

Parity Lights

Dennis Epple

Martin Kutz*

Mathematisches Institut II
Freie Universität Berlin
14195 Berlin, Germany
{epple,kutz}@math.fu-berlin.de

1 Introduction

Almost everybody has a light in their house that can be switched on and off from two or more different places. If your electrician set up the wiring properly, you are able to toggle the light with every single switch, independent of the others' current configuration.

We may think of the switches as implementing a parity function. If we label the two positions of a switch by 0 and 1, the light is on if the sum of switch positions is even and off if this sum is odd (or vice versa). So in the common case of only two switches, 'light on' is just the binary relation `xor`. (Not all real-life installations are of this kind. We also find the binary `and` implemented in a few houses; which can be very annoying—and even dangerous, at nights.)

We present a mathematical model for such electrical light switch circuits in which a single switch consists of several elementary switching units. Then we ask for an optimal setup for the general n -switch problem; *optimal* here means using as few of these components as possible.

We employ elementary techniques from graph theory to obtain the surprising result that, with respect to our formalism, the standard implementation of the n -switch problem installed all over the world is optimal.

2 The standard setup

Before delving into mathematical formalism, let us get some feel for light switches. For a start, operating one light from two switches is quite simple. The configuration shown in Figure 1 does the job. The left switch toggles between the connections (s, a) and (s, b) , the right between (c, t) and (d, t) .

Published in *The Mathematical Gazette*, vol. 88, no. 511, pp. 46–56, March 2004.

*Member of the European graduate school 'Combinatorics, Geometry, and Computation' supported by the Deutsche Forschungsgemeinschaft, grant GRK 588/1.




Figure 1: One light with two switches.




Figure 2: A four-way switch.

There are different names around for this circuit. In the United States most call it a ‘three-way switched circuit’, and the dashed boxes are then respectively called ‘three-way switches’. The origin of the term ‘three-way’ is not quite clear. Some attribute it to the number of wires that run between the power source and the load (two from the power source to the light and one back), others suggest that ‘three’ count the number of devices from the power supply to the load (two switches, one light).

An alternative name for the configuration in Figure 1 is ‘two-way circuit’, maybe because there are exactly two wires running between the two switches. We prefer the latter view and consequently call those dashed boxes *two-way switches* since they connect one wire to one out of *two* wires on the other side.

When $n > 2$ inputs are to be wired, so-called *four-way switches* shown in Figure 2 are used. Such a switch is made up of two two-way switches that either connect the pairs (x, u) and (y, v) , or the pairs (x, v) and (y, u) . That are four possible connections altogether, which somehow justifies the name. But historically, there again seem to be different reasons for the term ‘four-way’.

If we now put $n - 2$ four-way switches in line between two terminal two-way switches, as shown Figure 3, we get the desired switching property: any single switch can toggle the light, independent of the position of the other components. Figure 3 shows the *standard setup* for the n -switch problem. You will find it in almost any house with electric power supply—provided they did not use relays instead of ordinary switches.




Figure 3: The standard setup for more than two switches.

3 Modelling circuits

We want to investigate the efficiency of electrical light switch circuits with respect to the number of elementary switching units needed. The wiring itself shall be neglected. Assume that we have good connections to some major copper company with an almost infinite supply of wire so that we may run hundreds of kilometres of cable at no cost at all.

The cost of a switch. So what should be the price for a switch? Obviously, a complex structure like a four-way switch should not count as a single unit and even a two-way switch can be divided into more elementary components.

We let an *interruptor* be a device that can connect and disconnect a single pair of wire ends. Hence, a two-way switch consists of two interruptors, a four-way switch of four. And thus the naming conventions we picked for those switches turn out very sound.

Circuits as graphs. We will model light-switching circuits in the language of graph theory. The employed concepts are very basic and we briefly present required notions, so that even the reader who encounters them for the first time should have no great problems to follow our constructions. The only graph theoretic concept that might deserve a more careful treatment is that of *contraction*, which plays an important role in the central proof of this paper, so we will address it in greater detail when needed. The following definitions can be found in any textbook on graph theory, such as [1, 2, 3].

A *graph* $G = (V, E)$ consists of a set V of *vertices* and a set E of *edges*. An edge $e \in E$ is an unordered pair $\{a, b\}$ of vertices in V and we say that e *connects* the vertices a and b and that it is *incident* to either of them. A *path* in a graph G is a sequence v_0, v_1, \dots, v_r of pairwise distinct vertices with $\{v_i, v_{i+1}\} \in E$ for each $i < r$.

