Mathematics for Computer Scientists

Lecture notes for the module G51MCS

Venanzio Capretta
University of Nottingham
School of Computer Science

Chapter 4

Sets, Functions and Relations

4.1 Sets

Sets are such a basic notion in mathematics that the only way to 'define' them is by synonyms like *collection, class, grouping* and so on. A set is completely characterised by the elements it contains.

There are two main ways of defining a set: (1) by explicitly listing all its elements or (2) by giving a property that all elements must satisfy. Here is a couple of examples, the first is a set of fruits given by listing its elements, the second is a set of natural numbers given by specifying a property:

$$A = \{\mathsf{apple}, \mathsf{banana}, \mathsf{cherry}, \mathsf{peach}\}$$
$$E = \{n \in \mathbb{N} \mid 2 \text{ divides } n\}$$

The definition of E can be read: *E is the set of those natural numbers that are divisible by* 2, that is, it's the set of all even naturals.

The first definition method can be used only if the set has a finite number of elements. In some cases both methods can be used to define the same set, as in this example:

$$\{n \in \mathbb{N} \mid n \text{ is odd} \land n^2 + n \leq 100\} = \{1, 3, 5, 7, 9\}$$

To say that a certain object x is an element of a set S, we write $x \in S$. To say that it isn't an element of S we write $x \notin S$. For example:

$$\begin{array}{rl} \mathsf{apple} \in A & \quad 7 \notin E \\ \mathsf{strawberry} \notin A & \quad 8 \in E \end{array}$$

If every element of a set X is also an element of another set Y, then we say that X *is a subset of* Y and we write this symbolically as $X \subseteq Y$. Formally, the subset relation is defined as follows:

$$X \subseteq Y \quad \Leftrightarrow \quad \text{for every } x,\ x \in X \Rightarrow x \in Y.$$

As example, here's a couple of subsets of the sets A and E from above:

$$\{\mathsf{apple}, \mathsf{peach}\} \subseteq A$$
$$\{x \in \mathbb{N} \mid n \text{ is a multiple of } 4\} \subseteq E.$$

There is a set that is contained in any other set: the *empty set*, that is, the set with no elements. We use the symbol $\varnothing$ for it:

$$\varnothing = \{\ \}.$$

It is always trivially true that $\varnothing \subseteq X$ and also that $X \subseteq X$.

When we define a set by a property, we should also clarify in advance what kind of objects we are talking about: in the examples above, we wrote $n \in \mathbb{N}$ to specify that we are talking about natural numbers. This larger set, containing all the objects that we are interested in, is called the *universal set* or just the *universe*. We will be using the letter U to denote the universal set.

Sets can be combined and manipulated by using the operations of intersection, union, difference, complement. Here are their intuitive meaning and their rigorous mathematical definitions, assuming that S and T are any two sets:

Intersection $S \cap T$: the elements that belong both to S and to T.

$$S \cap T = \{x \in U \mid x \in S \wedge x \in T\}$$

Union $S \cup T$: the elements that belong either to S or to T (or both).

$$S \cup T = \{x \in U \mid x \in S \vee x \in T\}$$

Difference $S \setminus T$: the elements that belong to S but not to T.

$$S \setminus T = \{x \in U \mid x \in S \wedge x \notin T\}$$

Complement $\overline{S}$: elements (of the universe) that don't belong to S.

$$\overline{S} = \{x \in U \mid x \notin S\}$$

4.2 Venn Diagrams

We can represent arbitrary sets pictorially by some drawings called *Venn diagrams*. Sets are blobs that overlap each other. Any region of the drawing can be characterised by some expression obtained by combining the sets by set operations.

