

Reasoning About Lock Placements

Peter Hawkins, Alex Aiken*, Kathleen Fisher**, Martin Rinard, and Mooly Sagiv

Stanford University, Tufts University, MIT, Tel Aviv University

Abstract. A *lock placement* describes, for each heap location, which lock guards the location, and under what circumstances. We formalize methods for reasoning about lock placements, making precise the interactions between the program, the heap structure, and the lock placement.

1 Introduction

Most concurrent software uses *locks* as a primitive for ensuring mutual exclusion between threads. While it is correct to say that the key characteristic of a lock is that it may be held by only one thread at a time, such a description fails to capture the higher-level purposes for which programmers use locks. Universally, locks are used to protect data, guaranteeing that only one thread operates on particular parts of the store at a time. The association between locks and the data they protect is, however, implicit, and in the presence of mutable data structures it is not even clear how to describe the relationship between a possibly changing set of locks and the changing heap the locks protect.

This paper investigates what it means for locks to protect data. So far as we are aware, there are no proposals in the literature for even stating the relationship between locks and the data they protect that capture the range of ways in which locks are used in practice. In particular, we are interested in explaining *speculative locks* and the common case in which updates to the heap change which data locks protect. We believe ours is the first proposal to address these issues.

To explain our results, we begin with a slightly informal, simple, obviously correct, but impractical locking protocol. We assume the heap consists of a graph of *objects* (nodes), each of which has a set of *fields* (edges) that point to other objects. We also assume that concurrent operations are expressed as transactions that execute atomically (e.g., atomic blocks). Every heap edge has a *logical lock*. Each transaction t must obey a standard two-phase locking protocol:

- Acquire all logical locks of every edge read or written by t .
- Perform the reads and writes of t .
- Release all of t 's logical locks.

It is a classic result [8] that any interleaving of such transactions is *serializable* (equivalent to some sequential schedule of the transactions). However in practice

* This work was supported by NSF grants CCF-0702681 and CNS-050955.

** The views expressed are those of the author and do not reflect the official policy or position of the Department of Defense or U.S. Government. Distribution Statement A (Approved for Public Release, Distribution Unlimited)

acquiring a separate lock for every field of every object touched by a transaction is exorbitantly expensive. Thus, practical locking protocols use fewer locks. For example, a tree data structure might have a single lock at the root, or a hash table may have one lock per hash bucket, with no locks on the bucket contents.

The key insight is that the programmer has made an optimization: many logical locks are represented by a single *physical lock*. We can still think of a transaction as acquiring all of the logical locks required, but now instead of acquiring the lock on the actual edge e it must instead acquire the physical lock $\psi(e)$ assigned to the edge by a *lock placement* ψ , which is a mapping from logical locks to the physical locks that implement them. For example, in the tree case $\psi(e) = \rho$ for every edge in the tree, where ρ is the tree’s root. For the hash table, $\psi(e) = l_i$, where the i -th bucket has an associated physical lock l_i for every field e in the i -th bucket. If the same physical locks represents multiple logical locks then transactions need only acquire the single physical lock to obtain access to multiple heap locations.

Lock placements capture common idioms for programming with locks:

- Locking at different granularities corresponds to different lock placements. For example, each element of a tree may have its own lock, or there may be a single coarse-grained lock. Lock placements make explicit which locations are guarded by the same lock, and where that lock is placed.
- It is sometimes beneficial to place the lock guarding an object o in a field of o itself, which means that o cannot be locked without first accessing o in an unlocked state. Lock placements can describe such *speculative* locking, allowing us to reason about transactions that make use of it.
- Which locks guard which fields often changes over time. As a simple example, consider a heap in which all `nil` fields are guarded by a global lock, and all non-`nil` fields are guarded by a speculative lock in the object the field points to. When a `nil` field is assigned an object the global lock is *split* and no longer guards the field, and when a pointer field is assigned `nil` that field is *merged* into the global lock. Lock placements can depend on the state of the heap and so naturally capture lock splitting and merging.

We develop our results incrementally, beginning with flat “heaps” that are just a set of global variables with no pointers (Section 2). In this simple setting we formalize the key notions of lock placements and *stability*, we give a proof system for showing that transaction traces are *well-locked*, and we prove that well-locked transactions are serializable. We then consider heaps that are mutable trees (Section 3), where the main complication is that logical locks are now named by heap paths, which may be updated concurrently. Finally, we consider a class of mutable DAG heaps (Section 4) based on *decompositions* [13]; sharing complicates lock placements as there may be multiple access paths to an object.

For space reasons, all proofs are in the technical report [12]. Because our focus is on formalizing lock placements, we do not consider liveness properties, such as deadlock, or optimizations, such as early release, since these issues are orthogonal to the ones we explore. The standard techniques for ensuring deadlock-freedom apply, including both static techniques (imposing a total ordering on locks) and dynamic techniques (using a contention manager to resolve deadlocks at runtime).

$m \in \mathcal{M}$ memory locations $b ::= \text{F} \mid \text{T}$ booleans $\psi \subseteq \mathcal{M} \rightarrow 2^{\mathcal{L} \times \Phi}$ lock placements $t ::= \text{wr}(m, b) \mid \text{obs}(m) = b \mid \text{rd}(m) = b \mid \text{lock } l \mid \text{unlock } l$	$l \in \mathcal{L}, L \subseteq \mathcal{L}$ locks, lock sets $\omega ::= m \mapsto b$ heap assertions $\Phi \ni \phi ::= b \mid \omega \mid \phi \vee \phi \mid \phi \wedge \phi$ guards transaction ops.
---	---

Fig. 1. Locations, Lock Placements, Transaction Operations

2 Flat Maps

We first consider a simple class of heaps defined over a fixed set of *memory locations* \mathcal{M} . A *flat map heap* is a set of mappings $\{m \mapsto b\}_{m \in \mathcal{M}}$ from each location $m \in \mathcal{M}$ to a boolean value b . Let \mathcal{L} be a fixed set of *physical locks*; in this section we assume that memory locations and locks are disjoint. For ease of exposition we consider only exclusive locks — that is, if a transaction holds a lock then no other transaction may acquire concurrent access to the same lock.

A common correctness criterion for concurrent transactions is *serializability*. Informally a concurrent execution of a set of transactions is serializable if the reads and writes transactions make to the heap are equivalent to the reads and writes in some serial schedule of the same transactions. Serializability ensures we can reason about programs as if only one transaction executes at a time.

A *transaction* \mathbf{T} is a sequence $t^1 t^2 \dots$ of the atomic *transaction operations* given in Figure 1: a possibly unstable read of location m yielding b ($\text{rd}(m) = b$), a logical observation of location m yielding b ($\text{obs}(m) = b$), a write of b to location m ($\text{wr}(m, b)$), a lock of a physical lock l ($\text{lock } l$), or an unlock of physical lock l ($\text{unlock } l$). With the exception of the rd and obs operations the concrete semantics of transaction operations are standard; the details are in the technical report [12]. We assume the execution of operations is sequentially consistent.

The transaction language distinguishes between high-level obs operations, which are observations of the state of memory that affect the outcome of a transaction and for which the locking protocol must ensure serializability, and low-level rd operations, which do not directly affect the outcome of a transaction and need not be serializable. A transaction may freely perform a rd operation on any location at any time, regardless of the locks that it holds, however there is no guarantee that the value read will remain *stable*; a read is *stable* only if no concurrent transaction may write to the same location and invalidate the value that was read. If a transaction holds locks that ensure that the value returned by a rd operation is stable and cannot be altered by concurrent transactions, then a transaction may logically obs the result of the read operation and use that value to perform computation. The distinction between stable and unstable reads is key to reasoning about *speculative locking* (Section 2.1).

2.1 Lock Placements

We associate a *logical lock* with every heap location $m \in \mathcal{M}$. Whenever a transaction observes or changes the value of a memory location it must hold the associated logical lock. It is inefficient to attach a distinct lock to every memory

1: lock l	1: lock l_1	9: unlock l_1	1: lock l_1	9: obs(m_4) = F
2: rd(m_1) = T	2: rd(m_1) = T	10: unlock l_0	2: rd(m_1) = T	10: unlock l_4
3: obs(m_1) = T	3: obs(m_1) = T		3: obs(m_1) = T	11: unlock l_3
4: rd(m_3) = F	4: rd(m_3) = F		4: lock l_3	12: unlock l_1
5: obs(m_3) = F	5: obs(m_3) = F		5: rd(m_3) = F	
6: rd(m_4) = F	6: lock l_0		6: obs(m_3) = F	
7: obs(m_4) = F	7: rd(m_4) = F		7: lock l_4	
8: unlock l (a)	8: obs(m_4) = F (b)		8: rd(m_4) = F	(c)

Fig. 2. Transaction traces that observe the values of locations m_1 , m_3 , and m_4 under (a) coarse, (b) intermediate, and (c) fine-grained lock placements.

location; instead we use a smaller set of *physical locks* (or simply *locks*) \mathcal{L} to implement logical locks; a *lock placement* maps logical locks to physical locks. Different placement functions describe different granularities of locking.

