An NP-Complete Problem in Grid Coloring
by William Gasarch and Kevin Lawler

1 Introduction

On November 30, 2009 the following challenge was posted on Complexity Blog [2].

BEGIN EXCERPT OF POST

The 17×17 challenge: worth \$289.00. I am not kidding.

Def 1.1 The $n \times m$ grid is *c-colorable* if there exists a way to c-color the vertices of the $n \times m$ grid so that there is no rectangle with all four corners the same color. (The rectangles I care about have sides parallel to the x and y axis.)

The 17×17 **challenge:** The first person to email me a 4-coloring of 17×17 in LaTeX will win \$289.00.
END EXCERPT OF POST

(The reader can see the post or the paper by Fenner, Gasarch, Purewall [1] for the motivation for this problem.)

Initially there was a lot of activity on the problem. Some used SAT solvers, some used linear programming, and one person offered an exchange: *buy me a \$5000 computer and I'll solve it*. Finally in 2012 Bernd Steinbach and Christian Posthoff [4] solved it to get the money and the glory! They used a rather clever algorithm with a SAT solver. They also commented that this may be close to the limits of their techniques in terms of size of grids and number of colors.

Between the problem being posed and resolved the following challenge was posted [3] though with no cash prize. We paraphrase the post

BEGIN PARAPHRASE OF POST

Def 1.2 Let $c, N, M \in \mathbb{N}$.

1. A mapping χ of $N \times M$ to $\{1, \ldots, c\}$ is a *c-coloring* if there are no monochromatic rectangles.

2. A partial mapping χ of $N \times M$ to $\{1, \ldots, c\}$ is a *extendable to a c-coloring* if there is an extension of χ to a total mapping which is a c-coloring of $N \times M$. We will use the term *extendable* if the c is understood.

Def 1.3 Let

$$GCE = \{(N, M, c, \chi) \mid \chi \text{ is extendable}\}.$$

GCE stands for *Grid Coloring Extension.*

CHALLENGE: Prove that GCE is NP-complete.
END PARAPHRASE OF POST

In Section 2 we show that GCE is indeed NP-complete!! This result may explain why the original 17×17 challenge was so difficult. Then again—it may not. In Section 4 we show that GCE is fixed-parameter tractable. Hence, for a fixed c, the problem might not be that hard. In Section 5 we state some open problems.

There is another reason the results obtained may not be the reason why the 17×17 challenge was hard. The 17×17 challenge can be rephrased as proving that $(17, 17, 4, \chi) \in GCE$ where χ is the empty partial coloring. This is a rather special case of GCE since *none* of the spots are pre-colored. It is possible that GCE in the sppecial case where χ is the emtpy coloring is easy. While we doubt this is true, we note that we have not eliminated this possibility.

2 GCE is NP-complete

Theorem 2.1 *GCE is NP-complete.*

Proof:

Clearly $GCE \in$ NP.

Let $\phi(x_1, \ldots, x_n) = C_1 \wedge \cdots \wedge C_m$ be a 3-CNF formula. We determine N, M, c and a partial c-coloring χ of $N \times M$ such that

$$(N, M, c, \chi) \in GCE \text{ iff } \phi \in \text{3-SAT}.$$

The grid will be thought of as a main grid and some stuff at the left side and below, which is only there to enforce that some of the colors in the main grid occur only once. The colors will be T, F, and some of the $(i, j) \in N \times M$.

The construction is in four parts. We summarize the four parts here before going into details.

1. We will often need to define $\chi(i, j)$ to be (i, j) and to then never have the color (i, j) appear in any other cell of the main grid. We show how to color the cells that are not in the main grid to achieve this. While we show this first, it is actually the last step of the construction.

2. The main grid will have $2nm + 1$ rows. In the first column we have $2nm$ blank spaces and then space $(1, 2nm + 1)$ colored with $(1, 2nm + 1)$. The $2nm$ blank spaces will be forced to be colored T and F. We think of the column as being in n blocks of $2m$ each.

In the ith block the coloring will be forced to be

$$\begin{array}{c} T \\ F \\ \vdots \\ T \\ F \end{array}$$

if x_i is to be set to T, and

$$\begin{array}{c} F \\ T \\ \vdots \\ F \\ T \end{array}$$

if x_i is to be set to F.

3. For each clause C there will be two columns. The coloring χ will be defined on most of the cells in these columns. However, the coloring will extend to these two columns iff one of the literals in C is colored T in the first column.

