\cap(B\backslash A)} q_{ij}\Bigg) > 0,$$

which is clearly false as both terms are non-positive. Here we used that $A \,\Delta\, S_i \supseteq B \,\Delta\, S_i$ implies $B\backslash A \subset S_i$ and $A\backslash B \subset S_i^c$. □

Geometrically, if we used a coordinate system with S_i at the origin via the Q_n automorphism $C \mapsto C \,\Delta\, S_i$, $x_j \mapsto 1 - x_j$ for $j \in S_i$, then the normal direction to θ_i is in the negative orthant, and $A \supseteq B$. Hence $q_{ij} \leqslant 0$ for $j \neq i$ (i fixed), and we get a contradiction from $2\sum_{i\neq j\in A\backslash B} q_{ij} > 0$.

2.2. Proof of Theorem 1.3

For the remainder of the proof, we view Q_n as the family of subsets of $[n]$. $\mathcal{F}$ is weakly Sperner with respect to $S_1, \ldots, S_n \in Q_n$, and we need to show $|\mathcal{F}| \leqslant \binom{n}{\lfloor n/2\rfloor}$.

Inspired by Kleitman's proof of the Littlewood–Offord problem [4] (or see [1, Ch. 4] for general background), we seek an inductive decomposition of Q_n (described below) into $\binom{n}{\lfloor n/2\rfloor}$ parts (which we will call quasichains), such that we can guarantee that at most one element of $\mathcal{F}$ lies inside each part. The heart of this proof is the definition of a 'quasichain', which allows Kleitman's symmetric chain decomposition method to go through. The definition is rather surprising and seems contrived; however, as we will see, once we have this definition the proof is straightforward.

Recall that a *tournament* is a complete directed graph. For sets $B, C \subseteq [n]$, we write $B \supseteq_{S_i} C$ to mean $B \,\Delta\, S_i \supseteq C \,\Delta\, S_i$.

Definition. If A, B are vertices of Q_n, $i \in [n]$, $S_1, \ldots, S_n \in Q_n$, we say $A \xrightarrow{i} B$ is an admissible arrow with respect to $S_1, \ldots, S_n$ if $i \in A \,\Delta\, S_i$, $i \notin B \,\Delta\, S_i$, and $A \supseteq_{S_i} B$.

Definition. We define a *quasichain* $\mathcal{C}$ (with respect to $S_1, \ldots, S_n$) to be a coloured tournament (with colours in $[n]$) with vertex set a family of subsets $\mathcal{G} = \{G_1, \ldots, G_k\}$ of $[n]$, such that:

(i) whenever there is a directed edge from G_x to G_y of colour i, then $G_x \xrightarrow{i} G_y$ is an admissible arrow with respect to $S_1, \ldots, S_n$;
(ii) for any subset of the colours (including the empty set), if we swap the direction of edges associated with those colours, then the resulting tournament is acyclic (or equivalently transitive).

It is easy to check that the acyclicity condition is equivalent to saying that no triangle has three distinct colours, any monochromatic triangle is acyclic, and any triangle with two distinct colours has the two edges with the same colour either both leaving, or both entering, the same vertex.

Note that a quasichain does *not* contain all possible information about $\supseteq_{S_i}$ containment between its various members, but rather remembers only one such containment for every pair (just enough information to guarantee that at most one element from the pair can lie in $\mathcal{F}$, and hence at most one element from each quasichain can lie in $\mathcal{F}$).

Definition. Define a *symmetric quasichain decomposition* of Q_n (with respect to $S_1, \ldots, S_n$) to be a partition of the n-cube into families of sets each of which can support at least one structure of a quasichain with respect to $S_1, \ldots, S_n$, inductively built up from the trivial decomposition of the 0-cube as follows. Suppose we have decomposed Q_{k-1} into sets which support quasichain structures (with respect to $S_1 \cap [k-1], \ldots, S_{k-1} \cap [k-1]$). Let $\{\mathcal{C}_1, \ldots, \mathcal{C}_r\}$ be some fixed choices of these structures. Then for each $\mathcal{C}_j$ with underlying vertex set $\{A_1, \ldots, A_s\}$, we choose some $t \in [s]$ in such a way that $\{A_1, A_2, \ldots, A_s, A_t \cup \{k\}\}$ and $\{A_1 \cup \{k\}, \ldots, A_{t-1} \cup \{k\}, A_{t+1} \cup \{k\}, \ldots, A_s \cup \{k\}\}$ can both be given a quasichain structure with respect to $S_1 \cap [k], \ldots, S_k \cap [k]$, creating quasichains $\mathcal{C}'_j, \mathcal{C}''_j$ respectively. Discarding any quasichains with empty vertex set from the collection of $\mathcal{C}'_j$ and $\mathcal{C}''_j$, we now have the desired partition of Q_k supporting quasichain structures with respect to $S_1 \cap [k], \ldots, S_k \cap [k]$.

