

modelling at a *low* level of abstraction, akin to modelling an algorithm in a programming language. (See the answer to question 8 of Exercises 8.1 for a suggested alternative, more abstract, representation for the concept of a video copy for this specification.) This type of modelling may lead to a specification which is easier to implement in a programming language, but may also lead to a lack of clarity in the specification, and possibly less flexibility if it later becomes necessary to modify the specification to accommodate new requirements.

In contrast, the project allocation specification has a higher level of abstraction, with some use of non-determinism as already discussed. It may not be quite so obvious how to produce an implementation from such a specification, as it very much specifies *what* the system must do and not *how* it is to do it. However, the greater level of abstraction means that we can study the problem in its own right without having to be concerned about the paraphernalia of how we are eventually going to produce a program. So is it better to write specifications with a low level of abstraction, or a high level of abstraction? As always in the field of science, the answer is 'it depends'! The important thing is that the specification correctly represents the requirements, and is clear and understandable to all those who must read, discuss and use it. There are techniques in Z for writing specifications at a high level of abstraction, and then refining them to a lower level of abstraction, while proving that the properties of the original specification are preserved, thus getting the best of both worlds. However, these techniques are beyond the scope of this book.

Two outline Z specifications

Aims

To apply the preceding material to some more demanding specifications, and to illustrate the use of generic constants.

Learning objectives

When you have completed this chapter, you should be able to:

- apply the Z notation to modelling the state of more complex systems using combinations of mathematical structures;
- specify more challenging operations;
- appreciate the need to validate your specifications to ensure they are a true description of the required system;
- recognise situations in which it is desirable to define a new operator, or to name subexpressions using **let**, to enhance the readability of a specification.

11.1 Introduction

In this chapter, we will outline the construction of two Z specifications. The specifications are incomplete and not very large, but they are a little more challenging than the examples we have encountered so far, and should illustrate the value of Z in expressing complex ideas clearly and precisely.

11.2 A timetabling system

A university requires a timetabling system to keep track of where and when its degree modules are scheduled, and which students are registered for each module. Sometimes a module can be scheduled in several rooms at the same

time, perhaps for tutorial sessions in smaller groups, in which case we wish to know which room any given student should be in. We could visualise the final system as an on-line facility by which students can find out where they should be at any given time, register for modules, find out where their friends will be, etc. For simplicity, we will say that a given student may attend any number of modules provided none of their chosen modules clash on the timetable, and we will assume that there is no limit on the capacity of rooms or the number of students permitted to register for a module. For the time being, we will also ignore the role of lecturers and tutorial staff in the specification!

For our preliminary analysis, we identify the basic types (given sets) necessary for the specification. We will assume that for the purpose of time-tabling, the week is divided into fixed-size time slots in which modules may be scheduled. This suggests the following basic types:

[*STUDENT*, *MODULE*, *TIMESLOT*, *ROOM*]

which are respectively the set of all students, the set of all modules, the set of all time slots and the set of all teaching rooms.

The time slots could be units of one hour, two hours, half a day, or whatever. Making *TIMESLOT* a basic type means that we do not need to concern ourselves with the precise nature of a time slot at our chosen level of abstraction.

We must now describe the abstract state of the system using appropriate mathematical structures. Clearly, there is more than one way of doing this, but the one we will choose is based on the idea that each individual student and each individual module will be associated with their/its own individual schedule. This models the real-world situation, where each student will carry around their personal schedule, and a module lecturer might issue a schedule for his/her module. We can model a schedule as a relation between time slots and rooms. However, for any given time slot, a given student can only be in a maximum of one room (one cannot be in two places at the same time!), so therefore a student's schedule will be a partial function from time slots to rooms. Therefore, the type of a student's schedule will be

TIMESLOT → *ROOM*

partial function

and the type of a module's schedule will be

module can occupy more than one room for a time slot

TIMESLOT ↔ *ROOM*

binary relations

Now, we need a way to represent the schedules for *all* students and *all* modules on the course. We can do this using two partial functions, one of which maps individual students to their schedules, and the other of which maps individual

for a specific student

modules to their schedules. These functions are called *studentTT* and *moduleTT* respectively.

studentTT: *STUDENT* → (*TIMESLOT* → *ROOM*)

moduleTT: *MODULE* → (*TIMESLOT* ↔ *ROOM*)

*studentTT**s* is the schedule for student *s*, and *moduleTT**m* is the schedule for module *m*.

dom *studentTT*

is the set of all students on the course and

dom *moduleTT*

= course! in the English sense! (= program)

is the set of all modules which may be offered on the course. However, not all students are registered for modules and not all modules are offered at any given time. For a student *s* not registered for any modules,

*studentTT**s* = ∅

and for a module *m* not currently offered,

*moduleTT**m* = ∅

that is, *s* and *m* are mapped to empty schedules.

