

TRABAJO FIN DE MÁSTER EN PROGRAMACIÓN Y TECNOLOGÍA SOFTWARE

MÁSTER EN INVESTIGACIÓN EN INFORMÁTICA
FACULTAD DE INFORMÁTICA
UNIVERSIDAD COMPLUTENSE DE MADRID



Decision Procedures for the Temporal Verification of Concurrent Data Structures

Alejandro Sánchez

Director: Miguel Palomino Tarjuelo
Colaborador externo: César Sánchez

2010/2011

Calificación obtenida: 10

Autorización de difusión

El abajo firmante, matriculado en el Máster en Investigación en Informática de la Facultad de Informática, autoriza a la Universidad Complutense de Madrid (UCM) a difundir y utilizar con fines académicos, no comerciales y mencionando expresamente a su autor, el presente Trabajo Fin de Máster: *Decision Procedures for the Temporal Verification of Concurrent Data Structures*, realizado durante el curso académico 2010-2011 bajo la dirección de Miguel Palomino Tarjuelo y con la colaboración externa de dirección de César Sánchez en el Departamento de Sistemas Informáticos y Computación, y a la Biblioteca de la UCM a depositarlo en el Archivo Institucional E-Prints Complutense con el objeto de incrementar la difusión, uso e impacto del trabajo en Internet y garantizar su preservación y acceso a largo plazo.

En Madrid, a los 7 días del mes de julio de 2011,

Alejandro Sánchez
Y0403168-S

Miguel Palomino Tarjuelo
Director

Resumen

Los tipos de datos concurrentes básicamente son implementaciones concurrentes de abstracciones de datos clásicas, diseñados específicamente para aprovechar el gran paralelismo disponible en arquitecturas multiprocesador y multinúcleo. La corrección de los tipos de datos concurrentes resulta esencial para demostrar la corrección de sistemas concurrentes en su conjunto. En este trabajo estudiamos el problema de asistir en la automatización de la verificación de propiedades temporales en tipos de datos concurrentes.

La principal dificultad en el razonamiento sobre estos tipos de datos proviene de la interacción entre la alta concurrencia que éstos poseen y la necesidad de manipular al mismo tiempo la memoria dinámica. La mayoría de los enfoques utilizados hasta el momento para la verificación de tipos de datos concurrentes intentan enriquecer *separation logic* para poder hacer frente a la concurrencia, valiéndose del éxito alcanzado por *separation logic* en el razonamiento de algoritmos secuenciales y que manipulan el *heap*.

Este trabajo contiene dos contribuciones. En primer lugar, presentamos un enfoque complementario para la verificación de estructuras de datos concurrentes: partimos de verificación temporal deductiva, una poderosa técnica para razonar sobre sistemas concurrentes, y la enriquecemos para poder lidiar con memoria dinámica. El proceso de verificación utiliza diagramas de verificación y anotación explícita de regiones. Al final, cada prueba se descompone en una secuencia de condiciones de verificación. Los literales que forman parte de cada una de estas condiciones de verificación dependen casi exclusivamente de las estructuras de datos que vayan a ser verificadas.

La segunda, y principal, contribución consiste en dos procedimientos de decisión para tipos de datos específicos: listas concurrentes simplemente enlazadas con cerrojos en cadena y listas concurrentes con saltos. Estos procedimientos de decisión son capaces de razonar sobre regiones, punteros, listas del estilo de *Lisp* o listas ordenadas, permitiendo la verificación automática de las condiciones de verificación generadas previamente.

Aquí demostramos cómo, utilizando nuestra técnica, resulta posible demostrar no solo propiedades de *seguridad*, sino también *viveza* sobre una versión de listas concurrentes, además de la preservación de la forma de listas con saltos por parte de una estructura de datos. Así mismo, el enfoque presentado puede ser fácilmente extendido para razonar sobre un amplio espectro de tipos de datos concurrentes, incluyendo tablas *hash* y grafos.

Palabras claves:

Concurrencia, Verificación Formal, Estructuras de Datos, Procedimientos de Decisión, Diagramas de Verificación

Abstract

Concurrent datatypes are concurrent implementation of classical data abstractions, specifically designed to exploit the great deal of parallelism available in multiprocessor and multicore architectures. The correctness of concurrent datatypes is essential for the overall correctness of the system. In this work we study the problem of aiding in the automation of temporal verification of concurrent datatypes.

The main difficulty to reason about these datatypes comes from the combination of their inherently high concurrency and the manipulation of dynamic memory. Most previous approaches to verification of concurrent datatypes try to enrich separation logic to deal with concurrency, leveraging on the success of separation logic in reasoning about sequential heap algorithms.

This work contains two contributions. First, we present a complementary approach to the verification of concurrent data structures: we start from deductive temporal verification, a very powerful technique to reason about concurrent systems, and enrich it to cope with dynamic memory. The verification process uses verification diagrams and explicit region annotations. In the end, each proof is decomposed into a sequence of verification conditions. The literals involved in these verification conditions depend mainly on the data structured being verified. The second, and main, contribution consists in two decision procedures for specific data-types: concurrent lock-coupling singly-linked lists and concurrent skiplists. These decision procedures are capable of reasoning about regions, pointers, lisp-like lists and ordered lists allowing automatic verification of generated verification conditions.

We show how using our technique we are able to prove not only safety but also liveness properties of a version of concurrent lists and express the preservation of skiplist shape by a data structure. Moreover, the approach we present can be easily extended for using it in the verification of a wide range of similar concurrent datatypes including hash maps and graphs.

Keywords:

Concurrency, Formal Verification, Data Structures, Decision Procedures, Verification Diagrams

Acknowledgments

First of all, I would like to thank César Sanchez for being my advisor, my scientific guide, my inspiration source and the person who drove me into the research world. I am sure that nothing of this work could be possible without his help.

I would also like to thank Miguel Palomino for being the director of this work and for providing me with very useful hints on how to improve this presentation.

Special thanks to my family, who always supported me. Thanks to my parents and my sister who are always with me and, in special, to my grandmother Elena, who taught me my very first letters and has been my early mentor to walk this path of knowledge and science.

Thanks to IMDEA Software for providing me the tools, space and commodities to carry out the current work. In special to all people working there, for their friendship and support.

Finally, I would like to express my special gratitude towards Carolina, Javier and Julian, whose unexpected combined travel abroad gave me the necessary time to finish this work. 😊

Thank you all,
Ale.

Contents

Resumen	v
Abstract	vii
Acknowledgments	ix
1 Introduction	1
1.1 Designing Concurrent Data Structures	3
1.1.1 Performance Concerns	4
1.1.2 Synchronization Techniques	4
1.1.3 Verification Techniques	6
1.1.4 Synchronization Elements	8
1.1.5 Some Concurrent Data Structures	9
1.2 Decision Procedures for Pointer Based Data Structures	14
2 Preliminaries	15
2.1 Regional Logic	16
2.2 Verification Diagrams	16
3 Concurrent Lists and Skiplists	23
3.1 Concurrent Lock-Coupling Lists	23
3.2 Concurrent Skiplists	26

4	TLL3: A Decision Procedure for Concurrent Lock-Coupling Lists	37
4.1	Decidability of TLL3	42
4.2	A Combination-based Decision Procedure for TLL3	47
4.3	Verifying Some Properties Over Concurrent Lists	49
4.3.1	Termination of Concurrent Lock-Coupling Lists	50
4.3.2	No Thread Overtakes	53
5	TSL_K: A Decision Procedure for Concurrent Skiplists	55
5.1	Decidability of TSL _K	64
5.2	A Combination-based Decision Procedure for TSL _K	73
5.3	Extending TSL _K to Reason about (L, U, H)	78
5.4	Verifying Some Properties Over Concurrent Skiplists	82
5.4.1	Skiplist Preservation	82
5.4.2	Termination of an Arbitrary Thread	85
6	Conclusion	87
	Bibliography	89

Introduction

Concurrent data structures [HS08] are an efficient approach to exploit the parallelism of multiprocessor architectures. In contrast with sequential implementations, concurrent datatypes allow the simultaneous access of many threads to the memory representing the data value of the concurrent datatype. Concurrent data structures are hard to design, challenging to implement correctly and even more difficult to formally prove correct.

The main difficulty in reasoning about concurrent datatypes comes from the interaction of concurrency and heap manipulation. The most popular technique to reason about the structures in the heap is separation logic [Rey02]. Leveraging on this success, some researchers [VHHS06, HAN08] have extended this logic to deal with concurrent programs. However, in separation logic disjoint regions are implicitly declared (hidden in the separation conjunction), which makes the reasoning about unstructured concurrency more cumbersome.

Explicit regions allow the use of a classical first-order assertion language to reason about heaps, including mutation and disjointness of memory regions. Regions correspond to finite sets of object references. Unlike separation logic, the theory of sets [WPK09] can be easily combined with other classical theories to build more powerful decision procedures. Classical theories are also amenable of integration into SMT solvers [BSST08]. Moreover, being a classical logic one can use classical Assume-Guarantee reasoning, for example McMillan proof rules [McM99], for reasoning compositionally about liveness properties. In practice, using explicit regions requires the annotation and manipulation of ghost variables of type *region*, but adding these annotations is usually straightforward.

Most of the work in formal verification of pointer programs follows program logics in the Hoare tradition, either using separation logic or with specialized logics to deal with the heap and pointer structures [LQ08, YRS⁺06, BDES09]. However, extending these logics to deal with concurrent programs is hard, and though some success has been accomplished it is still an open area of research, particularly for liveness.

We propose a complementary approach. We start from temporal deductive verification in the style of Manna-Pnueli [MP95], in particular using general verification diagrams [BMS95, Sip99] to deal with concurrency. This style of reasoning allows a clean separation in a proof between the temporal part (why the interleavings of actions that a set of threads can perform satisfy a certain property) with the underlying data being manipulated. Then, inspired by regional logic [BNR08], we enrich the state predicate language to reason about the different regions in the heap that a program manipulates. Finally, we build decision procedures capable of checking all generated verification conditions generated during our proofs, to aid in the automation of the verification process.

Most previous approaches to verifying concurrent datatypes are restricted to safety properties. In comparison, the method we propose can be used to prove *all liveness* properties, relying on the completeness of verification diagrams.

Verification diagrams can be understood as an intuitive way to abstract the specific aspect of a program which illustrates why the program satisfies a given temporal property. Besides, verification diagrams have been proved to be complete, in the sense that given a program and a temporal formula satisfied by the program then a verification diagram does in fact exist, and sound in the sense that a diagram connecting a program with a formula is in fact a proof that such temporal property is held by the system. We propose the following verification process to show that a datatype satisfies a property expressed in linear temporal logic. First, we build the *most general client* of the datatype, parametrized by the number of threads. Then, we annotate the client and datatype with ghost fields and ghost code to support the reasoning, if necessary. Second, we build a verification diagram that serves as a witness of the proof that all possible parallel executions of the program satisfy the given temporal formula.

The proof is checked in two phases. First, we check that all executions abstracted by the diagram satisfy the property, which can be solved through a fully-automatic finite state model checking method. Second, we must check that the diagram does in fact abstract the program. A verification diagram decomposes a formal proof into a finite collection of verification conditions (VC), each of which corresponds to the effect that a small step in the program has in the data. Then, the verification reduces to verifying this collection of verification conditions. Each concurrent datatype maintains in memory a collection of nodes and pointers with a particular layout. Based on this fact, we propose to use an assertion language whose terms include predicates in specific theories for each layout. To automatize the process of checking the proof represented by a verification diagram it is necessary to use decision procedures for the kind of data structures manipulated. For instance, in the case of singly linked lists, we use a decision procedure capable of reasoning about ideal lists as well as pointers representing lists in memory.

The construction of a verification diagram is a manual task, but it often follows the programmer’s intuitive explanation of why the property holds. The activity that we want to automate is checking that the diagram indeed proves the property. Following this idea, in this work we study the automatic verification of VCs for the case of lists and skiplists extending the theory of linked lists presented in [RZ06a]. We present two theories for concurrent singly-linked lists and concurrent skiplists. Besides, we describe a decision procedure for each of them, applying a many-sorted variant of Nelson-Oppen [NO79] combination method.

Logics like those in [LQ08, YRS⁺06, BDES09] are very powerful to describe pointer structures, but they require the use of quantifiers to reach their expressive power. Hence, these logics preclude a combination a-la Nelson-Oppen or BAPA [KNR05] with other aspects of the program state.

Instead, our solution starts from a quantifier-free theory of single-linked lists [RZ06a], and extends it in a non trivial way with locks, order and sublists of ordered lists. The logics obtained can express list and skiplist-like properties without using quantifiers, allowing the combination with other theories. Proofs for an unbounded number of threads are achieved by parameterizing verification diagrams, splitting cases for interesting threads and producing a single verification condition to generalize the remaining cases. However, in this work we mainly focus on the decision procedure.

The decision procedures that we present here support the manipulation of explicit regions, as in regional logic [BNR08] equipped with *masked regions*, which enables reasoning about disjoint portions of the same memory cell. In this aspect, we use masked regions to “separate” different levels of the same skiplist node.

As we said, we present two theories: TLL₃ for concurrent singly-linked lists and TSL_K for concurrent skiplists of height at most K. To illustrate the use of these theories, we sketch the proof of termination of every invocation of an implementation of a lock-coupling concurrent list and the preservation of the skiplist shape of a data structure. We also prove the decidability of TLL₃ and TSL_K by showing that they enjoy the finite model property. Moreover, we propose an even more efficient decision procedure for them.

The rest of this work is structured as follows. We conclude Chapter 1 introducing the basic notions of concurrent data structures and decision procedures as well as a description of the progress achieved in these fields. Chapter 2 introduces some basic notions about regional logic and verification diagrams, the building blocks of our verification method. Chapter 3 describes the concurrent data structures for which we have developed decision procedures: concurrent lock-coupling lists and concurrent skiplists. In Chapter 4 we explain the decision procedure for concurrent lock-coupling singly-linked lists. Chapter 5

details the decision procedure for concurrent skiplist. Finally, Chapter 6 gives the conclusions obtained from this work.

1.1 Designing Concurrent Data Structures

Multiprocessor shared-memory systems are capable of running multiple threads, where all threads communicate and synchronize through data structures that reside in shared memory. As one may imagine, concurrent data structures are much more difficult to design than sequential ones due to the way in which involved threads may interleave. In addition, each possible interleave may lead to a different, and possibly unexpected, behavior. In this section we elaborate on some of the aspects described in [MS07] on concurrent systems. We provide an overview of the challenges involved in the design of concurrent data structures as well as a brief description of the most common concurrent datatypes present in the literature.

One of the main sources of difficulty when designing this kind of data structures comes from concurrency. As threads may run concurrently on different processors, they are subject to operating system scheduling policies, page fault, interruptions, etc. Hence, we should reason about the computation as a completely asynchronous process, where steps of different threads can be arbitrarily interleaved.

As an example, imagine we want to implement a *shared counter*. A *shared counter* consists of a shared variable that is concurrently incremented by many threads. A sequential version of such algorithm is depicted in Fig. 1.1(a). This version just fetches the value from counter c and increments it by one, returning the previous value. However, if many threads run this implementation concurrently, then the result may not be the expected one. For instance, consider the scenario at which c is 0 and two threads execute this implementation concurrently. Then, it exists the possibility that both threads read the value 0 from the memory and thus both return 0 as results, which is clearly wrong.

To avoid this problem, a common solution is the use of mutual exclusion locks (usually simply called *mutex* or *lock*). A lock is a construction with the property that, at any given time, nobody owns it or it is owned by a single thread. If a thread T_1 wants to acquire a lock owned by thread T_2 , then it must wait until T_2 releases it.

Now, using locks, a correct implementation of the *shared counter* program can be obtained, as shown in Fig. 1.1(b). This new version uses a lock to allow access to the critical section of a single thread at a time. We refer to a program section as *critical* when it provides access to a shared data structure which must not be concurrently executed by more that one thread simultaneously. In the example, the lock clearly prevents two threads to access the critical section. While this is a correct solution, it lacks of good performance. It is easy to achieve a correct implementation just by surrounding the whole code with a lock. However, the performance of such implementation may not differ too much from a sequential one, losing all power provided by concurrency. In the following, we describe some aspects to bear in mind in order to improve performance while preserving correctness, as well as some techniques for verifying the correctness of such data structures.

<p style="text-align: center;">Procedure SEQ-COUNTER</p> <p>1: $oldValue := c$ 2: $c := oldValue + 1$ 3: return $oldValue$</p> <p style="text-align: center;">End Procedure</p> <p style="text-align: center;">(a) Sequential</p>	<p style="text-align: center;">Procedure CONC-COUNTER</p> <p>1: acquire($lock$) 2: $oldValue := c$ 3: $c := oldValue + 1$ 4: release($lock$) 5: return $oldValue$</p> <p style="text-align: center;">End Procedure</p> <p style="text-align: center;">(b) Concurrent</p>
---	--

Figure 1.1: Possible implementation of a *shared variable*

1.1.1 Performance Concerns

The *speedup* of a program when run on P different processors is the relation between its execution time on a single processor and its execution time on P processors, concurrently. It can be considered as a measure on how efficient a program uses the processor on which it is running. Ideally, we would like to have linear speedup, however this is not always possible. Data structures whose speedup grows with P are called *scalable*. When designing concurrent data structures, scalability must be taken into account. Naive implementations, as `CONC-COUNTER` will surely undermine scalability.

Bearing in mind the example of a shared counter, it is evident that the lock introduces a sequential bottleneck: at every instant, a single thread is allowed to perform useful work. Hence, reducing the size of code sections being sequentially executed is crucial in order to achieve good performance. In the context of locking, we are interested in:

- reducing the number of acquired locks, and
- reducing the lock granularity.

Lock *granularity* measures the number of instructions executed while holding a lock. The fewer instructions, the finer lock granularity. Implementations like the shared counter's one showed above represent an example of a coarse-grain solution.

Another aspect to bear in mind is *memory contention*. Memory contention refers to the overhead in traffic as a result of multiple threads concurrently trying to access the same memory location. For instance, if the lock protecting a critical section is implemented in a single memory location, then a thread trying to access that section must continuously try to access the memory portion where the lock is stored. This problem may be solved if we consider *cache-coherent* multiprocessors. However, using this technique may lead to long waiting times each time a location needs to be modified. As a result, the exclusive ownership of the cache line containing the lock must be repeatedly transferred from one processor to the other.

A final problem with lock-based solutions is that, if a thread holding a lock delays in releasing the lock then all other threads with the intention to acquire the same lock are also delayed. This phenomenon is known as *blocking* and it is quite common on systems with many threads per processor. A possible solution is provided by *non-blocking* algorithms. Non-blocking programs do not use locks and thus the delay of a thread does not imply the delay of other threads. In the following section we describe the main characteristics of lock-based and non-blocking systems.

1.1.2 Synchronization Techniques

`CONC-COUNTER` used a lock to prevent many threads accessing the critical section simultaneously. However, sometimes not using locks at all is the best solution for getting concurrency on a system. Here, we describe the two major techniques for accomplishing mutual exclusion on modern concurrent systems: blocking techniques and non-blocking techniques.

Blocking Techniques

As we said, in many cases memory contention and sequential bottleneck can be reduced by reducing the granularity of locking schemes. In fine-grained locking approaches, multiple locks of small granularity are used to protect the critical sections modifying the data structure. The idea is to maximize the amount of time threads are allowed to run in parallel, as far as they do not require to access the same sections of the data structure. This approach can be also used to reduce the contention for individual memory locations. In some cases, this sound natural, as in the case of hash maps. In hash maps, values are hashed to different buckets, one independent of the other, reducing the number of possible conflicts. Hence, placing individual locks on each bucket sounds like a good approach to reduce granularity.

On the other hand, in cases as the shared counter, it is not quite intuitive how contention and sequential bottleneck can be reduced since, abstractly, all operations modify the same part of the data structure. One approach to deal with contention results in spreading each operation accessing the counter in a separate

time interval from the others. A widely used technique to accomplish this is *backoff* [AC89]. However, even reducing the contention, the lock-based implementation of the shared counter example still lacks parallelism and hence it is not scalable. Therefore, more powerful techniques are required.

A possible solution lies on the use of a method known as *combining trees* [GVW89,GGK⁺83,HLS95,YTL87]. *Combining trees* employ a binary tree with one leaf for each thread. The root of the tree keeps the current value of the counter and intermediate nodes are used to coordinate the access to the root. The idea is that while a thread climbs up to the root of the tree, it combines its effort with the one performed by the other threads. Hence, every time two threads meet on an internal node, their operations are combined into a single operation. Then, one thread waits in the node until a return value is delivered to it while the other one proceeds to the root carrying the operation obtained from the combination.

In the particular case of our shared counter, a winner thread reaching the tree's root modifies the counter in a single atomic operation, performing the action of all combined operations. Then, it goes down once again, delivering the return value to each thread waiting on an internal node. The return values are distributed in such a way that the final effect is as if all operations were executed one after the other at the moment the counter in the root was modified.

Threads waiting on internal nodes repeatedly read on a memory location, waiting for the return value. This sort of waiting is known as *spinning*. An important consequence in the case of a cache-coherent multiprocessor is that the accessed location resides in the local cache of the processor of the waiting thread. Following this approach, no extra traffic is generated. This waiting technique, known as *local spinning*, is very important for scalable performance [MCS91a].

In order to increase performance, *non-uniform memory access* (NUMA) architectures can be used. In NUMA, processors have a mechanism to access their local portions of shared memory faster than they can access the shared memory locations of other processors. A good organization of the data layout on such architectures has a significant impact on performance.

Some techniques, such as combining trees, have the drawback that the total number of threads involved in the system must be known in advance. Moreover, the required space in order to keep the combining tree is proportional to the total number of processors. This way, despite the fact that the technique provides a good through out when the structure is accessed by a significant number of threads, its best case performance under low loads is quite poor, as a result of the fact that until a thread reaches the root, it must traverse $O(\log P)$ tree nodes. On the other hand, the same operation in the lock-based implementation of Fig. 1.1 completes in constant time.

An extra drawback in the combining tree method is that, if coordination between threads going up and down the tree is done incorrectly, it may lead to deadlocks. A *deadlock* is a situation at which two or more threads are executing tasks such that all of them are blocked in a circular fashion, so none of them can progress. Deadlock avoidance is a crucial aspect to bear in mind when designing correct and efficient blocking concurrent data structures and algorithms. In fact, when designing blocking implementations, the number of locks to be taken is a key factor to have in consideration. Enough locks should be used as to ensure correctness while minimizing blocking, allowing other threads to perform concurrent operations in parallel.

Non-blocking Techniques

Non-blocking implementations are designed in order to solve some of the inconveniences present on blocking implementations, avoiding the use of locks. To formalize the idea of non-blocking algorithms, some non-blocking progress conditions are widely described in the literature. Such conditions are known as *wait-freedom*, *lock-freedom* and *obstruction-freedom*. A wait-free [Her91, Lam74] operation must terminate on its own, after a finite number of steps, no matter the behavior of the remaining operations. A lock-free [Her91] operation guarantees that after a finite number of its own steps, some operation terminates: maybe its own operation, or maybe the operation performed by any other thread. Finally, an obstruction-free [HLM03] operation guarantees to terminate after a finite number of its own steps, assuming no interference with other operations.

There exists an evident relation between all these properties. Wait-freedom is stronger than lock-freedom, while lock-freedom is stronger than obstruction-freedom. Clearly, strong progress conditions

are desired over weaker ones. However, weaker conditions are in general simpler and more efficient to design and verify them correct.

Remember that in non-blocking algorithms, no lock can be used. Then, we need to find a new way to implement the concurrent version of the shared counter. [FLP85] (extended to shared memory by [Her91] and [LAA87]) shows that it is not possible to implement a concurrent version using just *load* and *store* instructions. The problem can be solved using a hardware operation that atomically combines both, a *load* and a *store*. In fact, all modern multiprocessors provide one of such synchronization primitives. The most well known are *compare-and-swap* (CAS) [IBM,Int94,S194] and *load-linked/store-conditional* (LL/SC) [IBM03, Kan89, Sit92]. Operation CAS is described as an algorithm in Fig. 1.2. It receives as arguments a memory location L and two values: E and N . The operation is performed atomically: when called, if the value stored at location L matches E , then it replaces it with the new value N and returns *true*. On the other hand, if the value at L does not match E , *false* is returned.

```

Procedure CAS ( $L, E, N$ ) : Bool
1: if  $*L = E$  then
2:    $*L := N$ 
3:   return true
4: else
5:   return false
6: end if
End Procedure

```

Figure 1.2: Description of a CAS operation

In [Her91] it is shown that instructions such as CAS and LL/SC are universal. This means that, for any data structure, it exists a wait-free implementation in a system that supports such instructions.

Using the CAS instruction it is now possible to implement a non-blocking version of the shared counter program. We just require to load the value of counter c and then use CAS to atomically increment the read value by one. In the case the read value does not coincide with the value read when performing the CAS, the operation returns performing no modification to counter c . Hence, we would need to retry until a successful call to CAS is accomplished. Fig. 1.3 shows a non-blocking implementation of the shared counter using CAS. Because the CAS can only fail due to another succeeding *fetch and increment* operation, the implementation is lock-free. However, it is not wait-free as other's thread *fetch and increment* operation can continuously prevent a CAS to succeed.

```

Procedure NONBLOCK-COUNTER
1: repeat
2:    $oldValue := c$ 
3:    $newValue := oldValue + 1$ 
4: until CAS(& $c, oldValue, newValue$ )
End Procedure

```

Figure 1.3: Non-blocking shared counter implementation using CAS

Despite the example shown here being simple, in general, designing non-blocking algorithms is more complex than blocking ones, since a thread can use a lock to prevent other threads from interfering while it performs some action. If locks are not used, then the algorithm must be designed in order to be correct despite the actions performed by the other concurrent threads. Currently, in modern architectures, this requirement leads to the use of complicated and costly techniques.

1.1.3 Verification Techniques

In our case, it is quite easy to see that the lock-based implementation of the shared counter behaves exactly the same as the sequential implementation. However, if we consider a more complicated data structure, as a binary tree for instance, then the verification is significantly more difficult. Because of the inherent

difficulty in the design of concurrent datatypes, it is quite easy to fall into an incorrect implementation. Hence, it becomes imperative to rigorously prove that a particular design correctly implements the desired concurrent datatype.

In general, sequential implementations are easier to verify than concurrent ones. For example, the semantics of a sequential data structure can be specified using a set of states and a transition function that given a state, an operation and a set of arguments, returns a new state in addition to the result of applying the operation with the arguments. In fact, if a valid initial state is given, then it is possible to describe all acceptable sequences of operations over the data structure.

Operations on a sequential data structure are executed one after the other, in order. Then, we simply require that the resulting sequence of operations respects the sequential semantic denoted by the set of states and transitions as described above. On the contrary, in the case of concurrent data structures, operations do not really need to be totally ordered. However, correctness for concurrent implementations usually requires the existence of some total order of operations that respects the sequential semantics.

A common condition is Lamport's *sequential consistency* [Lam79]. This condition requires that the total order preserves the order of the operations run by each thread. However, it has a drawback: a data structure constructed by sequential consistent components may not be sequential consistent at all.

Another widely used concept is *linearizability* [HW90], a variation of the concept of *serializability* [BHG87] used in database transactions. Linearizability requires:

1. that the data structure is sequentially consistent, and
2. that the total ordering, which makes it sequentially consistent, respects the real-time ordering between the operations in the execution.

A way of thinking about linearizability is that it requires the user to identify specific points within the algorithms, called *linearization points*, such that if we order the operations according to their linearization points, the resulting order satisfies the desired sequential semantics of the data structure.

It is quite easy to see that `CONC-COUNTER` is linearizable. We just need to define the point after the incrementation on c is performed as the linearization point. In the case of `NONBLOCK-COUNTER`, the argument is similar, except that we must define the linearization point considering the semantics of `CAS`.

The intuitive appeal and modularity of linearizability makes it a very popular correctness condition. Although, most of the concurrent data structures can be shown to be linearizable, on some situations, better performance and scalability can be achieved by considering a weaker condition: *quiescent consistency*. Quiescent consistency condition [AHS94] eliminates the restriction that the total order of operations needs to respect the real-time order of the executed operations, but it requires that every operation executing after a quiescent state must be ordered after every operation executed before the quiescent state. A state is said quiescent if no operations are in progress.

In general, obtaining a formal proof of the correctness for the implementation of a data structure requires:

- a mathematical method for specifying correctness requirements,
- being able to come up with an accurate model of the data structure implementation, and
- ensuring that the proof of the implementation correctness is complete and accurate.

For instance, most linearizability arguments in the literature treat some of these aspects in an informal way. This makes proofs easier to follow and understand. However, the use of some informal description is prone to introduce some error, miss some cases or make some incorrect inferences. On the other hand, rigorous proofs usually contain a great amount of details regarding trivial properties that makes them difficult to write and tedious to read. Hence, computer assisted methods are desired for the formal verification of concurrent implementations. One approach is based on the use of theorem provers to aid in the verification. Another approach consists in the use of model checking. Model checking tools exhaustively verify all possible executions of an implementation, to ensure that all reachable states meet specified correctness conditions. However, some limitations exists on both approaches. Theorem provers usually require significant human insight, while model checking is generally limited by the number of

states it can consider. Therefore, a verification method involving as few human interaction as possible, while being able to verify systems with a possible infinite number of states is desired.

1.1.4 Synchronization Elements

In this section, we describe some basic mechanisms commonly used to achieve correct concurrent implementations: locks, barriers and transactional synchronization mechanisms. Locks and barriers are traditional low level synchronization mechanisms used to prevent some interleaving to happen. For instance, preventing two different threads to access the same section of code at the same time. On the other hand, transactional synchronization mechanisms are used to hide the complex reasoning required to design concurrent data algorithms, letting programmers to think in a more sequential fashion.

As said, locks are low level mechanisms used to prevent processes access the same region of program concurrently. A key issue to bear in mind when designing locks is what to do when a thread tries to get a lock that is already owned by other thread. A possibility is to let threads keep on trying to get the lock. Locks based on this technique are called *spinlocks*. A simple spinlock may use a primitive such as CAS to atomically change the value of a lock from unowned to owned. However, such spinning may cause heavy contention for the lock. An alternative to avoid contention is the use of exponential backoff. In exponential backoff [AC89] a thread that fails in its attempt of getting a lock waits for some time before a new attempt. With every failed attempt, the waiting time is increased. The idea is that threads will spread themselves out in time, resulting in a reduction of contention and memory traffic.

A disadvantage of exponential backoff is that it may happen that a lock remains unlocked for a long time, since all interested threads have been backed-off too much in time. A possible solution to this problem may consist in the use of a queue of interested threads. Locks based on this approach are known as *queuelocks*. Some implementations of queuelocks based on arrays are introduced in [And89, GT90] and then improved using list-based MCS queue locks [MCS91a] and CLH queuelocks [Cra93, MLH94].

Queuelocks also comes in many flavors. There exists abortable implementations of queuelocks where a thread can give up if it is delayed beyond some time limit [Sco02, SI01] or if they just fall into deadlock. On the other hand, preemptive-safe locks [MS98] ensure that an enqueued preempted thread does not prevent the lock to be taken by another running thread.

In some cases, we would like to have locks letting multiple readers access the concurrent data structure. A reader is a process that only extracts information from the data structure, without modifying it at all. Such locks are known as *reader-writer* locks. There exists many kinds of these locks. For instance, reader-writer queuelock algorithms [MCS91b] use a MCS queuelock, a counter for reads and a special pointer for writes. In [KSUH93] readers remove themselves from the lock's queue, keeping a double-linked list and some special locks on the list's nodes. In this case, when a thread removes itself from the list, it acquires a small lock on its neighbor nodes, and redirects the pointers removing itself from the list.

The reader-writer approach can be also generalized to *group mutual exclusion* or *room synchronization*. Under this approach, operations are divided into groups. Operations within the same group can be performed simultaneously with each other, while operations belonging to different groups cannot be executed concurrently. An application of such technique, for instance, could classify the *push* and *pop* operations over stacks on different groups [BCG01]. Group mutual exclusion is introduced in [Jou02] and implementations for *fetch and increment* counters (as the example shown at the beginning of this chapter) are described in [BCG01, KM99].

Another mechanism are barriers. A barrier stops all threads at a given point, allowing them to proceed only after all threads have achieved such point. Barriers are used when the access to the data structure is layered, preventing layer overlapping of different phases.

A barrier can simply be implemented using a counter to keep track of the number of threads that have achieved the barrier position. The counter is initialized with the total number of threads to wait for. Then, every time a thread reaches the barrier, it decrements the counter. Once the counter has reached zero, it let all threads to proceed. This approach still displays the problem of contention, as many threads may be spinning, waiting for the barrier to let them go through. Therefore, special implementations of barriers exist to attack this problem, making threads spin on different locations [III86, HFM88, SMC92]. An alternative approach consists in implementing barriers using diffusing computation trees [DS80]. In this model each thread owns a node in a binary tree. The idea is as follows. Each thread waits in its node

until all its children have arrived. At this moment, the thread communicates its parent that all threads on its branch have arrived. Once all threads have arrived, the root node releases all threads in the tree by disseminating the release information down the tree.

The main purpose for using locks in concurrent programming is to let threads to modify multiple memory locations atomically in such way that no partial result of the computation can be observed by other threads. In this aspect, transactional memory is a mechanism that lets programmers model sections of the code accessing multiple memory locations as a single atomic operation. The use of transactional mechanisms is inspired on the idea of transactions in databases, despite the fact that the problem in memory management is slightly different to the one existing on databases.

An example of transactional memory mechanism for concurrent data structures is *optimistic concurrency control* [KR81]. This approach uses a global lock which is held for a short time at the end of the transaction. However, as one may imagine, such lock is a cause of sequential bottleneck. Ideally, transaction synchronization should be accomplished without the use of locks. Moreover, transactions accessing disjoint sections of the memory should not synchronize with each other at all. A hardware-based transactional memory mechanism was first proposed in [HM93]. An extension to this idea is *lock elision* [RG01, RG02] where the hardware can automatically translate accesses to critical section into transactions that can be executed in parallel.

Despite the effort, up to this moment no hardware support for transactional memory has been developed. Nevertheless, many software based transactional memory approaches have been proposed [ST97b, HFP02, HLMI03, HF03].

1.1.5 Some Concurrent Data Structures

In this section we introduce the basic idea behind some of the most widely spread concurrent data structures.

Stacks and Queues

Stacks and queues are one of the simplest data structures. A concurrent stack is a data structure linearizable with a sequential stack, providing *push* and *pop* operations, respecting the usual LIFO semantics. As one may imagine, many implementations of such data structure exist in the literature. In [MS98] several lock-based linearizable concurrent stacks are presented. These implementations are based on sequential linked lists with a top pointer and a single global lock. As we saw before, structures with a single global lock scale poorly. Besides, if one succeeds in reducing contention on the lock, a sequential bottleneck still exists at the top of the stack.

A first non-blocking implementation of concurrent stacks was presented by Treiber in [Tre86]. In such implementation, a stack is a singly-linked list with a top pointer that is atomically modified through a CAS operation. In [MS98] this algorithm is compared to an optimized non-blocking algorithm based on [Her93] and several lock-based implementations (such as an MCS lock [MCS91a]) in low load scenarios. The results show that the original implementation reaches the best overall performance. However, as the pointer to the top element represents a sequential bottleneck, it offers little scalability as concurrency increases [HSY04].

In fact, in [HSY04] it is shown that any stack implementation can be made more scalable using the elimination technique [ST97a]. This technique allows pairs of operations with reverse semantics (such as *push* and *pop* over a stack) to complete without any central coordination. If this approach is applied using a collision array and adaptive backoff on the shared stack, then it can reach a high level of parallelism with little contention. Following this method, it is also possible to obtain a scalable lock-free implementation of a linearizable concurrent stack.

Using non-blocking algorithms has also its drawbacks. A common problem in many CAS-based implementation is the *ABA* problem [PLJ91]. To illustrate the problem, in Fig. 1.4 we provide a naive implementation of a non-blocking stack using CAS. We assume the stack is implemented using a singly-linked list. Each node in the list stores an element and keeps a *next* pointer to the following element in the stack. Pointer *topPtr* points to the top of the stack. Fig. 1.4 describes the implementation of operations *push* and *pop*. *push* receives as argument the address of the node to be inserted into the stack. The figure

```

Procedure PUSH (newPtr)
1: repeat
2:   nextPtr := topPtr
3:   newPtr.next := nextPtr
4: until CAS(topPtr, nextPtr, newPtr)
End Procedure

Procedure POP ()
1: repeat
2:   retPtr := topPtr
3:   if retPtr = null then
4:     return null
5:   end if
6:   nextPtr := retPtr.next
7: until CAS(topPtr, retPtr, nextPtr)
8: return retPtr
End Procedure

```




Figure 1.4: Representation of the ABA problem

also describes a scenario at which the ABA problem becomes evident. Fig. 1.4 a) shows the initial state of a stack. We assume two threads running concurrently: T_1 and T_2 . T_1 wants to perform a *pop*, while T_2 performs two *pop* followed by a *push* of A . For simplification, we use A to denote both the element and the address of the node storing such element.

Thread T_1 begins executing the *pop* action. However, before it can execute the CAS, T_2 begins to run. Fig. 1.4 b) shows the state of the stack just before T_1 can execute the CAS operation. Fig. 1.4 c), d) and e) depict the situation of the stack just after T_2 has executed *pop*, *pop* and *push*(A) respectively. Finally, Fig. 1.4 e) shows the state of the stack once T_1 finishes the original *pop* operation. Notice that according to this implementation, when both thread terminate, the top of the stack contains node B , which was eliminated by T_2 . This means that we could be accessing an invalid address through the *topPtr* pointer. However, this kind of problems can be solved. A possible solution to this problem consists in equipping the *topPtr* pointer with a version number that is incremented every time the *topPtr* pointer is modified.

As happened with stacks, a concurrent queue is a data structure that is linearizable to a sequential queue and respects the FIFO semantics. [MS98] presents an improved implementation of the single lock implementation of a queue by providing both, the head and tail pointer, with an independent lock. This modification allows the execution of an *enqueue* and a *dequeue* operation in parallel, provided both pointers are different one from each other.

It is possible to implement an array-based lock-free queue for just two processes, one performing only *enqueue* and the other one carrying out just *dequeues* operations, using only load and store operations [Lam83]. A linked-list version of this algorithm is presented in [HS02] and a lock-free array-based implementation assuming unbounded array size is described in [HW90]. Nevertheless, in general, it is quite difficult to implement a nonblocking queue with multiple enqueueers and dequeuers.

In [MS98] a lock-free implementation of a queue based on CAS operations capable of parallel access through both ends is presented. For such purpose it keeps a head and tail pointer, in addition to a dummy pointer. The *enqueue* operation, first uses CAS to add the new node to the end of the list and then uses CAS again to modify the tail pointer to point to the new node. If another *enqueue* operation tries to execute between the first and the second CAS, it will notice that the tail pointer has not been modified yet. A helping technique [Her93] is used to ensure the tail pointer is always at most one node of distance behind the end of the list. Although this implementation is quite simple and efficient, it has a drawback: operations can access nodes that have already been removed from the list. This implies that nodes cannot be freed. Instead, removed nodes are kept in a free-list for reusing them when a new enqueue operation is executed. Nevertheless, in [HLM02, Mic02b] non-blocking memory management techniques are presented in order to overcome this disadvantage.

A generalization of stacks and queues are deques. A deque [Knu68] is a double-ended queue that allows *push* and *pop* at both ends. Lock-based implementations of deques can be easily obtained following the two-lock approach for queues. However it is not easy to come up with a lock-free implementation of deques as shown in [MMS02], even making use of the two-word synchronization primitive *double-compare-and-swap* (DCAS) [Mot86]. One of the few non-blocking implementation of a deque supporting noninterfering operations on both ends is an obstruction-free CAS-based implementation given at [HLM03].

Much of the difficulty in implementing stacks and queues comes from the restrictions on when an inserted element can be removed, depending on the internal order of elements in the data structure. A concurrent pool [Man84] is a data structure which provides *insert* and *delete* operations and where the *delete* operation removes any occurrence of the given element from the structure. An efficient pool can be constructed using quiescently consistent counter implementations [AHS94, SZ96]. In such implementation, elements are kept in an array and a *fetch and increment* operation is used in order to find out the position where an element must be inserted or removed from. Alternatively, a stack-like pool can be also implemented.

Linked Lists

Consider now the implementation of a concurrent data structure supporting *insert*, *remove* and *search* operations. If order does not matter and the operation deals only with a key value, then we have a *multiset* or *set*, depending on whether repeated elements are allowed or not. Moreover, if a data value is associated with every key, then we have a *dictionary* [CLRS01].

In this section we focus on linked-lists used to implement sets. The following sections discuss other concurrent data structures, such as hash tables, search trees and skiplists, that can be used to implement sets.

Considering a single-linked list, the most common approach in order to manage concurrency consists in *lock-coupling* [BS77, Lea99]. In this approach, each node is associated with a different lock. This way, a thread traversing the list releases a node's lock only after it has acquired the lock of the next node in the list, preventing overtaking. Clearly, this approach reduces granularity but it also limits concurrency since insertions and deletions at different locations may delay each other.

A way to evict this problem is to employ ordered lock-free linked lists. The difficulty now resides in ensuring that insertion and deletion keep the validity of the list. The first CAS-based implementation of a list is given in [Val95]. Here, a special node is used in front of every regular node in order to prevent any undesired behavior while manipulating the list. In fact, this lock-free implementation is correct when combined with a memory management [MS95], but this solution is not practical at all. Other approaches use a special bit to mark nodes as deleted (as in [Har01]), but this scheme is only applicable in garbage collectors environments. The problem can be overcome using memory reclamation methods [HLM02, Mic02b] as described in [Mic02a].

Hash Tables

A hash table [CLRS01] can be seen as a resizable array of buckets, where each of them can hold an expected number of elements. Therefore, operations for insertion, deletion and search take constant time. On extensible hash tables, an extra operation for resizing the table is also required. Usually, this latter operation is the most costly of all.

Clearly, it is possible to implement a non-extensible hash table by placing a reader-writer lock on each bucket [Mic02a]. Nevertheless, to keep the good performance, a hash table needs to be extensible [FNPS79].

Already in the eighties, some extensible concurrent hash tables for distributed databases based on two-level locking schemes were introduced [Eil87, HY86, Kum90]. Besides, there exist some highly efficient and extensible hashing algorithms for non-multiprogrammed environments, based on a sequential version of an algorithm for linear hashing [Lit80]. These algorithms are based on a small set of high-level locks, instead of having one lock for each bucket. In addition, it allows concurrent searches while performing the resizing of the hash table. However, concurrent inserts and deletes are not allowed.

Lock-based implementations of concurrent hash tables suffer from the classical drawbacks of lock-based techniques, which becomes even more notorious due to the elaborated resizing methods they require.

Therefore, lock-free implementations of hash tables are desired. Some ideas for reaching a lock-free implementation consist in using an array of buckets, such that each bucket contains a lock-free list. However, this approach does not take into consideration how to deal with extensible arrays, as it is not obvious how to redistribute the elements in a lock-free manner through the buckets of the array. Moving an item from one bucket to another would require two CAS operations to be executed atomically.

Using DCAS it is possible to obtain a lock-free extensible hash table [Gre02]. Nevertheless, this solution has the inconvenient that not all architectures provide the primitive DCAS. In [SS03] an extensible lock-free hash table for concurrent architectures is presented. The idea is to keep a single lock-free list to store the elements instead of a list for each bucket. To preserve the fast access to the elements, the hash table keeps a resizable array of pointers to the list. Operations can then use these pointers to get fast access to internal sections of the list. To keep a constant number of steps per operation, new pointers must be added as the number of elements in the table grows.

Search Trees

As with every concurrent data structure, a concurrent implementation of a tree can be achieved just by placing a single lock protecting the whole structure. This approach can be improved by using a reader-writer lock, such that multiple threads are allowed to search information in the tree, while a single thread at the time is authorized to modify the tree. As expected, the exclusive lock for update operations creates a sequential bottleneck. Concurrency on lock-based implementations can be improved using some fine-locking strategy like placing a lock on each node.

