



Universiteit
Leiden
The Netherlands

Reasoning about object-oriented programs: from classes to interfaces

Bian, J.

Citation

Bian, J. (2024, May 21). *Reasoning about object-oriented programs: from classes to interfaces*. Retrieved from <https://hdl.handle.net/1887/3754248>

Version: Publisher's Version

License: [Licence agreement concerning inclusion of doctoral thesis in the Institutional Repository of the University of Leiden](#)

Downloaded from: <https://hdl.handle.net/1887/3754248>

Note: To cite this publication please use the final published version (if applicable).

Chapter 3

Class specification and verification: a step-by-step guide

This chapter is designed to serve as a comprehensive tutorial, illustrating the formal methods for specifying and reasoning about program executions in Java. As a specific case study for `java.util.LinkedList`, we provide the source code and video material to facilitate a deeper understanding of the entire proof process.

This chapter is based on the following publications and artifacts:

- Hiep, H. A., **Bian, J.**, de Boer, F. S., de Gouw, S. (2020). A Tutorial on Verifying LinkedList Using KeY. Deductive Software Verification: Future Perspectives: Reflections on the Occasion of 20 Years of KeY, 221-245.
- **Bian, J.**, Hiep, H. A.: A Tutorial on Verifying LinkedList using KeY: Part 1 (2020). <https://doi.org/10.6084/m9.figshare.11662824>
- **Bian, J.**, Hiep, H. A.: A Tutorial on Verifying LinkedList using KeY: Part 2 (2020). <https://doi.org/10.6084/m9.figshare.11673987>
- **Bian, J.**, Hiep, H. A.: A Tutorial on Verifying LinkedList using KeY: Part 3a (2020). <https://doi.org/10.6084/m9.figshare.11688816>
- **Bian, J.**, Hiep, H. A.: A Tutorial on Verifying LinkedList using KeY: Part3b(2020). <https://doi.org/10.6084/m9.figshare.11688858>
- **Bian, J.**, Hiep, H. A.: A Tutorial on Verifying LinkedList using KeY: Part 3c (2020). <https://doi.org/10.6084/m9.figshare.11688870>
- **Bian, J.**, Hiep, H. A.: A Tutorial on Verifying LinkedList using KeY: Part3d(2020). <https://doi.org/10.6084/m9.figshare.11688984>
- **Bian, J.**, Hiep, H. A.: A Tutorial on Verifying LinkedList using KeY: Part 3e (2020). <https://doi.org/10.6084/m9.figshare.11688891>
- **Bian, J.**, Hiep, H. A.: A Tutorial on Verifying LinkedList using KeY: Part 4a (2020). <https://doi.org/10.6084/m9.figshare.11699178>
- **Bian, J.**, Hiep, H. A.: A Tutorial on Verifying LinkedList using KeY: (2020). <https://doi.org/10.6084/m9.figshare.11699253>
- **Bian, J.**, Hiep, H. A.: A Tutorial on Verifying LinkedList using KeY: Video Material (2020). <https://doi.org/10.6084/m9.figshare.c.4826589.v2>
- Hiep, H. A., **Bian, J.**, de Boer, F.S., de Gouw, S.: A Tutorial on Verifying LinkedList using KeY: Proof Files (2020). <https://doi.org/10.5281/zenodo.3613711>

3.1 Introduction

Software libraries, such as the Java standard library, serve as the building blocks of many programs that run on the devices of many users every day. Given their widespread usage, even minor errors in these libraries can result in serious consequences, making it essential that such libraries be thoroughly tested and debugged before being deployed. Therefore, it is crucial to rigorously test and debug these libraries before deploying them. The focus should be on preempting issues, rather than merely addressing them after they occur.

This chapter intends to show how we take an existing Java program that is part of the Java standard library and study it closely to increase our understanding of it. If we are only interested in showing the presence of an issue with the program, e.g. that it lacks certain functionality, it suffices to show an example run that behaves unexpectedly. But to conclude that no unexpected behavior ever results from running the program first requires a precise specification of what behavior one expects, and further requires a convincing argument that all possible executions of the program exhibit that behavior.

We take a formal approach to both specification and reasoning about program executions, allowing us to increase the reliability of our reached conclusions to near certainty. In particular, the specifications we write are expressed in the JML, and our reasoning is tool-supported and partially automated by KeY. To the best of the authors' knowledge, KeY is the only tool that supports enough features of the Java programming language for reasoning about real programs, of which its run-time behavior crucially depends on the presence of features such as dynamic object creation, exception handling, integer arithmetic with overflow, `for` loops with early returns, nested classes (both static and non-static), inheritance, polymorphism, erased generics, etc.

As a demonstration of applying KeY to state-of-the-art, real software, we focus on Java's `LinkedList` class, for two reasons. First, a (doubly-)linked list is a well-known basic data structure for storing and maintaining unbounded data and has many applications: for example, in Java's secure sockets implementation.¹ Second, it has turned out that there is a 20-year-old bug lurking in its program, that might lead to security-in-depth issues on large memory systems, caused by the overflow of a field that caches the length of the list [38]. Our specification and verification effort is aimed at establishing the absence of this bug from a repaired program.

This chapter is based on the results as described in paper [38]. That paper provides a high-level overview of the specification and verification effort of the linked list class as a whole, for a more general audience. In the present chapter, more technical details on how we use the KeY theorem prover are given, and we give more detail concerning the production of the proofs. In particular, this tutorial consists of the online repository of proof files [39], and online video material that shows how to (re)produce the proofs [40]: these are video recordings of the interactive sessions in KeY that demonstrates exactly what steps one could take to complete

¹e.g. see the JVM internal class `sun.security.ssl.HandshakeStateManager`.

the correctness proofs of the proof obligations generated by KeY from the method contracts.

We see how to set up a project and configure the KeY tool (Section 3.2). We then study the source code of the `LinkedList` to gain an intuitive understanding of its structure: how the instances look like, and how the methods operate (Section 3.3). We formulate, based on previous intuition, a *class invariant* in JML that expresses a property that is true of every instance (Section 3.4). An interesting property that follows from the class invariant, that is used as a separation principle, is described next (Section 3.5). To keep this presentation reasonably short, we further focus on the methods that pose the main challenges to formal verification, `add` and `remove`. We give a *method contract* for the `add` method that describes its expected behavior and we verify that its implementation is correct (Section 3.6). The difficulty level increases after we specify the `remove` method (Section 3.7), as its verification requires more work than for `add`. We study a deeper method that `remove` depends on, and finally use *loop invariants* to prove the correctness of `remove`.

3.2 Preliminaries

First we set up the project files needed to use KeY. The project files are available on-line [39]: these can be downloaded and include the KeY software version that we use. After unpacking the project files, we end up with the directory structure of Table 3.1.

```

linkedlist-tutorial
├── key-2.6.3.zip
├── LinkedList.key
├── src
│   ├── java
│   │   └── util
│   │       ├── LinkedList.old
│   │       └── LinkedList.java
├── jre
│   └── java
│       └── ...
└── proof
    └── ...

```

Table 3.1: Directory structure of project files. The `src` directory contains the Java classes we want to specify and verify. The `jre` directory contains stub files, with specifications of unrelated classes. The `LinkedList.java` file is the source file we end up with after following this tutorial. The `proof` directory contains the completed proofs.

