CSCI3390-Lecture 9: Undecidability of the string-rewriting problem and the restricted tiling problem

September 28, 2018

Here we give fairly detailed proofs that the string-rewriting problem and the constrained tiling problem from the last lecture (tiling the first quadrant with a designated corner tile) are undecidable. As in all but one of the proofs of undecidability we've seen up until now, the proof is carried out by a reduction from a problem already known to be undecidable.

1 String rewriting

We show this by reduction from the first version of the halting problem. The idea itself is very simple: The sequence of configurations of a Turing machine calculation looks a lot like application of substitution rules in a string-rewriting system. For example, the following sequence (drawn from the first lecture notes)

$$\mathbf{q_0}001 \Rightarrow 0\mathbf{q_0}01 \Rightarrow 00\mathbf{q_0}1 \Rightarrow 001\mathbf{q_0}\square \Rightarrow 00\mathbf{q_1}1\# \Rightarrow 00X\mathbf{q_3}\# \Rightarrow 00x\#\mathbf{q_3}\square \Rightarrow 00x\mathbf{q_5}\#1$$

appears as though a substitution rule

$$\mathbf{q_0}0 \mapsto 0\mathbf{q}_1$$

was applied in the first two steps, and that

$$1\mathbf{q_0}\square \mapsto \mathbf{q_1}1\#$$

in the third step. So, roughly speaking, if you could tell everything about any string-rewriting system, you could tell everything about a Turing machine computation. We will flesh this idea out.

Deriving a string-rewriting system from a Turing machine specification is complicated by what happens with blank symbols at the end of the tape. To get around the issues this poses, given a Turing machine $\mathcal{M}$ we will first replace it by a very similar Turing machine $\mathcal{M}'$ that never writes a blank symbol, and thus in particular never erases a symbol. To do this, we introduce a new tape symbol B, which serves as an 'alternate blank'. We replace every transition of $\mathcal{M}$ of the form

$$\delta(q, \gamma) = (q', \square, D),$$

where $D = R$ or $D = L$, by

$$\delta(q, \gamma) = (q', B, D).$$

Further for every transition of the form

$$\delta(q, \square) = (q', \gamma, D),$$

we add a new transition

$$\delta(q, B) = (q', \gamma, D).$$

This modified TM $\mathcal{M}'$ does exactly what $\mathcal{M}$ does: In particular, if w is a word over the input alphabet Σ of $\mathcal{M}$ (and thus w does not contain any blank symbols), then $\mathcal{M}$ halts when started on w if and only if $\mathcal{M}'$ halts when started on w.

Having thus altered the machine $\mathcal{M}$, we now describe how to make substitution rules from the transitions of $\mathcal{M}'$:

- For each transition of the form

  $$\delta(q, \gamma) = (q', \gamma', R),$$

  where $\gamma \neq \square$, include the substitution rule

  $$\mathbf{q}\gamma \mapsto \gamma'\mathbf{q}'.$$

- For each transition of the form

  $$\delta(q, \square) = (q', \gamma', R),$$

  include the substitution rule

  $$\mathbf{q}\square \mapsto \gamma'\mathbf{q}'\square.$$

- For each transition of the form

$$\delta(q, \gamma) = (q', \gamma', L),$$

  where $\gamma \neq \Box$, and each tape symbol $\beta \neq \Box$ include the substitution rules

$$\beta \mathbf{q} \gamma \mapsto \mathbf{q}' \beta \gamma',$$

  as well as the rule

$$\Box \mathbf{q} \gamma \mapsto \Box \mathbf{q}' B \gamma'.$$

- For each transition rule of the form

$$\delta(q, \Box) = (q', \gamma', L),$$

  and each tape symbol $\beta \neq \Box$, include the substitution rules

$$\beta \mathbf{q} \Box \mapsto \mathbf{q}' \beta \gamma \Box,$$

  as well as the rule

$$\Box \mathbf{q} \Box \mapsto \Box \mathbf{q}' B \gamma' \Box.$$

- And, finally, if **h** is the halt state, for each tape symbol γ, add the substitution rules

$$\mathbf{h} \gamma \mapsto \mathbf{h}, \gamma \mathbf{h} \mapsto \mathbf{h}.$$

Our claim is that $\mathcal{M}$ halts when started on input $w \in \Sigma^*$, if and only if $\mathbf{h}$ can be derived from $\Box \mathbf{q}_0 w \Box$. Thus if we had an algorithm for the string-rewriting problem, we would have one for the halting problem.