Formally, a *lock placement* ψ for a boolean heap is a mapping from each location $m \in \mathcal{M}$ to a guarded set of locks that protect it. Each entry in $\psi(m)$ is a pair of a lock $l \in \mathcal{L}$ and a *guard* ϕ , which is a condition under which l protects m . A guard is a boolean combination of heap assertions $m \mapsto b$; for a given memory location each lock may only appear at most once on the left hand side of a guarded lock pair, and the set of guards must be mutually exclusive, and total, that is, exactly one guard is true for any given heap state.

For example, suppose $\mathcal{M} = \{m_0, \dots, m_{k-1}\}$. Different placements allow us to describe a range of different locking granularities:

- A coarse-grain locking strategy protects every memory location with the same lock, that is, set $L = \{l\}$ and set $\psi(m_i) = \{(l, \text{T})\}$ for all i . To observe or write to any memory location a transaction must hold lock l .
- An medium-grain locking strategy stripes memory locations across a small set of locks. Set $L = \{l_0, \dots, l_{p-1}\}$, and then set $\psi(m_i) = \{(l_{(i \bmod p)}, \text{T})\}$ for all i . To observe or write to memory location m_i , we must hold lock $l_{(i \bmod p)}$.
- A fine-grain strategy associates a distinct lock with every memory location. Set $L = \{l_0, \dots, l_{k-1}\}$ and set $\psi(m_i) = \{(l_i, \text{T})\}$ for all i . To observe or write to memory location m_i we must hold lock l_i .

Figure 2 shows three variants of a transaction that reads memory locations m_1 , m_3 and m_4 (chosen arbitrarily for the example), observing values T, F, and F respectively. The figure shows a variant of the transaction for each locking granularity, using $p = 2$ physical locks in the medium-grain case.

A *speculative lock placement* is a placement in which the identity of a lock that protects a memory location depends on the memory location itself. For example a simple speculative placement ψ_s is as follows. Let $L = \{l_f, l_t\}$ and $\mathcal{M} = \{m\}$. Set $\psi_s(m) = l_f$ if $m \mapsto \text{F}$, or l_t if $m \mapsto \text{T}$. Under this placement, lock l_f protects memory location m if m contains the value F, whereas lock l_t protects memory location m if m contains the value T.

A more realistic example of speculative lock placement is motivated by transactional predication [2] which uses a speculative placement of STM metadata.

(a) 1: $\text{rd}(m) = \text{T}$ 2: $\text{lock } l_t$ 3: $\text{rd}(m) = \text{T}$ 4: $\text{obs}(m) = \text{T}$ 5: $\text{unlock } l_t$	(b) 1: $\text{rd}(m) = \text{T}$ 7: $\text{obs}(m) = \text{F}$ 2: $\text{lock } l_t$ 8: $\text{unlock } l_f$ 3: $\text{rd}(m) = \text{F}$ 4: $\text{unlock } l_t$ 5: $\text{lock } l_f$ 6: $\text{rd}(m) = \text{F}$	(c) 1: $\text{lock } l_f$ 2: $\text{lock } l_t$ 3: $\text{rd}(m) = \text{T}$ 4: $\text{wr}(m, \text{F})$ 5: $\text{unlock } l_t$ 6: $\text{unlock } l_f$
--	--	--

Fig. 3. Traces that read and write location m under the speculative lock placement ψ_s .

We use a collection $\mathcal{M} = \{m_1, \dots, m_k\}$ of memory locations to model a concurrent set. Location m_i has value T if value i is present in the set. We use $L = \{l_\perp, l_1, \dots, l_k\}$ and the placement $\psi(m_i) = l_\perp$ if $m_i \mapsto \text{F}$, or l_i if $m_i \mapsto \text{T}$.

The speculative placement allows us to attach a distinct lock to every entry present in the set, without also requiring that we keep around a distinct lock for every entry that is absent from the set. Two transactions that operate on keys present in the set only contend on the same lock if they are accessing the same key. Transactions that operate on keys that are absent will however contend on l_\perp ; this strategy is effective if we expect sets to have at most a small fraction of all possible elements at any one time. If necessary, we can reduce contention on absent entries to arbitrarily low levels by striping the logical locks protecting absent entries across a set of physical locks $l_\perp^1, l_\perp^2, \dots$ as discussed earlier.

It may not be immediately obvious how to acquire a lock on a memory location when we do not know which lock to take without knowing the value of the memory location. The key to this apparent circularity is that a transaction can use unstable reads to guess the identity of the correct lock; once the transaction has acquired the lock it can redo the read to verify that its guess was correct. If the transaction guesses correctly, then the second read is stable. If the transaction guesses incorrectly it can release the lock and repeat the process. Figure 3(a) shows a transaction that observes the state of m under the speculative lock placement ψ_s . If another transaction had raced, we might have had to retry the read, as shown in Figure 3(b). Finally, to perform an update, we must hold both locks, as shown in Figure 3(c); otherwise by changing m we might implicitly release a lock that another transaction holds on the state of m .

2.2 Well-Locked Transactions

We represent the state of a transaction as two sets: the *observation set* Ω and a *lock set* L . The observation set Ω is a set of heap assertions $m \mapsto b$ that represent a transaction’s local view of the heap. The lock set L is a set of *physical locks* held by the transaction. Every heap assertion in the observation set must be *stable*; informally, the facts in the observation set are logically locked and cannot be invalidated by a concurrent interfering transaction. We write $\Omega[m \mapsto b]$ to denote the result of adding or updating the heap observation $m \mapsto b$ to Ω , replacing any existing observations about m . The predicate $\text{locked}_\psi(m, \Omega, L)$ holds for heap location m if a transaction with heap observations Ω and locks L has logically

$$\begin{array}{c}
\boxed{\text{FLOCK}} \quad \frac{l \notin L}{\Omega, L \vdash_{\psi} \text{lock } l; \Omega, L \cup \{l\}} \qquad \boxed{\text{FUNLOCK}} \quad \frac{l \in L \quad L' = L \setminus \{l\} \quad \Omega' = [\Omega \mid L'; \psi]}{\Omega, L \vdash_{\psi} \text{unlock } l; \Omega', L'} \\
\boxed{\text{FRDUNSTABLE}} \quad \frac{\Omega' = \Omega \cup \{m \mapsto b\} \quad \neg \text{locked}_{\psi}(m, \Omega', L)}{\Omega, L \vdash_{\psi} \text{rd}(m) = b; \Omega, L} \qquad \boxed{\text{FRDSTABLE}} \quad \frac{\Omega' = \Omega \cup \{m \mapsto b\} \quad \text{locked}_{\psi}(m, \Omega', L)}{\Omega, L \vdash_{\psi} \text{rd}(m) = b; \Omega', L} \\
\boxed{\text{FOBSERVE}} \quad \frac{(m \mapsto b) \in \Omega}{\Omega, L \vdash_{\psi} \text{obs}(m) = b; \Omega, L} \qquad \boxed{\text{FWRITE}} \quad \frac{m \in \text{dom } \Omega \quad \Omega' = \Omega[m \mapsto b] \quad (\forall m', l, \phi. (l, \phi) \in \psi(m') \wedge m \in \text{dom } \phi \implies l \in L)}{\Omega, L \vdash_{\psi} \text{wr}(m, b); \Omega', L}
\end{array}$$

Fig. 4. Well-locked transaction operations: $\Omega, L \vdash_{\psi} t; \Omega', L'$

locked location m under lock placement ψ , where $\Omega \vdash \phi$ denotes entailment:

$$\text{locked}_{\psi}(m, \Omega, L) = \exists(l, \phi) \in \psi(m). \quad l \in L \wedge \Omega \vdash \phi$$

The judgement $\Omega, L \vdash_{\psi} t; \Omega', L'$ defined in Figure 4 characterizes *well-locked operations*. The judgment holds if whenever a transaction with observations Ω and holding locks L executes operation t , then on completion of the operation the transaction has new observations Ω' and locks L' . Given the set of physical operations that a transaction performs, the well-lockedness judgement computes the set of stable observations of the transaction, and ensures that a transaction only logically observes and writes locations on which it holds logical locks.

The (FLOCK) rule allows a transaction to acquire a lock l if the transaction does not already have l in its set of locks L ; acquiring a lock has no affect on the observation set Ω . The (FUNLOCK) rule allows a transaction to release any lock l in its lock set L ; any facts in Ω that were protected by l are no longer stable, so the rule uses the *stabilization operator* to compute a new stable set of observations Ω' . The *stabilization* of a set of observations Ω_0 under locks L and placement ψ , written $[\Omega_0 \mid L; \psi]$, is the limit of the monotonic sequence:

$$\Omega_{i+1} = \{m \mapsto b \in \Omega_i \mid \text{locked}_{\psi}(m, \Omega_i, L)\}$$

Note that the limit always exists, because Ω_0 is finite (since it is constructed by a finite transaction execution) and the empty set is always a fixed point of the equation if no larger set is. A set of observations Ω is *stable* under locks L and placement ψ if Ω is its own stabilization, that is, $\Omega = [\Omega \mid L; \psi]$.