The original source file of `LinkedList.old` was obtained from OpenJDK version `jdk8-b132`. The original has been pre-processed: generic class parameters are removed, and all methods and nested classes irrelevant to this tutorial are removed. Both the removal of generics and the stub files in the `jre` folder were generated

Max. rule applications	1000	JavaCard	Off
Stop at	Default	Strings	On
Proof splitting	Delayed	Assertions	Safe
Loop treatment	Invariant	BigInt	On
Block treatment	Contract	Initialization	Disable static ...
Method treatment	Contract	Int Rules	Java semantics
Dependency contract	On	Integer Simpl. Rules	Full
Query treatment	Off	Join Generate	Off
Expand local queries	On	Model Fields	Treat as axiom
Arithmetic treatment	Basic	More Seq. Rules	On
Quantifier treatment	No splits	Permissions	Off
Class axiom rule	Off	Program Rules	Java
Auto induction	Off	Reach	On
User-defined taclets	All off	Runtime Exceptions	Ban
		Sequences	On
		Well-def. Checks	Off
		Well-def. Operator	L

Table 3.2: Proof search strategy and taclet options

automatically, using the Eclipse extensions for KeY. Repeating those steps is not necessary here.

Throughout the next sections, we modify the original source file and add annotations to formally specify its behavior, and helper methods for presenting intermediary lemmas. The annotations are usual Java comments and thus ignored when the file is read by a Java compiler. The helper methods introduce slight performance overhead (of calling a method that performs no operations, and immediately returning from it); it is clear that these do not change the original behavior of the program.

To produce proofs in KeY, the first step is to set up KeY’s proof strategy and taclet options. This has to be done only once, as these taclet settings are stored per computer user. Sometimes, KeY overwrites or corrupts these settings if different versions are used. To ensure KeY starts in a fresh state, one can remove the `.key` directory from the user’s home directory, and clean out preferences from the `.java/.userPrefs` directory by deleting the `de/uka/ilkd/key` hierarchy containing `prefs.xml` files.¹ Now start up KeY, and the example selection screen appears (if not, selecting File▷Load Example opens the same screen). Load any example, to enter proof mode.

First, we set up a proof strategy: this ensures the steps as done in the videos can be reproduced. On the left side of the window, change the settings in the Proof Search Strategy tab to match those of the left column in Table 3.2. We ensure to use particular taclet rules that correctly model Java’s integer overflow semantics.

¹On Windows, the preferences are instead stored in the Windows Registry. Use the `regedit` tool and clean out under `HKEY_CURRENT_USER\Software\JavaSoft\Prefs` or `HKEY_CURRENT_USER\Software\Wow6432Node\JavaSoft\Prefs` the same hierarchy. On Mac OS, open a terminal, change the directory to `~/Library/Preferences`, and delete `de.uka.ilkd.plist`

Select Options▷ Taclet Options, and configure the options as in the right column in Table 3.2. The taclet options become effective after loading the next problem. We do that now: the main proof file `LinkedList.key` can be loaded, and a Proof Management window opens up, showing a class hierarchy and its methods. A method is not shown when no specifications are present. After giving the specifications below, more methods can be selected in this window.

3.3 Structure and behavior of `java.util.LinkedList`

In this section we walk through part of the source code of Java’s linked list: see the `LinkedList.old` file in the artifact. Throughout this tutorial, we add annotations at the appropriate places. We finally obtain the `LinkedList.java` file.

Our `LinkedList` class has three attributes and a constructor (Listing 3.1): a `size` field, which stores a cached number of elements in the list, and two fields that store a reference to the first and last `Node`. The public constructor contains no statements: thus it initializes `size` to zero, and `first` and `last` to `null`.

```
package java.util;

public class LinkedList {

    transient int size = 0;
    transient Node first;
    transient Node last;

    public LinkedList() {}
```

Listing 3.1: The `LinkedList` class fields and constructor (begin of file).

The linked list fields are declared `transient` and package private. The `transient` flag is not relevant to our verification effort. The reason the fields are declared package private seems to prevent generating accessor methods by the Java compiler. However, in practice, the fields are treated as if they were private.

The `Node` class is defined as a private static nested class to represent the containers of items stored in the list (Listing 3.2). A static nested private class behaves like a top-level class, except that it is not visible outside the enclosing class. The nodes are doubly linked, that is, each node is connected to its preceding node (through field `prev`) and succeeding node (through field `next`). These fields contain `null` in case no preceding or succeeding node exists. The data itself is contained in the `item` field of a node.

```
private static class Node{
    Object item;
    Node next;
    Node prev;

    Node(Node prev, Object element, Node next) {
        this.item = element;
        this.next = next;
        this.prev = prev;
    }
}
```

Listing 3.2: The Node nested class fields and constructor (end of file).

The method `add` for adding new elements to the list, takes one argument, the item to add (Listing 3.3). According to Java informal documentation for the `Collection`, this method is designed to always return `true`. The implementation immediately calls `linkLast`.

```
// implements java.util.Collection.add
public boolean add(Object e) {
    linkLast(e);
    return true;
}
```

Listing 3.3: The method `add`.

The method `remove` for removing elements, also takes one argument, the item to remove (Listing 3.4). If that item was present in the list, then its first occurrence is removed and `true` is returned; otherwise, if the item was not present, then the list is not changed and `false` is returned.

```
// implements java.util.Collection.remove
public boolean remove(Object o) {
    if (o == null) {
        for (Node x = first; x != null; x = x.next) {
            if (x.item == null) {
                unlink(x);
                return true;
            }
        }
    } else {
        for (Node x = first; x != null; x = x.next) {
            if (o.equals(x.item)) {
                unlink(x);
                return true;
            }
        }
    }
}
```

```

    }
    }
    return false;
}

```

Listing 3.4: The method `remove`.

The presence of the item depends on whether the argument of the `remove` method is `null` or not. If the argument is `null`, then it searches for the first occurrence of a `null` item in the list. Otherwise, it uses the `equals` method, that every `Object` in Java has, to determine the equality of the argument with respect to the contents of the list. The first occurrence of an item that is considered equal by the argument is then returned. In both cases, the execution walks over the linked list, from the first node until it has reached the end. In the case that the node was found that contains the first occurrence of the argument, an internal method is called: `unlink`, and afterward `true` is returned.

Observe that the linked list is not modified if `unlink` is not called. Although not immediately obvious, this requires that the `equals` method of every object cannot modify our current linked list, for the duration of the call to `remove`. When the `remove` method is called with an item that is not contained in the list, either loop eventually exits, and `false` is returned.

The internal method `linkLast` changes the structure of the linked list (Listing 3.5). After performing this method, a new node has been created, and the `last` field of the linked list now points to it. To maintain structural integrity, also other fields change: if the linked list was empty, the `first` field now points to the new, and only, node. If the linked list was not empty, then the new last node is also reachable via the former last node's `next` field. It is always the case that all items of a linked list are reachable through the `first` field, then following the `next` fields, and also for the opposite direction.

```

void linkLast(Object e) {
    final Node l = last;
    final Node newNode = new Node(l, e, null);
    last = newNode;
    if (l == null) first = newNode;
    else l.next = newNode;
    size++;
}

```

Listing 3.5: The internal method `linkLast`.