These structures constitute the state declarations for the specification. We must now identify the invariant properties required to define the state predicate:

1. It must not be possible to have two modules scheduled in the same room at the same time; in other words, the relations (schedules) in the range of *moduleTT* must be disjoint.

$\forall r, s: \text{ran } moduleTT \bullet \text{disjoint } \langle r, s \rangle$

The predicate implies that no two module schedules have any maplets in common; that is, no two modules are scheduled in the same room at the same time.

2. A student can only be scheduled for a time and place where a module is scheduled. In other words, any maplet found in any of the functions in the range of *studentTT* must also be found in precisely one of the relations in the range of *moduleTT*.

To express this requirement in Z, we need a convenient way of referring to the set of *all* the maplets in the ranges of *studentTT* and of *moduleTT*

respectively. We will define a generic global constant function *allPairs*, which when applied to objects such as *studentTT* or *moduleTT* returns the distributed union of all the relations in the object's range. (See Section 9.4 for a description of generic constants.)

$$\begin{array}{l} \text{Types} \\ \hline [X, Y, Z] \\ \hline \text{allPairs} : (X \rightarrow (Y \leftrightarrow Z)) \rightarrow (Y \leftrightarrow Z) \\ \hline \forall f : (X \rightarrow (Y \leftrightarrow Z)) \bullet \\ \text{allPairs } f = \bigcup \{x : X \mid x \in \text{dom } f \bullet fx\} \end{array}$$

For example, say *studentTT* has the following value:

$$\text{studentTT} = \left\{ \begin{array}{l} \text{sally} \mapsto \{t_1 \mapsto r_4, t_3 \mapsto r_4, t_5 \mapsto r_1\}, \\ \text{helen} \mapsto \{t_2 \mapsto r_1, t_3 \mapsto r_4\}, \\ \text{john} \mapsto \{t_1 \mapsto r_1, t_5 \mapsto r_2\} \end{array} \right\}$$

then

$$\text{allPairs } \text{studentTT} = \{t_1 \mapsto r_4, t_3 \mapsto r_4, t_5 \mapsto r_1, t_2 \mapsto r_1, t_1 \mapsto r_1, t_5 \mapsto r_2\}$$

Condition 2 can now be expressed as

$$\text{allPairs } \text{studentTT} \subseteq \text{allPairs } \text{moduleTT}$$

Note that this predicate simply states that every maplet in every student's schedule also occurs in the schedule of *at least one* module. The additional constraint that each student schedule maplet is in the schedule of *precisely one* module is implied by the predicate for condition 1. *No two modules are scheduled in the same room at the same time.*

Here is our final condition.

- Any given student is either scheduled to attend a given module at every available time slot, or not at all. Although this is a valid requirement for our specification, it may not be reflected in the behaviour of real students!

$$\begin{array}{l} \forall s : \text{dom } \text{studentTT}; m : \text{dom } \text{moduleTT} \\ \bullet (\text{studentTT } s \cap \text{moduleTT } m) \neq \emptyset \Rightarrow \\ \text{dom}(\text{studentTT } s \cap \text{moduleTT } m) = \text{dom}(\text{moduleTT } m) \end{array}$$

The expression

$$(\text{studentTT } s \cap \text{moduleTT } m) \neq \emptyset$$

is true iff student *s* and module *m* are scheduled to be in the same place at the same time on at least one occasion. The expression

$$\text{dom}(\text{studentTT } s \cap \text{moduleTT } m) = \text{dom}(\text{moduleTT } m)$$

states that student *s* is scheduled to attend module *m* at all of the times when the module is available.

The state schema for our system is therefore

$$\begin{array}{l} \text{Timetable} \\ \hline \text{studentTT} : \text{STUDENT} \rightarrow (\text{TIMESLOT} \leftrightarrow \text{ROOM}) \\ \text{moduleTT} : \text{MODULE} \rightarrow (\text{TIMESLOT} \leftrightarrow \text{ROOM}) \\ \hline \forall r, s : \text{ran } \text{moduleTT} \bullet \text{disjoint}(r, s) \\ \hline \text{allPairs } \text{studentTT} \subseteq \text{allPairs } \text{moduleTT} \\ \hline \forall s : \text{dom } \text{studentTT}; m : \text{dom } \text{moduleTT} \\ \bullet (\text{studentTT } s \cap \text{moduleTT } m) \neq \emptyset \Rightarrow \\ \text{dom}(\text{studentTT } s \cap \text{moduleTT } m) = \text{dom}(\text{moduleTT } m) \end{array}$$

For the initial state, we will have no students and no modules in the system.

$$\begin{array}{l} \text{InitTimetable} \\ \hline \text{TimeTable}' \\ \hline \text{studentTT}' = \{\} \\ \text{moduleTT}' = \{\} \end{array}$$

Inspection of the system state shows that this initial state exists. (Strictly, we are obliged to verify this formally.)