To see how the substitution rules mirror the machine computation, consider the example of a Turing machine that moves to the right, reading past its inputs of a's and b's, until it hits the right end of the tape, at which point it writes a c. It then moves left, erasing all the original symbols on the tape, and halts. A typical run of the machine looks like this:

$$\mathbf{q}_0 ab \mapsto a \mathbf{q}_0 b \mapsto ab \mathbf{q}_0 \Box \mapsto a \mathbf{q}_1 bc \mapsto \mathbf{q}_1 a \Box c \mapsto \mathbf{q}_1 \Box\Box\Box c \mapsto \mathbf{h} \Box\Box c.$$

In the tweaked version of machine, the run looks like this:

$$\mathbf{q}_0 ab \mapsto a \mathbf{q}_0 b \mapsto ab \mathbf{q}_0 \Box \mapsto a \mathbf{q}_1 bc \mapsto \mathbf{q}_1 aBc \mapsto \mathbf{q}_1 \Box BBc \mapsto \mathbf{h} BBc.$$

We can use the resulting string-rewriting system to derive $\mathbf{h}$ from the word $\Box\mathbf{q}_0ab\Box$, as follows:

$$\begin{array}{rcl}
\Box\mathbf{q}_0ab\Box & \to & \Box a\mathbf{q}_0b\Box \\
& \to & \Box ab\mathbf{q}_0\Box \\
& \to & \Box a\mathbf{q}_1bc\Box \\
& \to & \Box\mathbf{q}_1aBc\Box \\
& \to & \Box\mathbf{q}_1BBBc\Box \\
& \to & \Box\mathbf{h}BBc\Box \\
& \to & \mathbf{h}BBc\Box \\
& \to & \mathbf{h}Bc\Box \\
& \to & \mathbf{h}c\Box \\
& \to & \mathbf{h}\Box \\
& \to & \mathbf{h}
\end{array}$$

You should observe that at each step (until the appearance of **h**) there was exactly one applicable substitution rule, determined by the Turing machine transition that was applied. Once **h** appears, we are able to erase all the other symbols. On the other hand, if the computation of the TM does not halt, then the symbol **h** will never appear, so we could not derive it.

2 Restricted Tiling problem

Here the undecidable problem from which we will reduce is that of determining whether a TM $\mathcal{M}$ with a one-way infinite tape eventually halts if it is started on a blank tape. (Undecidability of the halting problem with an initially blank tape was a homework problem, as was equivalence of one- and two-way infinite tapes.) That is, we will show, given a TM $\mathcal{M}$ with a one-way infinite tape, how to construct a set $T_{\mathcal{M}}$ of tiles with a designated corner tile $t \in T_{\mathcal{M}}$ such that $(T_{\mathcal{M}}, t)$ tiles the first quadrant if and only if $\mathcal{M}$ runs forever when started on an initially blank tape. Thus if we could decide this version of the tiling problem, we obtain the contradiction that we could decide an undecidable problem about Turing machines. (Strictly speaking, this is a reduction from the *complement* of the halting problem to the tiling problem, since the correspondence is 'doesn't halt $\leftrightarrow$ tiles the quadrant'.)

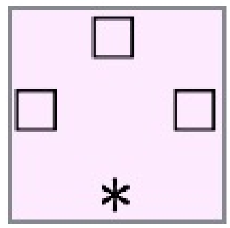

Figure 1: Type I tile.

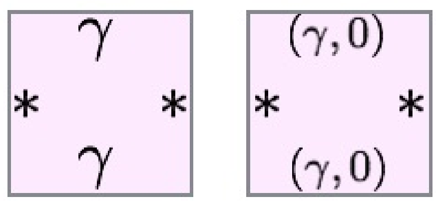

Figure 2: Type II tiles. There are two such tiles for every tape symbol γ.