The internal method `unlink` is among the most complex methods that alter the structure of a linked list (Listing 3.6). The method is used only when the linked list is not empty. Its argument is a node, necessarily one that belongs to the linked list. We first store the fields of the node in local variables: the old item, and the next and previous node references.

```
Object unlink(Node x) {  
    final Object element = x.item;  
    final Node next = x.next;  
    final Node prev = x.prev;  
    if (prev == null) {first = next;}  
    else {  
        prev.next = next;  
        x.prev = null;  
    }  
    if (next == null) {last = prev;}  
    else {  
        next.prev = prev;  
        x.next = null;  
    }  
    x.item = null;  
    size--;  
    return element;  
}
```

Listing 3.6: The internal method `unlink`.

After the method returns, the argument fields are all cleared, presumably to help the garbage collector. However, other fields also change: if the argument is the first node, the `first` field is updated; if the argument is the last node, the `last` field is updated. The predecessor or successor fields `next` and `prev` of other nodes are changed to maintain the integrity of the linked list: the successor of the unlinked node becomes the successor of its predecessor, and the predecessor of the unlinked node becomes the predecessor of its successor.

3.3.1 Expected and unexpected method behavior

We draw pictures of linked list instances to understand better how the structure looks like over time. In Figure 3.1, we see three linked list instances. The left-most linked list is an object without any items: its `size` is zero. When we perform `add` with some item (it is not important what item), a new node is created and the first and last pointers are changed to point to the new node. Now the `prev` and `next` fields of the new node are `null`, indicating that there is no other node before and there is no other node after it. Also, the `size` field is increased by one. Adding another item further creates another node, that is linked up to the previous node properly; the `last` field is then pointed to the newly created node. In the third instance, suppose we would perform `remove` with the first item. We would then have to unlink the node, see the code of `unlink` in Listing 3.6: the new value of `first` becomes the value of `next` which is the last node, and the value of `prev` of the succeeding node becomes the value of `prev` of the node that is unlinked, which is `null`. We thus end up in a similar situation as the second instance (except for the item that may be different). Removing the last item brings us back to the situation depicted by the first instance.




Figure 3.1: Three example linked lists: empty, with a chain of one node, and with a chain of two nodes. Items themselves are not shown.

An important aspect of the implementation of our linked list is the cached `size` field: it represents the number of nodes that form a chain between `first` and `last`. It turns out an overflow may happen under certain conditions [38]. Consider two facts: Java integer primitives are stored in signed 32-bit registers, and it is possible to create a chain that is larger than the maximum positive value that can be stored in such fields, $2^{31} - 1$. Now, the cached size and the actual size no longer correspond. In the methods we have seen above, this seems to be no issue. But another method of the linked list may be used to demonstrate the key problem: `toArray`. The intention of `toArray` is to give back an array containing all the items of the list (see Listing 3.7). There are two problems: after the overflow has occurred and the `size` is negative, the `toArray` throws an unexpected `NegativeArraySizeException`. Also, after adding more items that bring the size back to a positive integer, e.g. adding $2^{32} + 1$ items in total, the array is of the wrong (positive) size and cannot contain *all* items of the list, and an `IndexOutOfBoundsException` is thrown.

```
public Object[] toArray() {
    Object[] result = new Object[size];
    int i = 0;
    for (Node x = first; x != null; x = x.next)
        result[i++] = x.item;
    return result;
}
```

Listing 3.7: The method `toArray` has unexpected behavior.

3.3.2 Integer overflow bug

The `List` interface is part of the collection framework. The method `int size()` provided by class `LinkedList`, as shown in Listing 3.8.

```
/** Returns the number of elements in this list. */
private int size() {
    return size;
}
```

Listing 3.8: The `size` method.

Its (informal) specification states that it returns the number of elements in the list. This method is further specified by the interface `Collection`, where additionally it is stated that if the collection contains more than `Integer.MAX_VALUE` elements, then

`size()` returns `Integer.MAX_VALUE` (Java integers are bounded). Adding an element unconditionally increments the size, and thus can overflow, resulting in a negative size. This error breaks the contract of `size()` and propagates in various ways. For example, sorting a linked list first copies the list to a new array `Object[size]` by the method `toArray`, which thus may crash with a `NegativeArraySizeException`. Further, many methods of the `List` interface use an integer index to refer to an element. For very large lists with more than `Integer.MAX_VALUE` elements, it is unclear what the behavior should be. This is perhaps not surprising: triggering these bugs requires at least 2^{31} elements (i.e., at least 64 GiB of memory), and each element has 2^{32} values; the state space is too large to explore for fully automated analysis methods. Similarly, the smallest test cases that triggered the TimSort crash consist of inputs of 2^{26} (≈ 67 million) elements, and each element has 2^{32} possible values. Such cases are clearly out of range for fully automated methods, such as model checking, which are typically based on finite abstractions of limited size.

As this case study involves real software, we have performed an analysis of existing Java code to estimate the use of `LinkedList` and the potential impact of a bug. A basic analysis of at least 140,000 classes,¹ found by sampling Java packages on Maven Central and in software distributions such as Red Hat Linux, reveals 1,677 cases of direct use of the `LinkedList` class where a constructor of `LinkedList` is invoked. As an overestimation of the potential reach of such instances, we find at least 37,000 usage call sites where some method of the `Collection`, `List`, or `Queue` interface is called. It is infeasible for us to analyze for each constructor call site where its resulting instance will be used. However, some usage of the `LinkedList` class occurs in potentially security-sensitive contexts, such as the library for checking certificates used by the Java secure socket implementation and checking the authenticity of the source of dynamically loaded bytecode.

3.3.3 Verification goal

We thus revise the source code and add a method that implements an overflow check (see Listing 3.9). The intention is that the overflow is signaled before it occurs, by throwing an exception. This ensures that the integrity of the linked list is always maintained. We modify the `add` method and perform a call to this `checkSize()` method before invoking `linkLast`.

```
// new method, not in original LinkedList
private void checkSize() {
    if (size == Integer.MAX_VALUE)
        throw new IllegalStateException(...);
}
```

Listing 3.9: A new internal method `checkSize`.

Our aim in this tutorial is to keep the discussion general enough, without losing interesting particular details. We apply step-wise refinements to our arguments,

¹We have used the Rascal programming language [41] for loading `.jar`-files and analyzing their class file contents that revealed method call sites. We thank Thomas Degueule for his help in setting up this experiment.

where we start with a higher-level intuition and drill down on technical details as they become relevant. The reader can always see the video material; but, a high-level intuition seems essential for following along.

Our specification and verification goals comprise two points:

1. Specification captures the ‘intended’ behavior of the methods with respect to their structural properties: in particular, we abstract away from all properties about the contents of the linked list.
2. Verification ensures the overflow bug no longer happens in the revised linked list: the actual number of nodes and the cached size are always the same.

The first point depends on the aim of a verification attempt. Are we using the specification to verify the correctness of clients of the linked list? The properties of the contents of a linked list are essential. But, for our purpose of showing the absence of an overflow bug, we abstract away some properties of the contents of linked lists.