We should examine our specification to see whether it has any unforeseen undesirable properties. This check could take the form of a walkthrough with colleagues, in which the group tries to find inconsistencies in the specification as a representation of the user's requirements. For example, is it possible for a student to be scheduled for two different modules which clash on the timetable; that is, which take place at the same time? The answer is no, because every student's schedule is a function.

We should also check for any particular desirable properties which we would expect the specification to have. Other questions can give us more insight into our specification. For example, if two students are never in the same place at

the same time, must it be that they have no modules in common? A moment's consideration reveals that this is not necessarily true. It may be that a module has concurrent sessions in more than one room at every time slot in which it is scheduled, in which case two students could do the module without ever being in the same room.

Such checks can be made at any stage in the development of the specification, but it is a good idea to find any errors in the state specification before embarking on operation specifications.

The use of implication \Rightarrow is a common source of errors, and it is a good idea to check that the correct meaning has been captured. In condition 3, for example, what have we said about the following situation?

$$(studentTTs \cap moduleTTm) = \emptyset$$

In fact, $dom \emptyset = \emptyset$ (Spivey 1992), and so the consequent (right-hand side of \Rightarrow) will be true iff $moduleTTm = \emptyset$, that is module m is not timetabled, and false otherwise; either way the predicate will be true. *so why do we care?*

Validation exercises are not guaranteed to reveal all inconsistencies in the specification, but the fact that the specification is expressed in a formal language makes it easier to identify such inconsistencies. Verification by formal proof gives more confidence, but is much more expensive in time and effort.

Exercises 11.1

1. What assumption is implicit in condition 3 of the state invariant?
2. What issues would have to be considered in order to introduce the concept of *lecturers* for the modules? For example, what assumptions must you make about the case where a module is scheduled in more than one room at the same time?

Adding a student to the course

The student $s?$ must not already be on the course.

$$s? \notin dom\ studentTT$$

The student joins the course with an initially empty schedule.

$$studentTT' = studentTT \cup \{s? \mapsto \emptyset\}$$

The operation schema is as follows:

<i>AddStudent</i>
Δ <i>Timetable</i> $s?: STUDENT$
$s? \notin dom\ studentTT$ $studentTT' = studentTT \cup \{s? \mapsto \emptyset\}$ $moduleTT' = moduleTT$

Exercises 11.2

1. Write the operation schema *AddModule*, to add a module with an empty schedule to the system.
2. Write the operation schemas *RemoveStudent* and *RemoveModule*.

Scheduling a module

The module $m?$ must be a valid module for the course

$$m? \in dom\ moduleTT$$

and it must not already be scheduled:

$$moduleTTm? = \emptyset$$

The postcondition is non-deterministic in that more than one valid potential schedule for the module may be possible, and in these circumstances, the postcondition does not state which valid schedule is to be selected. This arbitrary decision is left to the implementors of the system. Thus the schema defines more than one possible final state.

$$\exists\ schedule: TIMESLOT \leftrightarrow ROOM \bullet$$

$$(allPairs\ moduleTT \cap schedule = \emptyset \quad \text{do not clash with another module}$$

$$\wedge\ moduleTT' = moduleTT \oplus \{m? \mapsto schedule\}) \quad \text{in terms of room-time slot}$$

The components of this predicate are explained as follows. Any valid schedule for this module must not clash with that of any other module.

$$allPairs\ moduleTT \cap schedule = \emptyset$$

The module is scheduled by overriding the *moduleTT* function.

$$moduleTT' = moduleTT \oplus \{m? \mapsto schedule\}$$

The operation schema is as follows:

<i>ScheduleModule</i>
Δ <i>Timetable</i> $m? : MODULE$
$m? \in \text{dom } moduleTT$ $moduleTT m? = \emptyset$ $\exists schedule : TIMESLOT \leftrightarrow ROOM \bullet$ $(allPairs \ moduleTT \cap \ schedule = \emptyset$ $\wedge \ moduleTT' = moduleTT \oplus \{m? \mapsto schedule\})$
$studentTT' = studentTT$

Descheduling a module

The module to be descheduled $m?$ must be a valid module for the course.