Of course, the difficult part is to ensure that we can always find an A_t which allows us to create the two new quasichains.

We now show (or see [1, pp. 17–20]) that this process will always result in $\binom{n}{\lfloor n/2 \rfloor}$ families of sets in Q_n each of which supports at least one quasichain structure.

Lemma 2.2. *A symmetric quasichain decomposition of Q_n has $\binom{n}{\lfloor n/2 \rfloor}$ parts.*

Proof. We inductively show that in any such decomposition of Q_n, there are $\binom{n}{i} - \binom{n}{i-1}$ parts of size $n+1-2i$ for $i = 0, 1, \ldots, \lfloor n/2 \rfloor$, where we set $\binom{n}{k} = 0$ if $k < 0$, and no parts of any other size. Indeed, this is true for $n = 0$; suppose it is now true for Q_{n-1}, $n \geqslant 1$. The number of parts of size r in Q_n is simply the sum of the number of parts of size $r-1$ and of size $r+1$ in Q_{n-1}. Hence only parts of sizes $n+1-2i$ for $i = 0, 1, \ldots, \lfloor n/2 \rfloor$ are possible, and

$$\left(\binom{n-1}{i-1} - \binom{n-1}{i-2}\right) + \left(\binom{n-1}{i} - \binom{n-1}{i-1}\right) = \binom{n}{i} - \binom{n}{i-1}$$

as desired (one can check separately that this calculation is valid for $i = 0, 1$ using the convention $\binom{n}{k} = 0$ if $k < 0$). Thus the total number of quasichains is

$$\binom{n}{0} + \left(\binom{n}{1} - \binom{n}{0}\right) + \cdots + \left(\binom{n}{\lfloor n/2 \rfloor} - \binom{n}{\lfloor n/2 \rfloor - 1}\right) = \binom{n}{\lfloor n/2 \rfloor}. \qquad \square$$

As a quasichain can intersect our weakly Sperner set $\mathcal{F}$ in at most one element, the above lemma shows that if we exhibit a symmetric quasichain decomposition for Q_n with respect to $S_1, \ldots, S_n$, then we have proved $|\mathcal{F}| \leqslant \binom{n}{\lfloor n/2 \rfloor}$ as desired.

Lemma 2.3. *Suppose* $i \in [n-1]$, *and* $A \xrightarrow{i} B$ *is an admissible arrow in* Q_{n-1} *with respect to* $S_1 \cap [n-1], \ldots, S_{n-1} \cap [n-1]$. *Then the following are admissible arrows with respect to* $S_1, \ldots, S_n$ *in* Q_n*:*

(i) $A \xrightarrow{i} B$,

(ii) $A \cup \{n\} \xrightarrow{i} B \cup \{n\}$,

(iii) $A \cup \{n\} \xrightarrow{i} B$ *if* $n \notin S_i$,

(iv) $A \xrightarrow{i} B \cup \{n\}$ *if* $n \in S_i$

Proof. For $X \in Q_{n-1}$, $i \in [n-1]$, we have $i \in X \,\Delta\, (S_i \cap [n-1])$ if and only if $i \in X \,\Delta\, S_i$. Thus to check admissibility, we only need to check the $\supseteq_{S_i}$ containment. But in all four cases this is trivial to check using $A \,\Delta\, (S_i \cap [n-1]) \supseteq B \,\Delta\, (S_i \cap [n-1])$ and splitting into two cases depending on whether $S_i = S_i \cap [n-1]$ or $S_i = (S_i \cap [n-1]) \cup \{n\}$. $\square$

Proof of Theorem 1.3. We will prove the theorem by using induction on n to show that Q_n admits a symmetric quasichain decomposition with respect to any collection of n sets $S_1, \ldots, S_n$. For $n = 0$, we simply take the whole of Q_0 as our quasichain, so now assume that $n \geqslant 1$.