The second point deserves an introduction: how can we be sure that in every linked list the number of nodes and the value stored in the `size` field are the same? Can we keep the number of nodes bounded by what a Java integer can represent? Keeping this number in a ghost field is not sufficient, since the number of nodes depends on the success of a `remove` call: removing an item not present in the linked list should not affect its size, while removing an item that is present decreases its size by one. We refine: we could keep track of the items that are stored by the linked list. The structure to collect these items cannot be a set of items, since we could have duplicate items in the list. A multiset of items, the contents of a linked list, seems right: the size field of the linked list must be the same as the size of its contents, and the remove method is only successful if its argument is contained before the call.

However, working with multisets is quite unnatural, as the `remove` method removes the first item in the list. That can be refined by specifying the contents as a sequence of items instead. Although this could work in principle, a major difficulty when verifying the `remove` method is to give an argument as to why the method terminates. This requires knowledge of the linking structure of the nodes. We could relate the sequence of items to a traversal over nodes, saying that the first item of a linked list is found in the node by traversing `first`, and the item at index $0 < i < n$ is found in the node by traversing `first` and then `next` for $i - 1$ times. Formalizing this seems quite difficult, and as we shall see, not even possible in first-order logic.

Hence, we end up with our last refinement: we keep a sequence of nodes in a ghost field. From this sequence, one obtains the sequence of items. We relate the sequence of nodes to the linked list instance and require certain structural properties to hold the nodes in the sequence. The length of this sequence is the actual number of nodes, that we show to be equal to the cached size.

3.4 Formulating a class invariant

We now formalize a class invariant, thereby characterizing all linked list instances. We focus on unbounded linked lists first, as these are the structures we intend to model: so most properties are expressed using KeY's unbounded integer type `\big-int`. Only at the latest, we restrict the size of each linked list to a maximum, as a limitation imposed by the implementation. The setting in which to do our characterization is multi-sorted first-order logic. This logic is presented in a simplified form, leaving out irrelevant details (such as the heap): the full logic used by KeY is described in Chapter 2 of [29].

Consider the following sorts or type symbols: *LinkedList* for a linked list, *Node* for a node, *Object* for objects, and *Null* for `null` values. We have a type hierarchy, where *Null* is related to *LinkedList* and *Node*, and *LinkedList* and *Node* are both related to *Object*. This means that any object of sort *Null* is also an object of sort *LinkedList* and *Node*. Moreover, every object that is a linked list is also of type *Object*, and similar for nodes. We have the following signature: $first : LinkedList \rightarrow Node$ and $last : LinkedList \rightarrow Node$ for the `first` and `last` fields of linked lists, and $prev : Node \rightarrow Node$, $item : Node \rightarrow Object$, and $next : Node \rightarrow Node$ for the `prev`, `item` and `next` fields of nodes. Further, we assume there is exactly one object of *Null* sort, which is the `null` constant, for which the above functions are left undefined: `null` is a valid object of the *LinkedList* and *Node* Java types, but one may not access its fields.

We search for an axiomatization that characterizes linked list instances. One can find these axioms by trial and error. We start listing some (obvious) axioms:

1. $\forall x^{LinkedList}; (x \neq \text{null} \rightarrow (first(x) \neq \text{null} \leftrightarrow last(x) \neq \text{null}))$
Every linked list instance either has both first and last set to `null`, or both point to some (possibly different) node.
2. $\forall x^{LinkedList}; (x \neq \text{null} \rightarrow (first(x) \neq \text{null} \rightarrow prev(first(x)) = \text{null}))$
The predecessor of the first node of a linked list is set to `null`.
3. $\forall x^{LinkedList}; (x \neq \text{null} \rightarrow (last(x) \neq \text{null} \rightarrow next(last(x)) = \text{null}))$
The successor of the last node of a linked list is set to `null`.
4. $\forall x^{Node}; (x \neq \text{null} \rightarrow (prev(x) \neq \text{null} \rightarrow next(prev(x)) = x))$
Every node that has a predecessor, must be the successor of that predecessor.
5. $\forall x^{Node}; (x \neq \text{null} \rightarrow (next(x) \neq \text{null} \rightarrow prev(next(x)) = x))$
Every node that has a successor, must be the predecessor of that successor.

These axioms are not yet sufficient: consider a linked list, in which its first and last nodes are different and both have neither a predecessor nor a successor. This linked list should not occur: intuitively, we know that the nodes between first and last are all connected and should form a doubly-linked 'chain'. Moreover, for every linked list, this chain is necessarily finite: one can traverse from first to last by following the next reference a finite number of times. This leads to a logical difficulty.

Proposition 1. It is not possible to define the reachability of nodes of a linked list

in first-order logic.

Proof 1. Let x be a linked list and y a node: there is no formula $\phi(x, y)$ that is true if and only if $\text{next}^i(\text{first}(x)) = y$ for some integer $i \geq 0$.¹ Suppose towards contradiction that there is such a formula $\phi(x, y)$. Now consider the infinite set Δ of first-order formulas $\{\phi(x, y)\} \cup \{\neg(\text{next}^i(\text{first}(x)) = y) \mid 0 \leq i\}$. Let Γ be an arbitrary finite subset of Δ . Consider that there must exist some j such that $\Gamma \subseteq \{\phi(x, y)\} \cup \{\neg(\text{next}^i(\text{first}(x)) = y) \mid 0 \leq i < j\}$, so we can construct a linked list with j nodes, and we interpret x as that linked list and y as the last node. Clearly $\phi(x, y)$ is true as the last node is reachable, and all $\neg(\text{next}^i(\text{first}(x)) = y)$ is true for all $0 \leq i < j$ because j is not reachable within i steps from the first node. Since Γ is arbitrary, we have established that all finite subsets of Δ have a model. By compactness, Δ must have a model too. However, that is contradictory: no such model for Δ can exist, as neither $\phi(x, y)$ and $\text{next}^i(\text{first}(x)) \neq y$ for all integers $0 \leq i$ can all be true.

We extend our signature to include other sorts: sequences and integers. These sorts are interpreted in the standard model. A schematic rule to capture integer induction is included in KeY (see [29, Section 2.4.2]). Sequences (see [29, Chapter 5.2]) have a non-negative integer length n and consist of an element at each position $0 \leq i < n$. We write $\sigma[i]$ to mean the i th element of sequence σ , and $\ell(\sigma)$ to mean its length n .