$$m? \in \text{dom } moduleTT$$

and the schedule for $m?$ must not already be empty.

$$moduleTT m? \neq \emptyset$$

To deschedule the module, we override $moduleTT$ to map the module to the empty schedule.

$$moduleTT' = moduleTT \oplus \{m? \mapsto \emptyset\}$$

We must also remove the module from the schedules of all students who are registered for it.

$$studentTT' = \bigcup \{s : \text{dom } studentTT \bullet \{s \mapsto (studentTT s \setminus moduleTT m?)\}\}$$

The operation schema is as follows:

<i>DescheduleModule</i>
Δ <i>Timetable</i> $m? : MODULE$
$m? \in \text{dom } moduleTT$ $moduleTT m? \neq \emptyset$ $moduleTT' = moduleTT \oplus \{m? \mapsto \emptyset\}$ $studentTT' = \bigcup \{s : \text{dom } studentTT \bullet \{s \mapsto (studentTT s \setminus moduleTT m?)\}\}$

Note that the precondition $moduleTT m? \neq \emptyset$ is not necessary for the operation as specified, but it may be important to identify it at implementation, to output an appropriate message if the precondition is not satisfied. This exception would be handled by a separate schema, and the total operation would be specified using schema calculus.

Registering a student for a module

The student $s?$ must be on the course, and the module $m?$ must be a valid module for the course.

$$s? \in \text{dom } studentTT$$

$$m? \in \text{dom } moduleTT$$

The module must be scheduled.

$$moduleTT m? \neq \emptyset$$

The student must be free at all the times when the module is scheduled.

$$\text{dom}(studentTT s?) \cap \text{dom}(moduleTT m?) = \emptyset$$

time-room combinations

We do not want to add all the module's slots to the student's schedule; where the module is scheduled in more than one room at the same time, the student must be allocated only one of these slots. The following postcondition predicate is non-deterministic in this respect, in that it does not state specifically which slot is to be allocated in these circumstances. (In a real system, issues such as room capacities and previous allocations would come into play.) The postcondition is as follows:

$$\exists newPairs : TIMESLOT \leftrightarrow ROOM$$

- $((\text{dom } newPairs = \text{dom } moduleTT m?)$ *match timeslots with those of the module*
- $\wedge (newPairs \subseteq moduleTT m?)$
- $\wedge (studentTT' = studentTT \oplus \{s? \mapsto studentTT s? \cup newPairs\})$

$newPairs$ is the set of those pairs from the module's schedule which are to be added to the student's schedule. Note that because students' schedules are functions, $newPairs$ must be a function. The constituent parts of this predicate are explained as follows.

The time slots for the pairs to be added to the student's schedule are precisely all the time slots when the module is available.

$$\text{dom } newPairs = \text{dom } moduleTT m?$$

The pairs to be added to the student's schedule are pairs from the module's schedule.

$$\text{newPairs} \subseteq \text{moduleTT}m?$$

The appropriate pairs are added to the student's schedule by overriding the *studentTT* function.

$$\text{studentTT}' = \text{studentTT} \oplus \{s? \mapsto \text{studentTT} s? \cup \text{newPairs}\}$$

The operation schema is as follows:

<i>RegForModule</i>
Δ <i>Timetable</i> $s? : \text{STUDENT}$ $m? : \text{MODULE}$
$s? \in \text{dom } \text{studentTT}$ $m? \in \text{dom } \text{moduleTT}$ $\text{moduleTT} m? \neq \emptyset$ $\text{dom}(\text{studentTT} s?) \cap \text{dom}(\text{moduleTT} m?) = \emptyset$
$\exists \text{newPairs} : \text{TIMESLOT} \leftrightarrow \text{ROOM}$ $\bullet ((\text{dom } \text{newPairs} = \text{dom } \text{moduleTT} m?)$ $\quad \wedge (\text{newPairs} \subseteq \text{moduleTT} m?)$ $\quad \wedge (\text{studentTT}' = \text{studentTT} \oplus \{s? \mapsto \text{studentTT} s? \cup \text{newPairs}\}))$
$\text{moduleTT}' = \text{moduleTT}$

Exercises 11.3

- Write a schema to 'deregister' a student from a module.
- How could we modify the state schema to specify that each module can use a maximum of one room in any given time slot?
- How could we modify the state schema to specify that each module is only allowed one time slot in the schedule?
- Write a Z expression for the set of all students registered for a module *m*.
- Write a Z expression for the set of all modules being taken by a student *s*.
- Write a Z expression for the set of all students in a room *r* at a time *t*.
- Write a Z expression for the set of all modules which student *p* and student *q* have in common.
- Write a Z expression for the set of all times when student *p* and student *q* are in the same room.

