
How to Turn All the Lights Out

Paulo Ventura Araújo*

Paulo Ventura Araújo was born in northern Portugal in July 1966 and, except for a 3-year period in England, has lived nowhere else. He completed his first degree in Mathematics at the University of Oporto (where he now teaches) and gained his Ph.D. in January 1993 from the University of Warwick. His thesis was on symbolic dynamics, but since then his mathematical interests have shifted somewhat: he is now a committed non-specialist with a leaning towards geometry. He has written two textbooks in Portuguese, one on differential geometry (published in Brazil) and the other on elementary geometry, and has also published a few research papers on the hyperbolic plane.

1 Introduction

A recent article [1] presents an analysis of a one-person game which consists of a square board divided into 25 smaller squares, each containing a light bulb attached to a button, and each button reversing the state of the corresponding bulb and of the neighbouring ones on the same row or column (thus each button-pushing affects from 3 to 5 bulbs). Given an initial position with some of the bulbs turned on, the object is to reach the position where all the bulbs are turned out (henceforth called the *zero position*). A complete description of the *solvable* positions (i.e., positions from which it is possible to reach the zero position), together with some generalizations for higher-order boards, is given in [1].

Im vorliegenden Beitrag von P.V. Araújo kann auf wenigen Seiten verfolgt werden, wie mit geschickt eingesetzten, elementaren Methoden der linearen Algebra über dem Körper mit zwei Elementen – wie man sie auch erfolgreich in der Kodierungstheorie anwendet – der Ausgang eines gewissen Spiels vorhergesagt werden kann. Für die vom Spielfieber erfassten Leserinnen und Leser lohnt es sich vielleicht, ein kleines Programm zum vorgetragenen Spiel zu schreiben, denn der Spielerfolg hängt von der Grösse und dem Anfangszustand des Spielfeldes ab. Am Ende des Beitrags wird vom Autor eine Spielstrategie vorgeschlagen. Jeder ist gefragt, ob es vielleicht noch eine bessere Strategie gibt. *jk*

*) Financially supported by FCT through the project Praxis 2/2.1/Mat/19/94.

Here we consider a variation of that game. The *switchboard of order n* (abbreviated \mathfrak{S}_n) is an $n \times n$ square of light bulbs and attached buttons where we assume that each button reverses the state of every bulb on the same row or on the same column (this adds up to $2n - 1$ bulbs); the object is again to turn all the lights out. We prove that the parity of n makes a great difference to our game: whereas for even n every position is *solvable* [Theorem 1], for odd n only a fraction of $1/2^{2n-2}$ of all possible positions are solvable [Theorem 3]. In the latter case, we give a simple characterization of the solvable positions [Theorem 5], and describe in our final section a strategy for reaching the zero position from any such position.

2 Even versus odd

We say that a sequence of button-pushings (or *moves*) is *elementary* if its net result is that of reversing the state of just one bulb, leaving all the others unchanged. It is clear, by symmetry, that if an elementary sequence exists for \mathfrak{S}_n , then we can specify one such sequence for reversing any given bulb, and therefore, by turning out each lighted bulb in turn, we see that every position of \mathfrak{S}_n is solvable. It is easy to exhibit elementary sequences when n is even: given any bulb, simply push once each button on the same row and each button on the same column. The given bulb changes state exactly $2n - 1$ times; since this number is odd, the final state is the reverse of the original one; all other bulbs suffer either 2 or n reversals, and therefore their final state is identical to the original one. Thus:

Theorem 1 *If n is even, then every position of \mathfrak{S}_n is solvable.*

		•
•		
•		

The above figure, where each dot corresponds to a lighted bulb, and each empty square to an unlighted bulb, gives an example of an unsolvable position for \mathfrak{S}_3 . Indeed, consider the set \mathfrak{T} of four squares belonging to either the first column or the first row *but not to both*. Every single move reverses the state of an even number of bulbs on \mathfrak{T} . Hence, starting with three lighted bulbs on \mathfrak{T} and applying an arbitrary sequence of moves, \mathfrak{T} will always end up with either one or three lighted bulbs, and therefore we never reach the zero position.