Intuitively, each linked list consists of a sequence of nodes between its first and last node. Let $\text{instanceof}_{\text{Node}} : \text{Object}$ be a built-in predicate that states that the object is not `null` and of sort *Node*. A *chain* is a sequence σ such that:

- (a) $\forall i^{\text{int}}; (0 \leq i < \ell(\sigma) \rightarrow \text{instanceof}_{\text{Node}}(\sigma[i]))$
All its elements are nodes and not `null`
- (b) $\forall i^{\text{int}}; (0 < i < \ell(\sigma) \rightarrow \text{prev}(\sigma[i]) = \sigma[i - 1])$
The predecessor of the node at position i is the node at position $i - 1$
- (c) $\forall i^{\text{int}} : (0 \leq i < \ell(\sigma) - 1 \rightarrow \text{next}(\sigma[i]) = \sigma[i + 1])$
The successor of the node at position i is the node at position $i + 1$

Let $\phi(\sigma)$ denote the above property that σ is a chain. If $\ell(\sigma) = 0$ then $\phi(\sigma)$ is vacuously true: the empty sequence is thus a chain. We now describe properties $\psi_1(\sigma, x)$ and $\psi_2(\sigma, x)$ that relate a chain σ to a linked list x . These denote the following intuitive properties: there is no first and last node and the chain is empty, or the chain is not empty and the first and last node are the first and last elements of the chain.

$$\begin{aligned} \psi_1(\sigma, x) &\equiv (\ell(\sigma) = 0 \wedge \text{first}(x) = \text{last}(x) = \text{null}) \\ \psi_2(\sigma, x) &\equiv (\ell(\sigma) > 0 \wedge \text{first}(x) = \sigma[0] \wedge \text{last}(x) = \sigma[\ell(\sigma) - 1]) \end{aligned}$$

6. $\forall x^{\text{LinkedList}}; (x \neq \text{null} \rightarrow \exists \sigma^{\text{sig}}; (\phi(\sigma) \wedge (\psi_1(\sigma, x) \vee \psi_2(\sigma, x))))$

Every linked list necessitates the existence of a chain of either property

¹ next^i is not a function symbol in first-order logic but an abbreviation of a finite term built by iteration of i times next , where $\text{next}^0(x) = x$ and $\text{next}^i(x) = \text{next}(\text{next}^{i-1}(x))$ for all $i > 0$.

3. CLASS SPECIFICATION AND VERIFICATION: A STEP-BY-STEP GUIDE

Further, we require that the size field of the linked list and the length of the chain are the same: this property is essential to our verification goal. The size field is modeled by the function $size : LinkedList \rightarrow int$, and we require its value (1) to equal the length of the chain, and (2) to be bounded by the maximum value stored in a 32-bit integer. In formulating the above properties in JML, we skolemize the existential quantifier using a ghost field: see Listing 3.10. This has the additional benefit that we can easily refer to the chain ghost field in specifications.

```
/*@ nullable @*/ transient Node first;
/*@ nullable @*/ transient Node last;
//@ private ghost \seq nodeList;
/*@ invariant
  @ nodeList.length == size \&\&
  @ nodeList.length <= Integer.MAX_VALUE \&\&
  @ (\forallall \bigint i; 0 <= i < nodeList.length;
  @ nodeList[i] instanceof Node) \&\&
  @ ((nodeList == \seq_empty \&\& first == null \&\& last == null)
  @ || (nodeList != \seq_empty \&\& first != null \&\&
  @ first.prev == null \&\& last != null \&\&
  @ last.next == null \&\& first == (Node)nodeList[0] \&\&
  @ last == (Node)nodeList[nodeList.length-1]) \&\&
  @ (\forallall \bigint i; 0 < i < nodeList.length;
  @ ((Node)nodeList[i]).prev == (Node)nodeList[i-1]) \&\&
  @ (\forallall \bigint i; 0 <= i < nodeList.length-1;
  @ ((Node)nodeList[i]).next == (Node)nodeList[i+1]);
  @*/
```

Listing 3.10: The class invariant of `LinkedList` expressed in JML.

The class invariant is implicitly required to hold for the `this` object when invoking methods on a linked list instance. In particular, for the constructor of the linked list, the class invariant needs to be established after it returns. In Listing 3.11, we state that the constructor always constructs a linked list instance for which its chain is empty. The proof of correctness follows easily: at construct time, the fields (including the ghost field) of the linked list instance are initialized with their default values. This means the size is zero, and the first and last references are `null`, and the ghost field is the empty sequence.

```
/*@
  @ public normal behavior
  @ ensures nodeList == \seq_empty;
  @*/
public LinkedList() {}
```

Listing 3.11: The method contract of the constructor of `LinkedList` in JML.

For verifying the constructor above, see the video [42, 0:23–0:53], where the relevant video material is between timestamps 0:23 and 0:53.

3.5 The acyclicity property

An interesting consequence of the class invariant is the property that traversal of only `next` fields is acyclic. In other words, following only `next` references of any node that is present in a chain never reaches itself. The acyclicity property implies there is a number of times to follow the `next` reference until the last node is reached. For the last node, this number is zero (the last node is already reached). A symmetric property holds for `prev` too.

We logically specify the acyclicity property as follows. Let σ be the chain of a non-empty linked list x for which the class invariant holds. The following holds:

$$\forall i^{\text{int}}; (0 \leq i < \ell(\sigma) - 1 \rightarrow \forall j^{\text{int}}; (i < j < \ell(\sigma) \rightarrow \sigma[i] \neq \sigma[j]))$$

Let n abbreviate $\ell(\sigma)$. By contradiction: assume there are two indices, $0 \leq i < j < n$, such that the nodes $\sigma[i]$ and $\sigma[j]$ are equal. Then it must hold that for all k such that $j \leq k < n$, the node $\sigma[k]$ is equal to the node $\sigma[k - (j - i)]$: by induction on k . Base case: if $k = j$, then node $\sigma[j]$ and node $\sigma[j - (j - i)]$ are equal by assumption, since $\sigma[j - (j - i)] = \sigma[i]$. Induction step: suppose node at $\sigma[k]$ is equal to the node at $\sigma[k - (j - i)]$. We must show if $k + 1 < n$ then node $\sigma[k + 1]$ equals node $\sigma[k + 1 - (j - i)]$. This follows from the fact that $\sigma[k + 1] = \text{next}(\sigma[k])$ and $\sigma[k + 1 - (j - i)] = \text{next}(\sigma[k - (j - i)])$ for $k < n - 1$, since σ is a chain and the chain property (c) of last section. Now we have established, for all $j \leq k < n$, node $\sigma[k]$ equals node $\sigma[k - (j - i)]$. In particular, this holds when k is $n - 1$, the index of the last node: so we have $\sigma[n - 1] = \sigma[n - 1 - (j - i)]$. Since the difference $(j - i)$ is positive, we know $\sigma[n - 1 - (j - i)]$ is not the last node. By the linked list property 3 we have $\text{next}(\text{last}(x)) = \text{null}$ and by $\psi_2(\sigma, x)$ we have $\text{last}(x) = \sigma[n - 1]$: so we have $\text{next}(\sigma[n - 1]) = \text{null}$. By the chain properties (c) and (a) we have $\text{next}(\sigma[n - 1 - (j - i)]) = \sigma[n - (j - i)]$ and $\text{instanceof}_{\text{Node}}(\sigma[n - (j - i)])$, respectively. From the latter we know $\sigma[n - (j - i)] \neq \text{null}$. So we have $\text{next}(\sigma[n - 1 - (j - i)]) \neq \text{null}$. But this is a contradiction: if nodes $\sigma[n - 1]$ and $\sigma[n - 1 - (j - i)]$ are equal then their `next` fields must also have equal values, but $\text{next}(\sigma[n - 1]) = \text{null}$ and $\text{next}(\sigma[n - 1 - (j - i)]) \neq \text{null}$!