- Write a Z expression for the set of all modules which clash with a module *m* on the timetable, that is which take place at the same time as module *m*.
- Write a Z expression for the set of all time/room maplets for which one or more modules are scheduled, but no students are scheduled.
- What would it mean if *allPairs moduleTT* was one-to-one? (See Section 7.4 for a definition of the term one-to-one for functions.)
- What would it mean if *allPairs studentTT* was one-to-one?
- How could you extend the specification to include concepts such as room capacities and maximum numbers of students allowed in a module.
- Why would it be difficult to extend the module registration example in Chapter 6 to include timetabling information?

11.3 A genealogical database

A database is required to keep track of genealogical relationships between people (family trees). It would be possible to represent the required relationships (parent, grandparent, aunt, cousin, etc.) separately, but this would limit the number of relationships, increase the complexity of the specification, and would make it necessary to carry out extensive integrity checks every time the database is updated. We will therefore represent only the minimum information necessary to be able to define operations to output any required relationships. The most fundamental genealogical relationship is that of parent to child, and this, together with the sex of all the individuals in the database, will be enough to enable us to specify all the operations we require. This suggests the following types:

$[PERSON]$ the set of all people
 $GENDER ::= \text{male} \mid \text{female}$

The parent/child relationship can be represented by the relation

$\text{parent} : PERSON \leftrightarrow PERSON$

where

$x \mapsto y \in \text{parent}$

represents the information that *y* is *x*'s parent.

The sex of all the people in the database can be represented by the function

$\text{sex} : PERSON \rightarrow GENDER$

where

$\text{dom } \textit{sex}$

is the set of all people in the database. A person can be in the database even if they do not occur in *parent*. It may be that information about their parents or children is not available or is incomplete.

$\textit{sex } p$

is the sex of person *p*.

We need an invariant predicate to state that the relation *parent* only holds information about people in the database.

$$\text{dom } \textit{parent} \cup \text{ran } \textit{parent} \subseteq \text{dom } \textit{sex}$$

We also need an invariant predicate to capture such requirements as:

1. If *y* is *x*'s parent, then *x* cannot be *y*'s parent.
2. If *y* is *x*'s parent, then *y* cannot also be an ancestor of *x* at a depth in the family tree greater than that of *parent*, that is *y* cannot be *x*'s grandparent, great-grandparent, etc.
3. A person cannot be their own parent.

We can neatly capture all of these requirements in a predicate which states that a person cannot be their own ancestor. (You may have to think about this to convince yourself that it is true.) Now the transitive closure of *parent*

\textit{parent}^+

will relate any person to their ancestors (parents, grandparents, great-grandparents, etc.). The required predicate is therefore simply

$$\forall p: \textit{PERSON} \bullet p \mapsto p \notin \textit{parent}^+$$

The final restriction is that anyone in the database can have a maximum of two parents, and if they have two, the parents must be of opposite sexes!

$$\forall p, q, r: \textit{PERSON} \bullet \{p \mapsto q, p \mapsto r\} \subseteq \textit{parent} \wedge q \neq r \Rightarrow \textit{sex } q \neq \textit{sex } r$$

Note that the condition restricts the number of parents to a maximum of two because the type *GENDER* only has two values.

The state schema is as follows:

<i>GenDB</i>
<i>parent</i> : <i>PERSON</i> \leftrightarrow <i>PERSON</i>
<i>sex</i> : <i>PERSON</i> \rightarrow <i>GENDER</i>
$\text{dom } \textit{parent} \cup \text{ran } \textit{parent} \subseteq \text{dom } \textit{sex}$
$\forall p: \textit{PERSON} \bullet p \mapsto p \notin \textit{parent}^+$
$\forall p, q, r: \textit{PERSON} \bullet \{p \mapsto q, p \mapsto r\} \subseteq \textit{parent} \wedge q \neq r \Rightarrow \textit{sex } q \neq \textit{sex } r$

You may notice that there is no requirement for a person's parents to come from the same generation as each other; nor does the database contain any information about whether the people it contains are alive or dead. Thus a person in the database could have Julius Caesar as one parent and Joan of Arc as the other. However, for simplicity, we will quietly ignore this.

For the initial state, we will have an empty database.

<i>InitGenDB</i>
<i>GenDB'</i>
$\textit{sex}' = \emptyset$
$\textit{parent}' = \emptyset$

The predicate $\textit{parent}' = \emptyset$ is not strictly necessary, as it is implied by the state invariant, but increases the clarity of the specification. The initial state satisfies the state invariant.