Rule (FOBSERVE) states that a transaction may logically observe any stable fact from its stable observation set Ω . The (FRDUNSTABLE) rule allows a transaction to perform a speculative read on a memory location on which the transaction does not hold a lock; however since the result may not be stable the rule does not update the set Ω . To enable reasoning about speculation, the determination whether the read is stable or not occurs in a context that includes the read of m ; since we assume that reads are atomic, there is an instant in time at which both the old stable facts in Ω and the newly read value of m hold, and it is in that context that we determine stability. The (FRDSTABLE) rule allows a transaction to read memory locations on which it holds a lock; since such a read is stable the rule updates the set of observations Ω with the newly read information about

the heap. Finally the (FWRITE) rule requires that a transaction can only update a location m if it holds the lock on m ; furthermore the lock for any location m' for which m appears in a guard must also be held by the transaction—hence no transaction can destabilize the observations of another transaction. The last condition together with $\text{locked}_\psi(m, \Omega, L)$ imply that $\text{locked}_\psi(m, \Omega', L)$ holds, which is why the latter is not listed as a precondition of the rule.

A transaction $\mathbf{T} = t^1 \dots t^k$ is *well-locked* if there exists a sequence of lock sets L^i and observation sets Ω^i such that

$$L^0 = L^k = \emptyset, \quad \Omega^0 = \Omega^k = \emptyset, \quad \text{and } \Omega^{i-1}, L^{i-1} \vdash_\psi t^i; \Omega^i, L^i \text{ for } 1 \leq i \leq k.$$

As an example of applying the rules, consider again the speculative read transaction shown in Figure 3(b). Let Ω^i and L^i denote the lock sets of the transaction after line i . Initially we have $\Omega^0 = \emptyset$ and $L^0 = \emptyset$. The read on line 1 is unstable, so $\Omega^1 = \emptyset$ and $L^1 = \emptyset$. The lock on line 2 adds an entry to the lock set l_t , so $\Omega^2 = \emptyset$ and $L^2 = \{l_t\}$. The read on line 3 yields $m \mapsto F$, however the read would only be stable if $\text{locked}_\psi(m, \{m \mapsto F\}, \{l_t\})$ holds, which it does not; once again we have $\Omega^3 = \emptyset$ and $L^3 = \{l_t\}$. Lines 4 and 5 update the lock set; we have $\Omega^4 = \Omega^5 = \emptyset$, $L^4 = \emptyset$, and $L^5 = \{l_f\}$. The read on line 6 once again yields $m \mapsto F$, but this time the predicate $\text{locked}_\psi(m, \{m \mapsto F\}, \{l_f\})$ holds and the read is stable, yielding $\Omega^6 = \{m \mapsto F\}$ and $L^6 = \{l_f\}$. The logical observation of $m \mapsto F$ on line 7 is permitted by the judgement since we know $m \mapsto F$ is part of the stable heap; the observation and lock sets are unchanged ($\Omega^7 = \Omega^6$, $L^7 = L^6$). Finally, line 8 releases lock l_f , so we have $L^8 = \emptyset$. The assertion about m in Ω^7 is no longer stable, so the stabilization operator removes it from the observation set, finally yielding $\Omega^8 = \emptyset$.

2.3 Serializability of Well-Locked Transactions

A *schedule* \mathbf{s} for a set of transactions $\mathbf{T}_1, \dots, \mathbf{T}_k$ is a permutation of the concatenation of all transactions in the set, such that each transaction \mathbf{T}_i is a subsequence of \mathbf{s} . Formally, a schedule is *valid* if it corresponds to an execution of the concrete semantics (see the technical report [12] for details). Informally, validity requires the execution respect the mutual exclusion property of locks, and memory accesses must accurately reflect the state of the global heap. A schedule is *serial* if operations of different transactions are not interleaved.

Lemma 1. *Let \mathbf{s} be a valid schedule of well-locked transactions $\{\mathbf{T}_1, \dots, \mathbf{T}_k\}$. Let Ω_i^j and L_i^j be the set of observations and locks of transaction i after schedule step j . Let h^j be the heap after schedule step j . Then for all time steps j the lock sets $\{L_i^j\}_{i=1}^k$ are disjoint, and the observation sets $\{\Omega_i^j\}_{i=1}^k$ are stable, have disjoint domains, and heap h^j is an extension of each $\{\Omega_i^j\}_{i=1}^k$.*

The disjointness of observation sets in Lemma 1 follows from the exclusivity of physical locks. If we allowed shared/exclusive locks, then we would also need to allow observation sets to overlap on values protected by shared locks.

A well-locked transaction $\mathbf{T} = (t^i)_{i=1}^k$ is *logically two-phase* if the domains of the observation sets of the transaction have a growing phase and a shrinking

phase, that is, there exists some j such that for all i where $1 \leq i \leq j$, we have $\text{dom } \Omega^{i-1} \subseteq \text{dom } \Omega^i$ and for all i where $j < i \leq k$ we have $\text{dom } \Omega^{i-1} \supseteq \text{dom } \Omega^i$.

A *logical schedule* $\hat{\mathbf{s}}$ is the subsequence of a schedule \mathbf{s} consisting of all the `obs` and `wr` operations. Two operations *conflict* if they access the same memory location m . Two schedules \mathbf{s}_1 and \mathbf{s}_2 are *conflict-equivalent* if the logical schedule $\hat{\mathbf{s}}_1$ can be turned into the logical schedule $\hat{\mathbf{s}}_2$ by a sequence of swaps of adjacent non-conflicting operations.

Lemma 2. *Any valid schedule of a set of well-locked, logically two-phase transactions $\{\mathbf{T}_1, \dots, \mathbf{T}_k\}$ is conflict-equivalent to a serial logical schedule.*

2.4 Shared/Exclusive Logical Locks

A limitation of the protocol just presented is that locks are exclusive — holding a lock gives a transaction sole access to an edge, even if the transaction only wants to read the edge. Lock placement is a separate issue from whether non-exclusive locks exist for reading. Exclusive locks are sufficient to illustrate all of the important features of our techniques and have the advantage of not introducing the extra and extraneous complications of supporting non-exclusive access. However, non-exclusive locks are important, and so we briefly illustrate how to extend our approach to locks providing shared read access.

To allow shared access to fields we relax the requirement that guards must be mutually exclusive, thereby allowing each logical lock to map to many physical locks at the same time. Under the relaxed definition of placement, a transaction has shared access to a memory location m if it holds at least one of the locks that protect m , whereas a transaction has exclusive access to m if it holds all of the locks that protect m . Formally, a transaction has shared access to a memory location m if $\text{locked}_\psi(m, \Omega, L)$ holds. We also define a new predicate

$$\text{exclusive}_\psi(m, \Omega, L) = \forall (l, \phi) \in \psi(m). \quad l \in L \vee \Omega \vdash \neg \phi$$

which holds for heap location m if a transaction with heap observations Ω and locks L has an exclusive logical lock on location m under lock placement ψ .

To show serializability, we need to add an additional precondition to the (FWRITE) rule requiring that a transaction have exclusive access to any memory location it writes. The statement of the proof of Lemma 1 must be altered since different observation sets may share fields on which they hold a shared lock—only the exclusively held fields must be disjoint between transactions. Finally we must update the definition of a two-phase transaction to ensure that transactions only release exclusive access to a field in the shrinking phase of a transaction.

3 Mutable Tree-Structured Heaps

In Section 2 we described a locking protocol for flat heaps with a fixed set of memory locations and locks. In this section we extend our results to dynamically allocated, mutable tree-shaped heaps with a placement function based on paths.

$f, \mathbf{f}, \mathcal{F}$ fields x, y, ρ object names $e ::= \text{nil} \mid x$ expressions
 $\omega ::= x.f \mapsto e$ heap assertions $\psi \subseteq 2^{\mathbf{f}} \rightarrow 2^{\mathbf{f}}$ placements
 $t ::= \text{wr}(x.f, e) \mid \text{obs}(x.f) = e \mid \text{rd}(x.f) = e \mid x = \text{new}() \mid \text{lock } x \mid \text{unlock } x$ trans. ops.

Fig. 5. Tree transactions

A *tree heap* h consists of a set of objects, each with a unique name, usually denoted x or y . Every object has a fixed set of fields \mathcal{F} . Each object field $x.f$ contains a pointer either to some object y or nil . The heap contains a distinguished *root object*, named ρ . In a quiescent state, that is, in the absence of running transactions, we require that the heap be a forest.

We associate a logical lock with every field of every object in the heap. Unlike the flat heaps of Section 2 we do not assume that we have a separate set of physical locks distinct from the set of memory locations; instead, following the practice of languages such as Java, we require that every heap object can function as a physical lock, and we use a lock placement function to describe a policy for mapping the logical locks attached to fields onto the physical locks (the objects). To define the lock placement, we use access paths from the root ρ to name both the fields we want to protect and the objects whose physical locks protect them.