Formally, our definition excludes multiple edges between pairs of vertices but it will be convenient to agree that a ‘set’ of edges may contain a single vertex pair several times so that we get *parallel edges*. This common technical sloppiness saves us some tedious formalism and will lead to no problems. Further we also allow *loops*, i.e., edges that connect a vertex to itself.

Our aim is to describe light-switching circuits as graphs. So let us begin with the vertices. They embody the static parts of the circuit, that is, the wires, the power supply, and the light. Since we are not interested in the routing of wires but




Figure 4: Graph representations of (a) the two-way circuit and (b) the four-way switch.

only in their function as contact points for interruptors, we may take the simplified view that coherent wire segments are just points, so they become the vertices of our graph. The direct connection between power source and light is irrelevant for our switches and can thus be ignored. For the circuit in Figure 1 we get four vertices: one for the left wire, one for the right wire, and two vertices between the switches. The position of power source and light are encoded by labelling the respective wire segments s ('source') and t ('target').

The edges of the graph are the dynamic components of the circuit—the interruptors. For our example circuit in Figure 1 this makes four edges altogether. Figure 4(a) shows a picture of the resulting graph.

Of course this description is still incomplete, it contains no information about the semantics of the interruptors. We have to specify which interruptor is operated by which switch, and how. To this end we associate two edge sets E and F with each switch. The edges in E are short when the switch is in position 0, those in F are short when it is in position 1. Figure 4(b) shows the graph representation of the four-way switch from Figure 2; the edges in E are drawn as continuous, those in F as dashed lines. Formally we have $E = \{\{x, u\}, \{y, v\}\}$ and $F = \{\{x, v\}, \{y, u\}\}$.

A single switch has only two different configurations. In a circuit we can manipulate several switches independently, so we need sets E and F for each switch. Here is the formal definition.

An n -input switching circuit is a tuple

$$C = (V, s, t, E_1, \dots, E_n, F_1, \dots, F_n),$$

where V is a finite set of vertices, s, t are vertices of V , and $E_1, \dots, E_n, F_1, \dots, F_n$ are sets of edges on V . We call

$$\text{size}(C) := \sum_{i=1}^n |E_i| + |F_i|$$

the *size* of the switching circuit C .

Circuits in action. It remains to specify which input configurations turn the light on and which turn it off. Consider therefore some binary input vector $d = (d_1, d_2, \dots, d_n) \in \{0, 1\}^n$. Each bit stands for the position of a switch, which in

turn ‘activates’ certain edges and ‘deactivates’ others. We define the d -graph of C to be the graph

$$G(C, d) := \left(V, \bigcup_{d_i=0} E_i \cup \bigcup_{d_i=1} F_i \right),$$

that is, the graph with vertex set V and with all those edges of C that correspond to active interruptors.

The light is on if and only if there is a connection between power source and light, formally, if there is a path from s to t in the d -graph. We define the *induced function* $\lambda_C : \{0, 1\}^n \rightarrow \{0, 1\}$ of the switching circuit C by

$$\lambda_C(d) := \begin{cases} 1 & \text{if there is a path from } s \text{ to } t \text{ in } G(C, d), \\ 0 & \text{if there is no path from } s \text{ to } t \text{ in } G(C, d). \end{cases}$$

Recall from the introduction that we are interested in precisely those circuits that realize the parity function. So we call an n -switching circuit C with

$$\lambda_C(d) \equiv \sum_{i=1}^n d_i \pmod{2} \quad \text{for all } d \in \{0, 1\}^n \quad (1)$$

an n -parity circuit.

The alert reader might remark that the above definition depends on the choice of the edge sets E_i and F_i . If we swap those two sets for an arbitrary index i , equation (1) is no longer satisfied, although toggling any input still changes the state of the light. But if we change the definition of parity circuits by adding 1 to either side of equation (1), the modified circuit becomes a parity circuit again. So there are actually two equitable notions of parity circuits. We chose ours just because it has the nice property that the light is off if all inputs are 0, formally, $\lambda_C(0) = 0$.¹ Fortunately, it will turn out that this choice only effects some intermediate results, our important statements hold for both notions of parity circuits.