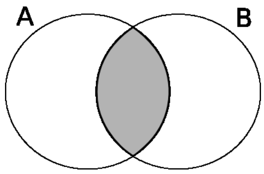

The shaded area represents $A \cap B$

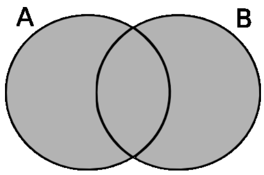

The shaded area represents $A \cup B$

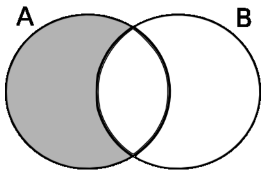

The shaded area represents $A \setminus B$

Whenever we want to show also the boundaries of the universal set, it is depicted as a big rectangle. For example, we need to show the universe when representing the complement:

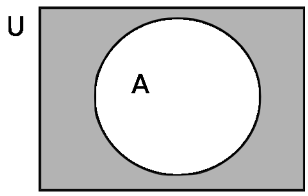

The shaded area represents $\overline{A}$

Venn diagram associated to an expression. Given any expression combining variable names for sets using the operators of intersection, union, difference and complement, we can draw a Venn diagram and identify on it the area associated with the given expression. For example, here is a Venn diagram with a shaded area associated to the expression $(A \cup C) \setminus B$:

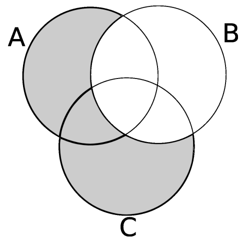

The shaded area represents $(A \cup C) \setminus B$

4.3 The algebra of sets

The expressions obtained by combining sets by set operations form a kind of algebra. To check what equalities hold in this algebra, we can use Venn diagrams. Remember, however, that the diagrams are only intuitive drawings and they are not considered a proper proof.

For example, we want to check if the following equality is true for all possible sets A, B and C:

$$A \cup (B \cap C) = (A \cup B) \cap (A \cup C).$$

In other words: does union distribute over intersection?

Let's construct two Venn diagrams depicting the left-hand and right-hand side of this equality, respectively:

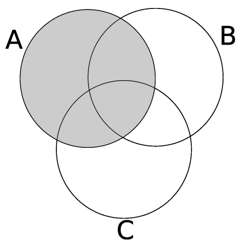

shaded area: A

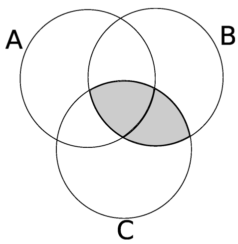

shaded area: $B \cap C$

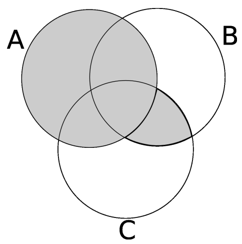

shaded area: $A \cup (B \cap C)$

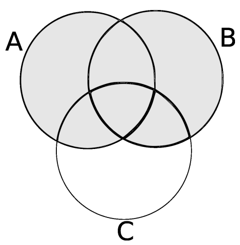

shaded area: $A \cup B$

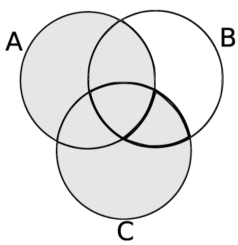

shaded area: $A \cup C$

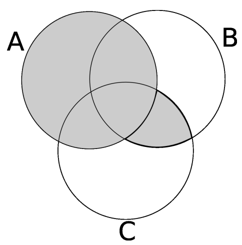

shaded area: $(A \cup B) \cap (A \cup C)$

We obtained the same area in the two diagrams for the two sides of the equality. This tells us that the equality is probably true.

This was not a proper proof: Venn diagrams are only an intuitive way to picture sets, they do not actually correspond to the real sets. If we want to be mathematically sure of the equality, we must prove it rigorously from the definitions.

Theorem 9 *Given any three sets A, B and C; the following equality holds:*

$$A \cup (B \cap C) = (A \cup B) \cap (A \cup C).$$

Proof. Let's unfold the definitions to check what it means to be an element of those two sets. For every element $x \in U$ we have that:

$$\begin{aligned} x \in A \cup (B \cap C) &\Leftrightarrow (x \in A) \vee (x \in B \cap C) \\ &\Leftrightarrow (x \in A) \vee ((x \in B) \wedge (x \in C)); \end{aligned}$$

$$\begin{aligned} x \in (A \cup B) \cap (A \cup C) &\Leftrightarrow (x \in A \cup B) \wedge (x \in A \cup C) \\ &\Leftrightarrow ((x \in A) \vee (x \in B)) \wedge ((x \in A) \vee (x \in C)). \end{aligned}$$