4. We set the number of colors properly so that the T and F will be forced to be used in all blank spaces.

1) **Forcing a color to appear only once in the main grid.**

Say we want the cell $(2,4)$ in the main grid to be colored $(2,4)$ and we do not want this color appearing anywhere else in the main grid. We can do the following: add a column of $(2,4)$'s to the left end (with one exception) and a row of $(2,4)$'s below. Here is what we get:

$(2,4)$								
$(2,4)$								
T		$(2,4)$						
$(2,4)$								
$(2,4)$								
$(2,4)$								
$(2,4)$	$(2,4)$	$(2,4)$	$(2,4)$	$(2,4)$	$(2,4)$	$(2,4)$	$(2,4)$	$(2,4)$

(The double lines are not part of the construction. They are there to separate the main grid from the rest.)

It is easy to see that in any coloring of the above grid the only cells that can have the color $(2,4)$ are those shown to already have that color. It is also easy to see that the color T we have will not help to create any monochromatic rectangles since there are no other T's in its column. The T we are using *is* the same T that will later mean TRUE. We could have used F. If we used a new special color we would need to be concerned whether there is a monochromatic grid of that color. Hence we use T.

What if some other cell needs to have a unique color? Lets say we also want to color cell $(5,3)$ in the main grid with $(5,3)$ and do not want to color anything else in the main grid $(5,3)$. Then we do the following:

$(5,3)$	$(2,4)$								
$(5,3)$	$(2,4)$								
$(5,3)$	T		$(2,4)$						
T	$(2,4)$					$(5,3)$			
$(5,3)$	$(2,4)$								
$(5,3)$	$(2,4)$								
$(5,3)$	$(2,4)$	$(2,4)$	$(2,4)$	$(2,4)$	$(2,4)$	$(2,4)$	$(2,4)$	$(2,4)$	$(2,4)$
$(5,3)$	$(5,3)$	$(5,3)$	$(5,3)$	$(5,3)$	$(5,3)$	$(5,3)$	$(5,3)$	$(5,3)$	$(5,3)$

It is easy to see that in any coloring of the above grid the only cells that can have the color $(2,4)$ or $(5,3)$ are those shown to already have those colors.

For the rest of the construction we will only show the main grid. If we denote a color as D (short for *Don't Care*) in the cell (i,j) then this means that (1) cell (i,j) is color (i,j) and (2) we have used the above gadget to make sure that (i,j) does not occur as a color in any other cell of the main grid. Note that we when we have D in the $(2,4)$ cell and in the $(5,3)$ cell they denote different colors.

2) **Forcing $(x,\overline{x})$ to be colored (T,F) or (F,T).**

There will be one column with cells labeled by literals. The cells are blank, uncolored. We will call this row *the literal column*. We will put to the left of the literal column, separated by a double line, the literals whose values we intend to set. These literals are not part of the construction; they are a visual aid for you, the reader. The color of the literal-labeled cells will be T or F. We need to make sure that all of the x_i have the same color and that it is different than that of $\overline{x}_i$.

Here is an example which shows how we can force $(x_1,\overline{x}_1)$ to be colored (T,F) or (F,T).

$\overline{x}_1$		T	F
x_1		T	F

We will actually need m copies of x_1 and m copies of $\overline{x}_1$. We will also put a row of D's on top which we will use later. We illustrate how to do this in the case of $m=3$.

	D	D	D	D	D	D	D	D	D	D	D
$\overline{x}_1$		D	D	D	D	D	D	D	D	T	F
x_1		D	D	D	D	D	D	T	F	T	F
$\overline{x}_1$		D	D	D	D	T	F	T	F	D	D
x_1		D	D	T	F	T	F	D	D	D	D
$\overline{x}_1$		T	F	T	F	D	D	D	D	D	D
x_1		T	F	D	D	D	D	D	D	D	D

We leave it as an exercise to prove that

- If the bottom x_1 cell is colored T then (1) all of the x_1 cells are colored T, and (2) all of the $\overline{x}_1$ cells are colored F.
- If the bottom x_1 cell is colored F then (1) all of the x_1 cells are colored F, and (2) all of the $\overline{x}_1$ cells are colored T.

Note that (1) if we want literal-pair (that is $x_1, \overline{x}_1$) then we use two columns, (2) if we want two literal-pairs then we use six columns, and (2) if we want three literal-pairs then we use ten columns, We leave it as an exercise to generalize the construction to m literal-pairs using $2 + 4(m-1)$ columns.

