Inherent limitations on disjoint-access parallel implementations of transactional memory

Hagit Attiya, Eshcar Hillel, Alessia Milani

► To cite this version:

Hagit Attiya, Eshcar Hillel, Alessia Milani. Inherent limitations on disjoint-access parallel implementations of transactional memory. SPAA, 2009, Unknown, pp.69-78, 10.1145/1583991.1584015 . hal-00992693

HAL Id: hal-00992693

<https://hal.inria.fr/hal-00992693>

Submitted on 19 May 2014

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Inherent Limitations on Disjoint-Access Parallel Implementations of Transactional Memory*

(Preliminary Version)

Hagit Attiya
Department of Computer
Science, Technion
Haifa, Israel
hagit@cs.technion.ac.il

Eshcar Hillel[†]
Department of Computer
Science, Technion
Haifa, Israel
eshcar@cs.technion.ac.il

Alessia Milani[‡]
Department of Computer
Science, Technion
Haifa, Israel
alessia@cs.technion.ac.il

ABSTRACT

Transactional memory (TM) is a promising approach for designing concurrent data structures, and it is essential to develop better understanding of the formal properties that can be achieved by TM implementations. Two fundamental properties of TM implementations are *disjoint-access parallelism*, which is critical for their scalability, and the *invisibility* of read operations, which reduces memory contention.

This paper proves an inherent tradeoff for implementations of transactional memories: they cannot be both disjoint-access parallel and have read-only transactions that are invisible and always terminate successfully. In fact, a lower bound of $\Omega(t)$ is proved on the number of writes needed in order to implement a read-only transaction of t items, which successfully terminates in a disjoint-access parallel TM implementation. The results assume *strict serializability* and thus hold under the assumption of *opacity*. It is shown how to extend the results to hold also for weaker consistency conditions, *serializability* and *snapshot isolation*.

Categories and Subject Descriptors

D.1.3 [Programming Techniques]: Concurrent Programming; F.2.2 [Analysis of Algorithms and Problems]: Nonnumerical algorithms and problems

General Terms

Theory, Algorithms, Design

*This research is partially supported by the *Israel Science Foundation* (grant number 953/06) and Intel Corporation.

[†]Supported in part by a scholarship from the Israel Ministry of Science.

[‡]On leave from Sapienza, Università di Roma, Dipartimento di Informatica e Sistemistica, “Antonio Ruberti”; supported in part by a fellowship from the Lady Davis Foundation and by a grant MUR, FIRB Italia-Israele RBIN047MH9

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

SPAA’09, August 11–13, 2009, Calgary, Alberta, Canada.
Copyright 2009 ACM 978-1-60558-606-9/09/08 ...\$5.00.

Keywords

Transactional memory, disjoint-access parallelism, partial snapshots, lower bound, impossibility result

1. INTRODUCTION

Transactional memory is an attractive paradigm for programming concurrent applications for multicores. A transaction encapsulates a sequence of operations, and it is guaranteed that if any operation takes place, they all do, and that if they do, they appear to other threads to do so atomically, as one indivisible operation. A transactional memory *implementation* translates high-level transaction operations to low-level primitive operations on base objects, containing the data and the meta-data needed for the implementation.

Transactional memory is seriously considered as part of software solutions and as a basis for novel hardware designs. It is therefore imperative to understand inherent tradeoffs in the design and implementation of transactional memory.

One property that is considered critical for the scalability of a transactional memory implementation is *disjoint-access parallelism*: operations on disconnected data should not interfere. Conceptualizing this notion is best done through the *conflict graph* of transactions that overlap in time. Informally, the vertices of the conflict graph correspond to data items, and there is an edge between data items if they are accessed by the same transaction. Consider, for example, four concurrent transactions: T_1 with the data set $\{v_1, v_2\}$, T_2 , T_3 , and T_4 with the data sets $\{v_2\}$, $\{v_3\}$, and $\{v_4\}$, respectively. Figure 1 depicts the conflict graph of the execution interval of these four transactions. This conflict graph contains only one edge connecting the two vertices representing the data items in the data set of T_1 (see the formal definition in Section 2).

Several transactional memories, e.g. [4, 13], guarantee that transactions access the same base object only if their data items are connected in the conflict graph. In particular, there is no concurrent access to the shared memory by transactions that access disjoint parts of the data. In such implementations, the transactions T_1 , T_3 , and T_4 in the example above, do not access the same base object, since no path connects their data items in the conflict graph.

Another important goal is to optimize *read-only transactions*, i.e., transactions that access the memory only through read operations. It is desirable that in their implementations, read-only transactions do not execute primitive write

operations to the memory, so as to reduce memory contention; implementations of read-only transactions that do not write to the memory are called *invisible*. Moreover, since read-only transactions do not write to data items, it seems plausible they should eventually be able to obtain a consistent view of the data, provided previous versions are kept (as is done in multi-version implementations [20, 22, 23]). Thus, read-only transactions should (eventually) terminate successfully, regardless of concurrent transactions; such transactions are called *wait-free*.