But now, by distributivity of disjunction over conjunction, we have that:

$$(x \in A) \vee ((x \in B) \wedge (x \in C)) \Leftrightarrow ((x \in A) \vee (x \in B)) \wedge ((x \in A) \vee (x \in C)).$$

[If you're not convinced of this step, go back to the rules of Boolean algebra. Check the rule of distributivity of disjunction over conjunction and make the following substitutions: replace A by $(x \in A)$, replace B by $(x \in B)$ and replace C by $(x \in C)$. You will obtain exactly the equivalence above.]

If we put all the equivalences together, we obtain:

$$x \in A \cup (B \cap C) \Leftrightarrow x \in (A \cup B) \cap (A \cup C).$$

This states that being an element of $A \cup (B \cap C)$ is equivalent to being an element of $(A \cup B) \cap (A \cup C)$. In conclusion: $A \cup (B \cap C) = (A \cup B) \cap (A \cup C)$. □

Notice that we exploited the Boolean law of distributivity of disjunction over conjunction to prove distributivity of union over intersection. This works because intersection was defined using conjunction and union was defined using disjunction. It is a general pattern: all the rules of Boolean algebra give corresponding rules of set algebra. Complement corresponds to negation. So if you take a Boolean equality, replace $\wedge$ by $\cap$, $\vee$ by $\cup$ and $\neg$ by $\bar{\cdot}$, you obtain a set equality.

For example, the first De Morgan law becomes:

$$\overline{A \cap B} = \overline{A} \cup \overline{B}.$$

Try to prove this equality formally, like we did above for distributivity.

4.4 Cartesian Product

Another important binary operation on sets is the *Cartesian Product*: given two sets A and B, their Cartesian product, indicated by $A \times B$ is the set of pairs of elements from them. If $a \in A$ and $b \in B$, then we indicate by $\langle a, b \rangle$ the pair that they form. So we have:

$$A \times B = \{\langle a, b \rangle \mid a \in A \wedge b \in B\}.$$

The order of the pair is important: the same two elements may form two different pairs in inverse orders. For example, take the two sets to be:

$$A = \{\mathsf{apple}, \mathsf{banana}, \mathsf{cherry}\}, \qquad B = \{\mathsf{peach}, \mathsf{banana}, \mathsf{apple}, \mathsf{strawberry}\}.$$

Then both $\langle \mathsf{apple}, \mathsf{banana} \rangle$ and $\langle \mathsf{banana}, \mathsf{apple} \rangle$ are elements of $A \times B$ and they are considered different:

$$\langle \mathsf{apple}, \mathsf{banana} \rangle \neq \langle \mathsf{banana}, \mathsf{apple} \rangle.$$

Notice, in passing, that a pair like $\langle \mathsf{peach}, \mathsf{cherry} \rangle$ is not an element of the Cartesian product, because peach is not an element of A and also because cherry is not an element of B:

$$\langle \mathsf{peach}, \mathsf{cherry} \rangle \notin A \times B.$$

On the other hand, the order of the elements is not important when we give a set by enumerating its elements. In that case we are only interested in what elements are in the set, not the way they are listed:

$$\{\mathsf{apple}, \mathsf{banana}\} = \{\mathsf{banana}, \mathsf{apple}\},$$
$$\{\mathsf{peach}, \mathsf{banana}, \mathsf{apple}, \mathsf{strawberry}\} = \{\mathsf{strawberry}, \mathsf{peach}, \mathsf{apple}, \mathsf{banana}\}.$$

4.5 Subsets

The subset relation, $\subseteq$, is a *partial order relation* on sets, that is, it satisfies the properties of reflexivity, antisymmetry and transitivity:

$$\begin{array}{rl} \textbf{Reflexivity:} & X \subseteq X \\ \textbf{Antisymmetry:} & (X \subseteq Y) \wedge (Y \subseteq X) \Rightarrow X = Y \\ \textbf{Transitivity:} & (X \subseteq Y) \wedge (Y \subseteq Z) \Rightarrow X \subseteq Z \end{array}$$

It is clearly not total: given two sets, it is not necessary that one of the two is contained in the other one.