[KL80] presents a concurrent implementation of binary search trees where update operations hold a constant number of locks. Besides, these locks have the property of excluding only other update operations, not blocking search operations. The drawback of this implementation is that it does not attempt to keep the tree balanced.

An alternative approach to allow concurrent modification of a tree based on fine-grained locking consists in forcing an operation to acquire a single lock protecting the subtree that is about to modify. This way, update operations working on disjoint sections of the tree can run in parallel.

For trees, in general, the maintenance operations such as splitting and merging nodes are performed as part of the update operations while inserting or removing elements from the tree. This tight relation between regular operations (*insert* or *delete*) with the maintenance operations is required to preserve the balancing properties of the tree. However, if we relax the balancing property, then maintenance operations can be done separate from regular operations. For instance, B^{link} -trees [Sag86] provide a *compress process* in charge of merging nodes, which can run concurrently with regular operations. By separating maintenance operations from regular ones, we let them run concurrently on different processors or in the background by specialized threads.

The idea of splitting maintenance operations from regular operations was initially described for red-black trees [GS78] and first implemented for AVL trees [AVL62] supporting *insert* and *search* operation in [Kes83]. An implementation supporting also *delete* is provided in [NSSW87]. These implementations improve concurrency by dividing balancing tasks in small transformations that can be applied locally and independently one from each other. It can be shown that the approach presented in [NSSW87] guarantees that each update operation needs at most $O(\log N)$ rebalancing operations for a N -node AVL tree. Similar results have been extended for B -trees [LF95, NSSW87] and red-black trees [BL94, NSS91]. On the non-blocking side we can find implementations based on dynamic software transactional memory mechanisms [Fra03, HLMI03].

Other existing implementations in the literature include studies about failure recovery [ML92] and implementations for *generalized search trees* (GiSTs) [KMH97]. GiSTs simplify the design of search trees without repeating the delicate work involving concurrency. Finally there exists some works [Lar01, Lar02] involving insertion and/or deletion of group of values.

Skiplists

A skiplist [Pug90] is a data structure that implements sets, maintaining several sorted singly-linked lists in memory. In fact, skiplists can be seen as virtual tree structured in multiple levels, where each level consists of a single linked list. The skiplist property establishes that the list at level $i + 1$ is a sublist of the list at level i . Each node in a skiplist stores a value and at least the pointer corresponding to the lowest level list. Some nodes also contain pointers at higher levels, pointing to the next element present at that level. The advantage of skiplists is that they are simpler and more efficient to implement than search trees, and search is still (probabilistically) logarithmic.




Figure 1.5: A skiplist with 4 levels

Consider for instance the skiplist shown in Fig. 1.5. Contrary to single-linked lists implementations, higher-level pointers allow to *skip* many elements during the search. A search is performed from left to right in a top down fashion, progressing as much as possible in a level before descending. For instance, in Fig. 1.5 a search for value 16 starts at level 3 of node *head*. From *head*, considering level 3, we can reach *tail* in a single jump, which is beyond the possible position of 16. Therefore, the search algorithm moves down one level, to level 2 of node *head*. The *head* successor at level 2 contains value 8, which is smaller than 16, so it is safe to go to node 8 considering level 2. This way, the search continues at level 2 until a node containing a greater value is found. At that moment, the search moves down one level again and so on. The expected logarithmic search follows because the probability of any given node occurring at a certain level decreases by 1/2 as a level increases (see [Pug90] for an analysis of the running time of skiplists).

Priority Queues

A concurrent priority queue is a data structure linearizable to a sequential priority queue, providing operations *insert* and *delete-min*. Many concurrent implementations of priority queues are based on heaps. When using heaps, insertion is done in bottom-up fashion while deletion of the minimum element is carried out from the root down to the leaves. As one operation proceeds from the leaves and the other from the root, there exists a high probability of achieving deadlock. A possibility to overcome this problem is the use of specialized threads in charge of cleaning up the data structure [BB87]. An alternative consists in implementing *insert* and *delete-min* operations both in a top down fashion [RK88]. An improvement of this algorithm consists in performing consecutive insertions on opposite sides of the heap [Aya90].

In [HMPS96] a heap-based algorithm is presented which does not require the acquisition of multiple locks while traversing the heap. For such purpose, it proceeds by locking a variable that stores the size of the heap, plus a lock over the first and last element in the heap. Other existing implementations [HW91] consist in concurrent priority queues based on concurrent versions of *fibonacci heaps* [FT87].

Non-blocking linearizable heap-based priority queues have been proposed by several authors [Her93, Bar94, IR93]. Even skiplists have been considered for implementing concurrent lock-free versions priority queues [ST03].

Other implementations include priority pools [HW91], where *delete-min* operation is not guaranteed to be linearizable. [Joh91] shows the implementation of a *delete bin* mechanism in charge of keeping elements deleted through *delete-min*, and thus, reducing the load when performing concurrent *delete-min* operations. In [SZ99] priority pools are implemented using skiplist. In general, the relaxed priority pools semantics allows for a significant increment of concurrency.

1.2 Decision Procedures for Pointer Based Data Structures

Pointer based data structures are currently widely used in computer science. In the previous section we have described some of them, including stacks, queues, lists and trees. Here we focus on decision procedures for the automatic verification of properties on these data structures. Many approaches have been suggested and studied in order to reason about them [BRS99, DN03, BPZ05, BR06, Nel83, LQ06, MN05, JJKS97] since the pioneer work of Burstall [Bur72]. However, most of them suffer from extreme difficulties to incorporate reasoning in a wealth of decidable theories over data and pointer values.

The key aspect in many of these approaches resides in the availability of decision procedures to reason about cells (the basic building blocks of the data structure), memories and a reachability notion induced by following pointers. As reachability is not a first-order concept, some features must be added in order to cope with it. Hence, despite the existence of precise and automatic techniques to reason about pointer reachability [JJKS97], not many approaches focus on the combination of such techniques with decision procedures capable of reasoning about data, pointers or locks. As a consequence, approximate solutions have come out and little is known about the combination of such logics with decidable first-order theories to reason about data and pointer values. In some cases, the information over data and pointers is abstracted away to make the use of reachability tools [DN03] possible. In other cases, a first-order approximation of reachability is used [Nel83] so that decision procedures for the theories of pointers and data can be used. In any case, both approaches mentioned above suffer from lack of precision when used.

In [BRS99], the decidability of a logic for pointer-based data structure is proved by showing that it enjoys from finite model property. That is, given a formula, a finite model can be constructed and used to verify the satisfiability of such formula. However, this logic can only reason about reachability on lists, trees and graphs, but it cannot reason about the data stored in the data structure.

A generalization of [BRS99] is presented in [YRS⁺06], but the emphasis once more is put on expressing complex shape constraints rather than considering the data stored in the data type. Regarding separation logic [Rey02] a decision procedure for lists based on finite model property is described in [BCO04], but it abstracts away the theories over data and pointers. Meanwhile, in combination with predicate abstraction, [BPZ05, BR06] describe decision procedures for logics abstracting away theories over data.

Other logics have been proposed for reasoning about complex data structures [BDES09], but they usually require the use of quantifiers, preventing them from being combined with other theories.

In [RZ06b] TLL is introduced. TLL is a logic for describing single-linked lists as well as a combination based decision procedure capable of reasoning about reachability. Moreover, it can be extended with available decision procedures for fragments of first-order logic. In TLL it is possible to reason about cells, memory and reachability. In addition, it has the property to be easily combined with other quantifier-free theories capable of reasoning, for instance, about data. In fact, the use of combination schemas [RRZ05] based on Nelson-Oppen [NO79] let them combine TLL with a wide range of available decision procedures for various decidable theories in first-order logic, such as BAPA [KNR05], for example.

The logics we present in this work are an extension of TLL, letting them reason about concurrent lock-based single-linked lists and concurrent skiplists. A key aspect is that both theories remain quantifier-free and with the same combination properties as TLL. Hence, they become perfect candidates for being combined with other theories in order to verify complex systems.

2

Preliminaries

As we said, our final goal is the verification of concurrent data structures. For such purpose, we use verification diagrams. Verification diagrams can be seen as a formula automaton with extra components. Then, the verification process consists in the following steps. First, we construct what we call the *most general client* of the datatype. That is, a program that non-deterministically performs calls to the operations provided by the data structure. Usually, those operations are INSERT, SEARCH and REMOVE. In addition, some non-executable ghost is added to aid the verification process. We can then proceed with the construction of a fair transition system representing the concurrent execution of M different instances of such *most general client*. The temporal property we want to verify is expressed as a formula in linear temporal logic and the proof that the system does in fact satisfy the property is represented by the verification diagram. Fig. 2.1 represents the main idea behind our method as well as the role played by decision procedures within this scheme. In this work, we just focus on the development of decision procedures for specific concurrent data structures.

The construction of a verification diagram is not an automatic task. Someone must sketch it, following the reasoning of the system under analysis. An advantage is that, after it has been constructed, verification can be done automatically. The verification that the diagram satisfies the temporal formula can be done

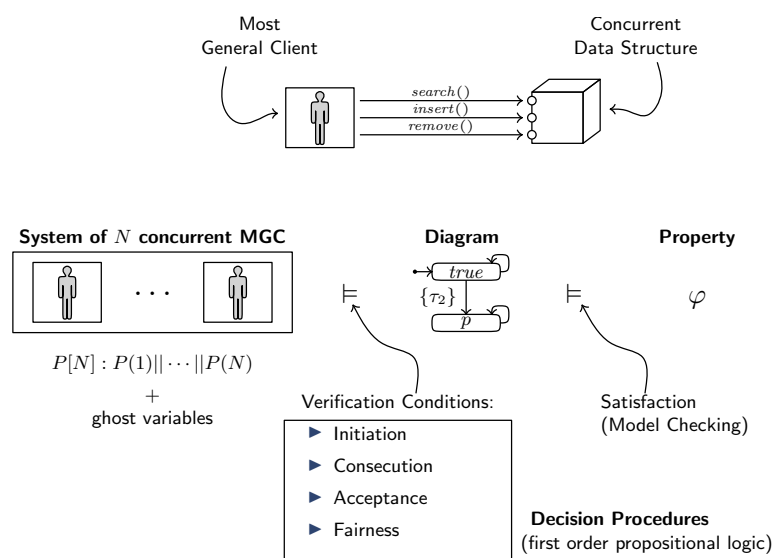


Figure 2.1: Sketch of our proposed approach

just by model checking. On the other hand, to verify that the diagram abstracts the system represented by the fair transition system, a number of verification conditions must be checked. If we count with a decision procedure capable of reasoning about the theories involved in such verification conditions, then this process can be done automatically. In Chapters 4 and 5 we focus on decision procedures for two specific data types: concurrent lists and skiplists. Meanwhile, in this chapter, we settle the basis of the verification process describing the use of regional logic (distinguishing our approach from the ones based on separation logic) and introducing the concept of verification diagrams, which will later be used in some motivating examples.

2.1 Regional Logic

We are interested in the formal verification of implementations of data structures, in particular in temporal verification (liveness and safety properties) of sequential and concurrent implementations. This verification activity requires to deal with unbounded mutable data. One popular approach to verification of heap programs is separation logic [Rey02].

Separation logic has been widely used up to now in the verification of a wide range of sequential programs. This power of expressiveness is not limited to sequential programs, as some advances on verification of concurrent programs show [GCPV09]. However, some data structures such as skiplists are problematic for separation-like approaches due to the aliasing and memory sharing between nodes at different levels. Based on the success of separation logic some researchers have extended this logic to deal with concurrent programs [Vaf07, HAN08], but concurrent datatypes follow a programming style in which the activities of concurrent threads are not structured according to critical regions with memory footprints. In these approaches based on separation logic memory regions are implicitly declared (hidden in the separation conjunction), which makes the reasoning about unstructured concurrency more cumbersome. Besides, despite the fact that some fragments of separation logic have been proved decidable [BCO04], it is still not completely clear how to integrate separation logic to current SMT solvers.

We use explicit regions to represent the manipulation of memory during the execution of the system. This reasoning is handled by extending the program code with ghost variables of type **rgn**, and ghost updates of these variables. Variables of type **rgn** represent finite sets of object references stored in the heap. Regional logic [BNR08] provides a rich set of language constructs and assertions. However, it is enough for our purposes to use only a small fragment of regional logic. The term **emp** denotes the empty region and $\langle x \rangle$ represents the singleton region whose only object is the one referenced by x . Traditional set-like operators such as \cup , \cap and \setminus are also provided and can be applied to **rgn** variables. The assertion language allows reasoning involving mutation and separation. Given two **rgn** expressions r_1 and r_2 we can assert whether they are equal ($r_1 = r_2$), one is contained into the other ($r_1 \subseteq r_2$) or they are completely disjoint ($r_1 \# r_2$).

2.2 Verification Diagrams

We now sketch the important notions from [BMS95, Sip99]. As we said, verification diagrams provide an intuitive way to abstract temporal proofs over fair transition systems (FTS). A FTS Φ is a tuple $\langle V, \Theta, \mathcal{T}, \mathcal{J} \rangle$ where V is a finite set of variables, Θ is an initial assertion, \mathcal{T} is a finite set of transitions and $\mathcal{J} \subseteq \mathcal{T}$ contains the fair transitions. Here we do not consider strong fairness. A *state* is an interpretation of V . We use \mathcal{S} to denote the set of all possible states. A transition $\tau \in \mathcal{T}$ is a function $\tau : \mathcal{S} \rightarrow 2^{\mathcal{S}}$, which is usually represented by a first-order logic formula $\rho_\tau(s, s')$ describing the relation between the values of the variables in a state s and in a successor state s' . Given a transition τ , the state predicate $En(\tau)$ denotes whether there exists a successor state s' such that $\rho_\tau(s, s')$.

A computation of Φ is an infinite sequence of states such that:

- (a) the first state satisfies Θ ;
- (b) any two consecutive states satisfy ρ_τ for some $\tau \in \mathcal{T}$;

(c) for each $\tau \in \mathcal{J}$, if τ is continuously enabled after some point, then τ is taken infinitely many times.

We use $\mathcal{L}(\Phi)$ to denote the set of computations of the FTS Φ . Given a formula φ , $\mathcal{L}(\varphi)$ denotes the set of sequences satisfying φ . A FTS Φ satisfies a temporal formula φ if all computations of Φ satisfy φ , i.e., $\mathcal{L}(\Phi) \subseteq \mathcal{L}(\varphi)$.

We use $SF(\varphi)$ to denote the set of all atomic subformulas of φ . A φ -propositional state, denoted by s_φ , assigns a truth-value $s_\varphi[q]$ to each $q \in SF(\varphi)$. Satisfaction of a formula φ over a φ -propositional model $\sigma_\varphi : s_{\varphi,0}, s_{\varphi,1}, \dots$ is defined inductively where for an atomic subformula q , $s_\varphi \models q$ iff $s_\varphi[q] = \text{true}$.

A σ -propositional model σ_φ is an infinite sequence of σ -propositional states. The propositional language of φ , $\mathcal{L}^p(\varphi)$, is the set $\{\sigma_\varphi \mid \sigma_\varphi \models \varphi\}$. A propositional projection s_φ^p of a state s is a φ -propositional state such that for each $q \in SF(\varphi)$, $s_\varphi^p \models q$ iff $s \models q$.

Then, a verification diagram (VD) $\Psi : \langle N, N_0, E, \mu, \eta, \mathcal{F}, \Delta, f \rangle$ is a formula automaton with extra components, abstracting the fair transition system Φ and still preserving the satisfaction of formula φ . All components of a verification diagram are:

- N is a finite set of nodes.
- $N_0 \subseteq N$ is the set of initial nodes.
- $E \subseteq N \times N$ is a set of edges.
- $\mu : N \rightarrow F(V)$ is a labeling function mapping nodes to assertions over V .
- $\eta : E \rightarrow 2^\tau$ is a labeling function assigning sets of transitions to edges.
- $\mathcal{F} \subseteq 2^{E \times E}$ is an edge acceptance set of the form $\{(P_1, R_1), \dots, (P_m, R_m)\}$.
- $\Delta \subseteq \{\delta \mid \delta : \mathcal{S} \rightarrow \mathcal{D}\}$ is a set of ranking functions from states to a well founded domain \mathcal{D} .
- f maps nodes into propositional formulas over atomic subformulas of φ .

If $n \in N$, we use $next(n)$ to denote the set $\{\tilde{n} \in N \mid (n, \tilde{n}) \in E\}$ and $\tau(n)$ for $\{\tilde{n} \in next(n) \mid \tau \in \eta(n, \tilde{n})\}$. For each $(P_j, R_j) \in \mathcal{F}$ and for each $n \in N$, Δ contains a ranking function $\delta_{j,n}$. An infinite sequence of nodes $\pi = n_0, n_1, \dots$ is a path if $n_0 \in N_0$ and for each $i > 0$, $(n_i, n_{i+1}) \in E$. A path π is accepted if for each pair $(P_j, R_j) \in \mathcal{F}$ some edge of R_j occurs infinitely often in π or all edges that occur infinitely often in π are also in P_j . An infinite path π is fair when, for any just transition τ , if τ is enabled on all nodes that appear infinitely often in π then τ is taken infinitely often.

Given a sequence of states $\sigma = s_0, s_1, \dots$ of Φ , a path $\pi = n_0, n_1, \dots$ is a trail of σ whenever $s_i \models \mu(n_i)$ for all $i \geq 0$. An infinite sequence of states σ is a computation of Ψ whenever there exists an accepting trail of σ such that is also fair. $\mathcal{L}(\Psi)$ is the set of computations of Ψ .

A computation $\sigma = s_{\varphi,0}, s_{\varphi,1}, \dots$ is a φ -propositional model of Ψ if there is a fair and accepting path n_0, n_1, \dots in Ψ such that for all $i \geq 0$, $s_{\varphi,i} \models f(n_i)$. Finally, $\mathcal{L}^p(\Psi)$ denotes the set of all φ -propositional models of Ψ .

A verification diagram shows that $\Phi \models \varphi$ via the inclusions $\mathcal{L}(\Phi) \subseteq \mathcal{L}(\Psi) \subseteq \mathcal{L}(\varphi)$. The map f is used to check $\mathcal{L}(\Psi) \subseteq \mathcal{L}(\varphi)$. To show $\mathcal{L}(\Phi) \subseteq \mathcal{L}(\varphi)$ it is enough to prove the following verification conditions:

Initiation: there is at least one initial node from N_0 satisfying the initial conditions of the fair transition system Φ .

Consecution: Any τ -successor of a state satisfying $\mu(n)$ does satisfy the labeling predicate of some successor node of n . This means that for every node $n \in N$ and transition $\tau \in \mathcal{T}$,

$$\mu(n)(s) \wedge \rho_\tau(s, s') \rightarrow \mu(next(n))(s')$$

Acceptance: For each pair (P_j, R_j) member of the acceptance list and for any $e = (n_1, n_2) \in E$ and $\tau \in \mathcal{T}$, when we take τ from an arbitrary state s , if $e \in P_j \setminus R_j$ then δ_j cannot increase, while if $e \notin P_j \cup R_j$ then function δ_j must decrease. That is:

1. if $(n_1, n_2) \in P_j \setminus R_j$ then

$$\rho_\tau(s, s') \wedge \mu(n_1)(s) \wedge \mu(n_2)(s') \rightarrow \delta_{j, n_1}(s) \succeq \delta_{j, n_2}(s')$$

2. if $(n_1, n_2) \notin P_j \cup R_j$ then

$$\rho_\tau(s, s') \wedge \mu(n_1)(s) \wedge \mu(n_2)(s') \rightarrow \delta_{j, n_1}(s) \succ \delta_{j, n_2}(s')$$

Fairness: For each $e = (n_1, n_2) \in E$ and $\tau \in \eta(e)$:

1. τ is guaranteed to be enabled in every $\mu(n_1)(s)$:

$$\mu(n_1)(s) \rightarrow \text{En}(\tau)$$

2. Any τ -successor of a state satisfying $\mu(n_1)$ satisfies the label of some node in $\tau(n_1)$:

$$\mu(n_1)(s) \wedge \rho_\tau(s, s') \rightarrow \mu(\tau(n_1))(s')$$

Satisfaction:

1. For all $n \in N$, if $s \models \mu(n)$ then $s_\varphi^p \models f(n)$
2. $\mathcal{L}^p(\Psi) \subseteq \mathcal{L}^p(\varphi)$

Some verification conditions such as *satisfaction* can be checked through model checking, just verifying whether the language accepted by the diagram is contained into the language of the temporal formula, or equivalently, checking if the intersection with the language of the negation of the formula is empty.

The remaining verification conditions can be automatically verified if we count with an appropriated decision procedure. In such case, the components involved in the formulas depend on the data structure we are verifying. In the following chapters, we present two concurrent data structures: concurrent lock-coupling lists and concurrent skiplists. Moreover, we present some examples of properties we would like to verify over these datatypes in addition to adequate decision procedures for each of them.

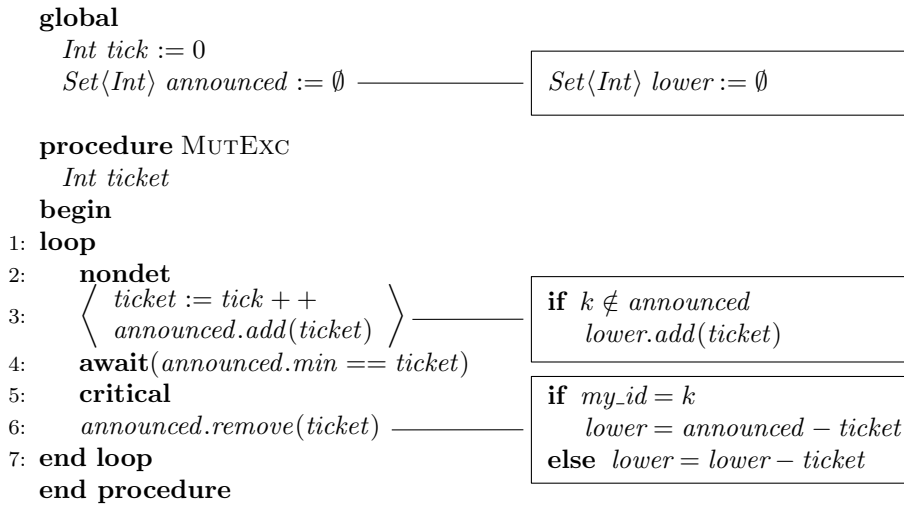


Figure 2.2: MUTEXC: A mutual exclusion algorithm (left), and annotations (right)

We now illustrate the use of generalized verification diagrams by showing the proof of a response property of a simple mutual exclusion algorithm. The algorithm is structured the following way. Each thread that wants to access the critical section acquires an increasing number (ticket) and announces its intention to enter the critical section by adding its ticket to a shared global set of tickets. Then, each thread waits until its ticket number becomes the lowest value in the set before entering the critical section. After a thread leaves the critical section it removes its ticket from the set. For the sake of simplicity, we assume that all operations are atomic.

Fig. 2.2 (left) shows `MUTEXC`, a program that implements this mutual exclusion protocol to protect access to the critical section at line 5, using two global variables. The `Int` variable `tick` stores the shared increasing counter. The `Set` variable `announced` stores the ticket numbers of all threads that are trying to access the critical section. We want to prove that if an arbitrary thread k requests access to the critical section, eventually it succeeds. To aid in the verification, ghost variable `lower` is used to keep the track of all requests that have been done before thread k . We use the cardinality of such set as a ranking function to prove that eventually, the request of thread k is taken into consideration. In general, ghost annotations are introduced in order to keep track of special conditions and values that aid in the verification process. Variable `ticket` stores the ticket number assigned to the thread while variable `pc` keeps the program counter. If S is a set of variables v_1, \dots, v_n , we use $pres(S)$ to denote that the value of all variables in S

$$\begin{aligned}
V &: \{tick, announced, ticket^{[1]}, ticket^{[2]}, pc^{[1]}, pc^{[2]}\} \\
\rho_{\tau_1^{[1]}} &: pc^{[1]} = 1 \wedge pc^{[1]} = 2 \wedge pres(V \setminus \{pc^{[1]}\}) \\
\rho_{\tau_2^{[1]}} &: pc^{[1]} = 2 \wedge pc^{[1]} = 3 \wedge pres(V \setminus \{pc^{[1]}\}) \\
\rho_{\tau_3^{[1]}}^T &: pc^{[1]} = 3 \wedge pc^{[1]} = 4 \wedge k \in announced \\
&\quad \wedge ticket'^{[1]} = tick + 1 \\
&\quad \wedge tick' = tick + 1 \\
&\quad \wedge announced' = announced \cup \{ticket^{[1]}\} \wedge pres(V \setminus \{pc^{[1]}, ticket^{[1]}, announced\}) \\
\rho_{\tau_3^{[1]}}^F &: pc^{[1]} = 3 \wedge pc^{[1]} = 4 \wedge k \notin announced \\
&\quad \wedge ticket'^{[1]} = tick + 1 \\
&\quad \wedge tick' = tick + 1 \\
&\quad \wedge announced' = announced \cup \{ticket^{[1]}\} \\
&\quad \wedge lower' = lower \cup \{ticket^{[1]}\} \wedge pres(V \setminus \{pc^{[1]}, ticket^{[1]}, announced, lower\}) \\
\rho_{\tau_4^{[1]}} &: pc^{[1]} = 4 \wedge pc^{[1]} = 5 \wedge announced.min = ticket^{[1]} \wedge pres(V \setminus pc^{[1]}) \\
\rho_{\tau_5^{[1]}} &: pc^{[1]} = 5 \wedge pc^{[1]} = 6 \wedge pres(V \setminus \{pc^{[1]}\}) \\
\rho_{\tau_6^{[1]}}^T &: pc^{[1]} = 6 \wedge pc^{[1]} = 7 \wedge my_id = k \\
&\quad \wedge announced' = announced \setminus \{ticket^{[1]}\} \\
&\quad \wedge lower' = announced \setminus \{ticket^{[1]}\} \wedge pres(V \setminus pc^{[1]}, announced, lower) \\
\rho_{\tau_6^{[1]}}^F &: pc^{[1]} = 6 \wedge pc^{[1]} = 7 \wedge my_id \neq k \\
&\quad \wedge announced' = announced \setminus \{ticket^{[1]}\} \\
&\quad \wedge lower' = lower \setminus \{ticket^{[1]}\} \wedge pres(V \setminus pc^{[1]}, announced, lower) \\
\rho_{\tau_7^{[1]}} &: pc^{[1]} = 7 \wedge pc^{[1]} = 1 \wedge pres(V \setminus \{pc^{[1]}\}) \\
\Theta &: tick = 0 \wedge announced = \emptyset \wedge lower = \emptyset \wedge pc^{[1]} = 1 \wedge pc^{[2]} = 1
\end{aligned}$$

Figure 2.3: Fair transition system for `MUTEXC`

is preserved by the transition. That is, that $v'_1 = v_1 \wedge \dots \wedge v'_n = v_n$.

For simplicity, we assume a system made just from two threads: T_1 and T_2 . We would like to prove that whenever T_1 shows it interest to enter the critical section, then it eventually succeeds.

Throughout this work we use $at_p_n^{[k]}$ to denote that thread k is at line n of program p . Similarly, we use $at_p_{n..m}^{[k]}$ to represent $\bigwedge_{i \in [n..m]} at_p_i^{[k]}$. If v is a local variable of a program, we use $v^{[k]}$ to denote the local instance of variable v owned by thread k . As usual, we use primed variables to describe the values of the variables after the transition is taken.

We can define the fair transition representing an instance of `MutExc` for two threads as shown in Fig. 2.3. Here we limit ourselves to show just the transition relations for thread T_1 . A set of analog transition relations are required for thread T_2 .

Before we describe the property and the diagram, we introduce some notation to ease the proof. We use $active(i)$ for $at_MutExc_{4,5,6}^{[i]}$, $wants(i)$ for $at_MutExc_3^{[i]}$, and $critical(i)$ for $at_MutExc_5^{[i]}$. Diagrams for the verification of safety properties are quite easy to construct. For instance, imagine we want to verify mutual exclusion of the system. Then, we require to verify a set of invariants. For instance, below $\varphi_1(i)$ describes the fact that every time a thread i has a valid ticket, the value of this ticket is smaller than $tick$. Formula $\varphi_2(i, j)$ establishes that two different threads cannot have the same ticket. Formula $\varphi_3(i)$ specifies that if a thread is in $active(k)$ then its ticket is in the announced set. Formula $\varphi_4(i)$ establishes that the thread in the critical section owns the lowest ticket. Finally, mutual exclusion is expressed as φ_{mutex} :

$$\begin{aligned} \varphi_1(i) &\triangleq \square(active(i) \rightarrow ticket(i) < tick) \\ \varphi_2(i, j) &\triangleq \square(i \neq j \wedge active(i) \wedge active(j) \rightarrow ticket(i) \neq ticket(j)) \\ \varphi_3(i) &\triangleq \square(active(i) \rightarrow ticket(i) \in announced) \\ \varphi_4(i) &\triangleq \square(critical(i) \rightarrow min(announced) = ticket(i)) \\ \varphi_{mutex}(i, j) &\triangleq \square(i \neq j \rightarrow \neg(critical(i) \wedge critical(j))) \end{aligned}$$

It is easy to see that these invariants hold. In this work we are not interested in describing parametric systems. Thus, given a system, we assume a fixed number of threads running concurrently. This way, all properties need to be verified for every thread in the system. If we would like to construct a diagram for, let say, $\varphi_1(T_1)$, then it would look like the one shown in Fig. 2.4. Despite verification can be done automatically, the construction of a diagram remains as a manual task. Hence, it is the user the one who must sketch it, following its intuition, trying to capture the behavior of the program.




Figure 2.4: Verification diagram for $\varphi_1(T_1)$

We show now how to verify liveness properties. The program `MutExc` satisfies that every thread that wants to enter the critical section eventually does, formally expressed in LTL using the following response property [MP95]:

$$\psi(k) \triangleq \square(wants(k) \rightarrow \diamond critical(k))$$

In the verification we use the ghost variables and ghost updates, shown in Fig. 2.2 (right). This variables allow to keep track of important aspects of the history of the computation, in this case, the set of lower tickets (with respect to $ticket^{[k]}$).

We show how to construct the diagram for verifying $\psi(T_1)$, that is, that the property holds for thread T_1 . Two diagrams are shown in Fig. 2.5. The first one labels nodes with an informal description of the state they represent. The second one is labelled with program positions and is the one used for the

verification. Formally, the diagram is defined by:

$$\begin{aligned}
 N &\triangleq \{n_i \mid 1 \leq i \leq 6\} \\
 N_0 &\triangleq \{n_1\} \\
 \mathcal{F} &\triangleq \{(P, R)\} \text{ with } R = \{(n_4, n_1)\} \text{ and } P = E - (R \cup \{(n_6, n_3)\}) \\
 f(n) &\triangleq \mu(n) \\
 \delta(n, s) &\triangleq |\text{lower}|
 \end{aligned}$$

The edges shown in Fig. 2.5 represent the set E . Similarly, μ and η are defined, respectively, as the functions that label nodes and edges with the formulas and transitions depicted in the same figure.

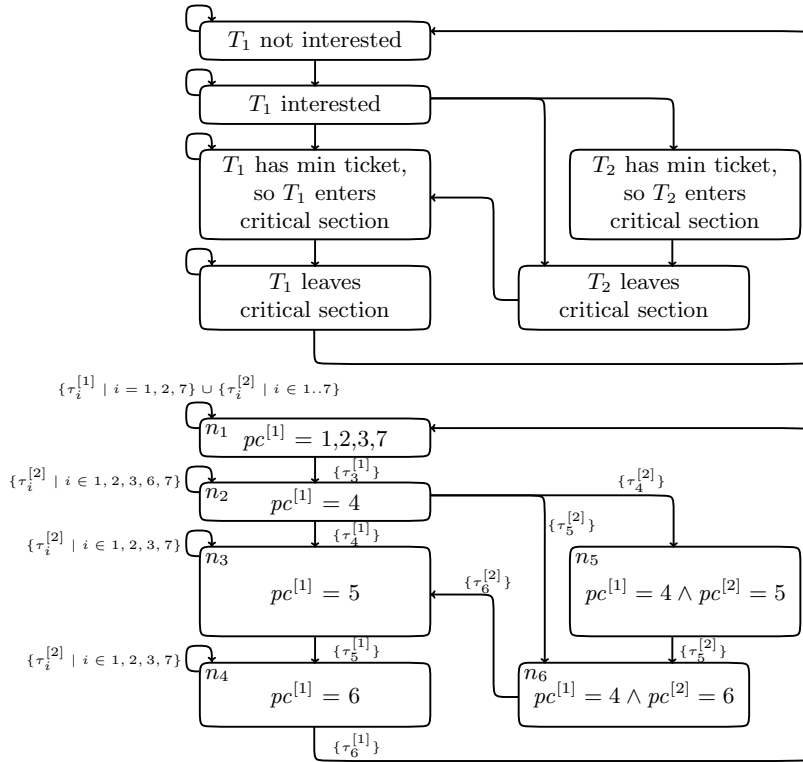


Figure 2.5: Verification diagram for $\psi(T_1)$

3

Concurrent Lists and Skiplists

In this work we focus on decision procedures for two pointer based data structures: concurrent lock-coupling singly-linked lists and concurrent skiplists. In this chapter, we introduce both data types, describe their structure and give the main operations that manipulate them. Besides, for each data type we sketch a proof requiring temporal reasoning. This will let us present the components we will require in our decision procedures for each data type.

3.1 Concurrent Lock-Coupling Lists

Lock-coupling concurrent lists [HS08, VHHS06] are ordered lists with non-repeating elements, in which each node is protected by a lock. A thread advances through the list acquiring the lock of the node it visits. This lock is only released after the lock of the next node has been acquired. The following *List* and *Node* structures are used to maintain the data of a concurrent list:

```
class List {Node* list; }
class Node {Value val; Node* next; Lock lock; }
```

A *List* contains one field pointing to the *Node* representing the head of the list. A *Node* consists of a value, a pointer to the next *Node* in the list and a lock. We assume that the operating system provides the operations *lock* and *unlock* to acquire and release a lock. Every list has two sentinel nodes, *Head* and *Tail*, with phantom values representing the lowest and highest possible values. For simplicity, we assume such nodes cannot be removed or modified. Fig 3.1 presents a concurrent lock-coupling skiplist.




Figure 3.1: A lock-coupling singly-linked list with elements 5 and 9 locked

Concurrent Lock-Coupling Lists are used to implement sets, so they offer three operations:

- **LOCATE**, shown as Algorithm 3.1, finds an element traversing the list. This operation returns the pair consisting of the desired node and the node that precedes it in the list. If the element is not found the *Tail* node is returned as the current node. A **SEARCH** operation, shown as Algorithm 3.2, that decides whether an element is in the list can be easily extended from **LOCATE**.

- INSERT, shown as Algorithm 3.3, inserts a new element in the list, using LOCATE to determine the position at which the element must be inserted. The operation *add* returns *true* upon success, otherwise it returns *false*.
- REMOVE, depicted as Algorithm 3.4, deletes a node from the list by redirecting the next pointer of the previous node appropriately.

Algorithm 3.1 Locate program for concurrent lock-coupling singly-linked lists

```

1: procedure LOCATE(Value e)
2:   prev := Head
3:   prev.lock()
4:   curr := prev.next
5:   curr.lock()
6:   while curr.val < e do
7:     prev.unlock()
8:     prev := curr
9:     curr := curr.next
10:    curr.lock()
11:  end while
12:  return (prev, curr)
13: end procedure

```

Algorithm 3.2 Search for concurrent lock-coupling singly-linked lists

```

1: procedure SEARCH(Value e)
2:   prev, curr := LOCATE(e)
3:   if curr.val = e then
4:     result := true
5:   else
6:     result := false
7:   end if
8:   curr.unlock()
9:   prev.unlock()
10:  return result
11: end procedure

```

Algorithm 3.3 Insertion for concurrent lock-coupling singly-linked lists

```

1: procedure INSERT(Value e)
2:   prev, curr := LOCATE(e)
3:   if curr.val ≠ e then
4:     aux := new Node(e)
5:     aux.next := curr
6:     prev.next := aux
7:     result := true
8:   else
9:     result := false
10:  end if
11:  prev.unlock()
12:  curr.unlock()
13:  return result
14: end procedure

```

Algorithm 3.4 Remove for concurrent lock-coupling singly-linked lists

```

1: procedure REMOVE(Value  $e$ )
2:    $prev, curr := LOCATE(e)$ 
3:   if  $curr.val = e$  then
4:      $aux := curr.next$ 
5:      $prev.next := aux$ 
6:      $result := true$ 
7:   else
8:      $result := false$ 
9:   end if
10:   $prev.unlock()$ 
11:   $curr.unlock()$ 
12:  return  $result$ 
13: end procedure

```

Algorithm 3.5 Most general client for concurrent lock-coupling singly-linked lists

```

1: procedure MGC
2:   while  $true$  do
3:      $e := NondetPickElem$ 
4:     nondet
5:      $\left[ \begin{array}{l} \text{call SEARCH}(e) \\ \text{or} \\ \text{call INSERT}(e) \\ \text{or} \\ \text{call REMOVE}(e) \end{array} \right]$ 
6:   end while
7: end procedure

```

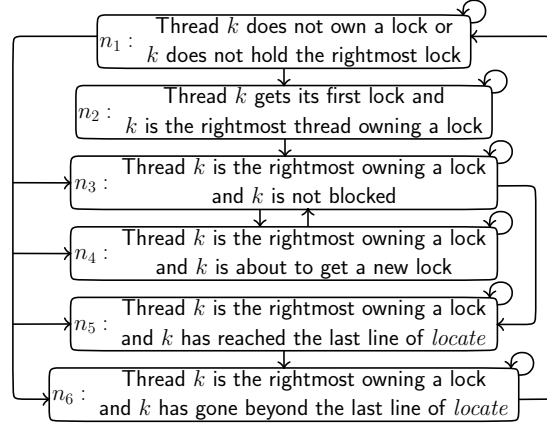
We can also construct the most general client of the concurrent-list datatype: the program MGC that repeatedly chooses non-deterministically a method and its parameters. Such client is shown as Algorithm 3.5.

Example 3.1 (Termination of rightmost thread)

Imagine we want to prove that on a lock-coupling concurrent list as the one introduced in this chapter, the thread owning the nearest lock to the end of the list eventually terminates. To prove it, we construct a fair transition system $\mathcal{S}[N]$ parametrized by the total number of threads N , in which all threads run MGC. Let $\psi(k)$ be the temporal formula that describes that if thread k owns the last lock in the list (i.e., the lock nearest to node *Tail*) then it eventually terminates. The verification problem is then casted as $\mathcal{S}[N] \models \psi(k)$, for all N . A sketch of a verification diagram is depicted in Fig. 3.2. We say that a thread is the rightmost owning a lock when there is no other thread owning a lock that protects a *Node* closer to the tail.

Each diagram node is labeled with a predicate. This predicate captures the set of states of the transition system that the node abstracts. Edges represent transitions between states abstracted by the nodes. For the sake of simplicity, in the diagram at Fig. 3.2 we have replaced the predicates labeling each node by an informal description of what they represent.

Node n_1 denotes all states in which thread k has not acquired a lock yet, or it does not own the rightmost lock. Node n_2 represents states in which thread k has acquired the rightmost lock in the list and it is about to execute a non-blocking statement. Node n_3 denotes the states at which thread k is the rightmost thread owning a lock, and no other thread is blocking it. Node n_4 represents the states on which thread k is about to acquire a lock so it could get blocked. Node n_5 denotes the states at which thread k

Figure 3.2: Sketched verification diagram for $\mathcal{S}[N] \models \psi(k)$

has reached the last line of program `LOCATE` meaning that no more locks needs to be acquired. Finally, node n_4 denotes the state at which thread k has gone beyond the last line of program `LOCATE`. At this point, the operations done by thread k consists only in modifying some pointers and releasing locks.

Since we assume that all threads are running program `MGC`, by "thread j terminates" we mean that thread j has reached the last program line of `LOCATE`. Notice that once thread j has arrived to the last line of `LOCATE` it is easy to show that j cannot be blocked.

Checking the proof represented by the verification diagram requires two activities. First, to show that all traces of the diagram satisfy the temporal formula $\psi(k)$, which can be performed by finite state model checking. Second, to prove that all computations of $\mathcal{S}[N]$ are traces of the verification diagram. This process involves the verification of formulas built from several theories. For instance, considering the execution of line 6 of program `INSERT` we should verify that the following condition holds:

$$at_INSERT_5^{[k]} \wedge IsLast(k) \wedge \left(r' = r \cup \langle aux^{[k]} \rangle \wedge \left(prev^{[k]}.next = aux^{[k]} \right) \right) \rightarrow at_INSERT_6^{[k]} \wedge IsLast'(k) \quad (3.1)$$

The predicate $prev^{[k]}.next = curr^{[k]}$ is in the theory of pointers, while $r' = r \cup \langle curr^{[k]} \rangle$ is in the theory of regions. Moreover, some predicates belong to a combination of theories, like $IsLast(k)$, which among other things establishes that $List(h, x, r)$ holds. $List(h, x, r)$ expresses that in heap h , starting from pointer x , the pointers form a list of elements following the `next` field, and that all nodes in this list form precisely the region r . *

Our intention here is not to show a full example or a complete verification condition, but just to give the reader the intuition of the construction our decision procedure should be able to deal with. At this point, it should be clear that we must deal with elements, addresses, thread identifiers, nodes, memory assignments, regions and locks. To accomplish the automatic verification, we must build a suitable decision procedure involving all these theories. Such decision procedure is described in Chapter 4.

3.2 Concurrent Skiplists

A skiplist [Pug90] is a data structure that implements sets, maintaining several sorted singly-linked lists in memory. Skiplists are structured in multiple levels, where each level consists of a single linked list. The skiplist property establishes that the list at level $i + 1$ is a sublist of the list at level i . Each node in a skiplist stores a value and at least the pointer corresponding to the lowest level list. Some nodes also contain pointers at higher levels, pointing to the next element present at that level. The advantage of skiplists is that they are simpler and more efficient to implement than search trees, and search is still (probabilistically) logarithmic.

In this section we present a simple concurrent implementation of skiplists using lock-coupling [HS08] to acquire and release locks. This implementation can be seen as an extension of concurrent lock-coupling lists [HS08, Vaf07] to multiple layers of pointers. This algorithm imposes a locking discipline, consisting of acquiring locks as the search progresses, and releasing a node's lock only after the lock of the next node in the search process has been acquired. A naïve implementation of this solution would equip each node with a single lock, allowing multiple threads to access simultaneously different nodes in the list, but protecting concurrent accesses to two different fields of the same node. The performance can be improved by carefully allowing multiple threads to simultaneously access the same node at different levels. We study here an implementation of the latter solution in which each node is equipped with a different lock at each level.

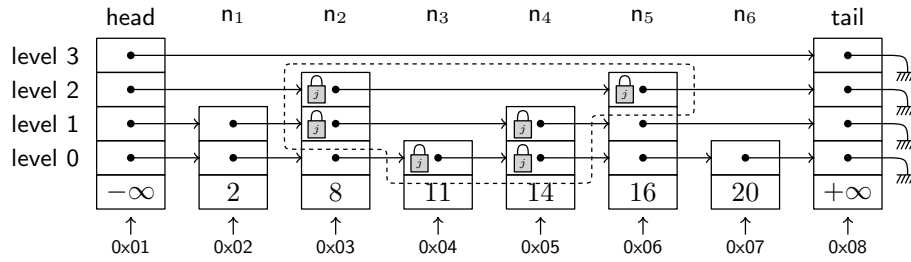


Figure 3.3: A skiplist with the masked region given by the fields locked by thread j

At execution time a thread uses locks to protect the access to only some fields of a given node. A precise reasoning framework needs to capture those portions of the memory protected by a set of locks, which may include only *parts* of a node. Approaches based on strict separation (separation logic [Rey02] or regional logic [BNR08]) do not provide the fine grain needed to reason about individual fields of shared objects. We now introduce the concept of *masked regions* to describe regions and the fields within. A masked region consists of a set of pairs formed by a region (*Node* cell) and a field (a skiplist level):

$$\text{mrng} \triangleq 2^{\text{Node} \times \mathbb{N}}$$

We call the field a mask, since it identifies which part of the object is relevant. For example, in Fig. 3.3 the region within dots represents the area of the memory that thread j is protecting. This portion of the memory is described by the masked region $\{(n_2, 2), (n_5, 2), (n_2, 1), (n_4, 1), (n_3, 0), (n_4, 0)\}$. As with regional logic, an empty set intersection denotes separation. In masked regions two memory nodes at different levels do not overlap. This notion is similar to data-groups [Lei98].