We will need m copies of x_2 and m copies of $\overline{x}_2$. We illustrate how to do this in the case of $m = 2$. We use double lines in the picture to clarify that the x_1 and the x_2 variables are not chained together in any way.

	D	D	D	D	D	D	D	D	D	D	D	D	D
$\overline{x}_2$		D	D	D	D	D	D	D	D	D	D	T	F
x_2		D	D	D	D	D	D	D	D	T	F	T	F
$\overline{x}_2$		D	D	D	D	D	D	T	F	T	F	D	D
x_2		D	D	D	D	D	D	T	F	D	D	D	D
$\overline{x}_1$		D	D	D	D	T	F	D	D	D	D	D	D
x_1		D	D	T	F	T	F	D	D	D	D	D	D
$\overline{x}_1$		T	F	T	F	D	D	D	D	D	D	D	D
x_1		T	F	D	D	D	D	D	D	D	D	D	D

We leave it as an exercise to prove that, for all $i \in \{1, 2\}$:

- If the bottom x_i cell is colored T then (1) all of the x_i cells are colored T, and (2) all of the $\overline{x}_1$ cells are colored F.
- If the bottom x_i cell is colored F then (1) all of the x_i cells are colored F, and (2) all of the $\overline{x}_1$ cells are colored T.

An easy exercise for the reader is to generalize the above to a construction with n variables, and for each one we get m literal-pairs. This will take $n(2 + 4(m-1))$ columns.

For the rest of the construction we will only show the literal column and the clause columns (which we define in the next part). It will be assumed that the D's and T's and F's are in place to ensure that all of the x_i cells are one of $\{T, F\}$ and the $\overline{x}_i$ cells are the other color.

3) How we force the coloring to satisfy ONE clause

Say one of the clauses is $C_1 = x_1 \vee x_2 \vee \overline{x}_3$. Pick an x_1 row, an x_2 row, and an $\overline{x}_3$ row. We will also use the top row, as we will see. For other clauses you will pick other rows. Since there are m copies of each variable and its negation this is easy to do.

The two T's in the top row in the next picture are actually in the very top row of the grid.

We put a C_1 over the columns that will enforce that C_1 is satisfied. These C_1 are not part of the grid. They are there as a visual aid for us.

		C_1	C_1
	D	T	T
$\overline{x}_3$		D	F
x_2			
x_1		F	D

Keep in mind that the cells labeled $x_1, x_2, \overline{x}_3$ are not colored— they are blank, uncolored.
Claim 1: If χ' is a 2-coloring of the blank spots in this grid (with colors T and F) then it CANNOT have the $x_1, x_2, \overline{x}_3$ spots all colored F.
Proof of Claim 1:

Assume, by way of contradiction, that that $x_1, x_2, \overline{x}_3$ are all colored F. Then this is what it looks like:

		C_1	C_1
	D	T	T
$\overline{x}_3$	F	D	F
x_2	F		
x_1	F	F	D

The two blank spaces are both FORCED to be T since otherwise you get a monochromatic rectangle of color F. Hence we have

		C_1	C_1
	D	T	T
$\overline{x}_3$	F	D	F
x_2	F	T	T
x_1	F	F	D

This coloring has a monochromatic rectangle which is colored T. This contradicts χ' being a 2-coloring of the blank spots.
End of Proof of Claim 1

We leave the proof of Claim 2 below to the reader.
Claim 2: If χ' colors $x_1, x_2, \overline{x}_3$ anything except F, F, F then χ' can be extended to a coloring of the grid shown.
Upshot: A 2-coloring of the grid is equivalent to a satisfying assignment of the clause.

Note that each clause will require 2 columns to deal with. So there will be $2m$ columns for this.

4) **Putting it all together**

Recall that $\phi(x_1, \ldots, x_n) = C_1 \wedge \cdots \wedge C_m$ is a 3-CNF formula. We can assume that no clause contains both a variable and its negation.

We first define the main grid and later define the entire grid and N, M, c.

The main grid will have $2nm + 1$ rows and $n(4m - 2) + 2m + 1$ columns. The first $n(4m - 2) + 1$ columns are partially colored using the construction in Part 2. This will establish the literal column. We will later set the number of colors so that the literal column must use the colors T and F.

For each of the m clauses we pick a set of its literal from the literals column. These sets-of-literals are all disjoint. We can do this since we have m copies of each literal-pair. We then do the construction in Part 3. Note that this uses two columns. Assuming that all of the D's are colored distinctly and that the only colors left are T and F, this will ensure that the main grid is c-colorable iff the formula is satisfiable.