A set can contain other sets, like a box containing smaller boxes. We can collect all the subsets of a given set X into one bigger set, called *the set of parts of* X, and denoted by $\mathcal{P}(X)$. Notice that the empty set $\varnothing$ and X itself are always subsets of X. For example:

$$\begin{array}{l} \mathcal{P}(\{\mathsf{apple}, \mathsf{banana}, \mathsf{cherry}\}) \\ \quad = \{ \ \varnothing, \{\mathsf{apple}\}, \{\mathsf{banana}\}, \{\mathsf{cherry}\}, \\ \qquad \{\mathsf{apple}, \mathsf{banana}\}, \{\mathsf{apple}, \mathsf{cherry}\}, \{\mathsf{banana}, \mathsf{cherry}\}, \\ \qquad \{\mathsf{apple}, \mathsf{banana}, \mathsf{cherry}\} \ \ \} \end{array}$$

Its elements are themselves sets. Be careful to keep this distinction in mind. For example we have:

$$\mathsf{banana} \notin \mathcal{P}(\{\mathsf{apple}, \mathsf{banana}, \mathsf{cherry}\})$$
$$\{\mathsf{banana}\} \in \mathcal{P}(\{\mathsf{apple}, \mathsf{banana}, \mathsf{cherry}\}).$$

The object banana is not an element of the set of parts, but the set containing just banana is.

4.6 Cardinality.

The number of elements of a set is called its *cardinality*. Given a set X, its cardinality is denoted by $|X|$. For example:

$$|\{\mathsf{apple}, \mathsf{banana}, \mathsf{cherry}\}| = 3$$
$$|\mathcal{P}(\{\mathsf{apple}, \mathsf{banana}, \mathsf{cherry}\})| = 8$$
$$|\{n \in \mathbb{N} \mid n \text{ is odd} \wedge n^2 + n \leq 100\}| = 5.$$

If a set has a finite number of elements, then its cardinality is a natural number. For infinite sets there are other "infinite numbers" that can be used to measure their cardinality, but we are not studying them in this course.

Suppose a set A has n elements, that is $|A| = n$, and another set B has m elements, that is $|B| = m$. What is the cardinality of $A \times B$. Remember that the elements of this Cartesian product are the pairs $\langle a, b\rangle$ with $a \in A$ and $b \in B$. We have n different possible choices for A; for each of these, we have m possible choices for b. So we have $n \times m$ different pairs:

$$|A \times B| = |A| \cdot |B|.$$

Suppose we know that a set X has n elements, that is, $|X| = n$. Can we compute the cardinality of $\mathcal{P}(X)$? Each subset of X is determined by choosing, for every element of X, whether it goes in the subset or not. So we have two choices for each object. We obtain all possible subsets by making all possible choices. Therefore, multiplying the possible choices for each element, we get 2^n subsets:

$$|\mathcal{P}(X)| = 2^{|X|}.$$

Another useful counting principle tells us about the relation of the cardinality of two sets A and B with their union and intersection. It is called *the inclusion-exclusion principle*:

$$|A \cup B| = |A| + |B| - |A \cap B|.$$

Convince yourself that this formula is correct by testing it on a few examples. The informal reasoning behind it is that if A and B have some elements in common, then when we compute $|A \cup B|$ we count those elements only once, but when we compute $|A|$ and $|B|$ we count them twice; so we need to subtract one of the double countings.