We extend the set of transaction operations of Section 2 to read from and write to fields of objects, to handle dynamic allocation of new objects, and to apply lock and unlock operations to objects rather than a separate set of locks. The transaction operations, shown in Figure 5, are: write an expression e (either an object y or nil) to field f of object x ($\text{wr}(x.f, e)$), a possibly unstable read of field f of object x yielding result e ($\text{rd}(x.f) = e$), a stable observation of field f of object x yielding e ($\text{obs}(x.f) = e$), allocation of a fresh object ($x = \text{new}()$), locking an object ($\text{lock } x$), and unlocking an object ($\text{unlock } x$).

3.1 Lock Placements

We name edges in the heap as a non-empty field path (a sequence of field names) $\mathbf{f} = f_1 f_2 \dots$ from the root, ending in the edge in question. Since the path names a field in the heap, the path must be non-empty. We also name objects using fields, except that the path ends at the object the field points to; note that in the case of objects the empty path names the root of the heap. A *lock placement* ψ is a function from non-empty paths to paths, which maps every edge in a heap to an object whose attached physical lock protects it.

Consider heaps with fields drawn from the set $\mathcal{F} = \{a, b\}$. We can protect every edge of the heap with a single coarse-grain lock by setting $\psi_1(\mathbf{f}) = \epsilon$ for all \mathbf{f} . If we want different locks for the a and b subtrees, we can use the lock placement

$$\psi_2(\mathbf{f}) = a \text{ if } a \prec \mathbf{f}, \quad b \text{ if } b \prec \mathbf{f}, \quad \text{and } \epsilon \text{ if } \mathbf{f} = a \text{ or } \mathbf{f} = b \quad (1)$$

where $\mathbf{g} \prec \mathbf{f}$ denotes that \mathbf{g} is a prefix of \mathbf{f} . For example, in Figure 6(a), under placement ψ_2 the lock at ρ protects the edges ρx and ρy , the lock at x protects the edges xz and xu , and the lock at y protects the edge yv .

(a) lock ρ	(b) lock ρ	unlock w	(c) rd($\rho.b$) = y	unlock w
rd($\rho.b$) = y	rd($\rho.b$) = y	unlock y	lock y	unlock y
obs($\rho.b$) = y	obs($\rho.b$) = y	unlock ρ	rd($\rho.b$) = y	
lock y	lock y		obs($\rho.b$) = y	
$w = \text{new } ()$	rd($y.a$) = nil		rd($y.a$) = nil	
wr($y.a, w$)	$w = \text{new } ()$		$w = \text{new } ()$	
unlock y	lock w		lock w	
unlock ρ	wr($y.a, w$)		wr($y.a, w$)	

Fig. 7. Three transaction traces that add a new outgoing edge labelled a from node y to the tree of Figure 6(b) under the lock placements (a) ψ_2 , (b) ψ_3 , and (c) ψ_4 .




Fig. 6. Examples of tree heaps. Nodes represent objects, whereas edges represent fields. Node ρ is the root object.

we use the lock on the longest non-nil prefix of $\psi_2(ab)$, namely ρ itself.

Modifications to the heap may implicitly alter the mapping from logical locks to physical locks. If a transaction updates an edge, then the transaction must hold all logical locks whose mapping to physical locks may change both before and after the update. For example consider again the lock placement ψ_2 in the context of the tree heap shown in Figure 6(b). According to the placement the lock on ρ protects the edge a from the root; however since edge a points to nil, edges on any path that begins with a are also protected by the lock at a . If a transaction were to set $\rho.a$ to point to a fresh vertex w , the lock at w would now protect the edges on paths that begin with a ; the transaction has *split* the logical roles of the lock at ρ before the write between the lock at ρ and the lock on new node w . Whenever a transaction splits or merges locks (e.g., by setting the field $\rho.a$ to nil again), it must hold every lock involved.

Figure 7(a) shows a trace of a transaction that adds a new edge labeled a from object y to a fresh object w to the heap of Figure 6(b) under placement ψ_2 . The transaction acquires two locks, namely the lock at ρ that protects the edge from ρ to y , and the lock at y that protects the entire subtree rooted at y . We need not hold a lock on w when adding w into the tree since no path in the range of the placement function is a suffix of the path to the updated edge ba .

If we desire finer-grained locking, we can use a lock attached to every object to protect the fields of that object by using the placement function $\psi_3(\mathbf{g}, f) = \mathbf{g}$ for any \mathbf{g}, f . The lock placement ψ_3 places the lock that protects each edge f on the object at the head of the edge. Figure 7(b) shows a trace of a transaction that again adds the edge labeled a from node y to a fresh node w to the heap of Figure 6(b), this time under lock placement ψ_3 . Unlike the transaction of

If for an edge \mathbf{f} the placement path $\psi(\mathbf{f})$ leads to nil in the heap, we use the lock on the object preceding the edge to nil in the placement path. For example, consider Figure 6(b) under the placement ψ_2 . According to the placement, the lock that protects the edge named by path ab is $\psi_2(ab) = a$, however edge a from the root ρ points to nil. Instead,

$$\begin{array}{c}
\boxed{\text{TLOCK}} \quad \frac{x \notin L}{\Omega, \Gamma, L \vdash_{\psi} \text{lock } x; \Omega, \Gamma, L \cup \{x\}} \qquad \boxed{\text{TUNLOCK}} \quad \frac{x \in L \quad L' = L \setminus \{x\} \quad (\Omega', \Gamma') = [\Omega; \Gamma \mid L'; \psi] \quad \text{forest}(\Omega, \Omega', \Gamma, \Gamma')}{\Omega, \Gamma, L \vdash_{\psi} \text{unlock } x; \Omega', \Gamma', L'} \\
\boxed{\text{TNEW}} \quad \frac{\Omega' = \Omega \cup \{x.f \mapsto \text{nil} \mid f \in \mathcal{F}\} \quad x \notin \text{dom } \Omega \quad x \notin \Gamma \quad \Gamma' = \Gamma \cup \{x\}}{\Omega, \Gamma, L \vdash_{\psi} x = \text{new}(); \Omega', \Gamma', L} \qquad \boxed{\text{TOBSERVE}} \quad \frac{(x.f \mapsto e) \in \Omega}{\Omega, \Gamma, L \vdash_{\psi} \text{obs}(x.f) = e; \Omega, \Gamma, L} \\
\boxed{\text{TRDUNSTABLE}} \quad \frac{x.f \notin \text{dom } \Omega \quad \Omega' = \Omega \cup \{x.f \mapsto e\} \quad \neg \text{locked}_{\psi}(x.f, \Omega', \Gamma, L)}{\Omega, \Gamma, L \vdash_{\psi} \text{rd}(x.f) = e; \Omega, \Gamma, L} \qquad \boxed{\text{TRDSTABLE}} \quad \frac{x.f \notin \text{dom } \Omega \quad \Omega' = \Omega \cup \{x.f \mapsto e\} \quad \text{locked}_{\psi}(x.f, \Omega', \Gamma, L) \quad \Gamma' = \begin{cases} \Gamma & \text{if } e = \text{nil} \\ \Gamma \cup \{y\} & \text{if } e = y \end{cases}}{\Omega, \Gamma, L \vdash_{\psi} \text{rd}(x.f) = e; \Omega', \Gamma', L} \\
\boxed{\text{TWRITE}} \quad \frac{x.f \in \text{dom } \Omega \quad \Omega' = \Omega[x.f \mapsto e] \quad (\forall \mathbf{g}, \mathbf{h}. (\Omega \vdash \mathbf{g} \sim x) \wedge \mathbf{g}f \preceq \psi(\mathbf{h}) \implies \text{pathlocked}_{\psi}(\mathbf{h}, \Omega, L) \wedge \text{pathlocked}_{\psi}(\mathbf{h}, \Omega', L))}{\Omega, \Gamma, L \vdash_{\psi} \text{wr}(x.f, e); \Omega', \Gamma, L}
\end{array}$$

Fig. 8. Well-locked tree operations: $\Omega, \Gamma, L \vdash_{\psi} t; \Omega', \Gamma', L'$

Figure 7(a), we need to ensure that by adding the new edge the write does not implicitly change the mapping from edges to locks. The well-lockedness conditions, which we introduce shortly, require that a transaction hold all physical locks which may map to different logical locks before and after a write. The operation $\text{rd}(y.a) = \text{nil}$ verifies that there is no existing subtree of y reachable via edge a . Before the update the lock at y protects every possible edge reachable from $y.a$, however after the write the lock y only protects the edge $y.a$ itself, whereas the lock at w protects everything reachable from node w . Hence we must hold lock w when performing the write, since adding the edge splits the lock at y . (In general one must hold locks when connecting objects into the heap, however in this specific case, since the write which links w to the heap is the last write in the transaction it would be possible to optimize away the lock and unlock.)

Finally, if we set $\psi_4(\mathbf{f}) = \mathbf{f}$ we obtain a speculative placement where each edge is protected by a lock at its target. Figure 7(c) once again shows a transaction that adds a fresh edge labeled a to node w , this time using placement ψ_4 . The transaction begins with a speculative read to guess that the identity of the object whose lock protects $\rho.b$ is y . After locking y , the transaction performs the read again; since the read still returns y , the read is stable since the transaction already holds lock y . The transaction then performs a read of $y.a$ which returns nil . The value of the placement function for edge $y.a$ is $\psi(ba) = ba$, however since edge ba points to nil , the lock on the longest non- nil prefix of ba protects ba , in this case path b (node y). Since we hold the lock on y already, we know that the read of $y.a$ is also stable. Finally, the transaction must hold the lock on w when adding it to the heap to maintain the invariant that a transaction must hold all physical locks whose logical/physical mapping changes as a consequence of a write.