The main grid is now complete. For every (i, j) that is colored (i, j) we do what what is in Part 1 to make sure that (i, j) is the only cell with color (i, j). Let the number of such (i, j) be C. The number of colors c is $C + 2$. ∎

3 An Example

We can make the construction slightly more efficient (and thus can actually work out an example). We took m pairs $\{x_i, \overline{x}_i\}$. We don't really need all m. If x_i appears in a clauses and $\overline{x}_i$ appears in b clauses then we only need $\max\{a, b\}$ literal-pairs. If $a \neq b$ then we only need $\max\{a, b\} - 1$ literal-pairs and one additional literal. (This will be the case in the example below.)

With this in mind we will do an example- though we will only show the main grid.

$$(x_1 \vee x_2 \vee \overline{x}_3) \wedge (\overline{x}_2 \vee x_3 \vee x_4) \wedge (\overline{x}_1 \vee \overline{x}_3 \vee \overline{x}_4)$$

We only need

- one $(x_1, \overline{x}_1)$ literal-pair,
- one $(x_2, \overline{x}_2)$ literal-pair,
- one $(x_3, \overline{x}_3)$ literal-pair,
- one additional $\overline{x}_3$,
- one $(x_4, \overline{x}_4)$ literal-pair.

												C_1	C_1	C_2	C_2	C_3	C_3
	D	D	D	D	D	D	D	D	D	D	D	T	T	T	T	T	T
$\overline{x}_4$		D	D	D	D	D	D	D	D	T	F	D	D	D	D	D	F
x_4		D	D	D	D	D	D	D	D	T	F	D	D	D	F	D	D
$\overline{x}_3$		D	D	D	D	D	D	T	F	D	D	D	D	D	D	D	D
x_3		D	D	D	D	T	F	T	F	D	D	D	D			D	D
$\overline{x}_3$		D	D	D	D	T	F	D	D	D	D	D	F	D	D		
$\overline{x}_2$		D	D	T	F	D	D	D	D	D	D	D	D	F	D	D	D
x_2		D	D	T	F	D	D	D	D	D	D			D	D	D	D
$\overline{x}_1$		T	F	D	D	D	D	D	D	D	D	D	D	D	D	F	D
x_1		T	F	D	D	D	D	D	D	D	D	F	D	D	D	D	D

4 Fixed Parameter Tractability

The 17×17 problem only involved 4-colorability. Does the result that GCE is NP-complete really shed light on the hardness of the 17×17 problem? What happens if the number of colors is fixed?

Def 4.1 Let $c \in \mathbb{N}$. Let

$$GCE_c = \{(N, M, \chi) \mid \chi \text{ can be extended to a } c\text{-coloring of } N \times M\}.$$

Clearly $GCE_c \in DTIME(c^{O(NM)})$. Can we do better? Yes. We will show that GCE is in time $O(N^2M^2 + c^{O(c^4)})$.

Notation 4.2 If $x \in \mathbb{N}$ then $[x]$ denotes the set $\{1, \ldots, x\}$. If X is a set then $\binom{X}{2}$ is the set of unordered pairs from X.

Lemma 4.3 *Assume $c + 1 \le N$ and $c\binom{c+1}{2} < M$. Then $N \times M$ is not c-colorable. Hence, for any χ, $(N, M, \chi) \notin GCE_c$.*

Proof: Assume, by way of contradiction, that there is a c-coloring of $N \times M$. Since every column has at least $c+1$ elements the following mapping is well defined: Map every column to the least $(\{i, j\}, a)$ such that the $\{i, j\} \in \binom{[c+1]}{2}$ and both the ith and the jth row of that column are colored a. The range of this function has $c\binom{c+1}{2}$ elements. Hence some element of the range is mapped to at least twice. This yields a monochromatic rectangle. ∎

Lemma 4.4 *Assume $N \le c$ and $M \in \mathbb{N}$. If χ is a partial c-coloring of $N \times M$ then $(N, M, \chi) \in GCE_c$.*

Proof: The partial c-coloring χ can be extended to a full c-coloring as follows: for each column use a different color for each blank spot, making sure that all of the new colors in that column are different from each other. ▮

Lemma 4.5 *There is an algorithm that will, given (N, M, χ), determine if χ is a partial c-coloring of $N \times M$ in time $O(N^2M^2)$.*

Proof: The algorithm looks at every $\{i_1, i_2\} \in \binom{[N]}{2}$ and every $\{j_1, j_2\} \in \binom{[M]}{2}$ to see if the rectangle $(i_1, j_1), (i_1, j_2), (i_2, j_1), (i_2, j_2)$ is monochromatic. ▮