We now present the pseudo-code declaration of the *Node* and *SkipList* classes:

```

class SkipList { Node* head; Node* tail; }
class Node { Value val;
             Key key;
             Array<Node*>(K) next;
             Array<Node*>(K) lock; }

```

Throughout the present work we use boxes to denote ghost code added for verification purposes. Note that the structure is parametrized by a value K , which determines the maximum possible level of any node in the modeled skiplist. The fields *val* and *key* in the class *Node* contain the value and the key of the element used to order them. Then, we can store key-value pairs, or use the skiplist as a set of arbitrary elements as long as the key can be used to compare. The *next* array stores the pointers to the next nodes at each of the possible K different levels of the skiplist. Finally, the *lock* array keeps the locks, one for each level, protecting the access to the corresponding *next* field.

The *SkipList* class contains two pointer fields: *head* and *tail* plus a ghost variable field *r*. Field *head* points to the first node of the skiplist, and *tail* to the last one. Variable *r*, only used for verification purposes, keeps the (masked) region represented by all nodes in the skiplist with all their levels. In this

Algorithm 3.6 Unfair implementation of concurrent skiplist operation INSERT

```

1: procedure INSERT(SkipList sl, Value newval)
2:   Vector⟨Node*⟩upd[0..K - 1]
3:   lvl := randomLevel(K)
4:   Node*pred := sl.head
5:   pred.locks[K - 1].lock()
6:   Node*curr := pred.next[K - 1]
7:   curr.locks[K - 1].lock()
8:   for i := K - 1 downto 0 do
9:     if i < K - 1 then
10:      pred.locks[i].lock()
11:      if i ≥ lvl then
12:        curr.locks[i + 1].unlock()
13:        pred.locks[i + 1].unlock()
14:      end if
15:      curr := pred.next[i]
16:      curr.locks[i].lock()
17:    end if
18:    while curr.val < newval do
19:      pred.locks[i].unlock()
20:      pred := curr
21:      curr := pred.next[i]
22:      curr.locks[i].lock()
23:    end while
24:    upd[i] := pred
25:  end for
26:  Bool valueWasIn := (curr.val = newval)
27:  if valueWasIn then
28:    for i := 0 to lvl do
29:      upd[i].next[i].locks[i].unlock()
30:      upd[i].locks[i].unlock()
31:    end for
32:  else
33:    x := CreateNode(lvl, newval)
34:    for i := 0 to lvl do
35:      x.next[i] := upd[i].next[i]
36:      upd[i].next[i] := x
37:      x.next[i].locks[i].unlock()
38:      upd[i].locks[i].unlock()
39:    end for
40:  end if
41:  return  $\neg$ valueWasIn
42: end procedure

```

$L := L\{me \leftarrow sl.head\}$
 $U := U\{me \leftarrow sl.tail\}$
 $H := H\{me \leftarrow K - 1\}$

$U := U\{me \leftarrow pred.next[i + 1]\}$

$H := H\{me \leftarrow i\}$

$\mathbf{if} (i \geq lvl)\{L := L\{me \leftarrow curr\}\}$

$L = L\{me \leftarrow null\}$

implementation, *head* and *tail* are sentinel nodes, with $key = -\infty$ and $key = +\infty$, respectively. For simplicity, these nodes are not eliminated during the execution and their *val* field remains unchanged.

When considering concurrent datatypes, a craft design must be taken in consideration. Apparently correct implementations may lead to undesired behaviors. For instance, consider the implementation for INSERT given as Algorithm 3.6. This implementation does not properly works under the assumption of

strong fairness, since a thread may prevent other to progress under the assumption of an unfair scheduler.

Example 3.2

We show that the implementation described as Algorithm 3.6 does not ensure termination of all threads under the assumption of weak-fairness. The program has been enriched with ghost arrays L and U from thread identifiers to addresses and an array H from thread identifiers to skiplist levels. These three global arrays are used to store for each thread the addresses and level that delimit the section of the skiplists that can be potentially be modified by a thread. We use $e = x[i]$ to denote that e is the value stored in array x at position i . We also use $y = x\{i \leftarrow e\}$ to denote that y is the array resulting from taking x and modifying position i to store element e . This means that $y[i] = e$ while $y[j] = x[j]$ for any other j different from i . This way, $L[t]$ stores the lowest address than can be modified by thread t . Similarly, $U[t]$ denotes the upper address of the region modifiable by t . Meanwhile, $H[t]$ is the highest level that can be modified by thread t . In any case, we assume any thread can use me to refer to its thread identifier.

Notice that skiplists levels are numerated starting from 0. Then, a node at level 0 is said to have height 1. Similarly, a node that goes up to level 1 is said to be of height 2 and so on.

Consider now the skiplist described in Fig. 3.4 and a system with two thread: T_1 and T_2 . Imagine T_1 wants to insert a node with value 14 and height 1, while T_2 tries to insert value 16 with height 2.




Figure 3.4: An example of skiplist

T_1 begins by grabbing the lock at level 2 on node $-\infty$ and node 18, as shown in Fig. 3.5(a). As 18 is beyond the position where 14 should be inserted, it decides then to go a level below. Afterwards, the algorithm proceeds as depicted in Fig. 3.5(b) and 3.5(c). In the figure, we use a dashed line to denote the

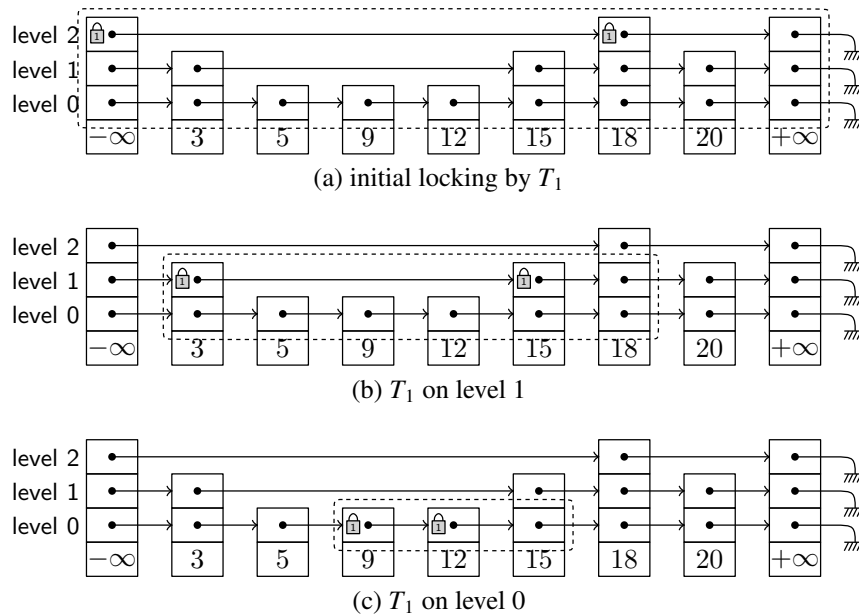


Figure 3.5: Progress of T_1 towards insertion of value 14

region of the skiplist that can potentially be modified by T_1 , i.e., the region limited by $L[T_1]$, $U[T_1]$ and $H[T_1]$.

At this moment, T_2 starts its execution. Fig. 3.6 shows the progress made by T_2 toward the insertion of a level 2 node with value 16. As before, in this case the heavily dashed region represents the masked region of the skiplist that can be modified by T_2 .




Figure 3.6: Progress of T_2 towards insertion of value 16

Notice that the potentially modifiable regions by thread T_1 and T_2 intersect, as shown in Fig. 3.6(c). In this case, it is quite easy to see that under the assumption of weak-fairness, if T_2 continuously perform the same insertion, it can prevent T_1 from progressing. However, no problem exists under the assumption of strong fairness, since T_1 is not continuously enabled. In fact, it is quite easy to see that T_1 becomes disabled every time T_2 gets the lock at level 0 on node 15. *

A correct implementation of the SEARCH, INSERT and REMOVE operations is shown as Algorithms 3.7, 3.8 and 3.9 respectively. We call these implementations *pessimistic*, since each thread keeps a lock a couple of levels above its current position, just to ensure that nothing bad happens (as the situation described in Example 3.2). In the pessimistic implementations, we extend the skiplist data structure with a ghost masked region field r . Then, if sl is a skiplist, we use $sl.r$ to refer to this new field. $sl.r$ keeps the masked region made by all the nodes that form the skiplist sl .

As we said before, boxes denote ghost code added to aid the verification. In this particular case, the ghost variable m_r stores a masked region containing all the nodes and fields currently locked by the running thread. The set operations \cup and $-$ are used for the manipulation of the corresponding sets of pairs. On the other hand, as we said, arrays L and U are used to keep an upper and lower bound of node's addresses that can be potentially modified by the thread performing the INSERT or REMOVE actions. In the case of SEARCH, these values denote the limits within the skiplist where the element to be searched can be located. Similarly, array H describes the upper bound for the levels involved in the operation.

The reader should convince himself that it is possible to show that thanks to the order in which locks are taken and released, a node cannot be removed until it has not been fully inserted. Similarly, it can be proved that if an element of the skiplist is being removed, then any INSERT program trying to insert a node with the same value into the skiplist will be prevented from progressing until the old node has been fully removed.

Algorithm 3.7 Pessimistic implementation of concurrent skiplist operation SEARCH

<pre> 1: procedure SEARCH(<i>SkipList</i> <i>sl</i>, <i>Value</i> <i>v</i>) 2: int <i>i</i> := <i>K</i> - 1 3: <i>Node</i>* <i>pred</i> := <i>sl.head</i> 4: <i>pred.locks</i>[<i>i</i>].lock() 5: <i>Node</i>* <i>curr</i> := <i>pred.next</i>[<i>i</i>] 6: <i>curr.locks</i>[<i>i</i>].lock() 7: <i>Node</i>* <i>cover</i> 8: while $0 \leq i \wedge \text{curr.val} \neq v$ do 9: if $i < K - 1$ then 10: <i>pred.locks</i>[<i>i</i>].lock() 11: <i>pred.locks</i>[<i>i</i> + 1].unlock() 12: if $i < K - 2$ then 13: <i>cover.locks</i>[<i>i</i> + 2].unlock() 14: end if 15: <i>cover</i> := <i>curr</i> 16: <i>curr</i> := <i>pred.next</i>[<i>i</i>] 17: <i>curr.locks</i>[<i>i</i>].lock() 18: end if 19: while <i>curr.val</i> < <i>v</i> do 20: <i>pred.locks</i>[<i>i</i>].unlock() 21: <i>pred</i> := <i>curr</i> 22: <i>curr</i> := <i>pred.next</i>[<i>i</i>] 23: <i>curr.locks</i>[<i>i</i>].lock() 24: end while 25: <i>i</i> := <i>i</i> - 1 26: end while 27: <i>Bool valueIsIn</i> := (<i>curr.val</i> = <i>v</i>) 28: if $i = K - 1$ then 29: <i>curr.locks</i>[<i>i</i>].unlock() 30: <i>pred.locks</i>[<i>i</i>].unlock() 31: else 32: if $i < K - 2$ then 33: <i>cover.locks</i>[<i>i</i> + 2].unlock() 34: end if 35: <i>curr.locks</i>[<i>i</i> + 1].unlock() 36: <i>pred.locks</i>[<i>i</i> + 1].unlock() 37: end if 38: return <i>valueIsIn</i> 39: end procedure </pre>	<table border="1" style="border-collapse: collapse; width: 100%;"> <tr> <td style="padding: 2px;"> $m_r := \emptyset$ $L := L\{me \leftarrow sl.head\}$ $U := U\{me \leftarrow sl.tail\}$ $H := H\{me \leftarrow K - 1\}$ </td> </tr> <tr> <td style="padding: 2px;"> $m_r := m_r \cup \{(pred, i)\}$ </td> </tr> <tr> <td style="padding: 2px;"> $m_r := m_r \cup \{(curr, i)\}$ </td> </tr> <tr> <td style="padding: 2px;"> $m_r := m_r \cup \{(pred, i)\}$ $U := U\{me \leftarrow pred.next[i + 1]\}$ </td> </tr> <tr> <td style="padding: 2px;"> $m_r := m_r - \{(pred, i + 1)\}$ $H := H\{me \leftarrow i\}$ </td> </tr> <tr> <td style="padding: 2px;"> $m_r := m_r \cup \{(curr, i)\}$ </td> </tr> <tr> <td style="padding: 2px;"> $m_r := m_r - \{(pred, i)\}$ $L := L\{me \leftarrow curr\}$ </td> </tr> <tr> <td style="padding: 2px;"> $m_r := m_r \cup \{(curr, i)\}$ </td> </tr> <tr> <td style="padding: 2px;"> $m_r := m_r - \{(curr, i)\}$ $U := U\{me \leftarrow L[me]\}$ </td> </tr> <tr> <td style="padding: 2px;"> $m_r := m_r - \{(pred, i)\}$ $L := L\{me \leftarrow null\}$ </td> </tr> <tr> <td style="padding: 2px;"> $m_r := m_r - \{(cover, i + 2)\}$ </td> </tr> <tr> <td style="padding: 2px;"> $m_r := m_r - \{(curr, i + 1)\}$ $U := U\{me \leftarrow L[me]\}$ </td> </tr> <tr> <td style="padding: 2px;"> $m_r := m_r - \{(pred, i + 1)\}$ $L := L\{me \leftarrow null\}$ </td> </tr> </table>	$m_r := \emptyset$ $L := L\{me \leftarrow sl.head\}$ $U := U\{me \leftarrow sl.tail\}$ $H := H\{me \leftarrow K - 1\}$	$m_r := m_r \cup \{(pred, i)\}$	$m_r := m_r \cup \{(curr, i)\}$	$m_r := m_r \cup \{(pred, i)\}$ $U := U\{me \leftarrow pred.next[i + 1]\}$	$m_r := m_r - \{(pred, i + 1)\}$ $H := H\{me \leftarrow i\}$	$m_r := m_r \cup \{(curr, i)\}$	$m_r := m_r - \{(pred, i)\}$ $L := L\{me \leftarrow curr\}$	$m_r := m_r \cup \{(curr, i)\}$	$m_r := m_r - \{(curr, i)\}$ $U := U\{me \leftarrow L[me]\}$	$m_r := m_r - \{(pred, i)\}$ $L := L\{me \leftarrow null\}$	$m_r := m_r - \{(cover, i + 2)\}$	$m_r := m_r - \{(curr, i + 1)\}$ $U := U\{me \leftarrow L[me]\}$	$m_r := m_r - \{(pred, i + 1)\}$ $L := L\{me \leftarrow null\}$
$m_r := \emptyset$ $L := L\{me \leftarrow sl.head\}$ $U := U\{me \leftarrow sl.tail\}$ $H := H\{me \leftarrow K - 1\}$														
$m_r := m_r \cup \{(pred, i)\}$														
$m_r := m_r \cup \{(curr, i)\}$														
$m_r := m_r \cup \{(pred, i)\}$ $U := U\{me \leftarrow pred.next[i + 1]\}$														
$m_r := m_r - \{(pred, i + 1)\}$ $H := H\{me \leftarrow i\}$														
$m_r := m_r \cup \{(curr, i)\}$														
$m_r := m_r - \{(pred, i)\}$ $L := L\{me \leftarrow curr\}$														
$m_r := m_r \cup \{(curr, i)\}$														
$m_r := m_r - \{(curr, i)\}$ $U := U\{me \leftarrow L[me]\}$														
$m_r := m_r - \{(pred, i)\}$ $L := L\{me \leftarrow null\}$														
$m_r := m_r - \{(cover, i + 2)\}$														
$m_r := m_r - \{(curr, i + 1)\}$ $U := U\{me \leftarrow L[me]\}$														
$m_r := m_r - \{(pred, i + 1)\}$ $L := L\{me \leftarrow null\}$														

Algorithm 3.8 Pessimistic implementation of concurrent skiplist operation INSERT

<pre> 1: procedure INSERT(<i>SkipList</i> <i>sl</i>, <i>Value</i> <i>newval</i>) 2: <i>Vector</i>⟨<i>Node</i>*⟩<i>upd</i>[0..<i>K</i> - 1] 3: <i>lvl</i> := <i>randomLevel</i>(<i>K</i>) 4: <i>Node</i>*<i>pred</i> := <i>sl.head</i> 5: <i>pred.locks</i>[<i>K</i> - 1].<i>lock</i>() 6: <i>Node</i>*<i>curr</i> := <i>pred.next</i>[<i>K</i> - 1] 7: <i>curr.locks</i>[<i>K</i> - 1].<i>lock</i>() 8: <i>Node</i>*<i>cover</i> 9: for <i>i</i> := <i>K</i> - 1 downto 0 do 10: if <i>i</i> < <i>K</i> - 1 then 11: <i>pred.locks</i>[<i>i</i>].<i>lock</i>() 12: if <i>i</i> ≥ <i>lvl</i> then 13: <i>pred.locks</i>[<i>i</i> + 1].<i>unlock</i>() 14: end if 15: if <i>i</i> < <i>K</i> - 2 ∧ <i>i</i> > <i>lvl</i> - 2 then 16: <i>cover.locks</i>[<i>i</i> + 2].<i>unlock</i>() 17: end if 18: <i>cover</i> := <i>curr</i> 19: <i>curr</i> := <i>pred.next</i>[<i>i</i>] 20: <i>curr.locks</i>[<i>i</i>].<i>lock</i>() 21: end if 22: while <i>curr.val</i> < <i>newval</i> do 23: <i>pred.locks</i>[<i>i</i>].<i>unlock</i>() 24: <i>pred</i> := <i>curr</i> 25: <i>curr</i> := <i>pred.next</i>[<i>i</i>] 26: <i>curr.locks</i>[<i>i</i>].<i>lock</i>() 27: end while 28: <i>upd</i>[<i>i</i>] := <i>pred</i> 29: end for 30: if <i>i</i> < <i>K</i> - 2 then 31: <i>cover.locks</i>[<i>i</i> + 2].<i>unlock</i>() 32: end if 33: <i>Bool value WasIn</i> := (<i>curr.val</i> = <i>newval</i>) 34: if <i>value WasIn</i> then 35: for <i>i</i> := <i>lvl</i> to 0 do 36: <i>upd</i>[<i>i</i>].<i>next</i>[<i>i</i>].<i>locks</i>[<i>i</i>].<i>unlock</i>() 37: <i>upd</i>[<i>i</i>].<i>locks</i>[<i>i</i>].<i>unlock</i>() 38: end for 39: else 40: <i>x</i> := <i>CreateNode</i>(<i>lvl</i>, <i>newval</i>) 41: for <i>i</i> := 0 to <i>lvl</i> do 42: <i>x.next</i>[<i>i</i>] := <i>upd</i>[<i>i</i>].<i>next</i>[<i>i</i>] 43: <i>upd</i>[<i>i</i>].<i>next</i>[<i>i</i>] := <i>x</i> 44: <i>x.next</i>[<i>i</i>].<i>locks</i>[<i>i</i>].<i>unlock</i>() 45: <i>upd</i>[<i>i</i>].<i>locks</i>[<i>i</i>].<i>unlock</i>() 46: end for 47: end if 48: return ¬<i>value WasIn</i> 49: end procedure </pre>	<div style="border: 1px solid black; padding: 2px; margin-bottom: 5px;"> $m_r := \emptyset$ $L := L\{me \leftarrow sl.head\}$ $U := U\{me \leftarrow sl.tail\}$ $H := H\{me \leftarrow K - 1\}$ </div> <div style="border: 1px solid black; padding: 2px; margin-bottom: 5px;"> $m_r := m_r \cup \{(pred, K - 1)\}$ </div> <div style="border: 1px solid black; padding: 2px; margin-bottom: 5px;"> $m_r := m_r \cup \{(curr, K - 1)\}$ </div> <div style="border: 1px solid black; padding: 2px; margin-bottom: 5px;"> $m_r := m_r \cup \{(pred, i)\}$ $U := U\{me \leftarrow pred.next[i + 1]\}$ </div> <div style="border: 1px solid black; padding: 2px; margin-bottom: 5px;"> $m_r := m_r - \{(pred, i + 1)\}$ $H := H\{me \leftarrow i\}$ </div> <div style="border: 1px solid black; padding: 2px; margin-bottom: 5px;"> $m_r := m_r \cup \{(curr, i)\}$ </div> <div style="border: 1px solid black; padding: 2px; margin-bottom: 5px;"> $m_r := m_r - \{(pred, i)\}$ if (<i>i</i> ≥ <i>lvl</i>) { $L := L\{me \leftarrow curr\}$ } </div> <div style="border: 1px solid black; padding: 2px; margin-bottom: 5px;"> $m_r := m_r \cup \{(curr, i)\}$ </div> <div style="border: 1px solid black; padding: 2px; margin-bottom: 5px;"> $m_r := m_r - \{(cover, i + 2)\}$ </div> <div style="border: 1px solid black; padding: 2px; margin-bottom: 5px;"> $m_r := m_r - \{(upd[i].next[i], i)\}$ </div> <div style="border: 1px solid black; padding: 2px; margin-bottom: 5px;"> $m_r := m_r - \{(upd[i], i)\}$ </div> <div style="border: 1px solid black; padding: 2px; margin-bottom: 5px;"> $sl.r := sl.r \cup \{(x, i)\}$ </div> <div style="border: 1px solid black; padding: 2px; margin-bottom: 5px;"> $m_r := m_r - \{(x.next[i], i)\}$ </div> <div style="border: 1px solid black; padding: 2px; margin-bottom: 5px;"> $m_r := m_r - \{(upd[i], i)\}$ </div> <div style="border: 1px solid black; padding: 2px; margin-bottom: 5px;"> $L := L\{me \leftarrow null\}$ </div>
---	---

Algorithm 3.9 Pessimistic implementation of concurrent skiplist operation REMOVE

<pre> 1: procedure REMOVE(<i>SkipList</i> <i>sl</i>, <i>Value</i> <i>v</i>) 2: <i>Vector</i> \langle <i>Node</i>\rangle <i>upd</i>[0..<i>K</i> - 1] 3: <i>Node</i>\ast <i>pred</i> := <i>sl.head</i> 4: <i>pred.locks</i>[<i>K</i> - 1].<i>lock</i>() 5: <i>Node</i>\ast <i>curr</i> := <i>pred.next</i>[<i>K</i> - 1] 6: <i>curr.locks</i>[<i>K</i> - 1].<i>lock</i>() 7: <i>cover</i> := <i>sl.tail</i> 8: <i>deleteFrom</i> := <i>K</i> - 1 9: for <i>i</i> := <i>K</i> - 1 downto 0 do 10: if <i>i</i> < <i>K</i> - 1 then 11: <i>pred.locks</i>[<i>i</i>].<i>lock</i>() 12: if <i>pred.next</i>[<i>i</i> + 1].<i>val</i> \neq <i>v</i> then 13: <i>deleteFrom</i> := <i>i</i> 14: <i>pred.locks</i>[<i>i</i> + 1].<i>unlock</i>() 15: end if 16: if <i>i</i> < <i>K</i> - 2 \wedge <i>cover.val</i> \neq <i>v</i> then 17: <i>cover.locks</i>[<i>i</i> + 2].<i>unlock</i>() 18: end if 19: <i>cover</i> := <i>curr</i> 20: <i>curr</i> := <i>pred.next</i>[<i>i</i>] 21: <i>curr.locks</i>[<i>i</i>].<i>lock</i>() 22: end if 23: while <i>curr.val</i> < <i>v</i> do 24: <i>pred.locks</i>[<i>i</i>].<i>unlock</i>() 25: <i>pred</i> := <i>curr</i> 26: <i>curr</i> := <i>pred.next</i>[<i>i</i>] 27: <i>curr.locks</i>[<i>i</i>].<i>lock</i>() 28: end while 29: <i>upd</i>[<i>i</i>] := <i>pred</i> 30: end for 31: if <i>i</i> < <i>K</i> - 2 then 32: <i>cover.locks</i>[<i>i</i> + 2].<i>unlock</i>() 33: end if 34: <i>Bool valueWasIn</i> := (<i>curr.val</i> = <i>v</i>) 35: for <i>i</i> := <i>deleteFrom</i> downto 0 do 36: if <i>upd</i>[<i>i</i>].<i>next</i>[<i>i</i>] = <i>curr</i> \wedge <i>curr.val</i> = <i>v</i> then 37: <i>upd</i>[<i>i</i>].<i>next</i>[<i>i</i>] := <i>curr.next</i>[<i>i</i>] 38: <i>curr.locks</i>[<i>i</i>].<i>unlock</i>() 39: else 40: <i>upd</i>[<i>i</i>].<i>next</i>[<i>i</i>].<i>locks</i>[<i>i</i>].<i>unlock</i>() 41: end if 42: <i>upd</i>[<i>i</i>].<i>locks</i>[<i>i</i>].<i>unlock</i>() 43: end for 44: if <i>valueWasIn</i> then 45: free (<i>curr</i>) 46: end if 47: return <i>valueWasIn</i> 48: end procedure </pre>	<div style="border: 1px solid black; padding: 2px; margin-bottom: 5px;"> $m_r := \emptyset$ $L := L\{me \leftarrow sl.head\}$ $U := U\{me \leftarrow sl.tail\}$ $H := H\{me \leftarrow K - 1\}$ </div> <div style="border: 1px solid black; padding: 2px; margin-bottom: 5px;"> $m_r := m_r \cup \{(pred, K - 1)\}$ </div> <div style="border: 1px solid black; padding: 2px; margin-bottom: 5px;"> $m_r := m_r \cup \{(curr, K - 1)\}$ </div> <div style="border: 1px solid black; padding: 2px; margin-bottom: 5px;"> $m_r := m_r \cup \{(pred, i)\}$ $U := U\{me \leftarrow curr\}$ </div> <div style="border: 1px solid black; padding: 2px; margin-bottom: 5px;"> $m_r := m_r - \{(pred, i + 1)\}$ $H := H\{me \leftarrow i\}$ </div> <div style="border: 1px solid black; padding: 2px; margin-bottom: 5px;"> $m_r := m_r \cup \{(curr, i)\}$ </div> <div style="border: 1px solid black; padding: 2px; margin-bottom: 5px;"> $m_r := m_r - \{(pred, i)\}$ if (<i>cover.val</i> \neq <i>e</i>) $L := L\{me \leftarrow curr\}$ } </div> <div style="border: 1px solid black; padding: 2px; margin-bottom: 5px;"> $m_r := m_r \cup \{(curr, i)\}$ </div> <div style="border: 1px solid black; padding: 2px; margin-bottom: 5px;"> $sl.r := sl.r - \{(curr, i)\}$ </div> <div style="border: 1px solid black; padding: 2px; margin-bottom: 5px;"> $m_r := m_r - \{(curr, i)\}$ </div> <div style="border: 1px solid black; padding: 2px; margin-bottom: 5px;"> $m_r := m_r - \{(upd[i].next[i], i)\}$ </div> <div style="border: 1px solid black; padding: 2px; margin-bottom: 5px;"> $m_r := m_r - \{(upd[i], i)\}$ $L := L\{me \leftarrow null\}$ </div>
--	--

Example 3.3 (Skiplist shape preservation)

Let sl be a pointer to an instance of the class *Skiplist* described before. The following predicate captures whether sl points to a well-formed skiplist of height 4 or less:

$$SkipList_4(h, head, tail) \hat{=} OList(h, head, 0) \wedge \quad (3.2)$$

$$\left(\begin{array}{l} tail.next[0] = null \wedge tail.next[1] = null \\ tail.next[2] = null \wedge tail.next[3] = null \end{array} \right) \wedge \quad (3.3)$$

$$\left(\begin{array}{l} SubList(h, head, tail, 1, head, tail, 0) \wedge \\ SubList(h, head, tail, 2, head, tail, 1) \wedge \\ SubList(h, head, tail, 3, head, tail, 2) \end{array} \right) \quad (3.4)$$

Assuming that $head$ (resp. $tail$) points to the first (resp. last) node of a skiplist, predicate $SkipList_4$ says that in fact between both nodes there is a structure with the shape of a skiplist of height 4. Considering the definition of $SkipList_4$, the predicate $OList$ in (3.2) describes that in heap h , the pointer $head$ points to an ordered linked-lists if we repeatedly follow the pointers at level 0. The predicate (3.3) indicates that all levels are *null* terminated, and (3.4) indicates that each level is in fact a sublist of its nearest lower level. Predicates of this kind also allow to express the effect of program statements via first order transition relations. *

Example 3.4 (Transition relation on skiplist INSERT operation)

Consider the statement at line 43 in program INSERT shown as Algorithm 3.8 on a skiplist of height 4, taken by thread with id j . This transition corresponds to a new node $x^{[j]}$ at level i being connected to the skiplist, as depicted in Fig. 3.7.

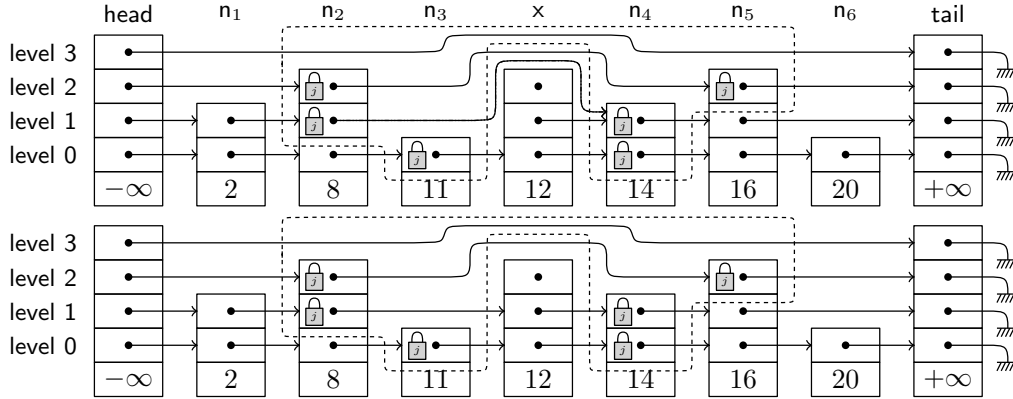


Figure 3.7: Skiplist modification at line 43 before and after inserting value 12 with $i = 1$

For the sake of simplicity, from this point on we use x to represent $x^{[j]}$. If the memory layout from pointer sl is that of a skiplist before the statement at line 43 is executed, then it is also a skiplist after the execution:

$$SkipList_4(h, sl.head, sl.tail) \wedge \varphi_{aux} \wedge \rho_{43}^{[j]}(V, V') \rightarrow SkipList_4(h', sl'.head, sl'.tail)$$

The effect of the statement at line 43 is represented by the first-order transition relation $\rho_{43}^{[j]}$. To ensure this property, i is required to be a valid level, and the key of the nodes that will be pointing to x must be lower than the key of node x . Moreover, the masked region of locked nodes remains unchanged. Predicate φ_{aux} contains support invariants. For simplicity, we use $prev$ for $upd^{[j]}[i]$. Then, the full verification condition is $SkipList_4(h, sl.head, sl.tail) \wedge \phi \rightarrow SkipList_4(h', sl'.head, sl'.tail)$, where ϕ is:

$$\left(\begin{array}{l} x.key = newval \quad \wedge \\ prev.key < newval \quad \wedge \\ x.next[i].key > newval \quad \wedge \\ prev.next[i] = x.next[i] \quad \wedge \\ (x, i) \notin sl.r \wedge 0 \leq i \leq 3 \end{array} \right) \wedge \left(\begin{array}{l} at_{43}^{[j]} \quad \wedge \\ prev'.next[i] = x \quad \wedge \\ at_{44}^{[j]} \quad \wedge \\ h' = h \wedge sl = sl' \quad \wedge \\ x' = x \quad \dots \end{array} \right) \quad *$$

Note that in the previous verification condition we do not explicitly indicate in all cases that the rest of the nodes, memory and local variables of other threads remain unchanged. This preservation predicates are generated as part of a verification condition based on those elements of the formula that are not modified by the statement. A full example of a verification condition is given in Chapter 5, where the decision procedure for concurrent skiplists is described. The examples we have seen should be enough to illustrate that to be able to automatically prove VCs for the verification of skiplist manipulating algorithms, we require a theory that allows to reason about heaps, addresses, nodes, masked regions, ordered lists and sublists.

4

TLL3: A Decision Procedure for Concurrent Lock-Coupling Lists

The automatic check of the proof represented by a verification diagram requires decision procedures to verify the generated verification conditions. These decision procedures must deal with formulas containing terms belonging to different theories. In particular, for concurrent lists the decision procedure must reason about pointer data structures with a list layout, regions and locks. To obtain a suitable decision procedure, we extend the Theory of Linked Lists (TLL) [RZ06a], a decidable theory including reachability of list-like structures. However, this theory lacks the expressivity to describe locked lists of cells, a fundamental component in our proofs. From the extension of TLL we get TLL3 [SS10]. In this chapter we focus on the definition, description and analysis of TLL3 as well as a decision procedure for this theory.

We begin with a brief description of the basic notation and concepts. A signature Σ is a triple (S, F, P) where S is a set of sorts, F is a set of function symbols and P is a set of predicate symbols constructed using the sorts in S . If $\Sigma_1 = (S_1, F_1, P_1)$ and $\Sigma_2 = (S_2, F_2, P_2)$ are two signatures, we define their union $\Sigma_1 \cup \Sigma_2 = (S_1 \cup S_2, F_1 \cup F_2, P_1 \cup P_2)$. Similarly we say that $\Sigma_1 \subseteq \Sigma_2$ when $S_1 \subseteq S_2$, $F_1 \subseteq F_2$ and $P_1 \subseteq P_2$.

Given a signature Σ , we assume the standard notions of Σ -term, Σ -literal and Σ -formula. A literal is flat if it is of the form $x = y$, $x \neq y$, $x = f(y_1, \dots, y_n)$, $p(y_1, \dots, y_n)$ or $\neg p(y_1, \dots, y_n)$, where x, y, y_1, \dots, y_n are variables, f is a function symbol and p is a predicate symbol. If $t(\varphi)$ is a term (resp. formula), then we denote with $V_\sigma(t)$ (resp. $V_\sigma(\varphi)$) the set of variables of sort σ occurring in t (resp. φ).

We assume the usual notion of Σ -interpretation over a set V of variables as a map which assigns a value to each symbol in Σ and V . Let \mathcal{A} be a Σ -interpretation over V . Then, for each $s \in S$, \mathcal{A}_s is a set of elements called the domain of s ; for each symbol $f : s_1 \times \dots \times s_n \rightarrow s$ in F , \mathcal{A}_f is a function that goes from $\mathcal{A}_{s_1} \times \dots \times \mathcal{A}_{s_n}$ to elements in \mathcal{A}_s and for each symbol $p : s_1 \times \dots \times s_n$ in P , \mathcal{A}_p is a relation over elements of $\mathcal{A}_{s_1} \times \dots \times \mathcal{A}_{s_n}$.

A Σ -structure is a Σ -interpretation over an empty set of variables. A Σ -formula over a set V of variables is satisfiable whenever it is true in some Σ -interpretation over V . Let Ω be a signature, \mathcal{A} a Ω -interpretation over a set V of variables, $\Sigma \subseteq \Omega$ and $U \subseteq V$. $\mathcal{A}^{\Sigma, U}$ denotes the interpretation obtained from \mathcal{A} restricting it to interpret only the symbols in Σ and the variables in U . We use \mathcal{A}^Σ to denote $\mathcal{A}^{\Sigma, \emptyset}$. A Σ -theory is a pair (Σ, \mathbf{A}) where Σ is a signature and \mathbf{A} is a class of Σ -structures. Given a theory $T = (\Sigma, \mathbf{A})$, a T -interpretation is a Σ -interpretation \mathcal{A} such that $\mathcal{A}^\Sigma \in \mathbf{A}$. Given a theory T , a Σ -formula φ over a set of variables V is T -satisfiable if it is true on a T -interpretation over V .

Signature:	Σ_{cell}
Sorts:	cell, elem, addr, thid
Functions:	<i>error</i> : cell <i>mkcell</i> : elem \times addr \times thid \rightarrow cell <i>_.data</i> : cell \rightarrow elem <i>_.next</i> : cell \rightarrow addr <i>_.lockid</i> : cell \rightarrow thid <i>_.lock</i> : cell \rightarrow thid \rightarrow cell <i>_.unlock</i> : cell \rightarrow cell
Signature:	Σ_{mem}
Sorts:	mem, addr, cell
Functions:	<i>null</i> : addr <i>[_]</i> : mem \times addr \rightarrow cell <i>upd</i> : mem \times addr \times cell \rightarrow mem
Signature:	Σ_{reach}
Sorts:	mem, addr, path
Functions:	ϵ : path <i>[_]</i> : addr \rightarrow path
Predicates:	<i>append</i> : path \times path \times path <i>reach</i> : mem \times addr \times addr \times path
Signature:	Σ_{set}
Sorts:	addr, set
Functions:	\emptyset : set $\{ _ \}$: addr \rightarrow set \cup, \cap, \setminus : set \times set \rightarrow set
Predicates:	\in : addr \times set \subseteq : set \times set
Signature:	Σ_{settid}
Sorts:	thid, settid
Functions:	\emptyset_{T} : settid $\{ _ \}_{\text{T}}$: thid \rightarrow settid $\cup_{\text{T}}, \cap_{\text{T}}, \setminus_{\text{T}}$: settid \times settid \rightarrow settid
Predicates:	\in_{T} : thid \times settid \subseteq_{T} : settid \times settid
Signature:	Σ_{bridge}
Sorts:	mem, addr, set, path
Functions:	<i>path2set</i> : path \rightarrow set <i>addr2set</i> : mem \times addr \rightarrow set <i>getp</i> : mem \times addr \times addr \rightarrow path <i>fstlock</i> : mem \times path \rightarrow addr

Figure 4.1: The signature of the TLL3 theory

signature:	Σ_{cell}
interpretation:	<ul style="list-style-type: none"> • $mkcell^A(e, a, k) = \langle e, a, k \rangle$ • $\langle e, a, t \rangle.data^A = e$ • $\langle e, a, t \rangle.next^A = a$ • $\langle e, a, t \rangle.lockid^A = t$ • $\langle e, a, t \rangle.lock^A(t') = \langle e, a, t' \rangle$ • $\langle e, a, t \rangle.unlock^A = \langle e, a, \emptyset \rangle$ • $error^A.next^A = null^A$ <p>for each $e \in \mathcal{A}_{\text{elem}}$, $a \in \mathcal{A}_{\text{addr}}$, $k \in \mathcal{A}_{\text{thid}}$ and $t, t' \in \mathcal{A}_{\text{thid}}$</p>
Signature:	Σ_{mem}
Interpretation:	<ul style="list-style-type: none"> • $m[a]^A = m(a)$ • $upd^A(m, a, c) = m_{a \rightarrow c}$ • $m^A(null^A) = error^A$ <p>for each $m \in \mathcal{A}_{\text{mem}}$, $a \in \mathcal{A}_{\text{addr}}$ and $c \in \mathcal{A}_{\text{cell}}$</p>
Signature:	Σ_{reach}
Interpretation:	<ul style="list-style-type: none"> • ϵ^A is the empty sequence • $[i]^A$ is the sequence containing $i \in \mathcal{A}_{\text{addr}}$ as the only element • $([i_1, \dots, i_n], [j_1, \dots, j_m], [i_1, \dots, i_n, j_1, \dots, j_m]) \in \text{append}^A$ iff i_k and j_l are all distinct • $(m, i, j, p) \in \text{reach}^A$ iff <ul style="list-style-type: none"> – $i = j$ and $p = \epsilon$, or – there exist addresses $i_1, \dots, i_n \in \mathcal{A}_{\text{addr}}$ such that: <ul style="list-style-type: none"> (a) $p = [i_1, \dots, i_n]$ (b) $i_1 = i$ (c) $m(i_r).next^A = i_{r+1}$, for $1 \leq r < n$ (d) $m(i_n).next^A = j$
Signature:	Σ_{set}
Interpretation:	<ul style="list-style-type: none"> • The symbols \emptyset, $\{_ \}$, \cup, \cap, \setminus, \in and \subseteq are interpreted according to their standard interpretation over sets of addresses.
Signature:	Σ_{settid}
Interpretation:	<ul style="list-style-type: none"> • The symbols \emptyset_{T}, $\{_ \}_{\text{T}}$, \cup_{T}, \cap_{T}, \setminus_{T}, \in_{T} and \subseteq_{T} are interpreted according to their standard interpretation over sets of thread identifiers.
Signature:	Σ_{bridge}
Interpretation:	<ul style="list-style-type: none"> • $addr2set^A(m, i) = \{j \in \mathcal{A}_{\text{addr}} \mid \exists p \in \mathcal{A}_{\text{path}} \text{ s.t. } (m, i, j, p) \in \text{reach}\}$ • $path2set^A(p) = \{i_1, \dots, i_n\}$ for $p = [i_1, \dots, i_n] \in \mathcal{A}_{\text{path}}$ • $getp^A(m, i, j) = \begin{cases} p & \text{if } (m, i, j, p) \in \text{reach}^A \\ \epsilon & \text{otherwise} \end{cases}$ <p>for each $m \in \mathcal{A}_{\text{mem}}$, $p \in \mathcal{A}_{\text{path}}$ and $i, j \in \mathcal{A}_{\text{addr}}$</p> <ul style="list-style-type: none"> • $fstlock^A(m, [a_1, \dots, a_n]) = \begin{cases} a_k & \text{if there is } 1 \leq k \leq n \text{ such that} \\ & \text{for all } 1 \leq j < k, m[a_j].lockid = \emptyset \\ & \text{and } m[a_k].lockid \neq \emptyset \\ null & \text{otherwise} \end{cases}$ <p>for each $m \in \mathcal{A}_{\text{mem}}$ and $a_1, \dots, a_n \in \mathcal{A}_{\text{addr}}$</p>

Figure 4.2: Characterization of a TLL3-interpretation \mathcal{A}

Formally, the theory of linked lists is defined as $\text{TLL} = (\Sigma_{\text{TLL}}, \mathbf{TLL})$, where

$$\Sigma_{\text{TLL}} := \Sigma_{\text{cell}} \cup \Sigma_{\text{mem}} \cup \Sigma_{\text{reach}} \cup \Sigma_{\text{set}} \cup \Sigma_{\text{bridge}}$$

The sorts, functions and predicates belonging to these signatures are described as part of Fig. 4.1, while \mathbf{TLL} is the class of Σ_{TLL} -structures satisfying the conditions shown in Fig. 4.2. In fact, these figures corresponds to the signatures and interpretations for TLL3 while TLL is contained into TLL3, as we will see next.