4.7 Functions

A function between two sets is a rule or a correspondence that associates to every element of the first set a unique element of the second set. For example, consider a correspondence between a set of three people and a set of fruit; it's a function that associates to every person her/his favourite fruit:

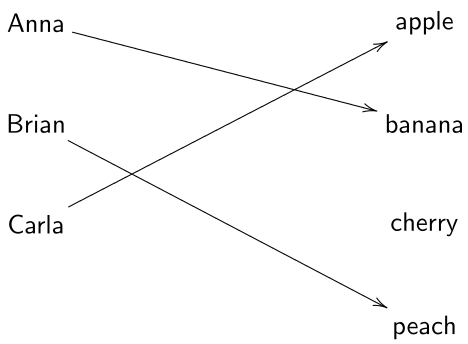

This defines a function, let's call it favourite, between two sets. We use the following notation to denote this fact:

$$\mathsf{favourite} : \{\mathsf{Anna}, \mathsf{Brian}, \mathsf{Carla}\} \to \{\mathsf{apple}, \mathsf{banana}, \mathsf{cherry}, \mathsf{peach}\}$$

$$\begin{aligned}&\mathsf{favourite}(\mathsf{Anna}) = \mathsf{banana}\\&\mathsf{favourite}(\mathsf{Brian}) = \mathsf{peach}\\&\mathsf{favourite}(\mathsf{Carla}) = \mathsf{apple}\end{aligned}$$

The set from which the function starts is called its *domain*, the one where it arrives is called its *codomain*.

When the domain is finite, as in the example above, we can define the function by just giving its values on every element, as we did. This is clearly impossible when the domain is an infinite set, for example the natural numbers. In that case the function needs to be defined by a formula or by some rule. Recursive definitions, which we studied a few lectures ago, are also a method to define a function.

Injective Functions. We say that a function is *injective* if every element of the domain is associated to a different result, that is, if no two elements share the same result. Formally we can define it as follows. Suppose $f : A \to B$:

$$f \text{ is injective} \quad \Leftrightarrow \quad x \neq y \Rightarrow f(x) \neq f(y) \text{ for all elements } x, y \in A.$$

The function favourite is injective because every person has a different favourite fruit. It is very useful to apply the definition in the contrapositive way: if two elements give the same result, then they must be equal:

$$f \text{ is injective} \quad \Leftrightarrow \quad f(x) = f(y) \Rightarrow x = y \text{ for all elements } x, y \in A.$$

Surjective Functions. We say that a function is *surjective* if every element of the codomain is the result of applying the function to some element of the domain, that is, if every element is the "target" of the function for some argument. Formally we can define it as follows.

$$f \text{ is surjective} \quad \Leftrightarrow \quad \text{for every } b \in B \text{ there is some } a \in A \text{ such that } f(a) = b.$$

The function $\mathsf{favourite}$ is NOT surjective because cherry is nobody's favourite fruit.

Consider instead the following function, going in the opposite direction, which associates to every fruit the person that owns it:

$$\mathsf{owner} : \{\mathsf{apple}, \mathsf{banana}, \mathsf{cherry}, \mathsf{peach}\} \to \{\mathsf{Anna}, \mathsf{Brian}, \mathsf{Carla}\}$$

$$\begin{aligned} &\mathsf{owner}(\mathsf{apple}) = \mathsf{Anna} \\ &\mathsf{owner}(\mathsf{banana}) = \mathsf{Carla} \\ &\mathsf{owner}(\mathsf{cherry}) = \mathsf{Anna} \\ &\mathsf{owner}(\mathsf{peach}) = \mathsf{Brian} \end{aligned}$$

This function is in fact surjective: every person owns at least one fruit. On the other hand it is not injective: when applied to apple and cherry it gives the same result, Anna.

Bijective Functions. A *bijective* function is one that is both injective and surjective. Neither of the two functions defined above is bijective: $\mathsf{favourite}$ isn't because it's not surjective and owner isn't because it's not injective.

Let's consider the following function f that associates a number smaller than 4 to fruit:

$$\mathsf{fn} : \{\mathsf{apple}, \mathsf{banana}, \mathsf{cherry}, \mathsf{peach}\} \to \{0, 1, 2, 3\}$$

$$\begin{aligned} &\mathsf{fn}(\mathsf{apple}) = 2 \\ &\mathsf{fn}(\mathsf{banana}) = 0 \\ &\mathsf{fn}(\mathsf{cherry}) = 3 \\ &\mathsf{fn}(\mathsf{peach}) = 1 \end{aligned}$$

This function is injective (no two elements give the same result) and surjective (every number in the codomain is the result for some argument), therefore it is bijective.