3.2 Well-Locked Transactions

We represent a transaction's state by three sets. As before, L is the set of locks that transaction holds, and Ω is a set of stable heap observations of the form

$x.f \mapsto e$. We do not require Ω be a forest; a transaction may create any heap shapes it desires within its local heap. However, the forest invariant must be restored when the transaction releases objects in its local heap back into the global heap. Enforcing this condition is the purpose of the set Γ . An object x is a member of Γ if the transaction has shown that there is no globally visible path from the root to x (i.e., the transaction has locked the edge to x). The well-lockedness rules for tree heaps ensure that there is at most one globally-visible edge to any node and hence the globally-visible part of the heap is a forest. At the start of every transaction Γ is the empty set. Transactions add entries to Γ by discovering global edges to nodes and transferring them into their local heap Ω ; entries are removed from Γ when pointers to objects are released from the stable heap Ω back into the global heap.

The *path alias* judgement $\Omega \vdash \mathbf{f} \sim x$ holds if \mathbf{f} is a path in Ω from the root to location x ; that is, if $|\mathbf{f}| = k$, then there is a sequence of vertices $\mathbf{v} = v_0 v_1 \dots$ such that $(\rho.f_0 \mapsto v_0) \in \Omega$, $(v_{i-1}.f_{i-1} \mapsto v_i) \in \Omega$ for all $1 < i < k - 1$, and $v_{k-1}.f_{k-1} \mapsto x$. We write $\mathbf{f} \in \Omega$ if the path \mathbf{f} from the root vertex exists in Ω , that is, $\Omega \vdash \mathbf{f} \sim x$ holds for some object x .

The *restriction of a path* \mathbf{f} to a local heap Ω , written $\mathbf{f}|_\Omega$, is defined as

$$\mathbf{f}|_\Omega = \mathbf{f} \text{ if } \mathbf{f} \in \Omega, \text{ or } \mathbf{g} \text{ if } \exists \mathbf{g}, h. \mathbf{g}h \preceq \mathbf{f} \wedge \Omega \vdash \text{nil} \sim \mathbf{g}h.$$

The restriction of path \mathbf{f} is either \mathbf{f} itself if present in the heap, or the longest prefix of the path present in the heap where no edge points to nil. The restriction of a path is undefined if the path \mathbf{f} leaves the stable local heap Ω .

We hold the lock on an edge reached via a path if we hold the corresponding lock placement, restricted to the heap:

$$\text{pathlocked}_\psi(\mathbf{f}, \Omega, L) ::= \exists x \in L. \Omega \vdash \psi(\mathbf{f})|_\Omega \sim x$$

We hold the lock on a field f of an object x under observations Ω , objects Γ and locks L , written $\text{locked}_\psi(x.f, \Omega, \Gamma, L)$, if we hold a lock on field f on every path in the local heap, and furthermore there are no paths to x outside the local heap. Formally,

$$\text{locked}_\psi(x.f, \Omega, \Gamma, L) ::= x \in \Gamma \wedge \forall \mathbf{g}. (\Omega \vdash x \sim \mathbf{g} \implies \text{pathlocked}_\psi(\mathbf{g}f, \Omega, L))$$

If the local heap Ω contains cycles, observe that there may be infinitely many paths \mathbf{g} and the predicate is well-defined in this case. To verify the absence of paths to x from outside the local heap, it is sufficient to check that $x \in \Gamma$, because any object $y \notin \Gamma$ is outside the local heap and thus has no stable fields and cannot form part of a path to x . Further, the definition of the *locked* predicate implies that if there is no path from the root ρ to node x , then the fields of x are locked for any transaction with $x \in \Gamma$; thus newly allocated objects can be added to Γ without taking a lock since they are disconnected from the global heap.

The judgement $\Omega, \Gamma, L \vdash_\psi t; \Omega', \Gamma', L'$ defined in Figure 8 captures the class of *well-locked tree operations*. If the judgement holds, then a transaction that executes operation t under stable observation set Ω , objects Γ , and lock set L yields a new stable observation set Ω' , objects Γ' and lock set L' . The (TNEW) rule states that all of the fields of a newly allocated object x point to nil, and since there can be no path to x in the heap all of x 's fields are stable and $x \in \Gamma'$.

As before, the (TLOCK) rule allows a transaction to acquire a lock it does not yet hold and has no affect on either Ω or Γ .

In the (TUNLOCK) rule, the stabilization operator is slightly more involved than in the case of flat heaps, because we must compute not just the stable set of heap facts, but also the set of objects for which the transaction has locked the incoming path: if an edge $x.f \mapsto y$ drops out of the stable observation set because a lock is released, the transaction can no longer assume it holds locks on all of the paths to object y . The stabilization (Ω', Γ') of a local heap Ω_0 and global heap Γ_0 under locks L and placement ψ , written $(\Omega', \Gamma') = \lceil \Omega_0; \Gamma_0 \mid L; \psi \rceil$, is the limit of the monotonically decreasing sequence:

$$\Omega_{i+1} = \{x.f \mapsto e \in \Omega_i \mid \text{locked}_\psi(x.f, \Omega_i, \Gamma_i, L)\} \quad \Gamma_{i+1} = \Gamma_i \setminus \{y \mid x.f \mapsto y \in \Omega_i \setminus \Omega_{i+1}\}$$

Rule (TUNLOCK) requires that transactions maintain the *forest condition*

$$\text{forest}(\Omega, \Omega', \Gamma, \Gamma') ::= \forall y. \left(|\{x.f \mid (x.f \mapsto y) \in \Omega \setminus \Omega'\}| = \begin{cases} 1 & \text{if } y \in \Gamma \setminus \Gamma' \\ 0 & \text{otherwise} \end{cases} \right).$$

The forest condition ensures that a transaction may only release a pointer to a node y into the global heap if there are no other references to y in the global heap ($y \in \Gamma$). Furthermore, the condition also ensures that a transaction cannot release two or more pointers to the same location y into the global heap.

The rules (TOBSERVE), (TRDUNSTABLE), and (TRDSTABLE) are similar to the rules in Section 2, updated to reflect that the heap now involves objects and fields. Note that (TRDSTABLE) adds the object that is the target of the read to Γ in the case that the field is not nil.

The most interesting rule is (TWRITE). Writing a field $x.f \mapsto y$ not only changes the paths to y , it changes the paths to every object reachable from y . Thus, as a result of a single field update, the placements may change for y and edges reachable from y . Furthermore, fields no longer reachable from $x.f$ after the update also may have altered lock placements. For this reason a transaction must hold locks on every edge reachable from $x.f$ both before and after the update. These conditions are necessary for safety, but need not be burdensome if the lock placement has a suitable granularity. For example, if the subtrees rooted at x and y are locked by the locks at x and y respectively, the update requires two locks.

A transaction $\mathbf{T} = t^1 \dots t^k$ is *well-locked* if there exists a sequence of lock sets L^i , observation sets Ω^i , and object sets Γ^i such that $L^0 = L^k = \emptyset$, $\Omega^0 = \emptyset$, $\Gamma^0 = \emptyset$, and $\Omega^{i-1}, \Gamma^{i-1}, L^{i-1} \vdash_\psi t^i; \Omega^i, \Gamma^i, L^i$ for $1 \leq i \leq k$. A well-locked transaction must begin with all three sets Ω , Γ , L empty. Furthermore at the end of the transaction the set of locks L must be empty again, and hence a transaction must release all of its locks. We do not require that Ω or Γ be empty at the conclusion of a transaction; however since the transaction may not hold any locks on termination, any part of the heap that is stable and in Ω with an empty lock set cannot be reachable from the global heap and is garbage.

Lemma 3. *Let \mathbf{s} be a valid schedule of a set of well-locked transactions $\{\mathbf{T}_1, \dots, \mathbf{T}_k\}$. Let Ω_i^j , Γ_i^j , and L_i^j be the set of observations, objects, and locks of each transaction after schedule step j . Let h^j be the heap after schedule step j , and suppose h^0 is a forest. Then for all time steps j :*

- the lock sets $\{L_i^j\}_{i=1}^k$ are disjoint, the sets $\{\Gamma_i^j\}_{i=1}^k$ are disjoint,
- the observation sets $\{\Omega_i^j\}_{i=1}^k$ are stable, have disjoint domains, and heap h^j is an extension of each $\{\Omega_i^j\}_{i=1}^k$, and
- the global heap h^j less edges present in the local heaps $\{\Omega_i^j\}_{i=1}^k$ is a forest. Further if $x \in \Gamma_i^j$ then all pointers to x are in some local heap $\Omega_{i'}^j$.