Given $S_1, \ldots, S_n \in Q_n$, we have by the induction hypothesis a symmetric chain decomposition for Q_{n-1} with respect to the $S_1 \cap [n-1], \ldots, S_{n-1} \cap [n-1]$. We will extend this decomposition to one of Q_n with respect to $S_1, \ldots, S_n$.

By Lemma 2.3, if we view a quasichain $\mathcal{C}$ with underlying vertex set $\{A_1, \ldots, A_r\}$ in Q_{n-1} with respect to $S_1 \cap [n-1], \ldots, S_{n-1} \cap [n-1]$ as a coloured tournament in Q_n, then all arrows remain admissible with respect to $S_1, \ldots, S_n$ by part (i) of Lemma 2.3, and the second quasichain condition is clearly preserved, so it remains a quasichain. Similarly,

$\mathcal{C}^n = \{A_1 \cup \{n\}, \ldots, A_r \cup \{n\}\}$ is also a quasichain with respect to $S_1, \ldots, S_n$ (with the same coloured arrows) by part (ii) of Lemma 2.3. To complete the symmetric quasichain decomposition, we need to exhibit an $A_t \cup \{n\}$ we can transfer from $\mathcal{C}^n$ to $\mathcal{C}$ in such a way that they can be made into quasichains with respect to $S_1, \ldots, S_n$.

Let $R = \{i \in [n-1] : n \text{ lies in } S_i\}$. By the second quasichain condition for $\mathcal{C}$, if all arrows with colours in R are reversed, then the resulting tournament is acyclic; let A be a maximal element for this new graph. This means that in $\mathcal{C}$, the colour of all arrows pointing into A belong to R, and the colour of all arrows out do not belong to R. We choose $A \cup \{n\}$ as the element to transfer from $\mathcal{C}^n$ to $\mathcal{C}$. If we remove a vertex from a quasichain, then all arrows remain admissible, and a subgraph of an acyclic graph is acyclic, so $\mathcal{C}^n - (A \cup \{n\})$ is a quasichain if we keep the same directed edges on the remaining vertices. Thus it suffices to show that $\mathcal{C} \cup (A \cup \{n\})$ is in fact a quasichain once we appropriately colour and direct edges containing $A \cup \{n\}$ (we retain the same edge colourings on all edges which were already in $\mathcal{C}$).

If $A \xrightarrow{i} X$, then $i \notin R$, so $n \notin S_i$, so $A \cup \{n\} \xrightarrow{i} X$ is admissible by part (iii) of Lemma 2.3. If $X \xrightarrow{i} A$, then $i \in R$, so $n \in S_i$, so $X \xrightarrow{i} A \cup \{n\}$ is admissible by part (iv) of Lemma 2.3. Thus if we duplicate the direction and colour of all arrows incident to A in $\mathcal{C}$ for $A \cup \{n\}$, then these new arrows will all be admissible. Finally, $A \cup \{n\} \xrightarrow{n} A$ is admissible if $n \notin S_n$, and $A \xrightarrow{n} A \cup \{n\}$ is admissible if $n \in S_n$, so one of the two directions for $\xrightarrow{n}$ between $A \cup \{n\}$ and A gives us the final admissible arrow for $\mathcal{C} \cup (A \cup \{n\})$.

To complete the proof, we must show that the newly constructed graph satisfies the acyclicity condition. After swapping some directions associated with a subset of the colours, we have a tournament H, with vertices x, y corresponding to A, $A \cup \{n\}$ respectively. For all z, then, we have either $x \to z$, $y \to z$, or $z \to x$, $z \to y$. If we have a cycle, then identifying x and y yields a cycle in the original graph (since x and y have the same incoming/outgoing edges, this identification is well-defined), which is a contradiction. □

3. Analysis of the quadratic form close to equality

Theorem 3.1. *If the number of local maxima of a quadratic function $\Theta : Q_n \to \mathbb{R}$ is greater than*

$$\left(1 - \frac{1}{n-c}\right)\binom{n}{\lfloor n/2 \rfloor},$$

where $c = 0, 1$ if n is odd/even respectively, then there is an automorphism of Q_n such that after applying it, the local maxima form an antichain.

Proof. Assume as before that $\Theta(\mathbf{x}) = \sum_{i,j} q_{ij} x_i x_j$ with $q_{ij} = q_{ji}$. Let

$$B = \{j \in [n] \backslash \{1\} : q_{1j} > 0\}.$$

Applying the Q_n automorphism $A \mapsto A \,\Delta\, B$, $x_j \mapsto 1 - x_j$ for $j \in B$, we can assume that $q_{1j} \leqslant 0$ for $j \in [n] \backslash \{1\}$.