This example can be generalized to arbitrary odd n . But first a bit of terminology. By the *T-set of type (i, j)* in \mathfrak{S}_n , denoted by \mathfrak{T}_{ij} , we mean the set of squares belonging to either row i or column j *but not to both*. Now let n be odd and consider any *T-set* \mathfrak{T}_{ij} : the crucial observation is that, if n is odd, then every move affects an even number (equal to 2, $n - 1$ or $2n - 2$) of bulbs on \mathfrak{T}_{ij} ; hence, if we have an initial position with an odd number of lighted bulbs on \mathfrak{T}_{ij} , then we can never reach the zero position. We can summarize this as follows:

Proposition 2 *If n is odd, then a necessary condition for a given position of \mathfrak{S}_n to be solvable is that every T-set contains an even number of lighted bulbs.*

Our main result, to be proved in the next section, is that the above condition is also sufficient. Thus, for odd n , we have a practical method for checking whether a given configuration is solvable; this involves n^2 checks, as many as the distinct T -sets. We will show, however, that these T -sets are not all independent and that, as a result, we only need to check $2n - 2$ suitably chosen T -sets (e.g., all the T -sets that include a given square – see Theorem 5).

3 mod 2 linear algebra

Following [1], we now introduce a method for coding all possible positions of \mathfrak{S}_n . Let $\mathbb{K} = \mathbb{Z}/2\mathbb{Z}$ be the field with the two elements 0 and 1, and let $V = \mathbb{K}^{n^2}$; it is also useful to think of V as $(\mathbb{K}^n)^n$. Each (v_1, \dots, v_n) in V represents a position of \mathfrak{S}_n as follows: row i is coded by $v_i = (v_{i1}, \dots, v_{in})$, so that $v_{ij} = 1$ if and only if the bulb on row i and column j is lighted. Reversing the state of a bulb means interchanging 0 and 1, which is equivalent in \mathbb{K} to addition by 1. For instance, the effect on position (v_1, \dots, v_n) of pushing the top-left button is that of adding $(1, 1, \dots, 1; 1, 0, \dots, 0; \dots; 1, 0, \dots, 0)$ to it. More generally, pushing the button on row i and column j is equivalent to adding the vector

$$X_{ij} = (e_j, \dots, e_j, X, e_j, \dots, e_j),$$

where $X = (1, 1, \dots, 1) \in \mathbb{K}^n$ appears in the i^{th} position, and e_j is the j^{th} vector of the canonical basis of \mathbb{K}^n .

A given position v is solvable if and only if there is a sequence of moves that transforms it into the zero position; but if we start from the zero position and perform the same moves in reverse order, we again obtain v . Thus we see that a position v is solvable if and only if it can be reached from zero; and this means that v can be written as a sum of vectors X_{ij} . Thus the solvable positions constitute a subspace U of V , more precisely the subspace spanned by the vectors X_{ij} .

We can reformulate Proposition 2 by using the ordinary inner product $v \cdot w$ of two vectors $v, w \in V$. Let each T -set \mathfrak{T}_{ij} be identified with the element in V whose entries corresponding to squares in \mathfrak{T}_{ij} are 1, and all others are 0. Then the following conditions are easily seen to be equivalent:

- (i) Position v is such that the number of lighted bulbs on \mathfrak{T}_{ij} is even;
- (ii) $v \cdot \mathfrak{T}_{ij} = 0$.

Let W be the subspace generated by all \mathfrak{T}_{ij} . Then we see that Proposition 2 asserts that, if n is odd, then $U \subseteq W^\perp$. (Incidentally, we also have $W \cap W^\perp = \{\vec{0}\}$, so that by Theorem 3 below U and W are complementary subspaces.) We now state our main result:

Theorem 3 *Let $n \geq 3$ be odd. Then the set U of all solvable positions is the orthogonal space to the subspace W generated by the T -sets. Furthermore, we have $\dim U = n^2 - 2n + 2$ and $\dim W = 2n - 2$.*

In our proof we use the fact that, for any subspace E of V , we have

$$\dim E + \dim E^\perp = \dim V = n^2, \quad (1)$$

and therefore also $(E^\perp)^\perp = E$. [For a proof of (1), pick up any $n^2 \times n^2$ matrix A over \mathbb{K} whose column vectors generate E , and let $f : V \rightarrow V$ be given by $f(v) = vA$; then (1) is equivalent to the well-known relation $\dim f(V) + \dim \text{Ker } f = \dim V$.] Since we have already shown that $U \subseteq W^\perp$, it is sufficient, to establish the theorem, to prove that

$$\dim U \geq n^2 - 2n + 2 \quad (2)$$

and

$$\dim W \geq 2n - 2. \quad (3)$$

Proof of (2). Let $B \in \mathcal{M}_{n \times n}(\mathbb{K})$ be the matrix all of whose entries are equal to 1, and let $I \in \mathcal{M}_{n \times n}(\mathbb{K})$ be the identity matrix. It is easy to check that the columns of the following $n^2 \times n^2$ matrix

$$M = \begin{bmatrix} B & I & \cdots & I & I \\ I & B & \cdots & I & I \\ & & \cdots & \cdots & \\ I & I & \cdots & B & I \\ I & I & \cdots & I & B \end{bmatrix}$$

are precisely the vectors X_{ij} defined above; hence the dimension of U equals the rank of M . The following lemma deals with general matrices of this type:

Lemma 4 *Let Y, Z be $p \times p$ matrices. Denote by $M_m(Y, Z)$ the $pm \times pm$ matrix whose diagonal blocks are equal to Y , and whose other blocks are equal to Z . Then we have*

$$\det M_m(Y, Z) = (\det[Y - Z])^{m-1} \cdot \det[Y + (m-1)Z]. \quad (4)$$

Proof. For a $pm \times pm$ matrix, we divide its rows into m disjoint sets, each consisting of p consecutive rows; we call each such set a *fat row*, and we define *fat columns* similarly. We perform elementary fat row and fat column operations on $M_m(Y, Z)$. First we replace, for each j , fat column j by the sum of the fat columns 1, 2, ..., j of $M_m(Y, Z)$; in the resulting matrix we subtract successively, for $i = 1, 2, \dots, m-1$, fat row $i+1$ from fat row i (thus we replace fat row 1 by the difference of fat rows 1 and 2, and so on). In the end we obtain

$$\begin{bmatrix} Y - Z & & & & \\ & Y - Z & & & \\ & & \ddots & & \\ & & & Y - Z & \\ Z & 2Z & \dots & (m-1)Z & Y + (m-1)Z \end{bmatrix}, \quad (5)$$

where all non-zero blocks belong either to the main diagonal or to the last fat row. Now (4) follows readily from (5). \square

In the notation of the lemma, we have $M = M_n(B, I)$, and therefore

$$\det M = (\det[B - I])^{n-1} \cdot \det[(n-1)I + B] = 0,$$

since $n-1 \equiv 0 \pmod{2}$, and so the second factor reduces to $\det B$, which is 0. Hence the rank of M is less than n^2 . Since $B - I = M_n(0, 1)$, we have by (4)

$$\det[B - I] = (-1)^{n-1} \cdot (n-1) \equiv 0 \pmod{2}, \quad (6)$$

and, denoting by B^* and I^* the $(n-1) \times (n-1)$ matrices analogous to B and I , also

$$\det[B^* - I^*] = (-1)^{n-2} \cdot (n-2) \equiv 1 \pmod{2}. \quad (7)$$

Let P be the lower triangular matrix of type (5) obtained from M as in Lemma 4: thus $\text{rank } M = \text{rank } P$, and P has, on the main diagonal, $n-1$ blocks $B - I$ and one block B . Let P^* be the sub-matrix of P obtained by removing rows and columns kn , for $k = 1, \dots, n-1$, and also removing the last $n-1$ rows and columns. Then P^* is a $(n^2 - 2n + 2) \times (n^2 - 2n + 2)$ triangular matrix where, on the main diagonal, we have $n-1$ consecutive blocks $B^* - I^*$ and one final entry equal to 1. Using (7), we have

$$\det P^* = (\det[B^* - I^*])^{n-1} \equiv 1 \pmod{2},$$

and it follows that $\text{rank } M \geq \text{rank } P^* = n^2 - 2n + 2$, which concludes the proof of (2). \square

Proof of (3). The subspace W is spanned by the T -sets

$$\mathfrak{T}_{ij} = (e_j, \dots, e_j, X - e_j, e_j, \dots, e_j) = X_{ij} - (0, \dots, 0, e_j, 0, \dots, 0),$$

and so it is sufficient to exhibit a set of $2n-2$ linearly independent \mathfrak{T}_{ij} . We prove that $\mathfrak{T}_{12}, \mathfrak{T}_{13}, \dots, \mathfrak{T}_{1n}, \mathfrak{T}_{21}, \mathfrak{T}_{31}, \dots, \mathfrak{T}_{n1}$ form just such a set. Indeed, let Q be the matrix whose columns are the given vectors in the same order, and let Q^* the $(2n-2) \times (2n-2)$ sub-matrix of Q made up of rows $2, 3, \dots, n$ and $n+1, 2n+1, \dots, (n-1)n+1$. Then we have

$$Q^* = \begin{bmatrix} B^* - I^* & 0 \\ 0 & B^* - I^* \end{bmatrix},$$

and therefore Q^* is non-singular by (7), and so $\text{rank } Q = 2n-2$. This completes the proof of (3) and also that of Theorem 3. \square

This proof gives an explicit basis for U ; and, in view of Theorem 3, to check whether a given position $v \in V$ is solvable, we only have to compute the inner product of v with each of the $2n-2$ vectors of any basis of U . The following theorem sums up these observations for basis of the type given in the proof of (3):


Theorem 5 *Let $n \geq 3$ be odd, and fix $k, l \in \{1, \dots, n\}$. The following conditions on a position $v \in V$ are equivalent:*

- (i) v is solvable;
- (ii) $v \cdot \mathfrak{T}_{kj} = 0$ for all $j \in \{1, \dots, n\} \setminus \{l\}$, and $v \cdot \mathfrak{T}_{il} = 0$ for all $i \in \{1, \dots, n\} \setminus \{k\}$.