Finally, we have a logical serializability lemma analogous to Lemma 2, which can be extended to shared/exclusive locks as for flat heaps:

Lemma 4. *Any valid schedule of a set of well-locked, logically two-phase tree transactions $\{\mathbf{T}_1, \dots, \mathbf{T}_k\}$ is conflict-equivalent to a serial logical schedule.*

4 Transactions on DAGs of Bounded Degree

At the core of the locking protocol of Section 3 is the invariant that the global heap is a forest. Since lock placements are defined using access paths, for soundness the locking protocol must show that a transaction holds locks that protect an edge on every possible path. In a forest there is at most one path between nodes.

In this section we show how to relax the forest restriction and apply lock placements to a class of directed acyclic graph heaps with a bounded number of paths to each node. The technical machinery developed so far remains almost unchanged, with the exception that the forest condition is replaced by a condition that allows for more paths to an object. One hurdle, however, is that we need some way to describe the aliasing patterns in the heap, for otherwise it is not possible to define what it means to be sound for any locking protocol. We use a recent proposal for describing a large class of heaps with sharing [13], which describe *decomposition heaps* whose shape matches a static description given by a *decomposition heap shape*. We stress that our results are not limited to the class of heaps described in [13]; our techniques could be applied analogously to any number of other methods for describing the possible shapes of the heap. The point is that we need some description of the sharing patterns in the heap, and one good choice is to use decomposition heaps.

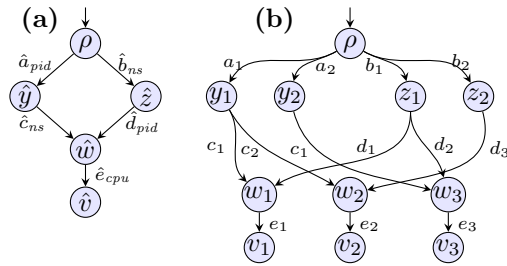


Fig. 9. (a): A decomposition heap shape, and **(b):** a decomposition heap that is an instance of decomposition heap shape (a).

A *decomposition heap shape* \hat{h} is a rooted, connected, directed acyclic graph (\hat{V}, \hat{E}) consisting of a set of vertices $\hat{V} = \{\hat{u}, \hat{v}, \dots\}$ and a set of edges $\hat{E} \subseteq \hat{V} \times \hat{F} \times \hat{V}$ labeled with field names drawn from a set \hat{F} . We require that every edge in a decomposition shape have a unique field label. Figure 9(a) gives a decomposition heap shape describing the data structures of a simple process scheduler. Every process has associated fields

$\psi \subseteq 2^{\hat{\mathbf{f}}} \rightarrow 2^{\hat{\mathbf{f}}}$	placements	u_i, v_j, \mathcal{V}	object names	\hat{u}, \hat{v}	vertices
$e ::= \text{nil} \mid v_i$	expressions	$\omega ::= u_i.f \mapsto e$	heap assertions	$\hat{f}, \hat{\mathbf{f}}, f_i, \mathbf{f}$	fields
$t ::= \text{wr}(u_i.f, e) \mid \text{obs}(u_i.f) = e \mid \text{rd}(u_i.f) = e \mid v_i = \text{new } \hat{v} \mid \text{lock } v_i \mid \text{unlock } v_i$	ops.				

Fig. 10. Decomposition transactions

pid (process id), ns (name space), and the process' assigned cpu ; a pair of a pid and a ns uniquely identify a process. To find the cpu of a particular process, we can first look up the the process id by following edge \hat{a}_{pid} and then the process' name space by following edge \hat{c}_{ns} , or we can first look up the name space by following edge \hat{b}_{ns} and then the process id by following edge \hat{d}_{pid} . For a given pair of process id and name space, the shared node \hat{v} in the decomposition shape assures us we will get the same result regardless of which path we take.

A decomposition shape is a static description of a class of heaps. Let $\text{in}(\hat{v})$ be the set of field names incoming to \hat{v} in a decomposition and let $\text{out}(\hat{v})$ be the set of outgoing field names. A heap (V, E) is an *instance* of a decomposition \hat{d} if

- every vertex in V is an instance v_i of some vertex $\hat{v} \in \hat{V}$,
- every edge $(u_i, f_j, v_k) \in E$ is an instance of some $(\hat{u}, \hat{f}, \hat{v}) \in \hat{E}$, and
- every vertex v_i has exactly one instance f_i of every incoming edge $\hat{f} \in \text{in}(\hat{v})$.

These conditions are a relaxation of usual definition of a valid instance [13], but suffice for our purposes. The last condition provides a bound on the in-degree of a vertex, which is the key to applying path-based lock placements to decomposition heaps. Figure 9(b) shows a heap that is an instance of the process scheduler decomposition shape of Figure 9(a). The nodes are objects in memory. Every edge \hat{f} from a vertex \hat{u} to a vertex \hat{v} of the decomposition shape has a corresponding set of edges $\{f_1, f_2, \dots\}$ outgoing from any instance u_i of \hat{u} in a decomposition heap. Intuitively, each vertex (object) u has a container data structure called f that contains references to a set of instances of \hat{v} . For example in Figure 9(b), the root object ρ has a set of process id's (the a_i) and a set of name spaces (the b_i). Note how the decomposition shape in Figure 9(a) is replicated across a number of different instances in Figure 9(b) with the stated sharing properties.

The well-lockedness rules defined below quantify over all suffixes of a path \mathbf{f} . To keep our transaction-language small, we require that the set of possible instances f_i of each abstract edge \hat{f} be drawn from a bounded set; that is $i \in \{1, \dots, k\}$ for some k . The bounded set restriction can be lifted by extending the transaction language with an iteration operation that allows a transaction to iterate over all instances of an edge from a vertex; the addition of iteration allows the rules to conclude the a fact holds for all instances of an abstract edge \hat{f} .

The transaction operations on DAGs are similar to those on trees (Section 3). The operations (Figure 10), are: write an expression e (either nil or some v_k to field f_j of object u_i ($\text{wr}(u_i.f_j, e)$)), a possibly unstable read of field f_j of object u_i yielding result e ($\text{rd}(u_i.f_j) = e$), a logical observation of field f_j of object u_i yielding e ($\text{obs}(u_i.f_j) = e$), allocation of a fresh object of type \hat{v} ($v_i = \text{new } \hat{v}$), locking an object ($\text{lock } v_i$), and unlocking an object ($\text{unlock } v_i$).

<p>(a) 1: lock ρ 9: wr($y_2.c_7, w_4$) 2: rd($\rho.a_2$) = y_2 10: wr($z_2.d_5, w_4$) 3: rd($\rho.b_2$) = z_2 11: unlock ρ 4: obs($\rho.a_2$) = y_2 5: obs($\rho.b_2$) = z_2 6: rd($y_2.c_7$) = nil 7: rd($z_2.d_5$) = nil 8: $w_4 = \text{new } \hat{w}$</p>	<p>(b) 1: rd($\rho.a_2$) = y_2 9: rd($y_2.c_7$) = nil 2: lock y_2 10: rd($z_2.c_5$) = nil 3: rd($\rho.a_2$) = y_2 11: $w_4 = \text{new } \hat{w}$ 4: obs($\rho.a_2$) = y_2 12: wr($y_2.c_7, w_4$) 5: rd($\rho.b_2$) = z_2 13: wr($z_2.d_5, w_4$) 6: lock z_2 14: unlock z_2 7: rd($\rho.b_2$) = z_2 15: unlock y_2 8: obs($\rho.b_2$) = z_2</p>
--	--

Fig. 11. Example transactions that add a new node w_4 with access paths a_2c_7 and b_2d_5 to the decomposition heap instance shown in Figure 9(b), under (a) lock placement $\psi_1(\mathbf{f}) = \epsilon$, and (b) lock placement $\psi_3(\mathbf{f}) = a_i$ if $a_i \preceq \mathbf{f}$, and b_j if $b_j \preceq \mathbf{f}$.

4.1 Lock Placements

Lock placements are defined exactly as in the tree case: ψ is a function from non-empty heap paths to paths, which maps every edge in a heap to an object whose lock protects it. Because edges may now have multiple paths that reach them, a transaction must hold locks on all paths to an edge to perform a stable read or to write the edge.

We now illustrate some of the possibilities for lock placements on decomposition heaps. For our standard first example, by setting $\psi_1(\mathbf{f}) = \epsilon$ for all \mathbf{f} we can use a single lock at the root of the heap to protect every edge in a decomposition instance. Figure 11(a) shows a well-locked transaction that adds a fresh instance of \hat{w} , namely w_4 , to the heap of Figure 9(b) under lock placement ψ_1 . Acquiring the lock on ρ protects the entire heap graph; the transaction then adds w_4 under both the access path a_2c_7 and b_2d_5 .

Another possibility is to use the placement $\psi_2(\mathbf{f}) = \epsilon$ if $\mathbf{f} \in \{a_i, b_i, a_i c_j, a_i d_j\}$, $a_i c_j$ if $\mathbf{f} = a_i c_j e_k$, and $b_i d_j$ if $\mathbf{f} = b_i d_j e_k$ which uses a lock at the root to protect instances of edges \hat{a} , \hat{b} , \hat{c} , and \hat{d} , and locks at instances of node \hat{w} to protect instances of edge \hat{e} . Instances of edge \hat{e} can be reached by two different paths, and thus to observe \hat{e} a transaction must acquire locks on both paths.