For verifying the lemma as formalized in Listing 3.12, see the video [43].

```

/*@ public normal behavior
  @ requires true;
  @ ensures ( $\backslash \text{forall } \backslash \text{bigint } i;$ 
    @  $0 \leq i < (\backslash \text{bigint})\text{nodeList.length} - (\backslash \text{bigint})1;$ 
    @ ( $\backslash \text{forall } \backslash \text{bigint } j; i < j < \text{nodeList.length}; \text{nodeList}[i] \neq \text{nodeList}[j]$ ));
  @*/
private /*@ strictly_pure @*/ void lemma_acyclic() {}

```

Listing 3.12: The method of a lemma added to `LinkedList` expressed in JML.

3.6 The add method

Due to the revision of the source code, the `add` method now calls `checkSize` first (see Listing 3.9) to ensure that the `size` field does not overflow when we add another

item. This means that the `add` method has two expected behaviors: the normal behavior when the length of the linked list is not yet at its maximum, and the exceptional behavior when the length of the linked list is at its maximum.

In the normal case, we expect the `add` method to add the given argument as an item to the linked list. Thus the sequence of nodes must become larger. We further specify the position where the item is added: at the end of the list. If `add` returns normally, it returns `true`. In the exceptional case, we expect that an exception is thrown. We formalize the contract for `add` in Listing 3.13.

```

/*@ private normal_behavior
  @ requires nodeList.length + (\bigint)1 <= Integer.MAX_VALUE;
  @ ensures
  @ nodeList == \seq_concat(\old(nodeList),
  @ \seq_singleton(nodeList[nodeList.length-1])) \&\&
  @ ((Node)nodeList[nodeList.length-1]).item == e \&\& \result;
  @ private exceptional_behavior
  @ requires nodeList.length == Integer.MAX_VALUE;
  @ signals_only IllegalStateException;
  @ signals (IllegalStateException e) true;
  @*/
public boolean add(/*@ nullable @*/ Object e) {
    checkSize(); // new
    linkLast(e);
    return true;
}

```

Listing 3.13: The `add` method with its method contract expressed in JML.

Since the `add` method calls the deeper methods `checkSize` and `linkLast`, we may employ their method contracts when verifying this method. So, before we verify `add`, we specify and verify these methods first.

We expect `checkSize` to throw an exception if the length of the linked list is too large to add another element, and it returns normally otherwise: see Listing 3.14 for its specification. Verification of `checkSize` in both normal and exceptional cases is done automatically by KeY, as can be seen in [42, 0:54–1:24].

```

/*@ private exceptional_behavior
  @ requires nodeList.length == Integer.MAX_VALUE;
  @ signals_only IllegalStateException;
  @ signals (IllegalStateException e) true;
  @ private normal_behavior
  @ requires nodeList.length != Integer.MAX_VALUE;
  @ ensures true;
  @*/

```

Listing 3.14: The method contract in JML of the `checkSize` method.

For the `linkLast` method, we assume that the length of the linked list is smaller than its maximum length, so we can safely add another node without causing an overflow

of the size field. When adding a new node, the resulting chain now is an extension of the previous chain, and additionally the class invariant holds afterwards—this is an implicit post-condition. Since we modify the chain, we need a **set** annotation that changes the ghost field.

```

/*@
  @ normal_behavior
  @ requires nodeList.length + (\bigint)1 <= Integer.MAX_VALUE;
  @ ensures nodeList == \seq_concat(\old(nodeList),
  @ \seq_singleton(nodeList[nodeList.length-1])) \&&
  @ ((Node)nodeList[nodeList.length-1]).item == e;
  @*/
void linkLast(/*@ nullable @*/ Object e) {
  final Node l = last;
  final Node newNode = new Node(l, e, null);
  last = newNode;
  if (l == null) first = newNode;
  else l.next = newNode;
  size++;
  //@ set nodeList = \seq_concat(nodeList, \seq_singleton(last));
}

```

Listing 3.15: The `linkLast` method with its method contract expressed in JML.

The verification of this method is no longer fully automatic, see [42, 1:25–6:52].

Observe that there are two different situations we have to deal with: either the linked list was empty, or it was not. If the linked list was empty, then `last` is `null`, and we not only set the `last` field but also the `first`. Otherwise, if the linked list was not empty, we update the former last node to set its `next` field. The challenge is to prove that the class invariant holds after these heap updates, knowing that the class invariant holds in the before heap. The main insight is that the creation of a new node does not alias with any of the existing nodes, and that the modification of the `next` field only affects the old last node. Intuitively, we have a proof situation with two heaps as depicted in Figure 3.2.

The properties (b) on page 33, that fixes `prev` fields to point to the previous node in the sequence, and (c), that fixes `next` fields to point to the next node in the sequence, of the chain are the remaining goals in [42, 3:58]. Proving (b) is straightforward if one makes a distinction between the old nodes and the new node. Proving (c) in the ‘heap after’ involves two cases: either the index is between 0 and less than $\ell(\sigma) - 2$, or it used to be the last node and now has index $\ell(\sigma) - 2$. In the former case, the heap update has no effect, as we can show that these nodes are separate from the old last node because they differ in the old value of the next field. In the latter case, the heap update can be used to prove the property directly.

Finally, we can verify the `add` method: see [42, 6:58] for the normal behavior case, and [42, 8:09] for the exceptional behavior case.




Figure 3.2: The heap before (H_b) consists of an arbitrary chain of nodes. In the heap after (H_a) the dashed lines show which objects are identical to the heap before. The old last node at $\sigma[\ell(\sigma) - 1]$ has a different value for its *next* field in the heap after: this must be the result of a heap update. The new last node is not created in the heap before: indeed, it is the result of creating a new node.

3.7 The remove method

The **remove** method takes as an argument an object; it searches the linked list for the first node which contains the argument as an item. If found, it unlinks the node from the linked list. Using this intuition, we specify the **remove** method contract. Like with the **add** method, there is a deeper method that is called, **unlink**, which we have to specify and verify first.

An immediate difficulty in specifying the **remove** method contract is that its intended behavior depends on the behavior of the **Object.equals** method. Namely, the informal Java documentation states that the first element occurrence in the list that is ‘equal to’ the argument must be removed. Equality can be user-defined by overriding the **equals** method! We solve this difficulty by assuming a method contract for the equality method, see Listing 3.16.

```

/*@ public normal_behavior
   @ requires true;
   @ ensures \result == self.equals(param0);
   @*/
public /*@ helper strictly_pure @*/
boolean equals(/*@ nullable @*/ Object param0);

```

Listing 3.16: The **equals** stub method with its method contract expressed in JML.

We declare the equality method to be strictly pure, which implies that it must be a side-effect-free and terminating method (see [29, Section 7.3.5]). Each strictly pure method is also directly accessible as an observer symbol (a function symbol) that can be used in specifications (see [29, Section 8.1.2]). However, no obvious relation between the possibly overridden equality method and its observer symbol is present. The intention of the contract given in Listing 3.16 is to relate the outcome of the method call of **equals** to the observer symbol *equals*, and this furthermore requires that the implementation is deterministic.