Theorem 4.6 *$GCE_c \in DTIME(N^2M^2 + c^{O(c^6)})$ time.*

Proof:

1. Input (N, M, χ).
2. If $N \le c$ or $M \le c$ then test if χ is a partial c-coloring of $N \times M$. If so then output YES. If not then output NO. (This works by Lemma 4.4.) This takes time $O(N^2M^2)$.
3. If $c + 1 \le N$ and $c\binom{c+1}{2} < M$ then output NO and stop. (This works by Lemma 4.3.)
4. The only case left is $c+1 \le N, M \le c\binom{c+1}{2}$. Look at all possible extensions of χ to see if one of them is a c-coloring. This will take $O(c^{NM}) = O(c^{O(c^6)})$ time.

▮

Can we do better? Yes, but it will require a result from [1].

Lemma 4.7 *Let $1 \le c' \le c - 1$.*

1. *If $N \ge c + c'$ and $M > \frac{c}{c'}\binom{c+c'}{2}$ then $G_{N,M}$ is not c-colorable.*
2. *If $N \ge 2c$ and $M > 2\binom{2c}{2}$ then $G_{N,M}$ is not c-colorable. (This follows from a weak version of the $c' = c - 1$ case of Part 1.)*

Theorem 4.8 *$GCE_c \in DTIME(N^2M^2 + c^{O(c^4)})$ time.*

Proof:

1. Input (N, M, χ).
2. If $N \le c$ or $M \le c$ then test if χ is a partial c-coloring of $N \times M$. If so then output YES. If not then output NO. (This works by Lemma 4.4.) This takes time $O(N^2M^2)$.

3. For $1 \leq c' \leq c-1$ we have the following pairs of cases.
   (a) $N = c + c'$ and $M > \frac{c}{c'}\binom{c+c'}{2}$ then output NO and stop. (This works by Lemma 4.7.)
   (b) $N = c+c'$ and $M \leq \frac{c}{c'}\binom{c+c'}{2}$ then look at all possible extensions of χ to see if one of them is a c-coloring. This will take $c^{O(NM)}$. Note that $MN \leq (c+c')\frac{c}{c'}\binom{c+c'}{2}$. On the interval $1 \leq c' \leq c-1$ this function achieves its maximum when $c' = 1$. Hence this case takes $c^{O(c^4)}$.
4. If $N \geq 2c$ and $M > 2\binom{2c}{2}$ then output NO and stop. (This works by Lemma 4.7.)
5. The only case left is $2c \leq N, M \leq 2\binom{2c}{2}$. Look at all possible extensions of χ to see if one of them is a c-coloring. This will take $c^{O(NM)} = c^{O(c^4)}$ time.

∎

Even for small c the additive term $c^{O(c^4)}$ is the real timesink. A cleverer algorithm that reduces this term is desirable. By Theorem 2.1 this term cannot be made polynomial unless P=NP.

5 Open Problems

In this paper we looked at colorings that avoid *rectangles*. What about other shapes? We give an equivalent definition of the rectangles problem in a way that can be generalized.

Def 5.1 Let $c, N, M \in \mathbb{N}$. A (full or partial) mapping of $N \times M$ to $\{1, \ldots, c\}$ is a *c-coloring* if there does not exists a set of points $\{(a,b), (a+t,b), (a,b+s), (a+t,b+s)\}$ that are all the same color.

Look at the points

$$\{(a,b), (a+t,b), (a,b+s), (a+t,b+s)\}.$$

We can view them as

$$\{(s\times 0, t\times 0)+(a,b), (s\times 1, t\times 0)+(a,b), (s\times 0, t\times 1)+(a,b), (s\times 1, t\times 1)+(a,b)\}.$$

Informally, the set of rectangles is generated by $\{(0,0),(0,1),(1,0),(1,1)\}$. Formally we can view the set of rectangles on the lattice points of the plane (upper quadrant) as the intersection of $\mathbb{N} \times \mathbb{N}$ with

$$\bigcup_{s,t,a,b\in\mathbb{Q}} \{\{(s\times 0, t\times 0)+(a,b), (s\times 1, t\times 0)+(a,b), (s\times 0, t\times 1)+(a,b), (s\times 1, t\times 1)+(a,b)\}\}$$

Note that the pair of curly braces is not a typo. We are looking at sets of 4-sets of points. We generalize rectangles.