We extend TLL into the theory for concurrent single linked lists with locks $\text{TLL3} := (\Sigma_{\text{TLL3}}, \mathbf{TLL3})$, where

$$\Sigma_{\text{TLL3}} = \Sigma_{\text{TLL}} \cup \Sigma_{\text{settid}} \cup \{\text{lockid}, \text{lock}, \text{unlock}, \text{fstlock}\}$$

Each sort σ in a Σ_{TLL3} -structure is mapped to a non-empty set \mathcal{A}_σ such that:

- (a) $\mathcal{A}_{\text{cell}} = \mathcal{A}_{\text{elem}} \times \mathcal{A}_{\text{addr}} \times \mathcal{A}_{\text{thid}}$
- (b) $\mathcal{A}_{\text{mem}} = \mathcal{A}_{\text{cell}}^{\mathcal{A}_{\text{addr}}}$
- (c) $\mathcal{A}_{\text{path}}$ is the set of all finite sequences of (pairwise) distinct elements of $\mathcal{A}_{\text{addr}}$
- (d) \mathcal{A}_{set} is the power-set of $\mathcal{A}_{\text{addr}}$
- (e) $\mathcal{A}_{\text{settid}}$ is the power-set of $\mathcal{A}_{\text{thid}}$

The sorts, functions and predicates of TLL3 correspond to the signatures shown in Fig. 4.1 and the interpretation of each function and predicate is depicted in Fig. 4.2. Informally, Σ_{cell} models *cells*, structures containing an element (data), an address (pointer) and a lock owner, which represents a node in a linked list. Σ_{mem} models the memory. Σ_{reach} models finite sequences of non-repeating addresses, to represent paths. Σ_{set} models sets of addresses. Finally, Σ_{bridge} is a *bridge theory* containing auxiliary functions that map paths of addresses to set of addresses or let us obtain the set of addresses reachable from a given address following a chain of *next* fields. The sort *thid* contains thread identifiers. The sorts *addr*, *elem* and *thid* are uninterpreted, except that $\circ : \text{thid}$ is different from all others thread ids. Otherwise, $\Sigma_{\text{addr}} = (\text{addr}, \emptyset, \emptyset)$, $\Sigma_{\text{elem}} = (\text{elem}, \emptyset, \emptyset)$ and $\Sigma_{\text{thid}} = (\text{thid}, \emptyset, \emptyset)$.

Example 4.1

Consider the list shown in Fig. 3.1 in Section 3.1. For that particular list and assuming that nodes are locked by let's say T_1 , we can construct a model \mathcal{A} such that for instance:

$$\begin{aligned} \mathcal{A}_{\text{addr}} &= \{0x00, 0x01, 0x02, 0x03, 0x04, 0x05, 0x06, 0x07\} \\ \mathcal{A}_{\text{ord}} = \mathcal{A}_{\text{elem}} &= \{-\infty, 3, 5, 9, 10, 17, +\infty\} \\ \mathcal{A}_{\text{thid}} &= \{T_1, T_2, \circ\} \\ \mathcal{A}_{\text{cell}} &= \mathcal{A}_{\text{elem}} \times \mathcal{A}_{\text{ord}} \times \mathcal{A}_{\text{addr}} \times \mathcal{A}_{\text{thid}} \\ \mathcal{A}_{\text{mem}} &= \{m : \mathcal{A}_{\text{addr}} \rightarrow \mathcal{A}_{\text{cell}}\} \end{aligned}$$

where

$$\begin{aligned} \text{null}^{\mathcal{A}} &= 0x00 \\ \text{error}^{\mathcal{A}} &= \langle -\infty, -\infty, \text{null}, \circ \rangle \end{aligned}$$

$$\begin{aligned} m(0x00) &= \langle -\infty, -\infty, \text{null}, \circ \rangle \\ m(0x01) &= \langle -\infty, +\infty, \text{null}, \circ \rangle \\ m(0x02) &= \langle 3, 3, 0x03, \circ \rangle \\ m(0x03) &= \langle 5, 5, 0x04, T_1 \rangle \\ m(0x04) &= \langle 9, 9, 0x05, T_1 \rangle \\ m(0x05) &= \langle 10, 10, 0x06, \circ \rangle \\ m(0x06) &= \langle 17, 17, 0x07, \circ \rangle \\ m(0x07) &= \langle +\infty, +\infty, 0x00, \circ \rangle \end{aligned}$$

*

We are interested in analyzing the satisfiability of quantifier-free first order formulas. Hence, if φ is a formula, we first write it into its disjunctive normal form, let's say $\varphi_1 \vee \dots \vee \varphi_n$. And then, we just verify the satisfiability of any of the φ_i , where each φ_i is a conjunction of TLL3 literals. We now classify the TLL3 literals into normalized and non-normalized ones. Non-normalized literals have the property that they can be written in terms of normalized ones. We now define the set of normalized TLL3-literals.

Definition 4.1 (TLL3-normalized literals).

A TLL3-literal is normalized if it is a flat literal of the form:

$$\begin{array}{lll}
e_1 \neq e_2 & a_1 \neq a_2 & \\
a = \text{null} & c = \text{error} & \\
c = \text{mkcell}(e, a, t) & c = \text{rd}(m, a) & m_2 = \text{upd}(m_1, a, c) \\
s = \{a\} & s_1 = s_2 \cup s_3 & s_1 = s_2 \setminus s_3 \\
p_1 \neq p_2 & p = [a] & p_1 = \text{rev}(p_2) \\
s = \text{path2set}(p) & \text{append}(p_1, p_2, p_3) & \neg \text{append}(p_1, p_2, p_3) \\
s = \text{addr2set}(m, a) & p = \text{getp}(m, a_1, a_2) & \\
t_1 \neq t_2 & a = \text{fstlock}(m, p) &
\end{array}$$

where e, e_1 and e_2 are elem-variables, a, a_1 and a_2 are addr-variables, c is a cell-variable, m, m_1 and m_2 are mem-variables, p, p_1, p_2 and p_3 are path-variables, and t, t_1 and t_2 are thid-variables. †

The remaining literals that are part of TLL3 but not detailed in the previous definition, can be written in terms of normalized literals considering the following equivalences:

$$\begin{array}{ll}
t = c.\text{lockid} & \leftrightarrow (\exists_{\text{elem}} e \exists_{\text{addr}} a) [c = \text{mkcell}(e, a, t)] \\
c_1 = c_2.\text{lock}(t) & \leftrightarrow c_2.\text{data} = c_1.\text{data} \wedge c_2.\text{next} = c_1.\text{next} \wedge t = c_1.\text{lockid} \\
c_1 = c_2.\text{unlock} & \leftrightarrow c_2.\text{data} = c_1.\text{data} \wedge c_2.\text{next} = c_1.\text{next} \wedge \emptyset = c_1.\text{lockid} \\
c_1 \neq_{\text{cell}} c_2 & \leftrightarrow c_1.\text{data} \neq c_2.\text{data} \vee c_1.\text{next} \neq c_2.\text{next} \vee c_1.\text{lockid} \neq c_2.\text{lockid} \\
m_1 \neq_{\text{mem}} m_2 & \leftrightarrow (\exists_{\text{addr}} a) [\text{rd}(m_1, a) \neq \text{rd}(m_2, a)] \\
s_1 \neq s_2 & \leftrightarrow (\exists_{\text{addr}} a) [a \in (s_1 \setminus s_2) \cup (s_2 \setminus s_1)] \\
s = \emptyset & \leftrightarrow s = s \setminus s \\
s_3 = s_1 \cap s_2 & \leftrightarrow s_3 = (s_1 \cup s_2) \setminus ((s_1 \setminus s_2) \cup (s_2 \setminus s_1)) \\
a \in s & \leftrightarrow \{a\} \subseteq s \\
s_1 \subseteq s_2 & \leftrightarrow s_2 = s_1 \cup s_2 \\
p = \epsilon & \leftrightarrow \text{append}(p, p, p) \\
\text{reach}(m, a_1, a_2, p) & \leftrightarrow a_2 \in \text{addr2set}(m, a_1) \wedge p = \text{getp}(m, a_1, a_2)
\end{array}$$

this means that we can rewrite such literals using the following table:

Flat:	$t = c.\text{lockid}$
Normalized:	$c = \text{mkcell}(e, a, t)$
Proviso:	e and a are fresh variables.
Flat:	$c_1 = c_2.\text{lock}(t)$
Normalized:	$c_1 = \text{mkcell}(e, a, t) \wedge c_2 = \text{mkcell}(e, a, t')$
Proviso:	e, a and t' are fresh variables.
Flat:	$c_1 = c_2.\text{lock}(t)$
Normalized:	$c_1 = \text{mkcell}(e, a, t) \wedge c_2 = \text{mkcell}(e, a, t')$
Proviso:	e, a and t' are fresh variables.
Flat:	$c_1 = c_2.\text{unlock}$
Normalized:	$c_1 = \text{mkcell}(e, a, \emptyset) \wedge c_2 = \text{mkcell}(e, a, t')$
Proviso:	e, a and t' are fresh variables.

Flat:	$c_1 \neq c_2$
Normalized:	$c_1.data \neq c_2.data \vee c_1.next \neq c_2.next \vee c_1.lockid \neq c_2.lockid$
Proviso:	—
Flat:	$m_1 \neq m_2$
Normalized:	$m[a] \neq m[a]$
Proviso:	a is fresh.
Flat:	$s_1 \neq s_2$
Normalized:	$s_{12} = s_1 \setminus s_2 \wedge s_{21} = s_2 \setminus s_1 \wedge s_3 = s_{12} \cup s_{21} \wedge s = s_3 \cup \{a\} \wedge \{a\} \subseteq s$
Proviso:	s_{12}, s_{21}, s_3, s and a are fresh.
Flat:	$s = \emptyset$
Normalized:	$s = s \setminus s$
Proviso:	—
Flat:	$s_3 = s_1 \cap s_2$
Normalized:	$s_{12} = s_1 \setminus s_2 \wedge s_{21} = s_2 \setminus s_1 \wedge s_{u_1} = s_1 \cup s_2 \wedge s_{u_2} = s_{12} \cup s_{21} \wedge s_3 = s_{u_1} \setminus s_{u_2}$
Proviso:	s_{12}, s_{21}, s_{u_1} and s_{u_2} are fresh.
Flat:	$a \in s$
Normalized:	$s = \{a\} \cup s$
Proviso:	—
Flat:	$s_1 \subseteq s_2$
Normalized:	$s_2 = s_1 \cup s_2$
Proviso:	—
Flat:	$p = \epsilon$
Normalized:	$append(p, p, p)$
Proviso:	—
Flat:	$reach(m, a_1, a_2, p)$
Normalized:	$a_2 \in addr2set(m, a_1) \wedge p = getp(m, a_1, a_2)$
Proviso:	—

4.1 Decidability of TLL3

TLL [RZ06a] enjoys the finite model property. We now show that TLL3 also has the finite model property with respect to domains `elem`, `addr` and `thid`. Hence, TLL3 is decidable because one can enumerate Σ_{TLL3} -structures up to a certain cardinality. Notice that a bound on the domain of these sorts it is enough to get finite interpretations for the remaining sorts (`cell`, `mem`, `path`, `set` and `settid`) as the elements in the domains of these latter sorts are constructed using the elements in the domains of `elem`, `addr` and `thid`.

Definition 4.2 (Finite Model Property).

Let $\Sigma = (S, F, P)$ be a signature, $S_0 \subseteq S$ be a set of sorts, and T be a Σ -theory. T has the finite model property with respect to S_0 if for every T -satisfiable quantifier-free Σ -formula φ there exists a T -interpretation \mathcal{A} satisfying φ such that for each sort $\sigma \in S_0$, \mathcal{A}_σ is finite. †

Lemma 4.1:

Deciding the TLL3-satisfiability of a quantifier-free TLL3-formula is equivalent to verifying the TLL3-satisfiability of the normalized TLL3-literals. ♠

Proof The proof is by cases on the shape of all possible TLL3-literals. Let φ be a quantifier-free TLL3 formula. As we are interested in the satisfiability problem, φ can be reduced to its disjunctive normal form $\varphi_1 \vee \dots \vee \varphi_n$, leading us to the problem of checking whether some φ_i is satisfiable. Each φ_i is a

conjunction of TLL3 normalized and non-normalized literals. Then, it just remains to see that if a model satisfies a non-normalized literal, then it also satisfies its representation made by normalized literals.

For instance, let's consider the non-normalized literal $c_1 = c_2.lock(t)$. According to the table we presented before, this literal can be written as a conjunction of normalized literals of the form $c_1 = mkcell(e, a, t) \wedge c_2 = mkcell(e, a, t')$ where e, a and t' are fresh variables. Let \mathcal{A} be a model satisfying $c_1 = c_2.lock(t)$. Then, we show that $\mathcal{A} \models c_1 = c_2.lock(t) \leftrightarrow \mathcal{A} \models c_1 = mkcell(e, a, t) \wedge c_2 = mkcell(e, a, t')$

$$\begin{aligned}
\mathcal{A} \models c_1 = c_2.lock(t) & \leftrightarrow \\
c_1^{\mathcal{A}} = c_2^{\mathcal{A}}.lock^{\mathcal{A}}(t^{\mathcal{A}}) & \leftrightarrow \\
c_1^{\mathcal{A}} = c_2^{\mathcal{A}}.lock^{\mathcal{A}}(t^{\mathcal{A}}) \wedge c_2^{\mathcal{A}} = \langle e^{\mathcal{A}}, a^{\mathcal{A}}, t'^{\mathcal{A}} \rangle & \leftrightarrow \\
c_1^{\mathcal{A}} = \langle e^{\mathcal{A}}, a^{\mathcal{A}}, t^{\mathcal{A}} \rangle \wedge c_2^{\mathcal{A}} = \langle e^{\mathcal{A}}, a^{\mathcal{A}}, t'^{\mathcal{A}} \rangle & \leftrightarrow \\
c_1^{\mathcal{A}} = mkcell^{\mathcal{A}}(e^{\mathcal{A}}, a^{\mathcal{A}}, t^{\mathcal{A}}) \wedge c_2^{\mathcal{A}} = mkcell^{\mathcal{A}}(e^{\mathcal{A}}, a^{\mathcal{A}}, t'^{\mathcal{A}}) & \leftrightarrow \\
\mathcal{A} \models c_1 = mkcell(e, a, t) \wedge c_2 = mkcell(e, a, t') &
\end{aligned}$$

The remaining cases can be proved in a similar way. \square

Consider an arbitrary TLL3-interpretation \mathcal{A} satisfying a conjunction of normalized TLL3-literals Γ . We show that if there are sets \mathcal{A}_{elem} , \mathcal{A}_{addr} and \mathcal{A}_{thid} (these sets are the domains interpreted by \mathcal{A} for the sorts `elem`, `addr` and `thid` respectively) then there are finite sets \mathcal{B}_{elem} , \mathcal{B}_{addr} and \mathcal{B}_{thid} with bounded cardinalities (the bound depending on Γ). \mathcal{B}_{elem} , \mathcal{B}_{addr} and \mathcal{B}_{thid} can in turn be used to obtain a finite interpretation \mathcal{B} satisfying Γ .

Before proceeding with the proof that TLL3 enjoys the finite model property, we define some auxiliary functions. We start by defining the function *first*. Let $X \subseteq \tilde{X}$, $m : \tilde{X} \rightarrow Z \times \tilde{X} \times Y$ and $a \in X$. The function *first*(m, a, X) is defined by

$$first(m, a, X) = \begin{cases} null & \text{if } (\forall r \geq 1) [m^r(a).next \notin X] \\ m^s(a).next & \text{if } (\exists s \geq 1) \left[m^s(a).next \in X \wedge \right. \\ & \left. (\forall r \geq 1) (r < s \rightarrow m^r(a).next \notin X) \right] \end{cases}$$

where $m^1(a).next$ stands for $m(a).next$ and $m^{n+1}(a).next$ for $m(m^n(a).next).next$ when $n > 1$.

Basically, given a set of addresses X , function *first* chooses the next address in X that can be reached from a given address following repeatedly the *next* pointer. It is easy to see, for example, that if $m(a).next \in X$ then $first(m, a, X) = m(a).next$. We will later filter out unnecessary intermediate nodes and use *first* to bypass properly the removed nodes, preserving the important connectivity properties.

Lemma 4.2:

Let $X \subseteq \tilde{X}$, $m_1, m_2 : \tilde{X} \rightarrow Z \times \tilde{X} \times Y$, $i, j \in X$, $c \in Z \times X \times Y$ and $i \neq j$. Then:

- (a) If $m_1(i).next \in X$ then $first(m_1, i, X) = m_1(i).next$
- (b) If $m_1 = upd(m_2, i, c)$ then $first(m_1, j, X) = first(m_2, j, X)$ \spadesuit

Proof (a) immediate.

- (b) Let $m_1 = upd(m_2, i, c)$ and assume first that $m_1^r(j).next \notin X$, for all $r \geq 1$. By induction it can be shown that $m_1^r(j) = m_2^r(j)$, for each $r \geq 1$. It follows that $first(m_1, j, X) = j$ and $first(m_2, j, X) = j$.

If instead $m_1^s(j).next \in X$, for some $s \geq 1$, assume without loss of generality that $first(m_1, j, X) = m_1^s(j).next$. By induction, it is possible to show that $m_1^r(j) = m_2^r(j)$, for each $1 \leq r \leq s$. It follows that $first(m_1, j, X) = m_1^s(j).next = m_2^s(j).next = first(m_2, j, X)$. \square

We also define the *compress* function which, given a path p and a set X of addresses, returns the path obtained from p by removing all the addresses that do not belong to X .

$$\text{compress}([i_1, \dots, i_n], X) = \begin{cases} \epsilon & \text{if } n = 0 \\ [i_1] \circ \text{compress}([i_2, \dots, i_n], X) & \text{if } n > 0 \text{ and } i_1 \in X \\ \text{compress}([i_2, \dots, i_n], X) & \text{otherwise} \end{cases}$$

We conclude by defining the function *diseq* [RZ06a] that outputs a set of addresses accountable for disequality of two given paths:

$$\text{diseq}([i_1, \dots, i_n], [j_1, \dots, j_m]) = \begin{cases} \emptyset & \text{if } n = m = 0 \\ \{i_1\} & \text{if } n > 0 \text{ and } m = 0 \\ \{j_1\} & \text{if } n = 0 \text{ and } m > 0 \\ \{i_1, j_1\} & \text{if } n, m > 0 \text{ and } i_1 \neq j_1 \\ \text{diseq}([i_2, \dots, i_n], [j_2, \dots, j_m]) & \text{otherwise} \end{cases}$$

and function *common* [RZ06a] that outputs an element common to two paths (an element that witnesses that $\text{path2set}(p) \cap \text{path2set}(q) \neq \emptyset$):

$$\text{common}([i_1, \dots, i_n], p) = \begin{cases} \emptyset & \text{if } n = 0 \\ \{i_1\} & \text{if } n > 0 \text{ and } i_1 \in \text{path2set}(p) \\ \text{common}([i_2, \dots, i_n], p) & \text{otherwise} \end{cases}$$

Lemma 4.3 (Finite Model Property):

Let Γ be a conjunction of normalized TLL3-literals. Let $\bar{e} = |V_{\text{elem}}(\Gamma)|$, $\bar{a} = |V_{\text{addr}}(\Gamma)|$, $\bar{m} = |V_{\text{mem}}(\Gamma)|$, $\bar{p} = |V_{\text{path}}(\Gamma)|$ and $\bar{t} = |V_{\text{thid}}(\Gamma)|$. Then the following are equivalent:

1. Γ is TLL3-satisfiable;
2. Γ is true in a TLL3 interpretation \mathcal{B} such that

$$\begin{aligned} |\mathcal{B}_{\text{addr}}| &\leq \bar{a} + 1 + \bar{m} \bar{a} + \bar{p}^2 + \bar{p}^3 \\ |\mathcal{B}_{\text{thid}}| &\leq \bar{t} + \bar{m} |\mathcal{B}_{\text{addr}}| + 1 \\ |\mathcal{B}_{\text{elem}}| &\leq \bar{e} + \bar{m} |\mathcal{B}_{\text{addr}}| \end{aligned} \quad \spadesuit$$

Proof (2 \rightarrow 1). Immediate.

(1 \rightarrow 2). As TLL3 is an extension of TLL, some normalized literals are shared between both theories. Hence, here we limit ourselves to prove the implication for the new TLL3-literals only. The proof for normalized literals shared between TLL and TLL3 can be seen in [RZ06b].

Let now \mathcal{A} be a TLL3-interpretation satisfying Γ . We will use \mathcal{A} to construct a finite TLL3-interpretation \mathcal{B} satisfying Γ .

$$\begin{aligned}
\mathcal{B}_{\text{addr}} = & V_{\text{addr}}^A \cup \{null^A\} && \cup \\
& \{m^A(v^A).next^A \mid m \in V_{\text{mem}} \text{ and } v \in V_{\text{addr}}\} && \cup \\
& \{v \in \text{diseq}(p^A, q^A) \mid \text{the literal } p \neq q \text{ is in } \Gamma\} && \cup \\
& \{v \in \text{common}(p_1^A, p_2^A) \mid \text{the literal } \neg \text{append}(p_1, p_2, p_3) \text{ is in } \Gamma \text{ and} \\
& \quad \text{path2set}^A(p_1^A) \cap \text{path2set}^A(p_2^A) \neq \emptyset\} && \cup \\
& \{v \in \text{common}(p_1^A \circ p_2^A, p_3^A) \mid \text{the literal } \neg \text{append}(p_1, p_2, p_3) \text{ is in } \Gamma \text{ and} \\
& \quad \text{path2set}^A(p_1^A) \cap \text{path2set}^A(p_2^A) = \emptyset\}
\end{aligned}$$

$$\mathcal{B}_{\text{thid}} = V_{\text{thid}}^A \cup \{\emptyset\} \cup \{m^A(v).lockid^A \mid m \in V_{\text{mem}} \text{ and } v \in X\}$$

$$\mathcal{B}_{\text{elem}} = V_{\text{elem}}^A \cup \{m^A(v).data^A \mid m \in V_{\text{mem}} \text{ and } v \in X\}$$

and let

$$\begin{aligned}
error^{\mathcal{B}} &= error^A \\
null^{\mathcal{B}} &= null^A \\
e^{\mathcal{B}} &= e^A && \text{for each } e \in V_{\text{elem}} \\
a^{\mathcal{B}} &= a^A && \text{for each } a \in V_{\text{addr}} \\
c^{\mathcal{B}} &= c^A && \text{for each } c \in V_{\text{cell}} \\
t^{\mathcal{B}} &= t^A && \text{for each } t \in V_{\text{thid}} \\
m^{\mathcal{B}}(v) &= (m^A(v).data^A, first(m^A, v, \mathcal{B}_{\text{addr}}), m^A(v).lockid^A) && \text{for each } m \in V_{\text{mem}}, v \in \mathcal{B}_{\text{addr}} \\
s^{\mathcal{B}} &= s^A \cap \mathcal{B}_{\text{addr}} && \text{for each } s \in V_{\text{set}} \\
g^{\mathcal{B}} &= g^A \cap \mathcal{B}_{\text{thid}} && \text{for each } g \in V_{\text{settid}} \\
p^{\mathcal{B}} &= \text{compress}(p^A, \mathcal{B}_{\text{addr}}) && \text{for each } p \in V_{\text{path}}
\end{aligned}$$

Clearly, by construction $\mathcal{B}_{\text{addr}}$, $\mathcal{B}_{\text{thid}}$ and $\mathcal{B}_{\text{elem}}$ satisfy the given cardinality constraints.

The proof that \mathcal{B} satisfies all TLL-literals in Γ is not shown here since, as we said, many TLL3-literals are shared with TLL. The proof for the shared literals can be found in [RZ06b]. Here we just limit ourselves to the new literals defined in TLL3. Hence, we consider the following cases:

Literals of the form $t_1 \neq t_2$: Immediate.

Literals of the form $c = \text{mkcell}(e, a, t)$: We know that

$$\begin{aligned}
c^{\mathcal{B}} &= c^A \\
&= \text{mkcell}^A(e, a, t) \\
&= \langle e^A, a^A, t^A \rangle \\
&= \langle e^{\mathcal{B}}, a^{\mathcal{B}}, t^{\mathcal{B}} \rangle \\
&= \text{mkcell}^{\mathcal{B}}(e, a, t)
\end{aligned}$$

Literals of the form $c = rd(m, a)$: In this case we have that

$$\begin{aligned}
[rd(m, a)]^{\mathcal{B}} &= m^{\mathcal{B}}(a^{\mathcal{B}}) \\
&= m^{\mathcal{B}}(a^A) \\
&= (m^A(a^A).data^A, first(m^A, a^A, \mathcal{B}_{\text{addr}}), m^A(a^A).lockid^A)
\end{aligned}$$

$$\begin{aligned}
&= (m^A(a^A).data^A, m^A(a^A).next^A, m^A(a^A).lockid^A) && \text{(Lemma 4.2(a))} \\
&= m^A(a^A) \\
&= c^A \\
&= c^B
\end{aligned}$$

Literals of the form $m = upd(\tilde{m}, a, c)$: In this particular case we want to prove that $m^B(a^B) = c^B$ while for any other $v \in \mathcal{B}_{\text{addr}}$, such that $v \neq a^B$, $m^B(v) = \tilde{m}^B(v)$. Since $m^A(a^A) = c^A$, then we have that $c^B = m^B(a^B)$. Let now $v \neq a^B$. Then,

$$\begin{aligned}
m^B(v) &= (m^A(v).data^A, first(m^A, v, \mathcal{B}_{\text{addr}}), m^A(v).lockid^A) \\
&= (\tilde{m}^A(v).data^A, first(\tilde{m}^A, v, \mathcal{B}_{\text{addr}}), \tilde{m}^A(v).lockid^A) && \text{(Lemma 4.2(b))} \\
&= \tilde{m}^B(v)
\end{aligned}$$

Literals of the form $a = fstlock(m, p)$: First consider the case $p = \epsilon$. Hence $fstlock^A(m^A, \epsilon^A) = null^A$. At the same time, we know that $\epsilon^B = compress(\epsilon^A, \mathcal{B}_{\text{addr}})$ and so $fstlock^B(m^B, \epsilon^B) = null^B$. Let's now consider the case in which $p = [a_1, \dots, a_n]$. There are two possible scenarios to consider.

- If for all $1 \leq k \leq n$, $m^A(a_k^A).lockid^A = \emptyset$, then we have that

$$fstlock^A(m^A, p^A) = null^A$$

Notice that function *compress* returns a subset of the path it receives with the property that all addresses in the returned path belong to the received path. Therefore, if we have $[\tilde{a}_1, \dots, \tilde{a}_m] = p^B = compress(p^A, \mathcal{B}_{\text{addr}})$, we can deduce that $\{\tilde{a}_1, \dots, \tilde{a}_m\} \subseteq \mathcal{B}_{\text{addr}}$ and hence for all $1 \leq j \leq m$, $m^B(\tilde{a}_j).lockid^B = \emptyset$. Then, we conclude that $fstlock^B(m^B, p^B) = null^B$.

- If there is a $1 \leq k \leq n$ s.t., for all $1 \leq j < k$, $m^A(a_j^A).lockid^A = \emptyset$ and $m^A(a_k^A).lockid^A \neq \emptyset$ then since by the construction of model \mathcal{B} , we have that $a^B = a^A$. Therefore, we can say that $a^B = a^A = x \in \mathcal{B}_{\text{addr}}$. It then remains to verify whether

$$x = fstlock^A(m^A, p^A) \quad \rightarrow \quad x = fstlock^B(m^B, compress(p^A, \mathcal{B}_{\text{addr}}))$$

By definition of *fstlock* we have that $x = a_k^A$ and by the construction of set $\mathcal{B}_{\text{addr}}$, we know that $a_k^A \in \mathcal{B}_{\text{addr}}$. Let $[\tilde{a}_1, \dots, \tilde{a}_i, \dots, \tilde{a}_m] = compress(p^A, \mathcal{B}_{\text{addr}})$ such that $\tilde{a}_i = a_k^A$. It is clear that $\tilde{a}_j \in \mathcal{B}_{\text{addr}}$ for all $1 \leq j \leq m$. Then, as *compress* preserves the order and for all $1 \leq j < k$, $m^A(a_j^A).lockid^A = \emptyset$, we have that for all $1 \leq j < i$, $m^B(\tilde{a}_j).lockid^B = \emptyset$. Besides $m^B(\tilde{a}_i).lockid^B \neq \emptyset$. Then:

$$\begin{aligned}
fstlock^B(m^B, compress(p^A, \mathcal{B}_{\text{addr}})) &= fstlock^B(m^B, [\tilde{a}_1, \dots, \tilde{a}_m]) \\
&= \tilde{a}_i \\
&= a^A \\
&= x
\end{aligned}$$

□

4.2 A Combination-based Decision Procedure for TLL3

Lemma 4.3 justifies a brute force method to automatically check TLL3 satisfiability of normalized TLL3-literals. However, such a method is not efficient in practice. To find a more efficient decision procedure we decompose TLL3 into a combination of theories, and apply a many-sorted variant of the Nelson-Oppen combination method [TZ04].

In [NO79], Nelson and Oppen show how a decision procedure for a theory built by many first-order theories can be derived as a combination of the existing decision procedures for each of these theories. The main idea is to combine the decision procedures by means of equality sharing, guessing an arrangement over the set of shared variables. This arrangement is used to build equalities and disequalities between variables, to constrain simultaneously the inputs of decision procedures for each of the component theories.

This method requires the theories to fulfill two conditions. First, each theory must have a decision procedure. Second, all involved theories must be stable infinite and share sorts only.

Definition 4.3 (stable-infiniteness).

A Σ -theory T is stably infinite if for every T -satisfiable quantifier-free Σ -formula φ there exists a T -interpretation \mathcal{A} satisfying φ whose domain is infinite. †

As we said, TLL [RZ06a] is a union of theories. All theories involved in TLL are stably-infinite, so the only missing theories are T_{settid} and the one defining *fstlock*. Clearly T_{settid} is stably infinite (following a similar reasoning as for T_{set}). This leaves us only with the problem of finding a decision procedure for Σ_{reach} and a way to define the functions and predicates in Σ_{bridge} . We begin by defining the theory T_{Base3} as follows:

$$T_{\text{Base3}} = T_{\text{addr}} \cup T_{\text{elem}} \cup T_{\text{cell}} \cup T_{\text{mem}} \cup T_{\text{path}} \cup T_{\text{set}} \cup T_{\text{settid}} \cup T_{\text{thid}}$$

where T_{path} extends the theory of finite sequences of addresses with the auxiliary functions and predicates depicted in Fig. 4.3. Moreover, we say that Σ_{Base3} is the signature obtained by the union of all the signatures for each theory that is part of T_{Base3} . The theory of finite sequences of addresses is defined by $T_{\text{fseq}} = (\Sigma_{\text{fseq}}, \text{TGen})$, where:

$$\Sigma_{\text{fseq}} = \left(\left\{ \begin{array}{l} \text{addr, fseq} \\ \{ \text{nil} : \text{fseq}, \\ \text{cons} : \text{addr} \times \text{fseq} \rightarrow \text{fseq}, \\ \text{hd} : \text{fseq} \rightarrow \text{addr}, \\ \text{tl} : \text{fseq} \rightarrow \text{fseq} \end{array} \right\}, \right. \\ \left. \{ \} \right)$$

and TGen is the class of multi-sorted term-generated structures that satisfy the axioms of T_{fseq} . These axioms are the standard for a theory of lists, such as distinctness, uniqueness and generation of sequences using the constructors *cons* and *nil*, as well as acyclicity of sequences (see, for example [BM07]). Let *PATH* be the set of axioms of T_{fseq} including all in Fig. 4.3. Using these definitions, we can formally define $T_{\text{path}} = (\Sigma_{\text{path}}, \text{ETGen})$ where ETGen is $\{ \mathcal{A}^{\Sigma_{\text{path}}} \mid \mathcal{A}^{\Sigma_{\text{path}}} \models \text{PATH} \text{ and } \mathcal{A}^{\Sigma_{\text{fseq}}} \in \text{TGen} \}$.

Next, we extend T_{Base3} defining the missing functions and predicates from T_{reach} and Σ_{bridge} . For *addr2set* and *reach* the definition is immediate from *fseq2set* and *isreachp_K* respectively. For the other

functions and predicates we can use the following definitions:

$$\begin{aligned}
nil &= \epsilon \\
cons(a, nil) &= [a] \\
ispath(p) \rightarrow fseq2set(p) &= path2set(p) \\
\left(\begin{array}{l} ispath(p_1) \wedge ispath(p_2) \\ fseq2set(p_1) \cap fseq2set(p_2) = \emptyset \\ app(p_1, p_2) = p_3 \end{array} \wedge \right) &\leftrightarrow append(p_1, p_2, p_3) \\
isreachp(m, a_1, a_2, p) &\rightarrow getp(m, a_1, a_2) = p \\
\neg isreachp(m, a_1, a_2, p) &\rightarrow getp(m, a_1, a_2) = nil \\
ispath(p) \wedge fstmark(m, p, i) &\leftrightarrow fstlock(m, p) = i
\end{aligned}$$

Let GAP be the set of axioms that define ϵ , $[\]$, $append$, $reach$, $path2set$, $addr2set$ and $getp$. We define $\widehat{TLL3} = (\Sigma_{\widehat{TLL3}}, \widehat{ETGen})$ where $\Sigma_{\widehat{TLL3}}$ is $\Sigma_{Base3} \cup \{getp, append, path2set, fstlock\}$ and \widehat{ETGen} is $\{\mathcal{A}^{\Sigma_{\widehat{TLL3}}} \mid \mathcal{A}^{\Sigma_{\widehat{TLL3}}} \models GAP \text{ and } \mathcal{A}^{\Sigma_{path}} \in \widehat{ETGen}\}$.

Using the definitions of GAP it is easy to prove that if Γ is a set of normalized TLL3-literals, then Γ is TLL3-satisfiable iff Γ is $\widehat{TLL3}$ -satisfiable. It is enough to carry out an analysis by cases. For all literals belonging to both theories, this verification is straightforward. On the other hand, for literals of the form involving $path2set$, $append$ or $getp$ for instance, we just need to apply the definitions introduced by GAP and then apply the definitions of $PATH$.

This way, $\widehat{TLL3}$ can be used in place of TLL3 for satisfiability checking. We reduce $\widehat{TLL3}$ into T_{Base3} in two steps. First we do the unfolding of the definition of auxiliary functions defined in $PATH$ and GAP ,

$app : fseq \times fseq \rightarrow fseq$
$app(nil, l) = l$ $app(cons(a, l), l') = cons(a, app(l, l'))$
$fseq2set : fseq \rightarrow set$
$fseq2set(nil) = \emptyset$ $fseq2set(cons(a, l)) = \{a\} \cup fseq2set(l)$
$ispath : fseq$
$ispath(nil)$ $ispath(cons(a, nil))$ $\{a\} \not\subseteq fseq2set(l) \wedge ispath(l) \rightarrow ispath(cons(a, l))$
$last : fseq \rightarrow addr$
$last(cons(a, nil)) = a$ $l \neq nil \rightarrow last(cons(a, l)) = last(l)$
$isreach : mem \times addr \times addr$
$isreach(m, a, a)$ $m[a].next = a' \wedge isreach(m, a', b) \rightarrow isreach(m, a, b)$
$isreachp : mem \times addr \times addr \times fseq$
$isreachp(m, a, a, nil)$ $m[a].next = a' \wedge isreachp(m, a', b, p) \rightarrow isreachp(m, a, b, cons(a, p))$
$fstmark : mem \times fseq \times addr$
$fstmark(m, nil, null)$ $p \neq nil \wedge p = cons(j, q) \wedge m[j].lockid \neq \emptyset \rightarrow fstmark(m, p, j)$ $p \neq nil \wedge p = cons(j, q) \wedge m[j].lockid = \emptyset \wedge fstmark(m, q, i) \rightarrow fstmark(m, p, i)$

Figure 4.3: Functions, predicates and axioms of T_{path}

$List : \text{mem} \times \text{addr} \times \text{set}$
$List(h, a, r) \leftrightarrow null \in \text{addr2set}(h, a) \wedge r = \text{path2set}(\text{getp}(h, a, null))$
$f_a : \text{mem} \times \text{addr} \rightarrow \text{path}$
$f_a(h, n) = \begin{cases} \epsilon & \text{if } n = \text{null} \\ \text{getp}(h, h[n].\text{next}, \text{null}) & \text{if } n \neq \text{null} \end{cases}$
$f_b : \text{mem} \times \text{addr} \times \text{addr} \times \text{path}$
$f_b(h, n, m, p) \leftrightarrow (n = \text{null} \wedge p = \epsilon) \vee$ $(n \neq \text{null} \wedge \text{reach}(h, n, m, \tilde{p}) \wedge m = \text{null} \rightarrow \text{append}(p, \text{null}, \tilde{p})$ $\wedge m \neq \text{null} \rightarrow \text{append}(p, [m], \tilde{p}))$
$LastMarked : \text{mem} \times \text{path} \rightarrow \text{addr}$
$LastMarked(m, p) = \text{fstlock}(m, \text{rev}(p))$
$NoMarks : \text{mem} \times \text{path}$
$NoMarks(m, p) \leftrightarrow \text{fstlock}(m, p) = \text{null}$
$SomeMark : \text{mem} \times \text{path}$
$SomeMark(m, p) \leftrightarrow \text{fstlock}(m, p) \neq \text{null}$

Figure 4.4: Auxiliary functions to reason about concurrent lists

getting rid of the extra functions, and obtaining a formula in T_{Base3} . Second, as we have proved that TSL_K enjoys finite model property then it is always possible to enumerate the finitely many ground terms. This means that we can always find a model with a finite number of elements that preserves satisfiability of a conjunction of TSL_K literals. Hence, by substituting all possible ground terms in the set of normalized literals and symbolically executing the definitions of the symbols in PATH it is possible to reduce the $\overline{\text{TLL3}}$ -satisfiability problem of normalized literals to the T_{Base3} -satisfiability of quantifier free formulas, setting a bound for the unfolding of recursive formulas. This way it is possible to obtain a decision procedure for reachability, based on the finite model property of the theory.

All theories involved in T_{Base3} share only sorts symbols, are stably-infinite and for all of them there is a decision procedure. Hence, the multi-sorted Nelson-Oppen combination method can be applied, obtaining a decision procedure for TLL3 .

We now define some auxiliary functions and predicates using TLL3 , that aid in the reasoning about concurrent linked-lists (see Fig. 4.4). For example, predicate $List(h, a, r)$ expresses that in heap h , starting from address a there is sequence of cells all of which form region r . Function f_a , given a heap h and an address n , returns the path that goes from the address pointed by the cell stored in n up to the end of the list (i.e., a null pointer is found). Similarly, predicate f_b holds when p is the path that connects the node stored at address n with the one at address m . Notice that m is not considered as part of the path. Function $LastMarked(h, p)$, on the other hand, returns the address of the last locked node in path p on memory h . All these functions can be used in verification conditions. Then, using the equivalences in Fig. 4.4 the predicates are removed, generating a pure $\overline{\text{TLL3}}$ formula whose satisfiability can be checked with the procedure described above.

4.3 Verifying Some Properties Over Concurrent Lists

In this section we sketch the proofs for some properties over lock-coupling singly-linked lists. The reader will notice that all literals involved in the formulas described from this point on belong to TLL3 . This way, the decision procedure presented in this chapter can be used to automatically verify them all.

Firstly, we prove that the thread owning the locks closest to the tail of the list, eventually terminates. We have already introduced this property on Example 3.1. Finally, we present the idea behind the verification of the property describing that no thread can overtake other thread that already owns some lock in the list.

4.3.1 Termination of Concurrent Lock-Coupling Lists

In this section we show the proof of a simple liveness property of concurrent lock-coupling lists: termination of the leading thread. To aid in the verification of this property we annotate the code with ghost fields and ghost updates. We use boxes to represent the annotations introduced in the code. We enrich *List* objects introduced in Section 3.1 with a ghost field *r* of type region that keeps track of all the nodes in the list:

```

class List {
  Node list;
  rgn r;
}

class Node {
  Value val;
  Node next;
  Lock lock;
}

```

The algorithms for INSERT and REMOVE are shown as Algorithms 4.1 and 4.2 respectively.

The predicate $c.lockid = \circ$ denotes that the lock of list node *c* is not taken. The predicate $c.lockid = k$ establishes that the lock at list node *c* is owned by thread *k*. As shown, the code for INSERT and REMOVE is extended with ghost updates to maintain *r*.

T_k denotes thread *k*. We want to prove that if a thread has acquired a lock at node *n* and no other thread holds a lock ahead of *n*, then thread *k* eventually terminates. The predicate $at_INSERT_n^{[k]}$ means that thread *k* is executing line *n* of program INSERT. Similarly, $at_INSERT_{n_1, \dots, n_m}^{[k]}$ is a short for thread *k* is running some of the lines n_1, \dots, n_m of program INSERT. We use $n..m$ to denote all correlative lines from *n* to *m*. The instance of a local variable *v* in thread *k* is represented by $v^{[k]}$. We define *DisjList* as an extension of *List* enriching it with the property that new nodes created during insertion are all disjoint from each other, including all nodes that are already part of the list:

$$\begin{aligned}
DisjList(h, a, r) &\hat{=} List(h, a, r) && \wedge \\
&\forall j : T_{ID}. at_INSERT_{5,6}^{[j]} \rightarrow \langle aux^{[j]} \rangle \# r && \wedge \\
&\forall i, j : T_{ID}. i \neq j \wedge at_INSERT_{5,6}^{[i]} \wedge at_INSERT_{5,6}^{[j]} \rightarrow \langle aux^{[i]} \rangle \# \langle aux^{[j]} \rangle \# r
\end{aligned}$$

We now define the following auxiliary predicate:

$$\begin{aligned}
IsLast(k) &\hat{=} DisjList(h, l.list, l.r) && \wedge \\
&SomeMark(h, getp(h, l.list, null)) && \wedge \\
&LastMarked(h, getp(h, l.list, null)) = a && \wedge \\
&h[a].lockid = k
\end{aligned}$$

Algorithm 4.1 Insertion for concurrent lock-coupling lists extended with ghost code

```

1: procedure INSERT( Value e)
2:   prev, curr := locate(e)
3:   if curr.val ≠ e then
4:     aux := new Node(e)
5:     aux.next := curr
6:     prev.next := aux ————— L.r := L.r ∪ {aux}
7:     result := true
8:   else
9:     result := false
10:  end if
11:  prev.unlock()
12:  curr.unlock()
13:  return result
14: end procedure

```

Algorithm 4.2 Remove for concurrent lock-coupling lists extended with ghost code

```

1: procedure REMOVE(Value  $e$ )
2:    $prev, curr := locate(e)$ 
3:   if  $curr.val = e$  then
4:      $aux := curr.next$ 
5:      $prev.next := aux$ 
6:      $result := true$ 
7:   else
8:      $result := false$ 
9:   end if
10:   $prev.unlock()$ 
11:   $curr.unlock()$ 
12:  return  $result$ 
13: end procedure

```

 $l.r := l.r - \{aux\}$

Notice how the ghost variable r of type region is used in the *List* predicate to assert that new nodes do not already belong to the list. The formula $IsLast(k)$ identifies whether T_k is the thread owning the last lock in the list (i.e., the closest node towards the end of the list). Using these predicates we define the parametrized temporal formula we want to verify as:

$$\psi(k) \triangleq \square \left(at_LOCATE_{4,11}^{[k]} \wedge IsLast(k) \rightarrow IsLast(k) \mathcal{U} at_LOCATE_{12}^{[k]} \right)$$

This temporal formula states that if thread k is running *LOCATE* and it owns the last locked node in the list, then thread T_k will still own the last locked node until T_k reaches the last line of *LOCATE*. Reachability of the last line of *LOCATE* implies termination of the invocation to the concurrent datatype because *LOCATE* is the only program containing potentially blocking operations.

We proceed with the construction of a verification diagram that proves that the parallel execution of all threads guarantees the satisfaction of formula $\psi(k)$. Given N , we build the transition system $\mathcal{S}[N]$, in which threads T_1, \dots, T_N run in parallel the program *MGC* and show that $\mathcal{S}[N] \models \psi(k)$. The verification diagram is depicted in Fig. 4.5.

We use $\tau_{p_n}^{[k]}$ to denote the transition taken by thread k on program p at line n , and $\tau_{p_n}^{[N]}$ to describe the set $\tau_{p_n}^{[j]}$, for all $j \in \{1..N\}$.