None of the existing transactional memory implementations is both disjoint-access parallel and has invisible, wait-free read-only transactions. Some are disjoint-access parallel and have invisible but not wait-free read-only transactions [4, 13], while others have invisible, wait-free read-only transactions but are not disjoint-access parallel [22].

Consider, for example, the four transactions above, and assume T_1 is a read-only transaction, while T_2 , T_3 , and T_4 all write to their data items. In the algorithm given in [13], which is disjoint-access parallel and has invisible read-only transactions, T_1 reads v_1 and v_2 , then T_2 writes to v_2 , finally T_1 validates its read set at commit time. The value of v_2 has changed since T_1 read it, and T_1 aborts. In the algorithm given in [22], which has invisible, wait-free read-only transactions, T_1 , T_3 , and T_4 all access a common counter (T_1 reads it, while T_3 and T_4 write to it), thus violating disjoint-access parallelism.

This paper shows that there is an inherent tradeoff—no transactional memory implementation can be disjoint-access parallel and have invisible, wait-free read-only transactions—and one of these desirable properties must always be compromised. In fact, we prove a stronger result, showing that in a disjoint-access parallel transactional memory implementation with wait-free read-only transactions, a transaction reading t data items must apply non-trivial primitives (e.g., writes) to at least $t - 1$ base objects. Thus, a read-only transaction must perform one low-level write essentially for each item in its *read set*.

The wait-freedom requirement might seem too restrictive for practical purposes; however, we can prove a similar result where a read-only transaction repeatedly aborts and never terminates successfully; see further discussion in Section 6. For read-dominated applications, this implies too much wasted work.

The consistency condition commonly used for transactional memory is *opacity* [10]; very roughly stated, opacity requires all transactions to appear to execute sequentially in an order that agrees with the order of non-overlapping transactions. This is similar to requiring *strict view serializability* [21] applied to all transactions (included aborted

ones), extended to allow operations other than reads and writes. Our results only assume *strict serializability* [21], and hence hold also under the assumption of opacity. In fact, the results also hold for weaker consistency conditions, *serializability* and *snapshot isolation*.

The rest of the paper is organized as follows: Section 2 introduces basic definitions and in particular, the notion of disjoint-access parallelism. Section 3.1 presents an impossibility result showing that in a disjoint-access parallel STM implementation with invisible read-only transactions, some read-only transaction may never terminate successfully; this result is proved using only three processes. Section 3.2 strengthens this result and shows that a read-only transaction on t items (in a disjoint-access parallel STM implementation with wait-free read-only transactions) must apply non-trivial primitives to $t - 1$ base objects; this result requires $t + 1$ processes. Section 4 extends the results to hold even with the weaker conditions of serializability and snapshot isolation. We discuss related work in Section 5, and conclude in Section 6.

2. PRELIMINARIES

A *transaction* is a sequence of operations executed by a single process on a set of *data items* shared with other transactions; all data items are initially 0. We assume data items are accessed by simple *read* and *write* operations; our impossibility results clearly hold for transactional memory that also supports other operations. A complete interface of transactional memory also includes *commit* and *abort* operations, which we do not model here, since they are not needed for our impossibility results.

The collection of data items accessed by a transaction is the transaction’s *data set*; in particular, the items written by the transaction are its *write set*, and the items read by the transaction are its *read set*. A transaction whose *write set* is empty, is said to be a read-only transaction. We assume the transaction’s read set and write set are provided at the start of the transaction, and do not elaborate further on the manner a transaction issues its operations; this only makes our impossibility results stronger.

An *implementation of software transactional memory* (abbreviated *STM*) provides data representation for transactions and data items using *base objects*, and algorithms, specified as *primitive operations* on the base objects, which asynchronous processes have to follow in order to execute the operations of transactions. In addition to ordinary read and write primitives, we allow arbitrary read-modify-write primitives, even those accessing several locations simultaneously. In particular, the implementation may use a *cas*(o , *exp*, *new*) that writes the value *new* to location o if its value is equal to *exp*, and returns a success or failure indication.

A primitive is *non-trivial* if it may change the value of the object, e.g., a write or *cas*; otherwise, it is *trivial*, e.g., a read.

An *event* is a computation *step* by a process consisting of local computation and the application of a primitive to base objects, followed by a change to the process’s state, according to the results of the primitive. A *configuration* is a complete description of the system at some point in time, i.e., the state of each process and the state of each shared base object. There is a unique *initial* configuration in which every process is in its initial state and every base object contains its initial value.