Finally, we can use a speculative lock placement. We could protect instances of edges \hat{a} and \hat{b} using speculative locks placed at their targets, and use locks at y and z to protect edges \hat{c} , \hat{d} , and \hat{e} , via the lock placement $\psi_3(\mathbf{f}) = a_i$ if $a_i \preceq \mathbf{f}$, and b_j if $b_j \preceq \mathbf{f}$. Figure 11(b) again shows a transaction that adds a fresh instance w_4 of node \hat{w} , this time under the speculative lock placement ψ_3 .

4.2 Well-Locked Transactions

As in the case of tree heaps we represent the state of a transaction using three sets: Ω (the local stable heap), L (the held set of locks), and Γ . Sets Ω and L are defined as for trees, but we extend the definition of Γ to DAGs with sharing.

The purpose of Γ is to track objects for which the transaction holds locks on incoming edges. In particular, if a transaction does not hold locks on some incoming edges to an object o , then there may be a path from the global heap to o and the transaction cannot rely on the stability of o 's fields. Thus Γ is

$$\begin{array}{c}
\boxed{\text{DNEW}} \quad \frac{\Omega' = \Omega \cup \{v_i.f_j \mapsto \text{nil} \mid \hat{f} \in \text{out}(\hat{v})\} \quad v_i \notin \text{dom } \Omega \quad \Gamma' = \Gamma[v_i \mapsto \text{in}(\hat{v})] \quad v_i \notin \text{dom } \Gamma}{\Omega, \Gamma, L \vdash_\psi v_i = \text{new } \hat{v}; \Omega', \Gamma', L} \quad \boxed{\text{DLOCK}} \quad \frac{v_i \notin L}{\Omega, \Gamma, L \vdash_\psi \text{lock } v_i; \Omega, \Gamma, L \cup \{v_i\}} \\
\boxed{\text{DUNLOCK}} \quad \frac{v_i \in L \quad L' = L \setminus \{v_i\} \quad (\Omega', \Gamma') = [\Omega; \Gamma \mid L'; \psi] \quad \text{balias}(\Omega, \Omega', \Gamma, \Gamma')}{\Omega, \Gamma, L \vdash_\psi \text{unlock } v_i; \Omega', \Gamma', L'} \quad \boxed{\text{DOBSERVE}} \quad \frac{(u_i.f_j \mapsto e) \in \Omega}{\Omega, \Gamma, L \vdash_\psi \text{obs}(u_i.f_j) = e; \Omega, \Gamma, L} \\
\boxed{\text{DRDUNSTABLE}} \quad \frac{u_i.f_j \notin \text{dom } \Omega \quad \Omega' = \Omega \cup \{u_i.f_j \mapsto e\} \quad \neg \text{locked}_\psi(u_i.f_j, \Omega', \Gamma, L)}{\Omega, \Gamma, L \vdash_\psi \text{rd}(u_i.f_j) = e; \Omega, \Gamma, L} \quad \boxed{\text{DRDSTABLE}} \quad \frac{u_i.f_j \notin \text{dom } \Omega \quad \Omega' = \Omega \cup \{u_i.f_j \mapsto e\} \quad \text{locked}_\psi(u_i.f_j, \Omega', \Gamma, L) \quad \Gamma' = \begin{cases} \Gamma & \text{if } e = \text{nil} \\ \Gamma[v_i \mapsto \Gamma(v_i) \cup \{\hat{f}\}] & \text{if } e = v_i \end{cases}}{\Omega, \Gamma, L \vdash_\psi \text{rd}(u_i.f_j) = e; \Omega', \Gamma', L} \\
\boxed{\text{DWRITE}} \quad \frac{u_i.f_j \in \text{dom } \Omega \quad \Omega' = \Omega[u_i.f_j \mapsto e] \quad (\forall \mathbf{g}, \mathbf{h}. (\Omega \vdash \mathbf{g} \sim u_i) \wedge \mathbf{g}f \preceq \psi(\mathbf{h}) \implies \text{pathlocked}_\psi(\mathbf{h}, \Omega, L) \wedge \text{pathlocked}_\psi(\mathbf{h}, \Omega', L))}{\Omega, \Gamma, L \vdash_\psi \text{wr}(u_i.f_j, e); \Omega', \Gamma, L}
\end{array}$$

Fig. 12. Well-locked decomposition operations: judgement $\Omega, \Gamma, L \vdash_\psi t; \Omega', \Gamma', L'$.

the transaction's view of the global heap and what other transactions might be able to do to objects of interest to the transaction. The global heap view Γ is a mapping from each vertex v_i in the heap to the subset of the incoming edge labels of the decomposition $\text{in}(\hat{v})$ known to be absent from the global heap (i.e., either non-existent or locked by the transaction). We maintain the invariant that in the global heap there is at most one edge to any instance of a decomposition vertex \hat{v} labeled with an instance of each $\hat{f} \in \text{in}(\hat{v})$. If $\Gamma(v_i) = \emptyset$, then v_i may have an instance of each incoming edge in $\text{in}(\hat{f})$ in the global heap. If $\Gamma(v_i) = \{\hat{f}\}$ then v has no incoming edge in the global heap labeled with an instance of \hat{f} . If $\Gamma(v) = \text{in}(\hat{v})$ then v has no incoming edges from the global heap.

As before, we hold the lock on an edge reached via a path if we hold the path's corresponding lock placement, restricted to the heap:

$$\text{pathlocked}_\psi(\mathbf{f}, \Omega, L) ::= \exists v_i \in L. \Omega \vdash \mathbf{g} \sim v_i \wedge \psi(\mathbf{f})|_\Omega = \mathbf{g},$$

where \mathbf{f}_Ω is the restriction of path \mathbf{f} to heap Ω , defined in Section 3.

The judgement $\Omega, \Gamma \vdash \text{exposed}(x)$ holds if there may be a path to vertex x in the heap that does not lie entirely in the stable observation set Ω ; the judgement is defined by the inference rules:

$$\frac{\Gamma(v_k) \neq \text{in}(\hat{v})}{\Omega, \Gamma \vdash \text{exposed}(v_k)} \quad \frac{\Omega, \Gamma \vdash \text{exposed}(u_i) \wedge (u_i.f_j \mapsto v_k) \in \Omega}{\Omega, \Gamma \vdash \text{exposed}(v_k)}$$

We hold the lock on a field $x.f$ if we hold a lock on that field on every path in the local heap, and there are no paths to x outside the local heap:

$$\text{locked}_\psi(v_i.f_j, \Omega, \Gamma, L) ::= \neg \text{exposed}(v_i) \wedge \forall \mathbf{g}. (\Omega \vdash \mathbf{g} \sim v_i \implies \text{pathlocked}_\psi(\mathbf{g}f_j, \Omega, L))$$

The judgement $\Omega, \Gamma, L \vdash_\psi t; \Omega', \Gamma', L'$ defined by the rules in Figure 12 describes the class of *well-locked decomposition operations*, analogous to the class of well-locked tree operations of Section 3. The judgement holds if a transaction executing operation t under local heap Ω , global heap approximation Γ , and

locks L yields an updated local heap Ω' , global heap approximation Γ' , and lock set L' . The (DNEW) rule states that the fields of a newly allocated object v_i point to nil; furthermore there can be no heap paths to a freshly allocated object so assertions about the fields of v_i are stable and $\Gamma(v_i) = \text{in}(\hat{v})$. The (DLOCK) rule allows a transaction to acquire a lock that it does not hold at any time.

The (DUNLOCK) rule allows a transaction to release any lock that it holds; the rule applies the stabilization operation to remove any newly unstable facts from Ω . Similar to the tree case, the *stabilization* (Ω', Γ') of a local heap Ω_0 and global heap Γ_0 under locks L and placement ψ , written $(\Omega', \Gamma') = [\Omega_0; \Gamma_0 \mid L; \psi]$, is the limit of the monotonically decreasing sequence:

$$\begin{aligned}\Omega_{i+1} &= \{u_j.f_k \mapsto e \in \Omega_i \mid \text{locked}_\psi(u_j.f_k, \Omega_i, \Gamma_i, L)\} \\ \Gamma_{i+1} &= \Gamma_i \setminus \{v_k \mapsto \hat{f} \mid u_i.f_j \mapsto v_k \in \Omega_i \setminus \Omega_{i+1}\}\end{aligned}$$

To ensure there is at most instance of any edge $\hat{f} \in \text{in}(\hat{v})$ in the global heap, the rule requires the *bounded alias* condition $\text{balias}(\Omega, \Omega', \Gamma, \Gamma')$, defined as

$$\forall v_k. |\{u_i.f_j \mid (u_i.f_j \mapsto v_k) \in \Omega \setminus \Omega'\}| = 1, \text{ if } \hat{f} \in \Gamma(v_k) \setminus \Gamma'(v_k) \text{ or } 0, \text{ otherwise.}$$

The bounded alias condition ensures that a transaction may only release an edge with abstract label \hat{f} to a node v_k into the global heap if there are no other edges to v_k labeled \hat{f} in the global heap ($\hat{f} \in \Gamma(v_k)$). The condition also forbids releasing two pointers with the same label \hat{f} to the same node v_k into the heap.