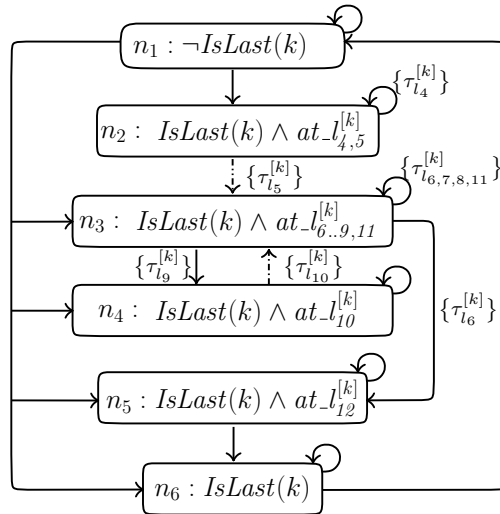


Figure 4.5: Verification diagram Ψ for $\parallel_{j < N} T_j \models \psi(k)$

Dashed arrows in the diagram denote transitions that strictly decrement the ranking function δ . Formally, the verification diagram is defined by:

- $N_0 = \{n_1\}$
- $\mathcal{F} = \{(P, R)\}$, with $P = \{(n_3, n_4), (n_3, n_5), (n_5, n_6), (n_6, n_1)\} \cup \{(n_1, n_j) | j \in 2..6\} \cup \{(n_j, n_j) | j \in 1..6\}$ and $R = \emptyset$
- $\delta(n, s) = \begin{cases} \{a \mid a \in \text{dom}(h)\} & n = n_1, n_2 \\ \text{path2set}(f_a(h, \text{LastMarked}(h, \text{getp}(h, \text{prev}^{[k]}, \text{null}))) & \text{otherwise} \end{cases}$
- $f(n) = \begin{cases} \emptyset & \text{if } n = n_1, n_6 \\ \text{at_LOCATE}_{4,5}^{[k]} & \text{if } n = n_2 \\ \text{at_LOCATE}_{6..9,11}^{[k]} & \text{if } n = n_3 \\ \text{at_LOCATE}_{10}^{[k]} & \text{if } n = n_4 \\ \text{at_LOCATE}_{12}^{[k]} & \text{if } n = n_5 \end{cases}$

We can now describe the verification conditions:

initialization Trivial, since in the initial state $l.\text{list}$ forms an empty list, and consequently $\neg \text{IsLast}(k)$.

consecution We will show, for illustration purposes, transition $\tau_{\text{LOCATE}_{10}}^{[j]}$ on node n_2 with $j \neq k$. The verification condition is:

$$\left(\begin{array}{c} \overbrace{\text{IsLast}(k)}^{\text{TLL3}} \wedge \overbrace{j \neq k}^{T_{\text{hid}}} \wedge \\ \text{at_l}_{4,5}^{[k]} \wedge \text{at_l}_{10}^{[j]} \wedge \\ \text{curr}^{[j]}.lockid = \emptyset \end{array} \right) \wedge \overbrace{\text{curr}^{[j]}.lock(j)}^{\text{TLL3}} \rightarrow \left(\begin{array}{c} \overbrace{\text{IsLast}(k') \wedge \text{at}'_{4,5}^{[k']}}^{\text{TLL3}} \wedge \\ \text{at}'_{11}^{[j']} \wedge \text{pres}(V - \text{curr}^{[j]}) \wedge \\ \text{curr}'^{[j']}.lockid = j' \end{array} \right)$$

where pres is the predicate denoting variable preservation. Note that all fragments of such verification condition belong to theories for which we have already defined a decision procedure, including propositional logic for the (finite) locations of the program counters.

acceptance The ranking function δ maps, at a given state, the set of list nodes accessible from the last node with an owned lock. This set remains identical for all transitions except $\tau_{l_5}^{[k]}$ and $\tau_{l_{10}}^{[k]}$, for which the set decrements (in the inclusion order on sets). The decision procedure presented in this chapter proves this automatically (using \subset operation and equality over sets of addresses).

fairness Only two conditions must be verified. First, all transitions labeling an edge are enabled since the only potentially blocking operation is $\tau_{l_{10}}^{[k]}$ and $\text{IsLast}(k)$ implies that $\tau_{l_{10}}^{[k]}$ is enabled. Second, for all nodes and labeled edges, starting from a state that satisfies the predicate of the incoming node satisfies the predicate of the outgoing node via taking the transition. Sequential progress of thread k is guaranteed by fairness, since all idling transitions for thread k are in fact a diagram idiom to represent the expansion of such nodes to a sequence of nodes with a single program position on each node.

satisfaction $\mathcal{L}(\Psi) \subseteq \mathcal{L}(\psi(k))$ is automatically checkable via a finite LTL model-checking problem. In this case, instead of proving that all accepting paths in the diagram are contained into the set of sequences satisfying formula $\psi(k)$, we show that the intersection with the negation of the formula is empty. We use $\langle v_1, v_2, v_3 \rangle$ to represent $\langle \text{at_l}_{4..11}^{[k]}, \text{at_l}_{12}^{[k]}, \text{IsLast}(k) \rangle$. In fact $\text{IsLast}(k)$ should be decomposed into all its atomic subformulas. However, for the sake of simplicity we assume that an

assignment to $IsLast(k)$ represents all necessarily assignments to its atomic subformulas in order to make $IsLast(k)$ predicate true. Then, since

$$\neg\psi(k) = \diamond \left(at_l_{4..11}^{[k]} \wedge IsLast(k) \wedge \left(\neg at_l_{12}^{[k]} \mathcal{W} \neg IsLast(k) \right) \right)$$

we have that:

$$\mathcal{L}^P(\neg\psi(k)) = \langle -, -, - \rangle^* \langle T, F, T \rangle (\langle -, F, - \rangle^\omega \cup \langle -, F, - \rangle^* \langle -, -, F \rangle \langle -, -, - \rangle^\omega)$$

while for our verification diagram we have

$$(\langle -, -, F \rangle^+ \langle T, F, T \rangle^* \langle F, T, T \rangle^+ \langle -, -, T \rangle^+)^\omega$$

Imagine we consider the sequence $\langle -, -, - \rangle^* \langle T, F, T \rangle \langle -, F, - \rangle^\omega$. Then, $\langle T, F, T \rangle$ should correspond to $\langle T, F, T \rangle^*$. Then, it is impossible to match $\langle -, F, - \rangle^\omega$ with any possible pattern in the accepting paths of the diagram. On the other hand, if we consider the sequence $\langle -, F, - \rangle^* \langle -, -, F \rangle \langle -, -, - \rangle^\omega$, notice that there is no possible way to match $\langle -, -, F \rangle$ with the accepting paths of the diagram. This way we have shown that both languages are disjoint.

4.3.2 No Thread Overtakes

In this section we sketch the proof that it is not possible in the presented implementation that, once a thread has acquired a lock, being overtaken by other thread. To carry out the proof, we first extend the code for concurrent lock-coupling singly-linked lists with two ghost variables: a global ghost variable *ticket* and a local ghost variable *myTicket*. Then, we modify LOCATE to consider the new ghost variables. The extended version of the LOCATE program with the corresponding modification is depicted in Algorithm 4.3.

Algorithm 4.3 Locate program for preventing overtake

```

1: procedure LOCATE( Value e)
2:   prev := Head
3:   prev.lock()
4:   curr := prev.next
5:   curr.lock()
6:   while curr.val < e do
7:     prev.unlock()
8:     prev := curr
9:     curr := curr.next
10:    curr.lock()
11:  end while
12:  return (prev, curr)
13: end procedure

```

myTicket := *ticket*
ticket := *ticket* + 1

The idea is that every time a thread gets the first lock in the list, it also gets a ticket number. Tickets are delivered by the data structure and they are strictly increasingly. This policy guarantees that no thread gets a duplicated ticket and threads are ordered according to the order in which they acquired their first lock. We want to assure that there is no overtake in the system. This means that all threads remain ordered as they own a ticket. To express the formula that describes this condition we first declare the *ahead* predicate defined over thread identifiers t_1 and t_2 and a memory configuration m :

$$ahead(t_1, t_2) \hat{=} haveLocks(t_1, t_2) \rightarrow \left(\begin{array}{l} fstlock(m, prev^{[t_1]}) \neq null \\ fstlock(m, prev^{[t_2]}) \neq null \\ fstlock(m, prev^{[t_2]}) \in addr2set(m, curr^{[t_1]}) \end{array} \wedge \right)$$

where

$$haveLocks(t_1, t_2) = at_LOCATE_{4..12}^{[t_1, t_2]} \vee at_INSERT_{3..12}^{[t_1, t_2]} \vee at_REMOVE_{3..11}^{[t_1, t_2]} \vee at_SEARCH_{3..9}^{[t_1, t_2]}$$

In the formula above, we use $at_p_n^{[t_1, t_2]}$ to denote $at_p_n^{[t_1]} \wedge at_p_n^{[t_2]}$. Roughly speaking, predicate *ahead* holds when both threads have at least one locked node and the leftmost node locked by thread T_2 (that is, the nearest to the head of the list) is between the position of t_1 and the tail of the list.

We can now use the *ahead* predicate to define the condition we want to verify. We use φ to represent such formula:

$$\varphi \triangleq \left(myTicket^{[t_1]} > myTicket^{[t_2]} \right) \rightarrow \square \left(ahead(t_1, t_2) \rightarrow \bigcirc \left(ahead(t_1, t_2) \right) \right)$$

Once more, notice how ghost variables *myTicket* and *ticket* are used and are part of formula φ . Now, φ is a system invariant; we just need to verify that:

- φ is satisfied by the system's initial condition, and
- all transitions preserve φ .

It is easy to see that the initial condition satisfies φ . At such moment, no ticket number has been assigned nor any thread has acquired any lock. Then, formula φ is trivially satisfied.

Now we analyze that every transition preserves φ . In fact, here we do not analyze all possible transitions, but we just focus on the possible offending ones. That is, transitions that get or release a lock, or that modifies the position of the *curr* pointer. Notice that all these transitions can potentially modify the validity of φ . All other transitions are not relevant as the modification they introduced cannot modify the value of φ at all or they are simply trivially valid because of propositional reasoning on the program location. Now, the transitions we analyze are:

- When a thread acquires its first lock. For an arbitrary thread t , this transition corresponds to $at_LOCATE_3^{[t]}$. This transition clearly satisfies φ' because the value of *ticket* is strictly increasing. This guarantees that for all other thread s with locks in the list, $myTicket^{[t]} > myTicket^{[s]}$ and *ahead*(t, s) hold, since $fstlock(m, s) \in addr2set(m, prev^{[t]})$.
- When a thread gets a new lock. This modification is accomplished by the transitions $at_LOCATE_5^{[t]}$ and $at_LOCATE_{10}^{[t]}$. If some of these transitions are taken by thread t_2 then φ is clearly preserved. If the transition is taken by t_1 , the ticket numbering and order is preserved. Besides, t_1 acquires its new lock over a node where no other thread has a lock. Hence, once more we conclude that $fstlock(m, t_2) \in addr2set(m, prev^{[t_1]})$.
- When a thread releases its last lock. This happens when transitions $at_INSERT_{12}^{[t]}$, $at_REMOVE_{11}^{[t]}$ or $at_SEARCH_9^{[t]}$ are taken. Notice that in this case one of the thread is moving outside of the range set by the *haveLocks* predicate. Hence, $\bigcirc ahead(t_1, t_2)$ trivially holds and therefore φ does.
- When the position of a thread advances through the nodes in the lists. This progress corresponds to transitions $at_LOCATE_4^{[t]}$ and $at_LOCATE_9^{[t]}$. In this case ticket numbers are not modified nor any lock is acquired or removed. Hence, φ clearly holds.

5

TSL_K: A Decision Procedure for Concurrent Skiplists

In this chapter, we present a theory of skiplists with a decidable satisfiability problem, and show its applications to the verification of concurrent skiplist implementations. As already described in Section 3.2, a skiplist is a data structure used to implement sets by maintaining several ordered singly-linked lists in memory, with a performance comparable to balanced binary trees. We define a theory capable of expressing the memory layout of a skiplist and show a decision procedure for the satisfiability problem of this theory. Concurrent lock-coupling skiplists introduced in Chapter 3, where every node contains a lock at each possible level, reduce granularity of mutual exclusion sections.

We now proceed with the description of the theory TSL_K . This theory is briefly described in [SS11]. TSL_K is a decidable theory capable of reasoning about list reachability, locks, ordered lists, and sublists of ordered lists. First, we show that TSL_K enjoys a finite model property and thus it is decidable. Finally, we show how to reduce the satisfiability problem of quantifier-free TSL_K formulas to a combination of theories for which a many-sorted version of Nelson-Oppen can be applied.

We build a decision procedure to reason about skiplist of height K combining different theories, aiming to represent pointer data structures with a skiplist layout, masked regions and locks. We extend the Theory of Concurrent Linked Lists (TLL3) presented in Chapter 4, a decidable theory that includes reachability of concurrent list-like structures, in the following way:

- each node is equipped with a *key* field, used to reason about node's order.
- the reasoning about single level lists is extended to all the K levels.
- we extend the theory of regions with masked regions.
- lists are extended to ordered lists and sub-paths of ordered lists.

Formally, the theory of skiplists of height K is defined as $TSL_K = (\Sigma_{TSL_K}, \mathbf{TSLK})$, where

$$\Sigma_{TSL_K} = \Sigma_{level_K} \cup \Sigma_{ord} \cup \Sigma_{thid} \cup \Sigma_{cell} \cup \Sigma_{mem} \cup \Sigma_{reach} \cup \Sigma_{set} \cup \Sigma_{settid} \cup \Sigma_{mrgn} \cup \Sigma_{aarr} \cup \Sigma_{larr} \cup \Sigma_{bridge}$$

Then, the signature of TSL_K is shown in Fig. 5.1 and 5.2. \mathbf{TSLK} is the class of Σ_{TSL_K} -structures satisfying the conditions depicted in Fig. 5.3 and 5.4. The symbols of Σ_{set} and Σ_{settid} follow their standard

Signature:	Σ_{level_K}
Sorts:	level_K
Functions:	$0, 1, \dots, K-1$: level_K
Predicates:	$<$: $\text{level}_K \times \text{level}_K$
Signature:	Σ_{ord}
Sorts:	ord
Functions:	$-\infty, +\infty$: ord
Predicates:	\preceq : $\text{ord} \times \text{ord}$
Signature:	Σ_{thid}
Sorts:	thid
Functions:	\emptyset : thid
Signature:	Σ_{cell}
Sorts:	$\text{cell}, \text{elem}, \text{ord}, \text{level}_K, \text{addr}, \text{thid}$
Functions:	error : cell mkcell : $\text{elem} \times \text{ord} \times \text{addr}^K \times \text{thid}^K \rightarrow \text{cell}$ $_.\text{data}$: $\text{cell} \rightarrow \text{elem}$ $_.\text{key}$: $\text{cell} \rightarrow \text{ord}$ $_.\text{next} [_]$: $\text{cell} \times \text{level}_K \rightarrow \text{addr}$ $_.\text{lockid} [_]$: $\text{cell} \times \text{level}_K \rightarrow \text{thid}$ $_.\text{lock} [_]$: $\text{cell} \times \text{level}_K \rightarrow \text{thid} \rightarrow \text{cell}$ $_.\text{unlock} [_]$: $\text{cell} \times \text{level}_K \rightarrow \text{cell}$
Signature:	Σ_{mem}
Sorts:	$\text{mem}, \text{addr}, \text{cell}$
Functions:	null : addr $_.[_]$: $\text{mem} \times \text{addr} \rightarrow \text{cell}$ upd : $\text{mem} \times \text{addr} \times \text{cell} \rightarrow \text{mem}$
Signature:	Σ_{reach}
Sorts:	$\text{mem}, \text{addr}, \text{level}_K, \text{path}$
Functions:	ϵ : path $[_]$: $\text{addr} \rightarrow \text{path}$
Predicates:	append : $\text{path} \times \text{path} \times \text{path}$ reach_K : $\text{mem} \times \text{addr} \times \text{addr} \times \text{level}_K \times \text{path}$
Signature:	Σ_{set}
Sorts:	addr, set
Functions:	\emptyset : set $\{ _ \}$: $\text{addr} \rightarrow \text{set}$ \cup, \cap, \setminus : $\text{set} \times \text{set} \rightarrow \text{set}$
Predicates:	\in : $\text{addr} \times \text{set}$ \subseteq : $\text{set} \times \text{set}$

Figure 5.1: The signature of the TSL_K theory - Part I

Signature:	Σ_{settid}
Sorts:	thid, settid
Functions:	\emptyset_T : settid $\{_ \}_T$: thid \rightarrow settid $\cup_T, \cap_T, \setminus_T$: settid \times settid \rightarrow settid
Predicates:	\in_T : thid \times settid \subseteq_T : settid \times settid
Signature:	Σ_{mrngn}
Sorts:	mrngn, addr, level _K
Functions:	\mathbf{emp}_{MR} : mrngn $\langle _, _ \rangle_{MR}$: addr \times level _K \rightarrow mrngn $\cup_{MR}, \cap_{MR}, \neg_{MR}$: mrngn \times mrngn \rightarrow mrngn
Predicates:	\in_{MR} : addr \times level _K \times mrngn \subseteq_{MR} : mrngn \times mrngn $\#_{MR}$: mrngn \times mrngn
Signature:	Σ_{aarr}
Sorts:	aarr, thid, addr
Functions:	$_ \lfloor _ \rfloor_A$: aarr \times thid \rightarrow addr $_ \{ _ \leftarrow _ \}_A$: aarr \times thid \times addr \rightarrow aarr
Signature:	Σ_{larr}
Sorts:	larr, thid, level _K
Functions:	$_ \lfloor _ \rfloor_L$: larr \times thid \rightarrow level _K $_ \{ _ \leftarrow _ \}_L$: larr \times thid \times level _K \rightarrow larr
Signature:	Σ_{bridge}
Sorts:	mem, addr, level _K , set, path
Functions:	$path2set$: path \rightarrow set $addr2set_K$: mem \times addr \times level _K \rightarrow set $getp_K$: mem \times addr \times addr \times level _K \rightarrow path $fstlock_K$: mem \times path \times level _K \rightarrow addr
Predicates:	$ordList$: mem \times path

Figure 5.2: The signature of the TSL_K theory - Part II

signature:	Σ_{cell}
interpretation:	<ul style="list-style-type: none"> • $\text{mkcell}^A(e, k, \vec{d}, \vec{t}) = \langle e, k, \vec{d}, \vec{t} \rangle$ • $\text{error}^A.\text{next}^A = \text{null}^A$ • $\langle e, k, \vec{d}, \vec{t} \rangle.\text{data}^A = e$ • $\langle e, k, \vec{d}, \vec{t} \rangle.\text{key}^A = k$ • $\langle e, k, \vec{d}, \vec{t} \rangle.\text{next}^A[j] = a_j$ • $\langle e, k, \vec{d}, \vec{t} \rangle.\text{lockid}^A[j] = t_j$ • $\langle e, k, \vec{d}, \dots, t_{j-1}, t_j, t_{j+1}, \dots \rangle.\text{lock}^A[j](t') = \langle e, k, \vec{d}, \dots, t_{j-1}, t', t_{j+1}, \dots \rangle$ • $\langle e, k, \vec{d}, \dots, t_{j-1}, t_j, t_{j+1}, \dots \rangle.\text{unlock}^A[j] = \langle e, k, \vec{d}, \dots, t_{j-1}, \emptyset, t_{j+1}, \dots \rangle$ <p>for each $e \in \mathcal{A}_{\text{elem}}, k \in \mathcal{A}_{\text{ord}}, t_0, \dots, t_j, t_{j+1}, t_{j-1}, t' \in \mathcal{A}_{\text{thid}}, \vec{d} \in \mathcal{A}_{\text{addr}}^K,$ $\vec{t} \in \mathcal{A}_{\text{thid}}^K$ and $j \in \mathcal{A}_{\text{level}_k}$</p>
signature:	Σ_{mem}
interpretation:	<ul style="list-style-type: none"> • $m[a]^A = m(a)$ • $\text{upd}^A(m, a, c) = m_{a \rightarrow c}$ • $m^A(\text{null}^A) = \text{error}^A$ <p>for each $m \in \mathcal{A}_{\text{mem}}, a \in \mathcal{A}_{\text{addr}}$ and $c \in \mathcal{A}_{\text{cell}}$</p>
signature:	Σ_{reach}
interpretation:	<ul style="list-style-type: none"> • ϵ^A is the empty sequence • $[i]^A$ is the sequence containing $i \in \mathcal{A}_{\text{addr}}$ as the only element • $([i_1 \dots i_n], [j_1 \dots j_m], [i_1 \dots i_n, j_1 \dots j_m]) \in \text{append}^A$ iff $i_k \neq j_l$ • $(m, a_{\text{init}}, a_{\text{end}}, l, p) \in \text{reach}_K^A$ iff $a_{\text{init}} = a_{\text{end}}$ and $p = \epsilon$, or there exist addresses $a_1, \dots, a_n \in \mathcal{A}_{\text{addr}}$ such that: <ul style="list-style-type: none"> (a) $p = [a_1 \dots a_n]$ (b) $a_1 = a_{\text{init}}$ (c) $m(a_r).\text{next}^A[l] = a_{r+1}$, for $r < n$ (d) $m(a_n).\text{next}^A[l] = a_{\text{end}}$
signature:	Σ_{mrgn}
interpretation:	<ul style="list-style-type: none"> • $\text{emp}_{\text{MR}}^A = \emptyset$ • $r \cup_{\text{MR}}^A s = r \cup s$ • $(a, j) \in_{\text{MR}}^A r \leftrightarrow (a, j) \in r$ • $\langle a, j \rangle_{\text{MR}}^A = \{(a, j)\}$ • $r \cap_{\text{MR}}^A s = r \cap s$ • $r \subseteq_{\text{MR}}^A s \leftrightarrow r \subseteq s$ • $r -_{\text{MR}}^A s = r \setminus s$ • $r \#_{\text{MR}}^A s \leftrightarrow r \cap_{\text{MR}}^A s = \text{emp}_{\text{MR}}^A$ <p>for each $a \in \mathcal{A}_{\text{addr}}, j \in \mathcal{A}_{\text{level}_k}$ and $r, s \in \mathcal{A}_{\text{mrgn}}$</p>
signature:	Σ_{aarr}
interpretation:	<ul style="list-style-type: none"> • $x[t]_{\text{A}}^A = x(t)$ • $x\{t \leftarrow a\}_{\text{A}}^A = x_{\text{new}}$, such that $x_{\text{new}}(i) = \begin{cases} a & \text{if } i = t \\ x(i) & \text{otherwise} \end{cases}$ <p>for each $x, x_{\text{new}} \in \mathcal{A}_{\text{aarr}}, t \in \mathcal{A}_{\text{thid}}$ and $a \in \mathcal{A}_{\text{addr}}$</p>
signature:	Σ_{larr}
interpretation:	<ul style="list-style-type: none"> • $z[t]_{\text{L}}^A = x(t)$ • $z\{t \leftarrow l\}_{\text{L}}^A = z_{\text{new}}$, such that $z_{\text{new}}(i) = \begin{cases} l & \text{if } i = t \\ z(i) & \text{otherwise} \end{cases}$ <p>for each $z, z_{\text{new}} \in \mathcal{A}_{\text{larr}}, t \in \mathcal{A}_{\text{thid}}$ and $l \in \mathcal{A}_{\text{level}_k}$</p>

Figure 5.3: Characterization of a TSL_K-interpretation \mathcal{A}

signature:	Σ_{bridge}
interpretation:	<ul style="list-style-type: none"> • $\text{path2set}^A(p) = \{a_1, \dots, a_n\}$ for $p = [a_1, \dots, a_n] \in \mathcal{A}_{\text{path}}$ • $\text{addr2set}_K^A(m, a, l) = \{a' \in \mathcal{A}_{\text{addr}} \mid \exists p \in \mathcal{A}_{\text{path}} \cdot (m, a, a', l, p) \in \text{reach}_K\}$ • $\text{getp}_K^A(m, a_{\text{init}}, a_{\text{end}}, l) = \begin{cases} p & \text{if } (m, a_{\text{init}}, a_{\text{end}}, l, p) \in \text{reach}_K^A \\ \epsilon & \text{otherwise} \end{cases}$ for each $m \in \mathcal{A}_{\text{mem}}$, $p \in \mathcal{A}_{\text{path}}$, $l \in \mathcal{A}_{\text{level}_K}$ and $a_{\text{init}}, a_{\text{end}} \in \mathcal{A}_{\text{addr}}$ • $\text{fstlock}^A(m, [a_1 \dots a_n], l) = \begin{cases} a_k & \text{if there is } k \leq n \text{ such that} \\ & \text{for all } j < k, m[a_j].\text{lockid}[l] = \emptyset \\ & \text{and } m[a_k].\text{lockid}[l] \neq \emptyset \\ \text{null} & \text{otherwise} \end{cases}$ • $\text{ordList}^A(m, p)$ iff <ul style="list-style-type: none"> – $p = \epsilon$, or – $p = [a]$, or – $p = [a_1 \dots a_n]$ with $n \geq 2$ and $m(a_i).\text{key}^A \preceq m(a_{i+1}).\text{key}^A$ for all $1 \leq i < n$, for any $m \in \mathcal{A}_{\text{mem}}$

Figure 5.4: Characterization of a TSL_K-interpretation \mathcal{A}

interpretation over sets of addresses and thread identifiers respectively. Each sort σ in Σ_{TSL_K} is mapped to a non-empty set \mathcal{A}_σ such that:

- (a) $\mathcal{A}_{\text{addr}}$ and $\mathcal{A}_{\text{elem}}$ are discrete sets.
- (b) $\mathcal{A}_{\text{thid}}$ is a discrete set containing \emptyset .
- (c) $\mathcal{A}_{\text{level}_K}$ is the finite collection $0, \dots, K - 1$.
- (d) \mathcal{A}_{ord} is a total ordered set.
- (e) $\mathcal{A}_{\text{cell}} = \mathcal{A}_{\text{elem}} \times \mathcal{A}_{\text{ord}} \times \mathcal{A}_{\text{addr}}^K \times \mathcal{A}_{\text{thid}}^K$.
- (f) $\mathcal{A}_{\text{mem}} = \mathcal{A}_{\text{cell}}^{\mathcal{A}_{\text{addr}}}$.
- (g) $\mathcal{A}_{\text{path}}$ is the set of all finite sequences of (pairwise) distinct elements of $\mathcal{A}_{\text{addr}}$.
- (h) \mathcal{A}_{set} is the power-set of $\mathcal{A}_{\text{addr}}$.
- (i) $\mathcal{A}_{\text{settid}}$ is the power-set of $\mathcal{A}_{\text{thid}}$.
- (j) $\mathcal{A}_{\text{mrgn}}$ is the power-set of $\mathcal{A}_{\text{addr}} \times \mathcal{A}_{\text{level}_K}$.
- (k) $\mathcal{A}_{\text{aarr}} = \mathcal{A}_{\text{addr}}^{\mathcal{A}_{\text{thid}}}$.
- (l) $\mathcal{A}_{\text{larr}} = \mathcal{A}_{\text{level}_K}^{\mathcal{A}_{\text{thid}}}$.

Informally, sort `addr` represents addresses; `elem` the universe of elements that can be stored in the skiplist; `ord` the ordered keys used to preserve a strict order in the skiplist; `thid` thread identifiers; `levelK` the levels of a skiplist; `cell` models *cells* representing a node in a skiplist; `mem` models the heap, mapping addresses to cells; `path` describes finite sequences of non-repeating addresses to model non-cyclic list paths; `set` models sets of addresses – also known as regions –, while `settid` models sets of thread identifiers and `mrgn` masked regions; `aarr` models arrays of addresses indexed by thread identifiers and finally `larr` models arrays of skiplist levels indexed by thread identifiers.

Σ_{level_K} contains symbols for level identifiers $0, 1, \dots, K - 1$ and their conventional order. Σ_{ord} contains two special elements $-\infty$ and ∞ for the lowest and highest values in the order \preceq . Σ_{thid} only contains,

besides = and ≠ as for all the other theories, a special constant \emptyset to represent the absence of a thread identifier. Σ_{cell} contains the constructors and selectors for building and inspecting cells, including *error* for incorrect dereferences. Σ_{mem} is the signature for heaps, with the usual memory access and single memory mutation functions. Σ_{set} and Σ_{settid} are the signatures of theories of sets of addresses and thread ids resp. Σ_{mrgn} is the signature of the theory of masked regions while Σ_{aarr} and Σ_{larr} are the signatures for the theory of arrays over addresses and skiplist levels respectively, indexed in both cases by thread identifiers. The signature Σ_{reach} contains predicates to check reachability of address using paths at different levels, while Σ_{bridge} contains auxiliary functions and predicates to manipulate and inspect paths and locks.

Example 5.1

Consider the skiplist shown in Fig. 3.3 in Section 3.2. For that particular list and assuming that thread id variable j is assigned to value T_1 , we can construct a model \mathcal{A} such that for instance:

$$\begin{aligned} \mathcal{A}_{\text{addr}} &= \{0x00, 0x01, 0x02, 0x03, 0x04, 0x05, 0x06, 0x07, 0x08\} \\ \mathcal{A}_{\text{ord}} = \mathcal{A}_{\text{elem}} &= \{-\infty, 2, 8, 11, 14, 16, 20, +\infty\} \\ \mathcal{A}_{\text{thid}} &= \{T_1, T_2, \emptyset\} \\ \mathcal{A}_{\text{level}_k} &= \{0, 1, 2, 3\} \\ \mathcal{A}_{\text{cell}} &= \mathcal{A}_{\text{elem}} \times \mathcal{A}_{\text{ord}} \times \mathcal{A}_{\text{addr}}^4 \times \mathcal{A}_{\text{thid}}^4 \\ \mathcal{A}_{\text{mem}} &= \{m : \mathcal{A}_{\text{addr}} \rightarrow \mathcal{A}_{\text{cell}}\} \end{aligned}$$

where

$$\begin{aligned} \text{null}^{\mathcal{A}} &= 0x00 \\ \text{error}^{\mathcal{A}} &= \langle -\infty, -\infty, \text{null}, \text{null}, \text{null}, \text{null}, \emptyset, \emptyset, \emptyset, \emptyset \rangle \\ m(0x00) &= \langle -\infty, -\infty, \text{null}, \text{null}, \text{null}, \text{null}, \emptyset, \emptyset, \emptyset, \emptyset \rangle \\ m(0x01) &= \langle -\infty, -\infty, 0x02, 0x02, 0x03, 0x08, \emptyset, \emptyset, \emptyset, \emptyset \rangle \\ m(0x02) &= \langle 2, 2, 0x03, 0x03, \text{null}, \text{null}, \emptyset, \emptyset, \emptyset, \emptyset \rangle \\ m(0x03) &= \langle 8, 8, 0x04, 0x05, 0x06, \text{null}, \emptyset, T_1, T_1, \emptyset \rangle \\ m(0x04) &= \langle 11, 11, 0x05, \text{null}, \text{null}, \text{null}, T_1, \emptyset, \emptyset, \emptyset \rangle \\ m(0x05) &= \langle 14, 14, 0x06, 0x06, \text{null}, \text{null}, T_1, T_1, \emptyset, \emptyset \rangle \\ m(0x06) &= \langle 16, 16, 0x07, 0x08, 0x08, \text{null}, \emptyset, \emptyset, T_1, \emptyset \rangle \\ m(0x07) &= \langle 20, 20, 0x08, \text{null}, \text{null}, \text{null}, \emptyset, \emptyset, \emptyset, \emptyset \rangle \\ m(0x08) &= \langle +\infty, +\infty, \text{null}, \text{null}, \text{null}, \text{null}, \emptyset, \emptyset, \emptyset, \emptyset \rangle \end{aligned}$$

and where for instance the masked region depicted in the figure is described by the set

$$\{(0x03, 1), (0x03, 2), (0x04, 0), (0x05, 0), (0x05, 1), (0x06, 2)\} \quad *$$

Once more, we are interested in analyzing the satisfiability of quantifier-free first order formulas. Hence, if φ is a formula, we first proceed to write it into its disjunctive normal form, let's say $\varphi_1 \vee \dots \vee \varphi_n$. Then, we just verify the satisfiability of any of the φ_i , where each φ_i is a conjunction of TSL_K literals. We now classify the TSL_K literals into normalized and non-normalized ones. Non-normalized literals have the property that they can be written in terms of normalized ones. We now define the set of normalized TSL_K-literals.

Definition 5.1 (TSL_K-normalized literals).

A TSL_K-literal is normalized if it is a flat literal of the form:

$e_1 \neq e_2$	$a_1 \neq a_2$	$l_1 \neq l_2$
$a = \text{null}$	$c = \text{error}$	$c = \text{rd}(m, a)$
$k_1 \neq k_2$	$k_1 \preceq k_2$	$m_2 = \text{upd}(m_1, a, c)$
$c = \text{mkcell}(e, k, a_0, \dots, a_{K-1}, t_0, \dots, t_{K-1})$		$l_1 < l_2$
$s = \{a\}$	$s_1 = s_2 \cup s_3$	$s_1 = s_2 \setminus s_3$
$g = \{t\}_{\text{T}}$	$g_1 = g_2 \cup_{\text{T}} g_3$	$g_1 = g_2 \setminus_{\text{T}} g_3$
$r = \langle a, l \rangle_{\text{MR}}$	$r_1 = r_2 \cup_{\text{MR}} r_3$	$r_1 = r_2 \text{---}_{\text{MR}} r_3$
$a = x[t]_{\text{A}}$	$x_{\text{new}} = x\{t \leftarrow a\}_{\text{A}}$	
$l = z[t]_{\text{L}}$	$z_{\text{new}} = z\{t \leftarrow l\}_{\text{L}}$	
$p_1 \neq p_2$	$p = [a]$	$p_1 = \text{rev}(p_2)$
$s = \text{path2set}(p)$	$\text{append}(p_1, p_2, p_3)$	$\neg \text{append}(p_1, p_2, p_3)$
$s = \text{addr2set}_K(m, a, l)$	$p = \text{getp}_K(m, a_1, a_2, l)$	
$t_1 \neq t_2$	$a = \text{fstlock}(m, p, l)$	$\text{ordList}(m, p)$

where e, e_1 and e_2 are elem-variables; $a, a_0, a_1, a_2, \dots, a_{K-1}$ are addr-variables; c is a cell-variable; m, m_1 and m_2 are mem-variables; p, p_1, p_2 and p_3 are path-variables; s, s_1, s_2 and s_3 are set-variables; g, g_1, g_2 and g_3 are settim-variables; r, r_1, r_2 and r_3 are mrgn-variables; x and x_{new} are aarr-variables; z and z_{new} are Σ_{Iarr} -variables; k, k_1 and k_2 are ord-variables; l, l_1 and l_2 are level_K-variables and $t, t_0, t_1, t_2, \dots, t_{K-1}$ are thid-variables. †

The remaining literals can be rewritten from the normalized ones using the following equivalences:

$e = c.\text{data}$	$\leftrightarrow (\exists_{\text{ord}} k \exists_{\text{addr}} a_0, \dots, a_{K-1} \exists_{\text{thid}} t_0, \dots, t_{K-1})$ $[c = \text{mkcell}(e, k, a_0, \dots, a_{K-1}, t_0, \dots, t_{K-1})]$	
$k = c.\text{key}$	$\leftrightarrow (\exists_{\text{elem}} e \exists_{\text{addr}} a_0, \dots, a_{K-1} \exists_{\text{thid}} t_0, \dots, t_{K-1})$ $[c = \text{mkcell}(e, k, a_0, \dots, a_{K-1}, t_0, \dots, t_{K-1})]$	
$a = c.\text{next}[l]$	$\leftrightarrow (\exists_{\text{elem}} e \exists_{\text{ord}} k \exists_{\text{addr}} a_0, \dots, a_{l-1}, a_{l+1}, \dots, a_{K-1} \exists_{\text{thid}} t_0, \dots, t_{K-1})$ $[c = \text{mkcell}(e, k, a_0, \dots, a_{l-1}, a, a_{l+1}, \dots, a_{K-1}, t_0, \dots, t_{K-1})]$	
$t = c.\text{lockid}[l]$	$\leftrightarrow (\exists_{\text{elem}} e \exists_{\text{ord}} k \exists_{\text{addr}} a_0, \dots, a_{K-1} \exists_{\text{thid}} t_0, \dots, t_{l-1}, t_{l+1}, \dots, t_{K-1})$ $[c = \text{mkcell}(e, k, a_0, \dots, a_{K-1}, t_0, \dots, t_{l-1}, t, t_{l+1}, \dots, t_{K-1})]$	
$c_1 = c_2.\text{lock}[l](t)$	$\leftrightarrow c_2.\text{data} = c_1.\text{data} \wedge c_2.\text{key} = c_1.\text{key}$	\wedge
	$c_2.\text{next}[0] = c_1.\text{next}[0] \wedge \dots \wedge c_2.\text{next}[K-1] = c_1.\text{next}[K-1]$	\wedge
	$c_2.\text{lockid}[0] = c_1.\text{lockid}[0] \wedge \dots \wedge c_2.\text{lockid}[l-1] = c_1.\text{lockid}[l-1]$	\wedge
	$c_2.\text{lockid}[l+1] = c_1.\text{lockid}[l+1] \wedge \dots \wedge c_2.\text{lockid}[K-1] = c_1.\text{lockid}[K-1]$	\wedge
	$t = c_1.\text{lockid}[l]$	
$c_1 = c_2.\text{unlock}(l)$	$\leftrightarrow c_2.\text{data} = c_1.\text{data} \wedge c_2.\text{key} = c_1.\text{key}$	\wedge
	$c_2.\text{next}[0] = c_1.\text{next}[0] \wedge \dots \wedge c_2.\text{next}[K-1] = c_1.\text{next}[K-1]$	\wedge
	$c_2.\text{lockid}[0] = c_1.\text{lockid}[0] \wedge \dots \wedge c_2.\text{lockid}[l-1] = c_1.\text{lockid}[l-1]$	\wedge
	$c_2.\text{lockid}[l+1] = c_1.\text{lockid}[l+1] \wedge \dots \wedge c_2.\text{lockid}[K-1] = c_1.\text{lockid}[K-1]$	\wedge
	$\odot = c_1.\text{lockid}[l]$	
$c_1 \neq_{\text{cell}} c_2$	$\leftrightarrow c_1.\text{data} \neq c_2.\text{data} \vee c_1.\text{key} \neq c_2.\text{key}$	\vee
	$c_1.\text{next}[0] \neq c_2.\text{next}[0] \vee \dots \vee c_1.\text{next}[K-1] \neq c_2.\text{next}[K-1]$	\vee
	$c_1.\text{lockid}[0] \neq c_2.\text{lockid}[0] \vee \dots \vee c_1.\text{lockid}[K-1] \neq c_2.\text{lockid}[K-1]$	
$m_1 \neq_{\text{mem}} m_2$	$\leftrightarrow (\exists_{\text{addr}} a) [\text{rd}(m_1, a) \neq \text{rd}(m_2, a)]$	

$s_1 \neq s_2$	$\leftrightarrow (\exists_{\text{addr}a} [a \in (s_1 \setminus s_2) \cup (s_2 \setminus s_1)])$
$s = \emptyset$	$\leftrightarrow s = s \setminus s$
$s_3 = s_1 \cap s_2$	$\leftrightarrow s_3 = (s_1 \cup s_2) \setminus ((s_1 \setminus s_2) \cup (s_2 \setminus s_1))$
$a \in s$	$\leftrightarrow \{a\} \subseteq s$
$s_1 \subseteq s_2$	$\leftrightarrow s_2 = s_1 \cup s_2$
$g_1 \neq_{\text{setid}} g_2$	$\leftrightarrow (\exists_{\text{thid}t} [t \in (g_1 \setminus_{\text{T}} g_2) \cup_{\text{T}} (g_2 \setminus_{\text{T}} g_1)])$
$g = \emptyset_{\text{T}}$	$\leftrightarrow g = g \setminus_{\text{T}} g$
$g_3 = g_1 \cap_{\text{T}} g_2$	$\leftrightarrow g_3 = (g_1 \cup_{\text{T}} g_2) \setminus_{\text{T}} ((g_1 \setminus_{\text{T}} g_2) \cup_{\text{T}} (g_2 \setminus_{\text{T}} g_1))$
$t \in_{\text{T}} g$	$\leftrightarrow \{t\}_{\text{T}} \subseteq_{\text{T}} g$
$g_1 \subseteq_{\text{T}} g_2$	$\leftrightarrow g_2 = g_1 \cup_{\text{T}} g_2$
$r_1 \neq_{\text{setid}} r_2$	$\leftrightarrow (\exists_{\text{addr}a} \exists_{\text{level}l} [(a, l) \in (r_1 -_{\text{MR}} r_2) \cup_{\text{MR}} (r_2 -_{\text{MR}} r_1)])$
$r = \mathbf{emp}_{\text{MR}}$	$\leftrightarrow r = r -_{\text{MR}} r$
$r_3 = r_1 \cap_{\text{MR}} r_2$	$\leftrightarrow r_3 = (r_1 \cup_{\text{MR}} r_2) -_{\text{MR}} ((r_1 -_{\text{MR}} r_2) \cup_{\text{MR}} (r_2 -_{\text{MR}} r_1))$
$(a, l) \in_{\text{MR}} r$	$\leftrightarrow \langle a, l \rangle_{\text{MR}} \subseteq_{\text{MR}} r$
$r_1 \subseteq_{\text{MR}} r_2$	$\leftrightarrow r_2 = r_1 \cup_{\text{MR}} r_2$
$r_1 \#_{\text{MR}} r_2$	$\leftrightarrow \mathbf{emp}_{\text{MR}} = (r_1 \cup_{\text{MR}} r_2) -_{\text{MR}} ((r_1 -_{\text{MR}} r_2) \cup_{\text{MR}} (r_2 -_{\text{MR}} r_1))$
$p = \epsilon$	$\leftrightarrow \text{append}(p, p, p)$
$\text{reach}_K(m, a_1, a_2, l, p)$	$\leftrightarrow a_2 \in \text{addr2set}_K(m, a_1, l) \wedge p = \text{getp}_K(m, a_1, a_2, l)$

This means that we can rewrite such literals using:

Flat:	$e = c.\text{data}$
Normalized:	$c = \text{mkcell}(e, k, a_0, \dots, a_{K-1}, t_0, \dots, t_{K-1})$
Proviso:	$k, a_0, \dots, a_{K-1}, t_0, \dots, t_{K-1}$ are fresh.
Flat:	$k = c.\text{key}$
Normalized:	$c = \text{mkcell}(e, k, a_0, \dots, a_{K-1}, t_0, \dots, t_{K-1})$
Proviso:	$e, a_0, \dots, a_{K-1}, t_0, \dots, t_{K-1}$ are fresh.
Flat:	$a = c.\text{next}[l]$
Normalized:	$c = \text{mkcell}(e, k, a_0, \dots, a_{l-1}, a, a_{l+1}, \dots, a_{K-1}, t_0, \dots, t_{K-1})$
Proviso:	$e, k, a_0, \dots, a_{l-1}, a_{l+1}, a_{K-1}, t_0, \dots, t_{K-1}$ are fresh.
Flat:	$t = c.\text{lockid}[l]$
Normalized:	$c = \text{mkcell}(e, k, a_0, \dots, a_{K-1}, t_0, \dots, t_{l-1}, t, t_{l+1}, \dots, t_{K-1})$
Proviso:	$e, k, a_0, \dots, a_{K-1}, t_0, \dots, t_{l-1}, t_{l+1}, \dots, t_{K-1}$ are fresh.
Flat:	$c_1 = c_2.\text{lock}[l](t)$
Normalized:	$c_1 = \text{mkcell}(e, k, a_0, \dots, a_{K-1}, t_0, \dots, t_{l-1}, t, t_{l+1}, \dots, t_{K-1}) \wedge$ $c_2 = \text{mkcell}(e, k, a_0, \dots, a_{K-1}, t_0, \dots, t_{l-1}, \tilde{t}, t_{l+1}, \dots, t_{K-1})$
Proviso:	$e, k, a_0, \dots, a_{K-1}, t_0, \dots, t_{l-1}, \tilde{t}, t_{l+1}, \dots, t_{K-1}$ are fresh.
Flat:	$c_1 = c_2.\text{unlock}[l]$
Normalized:	$c_1 = \text{mkcell}(e, k, a_0, \dots, a_{K-1}, t_0, \dots, t_{l-1}, \emptyset, t_{l+1}, \dots, t_{K-1}) \wedge$ $c_2 = \text{mkcell}(e, k, a_0, \dots, a_{K-1}, t_0, \dots, t_{l-1}, \tilde{t}, t_{l+1}, \dots, t_{K-1})$
Proviso:	$e, k, a_0, \dots, a_{K-1}, t_0, \dots, t_{l-1}, \tilde{t}, t_{l+1}, \dots, t_{K-1}$ are fresh.
Flat:	$c_1 \neq c_2$
Normalized:	$c_1.\text{data} \neq c_2.\text{data} \vee c_1.\text{key} \neq c_2.\text{key} \vee$ $c_1.\text{next}[0] \neq c_2.\text{next}[0] \vee \dots \vee c_1.\text{next}[K-1] \neq c_2.\text{next}[K-1] \vee$ $c_1.\text{lockid}[0] \neq c_2.\text{lockid}[0] \vee \dots \vee c_1.\text{lockid}[K-1] \neq c_2.\text{lockid}[K-1]$
Proviso:	-

Flat:	$m_1 \neq m_2$
Normalized:	$m[a] \neq m[a]$
Proviso:	a is fresh.
Flat:	$s_1 \neq s_2$
Normalized:	$s_{12} = s_1 \setminus s_2 \wedge s_{21} = s_2 \setminus s_1 \wedge s_3 = s_{12} \cup s_{21} \wedge s = s_3 \cup \{a\} \wedge \{a\} \subseteq s$
Proviso:	s_{12}, s_{21}, s_3, s and a are fresh.
Flat:	$s = \emptyset$
Normalized:	$s = s \setminus s$
Proviso:	-
Flat:	$s_3 = s_1 \cap s_2$
Normalized:	$s_{12} = s_1 \setminus s_2 \wedge s_{21} = s_2 \setminus s_1 \wedge s_{u_1} = s_1 \cup s_2 \wedge s_{u_2} = s_{12} \cup s_{21} \wedge s_3 = s_{u_1} \setminus s_{u_2}$
Proviso:	s_{12}, s_{21}, s_{u_1} and s_{u_2} are fresh.
Flat:	$a \in s$
Normalized:	$s = \{a\} \cup s$
Proviso:	-
Flat:	$s_1 \subseteq s_2$
Normalized:	$s_2 = s_1 \cup s_2$
Proviso:	-
Flat:	$g_1 \neq g_2$
Normalized:	$g_{12} = g_1 \setminus g_2 \wedge g_{21} = g_2 \setminus g_1 \wedge g_3 = g_{12} \cup g_{21} \wedge g = g_3 \cup \{t\} \wedge \{t\} \subseteq g$
Proviso:	g_{12}, g_{21}, g_3, g and t are fresh.
Flat:	$g = \emptyset_{\tau}$
Normalized:	$g = g \setminus_{\tau} g$
Proviso:	-
Flat:	$g_3 = g_1 \cap_{\tau} g_2$
Normalized:	$g_{12} = g_1 \setminus_{\tau} g_2 \wedge g_{21} = g_2 \setminus_{\tau} g_1 \wedge g_{u_1} = g_1 \cup_{\tau} g_2 \wedge g_{u_2} = g_{12} \cup_{\tau} g_{21} \wedge g_3 = g_{u_1} \setminus_{\tau} g_{u_2}$
Proviso:	g_{12}, g_{21}, g_{u_1} and g_{u_2} are fresh.
Flat:	$t \in_{\tau} g$
Normalized:	$g = \{t\} \cup_{\tau} g$
Proviso:	-
Flat:	$g_1 \subseteq_{\tau} g_2$
Normalized:	$g_2 = g_1 \cup_{\tau} g_2$
Proviso:	-
Flat:	$r_1 \neq r_2$
Normalized:	$r_{12} = r_1 \text{ }_{\text{MR}} \text{ } r_2 \wedge r_{21} = r_2 \text{ }_{\text{MR}} \text{ } r_1 \wedge r_3 = r_{12} \cup_{\text{MR}} r_{21} \wedge r = r_3 \cup_{\text{MR}} \{(a, l)\} \wedge \{(a, l)\} \subseteq_{\text{MR}} r$
Proviso:	$r_{12}, r_{21}, r_3, r, a$ and l are fresh.
Flat:	$r = \mathbf{emp}_{\text{MR}}$
Normalized:	$r = r \text{ }_{\text{MR}} \text{ } r$
Proviso:	-
Flat:	$r_3 = r_1 \cap_{\text{MR}} r_2$
Normalized:	$r_{12} = r_1 \text{ }_{\text{MR}} \text{ } r_2 \wedge r_{21} = r_2 \text{ }_{\text{MR}} \text{ } r_1 \wedge r_{u_1} = r_1 \cup_{\text{MR}} r_2 \wedge r_{u_2} = r_{12} \cup_{\text{MR}} r_{21} \wedge r_3 = r_{u_1} \text{ }_{\text{MR}} \text{ } r_{u_2}$
Proviso:	r_{12}, r_{21}, r_{u_1} and r_{u_2} are fresh.
Flat:	$(a, l) \in_{\text{MR}} r$
Normalized:	$r = \{(a, l)\} \cup_{\text{MR}} r$
Proviso:	-
Flat:	$r_1 \subseteq_{\text{MR}} r_2$
Normalized:	$r_2 = r_1 \cup_{\text{MR}} r_2$
Proviso:	-

Flat:	$r_1 \#_{MR} r_2$
Normalized:	$r_{12} = r_1 \text{---}_{MR} r_2 \wedge r_{21} = r_2 \text{---}_{MR} r_1 \wedge r_{u_1} = r_1 \cup_{MR} r_2 \wedge$ $r_{u_2} = r_{12} \cup_{MR} r_{21} \wedge r_3 = r_{u_1} \text{---}_{MR} r_{u_2} \wedge r_3 = r_3 \text{---}_{MR} r_3$
Proviso:	$r_{12}, r_{21}, r_{u_1}, r_{u_2}$ and r_3 are fresh.
Flat:	$p = \epsilon$
Normalized:	$append(p, p, p)$
Proviso:	-
Flat:	$reach_K(m, a_1, a_2, l, p)$
Normalized:	$a_2 \in addr2set_K(m, a_1, l) \wedge p = getp_K(m, a_1, a_2, l)$
Proviso:	-

5.1 Decidability of TSL_K

We show that TSL_K is decidable by proving that it enjoys the finite model property with respect to its sorts, and exhibiting upper bounds for the sizes of the domains of a small interpretation of a satisfiable formula.