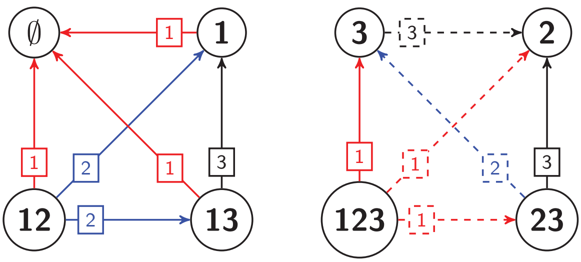

Figure 1. Quasichain decomposition at stage Q_3

Recall that

$$\mathcal{F} = \{A \subseteq [n] : A \text{ is a local maximum of } \Theta\}$$

was shown to be weakly Sperner with respect to

$$S_i = \{j \in [n] \backslash \{i\} : q_{ij} > 0\}.$$

Note that $S_1 = \emptyset$. We have $i \notin S_i$, and because $q_{ij} = q_{ji}$, $i \in S_j$ if and only if $j \in S_i$. We will show that if not all S_i are empty (*i.e.*, there is a $q_{ij} > 0$ with $i \neq j$), then we are bounded away from the optimal solution by the desired factor. Indeed, suppose without loss of generality that S_2 contains 3. Then S_3 contains 2, so the intersections with $\{1, 2, 3\}$ of S_1, S_2, S_3 are $\emptyset, \{3\}, \{2\}$ respectively. The first three stages yield the quasichain decomposition shown in Figure 1 (in bold lines).

Using the dotted lines, we see that we can 'glue' together two of these quasichains to make a single quasichain. We want to understand the 'evolution' of this quasichain as it goes through the symmetric chain algorithm. We know that a 2-quasichain after k steps becomes $\binom{k+1}{\lfloor (k+1)/2 \rfloor}$ quasichains. As a 2-quasichain evolves into a 4-quasichain plus two 2-quasichains after two steps, if $g(k)$ is the number of quasichains the 4-quasichain evolves into after k steps, then we have

$$\binom{k+3}{\lfloor (k+3)/2 \rfloor} = 2\binom{k+1}{\lfloor (k+1)/2 \rfloor} + g(k).$$

After we reach n dimensions, therefore, these two 4-quasichains will have evolved into

$$2g(n-3) = 2\binom{n}{\lfloor n/2 \rfloor} - 4\binom{n-2}{\lfloor (n-2)/2 \rfloor}$$

quasichains. The difference between the actual bound and this is

$$4\binom{n-2}{\lfloor (n-2)/2 \rfloor} - \binom{n}{\lfloor n/2 \rfloor},$$

which is equal to

$$\frac{1}{n-c}\binom{n}{\lfloor n/2 \rfloor},$$

where c is 0 or 1 depending on whether n is odd or even, respectively.

If every S_i is in fact empty, then the condition the S_i create is precisely the normal antichain condition, so the family creates an actual antichain. □

We remark here that in the construction of the quasichain decomposition, we can choose whether to transfer the largest member (after arrow reversing) of $\mathcal{C}^n$ to $\mathcal{C}$, or the smallest member of $\mathcal{C}$ to $\mathcal{C}^n$. By combining this choice with the above gluing, and keeping track of multiple different quasichain structures on the same set, we can improve the $1/(n-c)$ constant slightly.

Acknowledgements

I would like to thank Imre Leader for helpful conversations, and the referees for their helpful comments on the presentation. I would also like to thank Trinity College, Cambridge for providing me with financial support.

References

[1] Bollobás, B. (1986) *Combinatorics: Set Systems, Hypergraphs, Families of Vectors, and Combinatorial Probability*, Cambridge University Press.
[2] Erdős, P. (1945) On a lemma of Littlewood and Offord. *Bull. Amer. Math. Soc.* **51** 898–902.
[3] Gotsman, C. and Linial, N. (1994) Spectral properties of threshold functions. *Combinatorica* **14** 35–50.
[4] Kleitman, D. J. (1970) On a lemma of Littlewood and Offord on the distribution of linear combinations of vectors. *Adv. Math.* **5** 251–259.
[5] O'Donnell, R. (2012) Open problems in analysis of boolean functions. arXiv:1204.6447v1
[6] Sperner, E. (1928) Ein Satz über Untermenge einer endlichen Menge. *Math. Z.* 544–548.