Minimal circuits. The aim of this article is to investigate how many interruptors are needed to implement an n -parity circuit. Hence we define

$$p(n) := \min \{ \text{size}(C) \mid C \text{ } n\text{-parity circuit} \}.$$

Obviously $p(1) = 1$. For $n \geq 2$ consider Figure 5, which shows the standard setup translated into graph theory. A letter at an edge indicates to which set E_i or F_i it belongs. From this scheme we get the upper bound

$$p(n) \leq 4n - 4, \quad n \geq 2. \quad (2)$$

The rest of the paper is dedicated to showing equality in (2):

Theorem 1. *The standard setup for n -parity circuits is optimal! We have*

$$p(n) = \begin{cases} 1 & \text{for } n = 1, \\ 4n - 4 & \text{for } n \geq 2. \end{cases}$$

¹Here 0 also denotes the zero vector.




Figure 5: The standard setup in graph notation.

4 Basic properties of parity circuits

Parity circuits inherit a simple but important property from the parity function on bit vectors. Flipping an even number of bits in a vector does not, by definition, change its parity. And flipping an even number of switches in a parity circuit does not change its induced function. By *flipping* the i -th input of a switching circuit, we mean reversing that switch, i.e., exchanging the edge sets E_i and F_i .

Lemma 1 (Exchange Lemma). *For an n -parity circuit and two indices i, j , $1 \leq i < j \leq n$, exchanging E_i with F_i and E_j with F_j yields a parity circuit again.*

Proof. Let C denote the original circuit and C' the modified one. Consider any n -bit vector d' and let d denote the vector obtained from d' by flipping bits i and j . We have $G(C', d') = G(C, d)$ and hence $\lambda_{C'}(d') = \lambda_C(d)$. The bit vectors d' and d have the same parity, so C' is a parity circuit because C is. \square

We now turn to our actual task, to find a lower bound on the size of an n -parity circuit. Looking at the definition of the size of a circuit, we ask the natural question whether in a parity circuit any of the sets E_i and F_i can be empty. It turns out that, except for the minimal 1-parity circuit on two vertices s, t with $E_1 = \emptyset$ and $F_1 = \{\{s, t\}\}$, this cannot happen.

Lemma 2. *For any n -parity circuit, $n \geq 2$, none of the edge sets $E_1, \dots, E_n, F_1, \dots, F_n$ is empty.*

Proof. Assume for contradiction that $F_j = \emptyset$ for some index j . Let I_j denote the n -bit vector containing a 1 in position j and zeros elsewhere. Then $\lambda_C(I_j) = 1$ by parity. Hence the graph $G(C, I_j)$ contains an s - t -path. But since F_j is empty, the graph $G(C, 0)$ also contains this path. Thus we arrive at the contradiction $\lambda_C(0) = 1$.

The case $E_j = \emptyset$ reduces to the above case by application of the exchange lemma. Note that we need $n \geq 2$ here. \square

From Lemma 2 we can already deduce $p(2) = 4$ since Figure 1 shows a parity circuit of size 4.




Figure 6: A 3-parity circuit of size 9.

5 The lower bound

We learnt that except for the trivial case $n = 1$, a parity circuit has no empty edge sets. Let us follow this thread of thought and investigate whether we can improve upon this bound.

The basic idea. Let us find out whether any of the sets E_i or F_i in a parity circuit can be singleton. That is, can there be an edge set containing only one edge?

Of course, the standard setup shows there can; it has $|E_1| = |F_1| = |E_n| = |F_n| = 1$. But maybe this is just best possible. Perhaps we can show that in any n -parity circuit, $n \geq 3$, all but two inputs have edge sets of size at least two. Together with Lemma 2 this would directly imply optimality of the standard setup.

Unfortunately we cannot. Figure 6 shows a 3-parity circuit with all $|F_i| = 1$. But this example does not ultimately thwart our ambitions. It turns out that the case $n = 3$ is somewhat degenerate. In fact, restricting ourselves to $n \geq 4$, we can show the following:

Proposition 1. *In any n -parity circuit, $n \geq 4$, there are at most two indices i with $|E_i| = 1$ or $|F_i| = 1$.*

Although Proposition 1 excludes the case $n = 3$, our proof essentially relies on a fact about 3-parity circuits. So let us consider those first.

A closer look at 3-parity circuits. We consider again the circuit from Figure 6. The singleton sets are all F -sets, that means, they all correspond to a switch in 1-position. Is there a similar circuit with all E -sets of size one? The answer is ‘no’.