Definition 5.2 (Finite Model Property).

Let Σ be a signature, $S_0 \subseteq S$ be a set of sorts, and T be a Σ -theory. T has the finite model property with respect to S_0 if for every T -satisfiable quantifier-free Σ -formula φ there exists a T -interpretation \mathcal{A} satisfying φ such that for each sort $\sigma \in S_0$, \mathcal{A}_σ is finite. †

The fact that TSL_K has the finite model property with respect to domains $elem$, $addr$, ord , $level_K$ and $thid$, implies that TSL_K is decidable by enumerating all possible Σ_{TSL_K} -structures up to a certain cardinality. Notice that a bound on the domain of these sorts it is enough to get finite interpretations for the remaining sorts ($cell$, mem , $path$, set , $settid$, $mrgn$, arr and $larr$) as the elements in the domains of these latter sorts are constructed using the elements in the domains of $elem$, $addr$, ord , $level_K$ and $thid$.

Lemma 5.1:

Deciding the TSL_K -satisfiability of a quantifier-free TSL_K -formula is equivalent to verifying the TSL_K -satisfiability of the normalized TSL_K -literals. ♠

Proof First, transform a formula in disjunctive normal form. Then each conjunct can be normalized introducing auxiliary fresh variables when necessary. For instance, consider the non-normalized literal $c_1 = c_2.lock[l](t)$. According to the table we presented before, this literal can be written as a conjunction of normalized literals of the form

$$c_1 = mkcell(e, k, \vec{a}, t_0, \dots, t_{l-1}, t, t_{l+1}, \dots, t_{K-1}) \wedge c_2 = mkcell(e, k, \vec{a}, t_0, \dots, t_{K-1})$$

where $e, k, \vec{a} = a_0 \dots a_{K-1}, t_0, \dots, t_{K-1}$ and t are fresh variables. Let \mathcal{A} be a model satisfying the literal $c_1 = c_2.lock[l](t)$. Then, we show that

$$\mathcal{A} \models c_1 = c_2.lock[l](t)$$

iff

$$\mathcal{A} \models c_1 = mkcell(e, k, \vec{a}, t_0, \dots, t_{l-1}, t, t_{l+1}, \dots, t_{K-1}) \wedge c_2 = mkcell(e, k, \vec{a}, t_0, \dots, t_{K-1})$$

The proof is as follows:

$$\begin{aligned}
\mathcal{A} \models c_1 &= c_2.lock[l](t) && \leftrightarrow \\
c_1^{\mathcal{A}} &= c_2^{\mathcal{A}}.lock^{\mathcal{A}}[l^{\mathcal{A}}](t^{\mathcal{A}}) && \leftrightarrow \\
c_1^{\mathcal{A}} &= c_2^{\mathcal{A}}.lock^{\mathcal{A}}[l^{\mathcal{A}}](t^{\mathcal{A}}) \wedge c_2^{\mathcal{A}} = \langle e^{\mathcal{A}}, k^{\mathcal{A}}, \vec{d}^{\mathcal{A}}, t_0^{\mathcal{A}}, \dots, t_{K-1}^{\mathcal{A}} \rangle && \leftrightarrow \\
c_1^{\mathcal{A}} &= \langle e^{\mathcal{A}}, k^{\mathcal{A}}, \vec{d}^{\mathcal{A}}, t_0^{\mathcal{A}}, \dots, t_{l-1}^{\mathcal{A}}, t^{\mathcal{A}}, t_{l+1}^{\mathcal{A}}, \dots, t_{K-1}^{\mathcal{A}} \rangle && \wedge \\
c_2^{\mathcal{A}} &= \langle e^{\mathcal{A}}, k^{\mathcal{A}}, \vec{d}^{\mathcal{A}}, t_0^{\mathcal{A}}, \dots, t_{K-1}^{\mathcal{A}} \rangle && \leftrightarrow \\
c_1^{\mathcal{A}} &= mkcell^{\mathcal{A}}(e^{\mathcal{A}}, k^{\mathcal{A}}, \vec{d}^{\mathcal{A}}, t_0^{\mathcal{A}}, \dots, t_{l-1}^{\mathcal{A}}, t^{\mathcal{A}}, t_{l+1}^{\mathcal{A}}, \dots, t_{K-1}^{\mathcal{A}}) && \wedge \\
c_2^{\mathcal{A}} &= mkcell^{\mathcal{A}}(e^{\mathcal{A}}, k^{\mathcal{A}}, \vec{d}^{\mathcal{A}}, t_0^{\mathcal{A}}, \dots, t_{K-1}^{\mathcal{A}}) && \leftrightarrow \\
\mathcal{A} \models c_1 &= mkcell(e, k, \vec{d}, t_0, \dots, t_{l-1}, t, t_{l+1}, \dots, t_{K-1}) && \wedge \\
c_2 &= mkcell(e, k, \vec{d}, t_0, \dots, t_{K-1}) &&
\end{aligned}$$

The remaining cases can be proved in a similar way. \square

The phase of normalizing a formula is commonly known [RRZ05] as the ‘‘variable abstraction phase’’. Note that each normalized literal belongs to just one of the theories that are part of TSL_K .

Consider an arbitrary TSL_K -interpretation \mathcal{A} satisfying a conjunction of normalized TSL_K -literals Γ . We show that if \mathcal{A} consists of domains \mathcal{A}_{elem} , \mathcal{A}_{addr} , \mathcal{A}_{thid} , \mathcal{A}_{level_k} and \mathcal{A}_{ord} then there are finite sets \mathcal{B}_{elem} , \mathcal{B}_{addr} , \mathcal{B}_{thid} , \mathcal{B}_{level_k} and \mathcal{B}_{ord} with bounded cardinalities, where the finite bound on the sizes can be computed from Γ . Such sets can in turn be used to obtain a finite interpretation \mathcal{B} satisfying Γ , since all the other sorts are bounded by the sizes of these sets.

Before proving that TSL_K enjoys the finite model property, we define some auxiliary functions. We start by defining the function

$$first_K : mem \times addr \times level_k \times set \rightarrow addr$$

Let $\mathcal{B}_{addr} \subseteq \mathcal{A}_{addr}$, $m : \mathcal{A}_{addr} \rightarrow \mathcal{B}_{elem} \times \mathcal{B}_{ord} \times \mathcal{A}_{addr}^K \times \mathcal{B}_{thid}^K$, $a \in \mathcal{B}_{addr}$ and $l \in \mathcal{B}_{level_k}$. The function $first_K(m, a, l, \mathcal{B}_{addr})$ is defined by

$$first_K(m, a, l, \mathcal{B}_{addr}) = \begin{cases} null & \text{if for all } r \geq 1 \ m^r(a).next(l) \notin \mathcal{B}_{addr} \\ m^s(a).next(l) & \text{if for some } s \geq 1 \ m^s(a).next(l) \in \mathcal{B}_{addr}, \\ & \text{and for all } r < s \ m^r(a).next(l) \notin \mathcal{B}_{addr} \end{cases}$$

where

- $m^1(a).next(l)$ stands for $m(a).next(l)$, and
- $m^{n+1}(a).next(l)$ stands for $m(m^n(a).next(l)).next(l)$ when $n > 0$.

Basically, given the original model \mathcal{A} and a subset of addresses $\mathcal{B}_{addr} \subseteq \mathcal{A}_{addr}$, function $first_K$ chooses the next address in \mathcal{B}_{addr} that can be reached from a given address following repeatedly the $next(l)$ pointer. It is easy to see, for example, that if $m(a).next(l) \in \mathcal{B}_{addr}$ then $first_K(m, a, l, \mathcal{B}_{addr}) = m(a).next(l)$. We will later filter out unnecessary intermediate nodes and use $first_K$ to bypass properly the removed nodes, preserving the important connectivity properties.

Lemma 5.2:

Let $\mathcal{B}_{addr} \subseteq \mathcal{A}_{addr}$, $m : \mathcal{A}_{addr} \rightarrow \mathcal{B}_{elem} \times \mathcal{B}_{ord} \times \mathcal{A}_{addr}^K \times \mathcal{B}_{thid}^K$, $a \in \mathcal{B}_{addr}$ and $l \in \mathcal{B}_{level_k}$. Then, if $m(a).next(l) \in \mathcal{B}_{addr}$, we have that $first_K(m, a, l, \mathcal{B}_{addr}) = m(a).next(l)$. \spadesuit

Proof Immediate from the definition of $first_K$. \square

Secondly, the function $unordered$ that given a memory m and a path p , returns a set containing two addresses that witness the failure to preserve the *key* order of elements in p :

$$unordered(m, [i_1, \dots, i_n]) = \begin{cases} \emptyset & \text{if } n = 0 \text{ or } n = 1 \\ \{i_1, i_2\} & \text{if } m(i_2).key \preceq m(i_1).key \text{ and} \\ & m(i_2).key \neq m(i_1).key \\ unordered(m, [i_2, \dots, i_n]) & \text{otherwise} \end{cases}$$

If two such addresses exist, $unordered$ returns the first two consecutive addresses whose keys violate the order.

Lemma 5.3:

Let p be a path such that $p = [a_1, \dots, a_n]$ with $n \geq 2$ and let m be a memory. If it exists a_i , with $1 \leq i < n$, such that $m(a_{i+1}).key \preceq m(a_i).key$ and $m(a_{i+1}).key \neq m(a_i).key$, then $unordered(m, p) \neq \emptyset$ \spadesuit

Proof By induction. Let's consider $n = 2$ and let $p = [a_1, a_2]$ such that $m(a_2).key \preceq m(a_1).key$ and $m(a_2).key \neq m(a_1).key$. Then, by definition of $unordered$, we have that $unordered(m, p) = \{a_1, a_2\} \neq \emptyset$.

Now assume $n > 2$ and let $p = [a_1, \dots, a_n]$. If $m(a_2).key \preceq m(a_1).key$ and $m(a_2).key \neq m(a_1).key$, then we have that $unordered(m, p) = \{a_1, a_2\} \neq \emptyset$. On the other hand, if $m(a_1).key \preceq m(a_2).key$, we still know that there is a_i , with $2 \leq i < n$, s.t., $m(a_{i+1}).key \preceq m(a_i).key$ and $m(a_{i+1}).key \neq m(a_i).key$. Therefore, by induction we have that $unordered(m, [a_2, \dots, a_n]) \neq \emptyset$ and by definition of $unordered$, $unordered(m, p) = unordered(m, [a_2, \dots, a_n]) \neq \emptyset$. \square

Finally, we define the function $knownTID$ that given a thread identifier and a set of threads, returns the same thread id if it belongs to the given set. Otherwise, it returns \emptyset :

$$knownTID(t, T) = \begin{cases} t & \text{if } t \in T \\ \emptyset & \text{otherwise} \end{cases}$$

In this section, we also make use of the functions $compress$, $diseq$ and $common$ defined in Section 4.1. We now show that given a conjunction of TSL_K literals, if there exists a model satisfying such formula, then it is possible to construct a new model with a finite number of elements satisfying the same formula. This property is described in the following Lemma.

Lemma 5.4 (Finite Model Property):

Let Γ be a conjunction of normalized TSL_K -literals. Let $\bar{e} = |V_{elem}(\Gamma)|$, $\bar{a} = |V_{addr}(\Gamma)|$, $\bar{m} = |V_{mem}(\Gamma)|$, $\bar{p} = |V_{path}(\Gamma)|$, $\bar{t} = |V_{thid}(\Gamma)|$, $\bar{o} = |V_{ord}(\Gamma)|$ and $\bar{x} = |V_{aarr}(\Gamma)|$. Then the following are equivalent:

1. Γ is TSL_K -satisfiable;
2. Γ is true in a TSL_K interpretation \mathcal{B} such that

$$\begin{aligned} |\mathcal{B}_{level_K}| &\leq K \\ |\mathcal{B}_{thid}| &\leq \bar{t} + 1 + K \bar{m} \bar{a} \\ |\mathcal{B}_{addr}| &\leq \bar{a} + 1 + K \bar{m} \bar{a} + \bar{p}^2 + \bar{p}^3 + 2 \bar{m} \bar{p} + \bar{x} |\mathcal{B}_{thid}| \\ |\mathcal{B}_{elem}| &\leq \bar{e} + \bar{m} |\mathcal{B}_{addr}| \\ |\mathcal{B}_{ord}| &\leq \bar{o} + \bar{m} |\mathcal{B}_{addr}| \end{aligned} \quad \spadesuit$$

Proof (2 \rightarrow 1) is immediate.

(1 \rightarrow 2). We prove this implication only for the new TSL_K -literals since the ones shared with theory TLL3 have already been proved in Lemma 4.3.

Bearing in mind the auxiliary functions we have defined, let \mathcal{A} be a TSL_K -interpretation satisfying a set of normalized TSL_K -literals Γ . We use \mathcal{A} to construct a finite TSL_K -interpretation \mathcal{B} that satisfies Γ .

$$\begin{aligned}
\mathcal{B}_{\text{level}_k} &= \mathcal{A}_{\text{level}_k} = [0 \dots K - 1] \\
\mathcal{B}_{\text{thid}} &= V_{\text{thid}}^{\mathcal{A}} \cup \{ \emptyset \} \cup \{ m^{\mathcal{A}}(a^{\mathcal{A}}).lockid^{\mathcal{A}}(l) \mid m \in V_{\text{mem}}, a \in V_{\text{addr}} \text{ and } l \in \mathcal{B}_{\text{level}_k} \} \\
\mathcal{B}_{\text{addr}} &= V_{\text{addr}}^{\mathcal{A}} \cup \{ null^{\mathcal{A}} \} && \cup \\
&\quad \{ m^{\mathcal{A}}(a^{\mathcal{A}}).next^{\mathcal{A}}(l) \mid m \in V_{\text{mem}}, a \in V_{\text{addr}} \text{ and } l \in \mathcal{B}_{\text{level}_k} \} && \cup \\
&\quad \{ v \in \text{diseq}(p^{\mathcal{A}}, q^{\mathcal{A}}) \mid \text{the literal } p \neq q \text{ is in } \Gamma \} && \cup \\
&\quad \{ v \in \text{common}(p_1^{\mathcal{A}}, p_2^{\mathcal{A}}) \mid \text{the literal } \neg \text{append}(p_1, p_2, p_3) \text{ is in } \Gamma \text{ and} \\
&\quad \quad \quad \text{path2set}^{\mathcal{A}}(p_1^{\mathcal{A}}) \cap \text{path2set}^{\mathcal{A}}(p_2^{\mathcal{A}}) \neq \emptyset \} && \cup \\
&\quad \{ v \in \text{common}(p_1^{\mathcal{A}} \circ p_2^{\mathcal{A}}, p_3^{\mathcal{A}}) \mid \text{the literal } \neg \text{append}(p_1, p_2, p_3) \text{ is in } \Gamma \text{ and} \\
&\quad \quad \quad \text{path2set}^{\mathcal{A}}(p_1^{\mathcal{A}}) \cap \text{path2set}^{\mathcal{A}}(p_2^{\mathcal{A}}) = \emptyset \} && \cup \\
&\quad \{ v \in \text{unordered}(m^{\mathcal{A}}, p^{\mathcal{A}}) \mid \neg \text{ordList}(m, p) \text{ is in } \Gamma \} && \cup \\
&\quad \{ x^{\mathcal{A}}(i) \mid x \in V_{\text{arr}} \text{ and } i \in \mathcal{B}_{\text{thid}} \} \\
\mathcal{B}_{\text{elem}} &= V_{\text{elem}}^{\mathcal{A}} \cup \{ m^{\mathcal{A}}(v).data^{\mathcal{A}} \mid m \in V_{\text{mem}} \text{ and } v \in \mathcal{B}_{\text{addr}} \} \\
\mathcal{B}_{\text{ord}} &= V_{\text{ord}}^{\mathcal{A}} \cup \{ m^{\mathcal{A}}(v).key^{\mathcal{A}} \mid m \in V_{\text{mem}} \text{ and } v \in \mathcal{B}_{\text{addr}} \}
\end{aligned}$$

For each sort σ , $V_{\sigma}^{\mathcal{A}}$ denotes the set of values obtained by the interpretation of all variables of type σ appearing in Γ . Notice that as a finite number of variables can appear in Γ , $V_{\sigma}^{\mathcal{A}}$ is also a finite set.

The domains described above satisfy the cardinality constraints expressed in the statement of the theorem. The interpretations of the symbols are:

$$\begin{aligned}
error^{\mathcal{B}} &= error^{\mathcal{A}} \\
null^{\mathcal{B}} &= null^{\mathcal{A}} \\
e^{\mathcal{B}} &= e^{\mathcal{A}} && \text{for each } e \in V_{\text{elem}} \\
a^{\mathcal{B}} &= a^{\mathcal{A}} && \text{for each } a \in V_{\text{addr}} \\
c^{\mathcal{B}} &= c^{\mathcal{A}} && \text{for each } c \in V_{\text{cell}} \\
t^{\mathcal{B}} &= t^{\mathcal{A}} && \text{for each } t \in V_{\text{thid}} \\
k^{\mathcal{B}} &= k^{\mathcal{A}} && \text{for each } k \in V_{\text{ord}} \\
l^{\mathcal{B}} &= l^{\mathcal{A}} && \text{for each } l \in V_{\text{level}_k} \\
m^{\mathcal{B}}(v) &= \left(m^{\mathcal{A}}(v).data^{\mathcal{A}}, m^{\mathcal{A}}(v).key^{\mathcal{A}}, \right. && \text{for each } m \in V_{\text{mem}} \\
&\quad \text{first}_K(m^{\mathcal{A}}, v, 0, \mathcal{B}_{\text{addr}}), \dots, \text{first}_K(m^{\mathcal{A}}, v, K - 1, \mathcal{B}_{\text{addr}}), && \text{and } v \in \mathcal{B}_{\text{addr}} \\
&\quad \text{knownTID}(m^{\mathcal{A}}(v).lockid^{\mathcal{A}}(0), \mathcal{B}_{\text{thid}}), \\
&\quad \dots, \\
&\quad \left. \text{knownTID}(m^{\mathcal{A}}(v).lockid^{\mathcal{A}}(K - 1), \mathcal{B}_{\text{thid}}) \right) \\
s^{\mathcal{B}} &= s^{\mathcal{A}} \cap \mathcal{B}_{\text{addr}} && \text{for each } s \in V_{\text{set}} \\
g^{\mathcal{B}} &= g^{\mathcal{A}} \cap \mathcal{B}_{\text{thid}} && \text{for each } g \in V_{\text{settid}} \\
r^{\mathcal{B}} &= r^{\mathcal{A}} \cap (\mathcal{B}_{\text{addr}} \times \mathcal{B}_{\text{level}_k}) && \text{for each } r \in V_{\text{mrgn}} \\
x^{\mathcal{B}}(i) &= x^{\mathcal{A}}(i) && \text{for each } x \in V_{\text{arr}}, i \in \mathcal{B}_{\text{thid}} \\
z^{\mathcal{B}}(i) &= z^{\mathcal{A}}(i) && \text{for each } z \in V_{\text{larr}}, i \in \mathcal{B}_{\text{thid}} \\
p^{\mathcal{B}} &= \text{compress}(p^{\mathcal{A}}, \mathcal{B}_{\text{addr}}) && \text{for each } p \in V_{\text{path}}
\end{aligned}$$

Essentially, all variables and constants in \mathcal{B} are interpreted as in \mathcal{A} except that *next* pointers use *first_K* to point to the next reachable element that has been preserved in $\mathcal{B}_{\text{addr}}$, and paths filter out all elements except those in $\mathcal{B}_{\text{addr}}$. It can be routinely checked that \mathcal{B} is an interpretation of Γ . So it remains to be seen that \mathcal{B} satisfies all literals in Γ assuming that \mathcal{A} does, concluding that \mathcal{B} is indeed a model of Γ . This check is performed by cases. The proof that \mathcal{B} satisfies all TSL_K-literals in Γ is not shown here. We just focus on the new functions and predicates that are not part of TLL. The proof for literals of TLL can be found in [RZ06a] and for TLL3 in Chapter 4. For TSL_K-literals we must consider the following cases:

Literals of the form $l_1 \neq l_2, k_1 \neq k_2, l_1 < l_2$ and $k_1 \preceq k_2$: Immediate.

Literals of the form $c = \text{mkcell}(e, k, a_0, \dots, a_{K-1}, t_0, \dots, t_{K-1})$:

$$\begin{aligned}
c^{\mathcal{B}} &= c^{\mathcal{A}} \\
&= \text{mkcell}^{\mathcal{A}}(e, k, a_0, \dots, a_{K-1}, t_0, \dots, t_{K-1}) \\
&= \langle e^{\mathcal{A}}, k^{\mathcal{A}}, a_0^{\mathcal{A}}, \dots, a_{K-1}^{\mathcal{A}}, t_0^{\mathcal{A}}, \dots, t_{K-1}^{\mathcal{A}} \rangle \\
&= \langle e^{\mathcal{B}}, k^{\mathcal{B}}, a_0^{\mathcal{B}}, \dots, a_{K-1}^{\mathcal{B}}, t_0^{\mathcal{B}}, \dots, t_{K-1}^{\mathcal{B}} \rangle \\
&= \text{mkcell}^{\mathcal{B}}(e, k, a_0, \dots, a_{K-1}, t_0, \dots, t_{K-1})
\end{aligned}$$

Literals of the form $c = \text{rd}(m, a)$:

$$\begin{aligned}
[\text{rd}(m, a)]^{\mathcal{B}} &= m^{\mathcal{B}}(a^{\mathcal{B}}) \\
&= m^{\mathcal{B}}(a^{\mathcal{A}}) \\
&= \left(m^{\mathcal{A}}(a^{\mathcal{A}}).data^{\mathcal{A}}, m^{\mathcal{A}}(a^{\mathcal{A}}).key^{\mathcal{A}}, \right. \\
&\quad \left. \text{first}_{t_K}(m^{\mathcal{A}}, a^{\mathcal{A}}, 0, \mathcal{B}_{\text{addr}}), \dots, \text{first}_{t_K}(m^{\mathcal{A}}, a^{\mathcal{A}}, K-1, \mathcal{B}_{\text{addr}}), \right. \\
&\quad \left. \text{knownTID}(m^{\mathcal{A}}(v).lockid^{\mathcal{A}}(0), \mathcal{B}_{\text{thid}}), \dots, \text{knownTID}(m^{\mathcal{A}}(v).lockid^{\mathcal{A}}(K-1), \mathcal{B}_{\text{thid}}) \right) \\
&= \left(m^{\mathcal{A}}(a^{\mathcal{A}}).data^{\mathcal{A}}, m^{\mathcal{A}}(a^{\mathcal{A}}).key^{\mathcal{A}}, \right. \\
&\quad \left. m^{\mathcal{A}}(a^{\mathcal{A}}).next^{\mathcal{A}}(0), \dots, m^{\mathcal{A}}(a^{\mathcal{A}}).next^{\mathcal{A}}(K-1), \right. \\
&\quad \left. m^{\mathcal{A}}(a^{\mathcal{A}}).lockid^{\mathcal{A}}[0], \dots, m^{\mathcal{A}}(a^{\mathcal{A}}).lockid^{\mathcal{A}}[K-1] \right) \quad (\text{Lemma 5.2}) \\
&= m^{\mathcal{A}}(a^{\mathcal{A}}) \\
&= c^{\mathcal{A}} \\
&= c^{\mathcal{B}}
\end{aligned}$$

Literals of the form $g = \{t\}_{\text{T}}$:

$$g^{\mathcal{B}} = g^{\mathcal{A}} \cap \mathcal{B}_{\text{thid}} = \{t^{\mathcal{A}}\}_{\text{T}} \cap \mathcal{B}_{\text{thid}} = \{t^{\mathcal{B}}\}_{\text{T}} \cap \mathcal{B}_{\text{thid}} = \{t^{\mathcal{B}}\}_{\text{T}}$$

Literals of the form $g_1 = g_2 \cup_{\text{T}} g_3$:

$$\begin{aligned}
g_1^{\mathcal{B}} &= g_1^{\mathcal{A}} \cap \mathcal{B}_{\text{thid}} \\
&= (g_2^{\mathcal{A}} \cup_{\text{T}} g_3^{\mathcal{A}}) \cap \mathcal{B}_{\text{thid}} \\
&= (g_2^{\mathcal{A}} \cap \mathcal{B}_{\text{thid}}) \cup_{\text{T}} (g_3^{\mathcal{A}} \cap \mathcal{B}_{\text{thid}}) \\
&= g_2^{\mathcal{B}} \cup_{\text{T}} g_3^{\mathcal{B}}
\end{aligned}$$

Literals of the form $g_1 = g_2 \setminus_{\tau} g_3$:

$$\begin{aligned}
 g_1^B &= g_1^A \cap \mathcal{B}_{\text{thid}} \\
 &= (g_2^A \setminus_{\tau} g_3^A) \cap \mathcal{B}_{\text{thid}} \\
 &= (g_2^A \cap \mathcal{B}_{\text{thid}}) \setminus_{\tau} (g_3^A \cap \mathcal{B}_{\text{thid}}) \\
 &= g_2^B \setminus_{\tau} g_3^B
 \end{aligned}$$

Literals of the form $r = \langle a, l \rangle_{\text{MR}}$:

$$\begin{aligned}
 r^B &= r^A \cap (\mathcal{B}_{\text{addr}} \times \mathcal{B}_{\text{level}_k}) \\
 &= \langle a^A, l^A \rangle_{\text{MR}} \cap (\mathcal{B}_{\text{addr}} \times \mathcal{B}_{\text{level}_k}) \\
 &= \langle a^B, l^B \rangle_{\text{MR}} \cap (\mathcal{B}_{\text{addr}} \times \mathcal{B}_{\text{level}_k}) \\
 &= \langle a^B, l^B \rangle_{\text{MR}}
 \end{aligned}$$

Literals of the form $r_1 = r_2 \cup_{\text{MR}} r_3$:

$$\begin{aligned}
 r_1^B &= r_1^A \cap (\mathcal{B}_{\text{addr}} \times \mathcal{B}_{\text{level}_k}) \\
 &= (r_2^A \cup_{\text{MR}} r_3^A) \cap (\mathcal{B}_{\text{addr}} \times \mathcal{B}_{\text{level}_k}) \\
 &= (r_2^A \cap (\mathcal{B}_{\text{addr}} \times \mathcal{B}_{\text{level}_k})) \cup_{\text{MR}} (r_3^A \cap (\mathcal{B}_{\text{addr}} \times \mathcal{B}_{\text{level}_k})) \\
 &= r_2^B \cup_{\text{MR}} r_3^B
 \end{aligned}$$

Literals of the form $r_1 = r_2 \text{--}_{\text{MR}} r_3$:

$$\begin{aligned}
 r_1^B &= r_1^A \cap (\mathcal{B}_{\text{addr}} \times \mathcal{B}_{s\text{Level}_k}) \\
 &= (r_2^A \text{--}_{\text{MR}} r_3^A) \cap (\mathcal{B}_{\text{addr}} \times \mathcal{B}_{s\text{Level}_k}) \\
 &= (r_2^A \cap (\mathcal{B}_{\text{addr}} \times \mathcal{B}_{s\text{Level}_k})) \text{--}_{\text{MR}} (r_3^A \cap (\mathcal{B}_{\text{addr}} \times \mathcal{B}_{s\text{Level}_k})) \\
 &= r_2^B \text{--}_{\text{MR}} r_3^B
 \end{aligned}$$

Literals of the form $a = x[t]_{\Lambda}$:

$$\begin{aligned}
 x[t]_{\Lambda}^B &= x^B(t^B) \\
 &= x^A(t^B) \\
 &= x^A(t^A) \\
 &= x[t]_{\Lambda}^A \\
 &= a^A \\
 &= a^B
 \end{aligned}$$

Literals of the form $x_{\text{new}} = x\{t \leftarrow a\}_{\Lambda}$: In this case, since $x_{\text{new}}^A(t^A) = a^A$, it follows that $x_{\text{new}}^B(t^B) = a^B$. Let us now consider an arbitrary $i \in \mathcal{B}_{\text{thid}}$ such that $i \neq t^A$. Then, $x_{\text{new}}^A(i) = x^A(i) = x^B(i)$.

Literals of the form $l = z[t]_L$:

$$\begin{aligned}
z[t]_L^B &= z^B(t^B) \\
&= z^A(t^B) \\
&= z^A(t^A) \\
&= z[t]_L^A \\
&= l^A \\
&= l^B
\end{aligned}$$

Literals of the form $z_{new} = z\{t \leftarrow l\}_L$: In this case, since $z_{new}^A(t^A) = l^A$, it follows that $z_{new}^B(t^B) = l^B$.
Let us now consider an arbitrary $i \in \mathcal{B}_{\text{thid}}$ such that $i \neq t^A$. Then, $z_{new}^A(i) = z^A(i) = z^B(i)$.

Literals of the form $s = \text{addr2set}_K(m, a, l)$: Let $x = a^B = a^A$. Then,

$$\begin{aligned}
s^B &= s^A \cap \mathcal{B}_{\text{addr}} \\
&= \{y \in \mathcal{A}_{\text{addr}} \mid \exists p \in \mathcal{A}_{\text{path}} \text{ s.t.}, (m^A, x, y, l, p) \in \text{reach}_K^A\} \cap \mathcal{B}_{\text{addr}} \\
&= \{y \in \mathcal{B}_{\text{addr}} \mid \exists p \in \mathcal{A}_{\text{path}} \text{ s.t.}, (m^A, x, y, l, p) \in \text{reach}_K^A\} \\
&= \{y \in \mathcal{B}_{\text{addr}} \mid \exists p \in \mathcal{B}_{\text{path}} \text{ s.t.}, (m^B, x, y, l, p) \in \text{reach}_K^B\}
\end{aligned}$$

It just remains to see that the last equality holds. Let

- $S_B = \{y \in \mathcal{B}_{\text{addr}} \mid \exists p \in \mathcal{B}_{\text{path}} \text{ s.t.}, (m^B, x, y, l, p) \in \text{reach}_K^B\}$, and
- $S_A = \{y \in \mathcal{B}_{\text{addr}} \mid \exists p \in \mathcal{A}_{\text{path}} \text{ s.t.}, (m^A, x, y, l, p) \in \text{reach}_K^A\}$

We first show that $S_A \subseteq S_B$. Let $y \in S_A$. Then it exists $p \in \mathcal{A}_{\text{path}}$ such that $(m^A, x, y, l, p) \in \text{reach}_K^A$. Then, by definition of reach_K there are two possible cases.

- If $p = \epsilon$ and $x = y$, then $(m^B, x, y, l, \epsilon^B) \in \text{reach}_K^B$ and therefore $y \in S_B$.
- Otherwise, there exists $a_1, \dots, a_n \in \mathcal{A}_{\text{addr}}$ s.t.,

$$\begin{array}{ll}
i) & p = [a_1, \dots, a_n] & iii) & m^A(a_r).\text{next}^A(l) = a_{r+1}, \text{ for } 1 \leq r < n \\
ii) & x = a_1 & iv) & m^A(a_n).\text{next}^A(l) = y
\end{array}$$

Then, we only need to find $\tilde{a}_1, \dots, \tilde{a}_m \in \mathcal{B}_{\text{addr}}$ s.t.,

$$\begin{array}{ll}
i) & q = [\tilde{a}_1, \dots, \tilde{a}_m] & iii) & m^B(\tilde{a}_r).\text{next}^B(l) = \tilde{a}_{r+1}, \text{ for } 1 \leq r < m \\
ii) & x = \tilde{a}_1 & iv) & m^B(\tilde{a}_m).\text{next}^B(l) = y
\end{array}$$

We define $\tilde{a}_1 = a_1 = x$ and $\tilde{a}_2 = \text{first}_K(m^A, \tilde{a}_1, l, \mathcal{B}_{\text{addr}})$. Then, $\tilde{a}_2 = m^B(\tilde{a}_1).\text{next}^B(l)$ and $\tilde{a}_2 \in \mathcal{B}_{\text{addr}}$ and thus $\tilde{a}_2 \in \mathcal{B}_{\text{addr}}$. If $\tilde{a}_2 = y$ there is nothing else to prove. On the other hand, if $\tilde{a}_2 \neq y$ then we proceed in the same way to define \tilde{a}_3 and so on until $\tilde{a}_{m+1} = y$. Notice that this way, y is guaranteed to be found in at most n steps.

To show that $S_B \subseteq S_A$ we proceed in a similar way. Let $y \in S_B$. Then $x = y$ and $p = \epsilon$ and thus $(m^A, x, y, l, \epsilon^A) \in \text{reach}_K^A$, or exists $a_1, \dots, a_n \in \mathcal{B}_{\text{addr}}$ such that

$$\begin{array}{ll}
i) & p = [a_1, \dots, a_n] & iii) & m^B(a_r).\text{next}^B(l) = a_{r+1}, \text{ for } 1 \leq r < n \\
ii) & x = a_1 & iv) & m^B(a_n).\text{next}^B(l) = y
\end{array}$$

As we know that $a_1, \dots, a_n, y \in \mathcal{B}_{\text{addr}}$, by definition of first_K we know that there exists $s \geq 1$ s.t.,

$$m^A \left(\underbrace{\dots (m^A(a_1).next^A(l)) \dots}_{s} \right).next^A(l) = a_2$$

Let then $a_1^1, \dots, a_1^{s-1} \in \mathcal{A}_{\text{addr}}$ such that

$$\begin{aligned} m^A(a_1).next^A(l) &= a_1^1 \\ m^A(a_1^1).next^A(l) &= a_1^2 \\ &\vdots \\ m^A(a_1^{s-1}).next^A(l) &= a_2 \end{aligned}$$

We then use a_1^1, \dots, a_1^{s-1} to construct the section of a path q that goes from a_1 up to a_2 . Finally we use the same approach to finish the construction of such a path in \mathcal{A} . Then we have that $(m^A, x, y, l, q^A) \in reach_{\mathcal{K}}^A$. Hence, $y \in \mathcal{S}_A$.

Literals of the form $p = getp_{\mathcal{K}}(m, a, b, l)$: We consider two possible cases.

- Case $b^A \in addr2set_{\mathcal{K}}(m^A, a^A, l)$.
Since $(m^A, a^A, b^A, l, p^A) \in reach_{\mathcal{K}}^A$, it is enough to prove:

$$\text{if } (m^A, x, y, l, q) \in reach_{\mathcal{K}}^A \text{ then } (m^B, x, y, l, compress(q, \mathcal{B}_{\text{addr}})) \in reach_{\mathcal{K}}^B$$

for each $x, y \in \mathcal{B}_{\text{addr}}$ and $q \in \mathcal{A}_{\text{path}}$. If $(m^A, x, y, l, q) \in reach_{\mathcal{K}}^A$, $x = y$ and $q = \epsilon$, then $(m^B, x, y, l, compress(q, \mathcal{B}_{\text{addr}})) \in reach_{\mathcal{K}}^B$. Otherwise, there exist $a_1, \dots, a_n \in \mathcal{A}_{\text{addr}}$ such that:

$$\begin{array}{ll} \text{i)} & q = [a_1, \dots, a_n] \\ \text{ii)} & x = a_1 \\ \text{iii)} & m^A(a_r).next^A(l) = a_{r+1}, \text{ for } 1 \leq r < n \\ \text{iv)} & m^A(a_n).next^A(l) = y \end{array}$$

Then, we proceed by induction on n .

- If $n = 1$, then $q = [a_1]$ and therefore $compress(q, \mathcal{B}_{\text{addr}}) = [a_1]$, as $x = a_1 \in \mathcal{B}_{\text{addr}}$. Besides, $m^A(a_1).next^A(l) = y$ which implies that $m^B(a_1).next^B(l) = y$. Therefore, $(m^B, x, y, l, compress(q, \mathcal{B}_{\text{addr}})) \in reach_{\mathcal{K}}^B$.
- If $n > 1$, then let $a_i = first_{\mathcal{K}}(m^A, x, l, \mathcal{B}_{\text{addr}})$. As

$$q = [x = a_1, a_2, \dots, a_i, a_{i+1}, \dots, a_n]$$

we have that

$$compress(q, \mathcal{B}_{\text{addr}}) = [x = a_1] \circ compress([a_i, a_{i+1}, \dots, a_n], \mathcal{B}_{\text{addr}})$$

Besides, as $(m^A, a_i, y, l, [a_i, a_{i+1}, \dots, a_n]) \in reach_{\mathcal{K}}^A$, by induction we have that

$$(m^B, a_i, y, l, compress([a_i, a_{i+1}, \dots, a_n], \mathcal{B}_{\text{addr}})) \in reach_{\mathcal{K}}^B$$

Moreover $m^B(x).next^B(l) = a_i$ and therefore

$$(m^B, x, y, l, compress(q, \mathcal{B}_{\text{addr}})) \in reach_{\mathcal{K}}^B$$

- Case $b^A \notin addr2set_{\mathcal{K}}(m^A, a^A, l)$.

In such case we have that $p^A = \epsilon$, which implies that $p^B = \epsilon$. Then using a reasoning similar to the previous case we can deduce that $b^B \notin addr2set_{\mathcal{K}}(m^B, a^B, l)$.

Literals of the form $a = fstlock_K(m, p, l)$: Consider the case $p = \epsilon$: then $fstlock_K^A(m^A, \epsilon^A, l^A) = null^A$. At the same time, we know that $\epsilon^B = compress(\epsilon^A, \mathcal{B}_{addr})$ and so $fstlock_K^B(m^B, \epsilon^B, l^B) = null^B$. Let's now consider the case at which $p = [a_1, \dots, a_n]$. There are two scenarios to consider.

