Problem 1

We use the following notation:

S_i	are the symptom variables, e.g., `cough, fever, painful toe`
D_i	are the disease variables, with values as given in the problem
`May-cause`(D_i, S_j)	disease i may cause symptom j

Part A

The domain of values of each symptom variable is the set of diseases that may cause it.

$$S_j = \{D_i \mid \text{may-cause}(D_i, S_j)\}$$

Intuitively, the value of a symptom variable is that disease that is said to be causing it, in this representation.

Part B

Constraints among disease variables must assure that the same "disease slot" is not assigned to more than one disease variable.

$$\text{Constraint}(D_i, D_j) = \begin{cases} \text{true} & \text{if } D_i = \texttt{NOTPRESENT} \text{ or } D_j = \texttt{NOTPRESENT} \text{ or } D_i \neq D_j \\ \text{false} & \text{otherwise} \end{cases}$$

Part C

There are no constraints among symptom variables.

Part D

To assure that each symptom has an explanation, we insist that if the value of a symptom variable is one of the diseases (which must be linked to it by a `may-cause` relationship), then the value of the corresponding disease variable may not be `NOTPRESENT`.

Note that this is a bit subtle, but has the right behavior. We cannot insist that *every* cause of a symptom be present, but this constraint assures that whichever one is chosen (by the value of the symptom variable) as the cause is in fact not `NOTPRESENT`—in other words, will have one of the DX_i values assigned to it.

The representation that we have chosen here allows us to express the constraints among symptoms and diseases as purely binary constraints.