Lemma 3. *There is no 3-parity-circuit with $|E_1| = |E_2| = |E_3| = 1$.*

Proof. We assume for contradiction that there exists a 3-parity circuit C with $E_1 = \{e_1\}$, $E_2 = \{e_2\}$, and $E_3 = \{e_3\}$.

If two of these edges were identical, say $e_1 = e_2$, then the graph $G(C, (0, 0, 1)) = (V, E_1 \cup E_2 \cup F_3)$ would be a subgraph of $G(C, (0, 1, 1)) = (V, E_1 \cup F_2 \cup F_3)$. Then from $\lambda_C((0, 0, 1)) = 1$ we deduce the contradiction $\lambda_C((0, 1, 1)) = 1$. Thus, e_1, e_2 , and e_3 must be pairwise different.

We now consider an s - t -path p in the graph $G(C, (0, 0, 1))$. Obviously p must contain e_1 and e_2 since otherwise $\lambda_C((1, 0, 1)) = 1$ or $\lambda_C((0, 1, 1)) = 1$. We may assume without loss of generality that p is of the form

$$sF_3^*e_1F_3^*e_2F_3^*t \quad (3)$$

where F^* stands for an arbitrary number of edges from the set F .

Analogously, the graph $G(C, (0, 1, 0))$ contains an s - t -path q of the form

$$sF_2^*e_1F_2^*e_3F_2^*t \quad (4)$$

or

$$sF_2^*e_3F_2^*e_1F_2^*t. \quad (5)$$

Combining (3) and (5) yields an s - t -path in $(V, E_1 \cup F_2 \cup F_3)$, a contradiction to $\lambda_C((0, 1, 1)) = 0$. Therefore q must be of the form (4).

Finally there is also an s - t -path in $G(C, (1, 0, 0))$, which looks like

$$sF_1^*e_2F_1^*e_3F_1^*t \quad (6)$$

or

$$sF_1^*e_3F_1^*e_2F_1^*t. \quad (7)$$

Now an argument analogous to that for (3) and (5) shows that (6) conflicts with (3) and (7) with (5). A contradiction. \square

Comparison of Figure 6 and Lemma 3 points out again the discussed asymmetry between E -sets and F -sets. Had we chosen the alternative notion of parity in the defining relation (1), the roles of E -sets and F -sets in Lemma 3 and the Figure would of course be interchanged. However, though the proof of Proposition 1 will be based on Lemma 3, the proposition itself is obviously independent of the precise notion of parity because the statement is symmetric in the F_i and E_i .

Hardwired switches. The crucial idea for the proof of Proposition 1 is to fix several inputs of an n -switching circuit to 0-position so that they cannot be turned on again. The resulting circuit can then be considered as one with fewer inputs. Fixing all but three inputs allows application of Lemma 3, from which we will gain sufficient information about the original n -circuit.

This procedure demands that we define precisely how to *fix* an input of a switching circuit. We have to turn all active interruptors of that input into invariable wire. But then the two adjacent wire segments—which are simply vertices in our model—have to be merged. Phrased in terms of graph theory, we have to *contract* edges.

Contracting an edge $e = \{a, b\}$ of a graph means to merge the two vertices a and b into a single vertex and to remove the edge e . All edges that were incident to either a or b become incident to the new vertex (see [3, §12]). One can visualise




Figure 7: Hardwiring inputs.

contraction as a continuous process in which we shrink the edge $\{a, b\}$ more and more until both vertices amalgamate to a new vertex. Note that such contractions may create parallel edges and loops, which was our main reason to introduce these notions.

It is an easy exercise to show that a sequence of contraction steps on a graph is independent of the order of those operations. So we may legitimately speak of contracting whole sets of edges, not worrying about the precise order in which the individual steps are performed.

It should be clear now how to fix inputs of a switching circuit C . We *hardwire* the j th input of an n -switching circuit to 0 by contracting all edges in the set E_j . More precisely, we contract all E_j -edges in the graph $(V, \bigcup_{i=1}^n E_i \cup \bigcup_{i=1}^n F_i)$, always remembering for each edge to which set E_i or F_i it belongs. Afterwards we simply forget about the edge sets E_j and F_j (the former of which is now empty). This leaves us with a modified $(n - 1)$ -circuit which we denote by C_j . We restrict our attention to switches fixed in 0-position because we shall only need this case, hardwiring to 1 works just the same.

In analogy to single contractions, hardwiring several inputs is independent of the order of the inputs, so we may write C_I when we hardwire a whole set $I \subseteq \{1, \dots, n\}$ of inputs. Figure 7 shows two successive hardwiring steps performed on the circuit T from Figure 6.