The ramifications of adding this assumed contract are not clear. We note that there are Java classes for which equality is not terminating under certain circumstances. Even **LinkedList** itself does not have terminating equality, where two linked lists that contain each other may lead to a **StackOverflowError** when testing their

equality. This example is described in the Javadoc [44] of the linked list: “Some collection operations which perform recursive traversal of the collection may fail with an exception for self-referential instances where the collection directly or indirectly contains itself.” Another approach is to specify the outcome of equality as referential equality only, e.g. see [10, Section 4.4].

Now we can specify the behavior of `remove`. It can be seen as consisting of two cases: either its result is `true` and it has removed the first item equal to its argument from the list, or its result is `false` and the argument was not found and thus not removed. In the first case, the number of elements decreases by one. In the second case, the number of elements in the linked list remains unchanged. See Listing 3.17.

```

/*@
  @ public normal_behavior
  @ requires true;
  @ ensures \result == false ==>
  @ (\forallall \bigint i; 0 <= i < \old(nodeList.length);
  @ (o==null ==> \old(((Node)nodeList[i]).item) != null) \&\&
  @ (o!=null ==> !\old(o.equals(((Node)nodeList[i]).item)))) \&\&
  @ nodeList == \old(nodeList);
  @ ensures \result == true ==>
  @ (\exists \bigint j; 0 <= j < \old(nodeList.length);
  @ (\forallall \bigint i; 0 <= i < j;
  @ (o==null ==> \old(((Node)nodeList[i]).item) != null) \&\&
  @ (o!=null ==> !\old(o.equals(((Node)nodeList[i]).item)))) \&\&
  @ nodeList == \seq_concat(\old(nodeList)[0..j],
  @ \old(nodeList)[j+1..\old(nodeList.length)]) \&\&
  @ (o==null ==> \old(((Node)nodeList[j]).item) == null) \&\&
  @ (o!=null ==> \old(o.equals(((Node)nodeList[j]).item)))));
/*@

```

Listing 3.17: The `remove` method contract expressed in JML.

It is important to note that we make use of JML’s `\old` operator to refer to the equality observer symbol in the old heap. Using equality in the new heap is a different observation, and it should not be possible to verify the `remove` method in this case. To see why, consider two linked list instances x and y : we add x to itself, and to y we add x and then y . Now we perform the `remove` operation on y with y as argument. Clearly, x and y are not equal, because they have a different length. But the second item is y itself, and y equals y , so it is removed: see Figure 3.3. In the resulting heap, both x and y contain x as the only item: thus, x and y are equal. If we would observe equality in the new heap, then the implementation is incorrect: the item to remove should not be the second but the first!

Before we can verify the `remove` method, we must specify and verify its deeper method: `unlink`. Within the method of `unlink`, we have to update the chain ghost field as well, to remove a node from the sequence, so we add a set annotation to the method body. Additionally, we make use of the lemma `lemma_acyclic` by calling it as the first statement of the method. See Listing 3.18. This method call is also




Figure 3.3: The situation (H_b) before `remove` is invoked on y with argument y , and after (H_a). The result of this operation is that y 's second node n_3 is unlinked, hence the first node n_2 becomes the last node and its next pointer is cleared: now x and y are equal because they have the same length, and they have the same item, namely x .

not present in the original definition for `unlink`, but we already argued that it does not affect behavior.

```

/*@ normal_behavior
  @ requires nodeList != \seq_empty &&
  @ 0 <= nodeIndex < nodeList.length &&
  @ (Node)nodeList[nodeIndex] == x;
  @ ensures \result == \old(x.item) &&
  @ nodeList == \seq_concat(\old(nodeList)[0..nodeIndex],
  @ \old(nodeList)[nodeIndex+1.. \old(nodeList).length]) &&
  @ nodeIndex == \old(nodeIndex);
  @*/
/*@ nullable @*/ Object unlink(Node x) {
  lemma_acyclic(); // new
  //@ set nodeList = \seq_concat(\dl_seqSub(nodeList,0,nodeIndex),
  \dl_seqSub(nodeList,nodeIndex+1,\dl_seqLen(nodeList)));
  // rest of method body
  ...
}
    
```

Listing 3.18: The first part of method `unlink` and its method contract. Note that the `@set` annotation must not contain a new line, but kept on a single line: otherwise KeY 2.6.3 cannot load the source file. Here, `/*@ @*/` does not work.

An interesting aspect of the specification of `unlink` is its use of a *ghost parameter*. Although KeY does not directly support ghost parameters, we can work around that by adding the parameter as a ghost field to our class:

```

//@ private ghost \bigint nodeIndex;
    
```

Its value is left undefined for most of the lifetime of the linked list until we are about to invoke `unlink`. In particular, the ghost parameter contains the index of the node argument, thereby requiring that the node object passed in is part of the chain. In the following discussion, let I be the node index ghost parameter and σ the chain of the linked list: then $\sigma[I]$ is assumed to be the node argument of the method `unlink`.

The verification of the `unlink` method is not fully automatic, see the five videos

[45, 46, 47, 48, 49].

Verifying unlink consists of four main cases: these correspond to the possible branches of the two if-statements (see Listing 3.6). The challenge again is to reestablish the class invariant in the heap after the method completes. The main insight is that, by the acyclicity property, all the nodes are separate: this allows us to distinguish the heap updates to apply only to the node that is actually affected while leaving the other nodes equal to the situation in the heap before. The three important cases are depicted in Figure 3.4, Figure 3.5, Figure 3.6 (compare with Figure 3.2).

1. Suppose the test of both if-statements evaluate to true: for node x , it holds that $\text{next}(x) = \text{null}$ and $\text{prev}(x) = \text{null}$. Then we know the list consists of exactly one node, as the node we are unlinking is the first and the last node. So I cannot be larger than 0. In the case the node index is zero, the class invariant is proven automatically [45, 7:25].




Figure 3.4: The heap before (H_b) consists of a chain σ with $\ell(\sigma) \geq 2$. The dashed lines show which objects are identical in the heaps. Interpreting $\sigma[I]$ in the new heap gives $\sigma[I + 1]$ in the old heap, as I does not change. The **first** field of the linked list has changed (not shown), and the second node in the old heap now has **null** as **prev**. Moreover, the **next** and **prev** fields of the unlinked node have been set to **null** (not shown).

2. Suppose the test of the first if-statement evaluates to true, but the test of the second if-statement evaluates to false: for node x , it holds that $\text{next}(x) \neq \text{null}$ and $\text{prev}(x) = \text{null}$. We thus know that the list consists of at least two nodes, and it is the first node we are unlinking. Thus, I cannot be larger than 0. If the node index is zero, the class invariant is not proven automatically [46, 2:40], but we have two open goals corresponding to the chain properties (b) and (c), cf. proof of **linkLast**. Our situation is different now, see Figure 3.4. Here our insight applies: because of acyclicity, we know all nodes are different. Thus, an update of $\sigma[I]$'s fields does not affect the other nodes. When proving (c) this is sufficient as no next field of nodes in the new chain is changed compared to the old heap [46, 10:14-15:28]. When proving (b), we furthermore make a case distinction between the new first node and the other nodes: the former follows from the heap update, the latter from the old invariant [46, 3:27-10:13].
3. Suppose the test of the first if-statement evaluates to false, and the test of the second to true: for node x , it holds that $\text{next}(x) = \text{null}$ and $\text{prev}(x) \neq \text{null}$. This means that the list consists of at least two nodes, and it is the last node we are unlinking. The proof is similar to the previous case: see Figure 3.5 and [47].