Let's look at three numerical functions now and determine which of the properties of injectivity, surjectivity and bijectivity they satisfy.

$$\begin{aligned} &\mathsf{f} : \mathbb{N} \to \mathbb{N} \\ &\mathsf{f}(n) = 2 \times n + 1 \end{aligned}$$

This function is injective: suppose $\mathsf{f}(n) = \mathsf{f}(m)$, that is, $2 \times n + 1 = 2 \times m + 1$; simple Arithmetic then tells us that $n = m$. On the other hand it is not surjective: the values 0 and 2 (and all other even numbers) are not results of f.

$$\begin{aligned} &\mathsf{half} : \mathbb{N} \to \mathbb{N} \\ &\mathsf{half}(n) = \lfloor n/2 \rfloor \end{aligned}$$

This function is not injective: $\mathsf{half}(0) = 0$ and $\mathsf{half}(1) = 0$, so two distinct arguments give the same result. But it is surjective: every number m can be obtained as the result of this function on a certain argument, by taking $n = 2 \times m$; in fact, $\mathsf{half}(2 \times m) = m$.

$$\begin{array}{ll} \mathsf{swap} : \mathbb{N} \to \mathbb{N} & \\ \mathsf{swap}(n) = n + 1 & \text{if } n \text{ is even} \\ \mathsf{swap}(n) = n - 1 & \text{if } n \text{ is odd} \end{array}$$

This function is both injective and surjective (I leave it to you to prove it). This fact can be clearly seen if we draw it using arrows:

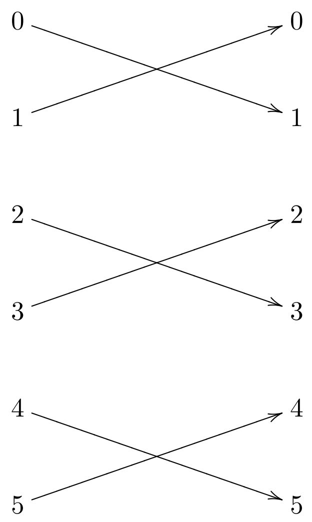

$\ldots$

Composition. Suppose we have two functions such that the codomain of the first coincides with the domain of the second: $f : A \to B$ and $g : B \to C$. We can *compose* them by applying one after the other: starting with an element of A we first compute f on it and then we compute g on the result that we obtained from the first step:

$$\begin{array}{ccccc} A & \xrightarrow{f} & B & \xrightarrow{g} & C \\ x & \mapsto & f(x) & \mapsto & g(f(x)) \end{array}$$

The result is a function from A to C that we denote by $g \circ f$. Attention: the first function to be applied, f, is written to the right of the second to be applied, g.

$$\begin{array}{l} g \circ f : A \to C \\ (g \circ f)(x) = g(f(x)) \end{array}$$

As an example, let's compute the composition of the *favourite fruit* and *owner* functions from above. The clearest way to do it is to represent them using arrows to show the associations and then "follow the arrows" to find the result of the composition. In our case we have: $A = \{\mathsf{Anna}, \mathsf{Brian}, \mathsf{Carla}\}$, $B = \{\mathsf{apple}, \mathsf{banana}, \mathsf{cherry}, \mathsf{peach}\}$ and $C = \{\mathsf{Anna}, \mathsf{Brian}, \mathsf{Carla}\}$.

$$A \xrightarrow{\mathsf{favourite}} B \qquad\qquad B \xrightarrow{\mathsf{owner}} C$$