Def 5.2 Let
$$S = \{(x_1, y_1), \ldots, (x_L, y_L)\}$$
be a set of lattice points in the plane. Let
$$\text{stretch}(S) = \bigcup_{s,t \in \mathsf{Q}} \{\{(sx_1, ty_1), \ldots, (sx_L, ty_L)\}\}$$
and
$$\text{translate}(S) = \bigcup_{a,b \in \mathsf{Q}} \{\{(x_1 + a, y_1 + b), \ldots, (x_L + a, y_L + b)\}\}.$$
These are the sets of points we will be trying to avoid making monochromatic. Hence let
$$\text{avoid}(S) = \text{translate}(\text{stretch}(S)).$$

We can now generalize the rectangle problem.

Def 5.3 Let $N, M \in \mathbb{N}$ and S be a set of lattice points. A (partial or full) mapping χ from $N \times M$ into $\{1, \ldots, c\}$ is a (c, S)-*coloring* if there are no monochromatic sets in avoid(S).

Open Problem 1: For which sets of lattice points S is the following problem NP-complete?
$$\{(N, M, c, \chi) \mid \chi \text{ can be extended to a } (c, S)\text{-coloring of } N \times M\}$$

One could look at other ways to move the points in S around. There is one we find particular interesting. We motivate our definition.

What if we wanted to look at colorings that avoided a monochromatic *square*? The square
$$\{(a, b), (a + s, b), (a, b + s), (a + s, b + s)\}$$
can be viewed as
$$\{(0, 0) + (a, b), (s, 0) + (a, b), (0, s) + (a, b), (s, s) + (a, b)\}.$$
We generalize this.

Def 5.4 Let
$$S = \{(x_1, y_1), \ldots, (x_L, y_L)\}$$
be a set of lattice points in the plane. Let
$$\text{halfstretch}(S) = \bigcup_{s \in \mathsf{Q}} \{\{(sx_1, sy_1), \ldots, (sx_L, sy_L)\}\}$$
These are the sets of points we will be trying to avoid making monochromatic. We would like to call it avoid but that name has already been taken; hence we call it avoid$_2$.
$$\text{avoid}_2(S) = \text{translate}(\text{halfstretch}(S)).$$
(Note that the 2 has no significance. It is just there to distinguish avoid and avoid$_2$.)

Def 5.5 Let $N, M \in \mathbb{N}$ and S be a set of lattice points. A (partial or full) mapping χ from $N \times M$ into $\{1, \ldots, c\}$ is a $(c, S)_2$-*coloring* if there are no monochromatic sets in $\text{avoid}_2(S)$. (Note that the 2 has no significance. It is just there to distinguish (c, S) and $(c, S)_2$.)

Def 5.6 Let $N, M \in \mathbb{N}$ and S be a set of lattice points. A (partial or full) mapping χ from $N \times M$ into $\{1, \ldots, c\}$ is a $(c, S)_2$-*coloring* if there are no monochromatic sets in $\text{avoid}_2(S)$.

Open Problem 2: For which sets of lattice points S is the following problem NP-complete?

$$\{(N, M, c, \chi) \mid \chi \text{ can be extended to a } (c, S)_2\text{-coloring of } N \times M\}$$

Open Problem 3: Improve our FPT algorithm. Develop an FPT algorithm for the variants we have discussed.

Open Problem 4: Prove that grid coloring problems starting with the empty grid are hard. This may need a new formalism.

6 Acknowledgments

We would like to thank Daniel Apon and Clyde Kruskal for proofreading and discussion. We would also like to thank Richard Chang, Peter Fontana, David Harris, and Marius Zimand for listening to Bill present the NP-completeness result and hence clarifying it.

References

[1] S. Fenner, W. Gasarch, C. Glover, and S. Purewal. Rectangle free colorings of grids, 2009. `http://www.cs.umd.edu/~gasarch/papers/papers.html`.

[2] W. Gasarch. The 17×17 challenge. Worth $289.00. This is not a joke, `http://blog.computationalcomplexity.org/2009/11/17x17-challenge-worth-28900-this-is-not.html`, 2009.

[3] W. Gasarch. A possible NP-intermediary problem. `http://blog.computationalcomplexity.org/2010/04/possible-np-intermediary-problem.html`, 2010.

[4] B. Steinbach and C. Posthoff. Most complex four-colored rectangle-free grids– solution of an open multiple-valued problem. In *Proceedings of the Forty-Second IEEE International Symposia on Multiple-Valued Logic*, 2012.