Operations to change the database

The users of the database will be able to modify the information it contains using operations to add a person to the database, remove a person from the database, add or remove a parent/offspring relationship to/from the database, change the name of a person in the database, and change the sex of a person in the database.

Adding a person to the database

The inputs are a person and their sex.

$\textit{name}?: \textit{PERSON}$
 $\textit{morf}?: \textit{GENDER}$

The person *name?* must not already be in the database.

$\textit{name}? \notin \text{dom } \textit{sex}$

We add them to the *sex* function.

$$sex' = sex \cup \{name? \mapsto morf?\}$$

The operation schema is as follows:

<i>AddPerson</i>
$\Delta GenDB$ $name? : PERSON$ $morf? : GENDER$
$name? \notin \text{dom } sex$ $sex' = sex \cup \{name? \mapsto morf?\}$ $parent' = parent$

Exercise 11.4

Write a schema for an operation to remove a person from the database.

Adding a parent/offspring relationship to the database

The inputs are a potential offspring and a potential parent.

$$off?, par? : PERSON$$

The potential offspring and parent must be in the database

$$\{off?, par?\} \subseteq \text{dom } sex$$

and must not already be a maplet in *parent*, in either order.

$$off? \mapsto par? \notin parent$$

$$par? \mapsto off? \notin parent$$

There must not be more than one parent already in the database for the potential offspring

$$\#(\{off?\} \triangleleft parent) \leq 1$$

and if there is one such parent, their sex must not be the same as the new potential parent.

$$\forall x : PERSON \bullet off? \mapsto x \in parent \Rightarrow sex\ x \neq sex\ par?$$

For the postcondition, we simply add the new maplet to the relation *parent*.

$$parent' = parent \cup \{off? \mapsto par?\}$$

The operation schema is as follows:

<i>AddRel</i>
$\Delta GenDB$ $off?, par? : PERSON$
$\{off?, par?\} \subseteq \text{dom } sex$ $off? \mapsto par? \notin parent$ $par? \mapsto off? \notin parent$ $\#(\{off?\} \triangleleft parent) \leq 1$ $\forall x : PERSON \bullet off? \mapsto x \in parent \Rightarrow sex\ x \neq sex\ par?$ $parent' = parent \cup \{off? \mapsto par?\}$ $sex' = sex$

Exercise 11.5

Write a schema for an operation to remove a parent/offspring relationship from the database.

Changing the name of a person in the database

This is a rather artificial example, as we will now have to describe the type *PERSON* as the set of all people's names, which implies that all names are unique. In a real system, *PERSON* would have to be implemented as a set of unique identifiers of some sort, and for simplicity, when we refer to a 'name', we will take it to mean one of these identifiers. The inputs are the old name and the new name.

$$old?, new? : PERSON$$

The old name must be in the database, and the new one must not.

$$old? \in \text{dom } sex$$

$$new? \notin \text{dom } sex$$

The old name must be replaced by the new one in the *sex* function

$$sex' = (\{old?\} \triangleleft sex) \cup \{new? \mapsto sex\ old?\}$$

and all relationships involving the old name must be modified to use the new name.

$$\begin{aligned} \text{parent}' &= (\{old?\} \triangleleft \text{parent} \triangleright \{old?\}) \\ &\cup \{x: PERSON \mid x \in \text{parent} (\{old?\}) \bullet \text{new?} \mapsto x\} \\ &\cup \{x: PERSON \mid x \in \text{parent}^{-1} (\{old?\}) \bullet x \mapsto \text{new?}\} \end{aligned}$$

functional notation used for relations!

The operation schema is as follows:

ChangeName
$\Delta \text{ GenDB}$ $old?, new?: PERSON$
$old? \in \text{dom } sex$ $new? \notin \text{dom } sex$ $sex' = (\{old?\} \triangleleft sex) \cup \{new? \mapsto sex \text{ old?}\}$ $\text{parent}' = (\{old?\} \triangleleft \text{parent} \triangleright \{old?\})$ $\cup \{x: PERSON \mid x \in \text{parent} (\{old?\}) \bullet \text{new?} \mapsto x\}$ $\cup \{x: PERSON \mid x \in \text{parent}^{-1} (\{old?\}) \bullet x \mapsto \text{new?}\}$

Changing the sex of a person in the database

The person $p?$ must be in the database.

At first sight, this operation would appear to require a simple overriding of the sex function. However, the tricky part in specifying the operation is that the person may be recorded as a $parent$ in the database, that is they may be a member of $\text{ran } parent$. This means that the sex recorded for other people may have to change in order to maintain the integrity of the database, that is so that the database will not contain children with two mothers or two fathers.