- If for all $1 \leq k \leq n$, $m^A(a_k^A).lockid(l) = \emptyset$, then we have that

$$fstlock_K^A(m^A, p^A, l^A) = null^A$$

Notice that function *compress* returns a subset of the path it receives with the property that all addresses in the returned path belong to the received set. Hence, if $[\tilde{a}_1, \dots, \tilde{a}_m] = p^B = compress(p^A, \mathcal{B}_{addr})$, we know that $\{\tilde{a}_1, \dots, \tilde{a}_m\} \subseteq \mathcal{B}_{addr}$ and therefore for all $1 \leq j \leq m$, $m^B(\tilde{a}_j).lockid(l^B) = \emptyset$. Then, we can conclude that in fact $fstlock_K^B(m^B, p^B, l^B) = null^B$.

- If there is a k , with $1 \leq k \leq n$, such that for all $1 \leq j < k$, $m^A(a_j^A).lockid(l) = \emptyset$ and $m^A(a_k^A).lockid(l) \neq \emptyset$ then since $a^B = a^A$, we can say that $a^B = a^A = x \in \mathcal{B}_{addr}$. It then remains to verify whether

$$\text{if } x = fstlock_K^A(m^A, p^A, l^A) \text{ then } x = fstlock_K^B(m^B, compress(p^A, \mathcal{B}_{addr}), l^B)$$

By definition of *fstlock_K* we have that $x = a_k^A$ and by construction of \mathcal{B}_{addr} we know that $a_k^A \in \mathcal{B}_{addr}$. Let $[\tilde{a}_1, \dots, \tilde{a}_i, \dots, \tilde{a}_m] = compress(p^A, \mathcal{B}_{addr})$ such that $\tilde{a}_i = a_k^A$. We had that $\tilde{a}_j \in \mathcal{B}_{addr}$ for all $1 \leq j \leq m$. As *compress* preserves the order and for all $1 \leq j < k$, $m^A(a_j^A).lockid(l^A) = \emptyset$, we have that for all $1 \leq j < i$, $m^B(\tilde{a}_j).lockid(l^B) = \emptyset$. Besides $m^B(\tilde{a}_i).lockid(l^B) \neq \emptyset$. Then:

$$\begin{aligned} fstlock_K^B(m^B, compress(p^A, l^B)) &= fstlock_K^B(m^B, [\tilde{a}_1, \dots, \tilde{a}_m], l^B) \\ &= \tilde{a}_i \\ &= a_k^A \\ &= x \end{aligned}$$

Literals of the form $ordList(m, p)$: Assume that $(m^A, p^A) \in ordList^A$. We would like to verify whether $(m^B, p^B) \in ordList^B$ i.e., $(m^B, compress(p^A, \mathcal{B}_{addr})) \in ordList^B$. We proceed by induction on p .

- If $p = \epsilon$, by definition of *compress* and *ordList*, we have that $(m^B, \epsilon^B) \in ordList^B$.
- If $p = [a_1]$, we know that $(m^A, [a_1]^A) \in ordList^A$ and that $p^B = compress(p^A, \mathcal{B}_{addr})$. Then, if $a_1^A \in \mathcal{B}_{addr}$, we have that $p^B = [a_1]^B$ and then clearly $(m^B, p^B) \in ordList^B$ holds. On the other hand, if $a_1^A \notin \mathcal{B}_{addr}$, then $p^B = \epsilon^B$ and once more $(m^B, p^B) \in ordList^B$ holds.
- If $p = [a_1, \dots, a_{n+1}]$ with $n \geq 1$, then we have two possible cases to bear in mind. If we consider the case at which $a_1^A \notin \mathcal{B}_{addr}$ then

$$compress(p^A, \mathcal{B}_{addr}) = compress([a_2, \dots, a_{n+1}]^A, \mathcal{B}_{addr})$$

and, as by induction we have that $(m^B, compress([a_2, \dots, a_{n+1}]^A, \mathcal{B}_{addr})) \in ordList^B$, we conclude that $(m^B, compress([a_1, a_2, \dots, a_{n+1}]^A, \mathcal{B}_{addr})) \in ordList^B$. On the other hand, if we assume $a_1^A \in \mathcal{B}_{addr}$, by induction, $(m^B, compress([a_2, \dots, a_{n+1}]^A, \mathcal{B}_{addr})) \in ordList^B$. Besides, as $m^A(a_1^A).key^A \preceq m^A(a_2^A).key^A$ we can deduce that $m^B(a_1^A).key^B \preceq m^B(a_2^A).key^B$. And so, $(m^B, compress([a_1, a_2, \dots, a_{n+1}]^A, \mathcal{B}_{addr})) \in ordList^B$.

Literals of the form $\neg ordList(m, p)$: Let's assume that $(m^A, p^A) \notin ordList^A$. We want to see that $(m^B, p^B) \notin ordList^B$. If $(m^A, p^A) \notin ordList^A$, then it means that $p = [a_1, \dots, a_n]$ with $n \geq 2$ and $m^A(a_{i+1}).key^A \preceq m^A(a_i).key^A$ and $m^A(a_{i+1}).key^A \neq m^A(a_i).key^A$ for some $i \in 1, \dots, n-1$. Let that i be the one such that for all $j < i$, $m^A(a_j).key^A \preceq m^A(a_{j+1}).key^A$. Then, by Lemma 5.3 we have $unordered(m^A, [a_1, \dots, a_n]^A) \neq \emptyset$ and besides $\{a_i^A, a_{i+1}^A\} \subseteq unordered(m^A, [a_1, \dots, a_n]^A) \subseteq \mathcal{B}_{addr}$. This means that

$$\text{compress}([a_1, \dots, a_n]^A, \mathcal{B}_{\text{addr}}) = [\tilde{a}_1, \dots, a_i, a_{i+1}, \dots, \tilde{a}_m]^B$$

Therefore, since $m^B(a_{i+1}).\text{key}^B \preceq m^B(a_i).\text{key}^B$ and $m^B(a_{i+1}).\text{key}^B \neq m^B(a_i).\text{key}^B$, we have that

$$(m^B, \text{compress}([a_1, \dots, a_n]^A, \mathcal{B}_{\text{addr}})) \notin \text{ordList}^B. \quad \square$$

5.2 A Combination-based Decision Procedure for TSL_K

Lemma 5.4 enables a brute force method to automatically check whether a set of normalized TSL_K -literals is satisfiable. However, such a method is not efficient in practice. We describe now how to obtain a more efficient decision procedure for TSL_K applying a many-sorted variant [TZ04] of the Nelson-Oppen combination method [NO79], by combining the decision procedures for the underlying theories. This combination method requires that the theories fulfill some conditions. First, each theory must have a decision procedure. Second, two theories can only share sorts (but not functions or predicates). Third, when two theories are combined, either both theories are stable infinite or one of them is polite with respect to the underlying sorts that it shares with the other. The stable infinite condition for a theory establishes that if a formula has a model then it has a model with infinite cardinality. In our case, some theories are not stable infinite. For example, T_{level_K} is not stably infinite, T_{ord} , and T_{thid} need not be stable infinite in some instances. The observation that the condition of stable infinity may be cumbersome in the combination of theories for data structures was already made in [RZ06a] where they suggest the condition of *politeness*:

Definition 5.3 (Smoothness).

Let $\Sigma = (S, F, P)$ be a signature, $S_0 = \{\sigma_1, \dots, \sigma_n\} \subseteq S$ be a set of sorts, and T be a Σ -theory. T is smooth with respect to S_0 if:

- (i) for every T -satisfiable quantifier-free Σ -formula φ ,
- (ii) for every T -interpretation \mathcal{A} satisfying φ ,
- (iii) for every cardinal number k_1, \dots, k_n such that $k_i \geq |\mathcal{A}_{\sigma_i}|$

there exists a T -interpretation \mathcal{B} satisfying φ such that $|\mathcal{B}_{\sigma_i}| = k_i$, for $i = 1, \dots, n$. †

Definition 5.4 (Finite witnessability).

Let $\Sigma = (S, F, P)$ be a signature, $S_0 \subseteq S$ be a set of sorts, and T be a Σ -theory. T is finitely witnessable with respect to S_0 if there exists a computable function *witness* that for every quantifier-free Σ -formula φ returns a quantifier-free Σ -formula $\psi = \text{witness}(\varphi)$ such that

- (i) φ and $(\exists \bar{v})\psi$ are T -equivalent, where $\bar{v} = \text{vars}(\psi) \setminus \text{vars}(\varphi)$, and
- (ii) if ψ is T -satisfiable then there exists a T -interpretation \mathcal{A} satisfying ψ such that the domain \mathcal{A}_σ interpreting the sort σ in \mathcal{A} is the (finite) set $[\text{vars}_\sigma(\psi)]^{\mathcal{A}}$ of elements in \mathcal{A} interpreting the variables of sort σ in ψ (in symbols, $\mathcal{A}_\sigma = [\text{vars}_\sigma(\psi)]^{\mathcal{A}}$), for each $\sigma \in S_0$. †

Definition 5.5 (Politeness).

Let $\Sigma = (S, F, P)$ be a signature, $S_0 \subseteq S$ be a set of sorts and T be a Σ -theory. T is polite with respect to S_0 if it is both smooth and finitely witnessable with respect to S_0 . †

Smoothness guarantees that interpretations can be enlarged as needed. Finite witnessability gives a procedure to produce a model in which every element is represented by a variable. The Finite Model Property, Lemma 5.4 above, guarantees that every sub-theory of TSL_K is finite witnessable since one can add as many fresh variables as the bound for the corresponding sort in the lemma, even one for each element in the domain of the sorts. The smoothness property can be shown for:

$$T_{\text{cell}} \cup T_{\text{mem}} \cup T_{\text{path}} \cup T_{\text{set}} \cup T_{\text{settid}} \cup T_{\text{mrgn}} \cup T_{\text{aarr}} \cup T_{\text{larr}}$$

with respect to sorts `addr`, `levelK`, `elem`, `ord` and `thid`. Moreover, these theories can be combined because all of them are stably infinite. The following can also be combined:

$$T_{\text{level}_k} \cup T_{\text{ord}} \cup T_{\text{thid}}$$

because they do not share any sorts, so combination is trivial. The many-sorted Nelson-Oppen method allows to combine the first collection of theories with the second. Regarding the decision procedures for each individual theory, T_{level_k} is trivial since it is just a finite set of naturals with order. For T_{ord} we can adapt a decision procedure for dense orders as the reals [Tar51], or other appropriate theory. For instance, if no dense order is required, we can consider Presburger Arithmetic equipped only with 0, *succ* and *<*. For T_{cell} we can use a decision procedure for recursive data structures [Opp80]. T_{mem} is the theory of arrays [ARR03]. T_{set} , T_{settid} and T_{mrgn} are theories of (finite) sets for which there are many decision procedures [Zar03, KNR05]. For theories T_{aarr} and T_{larr} we can use a decision procedure for arrays [BMS06]. The remaining theories are T_{reach} and T_{bridge} . Following the approaches in [RZ06a, SS10] we extend a decision procedure for the theory T_{path} of finite sequences of (non-repeated) addresses with the auxiliary functions and predicates shown in Fig. 5.5, and combine this theory with the theories that are part of TSL_K (except from T_{reach}) to obtain:

$$T_{\text{SLKBase}} = T_{\text{addr}} \cup T_{\text{ord}} \cup T_{\text{thid}} \cup T_{\text{level}_k} \cup T_{\text{cell}} \cup T_{\text{mem}} \cup T_{\text{path}} \cup T_{\text{set}} \cup T_{\text{settid}} \cup T_{\text{mrgn}} \cup T_{\text{aarr}} \cup T_{\text{larr}}$$

Using T_{path} all symbols in T_{reach} can be easily defined. Once more, we pick the theory of finite sequences of addresses defined by $T_{\text{fseq}} = (\Sigma_{\text{fseq}}, \text{TGen})$, where:

$$\Sigma_{\text{fseq}} = \left(\begin{array}{l} \{ \text{addr, fseq} \\ \{ \text{nil} : \text{fseq}, \\ \text{cons} : \text{addr} \times \text{fseq} \rightarrow \text{fseq}, \\ \text{hd} : \text{fseq} \rightarrow \text{addr}, \\ \text{tl} : \text{fseq} \rightarrow \text{fseq} \\ \{ \end{array} \right)$$

and TGen as the class of term-generated structures that satisfy the axioms of distinctness, uniqueness and generation of sequences using constructors, as well as acyclicity (see, for example [BM07]). Let Σ_{path} be Σ_{fseq} extended with the symbols of Fig. 5.5 and let *PATH* be the set of axioms of T_{fseq} including the ones in Fig. 5.5. Then, we can formally define $T_{\text{path}} = (\Sigma_{\text{path}}, \text{ETGen})$ where

$$\text{ETGen} = \left\{ \mathcal{A}^{\Sigma_{\text{path}}} \mid \mathcal{A}^{\Sigma_{\text{path}}} \models \text{PATH} \text{ and } \mathcal{A}^{\Sigma_{\text{fseq}}} \in \text{TGen} \right\}$$

Next, we extend T_{SLKBase} with definitions for translating all missing functions and predicates from Σ_{reach} and Σ_{bridge} appearing in normalized TSL_K-literals by definitions from T_{SLKBase} . Let *GAP* be the set

of axioms that define ϵ , $[_]$, $append$, $reach_K$, $path2set$, $getp_K$, $fstlock_K$ and $ordList$. Some of the equivalences for GAP were given in Section 4.2. Hence, here we limit ourselves only to the new equivalences. For instance:

$$\begin{aligned} isreachp_K(m, a_1, a_2, l, p) &\rightarrow getp_K(m, a_1, a_2, l) = p \\ \neg isreachp(m, a_1, a_2, l, p) &\rightarrow getp_K(m, a_1, a_2, l) = nil \\ ispath(p) \wedge fstmark(m, p, l, i) &\leftrightarrow fstlock_K(m, p, l) = i \\ ispath(p) \wedge ordPath(m, p) &\leftrightarrow ordList(m, p) \end{aligned}$$

We now define $\widehat{TSL}_K = (\Sigma_{\widehat{TSL}_K}, \widehat{ETGen})$ where

$$\Sigma_{\widehat{TSL}_K} = \Sigma_{T_{SLKBase}} \cup \{append, reach_K, path2set, getp_K, fstlock_K, ordList\}$$

and $\widehat{ETGen} := \{ \mathcal{A}^{\Sigma_{\widehat{TSL}_K}} \mid \mathcal{A}^{\Sigma_{\widehat{TSL}_K}} \models GAP \text{ and } \mathcal{A}^{\Sigma_{T_{SLKBase}}} \in ETGen \}$.

Using the definitions of GAP it is easy to prove that if Γ is a set of normalized TSL_K -literals, then Γ is TSL_K -satisfiable iff Γ is \widehat{TSL}_K -satisfiable. Therefore, \widehat{TSL}_K can be used in place of TSL_K for satisfiability checking. The reduction from \widehat{TSL}_K into $T_{SLKBase}$ is performed in two steps. First, by the finite model theorem (Lemma 5.4), it is always possible to calculate an upper bound in the number of elements of

$app : fseq \times fseq \rightarrow fseq$
$app(nil, l) = l$
$app(cons(a, l), l') = cons(a, app(l, l'))$
$fseq2set : fseq \rightarrow set$
$fseq2set(nil) = \emptyset$
$fseq2set(cons(a, l)) = \{a\} \cup fseq2set(l)$
$ispath : fseq$
$ispath(nil)$
$ispath(cons(a, nil))$
$\{a\} \notin fseq2set(l) \wedge ispath(l) \rightarrow ispath(cons(a, l))$
$last : fseq \rightarrow addr$
$last(cons(a, nil)) = a$
$l \neq nil \rightarrow last(cons(a, l)) = last(l)$
$isreach_K : mem \times addr \times addr \times level_K$
$isreach_K(m, a, a, l)$
$m[a].next[l] = a' \wedge isreach_K(m, a', b, l) \rightarrow isreach_K(m, a, b, l)$
$isreachp_K : mem \times addr \times addr \times level_K \times fseq$
$isreachp_K(m, a, a, l, nil)$
$m[a].next[l] = a' \wedge isreachp(m, a', b, l, p) \rightarrow isreachp(m, a, b, l, cons(a, p))$
$fstmark : mem \times fseq \times level_K \times addr$
$fstmark(m, nil, l, null)$
$p \neq nil \wedge p = cons(a, q) \wedge m[a].lockid[l] \neq \emptyset \rightarrow fstmark(m, p, l, a)$
$p \neq nil \wedge p = cons(a, q) \wedge m[a].lockid[l] = \emptyset \wedge fstmark(m, q, l, b) \rightarrow fstmark(m, p, l, b)$
$ordPath : mem \times fseq$
$ordPath(h, nil)$
$ordPath(h, cons(a, nil))$
$h[a].next[0] = a' \wedge h[a].key \preceq h[a'].key \wedge p = cons(a, q) \wedge ordPath(h, q) \rightarrow ordPath(h, p)$

Figure 5.5: Functions, predicates and axioms of T_{path}

sort $addr$, $elem$, $thid$, ord and $level$ in a model (if there is one model), based on the input formula. Therefore, one can introduce one variable per element of each of these sorts and unfold all definitions in $PATH$ and GAP , by symbolic expansion, leading to terms in Σ_{fseq} , and thus, in $T_{SLKBase}$. This way, it is always possible to reduce a \widehat{TSL}_K -satisfiability problem of normalized literals into a $T_{SLKBase}$ -satisfiability problem. Hence, using a decision procedure for $T_{SLKBase}$ we obtain a decision procedure for \widehat{TSL}_K , and thus, for TSL_K . Notice, for instance, that the predicate $subPath : path \times path$ for ordered lists can be defined using only $path2set$ as:

$$subPath(p_1, p_2) \hat{=} path2set(p_1) \subseteq path2set(p_2)$$

Example 5.2

We now illustrate how reduction from \widehat{TSL}_K to $T_{SLKBase}$ is performed. For the sake of simplicity, let just consider the formula $\varphi = SkipList_4(h, sl.head, sl.tail)$. According to the definition we gave in Example 3.3,

$$\varphi \hat{=} OList(h, sl.head, 0) \wedge \tag{5.1}$$

$$\left(\begin{array}{l} sl.tail.next[0] = null \wedge sl.tail.next[1] = null \\ sl.tail.next[2] = null \wedge sl.tail.next[3] = null \end{array} \right) \wedge \tag{5.2}$$

$$\left(\begin{array}{l} SubList(h, sl.head, sl.tail, 1, sl.head, sl.tail, 0) \wedge \\ SubList(h, sl.head, sl.tail, 2, sl.head, sl.tail, 1) \wedge \\ SubList(h, sl.head, sl.tail, 3, sl.head, sl.tail, 2) \end{array} \right) \tag{5.3}$$

For simplicity, from now on, we use $head$ for $sl.head$ and $tail$ for $sl.tail$. Notice that φ is not in TSL_K . In fact, it contains some functions like $OList$ and $SubList$ which do not belong to TSL_K but give us the intuition of what a skiplist should look like. Hence, we first rewrite φ using terms of TSL_K :

$$\begin{array}{l} (5.1) \quad \left[\begin{array}{l} \overbrace{p = getp_K(h, head, null, 0)}^{\Sigma_{bridge}} \wedge \overbrace{ordList(h, p)}^{\Sigma_{bridge}} \wedge \\ \overbrace{c_{tail} = h[tail]}^{\Sigma_{mem}} \wedge \overbrace{a_{tail_0} = c_{tail}.next[0]}^{\Sigma_{cell}} \wedge \overbrace{a_{tail_0} = null}^{\Sigma_{addr}} \end{array} \right. \\ (5.2) \quad \left[\begin{array}{l} \wedge \vdots \wedge \vdots \\ \wedge \overbrace{a_{tail_3} = c_{tail}.next[3]}^{\Sigma_{cell}} \wedge \overbrace{a_{tail_3} = null}^{\Sigma_{addr}} \wedge \end{array} \right. \\ (5.3) \quad \left[\begin{array}{l} \overbrace{p_{tail_0} = getp_K(h, head, tail, 0)}^{\Sigma_{bridge}} \wedge \dots \wedge \overbrace{p_{tail_3} = getp_K(h, head, tail, 3)}^{\Sigma_{bridge}} \wedge \\ \overbrace{s_{tail_0} = path2set(p_{tail_0})}^{\Sigma_{bridge}} \wedge \dots \wedge \overbrace{s_{tail_3} = path2set(p_{tail_3})}^{\Sigma_{bridge}} \wedge \\ \overbrace{s_{tail_1} \subseteq s_{tail_0}}^{\Sigma_{set}} \wedge \dots \wedge \overbrace{s_{tail_3} \subseteq s_{tail_2}}^{\Sigma_{set}} \end{array} \right. \end{array}$$

Notice that after been rewritten, we already have a conjunction of literals belonging to \widehat{TSL}_K . We can go a little further, and rewrite φ using normalized literals. Some of the literals are already normalized, like $ordList(h, p)$, while other must be rewritten, like $a_{tail_i} = c_{tail}.next[i]$ or $s_{tail_j} \subseteq s_{tail_{j-1}}$ to:

$$\begin{array}{l} c_{tail} = mkcell(e_0, k_0, a_{tail_0}, a_0^1, a_0^2, a_0^3, t_0^0, t_0^1, t_0^2, t_0^3) \wedge a_{tail_0} = null \\ c_{tail} = mkcell(e_1, k_1, a_1^0, a_{tail_1}, a_1^2, a_1^3, t_1^0, t_1^1, t_1^2, t_1^3) \wedge a_{tail_1} = null \\ c_{tail} = mkcell(e_2, k_2, a_2^0, a_2^1, a_{tail_2}, a_2^3, t_2^0, t_2^1, t_2^2, t_2^3) \wedge a_{tail_2} = null \\ c_{tail} = mkcell(e_3, k_3, a_3^0, a_3^1, a_3^2, a_{tail_3}, t_3^0, t_3^1, t_3^2, t_3^3) \wedge a_{tail_3} = null \end{array}$$

and

$$s_{tail_0} = s_{tail_0} \cup s_{tail_1} \wedge s_{tail_1} = s_{tail_1} \cup s_{tail_2} \wedge s_{tail_2} = s_{tail_2} \cup s_{tail_3}$$

respectively. So far, we have

$$\begin{aligned} V_{\text{elem}}(\varphi) &= \{e_0, e_1, e_2, e_3\} \\ V_{\text{addr}}(\varphi) &= \{\text{head}, \text{tail}\} \cup \{a_{tail_i} \mid 1 \in 0..3\} \cup (\{a_i^j \mid i, j \in 0..3 \wedge i \neq j\}) \\ V_{\text{mem}}(\varphi) &= \{h\} \\ V_{\text{path}}(\varphi) &= \{p\} \cup \{p_{tail_i} \mid i \in 0..3\} \\ V_{\text{thid}}(\varphi) &= \emptyset \\ V_{\text{ord}}(\varphi) &= \emptyset \\ V_{\text{aarr}}(\varphi) &= \emptyset \end{aligned}$$

Therefore, assuming φ holds in a model \mathcal{A} , when constructing a finite model \mathcal{B} , we can ensure the following bounds for the domains of each sort in \mathcal{B} :

$$\begin{aligned} |\mathcal{B}_{\text{level}_K}| &\leq K = 4 \\ |\mathcal{B}_{\text{thid}}| &\leq \bar{i} + 1 + K \bar{m} \bar{a} = 0 + 1 + 4 \times 1 \times 18 = 73 \\ |\mathcal{B}_{\text{addr}}| &\leq \bar{a} + 1 + K \bar{m} \bar{a} + \bar{p}^2 + \bar{p}^3 + 2 \bar{m} \bar{p} + \bar{x} |\mathcal{B}_{\text{thid}}| \\ &= 18 + 1 + 4 \times 1 \times 18 + 5^2 + 5^3 + 2 \times 1 \times 5 + 0 \times 73 = 251 \\ |\mathcal{B}_{\text{elem}}| &\leq \bar{e} + \bar{m} |\mathcal{B}_{\text{addr}}| = 4 + 1 \times 251 = 255 \\ |\mathcal{B}_{\text{ord}}| &\leq \bar{o} + \bar{m} |\mathcal{B}_{\text{addr}}| = 0 + 1 \times 251 = 251 \end{aligned}$$

In fact, when we construct the domains for \mathcal{B} , we find out that we only require the following sets, which clearly are under the bounds we have calculated:

$$\begin{aligned} |\mathcal{B}_{\text{level}_K}| &= \{0, 1, 2, 3\} \\ |\mathcal{B}_{\text{thid}}| &= \{\emptyset\} \\ |\mathcal{B}_{\text{addr}}| &= \{0x00, 0x01, 0x08\} \\ |\mathcal{B}_{\text{elem}}| &= \{-\infty, +\infty\} \\ |\mathcal{B}_{\text{ord}}| &= \{-\infty, +\infty\} \end{aligned}$$

Then, it just remains to reduce literals from Σ_{bridge} using the definitions of *GAP* and *PATH*, to end up with terms in T_{SLKBase} . This way, we get:

Σ_{bridge} -term	T_{SLKBase} -term
$p = \text{getp}_K(h, \text{head}, \text{tail}, 0)$	$h[\text{head}].\text{next}[0] = \text{tail} \quad \wedge \quad h[\text{tail}].\text{next}[0] = \text{null} \quad \wedge$ $\text{isreachp}_K(h, \text{null}, \text{null}, 0, \text{nil})$
$\text{ordList}(h, p)$	$h[\text{head}].\text{next}[0] = \text{tail} \quad \wedge \quad h[\text{head}].\text{key} \preceq h[\text{tail}].\text{key} \quad \wedge$ $p = \text{cons}(\text{head}, \text{cons}(\text{tail}, \text{nil})) \quad \wedge \quad \text{ordPath}(h, \text{cons}(\text{tail}, \text{nil})) \quad \wedge$ $\neg\{\text{head}\} \subseteq \{\text{tail}\} \quad \wedge \quad \text{ispath}(\text{cons}(\text{tail}, \text{nil}))$
$p_{tail_i} = \text{getp}_K(h, \text{head}, \text{tail}, i)$	$h[\text{head}].\text{next}[i] = \text{tail} \quad \wedge \quad \text{isreachp}_K(h, \text{tail}, \text{tail}, i, \text{nil})$
$s_{tail_i} = \text{path2set}(p_{tail_i})$	$s_{tail_i} = \{\text{head}\}$

Once we have achieved this point, we only need to call the decision procedures for each of the involved theories in $T_{SLKBase}$. *

5.3 Extending TSL_K to Reason about (L, U, H)

In Section 3.2 we presented the algorithms for `INSERT`, `REMOVE` and `SEARCH` over a concurrent skiplist. In that moment, we also annotated the code with ghost arrays L , U and H . L and U are arrays from thread identifiers to addresses while H is an array from thread identifiers to skiplist levels. The idea is that the triplet $(L[t], U[t], H[t])$ denotes the bounds of the section of the skiplist that thread t can modify. This is not a restriction imposed to thread t , but just a consequence of the progress of t . We usually refer to such section of the skiplist as the (L, U, H) of thread t .

Following these algorithms, in this section we show how we can extend TSL_K with some extra functions to let us reason about (L, U, H) inclusion. When we presented the pessimistic version of programs `INSERT`, `REMOVE` and `SEARCH` we stated that such implementation enjoys fairness. In fact, we can verify that any thread whose (L, U, H) does not strictly contain the (L, U, H) of another thread is guaranteed to terminate.

To begin with, we present two new theories we will use. A theory of pairs of thread identifiers (T_{pair}) and a theory of sets of pairs ($T_{setpair}$). The signature of these two new theories is defined as one may expect:

$$\begin{aligned} \Sigma_{pair} &= \left(\begin{array}{l} \{ \text{pair, thid} \\ \{ \langle _ , _ \rangle : \text{thid} \times \text{thid} \rightarrow \text{pair} \\ \{ \end{array} \right) \\ \\ \Sigma_{setpair} &= \left(\begin{array}{l} \{ \text{setpair, pair} \\ \{ \emptyset_p : \text{setpair}, \\ \{ _ \}_p : \text{pair} \rightarrow \text{setpair}, \\ \cup_p : \text{setpair} \times \text{setpair} \rightarrow \text{setpair}, \\ \cap_p : \text{setpair} \times \text{setpair} \rightarrow \text{setpair}, \\ \setminus_p : \text{setpair} \times \text{setpair} \rightarrow \text{setpair} \\ \{ \in_p : \text{pair} \times \text{setpair} \\ \subseteq_p : \text{setpair} \times \text{setpair} \end{array} \right) \end{aligned}$$

The interpretation for each new function and predicate is the expected one for pairs and sets. $\mathcal{A}_{pair} = \mathcal{A}_{thid} \times \mathcal{A}_{thid}$ while $\mathcal{A}_{setpair}$ is the power-set of \mathcal{A}_{pair} . Besides, we introduce a new theory, T_{LUH} , to reason about (L, U, H) . The signature of T_{LUH} , Σ_{LUH} , is given by:

$$\begin{aligned} \Sigma_{LUH} &= \left(\begin{array}{l} \{ \text{mem, aarr, larr, thid, settid, pair, setpair} \\ \{ \text{subset}_{LUH} : \text{mem} \times \text{aarr} \times \text{aarr} \times \text{larr} \times \text{addr} \times \text{addr} \times \text{level}_K \rightarrow \text{settid} \\ \text{inter}_{LUH} : \text{mem} \times \text{aarr} \times \text{aarr} \times \text{larr} \times \text{settid} \rightarrow \text{setpair} \\ \text{empty}_{LUH} : \text{mem} \times \text{aarr} \times \text{aarr} \times \text{larr} \times \text{addr} \times \text{addr} \times \text{level}_K \rightarrow \text{settid} \\ \text{min}_{LUH} : \text{mem} \times \text{aarr} \times \text{settid} \rightarrow \text{thid} \\ \{ \end{array} \right) \end{aligned}$$

Function $subset_{LUH}$ receives as argument a memory, three arrays (representing the (L, U, H) regions for a skiplist), a couple of addresses and a level. It returns the set of thread identifiers whose (L, U, H) is strictly contained into the region bounded by the address and level taken as argument. Function $inter_{LUH}$

receives as argument a set of pairs and it returns the subset of pairs whose (L, U, H) intersects with each other. Function $empty_{LUH}$ is similar to $subset_{LUH}$, except that it returns the set of thread identifiers whose (L, U, H) is empty, that is, those whose (L, U, H) does not contain the (L, U, H) of any other thread. Finally, function min_{LUH} receives a set of thread id and returns the minimum of all of them. We say that the minimum is the one with its right bound closest to the tail of the skiplist. To simplify notation, if a_1 and a_2 are two variables of type address, considering a memory m , we use

$$a_1 \leq a_2 \quad \text{for} \quad a_2 \in \text{addr2set}_K(m, a_1, 0)$$

Similarly, if \mathcal{A} is a model, i_1 and i_2 belong to $\mathcal{A}_{\text{addr}}$ and $m \in \mathcal{A}_{\text{mem}}$, we abuse notation and write

$$i_1 \leq i_2 \quad \text{for} \quad i_2 \in \text{addr2set}_K^{\mathcal{A}}(m, i_1, 0)$$

Using the syntactic sugar we have just introduced, given a model \mathcal{A} , the functions of Σ_{LUH} are interpreted this way:

$$\begin{aligned} \text{subset}_{LUH}^{\mathcal{A}}(m, L, U, H, a_1, a_2, l) &= \left\{ t' \in \mathcal{A}_{\text{thid}} \mid a_1 \leq L(t') \text{ and } U(t') \leq a_2 \text{ and} \right. \\ &\quad \left. (H(t') < l \text{ or } H(t') = l) \right\} \\ \text{inter}_{LUH}^{\mathcal{A}}(m, L, U, H, g) &= \left\{ \langle t_1, t_2 \rangle \in \mathcal{A}_{\text{pair}} \mid t_1 \in g, t_2 \in g, \right. \\ &\quad \left. \text{intersects}(t_1, t_2) \text{ and } t_1 \neq t_2 \right\} \\ \text{empty}_{LUH}^{\mathcal{A}}(m, L, U, H, a_1, a_2, l) &= \left\{ t' \in \mathcal{A}_{\text{thid}} \mid t' \in \text{subset}_{LUH}^{\mathcal{A}}(m, L, U, H, a_1, a_2, l) \text{ and} \right. \\ &\quad \left. \text{subset}_{LUH}^{\mathcal{A}}(m, L, U, H, L(t'), U(t'), H(t')) = \emptyset \right\} \\ \text{min}_{LUH}^{\mathcal{A}}(m, U, g) &= \begin{cases} \emptyset & \text{if } g = \emptyset \\ t_i & \text{if } g = \{t_1, \dots, t_q\}, i \in 1..q \text{ and for all } j \in 1..q \ U(t_i) \leq U(t_j) \end{cases} \end{aligned}$$

for each $m \in \mathcal{A}_{\text{mem}}$; $L, U \in \mathcal{A}_{\text{aarr}}$; $H \in \text{larr}$; $a_1, a_2 \in \mathcal{A}_{\text{addr}}$; $l \in \mathcal{A}_{\text{level}_K}$; $w \in \mathcal{A}_{\text{setpair}}$ and $g \in \mathcal{A}_{\text{settid}}$. Besides, we define intersects by:

$$\begin{aligned} \text{intersects}(t_1, t_2) &= (L(t_2) \leq L(t_1) \wedge L(t_1) \leq U(t_2) \wedge H(t_1) \geq 1 \wedge H(t_2) \geq 1) \vee \\ &\quad (L(t_2) \leq U(t_1) \wedge U(t_1) \leq U(t_2) \wedge H(t_1) \geq 1 \wedge H(t_2) \geq 1) \vee \\ &\quad (L(t_1) \leq L(t_2) \wedge U(t_2) \leq U(t_1) \wedge H(t_2) > H(t_1)) \end{aligned}$$

We now define the theory TSL_{K+} as the one obtained from the union of TSL_K with T_{pair} , T_{setpair} and T_{LUH} . As before, we now define the set of TSL_{K+} -normalized literals.

Definition 5.6 (TSL_{K+} -normalized literals).

A TSL_{K+} -literal is normalized if it is a TSL_K -normalized literal or it is a flat literal of the form:

$$\begin{aligned} u &= \langle t_1, t_2 \rangle \\ w &= \{u\}_P & w_1 &= w_2 \cup_P w_3 \\ w_1 &= w_2 \setminus_P w_3 \\ g &= \text{subset}_{LUH}(m, L, U, H, a_1, a_2, l) & w &= \text{inter}_{LUH}(m, L, U, H, g) \\ g &= \text{empty}_{LUH}(m, L, U, H, a_1, a_2, l) & t &= \text{min}_{LUH}(m, U, g) \end{aligned}$$

where a_1 and a_2 are addr -variables; m is a mem -variable; g is a settid -variable; L and U are aarr -variables; H is a Σ_{larr} -variable; l is a level_K -variable; t_1 and t_2 are thid -variables; u is pair -variable and w, w_1, w_2 and w_3 are setpair -variables. †

The non normalized literals can be constructed from the normalized ones using an approach similar to the one employed in the definition of literals from sets of thread identifiers. We now show that TSL_{K+} still satisfies the finite model property.

Lemma 5.5 (Finite Model Property for TSL_{K+}):

Let Γ be a conjunction of normalized TSL_{K+} -literals. Let $\bar{e} = |V_{\text{elem}}(\Gamma)|$, $\bar{a} = |V_{\text{addr}}(\Gamma)|$, $\bar{m} = |V_{\text{mem}}(\Gamma)|$, $\bar{p} = |V_{\text{path}}(\Gamma)|$, $\bar{t} = |V_{\text{thid}}(\Gamma)|$, $\bar{o} = |V_{\text{ord}}(\Gamma)|$, $\bar{u} = |V_{\text{pair}}(\Gamma)|$ and $\bar{x} = |V_{\text{aarr}}(\Gamma)|$. Then the following are equivalent:

1. Γ is TSL_{K+} -satisfiable;
2. Γ is true in a TSL_{K+} interpretation \mathcal{B} such that

$$\begin{aligned}
|\mathcal{B}_{\text{level}_k}| &\leq K \\
|\mathcal{B}_{\text{thid}}| &\leq \bar{t} + 1 + K\bar{m}\bar{a} + 2\bar{u} \\
|\mathcal{B}_{\text{addr}}| &\leq \bar{a} + 1 + K\bar{m}\bar{a} + \bar{p}^2 + \bar{p}^3 + 2\bar{m}\bar{p} + \bar{x}|\mathcal{B}_{\text{thid}}| \\
|\mathcal{B}_{\text{elem}}| &\leq \bar{e} + \bar{m}|\mathcal{B}_{\text{addr}}| \\
|\mathcal{B}_{\text{ord}}| &\leq \bar{o} + \bar{m}|\mathcal{B}_{\text{addr}}|
\end{aligned}$$



Proof (2 \rightarrow 1) is immediate.

(1 \rightarrow 2).

We prove this implication only for the new TSL_{K+} -literals. Let \mathcal{A} be a TSL_K -interpretation satisfying a set of normalized TSL_{K+} -literals Γ . We use \mathcal{A} to construct a finite TSL_{K+} -interpretation \mathcal{B} . \mathcal{B} is constructed as for TSL_K , except for the set of thread identifiers, which is defined as

$$\begin{aligned}
\mathcal{B}_{\text{thid}} = & V_{\text{thid}}^{\mathcal{A}} \cup \{ \emptyset \} \cup \\
& \{ m^{\mathcal{A}}(a^{\mathcal{A}}).lockid^{\mathcal{A}}(l) \mid m \in V_{\text{mem}}, a \in V_{\text{addr}} \text{ and } l \in \mathcal{B}_{\text{level}_k} \} \cup \\
& \{ t_1, t_2 \in \mathcal{A}_{\text{thid}} \mid (t_1, t_2) = p^{\mathcal{A}} \text{ and } p \in V_{\text{pair}} \}
\end{aligned}$$

and the interpretation of the new variables is given by

$$\begin{aligned}
u^{\mathcal{B}} &= u^{\mathcal{A}} && \text{for each } u \in V_{\text{pair}} \\
w^{\mathcal{B}} &= w^{\mathcal{A}} \cap (\mathcal{B}_{\text{thid}} \times \mathcal{B}_{\text{thid}}) && \text{for each } w \in V_{\text{setpair}}
\end{aligned}$$

We now proceed to the analysis of the new literals:

Literals of the form $u = \langle t_1, t_2 \rangle$:

$$\begin{aligned}
u^{\mathcal{B}} &= u^{\mathcal{A}} \\
&= (t_1^{\mathcal{A}}, t_2^{\mathcal{A}}) \\
&= (t_1^{\mathcal{B}}, t_2^{\mathcal{B}})
\end{aligned}$$

Literals of the form $w = \{u\}_p$:

$$\begin{aligned}
w^{\mathcal{B}} &= w^{\mathcal{A}} \cap (\mathcal{B}_{\text{thid}} \times \mathcal{B}_{\text{thid}}) \\
&= \{u^{\mathcal{A}}\}_p \cap (\mathcal{B}_{\text{thid}} \times \mathcal{B}_{\text{thid}}) \\
&= \{u^{\mathcal{B}}\}_p \cap (\mathcal{B}_{\text{thid}} \times \mathcal{B}_{\text{thid}}) \\
&= \{u^{\mathcal{B}}\}_p
\end{aligned}$$

Literals of the form $w_1 = w_2 \cup_p w_3$:

$$\begin{aligned}
w_1^{\mathcal{B}} &= w_1^{\mathcal{A}} \cap (\mathcal{B}_{\text{thid}} \times \mathcal{B}_{\text{thid}}) \\
&= (w_2^{\mathcal{A}} \cup_p w_3^{\mathcal{A}}) \cap (\mathcal{B}_{\text{thid}} \times \mathcal{B}_{\text{thid}}) \\
&= (w_2^{\mathcal{A}} \cap (\mathcal{B}_{\text{thid}} \times \mathcal{B}_{\text{thid}})) \cup (w_3^{\mathcal{A}} \cap (\mathcal{B}_{\text{thid}} \times \mathcal{B}_{\text{thid}})) \\
&= w_2^{\mathcal{B}} \cup_p w_3^{\mathcal{B}}
\end{aligned}$$

Literals of the form $w_1 = w_2 \setminus_P w_3$:

$$\begin{aligned}
w_1^B &= w_1^A \cap (\mathcal{B}_{\text{thid}} \times \mathcal{B}_{\text{thid}}) \\
&= (w_2^A \setminus w_3^A) \cap (\mathcal{B}_{\text{thid}} \times \mathcal{B}_{\text{thid}}) \\
&= (w_2^A \cap (\mathcal{B}_{\text{thid}} \times \mathcal{B}_{\text{thid}})) \setminus (w_3^A \cap (\mathcal{B}_{\text{thid}} \times \mathcal{B}_{\text{thid}})) \\
&= w_2^B \setminus_P w_3^B
\end{aligned}$$

Literals of the form $g = \text{subset}_{LUH}(m, L, U, H, a_1, a_2, l)$:

$$\begin{aligned}
g^B &= g^A \cap \mathcal{B}_{\text{thid}} \\
&= \text{subset}_{LUH}^A(m^A, L^A, U^A, H^A, a_1^A, a_2^A, l^A) \cap \mathcal{B}_{\text{thid}} \\
&= \left\{ t' \in \mathcal{A}_{\text{thid}} \mid a_1^A \leq L^A(t') \text{ and } U^A(t') \leq a_2^A \text{ and } (H^A(t') < l^A \text{ or } H^A(t') = l^A) \right\} \cap \mathcal{B}_{\text{thid}} \\
&= \left\{ t' \in \mathcal{B}_{\text{thid}} \mid a_1^B \leq L^B(t') \text{ and } U^B(t') \leq a_2^B \text{ and } (H^B(t') < l^B \text{ or } H^B(t') = l^B) \right\} \\
&= \text{subset}_{LUH}^B(m^A, L^B, U^B, H^B, a_1^B, a_2^B, l^B)
\end{aligned}$$

Literals of the form $w = \text{inter}_{LUH}(m, L, U, H, g)$:

$$\begin{aligned}
w^B &= w^A \cap (\mathcal{B}_{\text{thid}} \times \mathcal{B}_{\text{thid}}) \\
&= \text{inter}_{LUH}^A(m^A, L^A, U^A, H^A, g_1^A) \cap (\mathcal{B}_{\text{thid}} \times \mathcal{B}_{\text{thid}}) \\
&= \left\{ \langle t_1, t_2 \rangle \in \mathcal{A}_{\text{pair}} \mid t_1 \in g^A, t_2 \in g^A, \text{intersects}(t_1, t_2) \text{ and } t_1 \neq t_2 \right\} \cap (\mathcal{B}_{\text{thid}} \times \mathcal{B}_{\text{thid}}) \\
&= \left\{ \langle t_1, t_2 \rangle \in \mathcal{B}_{\text{pair}} \mid t_1 \in g^B, t_2 \in g^B, \text{intersects}(t_1, t_2) \text{ and } t_1 \neq t_2 \right\} \\
&= \text{inter}_{LUH}^B(m^A, L^B, U^B, H^B, g^B)
\end{aligned}$$

Literals of the form $g = \text{empty}_{LUH}(m, L, U, H, a_1, a_2, l)$:

$$\begin{aligned}
g^B &= g^A \cap \mathcal{B}_{\text{thid}} \\
&= \text{empty}_{LUH}^A(m^A, L^A, U^A, H^A, a_1^A, a_2^A, l^A) \cap \mathcal{B}_{\text{thid}} \\
&= \left\{ t' \in \mathcal{A}_{\text{thid}} \mid t' \in \text{subset}_{LUH}^A(m^A, L^A, U^A, H^A, a_1^A, a_2^A, l^A) \text{ and } \right. \\
&\quad \left. \text{subset}_{LUH}^A(m^A, L^A, U^A, H^A, L^A(t'), U^A(t'), H^A(t')) = \emptyset \right\} \cap \mathcal{B}_{\text{thid}} \\
&= \left\{ t' \in \mathcal{B}_{\text{thid}} \mid t' \in \text{subset}_{LUH}^B(m^B, L^B, U^B, H^B, a_1^B, a_2^B, l^B) \text{ and } \right. \\
&\quad \left. \text{subset}_{LUH}^B(m^A, L^B, U^B, H^B, L^B(t'), U^B(t'), H^B(t')) = \emptyset \right\} \\
&= \text{empty}_{LUH}^B(m^B, L^B, U^B, H^B, a_1^B, a_2^B, l^B)
\end{aligned}$$

Literals of the form $t = \text{min}_{LUH}(m, U, g)$:

$$\begin{aligned}
t^B &= t^A \\
&= \text{min}_{LUH}^A(m^A, U^A, g^A)
\end{aligned}$$

Then, we have two cases to consider. If $g^A = \emptyset$, then $g^B = g^A \cap \mathcal{B}_{\text{thid}} = \emptyset \cap \mathcal{B}_{\text{thid}} = \emptyset$. Besides, $\emptyset = t^A = t^B = \text{min}_{LUH}^B(m^B, U^B, g^B)$. So, let's now consider the other possible case. Let's assume $g = \{t_1, \dots, t_q\}$. Let $t_i \in g$ be a thread such that for all $j \in 1..q$, $U(t_j) \leq U(t_i)$. At $t \in V_{\text{thid}}$, $t^A = t^B = t_i^B \in \mathcal{B}_{\text{thid}}$. Besides, $g^B = g^A \cap \mathcal{B}_{\text{thid}}$ and hence $t_i^B \in g^B$. Therefore, $t_i^B = \text{min}_{LUH}(m^B, U^B, g^B)$. \square

5.4 Verifying Some Properties Over Concurrent Skiplists

In this section we show how some of the properties we would like to describe about a skiplist are translated to TSL_K, so that the decision procedure described in this chapter can be used to solve them.