Rule (DOBSERVE) states that a transaction may logically observe stable heap facts. The (DRDUNSTABLE) rule allows a transaction to read a value speculatively at any time, however unstable reads do not update Ω or Γ . A transaction may perform a stable read of a pointer if it holds the appropriate lock, transferring the pointer from the global heap into Ω and updating Γ accordingly. Finally, a transaction may write to a field if it holds the associated lock and holds locks on any edges whose logical/physical mapping may change due to the update.

A transaction $\mathbf{T} = t^1 \dots t^k$ is *well-locked* if there exists a sequence of lock sets L^i , observation sets Ω^i , and global heap sets Γ^i such that $L^0 = L^k = \emptyset$, $\Omega^0 = \emptyset$, $\Gamma^0 = \emptyset$, and $\Omega^{i-1}, \Gamma^{i-1}, L^{i-1} \vdash_\psi t^i; \Omega^i, \Gamma^i, L^i$ for $1 \leq i \leq k$.

Lemma 5. *Let \mathbf{s} be a valid schedule of well-locked transactions $\{\mathbf{T}_1, \dots, \mathbf{T}_k\}$. Let Ω_i^j, Γ_i^j , and L_i^j be the set of observations, global heaps, and locks of each transaction after schedule step j . Let h^j be the heap after schedule step j , and suppose the part of h^0 reachable from the root is a tree. Then for all time steps j :*

- *the lock sets $\{L_i^j\}_{i=1}^k$ are disjoint, and the non-alias sets $\{\Gamma_i^j\}_{i=1}^k$ are disjoint,*
- *the observation sets $\{\Omega_i^j\}_{i=1}^k$ are stable, disjoint, and heap h^j is an extension of each $\{\Omega_i^j\}_{i=1}^k$, and*
- *Let heap h be the heap h^j less edges present in the local heaps $\{\Omega_i^j\}_{i=1}^k$. Then for every vertex $v \in h$ and edge label $\hat{f} \in \text{in}(\hat{v})$ either there is exactly one edge labeled with an instance of \hat{f} pointing to v in h , or $\hat{f} \in \Gamma_i^j$ for some i and there are no edges labeled with an instance of \hat{f} pointing to v in h .*

Finally, we have a logical serializability lemma similar to Lemma 2 and Lemma 4, which can be extended to shared/exclusive locks using the approach in Section 2.4.

Lemma 6. *Any valid schedule of well-locked, logically two-phase decomposition transactions $\{\mathbf{T}_1, \dots, \mathbf{T}_k\}$ is conflict-equivalent to a serial logical schedule.*

5 Related Work

Two-phase locking was originally introduced in the context of transactions operating over abstract entities, each with its own associated lock [8]. The core technical idea of this paper is that we can use two-phase locking to show serializability of a wide class of locking strategies by adding a layer of indirection between logical locks, which are the entities that are the subject of the original two-phase locking protocol, and the physical locks that implement them. Other authors have also advocated more logical notions of locking [5].

Various authors have investigated techniques for inferring locks to implement atomic sections [16,14,7,11,3,4,20]. A related problem is automatically optimizing programs with explicit locking by combining multiple locks into one [6]. A key part of this class of work is constructing a mapping from program objects to the locks that protect them, similar to our lock placement language. The lock placements we propose are much more flexible; in particular existing formalisms cannot handle the class of path placements we propose in this paper, such as speculative locks, or lock placements that vary with heap updates. A possible future application of our methods is extending lock inference techniques to take advantage of the additional expressive power of our techniques.

A novel feature of our proposal is that we can reason about speculative lock placements. Speculative locking is used in practice in highly concurrent libraries and has appeared in the literature in the context of software transactional memory [2]. Although we present our ideas in the context of pessimistic locks, the same idea can also be used to reason about speculative placements of STM metadata.

A variety of locking protocols have been proposed in the literature that extend two-phase locking to handle dynamically changing heaps and to allow early release. Examples include the dynamic tree and DAG locking protocols [1] and domination locking [9]. Existing protocols use the lock on each object to protect that object's fields, whereas our work investigates a more flexible space of mappings. We do not address early release as it is orthogonal to the issues of lock placement.

The concept of a stable set and stabilization is related to rely-guarantee logic [15] and its developments [19]. Concurrent extensions of separation logic, such as Concurrent Separation Logic [17], RGSep [18] and work on storable locks [10] allow local reasoning about programs with shared mutable state that is accessed concurrently. Our work complements work on direct reasoning about concurrent code; we propose a locking protocol, parameterized by a lock placement, by which we can show conflict-serializability for code that obeys the protocol.

6 Conclusion

We have formalized lock placements, showing that such diverse concepts as lock granularity, speculative locks, lock splitting and merging, and dynamically

changing lock assignments can all be expressed using a lock placement that maps each heap field to a lock that guards it. We described a series of proof systems for showing that transaction traces are well-locked and therefore serializable, applied to flat heaps, tree-structured heaps, and to DAG heaps with bounded degree.

References

1. Attiya, H., Ramalingam, G., Rinetzky, N.: Sequential verification of serializability. In: POPL. pp. 31–42. ACM, New York, NY, USA (2010)
2. Bronson, N.G., Casper, J., Chafi, H., Olukotun, K.: Transactional predication: high-performance concurrent sets and maps for STM. In: PODC. pp. 6–15. ACM, New York, NY, USA (2010)
3. Cherem, S., Chilimbi, T., Gulwani, S.: Inferring locks for atomic sections. In: PLDI. pp. 304–315. ACM, New York, NY, USA (2008)
4. Cunningham, D., Gudka, K., Eisenbach, S.: Keep off the grass: Locking the right path for atomicity. In: Hendren, L. (ed.) Compiler Construction, LNCS, vol. 4959, pp. 276–290. Springer Berlin / Heidelberg (2008)
5. Deshmukh, J., Ramalingam, G., Ranganath, V.P., Vaswani, K.: Logical concurrency control from sequential proofs. In: Programming Languages and Systems, LNCS, vol. 6012, pp. 226–245. Springer Berlin / Heidelberg (2010)
6. Diniz, P.C., Rinard, M.C.: Lock coarsening: Eliminating lock overhead in automatically parallelized object-based programs. JPDC 49(2), 218–244 (1998)
7. Emmi, M., Fischer, J.S., Jhala, R., Majumdar, R.: Lock allocation. In: POPL. pp. 291–296. ACM, New York, NY, USA (2007)
8. Eswaran, K.P., Gray, J.N., Lorie, R.A., Traiger, I.L.: The notions of consistency and predicate locks in a database system. Commun. ACM 19, 624–633 (Nov 1976)
9. Golan-Gueta, G., Bronson, N., Aiken, A., Ramalingam, G., Sagiv, M., Yahav, E.: Automatic fine-grained locking using shape properties. In: OOPSLA (2011)
10. Gotsman, A., Berdine, J., Cook, B., Rinetzky, N., Sagiv, M.: Local reasoning for storable locks and threads. In: APLAS. pp. 19–37 (2007)
11. Halpert, R.L., Pickett, C.J.F., Verbrugge, C.: Component-based lock allocation. In: PACT. pp. 353–364. IEEE Computer Society, Washington, DC, USA (2007)
12. Hawkins, P., Aiken, A., Fisher, K., Rinard, M., Sagiv, M.: Reasoning about lock placements (extended version). Tech. rep., Stanford University (2011), <http://theory.stanford.edu/~hawkinsp/papers/tr/lockplacements.pdf>
13. Hawkins, P., Aiken, A., Fisher, K., Rinard, M., Sagiv, M.: Data representation synthesis. In: PLDI. pp. 38–49. ACM, New York, NY, USA (2011)
14. Hicks, M., Foster, J.S., Pratikakis, P.: Lock inference for atomic sections. In: TRANSACT (2006)
15. Jones, C.: Development methods for computer programs including a notion of interference. Ph.D. thesis, Oxford University (1981)
16. McCloskey, B., Zhou, F., Gay, D., Brewer, E.: Autolocker: Synchronization inference for atomic sections. In: POPL. pp. 346–358. ACM, New York, NY, USA (2006)
17. O’Hearn, P.W.: Resources, concurrency, and local reasoning. Theoretical Computer Science 375(1–3), 271–307 (2007)
18. Vafeiadis, V., Parkinson, M.J.: A marriage of rely/guarantee and separation logic. In: CONCUR (2007)
19. Wickerson, J., Dodds, M., Parkinson, M.: Explicit stabilisation for modular rely-guarantee reasoning. LNCS, vol. 6012, pp. 610–629 (2010)
20. Zhang, Y., Sreedhar, V., Zhu, W., Sarkar, V., Gao, G.: Minimum lock assignment: A method for exploiting concurrency among critical sections. In: Languages and Compilers for Parallel Computing, LNCS, vol. 5335, pp. 141–155 (2008)