2.1 Construction of $T_{\mathcal{M}}$.

The 'colors' of the tiles are derived from the states and tape symbols of $\mathcal{M}$. There are five types of tiles.

2.1.1 Type I

Our set $T_{\mathcal{M}}$ of tiles includes the tile shown in Figure 1.

2.1.2 Type II

For each letter γ in the tape alphabet of $\mathcal{M}$, $T_{\mathcal{M}}$ includes the tiles shown in Figure 2. Thus the total number of Type II tiles is $2|\Gamma|$, where Γ is the tape alphabet.

2.1.3 Type III

For each right-moving transition

$$\delta(p, \gamma) = (q, \beta, R)$$

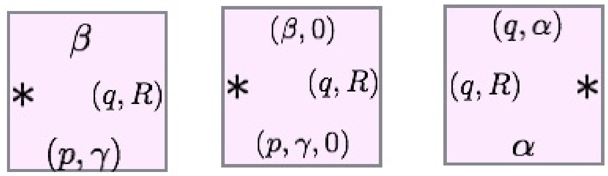

Figure 3: Type III tiles associated with a right-moving transition $\delta(p, \gamma) = (q, \beta, R)$. We have one copy of the third tile for each tape symbol α.

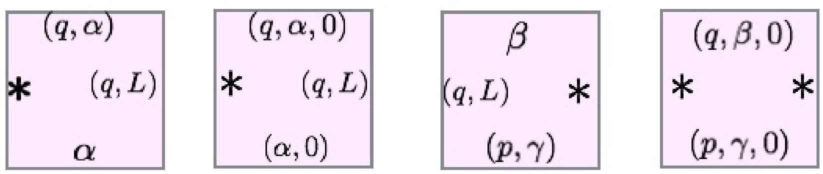

Figure 4: Type IV tiles associated with a left-moving-moving transition $\delta(p, \gamma) = (q, \beta, R)$.

of $\mathcal{M}$, and for every tape symbol α, $T_{\mathcal{M}}$ includes the tiles shown in Figure 3. Thus each transition right-moving transition gives rise to $2 + |\Gamma|$ tiles.

2.1.4 Type IV

For each left-moving transition

$$\delta(p, \gamma) = (q, \beta, L)$$

and each tape symbol α, we have the tiles depicted in Figure 4.

2.1.5 Type V

The designated corner tile is shown in Figure 5. Here q_0 denotes the initial state of $\mathcal{M}$.

2.2 Tiling with $T_{\mathcal{M}}$.

Now let's see what kind of tiling we can create with the tiles in this set. We are obliged to put the Type V tile in the lower left corner, and then only the type I tile

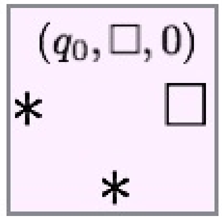

Figure 5: Type V tile—the designated corner tile.

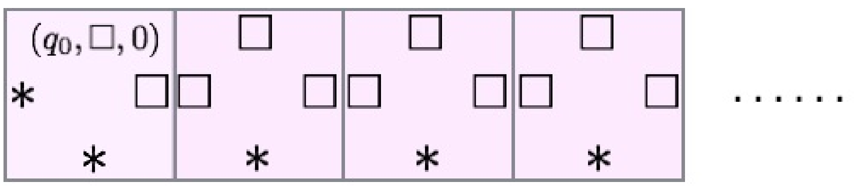

Figure 6: The only possible tiling of the first row.

can be placed to its right, because that is the only tile that has the symbol $\Box$ on its left edge. Thus the only possible tiling of the first row is the one shown in Figure 6.