The fact that all hardwired inputs of a circuit C_I behave like switches in 0-position can be expressed through the obvious equation

$$\lambda_{C_{\{k+1, k+2, \dots, n\}}}((d_1, d_2, \dots, d_k)) = \lambda_C((d_1, d_2, \dots, d_k, 0, \dots, 0)). \quad (8)$$

The concept of hardwiring switches allows us to prove Proposition 1 in a very elegant and succinct way.

Proof of Proposition 1. Assume for contradiction that there is a circuit with three indices i that satisfy $|E_i| = 1$ or $|F_i| = 1$. Without loss of generality we may assume that they are the first three.

For each index $i \leq 3$ with $|F_i| = 1$, we apply the exchange lemma to the pair $(i, 4)$. Then we have $|E_1| = |E_2| = |E_3| = 1$. Call the resulting circuit C . Now consider the switching circuit $C_{\{4, 5, \dots, n\}}$ obtained from C by fixing all but the first three inputs to 0. By (8) this is a parity circuit because C is. A contradiction to Lemma 3. \square

Conclusion. With Proposition 1 for the case $n \geq 4$, Lemma 3 for $n = 3$, and the fact that n -parity circuits have no empty edge sets for $n \geq 2$, we can now prove the standard setup optimal. Let us put things together.

Proof of Theorem 1. The case $n = 1$ was clear and $p(2) = 4$ followed from Lemma 2. Proposition 1 together with Lemma 2 shows $p(n) \geq 4n - 4$ for $n \geq 4$. So it only remains to show $p(3) \geq 8$. But this is quite simple now.

Assume for contradiction that there is a 3-parity circuit C of size less than 8. By Lemma 2 there is at most one E_i or one F_i with two elements, all other edge sets being singletons. The latter case is in direct contradiction to Lemma 3; and with one application of the exchange lemma the former also is. \square

6 Possible Extensions

Since our actual task has been solved, one might think about possible variations of parity circuits.

An immediate idea is to replace s - t -reachability in the definition of the induced function λ_C by some other graph property. For example, our light might depend on the existence of a Hamilton cycle in the d -graph. A *cycle* in a graph $G = (V, E)$ is a closed path in G , i.e., a sequence v_0, v_1, \dots, v_r of vertices with $\{v_i, v_{i+1}\} \in E$ for $0 \leq i < r$, and v_0 through v_{r-1} all distinct from each other, but $v_r = v_0$. A *Hamilton cycle* in G is a cycle that contains all vertices of G .

So we consider the function

$$\lambda_C^{\text{HAM}}(d) = \begin{cases} 1 & \text{if } G(C, d) \text{ contains a Hamilton cycle,} \\ 0 & \text{if } G(C, d) \text{ contains no Hamilton cycle} \end{cases}$$

and try to construct n -switching circuits of small size that induce the parity function under these changed conditions. (Here the vertices s and t have no meaning.)

We have to admit that the practical relevance of such circuits is rather questionable. Finding a Hamilton cycle in a graph is a hard problem. An elementary physical device that turns a light on if some graph contains a Hamilton cycle and turns it off if not, would represent a devastating scientific breakthrough. Hence today, Hamilton-cycle-sensitive lights are mere science fiction. Albeit, reality is no impediment for mathematical curiosity. So what can be said about the effect of alternative graph properties on parity circuits?

It is not clear whether such switching circuits exist at all. But, at least for our example they do. Together with Emo Welzl, we found n -parity circuits of size $4n$ for the Hamilton-cycle setting, making use of four-way switches again. And once having seen such a construction, one easily derives similar circuits for other properties like planarity or bipartiteness.

It remains an open question whether for each nontrivial graph property P and each $n \geq 1$, there exists an n -switching circuit C whose induced function λ_C^P is the parity function. If the answer is ‘no’, it might be worth seeking for a classification of those properties that have such circuits and of those who have not.

Acknowledgements

We thank Phil Kingery for his informations about the history of the terms ‘three-’ and ‘four-way switch’, and Mark de Longueville for helpful comments on the representation of our results.

References

- [1] Béla Bollobás. *Graph Theory*. Springer, 1979.
- [2] Frank Harary. *Graph Theory*. Addison-Wesley, second edition, 1971.
- [3] Robin J. Wilson. *Introduction to Graph Theory*. Academic Press, 1972.