5.4.1 Skiplist Preservation

In Example 3.4 we gave an idea of the functions and predicates we would require to verify skiplists presenting the verification condition for transition 43 of program INSERT. Such verification condition applied to a skiplist of height 4 was $SkipList_4(h, sl.head, sl.tail) \wedge \phi \rightarrow SkipList_4(h', sl'.head, sl'.tail)$, where ϕ is:

$$\left(\begin{array}{l} x.key = newval \quad \wedge \\ prev.key < newval \quad \wedge \\ x.next[i].key > newval \quad \wedge \\ prev.next[i] = x.next[i] \quad \wedge \\ (x, i) \notin sl.r \wedge 0 \leq i \leq 3 \end{array} \right) \wedge \left(\begin{array}{l} at_{43}^{[j]} \quad \wedge \\ prev'.next[i] = x \quad \wedge \\ at_{44}^{[j]} \quad \wedge \\ h' = h \wedge sl = sl' \quad \wedge \\ x' = x \quad \dots \end{array} \right)$$

We now show how to write this verification condition using TSL_K so that the decision procedure can be applied. We assume $x : \text{addr}$ and $prev : \text{addr}$. Besides, for the sake of simplicity, we assume variables $head : \text{addr}$, $tail : \text{addr}$ and $r : \text{mrgn}$ such that $head = sl.head$, $tail = sl.tail$ and $r = sl.r$.

Remember we are working with parametrized systems. Hence, when we use x , $prev$ or $newval$, in fact we mean the local copy of these variables for a particular thread, let's say, thread j . Then we are in fact working with $x^{[j]}$, $prev^{[j]}$ and $newval^{[j]}$ respectively. We are also assuming that proofs are done for a closed system, after the number of threads involved in the system has been fixed. This way, we assume reasoning over a finite and bounded number of threads. In more general parametrized systems, one can reason over an unbounded number of threads. In such case, a theory of arrays equipped with a decision

Original term	TSL _K term
$ordList(h, head, 0)$	$\overbrace{p = \text{getp}_K(h, head, null, 0)}^{\Sigma_{\text{bridge}}} \wedge \overbrace{ordList(h, p)}^{\Sigma_{\text{bridge}}}$
$tail.next[0] = null$ $\wedge \dots \wedge$ $tail.next[3] = null$	$\overbrace{c_{tail} = h[tail]}^{\Sigma_{\text{mem}}} \wedge \overbrace{a_{tail_0} = c_{tail}.next[0]}^{\Sigma_{\text{cell}}} \wedge \overbrace{a_{tail_0} = null}^{\Sigma_{\text{addr}}}$ $\wedge \quad \quad \quad \vdots \quad \quad \quad \wedge \quad \quad \quad \vdots$ $\overbrace{a_{tail_3} = c_{tail}.next[3]}^{\Sigma_{\text{cell}}} \wedge \overbrace{a_{tail_3} = null}^{\Sigma_{\text{addr}}}$
$SubList(h, head, tail, 1,$ $head, tail, 0)$ $\wedge \dots \wedge$ $SubList(h, head, tail, 3,$ $head, tail, 2)$	$\overbrace{p_{tail_0} = \text{getp}_K(h, head, tail, 0)}^{\Sigma_{\text{bridge}}} \wedge \dots \wedge \overbrace{p_{tail_3} = \text{getp}_K(h, head, tail, 3)}^{\Sigma_{\text{bridge}}} \wedge$ $\overbrace{stail_0 = \text{path2set}(p_{tail_0})}^{\Sigma_{\text{bridge}}} \wedge \dots \wedge \overbrace{stail_3 = \text{path2set}(p_{tail_3})}^{\Sigma_{\text{bridge}}} \wedge$ $\overbrace{stail_1 \subseteq stail_0}^{\Sigma_{\text{set}}} \wedge \dots \wedge \overbrace{stail_3 \subseteq stail_2}^{\Sigma_{\text{set}}}$

Figure 5.6: Translation of $SkipList_4(h, sl.head, sl.tail)$ into TSL_K

Original term	TSL _K term
$x.key = newval$	$\overbrace{c_x = h[x]}^{\Sigma_{mem}} \wedge \overbrace{c_x.key = newval}^{\Sigma_{cell}}$
$prev.key < newval$	$\overbrace{c_{prev} = h[prev]}^{\Sigma_{mem}} \wedge \overbrace{c_{prev}.key = k_{prev}}^{\Sigma_{cell}} \wedge$ $\overbrace{k_{prev} \preceq newval}^{\Sigma_{ord}} \wedge \overbrace{k_{prev} \neq newval}^{\Sigma_{ord}}$
$prev.next[i].key > newval$	$\overbrace{a_{next} = c_{prev}.next[i]}^{\Sigma_{cell}} \wedge \overbrace{c_{next} = h[a_{next}]}^{\Sigma_{mem}} \wedge \overbrace{k_{next} = c_{next}.key}^{\Sigma_{cell}} \wedge$ $\overbrace{newval \preceq k_{next}}^{\Sigma_{ord}} \wedge \overbrace{newval \neq k_{next}}^{\Sigma_{ord}}$
$prev.next[i] = x.next[i]$	$\overbrace{a_{next_x} = c_x.next[i]}^{\Sigma_{cell}} \wedge \overbrace{a_{next} = a_{next_x}}^{\Sigma_{addr}}$
$prev.lockid[i] = j$	$\overbrace{t_{prev} = c_{prev}.lockid[i]}^{\Sigma_{cell}} \wedge \overbrace{t_{prev} = j}^{\Sigma_{thid}}$
$x.next[i].lockid[i] = j$	$\overbrace{c_{next_x} = h[a_{next_x}]}^{\Sigma_{mem}} \wedge \overbrace{t_{next_x} = c_{next_x}.lockid[i]}^{\Sigma_{cell}} \wedge \overbrace{t_{next_x} = j}^{\Sigma_{thid}}$
$(x, i) \notin r$	$\overbrace{empty = \mathbf{emp}_{MR}}^{\Sigma_{mrgn}} \wedge \overbrace{x_{i_{set}} = \langle x, i \rangle_{MR}}^{\Sigma_{mrgn}} \wedge \overbrace{empty = r \cap_{MR} x_{i_{set}}}^{\Sigma_{mrgn}}$

We must generate four different scenarios to be combined with the conjunction of the remaining literals. These scenarios are:

- $$0 \leq i \leq 3$$
1. $\overbrace{0 < i}^{\Sigma_{level_K}} \wedge \overbrace{i < 3}^{\Sigma_{level_K}}$
 2. $\overbrace{0 < i}^{\Sigma_{level_K}} \wedge \overbrace{i = 3}^{\Sigma_{level_K}}$
 3. $\overbrace{0 = i}^{\Sigma_{level_K}} \wedge \overbrace{i < 3}^{\Sigma_{level_K}}$
 4. $\overbrace{0 = i}^{\Sigma_{level_K}} \wedge \overbrace{i = 3}^{\Sigma_{level_K}}$

Figure 5.7: Translation of enabling conditions into TSL_K

Original term	TSL_K term
$at_{43}^{[i]} \wedge at_{44}^{[i]}$	$\overbrace{at_{j_{43}}}^{\text{boolean}} \wedge \overbrace{\neg at_{j_{44}}}^{\text{boolean}} \wedge \overbrace{\neg at'_{j_{43}}}^{\text{boolean}} \wedge \overbrace{at'_{j_{44}}}^{\text{boolean}}$
$prev'.next[i] = x$	$\overbrace{a_{next_{new}} = c_{prev_{new}}.next[i]}^{\Sigma_{cell}} \wedge \overbrace{a_{next_{new}} = x}^{\Sigma_{addr}}$
$prev'.key = prev.key$	$\overbrace{k_{prev_{new}} = c_{prev_{new}}.key}^{\Sigma_{cell}} \wedge \overbrace{k_{prev} = k_{prev_{new}}}^{\Sigma_{ord}}$
$prev'.data = prev.data$	$\overbrace{e_{prev} = c_{prev}.data}^{\Sigma_{cell}} \wedge \overbrace{e_{prev_{new}} = c_{prev_{new}}.data}^{\Sigma_{cell}} \wedge \overbrace{e_{prev} = e_{prev_{new}}}^{\Sigma_{elem}}$
$prev'.next[0] = prev.next[0]$ $\wedge \dots \wedge$ $prev'.next[i-1] = prev.next[i-1]$ \wedge $prev'.next[i+1] = prev.next[i+1]$ $\wedge \dots \wedge$ $prev'.next[3] = prev.next[3]$	$\overbrace{a_{prev_0} = c_{prev}.next[0]}^{\Sigma_{cell}} \wedge \dots \wedge \overbrace{a_{prev_{i-1}} = c_{prev}.next[i-1]}^{\Sigma_{cell}} \wedge$ $\overbrace{a_{prev_{i+1}} = c_{prev}.next[i+1]}^{\Sigma_{cell}} \wedge \dots \wedge \overbrace{a_{prev_3} = c_{prev}.next[3]}^{\Sigma_{cell}} \wedge$ $\overbrace{a_{prev_{new_0}} = c_{prev_{new}}.next[0]}^{\Sigma_{cell}} \wedge \dots \wedge \overbrace{a_{prev_{new_{i-1}}} = c_{prev_{new}}.next[i-1]}^{\Sigma_{cell}} \wedge$ $\overbrace{a_{prev_{new_{i+1}}} = c_{prev_{new}}.next[i+1]}^{\Sigma_{cell}} \wedge \dots \wedge \overbrace{a_{prev_{new_3}} = c_{prev_{new}}.next[3]}^{\Sigma_{cell}} \wedge$ $\overbrace{a_{prev_{new_0}} = a_{prev_0}}^{\Sigma_{addr}} \wedge \dots \wedge \overbrace{a_{prev_{new_{i-1}}} = a_{prev_{i-1}}}^{\Sigma_{addr}} \wedge$ $\overbrace{a_{prev_{new_{i+1}}} = a_{prev_{i+1}}}^{\Sigma_{addr}} \wedge \dots \wedge \overbrace{a_{prev_{new_3}} = a_{prev_3}}^{\Sigma_{addr}} \wedge$ $\overbrace{h' = upd(h, prev, c_{prev_{new}})}^{\Sigma_{mem}}$

Figure 5.8: Translation of the modifications introduced by the transition relation into TSL_K

procedure [BMS06] can be used to keep the values of all local variables and program counters as arrays. Moreover, such theory is easily combined with TSL_K . However, this approach goes beyond the scope of this work.

Fig. 5.6 contains the translation of the $SkipList_4$ predicate into literals of TSL_K . The translation for predicate $SkipList_4(h', sl'.head, sl'.tail)$ is very similar to $SkipList_4(h, sl.head, sl.tail)$ and thus it is omitted. The difference between them resides only on the modifications introduced by the transition relation, stating $SkipList_4(h', sl'.head, sl'.tail)$ that the structure keeps the shape of a skiplist after the transition has been taken. Fig. 5.7 describes the translation of the enabling conditions for the presented verification condition into TSL_K . Similarly, Fig. 5.8 depicts the translation of the modifications produced by the transition. In all cases, it is straightforward to realize that the resulting literals are all in TSL_K and thus it is possible to check their satisfiability using the decision procedure described in this chapter. For the moment, this translation must be accomplished manually. Future work includes the construction of a

tool to automatically translate formulas into TSL_K .

When describing the new literals in TSL_K obtained from the verification condition, some fresh variables must be introduced. To describe them we use a for variables of sort `addr`, k for variables of sort `ord`, e for variables of sort `elem`, t for variables of sort `thid`, c for variables of sort `cell`, p for variables of sort `path` and s for variables of sort `set`.

5.4.2 Termination of an Arbitrary Thread

When we introduced concurrent skiplists in Chapter 3, we showed that some implementations do not guarantee termination of all threads if we do not assume a fair scheduler. Because of this, we introduced a new implementation we called *pessimistic*. Now, we sketch how the termination of an arbitrary thread running the pessimistic implementation of a concurrent skiplist (presented in Section 3.2) could be proved.

When we gave the implementations, we enriched the algorithms with ghost arrays L , U and H , to define the (L, U, H) section for each running thread. Remember that (L, U, H) represents the maximum portion of the skiplist we are sure a thread can potentially modify.

The pessimistic version left a locked node behind, to prevent threads to overtake other threads. First of all, we extend the definition of skiplist to capture the property that, in pessimistic skiplists, two (L, U, H) belonging to different threads are disjoint or one is strictly contained into the other. For such purpose, we define

$$\begin{aligned} \text{Skiplist}_K^{\text{pessi}}(h, sl : \text{Skiplist}) &\hat{=} \text{SkipList}_K(h, sl.\text{head}, sl.\text{tail}) \wedge \\ &\emptyset_v = \text{inter}_{LUH}(m, L, U, H, \text{subset}_{LUH}(m, L, U, H, sl.\text{head}, sl.\text{tail}, K - 1)) \end{aligned}$$

The new predicate $\text{Skiplist}_K^{\text{pessi}}(h, sl)$ says that in heap h , sl points to a structure with the shape of a skiplist and no (L, U, H) within the skiplist partially overlaps with other thread's (L, U, H) .

Imagine we want to prove that an arbitrary thread k terminates. Here we do not give the diagram, but just the intuition behind the proof. Let $(L, U, H)_k$ be the (L, U, H) of thread k . A key property is that the number of threads whose (L, U, H) is contained into $(L, U, H)_k$ does not increment. In fact, we can use the cardinality of $\text{subset}_{LUH}(m, L, U, H, L(k), U(k), H(k))$ as a ranking function to ensure that no new thread is entering $(L, U, H)_k$. Besides, we can ensure that it always exists a thread that has no contention when trying to progress. Such thread is the one with the minimum (L, U, H) within $(L, U, H)_k$, that is, thread t_{\min} where

$$t_{\min} = \text{min}_{LUH}(m, U, \text{empty}_{LUH}(m, L, U, H, L(k), U(k), H(k), K - 1)).$$

Conclusion

In this work, we have presented combinable decision procedures for concurrent lists and skiplists. The development of such decision procedures is motivated by the necessity of automating the verification process of temporal properties (safety and liveness) of an imperative implementation of concurrent datatypes. Although we have focused on concurrent lists and skiplists, the approach can be applied to the verification of more complex concurrent data structures.

As shown, the verification is performed using verification diagrams – a complete method to prove temporal properties of reactive systems – and explicit reasoning on memory regions. The verification process usually requires the aid of ghost variables. Verification is reduced to proving a finite number of verification conditions, which requires decision procedures in the appropriate theories, including regions, pointers, locks and specific theories for memory layouts (in this case single linked-lists and skiplists).

There are some key differences between our approach and other approaches in the literature. Building on the success of separation logic in proving sequential programs, the most popular approach has consisted in extending separation logic to concurrent programs. These extensions require adapting techniques like rely-guarantee that cannot be directly used with separation logic. Our decision to use explicit regions (finite sets of addresses) allows the direct use of classical techniques like assume-guarantee and the combination of decision procedures. Furthermore, in concurrent separation logic, it is critical to describe memory footprints of sections of code. This description becomes very cumbersome when the code is not organized in mutual exclusion regions, as in fine-grain synchronization algorithms. Moreover, the integration into SMT solvers is quite straightforward with classical logics, but it is still an open question with separation logic.

The technique we propose can be seen as a method to separate the reasoning about concurrency (with verification diagrams) from the reasoning about the memory (with decision procedures). The former is independent of the data structure under consideration. We are currently extending our approach to the verification of other pointer-based concurrent data structures like concurrent hash maps. Again, sharing in these data structures makes it very hard to reason using separation logic. For our approach, these extensions will require the design of suitable decision procedures.

We have also presented TLL3 and TSL_K, theories of concurrent singly linked lists and concurrent skiplists of height at most K respectively. We showed both of them are useful for automatically proving the VCs generated during the verification of concurrent list and skiplist implementations.

TLL3 is capable of reasoning about memory, cells, pointers and reachability. TSL_K can also reason about masked regions, enabling ordered lists and sublists, allowing the description of the skiplist property as well as the representation of memory modifications introduced by the execution of program statements.

We showed both theories decidable by proving their finite model property and exhibiting the minimal cardinality of a model if one such model exists. Moreover, we showed how to reduce the satisfiability problem of quantifier-free TLL3 and TSL_K formulas to a combination of theories using the many-sorted

version of Nelson-Oppen, allowing the use of well studied decision procedures. The complexity of the decision problem for TLL3 and TSL_K is easily shown to be NP-complete since they properly extend TLL [RZ06a].

Current work includes among other things the translation of formulas from T_{ord} , T_{level_k} , T_{set} , T_{setid} and T_{mrgn} into BAPA [KNR05]. In BAPA, arithmetic, sets and cardinality aid in the definition of skiplists properties. At the same time, paths can be represented as finite sequences of addresses. We are studying how to replace the recursive functions from T_{reach} and Σ_{bridge} by canonical set and list abstractions [SDK10], which would lead to a more efficient decision procedure, essentially encoding full TLL3 and TSL_K formulas into BAPA.

In the case of TSL_K , the family of theories presented in the work is limited to skiplists of a fixed maximum height. Typical skiplist implementations fix a maximum number of levels and this can be handled with TSL_K . Inserting more than 2^{levels} elements into a skiplist may slow-down the search of a skiplist implementation but this issue affects performance and not correctness, which is the goal pursued in this paper. Following this line, we are studying techniques to describe skiplists of arbitrary many levels. A promising approach consists of reducing a formula describing an unbounded number of levels to a formula grounded on a finite number of them, reducing the satisfiability problem to TSL_K . This approach, however, is still work in progress.

Future work also includes building a generic verification condition generator for verification diagrams, implementing an ad-hoc version of the decision procedure described here, and later integrating this decision procedure into state-of-the-art SMT solvers. We are also interested in the temporal verification of sequential and concurrent skiplists implementations, including one at the `java.concurrent` standard library. This can be accomplished by the design of verification diagrams that use the decision procedure presented in this paper.

Bibliography

- [AC89] Anant Agarwal and Mathews Cherian. Adaptive backoff synchronization techniques. In *Proc. of ISCA 1989*, pages 396–406, 1989.
- [AHS94] James Aspnes, Maurice Herlihy, and Nir Shavit. Counting networks. *Journal of the ACM*, 41(5):1020–1048, 1994.
- [And89] Thomas E. Anderson. The performance implications of spin-waiting alternatives for shared-memory multiprocessors. In *Proc. of ICPP 1989*, pages 170–174, 1989.
- [ARR03] Alessandro Armando, Silvio Ranise, and Michaël Rusinowitch. A rewriting approach to satisfiability procedures. *Information and Computation*, 183(2):140–164, 2003.
- [AVL62] G. Adel’son-Vel’skii and E. Landis. An algorithm for the organization of information. *Doklady Akademii Nauk USSR*, 146(2):263–266, 1962.
- [Aya90] Rassul Ayani. Lr-algorithm: concurrent operations on priority queues. In *Proc. of SPDP 1990*, pages 22–25, 1990.
- [Bar94] G. Barnes. Wait free algorithms for heaps. Technical report, University of Washington, 1994.
- [BB87] Jit Biswas and James C. Browne. Simultaneous update of priority structures. In *Proc. of ICPP 1987*, pages 124–131, 1987.
- [BCG01] Guy E. Blelloch, Perry Cheng, and Phillip B. Gibbons. Room synchronizations. In *Proc. of SPAA 2001*, pages 122–133, 2001.
- [BCO04] Berdine, Calcagno, and O’Hearn. A decidable fragment of separation logic. *FSTTCS: Foundations of Software Technology and Theoretical Computer Science*, 24, 2004.
- [BDES09] Ahmed Bouajjani, Cezara Dragoi, Constantin Enea, and Mihaela Sighireanu. A logic-based framework for reasoning about composite data structures. In *CONCUR’09*, pages 178–195, 2009.
- [BHG87] P. Bernstein, V. Hadzilacos, and N. Goodman. *Concurrency Control and Recovery in Database Systems*. Addison-Wesley, 1987.
- [BL94] Joan Boyar and Kim S. Larsen. Efficient rebalancing of chromatic search trees. *J. Comput. Syst. Sci.*, 49(3):667–682, 1994.
- [BM07] Aaron R. Bradley and Zohar Manna. *The Calculus of Computation*. Springer-Verlag, 2007.
- [BMS95] Anca Browne, Zohar Manna, and Henny B. Sipma. Generalized verification diagrams. In *Proc. of FSTTCS’95*, volume 1206 of *LNCS*, pages 484–498. Springer, 1995.

- [BMS06] Aaron R. Bradley, Zohar Manna, and Henny B. Sipma. What's decidable about arrays? In *Proc. of VMCAI 2006*, volume 3855 of *LNCS*, pages 427–442. Springer, 2006.
- [BNR08] Anindya Banerjee, David A. Naumann, and Stan Rosenberg. Regional logic for local reasoning about global invariants. In *Proc. of ECOOP'08*, pages 387–411. Springer, 2008.
- [BPZ05] Ittai Balaban, Amir Pnueli, and Lenore D. Zuck. Shape analysis by predicate abstraction. In *Proc. of VMCAI 2005*, pages 164–180, 2005.
- [BR06] Jesse D. Bingham and Zvonimir Rakamaric. A logic and decision procedure for predicate abstraction of heap-manipulating programs. In *Proc. of VMCAI 2006*, pages 207–221, 2006.
- [BRS99] Michael Benedikt, Thomas W. Reps, and Shmuel Sagiv. A decidable logic for describing linked data structures. In *Proc. of ESOP 1999*, pages 2–19, 1999.
- [BS77] Rudolf Bayer and Mario Schkolnick. Concurrency of operations on B-trees. *Acta Informatica*, 9:1–21, 1977.
- [BSST08] Clark Barrett, Roberto Sebastiani, Sanjit A. Seshia, and Cesar Tinelli. *Handbook of Satisfiability*, chapter Satisfiability Modulo Theories. IOS Press, 2008.
- [Bur72] R. M. Burstall. Some techniques for proving correctness of programs which alter data structures. *Machine Intelligence*, 7:23–50, 1972.
- [CLRS01] T. Cormen, C. Leiserson, R. Rivest, and C. Stein. *Introduction to Algorithms*. MIT Press, 2001.
- [Cra93] Travis Craig. Building fifo and priority-queuing spin locks from atomic swap. Technical report, University of Washington, Department of Computer Science, 1993.
- [DN03] Dennis Dams and Kedar S. Namjoshi. Shape analysis through predicate abstraction and model checking. In *Proc. of VMCAI 2003*, pages 310–324, 2003.
- [DS80] Edsger W. Dijkstra and Carel S. Scholten. Termination detection for diffusing computations. *Inf. Process. Lett.*, 11(1):1–4, 1980.
- [Ell87] Carla Schlatter Ellis. Concurrency in linear hashing. *ACM Trans. Database Syst.*, 12(2):195–217, 1987.
- [FLP85] Michael J. Fischer, Nancy A. Lynch, and Mike Paterson. Impossibility of distributed consensus with one faulty process. *Journal of the ACM*, 32(2):374–382, 1985.
- [FNPS79] Ronald Fagin, Jürg Nievergelt, Nicholas Pippenger, and H. Raymond Strong. Extendible hashing - a fast access method for dynamic files. *ACM Trans. Database Syst.*, 4(3):315–344, 1979.
- [Fra03] K. Fraser. *Practical Lock-Freedom*. PhD thesis, University of Cambridge, 2003.
- [FT87] Michael L. Fredman and Robert Endre Tarjan. Fibonacci heaps and their uses in improved network optimization algorithms. *J. ACM*, 34(3):596–615, 1987.
- [GCPV09] Alexey Gotsman, Byron Cook, Matthew J. Parkinson, and Viktor Vafeiadis. Proving that non-blocking algorithms don't block. In Zhong Shao and Benjamin C. Pierce, editors, *Proceedings of the 36th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2009, Savannah, GA, USA, January 21-23, 2009*, pages 16–28. ACM, 2009.
- [GGK⁺83] Allan Gottlieb, Ralph Grishman, Clyde P. Kruskal, Kevin P. McAuliffe, Larry Rudolph, and Marc Snir. The nyu ultracomputer - designing an mimd shared memory parallel computer. *IEEE Trans. Computers*, 32(2):175–189, 1983.

- [Gre02] Michael Greenwald. Two-handed emulation: how to build non-blocking implementation of complex data-structures using dcas. In *Proc. of PODC 2002*, pages 260–269, 2002.
- [GS78] Leonidas J. Guibas and Robert Sedgewick. A dichromatic framework for balanced trees. In *Proc. of FOCS 1978*, pages 8–21, 1978.
- [GT90] Gary Graunke and Shreekanth S. Thakkar. Synchronization algorithms for shared-memory multiprocessors. *IEEE Computer*, 23(6):60–69, 1990.
- [GVW89] James R. Goodman, Mary K. Vernon, and Philip J. Woest. Efficient synchronization primitives for large-scale cache-coherent multiprocessors. In *Proc. of ASPLOS 1989*, pages 64–75, 1989.
- [HAN08] Aquinas Hobor, Andrew W. Appel, and Francesco Zappa Nardelli. Oracle semantics for concurrent separation logic. In *Proc. of ESOP'08*, volume 4960 of *LNCS*, pages 353–367. Springer, 2008.
- [Har01] Timothy L. Harris. A pragmatic implementation of non-blocking linked-lists. In *Proc. of DISC 2001*, pages 300–314, 2001.
- [Her91] Maurice Herlihy. Wait-free synchronization. *ACM Trans. Program. Lang. Syst.*, 13(1):124–149, 1991.
- [Her93] Maurice Herlihy. A methodology for implementing highly concurrent objects. *ACM Trans. Program. Lang. Syst.*, 15(5):745–770, 1993.
- [HF03] Timothy L. Harris and Keir Fraser. Language support for lightweight transactions. In *Proc. of OOPSLA 2003*, pages 388–402, 2003.
- [HFM88] Debra Hensgen, Raphael Finkel, and Udi Manber. Two algorithms for barrier synchronization. *Int. J. Parallel Program.*, 17(1):1–17, 1988.
- [HFP02] Timothy L. Harris, Keir Fraser, and Ian A. Pratt. A practical multi-word compare-and-swap operation. In *Proc. of DISC 2002*, pages 265–279, 2002.
- [HLM02] Maurice Herlihy, Victor Luchangco, and Mark Moir. The repeat offender problem: A mechanism for supporting dynamic-sized, lock-free data structures. In *Proc. of DISC 2002*, pages 339–353, 2002.
- [HLM03] Maurice Herlihy, Victor Luchangco, and Mark Moir. Obstruction-free synchronization: Double-ended queues as an example. In *Proc. of ICDCS 2003*, pages 522–529, 2003.
- [HLMi03] Maurice Herlihy, Victor Luchangco, Mark Moir, and William N. Scherer III. Software transactional memory for dynamic-sized data structures. In *Proc. of PODC 2003*, pages 92–101, 2003.
- [HLS95] Maurice Herlihy, Beng-Hong Lim, and Nir Shavit. Scalable concurrent counting. *ACM Trans. Comput. Syst.*, 13(4):343–364, 1995.
- [HM93] Maurice Herlihy and J. Eliot B. Moss. Transactional memory: Architectural support for lock-free data structures. In *Proc. of ISCA 1993*, pages 289–300, 1993.
- [HMPS96] Galen C. Hunt, Maged M. Michael, Srinivasan Parthasarathy, and Michael L. Scott. Concurrent access of priority queues. *Inf. Process. Lett.*, 60(3):151–157, 1996.
- [HS02] Danny Hendler and Nir Shavit. Work dealing. In *Proc. of SPAA 2002*, pages 164–172, 2002.
- [HS08] Maurice Herlihy and Nir Shavit. *The Art of Multiprocessor Programming*. Morgan-Kaufmann, 2008.

- [HSY04] Danny Hendler, Nir Shavit, and Lena Yerushalmi. A scalable lock-free stack algorithm. In *Proc. of SPAA 2004*, pages 206–215, 2004.
- [HW90] Maurice Herlihy and Jeannette M. Wing. Linearizability: A correctness condition for concurrent objects. *ACM Trans. Program. Lang. Syst.*, 12(3):463–492, 1990.
- [HW91] Qin Huang and William E. Weihl. An evaluation of concurrent priority queue algorithms. In *Proc. of SPDP 1991*, pages 518–525, 1991.
- [HY86] Meichun Hsu and Wei-Pang Yang. Concurrent operations in extendible hashing. In *Proc. of VLDB 1986*, pages 241–247, 1986.
- [IBM] IBM. System/370 principles of operation. order number ga22-7000.
- [IBM03] IBM. Powerpc microprocessor family: Programming environments manual for 64 and 32-bit microprocessors, version 2.0, 2003.
- [III86] Eugene D. Brooks III. The butterfly barrier. *Int. J. Parallel Program.*, 15(4):295–307, 1986.
- [Int94] Intel. Pentium processor family user’s manual: Vol 3, architecture and programming manual, 1994.
- [IR93] Amos Israeli and Lihu Rappoport. Efficient wait-free implementation of a concurrent priority queue. In *Proc. of WDAG 1993*, pages 1–17, 1993.
- [JJKS97] Jakob L. Jensen, Michael E. Jørgensen, Nils Klarlund, and Michael I. Schwartzbach. Automatic verification of pointer programs using monadic second-order logic. In *Proc. of PLDI 1997*, pages 226–236, 1997.
- [Joh91] T. Johnson. A highly concurrent priority queue based on the b-link tree. Technical report, University of Florida, 1991.
- [Jou02] Yuh-Jzer Joung. Nonblocking algorithms and preemption-safe locking on multiprogrammed shared memory multiprocessors. *Distributed Computing*, 15(3):155–175, 2002.
- [Kan89] G. Kane. *MIPS RISC Architecture*. Prentice-Hall, Inc., New York, USA, 1989.
- [Kes83] Joep L. W. Kessels. On-the-fly optimization of data structures. *Commun. ACM*, 26(11):895–901, 1983.
- [KL80] H. T. Kung and Philip L. Lehman. Concurrent manipulation of binary search trees. *ACM Trans. Database Syst.*, 5(3):354–382, 1980.
- [KM99] Patrick Keane and Mark Moir. A simple local-spin group mutual exclusion algorithm. In *Proc. of PODC 1999*, pages 23–32, 1999.
- [KMH97] Marcel Kornacker, C. Mohan, and Joseph M. Hellerstein. Concurrency and recovery in generalized search trees. In *Proc. of SIGMOD 1997*, pages 62–72, 1997.
- [KNR05] Viktor Kuncak, Huu Hai Nguyen, and Martin C. Rinard. An algorithm for deciding BAPA: Boolean Algebra with Presburger Arithmetic. In *CADE’05*, pages 260–277, 2005.
- [Knu68] D. Knuth. *The Art of Computer Programming: Fundamental Algorithms*. Addison-Wesley, 1968.
- [KR81] H. T. Kung and John T. Robinson. On optimistic methods for concurrency control. *ACM Trans. Database Syst.*, 6(2):213–226, 1981.
- [KSUH93] Orran Krieger, Michael Stumm, Ronald C. Unrau, and Jonathan Hanna. A fair fast scalable reader-writer lock. In *Proc. of ICPP 1993*, pages 201–204, 1993.

- [Kum90] Vijay Kumar. Concurrent operations on extendible hashing and its performance. *Commun. ACM*, 33(6):681–694, 1990.
- [LAA87] M. Loui and H. Abu-Amara. Memory requirements for agreement among unreliable asynchronous processes. *Advances in Computing Research*, 4:163–183, 1987.
- [Lam74] Leslie Lamport. A new solution of dijkstra’s concurrent programming problem. *Commun. ACM*, 17(8):453–455, 1974.
- [Lam79] Leslie Lamport. How to make a multiprocessor computer that correctly executes multiprocess programs. *IEEE Trans. Computers*, 28(9):690–691, 1979.
- [Lam83] Leslie Lamport. Specifying concurrent program modules. *ACM Trans. Program. Lang. Syst.*, 5(2):190–222, 1983.
- [Lar01] Kim S. Larsen. Relaxed multi-way trees with group updates. In *Proc. of PODS 2001*, pages 93–101, 2001.
- [Lar02] Kim S. Larsen. Relaxed red-black trees with group updates. *Acta Inf.*, 38(8):565–586, 2002.
- [Lea99] D. Lea. *Concurrent Programming in JAVA(TM): Design Principles and Pattern*. Addison-Wesley, 1999.
- [Lei98] K. Rustan M. Leino. Data groups: Specifying the modification of extended state. In *OOP-SLA’98*, pages 144–153. ACM, 1998.
- [LF95] Kim S. Larsen and Rolf Fagerberg. B-trees with relaxed balance. In *Proc. of IPPS 1995*, pages 196–202, 1995.
- [Lit80] Witold Litwin. Linear hashing: A new tool for file and table addressing. In *Proc. of VLDB 1980*, pages 212–223, 1980.
- [LQ06] Shuvendu K. Lahiri and Shaz Qadeer. Verifying properties of well-founded linked lists. In *Proc. of POPL 2006*, pages 115–126, 2006.
- [LQ08] Shuvendu K. Lahiri and Shaz Qadeer. Back to the future: revisiting precise program verification using smt solvers. In *Proc. of POPL’08*, pages 171–182. ACM, 2008.
- [Man84] Udi Manber. On maintaining dynamic information in a concurrent environment. In *Proc. of STOC 1984*, pages 273–278, 1984.
- [McM99] Kenneth L. McMillan. Circular compositional reasoning about liveness. In *Proc. of CHARME’99*, volume 1703 of LNCS, pages 342–345. Springer, 1999.
- [MCS91a] John M. Mellor-Crummey and Michael L. Scott. Algorithms for scalable synchronization on shared-memory multiprocessors. *ACM Trans. Comput. Syst.*, 9(1):21–65, 1991.
- [MCS91b] John M. Mellor-Crummey and Michael L. Scott. Scalable reader-writer synchronization for shared-memory multiprocessors. In *Proc. of PPOPP 1991*, pages 106–113, 1991.
- [Mic02a] Maged M. Michael. High performance dynamic lock-free hash tables and list-based sets. In *Proc. of SPAA 2002*, pages 73–82, 2002.
- [Mic02b] Maged M. Michael. Safe memory reclamation for dynamic lock-free objects using atomic reads and writes. In *Proc. of PODC 2002*, pages 21–30, 2002.
- [ML92] C. Mohan and Frank E. Levine. Aries/im: An efficient and high concurrency index management method using write-ahead logging. In *Proc. of SIGMOD 1992*, pages 371–380, 1992.

- [MLH94] Peter S. Magnusson, Anders Landin, and Erik Hagersten. Queue locks on cache coherent multiprocessors. In *Proc. of IPPS 1994*, pages 165–171, 1994.
- [MMS02] P. Martin, M. Moir, and G. Steele. Dcas-based concurrent dequeues supporting bulk allocation. Technical report, Sun Microsystems Laboratories, 2002.
- [MN05] Scott McPeak and George C. Necula. Data structure specifications via local equality axioms. In *Proc. of CAV 2005*, pages 476–490, 2005.
- [Mot86] Motorola. Mc68020 32-bit microprocessor user’s manual, 1986.
- [MP95] Zohar Manna and Amir Pnueli. *Temporal Verification of Reactive Systems*. Springer, 1995.
- [MS95] M. Michael and M. Scott. Correction of a memory management method for lock-free data structures. Technical report, University of Rochester, 1995.
- [MS98] Maged M. Michael and Michael L. Scott. Nonblocking algorithms and preemption-safe locking on multiprogrammed shared memory multiprocessors. *J. Parallel Distrib. Comput.*, 51(1):1–26, 1998.
- [MS07] M. Moir and N. Shavit. Concurrent data structures. In *Handbook of Data Structures and Applications*, pages 47–14 — 47–30, 2007. Chapman and Hall/CRC Press.
- [Nel83] Greg Nelson. Verifying reachability invariants of linked structures. In *Proc. of POPL 1983*, pages 38–47, 1983.
- [NO79] Greg Nelson and Derek C. Oppen. Simplification by cooperating decision procedures. *ACM Trans. Program. Lang. Syst.*, 1(2):245–257, 1979.
- [NSS91] Otto Nurmi and Eljas Soisalon-Soininen. Uncoupling updating and rebalancing in chromatic binary search trees. In *Proc. of PODS 1991*, pages 192–198, 1991.
- [NSSW87] Otto Nurmi, Eljas Soisalon-Soininen, and Derick Wood. Concurrency control in database structures with relaxed balance. In *Proc. of PODS 1987*, pages 170–176, 1987.
- [Opp80] Derek C. Oppen. Reasoning about recursively defined data structures. *J. ACM*, 27(3):403–411, 1980.
- [PLJ91] Sundeep Prakash, Yann-Hang Lee, and Theodore Johnson. A non-blocking algorithm for shared queues using compare-and-swap. In *Proc. of ICPP 1991*, pages 68–75, 1991.
- [Pug90] William Pugh. Skip lists: A probabilistic alternative to balanced trees. *Commun. ACM*, 33(6):668–676, 1990.
- [Rey02] John C. Reynolds. Separation logic: A logic for shared mutable data structures. In *Proc. of LICS’02*, pages 55–74. IEEE CS Press, 2002.
- [RG01] Ravi Rajwar and James R. Goodman. Speculative lock elision: enabling highly concurrent multithreaded execution. In *Proc. of MICRO 2001*, pages 294–305, 2001.
- [RG02] Ravi Rajwar and James R. Goodman. Transactional lock-free execution of lock-based programs. In *Proc. of ASPLOS 2002*, pages 5–17, 2002.
- [RK88] V. Nageshwara Rao and Vipin Kumar. Concurrent access of priority queues. *IEEE Trans. Computers*, 37(12):1657–1665, 1988.
- [RRZ05] Silvio Ranise, Christophe Ringeissen, and Calogero G. Zarba. Combining data structures with nonstably infinite theories using many-sorted logic. In *FRODOS’05*, pages 48–64, 2005.
- [RZ06a] Silvio Ranise and Calogero G. Zarba. A theory of singly-linked lists and its extensible decision procedure. In *Proc. of SEFM 2006*. IEEE CS Press, 2006.

- [RZ06b] Silvio Ranise and Calogero G. Zarba. A theory of singly-linked lists and its extensible decision procedure. In *Technical report*, 2006.
- [Sag86] Yehoshua Sagiv. Concurrent operations on b*-trees with overtaking. *J. Comput. Syst. Sci.*, 33(2):275–296, 1986.
- [Sco02] Michael L. Scott. Non-blocking timeout in scalable queue-based spin locks. In *Proc. of PODC 2002*, pages 31–40, 2002.
- [SDK10] Philippe Suter, Mirco Dotta, and Viktor Kuncak. Decision procedures for algebraic data types with abstractions. In *Proc. of POPL'10*, pages 199–210. ACM, 2010.
- [SI94] CORPORATE SPARC International, Inc. *The SPARC architecture manual (version 9)*. Prentice-Hall, Inc., Upper Saddle River, NJ, USA, 1994.
- [SI01] Michael L. Scott and William N. Scherer III. Scalable queue-based spin locks with timeout. In *Proc. of PPOPP 2001*, pages 44–52, 2001.
- [Sip99] Henny B. Sipma. *Diagram-Based Verification of Discrete, Real-Time and Hybrid Systems*. PhD thesis, Stanford University, 1999.
- [Sit92] R. Sites. Alpha architecture reference manual, 1992.
- [SMC92] M. L. Scott and J. M. Mellor-Crummey. Fast, contention-free combining tree barriers. Technical report, University of Rochester, Rochester, NY, USA, 1992.
- [SS03] Ori Shalev and Nir Shavit. Split-ordered lists: lock-free extensible hash tables. In *Proc. of PODC 2003*, pages 102–111, 2003.
- [SS10] Alejandro Sánchez and César Sánchez. Decision procedures for the temporal verification of concurrent lists. In *Proc. of ICFEM'10*, volume 6447 of *LNCS*, pages 74–89. Springer, 2010.
- [SS11] Alejandro Sánchez and César Sánchez. A theory of skiplists with applications to the verification of concurrent datatypes. In *Proc. of NFM'11*, volume 6617 of *LNCS*, pages 343–358. Springer, 2011.
- [ST97a] Nir Shavit and Dan Touitou. Elimination trees and the construction of pools and stacks. *Theory Comput. Syst.*, 30(6):645–670, 1997.
- [ST97b] Nir Shavit and Dan Touitou. Software transactional memory. *Distributed Computing*, 10(2):99–116, 1997.
- [ST03] Håkan Sundell and Philippas Tsigas. Fast and lock-free concurrent priority queues for multi-thread systems. In *Proc. of IPDPS 2003*, page 84, 2003.
- [SZ96] Nir Shavit and Asaph Zemach. Diffracting trees. *ACM Trans. Comput. Syst.*, 45(4):385–428, 1996.
- [SZ99] Nir Shavit and Asaph Zemach. Scalable concurrent priority queue algorithms. In *Proc. of PODC 1999*, pages 113–122, 1999.
- [Tar51] Alfred Tarski. A decision method for elementary algebra and geometry. *University of California Press*, 1951.
- [Tre86] R. K. Treiber. Systems programming: Coping with parallelism. Technical report, IBM Almaden Research Center, 1986.
- [TZ04] Cesare Tinelli and Calogero G. Zarba. Combining decision procedures for sorted theories. In *JELIA'04*, volume 3229 of *LNCS*, pages 641–653. Springer, 2004.

-
- [Vaf07] Viktor Vafeiadis. *Modular fine-grained concurrency verification*. PhD thesis, University of Cambridge, 2007.
- [Val95] John D. Valois. Lock-free linked lists using compare-and-swap. In *Proc. of PODC 1995*, pages 214–222, 1995.
- [VHHS06] Viktor Vafeiadis, Maurice Herlihy, Tony Hoare, and Marc Shapiro. Proving correctness of highly-concurrent linearisable objects. In *Proc of PPOPP'06*, pages 129–136. ACM, 2006.
- [WPK09] Thomas Wies, Ruzica Piskac, and Viktor Kuncak. Combining theories with shared set operations. In *Proc. of Frontiers of Combining Systems (FroCoS'09)*, volume 5749 of LNCS, pages 366–382. Springer, 2009.
- [YRS⁺06] Greta Yorsh, Alexander Moshe Rabinovich, Mooly Sagiv, Antoine Meyer, and Ahmed Bouajjani. A logic of reachable patterns in linked data-structures. In *FOSSACS'06*, pages 94–110, 2006.
- [YTL87] Pen-Chung Yew, Nian-Feng Tzeng, and Duncan H. Lawrie. Distributing hot-spot addressing in large-scale multiprocessors. *IEEE Trans. Computers*, 36(4):388–395, 1987.
- [Zar03] Calogero G. Zarba. Combining sets with elements. In *Verification: Theory and Practice*, volume 2772 of LNCS, pages 762–782. Springer, 2003.