An implementation of the operation would probably ask the user whether s/he wished to proceed in these circumstances, as it is unlikely that s/he would wish to make such changes simply to maintain the integrity of the database. It is much more likely that the proposed change was not correct in the first place. However, we will specify a rather contrived operation which simply makes the necessary changes. Essentially, any people with whom our person has had children must change their sex, as must anyone who has had children with those people, and so on, to ensure that anyone with two parents in the database has one of each sex.

Now the relation

$$\text{parent}^{-1}; \text{parent}$$

composition

So marriage is not an issue!

relates together people who have had children with each other. The transitive closure of this relation

$$(\text{parent}^{-1}; \text{parent})^+$$

is the relation such that

$$p \mapsto q \in (\text{parent}^{-1}; \text{parent})^+$$

iff p has had children with q , or p has had children with someone who has had children with q , or p has had children with someone who has had children with someone who has had children with q , ... well, you get the idea. The relational image in this relation of the set consisting solely of our person $p?$ gives the set of all people who must have their sex changed by this operation.

$$(\text{parent}^{-1}; \text{parent})^+ (\{p?\})$$

We can now use this to specify a function mapping these people to the opposite sex from that given them by the sex function, and finally use this function to override the original sex function. The resulting postcondition is as follows:

$$\begin{aligned} sex' &= sex \oplus && \leftarrow \text{"enumerate" genders!} \\ \{q: PERSON; s: GENDER \mid (q \in (\text{parent}^{-1}; \text{parent})^+ (\{p?\})) \} \\ &\wedge (s \neq sex \ q) \bullet q \mapsto s \end{aligned}$$

Note that the relation $\text{parent}^{-1}; \text{parent}$ also relates every parent to themselves, which enables the postcondition to specify the sex change for $p?$ him/herself.

The operation schema is as follows:

ChangeSex
$\Delta \text{ GenDB}$ $p?: PERSON$
$p? \in \text{dom } sex$ $sex' = sex \oplus$ $\{q: PERSON; s: GENDER \mid (q \in (\text{parent}^{-1}; \text{parent})^+ (\{p?\}))$ $\wedge (s \neq sex \ q) \bullet q \mapsto s\}$
$\text{parent}' = \text{parent}$

We also require query operations to interrogate the database for information about various relationships. The next set of exercises gives you the opportunity to specify some, after which we conclude with an operation to find the common ancestors of two given people, and an operation to find the set of all cousins of specified type and removedness for a given person.

Exercises 11.6

1. Specify an operation to return the set of all people who have any one of the following relationships to a given person $x?$. The name of the required relationship should be an input to the operation.
 - (i) The parents of person $x?$
 - (ii) The grandparents of person $x?$
 - (iii) The grandchildren of person $x?$
 - (iv) The descendants of person $x?$
 - (v) The siblings of person $x?$
 - (vi) The aunts of person $x?$
2. Give a Z expression for the set of all people in the database who have no relatives in the database.
3. Give a Z expression for the set of all people in the database who have no siblings in the database.
4. Specify an operation to output the number of generations through which a given individual p can trace his/her family history in the database.

The common ancestors of two given people

This operation must return the set $cas!$ containing the common ancestors of two people, say $p?$ and $q?$. We will further stipulate that the set contains only the common ancestors of minimum degree, where the degree refers to the number of steps in the path up and/or down the family tree between the two people. For example, if $p?$ and $q?$ are siblings, the degree is 2. If $p?$ and $q?$ are first cousins, the degree is 4. If $p?$ is the grandparent of $q?$, the degree is 2.

The precondition is that all the people involved are in the database.

$$\{p?, q?\} \cup cas! \subseteq \text{dom } sex$$

But cas! is the computed result!

The postcondition characterises elements of the set of common ancestors of minimum degree as people for whom there are multiple compositions of $parent$ which map both $p?$ and $q?$ to them, and furthermore there are no other common ancestors with degree smaller than that of members of this set.

$$cas! = \{ca : PERSON \mid \exists m, n : \mathbb{N} \bullet ((p? \mapsto ca \in parent^m \wedge q? \mapsto ca \in parent^n) \wedge \neg \exists r : PERSON; x, y : \mathbb{N} \bullet ((x + y < m + n) \wedge p? \mapsto r \in parent^x \wedge q? \mapsto r \in parent^y))\}$$

Some rel^o : identity relation p. 74

The operation schema is as follows:

CommonAncestors

$\exists GenDB$
 $p?, q? : PERSON$
 $cas! : \mathbb{P} PERSON$

$\{p?, q?\} \cup cas! \subseteq \text{dom } sex$

$cas! = \{ca : PERSON \mid \exists m, n : \mathbb{N} \bullet ((p? \mapsto ca \in parent^m \wedge q? \mapsto ca \in parent^n) \wedge \neg \exists r : PERSON; x, y : \mathbb{N} \bullet ((x + y < m + n) \wedge p? \mapsto r \in parent^x \wedge q? \mapsto r \in parent^y))\}$

Note that the specification allows for the special case where one of the two people is a direct descendant of the other, in which case the common ancestor is the person higher in the family tree. m and n are natural numbers, and it is therefore possible for either or both to be zero, yielding the identity relation on $PERSON$. The only debate might be as to whether a person can be their own ancestor!