Figure 3.5: The heap before (H_b) consists of a chain σ with $\ell(\sigma) \geq 2$. The dashed lines show which objects are identical in the heaps. Interpreting $\sigma[I]$ in the new heap is invalid, as $I = \ell(\sigma)$ in the new heap. The **last** field of the linked list has changed (not shown), and the before last node in the old heap now has **null** as **next** field.

4. Suppose both tests of if-statements evaluate to false: for node x , it holds that $\text{next}(x) \neq \text{null}$ and $\text{prev}(x) \neq \text{null}$. This implies that the list consists of at least three nodes: where x is some ‘interior’ node. This part of the proof is the largest, as it involves many case distinctions. Up to the point where the class invariant is established in the heap after goes as before, except for (b) and (c). Keep in mind the situation as depicted in Figure 3.6, and see [48, 49].




Figure 3.6: The situation (H_b) consists of a chain σ with $\ell(\sigma) \geq 3$. The index I remains unchanged in the heaps, thus $\sigma[I]$ in the heap after is equal to $\sigma[I + 1]$ in the heap before. The following fields are updated in the new heap: the **next** field of the node at $\sigma[I - 1]$ in the old heap becomes the node at $\sigma[I + 1]$ in the old heap, and the **prev** field of the node at $\sigma[I + 1]$ in the old heap becomes the node at $\sigma[I - 1]$ in the old heap. In the new heap, these two nodes are present in succession in the chain, thus satisfying the chain properties (b) and (c).

We distinguish the two cases:

- (b) Establishing the **prev** field property of the chain involves the following observation: there are three cases. First case, for all nodes at an index $0 \leq i < I$ in the old heap, we know they are identical to the nodes at the same index in the new heap. We know the heap is updated to assign the **prev** field of $\sigma[I + 1]$ in the old heap, and by acyclicity we know this node is separate from the nodes all before $\sigma[I]$ in the old heap. Second case, $\sigma[I]$ interpreted in the new heap is identical to $\sigma[I + 1]$ in the old heap, and precisely for this node the **prev** field was updated to become $\sigma[I - 1]$ in the old heap (which is $\sigma[I]$ in the new heap). Third and last case, for all nodes at an index $I + 1 < i < \ell(\sigma)$ in the old heap, we know they are identical to the nodes at $\sigma[i - 1]$ in the new heap. Again, by acyclicity, we know that the node $\sigma[I + 1]$ in the old heap is separate from the nodes

with a higher index, so we know their `prev` field cannot be affected by the update.

- (c) Establishing the `next` fields property of the chain is very similar but with the index offset by one. Observe that the `next` field of $\sigma[I - 1]$ in the old heap is updated to $\sigma[I + 1]$ in the old heap. Thus the three cases are: first, for the nodes with index $0 \leq i < I - 1$ in the old heap, second, for the node $\sigma[I - 1]$ in the old heap, and third, for the nodes with index $I < i < \ell(\sigma)$ in the old heap.

Now that we have established that `unlink` removes a node from the chain while maintaining the class invariant, we can return to the verification of the `remove` method. The `remove` method iterates over the linked list until it has obtained a node to remove. However, the termination of this iteration is not obvious. Moreover, before invoking the `unlink` method, we need to specify the value of its ghost parameter: the index corresponding to the node. So, before we can verify the `remove` method, we add three kinds of annotations to its source: a ghost variable for maintaining the current index, a loop invariant that establishes termination and maintains the class invariant, and a set annotation before invoking the `unlink` method (see Listing 3.19).

```

public boolean remove(/*@ nullable @*/ Object o) {
  /*@ ghost \bigint index = -1;
  if (o == null) {
    /*@ maintaining 0 <= (index + 1) \&\&
    @ (index + 1) <= nodeList.length;
    @ maintaining (\forallall \bigint i; 0 <= i < (index + 1);
    @ ((Node)nodeList[i]).item != null);
    @ maintaining (index + 1) < nodeList.length ==>
    @ x == nodeList[index + 1];
    @ maintaining
    @ (index + 1) == nodeList.length <==> x == null;
    @ decreasing nodeList.length - (index + 1);
    @ assignable \strictly_nothing; */
    for (Node x = first; x != null; x = x.next) {
      /*@ set index = index + 1;
      if (x.item == null) {
        /*@ set nodeIndex = index;
        unlink(x);
        return true;
      }
    }
  } else {
    /*@ maintaining 0 <= (index + 1) \&\&
    @ (index + 1) <= nodeList.length;
    @ maintaining (\forallall \bigint i; 0 <= i < (index + 1);
    @ !o.equals(((Node)nodeList[i]).item));
    @ maintaining (index + 1) < nodeList.length ==>
    @ x == nodeList[index + 1];

```

```

    @ maintaining
    @ (index + 1) == nodeList.length <==> x == null;
    @ decreasing nodeList.length - (index + 1);
    @ assignable \strictly_nothing; */
    for (Node x = first; x != null; x = x.next) {
        //@ set index = index + 1;
        if (o.equals(x.item)) {
            //@ set nodeIndex = index;
            unlink(x);
            return true;
        } } }
    return false;
}

```

Listing 3.19: The JML annotations of method `remove`. We use a slightly unnatural initial value for the index ghost variable since the KeY 2.6.3 parser does not recognize the `@set` annotation if it appears after the if-statement.

The verification of the above method is not fully automatic, see [50, 51]. The proof consists of two parts, corresponding to the branches of the if-statement. In the proof, one shows (among other properties) that the loop invariant holds initially, after each iteration, and at the end of the loop. It is important to note that $(index + 1)$ is equal to the length of the chain precisely when the end of the loop has been reached. This holds since we use the `next` field to traverse the chain, and only the last node has a `null` successor. Moreover, the distance to the last node decreases each iteration, and this distance is bounded from below by zero: thus the loop must terminate. Moreover, the loop is *strictly pure*, as it never modifies the heap in any of its completed iterations. The exceptional case is the last iteration in which the `remove` method returns early. Due to the early return, the loop invariant no longer needs to be shown (and so also not its heap purity). For reasons of limited space, further examination of its proof is left as a challenge to the reader.

3.8 Summary

In the chapter, we have shown the verification of two essential methods of Java’s `LinkedList` class: `add` and `remove`. The original implementation contains an overflow bug (see Section 3.3.2), and we have looked at a revised version that imposes a maximum length of the list. Furthermore, we have set out to verify that the overflow bug indeed no longer occurs. Towards this end, we have formally specified a class invariant and method contracts, with two goals: establishing the absence of the overflow bug and capturing the ‘essential’ behavior of the methods with respect to the structural properties of the linked list. All methods have been formally verified [39] using the KeY theorem prover, and video material shows on [40].

This chapter aims to provide a comprehensive guide on using KeY and JML for the specification and verification of selected portions of Java libraries. We did a lot of work on recording the video materials, aiming to cater to both the beginning user,

the expert user, and the developer of KeY as a ‘benchmark’ for specification and (automatic) verification techniques.

