6(a) In any partial tiling of the board, each unit rhombus covers one triangle facing up and one facing down, so that the total number of covered triangles facing in each direction is equal. Since the board contains $\binom{n+1}{2}$ triangles facing up and $\binom{n}{2}$ facing down and we may only avail ourselves of $\binom{n+1}{2} - \binom{n}{2} = n$ holes, all of these holes must fall on triangles facing up.

Now, consider the bipartite graph G whose left vertex set is the set of all triangles facing down, whose right vertex set is the set of all triangles facing up, and with an edge between two vertices if and only if the corresponding triangles are adjacent. Then each edge of G corresponds to a possible placement of a unit rhombus in a partial triangle, and two rhombi are nonoverlapping just if their corresponding edges are non-incident. So a transversal in G corresponds to a placement of maximally many nonoverlapping rhombi on the triangular board. It is possible to fit $\binom{n}{2}$ rhombi on the board, leaving only n holes, for instance by placing the rhombi all in the same orientation. Therefore, those coverings by $\binom{n}{2}$ rhombi and n holes give transversals, and the bases in the transversal matroid $\mathcal{M}$ given by G are the sets of all triangles facing up that *are* covered in some tiling with n holes.

Finally, the good sets of holes are the sets of triangles left uncovered in a maximal tiling, so they're all the complements of the above bases. The set of good sets of holes is therefore a matroid, namely the dual $\mathcal{M}^*$ of $\mathcal{M}$.

(b) To generalize the result to an arbitrary board we have to decide what n should be. We abstract from part (a) as the relevant property that in a tiling of the triangular board using the maximum possible number of unit rhombi, there are n holes, and so on general boards will look the holes in tilings accommodating the maximum number of unit rhombi. For brevity we'll call these tilings *maximum tilings*.

We observe that the only property of the board we used in part (a) is that all triangles pointing down are covered in every maximum tiling, and therefore we can generalize the result of that part to every such board. We can of course do the same if every maximum tiling covers every triangle pointing up.

Now, take a general board B and a tiling T of B using the maximum number of rhombi which it accommodates. We label some of the rhombi in T 'up', such that every rhombus whose downward-pointing half is adjacent to a (upward-pointing) hole or to a rhombus labelled 'up' is also labelled 'up', and we label the smallest set possible subject to these conditions: therefore the set of rhombi labelled 'up' contains all those rhombi for which there is a path in B from the upward-pointing side to an upward-pointing hole, passing through both halves of any rhombus it touches. We do the same for the label 'down', switching the roles of up and down.

Now we claim that after this process, no single rhombus is labelled both 'up' and 'down'. For if there was such a rhombus r, then by our observation on paths, there is a path in B from an upward-pointing hole to a downward-pointing hole, passing through r, and passing through both halves of every rhombus it touches. But we can then produce a new tiling T' by removing all the rhombi on this path and filling it again, starting by filling the hole at one end; for parity reasons this covers all the triangles on this path including both holes. But that gives us a tiling of B including more rhombi than T, a contradiction.

Therefore we have our board B partitioned into areas tiled with rhombi with a single labelling, either 'up', 'down', or nothing. Transfer these labels to the underlying cells of B, and also label all holes of B with their orientation; this divides B into regions each labelled 'up', 'down', or nothing. Any region labelled 'up' is covered in T by rhombi and upward-pointing holes, and similarly for 'down', while any unlabelled region is completely covered by rhombi in T. Furthermore, by construction of the labelling, every triangle on the outside periphery of a region labelled 'up' (i.e. not inside the region but adjacent) must be pointing up, and similarly for 'down'.

Now look at any single one of these regions R of B independently. Suppose that in T it is tiled with n holes. We claim it's impossible to tile R with fewer than n holes even if we permit the use of rhombi crossing the boundary of R into the rest of B. If R is unlabelled, this is trivial, since $n = 0$. Otherwise

R is labelled 'up' or 'down'; in these cases we make the stronger claim that if any tile is laid across the boundary of R, we must leave strictly more than n holes. Assume wlog that R is labelled 'up' (the 'down' case will be identical up to exchange of up and down). Then all n holes in T are pointing up. Now if we lay tiles across the boundary of R, they cover downward-pointing triangles on the inside of R; but this makes the disparity between the number of upward and downward triangles greater, and means we cannot possibly leave as few as n holes.

As a consequence, no maximum tiling of B contains any rhombus crossing a region boundary such that the region on one side is labelled 'up' or 'down'; but every region boundary abuts an 'up' or 'down' region on one side, so that no maximum tiling of B contains any tiles crossing region boundaries at all. Therefore, the maximum tilings on each region of B are completely independent. Since the sets of holes on each region R_i form the bases of a matroid M_i by our first observation, the sets of holes in tilings of B give the set of bases of the direct sum of these matroids M_i. This answers the problem's question in the positive.