Now let's consider the second row. If a tile is placed in the leftmost column of this row, it must have $(q_0, \Box, 0)$ on its bottom edge, so it must be a Type III or a Type IV tile, depending on whether $\delta(q_0, \Box)$ is a left-moving or right-moving transition. Let's assume that it is a right-moving transition, and that it is

$$\delta(q_0, \Box) = (q_1, a, R).$$

This forces the placement of a second Type III tile in the next cell of the second row, since there is only one tile with (q_1, R) on its left edge and $\Box$ on its bottom edge. And this placement in turn forces the placement of Type II tiles bearing the tape symbol $\Box$ for the rest of the row. So the tiling of the second row looks like Figure 7. (There may actually be two possibilities for a tile placed in the third column of the second row, because we could choose a tile that has $\Box$ at the bottom, $(q, \Box)$ at top, and (q, L) at right, as long as q is a state from which there is a left transition. But this choice will not allow us to extend the tiling further along the row, because we would next be forced to place a tile that has (q, γ), for some

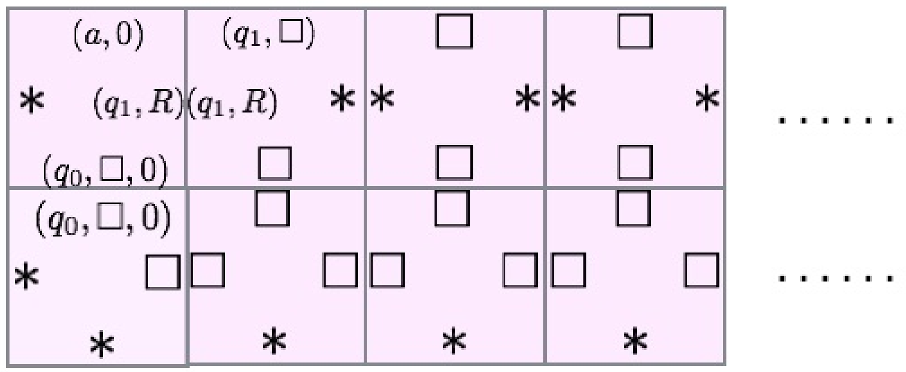

Figure 7: The only possible tiling of the first two rows, given the transition $\delta(q_0, \square) = (q_1, a, R)$.

state symbol γ, on its bottom edge, and this would not match the top edge $\square$ of the second row.)

Let's tile one more row—this should be enough to give you an idea of how the process unfolds. Any tile that goes directly above our second tile has to have $(q_1, \square)$ on its bottom edge. The only possibilities are a type III tile and a type IV tile, and which of these is available depends on whether $\delta(q_1, \square)$ is a left-moving or right-moving transition. Let us suppose that it is a left-moving transition (because we haven't done one of those yet):

$$\delta(q_1, \square) = (q_2, b, L).$$

This completely determines the tile placed in the second column of the third row (Figure 8).

There is then only one possibility–a tile of Type IV–that can be placed to the left of this tile. As we argued above, while there may be a choice for the tile to place immediately to the right, there is only one choice if we want to tile the entire third row. The result is shown in Figure 9.

You can now see what is going on in general: Look at the top edges of those first three rows, starting at the bottom. If we write out the colors on those edges, and ignore the zeros in the first column, we get

$$\mathbf{q_0}\square\square\square\square\cdots$$

$$a\mathbf{q_1}\square\square\square\cdots$$

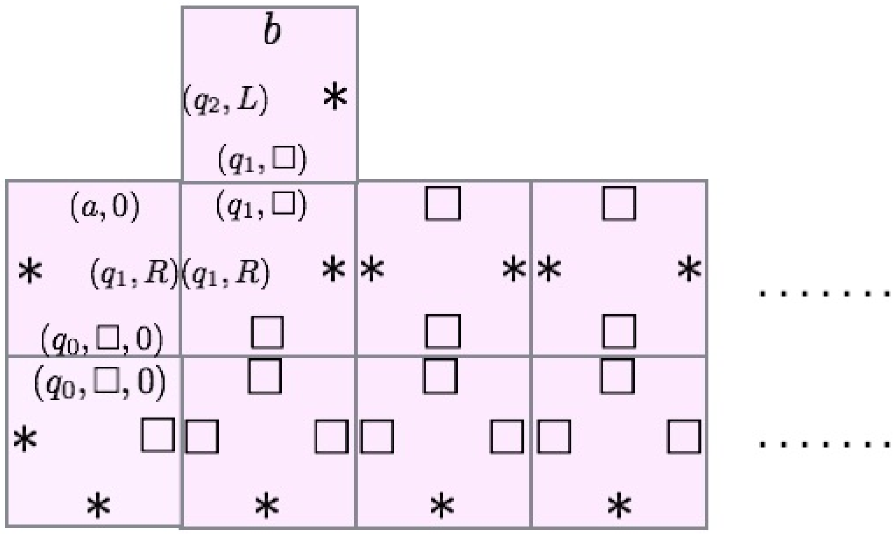

Figure 8: The only possible placement of a tile in the second column of the third row, assuming $\delta(q_1, \square)$ is a left-moving transition.