The cousins of a given person

This operation returns the set $cousins!$ containing all cousins of a given type and removedness for a given person, say $p?$. Cousins are people who have a common ancestor who is more distant than a parent, and who are not siblings. The type of cousinship is determined by the shortest path from the common ancestor to either one of the cousins. For example, for first cousins the common ancestor would be a grandparent of one cousin, for second cousins a great-grandparent, and so on. 'Removed' refers to the difference in the number of steps in the path from each cousin to the common ancestor. 'Once removed' means a difference of one step, 'twice removed' means a difference of two steps, etc. For example, my first cousins' children are my first cousins once removed, and their children are my first cousins twice removed.

This operation has an input $nth? : \mathbb{N}_1$, to represent the type of cousinship, and an input $rem? : \mathbb{N}$, to represent how far removed the relationship is. $nth?$ has type \mathbb{N}_1 because the minimum type of cousinship is first cousins, represented by the number 1.

The precondition simply states that everyone involved is in the database:

$$\{p?\} \cup cousins! \subseteq \text{dom } sex$$

For the postcondition we define a relation which maps a cousin to his/her cousins at the same or a lower level in the family tree. Remember that $nth?$ is defined with respect to the cousin with the shortest path to the common

ancestor, and $rem?$ represents the difference in length between the two paths. The required relation is

$$(parent^{nth?+1}; (parent^{-1})^{nth?+1+rem?}) \setminus (parent; parent^{-1})$$

sibling relation, including a person with himself as a sibling

The set subtraction removes from the relation all pairs of people who have the same parents. Therefore if $rem?$ is zero, the relation does not relate anyone to themselves or their siblings.

The image of our person in this relation will yield all the relevant cousins at the same or a lower level in the family tree as our person. However, if the value of $rem?$ is non-zero, our person may have some cousins with the same relationship at a higher level in the tree. We therefore also require the image of our person in the *inverse* of the above relation. To avoid repeating the expression for this relation, we use a Z construct called a *let predicate* which provides a method for naming subexpressions in complex predicates.

$$\text{let } n = e \bullet p$$

stands for the predicate p , but wherever the name n occurs in p , it represents the value e . We use a let predicate to give the name $cosrel$ to the above relation for use as a subexpression in specifying the set of all $nth?$ cousins $rem?$ removed, both above and below person $p?$ in the family tree.

The operation schema is as follows:

<i>Cousins</i>
$\exists GenDB$ $p?: PERSON$ $nth?: \mathbb{N}_1$ $rem?: \mathbb{N}$ $cousins!: \mathbb{P} PERSON$
$\{p?\} \cup cousins! \subseteq \text{dom } sex$
$\text{let } cosrel = (parent^{nth?+1}; (parent^{-1})^{nth?+1+rem?}) \setminus (parent; parent^{-1}) \bullet$ $cousins! = cosrel(\{p?\}) \cup cosrel^{-1}(\{p?\})$

11.4 In conclusion

In this chapter, we have looked at two outline specifications which are somewhat more complex than those we have met previously. However, we have not introduced much new notation; we have simply applied our previous knowledge to some harder problems. Hopefully these examples will serve to

demonstrate the power of the Z notation to model complex situations succinctly, clearly and precisely. If you are having difficulty understanding them, it may help to draw pictures of typical values of the state variables, to enable you to visualise what is going on. If you understand the examples at first viewing that's fine, but if not, don't be disheartened! Complex problems require a lot of study, and if it was easy it wouldn't be rewarding. When you can understand, modify and extend the above examples, you will be very well prepared for further study of the Z language.

Farklı yapılabildi:
 $temp = (parent^{nth?+1}; (parent^{-1})^{nth+1+rem?}) \setminus (parent; parent^{-1})$
 $cousins! = \{cousin : PERSON \mid P? \rightarrow cousin \in temp \bullet cousin\} \cup \{cousin : PERSON \mid cousin \rightarrow P? \in temp \bullet cousin\}$