4 Back to the switchboard

Throughout this section, $n \geq 3$ is a fixed odd number. In Section 3 we gave a method for checking whether a given position v of \mathfrak{S}_n is solvable (Theorem 5); in this section, starting with a solvable position v , we describe explicitly a sequence of moves that will lead from v to the zero position.

Our strategy is to identify a number of *basic configurations* which every non-zero position must contain. The basic configurations are of three types (see the figure below):



- Diagonal*: we say that v contains a diagonal if there exists a permutation σ of $\{1, \dots, n\}$ such that $v_{i\sigma(i)} = 1$ for all i .
- Corner*: we say that v contains a corner if there exist $k, l, m, p \in \{1, \dots, n\}$ such that $k \neq l, m \neq p$ and $v_{mk} = v_{pk} = v_{pl} = 1$.
- Cross*: v contains a cross if there exist l_1, l_2, l_3 , all distinct, and m_1, m_2, m_3 , also all distinct, such that $v_{l_1 m_1} = v_{l_2 m_1} = v_{l_3 m_2} = v_{l_3 m_3} = 1$.

We prove below (Proposition 6) that, if v is solvable and non-zero, then v contains one of these configurations. We now describe how, in each of the above cases, we can decrease the *size* of v (i.e., its total number of 1's) by an appropriate sequence of moves. By the results of the previous section the new position thus obtained is solvable; if non-zero, we can again decrease its size, and continue in this way until only 0's are left.

- For each $i = 1, \dots, n$, push once the button on $(i, \sigma(i))$. This reverses all the bulbs on these squares, leaving the others unchanged. The size of v is thus decreased by n .
- Push once every button on columns k and l , except those on rows m and p . Exactly four bulbs are reversed: those on (m, k) , (p, k) , (p, l) or (m, l) . The size of v is decreased by either 2 or 4, depending on whether we had $v_{ml} = 0$ or $v_{ml} = 1$.
- If $v_{l_2 m_2} = 1$ or $v_{l_2 m_3} = 1$ then v contains a corner, and we proceed as in *b*). Otherwise, we push every button on columns m_2 and m_3 , except those on rows l_2 and l_3 . The size of v is unchanged, but we now have two additional 1's on row l_2 . The new position contains a corner on (l_1, m_1) , (l_2, m_1) and (l_2, m_2) ; and now we may proceed as in *b*).

Proposition 6 *If v is a solvable non-zero position, then v contains at least one basic configuration.*

Proof. Assume first that no row or column of v contains two or more 1's. Pick up (k, l) such that $v_{kl} = 1$. By Theorem 5, we have $v \cdot \mathfrak{T}_{il} = 0$ for each $i \neq k$. Since v_{kl} is the only non-zero entry on column l , this implies that $\sum_{\{j:j \neq l\}} v_{ij}$ is odd, and therefore row i contains at least one (and therefore exactly one) 1. Similarly, each column contains exactly one 1, and we conclude that v is a *diagonal configuration*.

Now assume that some row or column contains at least two 1's. For definiteness, assume that row k contains two or more 1's. If, for some j such that $v_{kj} = 1$, there exists some other 1 on column j , then we have succeeded in finding a *corner configuration* in v . Otherwise, we have the situation where each 1 on row k is the only non-zero entry on the respective column. *We claim that then there is a column containing at least two 1's*; and it follows from this claim that v contains a *cross configuration*. Fix l such that $v_{kl} = 1$, and let $i \neq k$. Using the fact that $v \cdot \mathfrak{T}_{il} = 0$, we conclude as above that there exists $\sigma(i)$ such that $v_{i\sigma(i)} = 1$. But, since we must have $v_{k\sigma(i)} = 0$, the function $\sigma : \{1, \dots, n\} \setminus \{k\} \rightarrow \{1, \dots, n\}$ excludes at least two values and so is not injective; and this proves our claim. \square

References

- [1] Anderson, M. and Fell, T.: Turning lights out with linear algebra. *Math. Magazine* 71 (1998), 300–303.

Paulo Ventura Araújo
Centro de Matemática
Faculdade de Ciências do Porto
Praça Gomes Teixeira
4099–002 Porto, Portugal
e-mail: paraujo@fc.up.pt