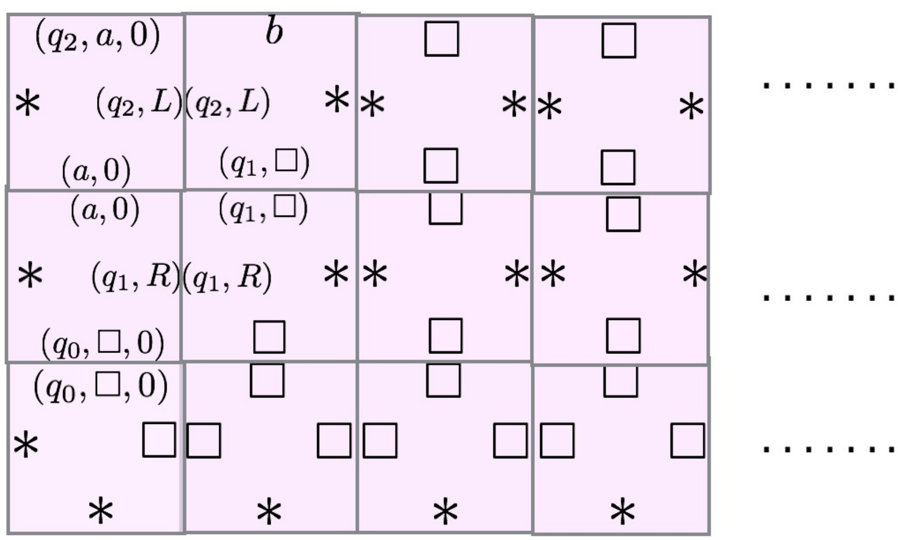

Figure 9: The tiling of the entire third row is then forced.

$$\mathbf{q_2}ab\Box\Box\cdots$$

These are the first three configurations of $\mathcal{M}$ when it is started on a blank tape. The tiles have been designed so that for any $k \geq 1$, there is at most one way to tile the first k rows, and the top edges of these rows spell out the first k configurations of the computation of $\mathcal{M}$ on a blank tape. (The zeros are present so that only tiles with a zero can be placed in the leftmost column, and these tiles cannot be placed anywhere else: This is because we have to handle left-moving transitions from the leftmost cell a little differently.) If $\mathcal{M}$ runs forever, the tiling can be continued indefinitely. If $\mathcal{M}$ halts after k steps, then one of the tiles in the k^{th} row will have its top edge labeled $(q_{\text{halt}}, \gamma)$. Since there is no transition from a halted state, the tiling cannot be continued. This establishes that the reduction has the desired property, and completes the proof of undecidability.

3 Non-periodic tilings

Let us now consider briefly the unrestricted tiling problem. Observe that the set of 'No' instances is Turing-recognizable: We start using our tiles to cover larger and larger sections of the plane. Eventually we will discover an n such that no way of placing our tiles on an $n \times n$ square correctly tiles it—*i.e,* there must be two adjacent tiles whose boundary colors don't match.

On the other hand, the way we discovered the tiling illustrated in the notes for the preceding lecture was to place tiles on larger and larger areas, until we discovered a repeated pattern that could be extended to tile the whole plane. So the set of tiles that yield such a 'periodic' tiling, formed by repeating the same pattern over and over again, is also Turing-recognizable. Thus if every infinite tiling of the plane were periodic, the tiling problem would be decidable: Just start placing tiles until you either find a square that you cannot tile correctly, or you find a repeated pattern. So undecidability of the general problem (which we have not proved) implies the existence of infinite tilings of the plane that are not periodic.