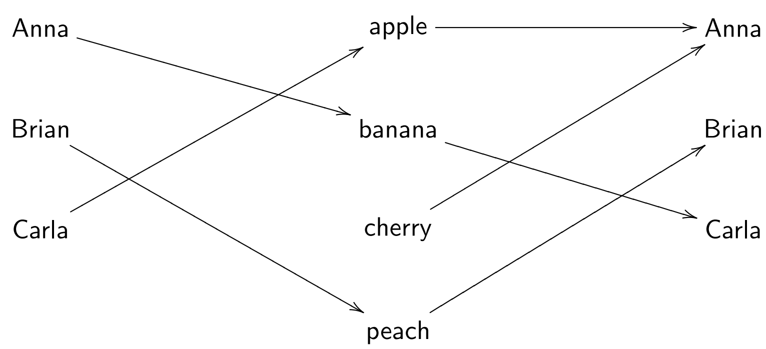

$$\mathsf{owner} \circ \mathsf{favourite} : \{\mathsf{Anna}, \mathsf{Brian}, \mathsf{Carla}\} \to \{\mathsf{Anna}, \mathsf{Brian}, \mathsf{Carla}\}$$

$$(\mathsf{owner} \circ \mathsf{favourite})(\mathsf{Anna}) = \mathsf{Carla}$$
$$(\mathsf{owner} \circ \mathsf{favourite})(\mathsf{Brian}) = \mathsf{Brian}$$
$$(\mathsf{owner} \circ \mathsf{favourite})(\mathsf{Carla}) = \mathsf{Anna}$$

For a numeric example, let's compose the two functions f and half on the natural numbers:

$$\mathsf{half} \circ \mathsf{f} : \mathbb{N} \to \mathbb{N}$$
$$(\mathsf{half} \circ \mathsf{f})(n) = \lfloor (2 \times n + 1)/2 \rfloor$$

In this case the expression for the composition can be simplified:

$$(\mathsf{half} \circ \mathsf{f})(n) = n.$$

The simplest of all functions is the one that doesn't do anything: it gives as result the argument itself. It is called the *identity function*:

$$\mathsf{id} : A \to A$$
$$\mathsf{id}(a) = a$$

Suppose we have two functions going in opposite directions: $f : A \to B$ and $g : B \to A$. We say that they are *inverse* of each other if both $g \circ f = \mathsf{id}$ and $f \circ g = \mathsf{id}$. Be careful: both compositions must be checked, in general they give different functions. In fact $g \circ f$ is a function from A to A, while $f \circ g$ is a function from B to B.

For example notice that, as we showed above, $\mathsf{half} \circ \mathsf{f} = \mathsf{id}$. But if we compose the functions the other way around we don't get the identity anymore: for example, $(\mathsf{f} \circ \mathsf{half})(2) = 3$. So half and f are not inverse of each other. (We may still say that half is a *left inverse* of f and that f is a *right inverse* of half).

If $f : A \to B$ has an inverse, this is denoted by f^{-1}.

The most important fact about bijections is that they are exactly those functions that can be inverted.

Theorem 10 *The following equivalence is true for every function $f : A \to B$:*

$$f \text{ is bijective} \quad \Leftrightarrow \quad f \text{ has an inverse.}$$

(We will not look at the proof of this theorem, but you may want to try to give it yourself.)

For example, we remarked earlier that the function fn is bijective. It is easy to compute its inverse by associating to each number the fruit that's mapped to it by fn:

$$\begin{aligned} \mathsf{fn}^{-1} &: \{0,1,2,3\} \to \{\mathsf{apple}, \mathsf{banana}, \mathsf{cherry}, \mathsf{peach}\} \\ \mathsf{fn}^{-1}(0) &= \mathsf{banana} \\ \mathsf{fn}^{-1}(1) &= \mathsf{peach} \\ \mathsf{fn}^{-1}(2) &= \mathsf{apple} \\ \mathsf{fn}^{-1}(3) &= \mathsf{cherry}. \end{aligned}$$

Cardinality of functions We indicate by $A \to B$ the set of all functions from A to B. If A and B are finite sets with cardinalities $n = |A|$ and $m = |B|$, we want to compute the number of functions between them, $|A \to B|$. To determine one function f we have to specify, for every element $a \in A$, what the value of $f(a)$ is, and we have m different choices for it. This is true for every element of A, and the choice that we make for each is independent from the others. So we have to chose one value out of m possibilities, n times. The total number of possible choices is then m^n:

$$|A \to B| = |B|^{|A|}.$$