Figure 1: Example of a simple conflict graph: v_1 and v_2 are the data items of T_1 ; v_2 , v_3 , and v_4 are the data items of T_2 , T_3 , and T_4 , respectively.

An *execution interval* α is a finite or infinite alternating sequence $C_0, \phi_0, C_1, \phi_1, C_2, \dots$, where C_k is a configuration, ϕ_k is an event and the application of ϕ_k to C_k results in C_{k+1} , for every $k = 0, 1, \dots$. An *execution* is an execution interval in which C_0 is the initial configuration.

Two executions α_1 and α_2 are *indistinguishable* to a process p , if p goes through the same sequence of state changes in α_1 and in α_2 ; in particular, this implies that it goes through the same sequence of events.

We point out that the model encompasses two levels of abstraction: The high level has transactions, each of which is a sequence of operations accessing data items. At the low level, these transactions are translated into executions in which a sequence of events apply primitive operations (or primitives) to base objects, containing the data and the meta-data needed for the implementation.

STM Properties. The *interval* of a transaction T is the execution interval that starts at the first event of T and ends at the last event of T , if there is one, taken by the process executing the algorithm for T . If T does not have a last event in the execution, then the interval of T is the (possibly infinite) execution interval starting at the first event of T . Two transactions *overlap* if their intervals overlap. A configuration C is *quiescent* if no transaction is pending in C , i.e., it is not inside the interval of any transaction.

An STM is *serializable* if transactions appear to execute sequentially, one after the other [21]; we assume that the serialization preserves the *per-process order*, i.e., transactions of the same process maintain their order. An STM is *strictly serializable* if this order preserves the order of non-overlapping transactions [21]; this notion is called *order-preserving serializability* in [25], and is the analogue of *linearizability* [15] for transactions. Note that strict serializability is implied by the *opacity* correctness condition, recently defined for transactional memory [10].

We assume that a transaction terminates successfully if it runs alone from a quiescent configuration. This property is satisfied by *obstruction-free* STM implementations, in which a process that eventually runs alone for long enough makes progress, i.e., transactions terminate successfully when eventually executing solo [13]. This property is also satisfied by STM implementations that are *weakly progressive* [11], in which a transaction that does not encounter conflicts has to terminate successfully; note that blocking STM implementations like TL2 [6] are weakly progressive.

Memory disjoint-access parallelism. An important property STM implementations have to provide is allowing unrelated transactions to progress independently, even if they are concurrent. Below, we formally define what it means for two transactions to be *unrelated* through a conflict graph that represents the relations between transactions. Then we define *disjoint-access parallelism*, a property that captures the intuition that an implementation should not cause two transactions, which are unrelated at the high-level, to simultaneously access the same low-level shared memory.

The *conflict graph* of an execution interval I is an undirected graph in which vertices represent data items and edges connect data items if they are accessed by the same transaction. If I overlaps the execution interval of a transaction T , and the data items v_1 and v_2 are in the data set of T , the graph includes an edge between the vertices representing

v_1 and v_2 .

Two transactions T_1 and T_2 are *disjoint-access* if there is no path between an item in the data set of T_1 and an item in the data set of T_2 , in the conflict graph of the minimal execution interval containing the intervals of T_1 and T_2 .

Two events *contend* on a base object o if they both access o , and at least one of them applies a non-trivial primitive to o . Two processes *concurrently contend* on a base object o if they have pending events at the same configuration that contend on o .

DEFINITION 1. An STM implementation is weakly disjoint-access parallel if two processes p_1 and p_2 , executing transactions T_1 and T_2 , concurrently contend on the same base object, only if T_1 and T_2 are not disjoint-access.

This definition captures the first condition of the disjoint-access parallelism property of Israeli and Rappoport [17], in accordance with most of the literature (cf. [14]). Our definition is weaker than their definition, as it allows two processes to apply a trivial primitive on the same base object, e.g., read, when executing two transactions even if they are disjoint-access. Moreover, our definition only prohibits concurrent contending accesses, allowing transactions to contend on a base object o at different points of the execution; we shall see in Lemma 2 that, under some conditions, these transactions can be made to concurrently contend on o .

The original definition [17] also restricts the impact of concurrent transactions on the *step complexity* of a transaction; our results do not rely on this additional condition, making them stronger.

3. STRICTLY SERIALIZABLE STMS

3.1 Impossibility of Invisible Read-Only Transactions

A read-only transaction is *invisible* if its algorithm only applies trivial primitives to base objects. We prove that in a disjoint-access parallel STM implementation with invisible read-only transactions, some read-only transaction will not terminate successfully in a finite number of steps; this is formally stated in Theorem 4.

Specifically, we construct an infinite execution of a read-only transaction. This execution consists of a single read-only transaction with one complete update transaction between any pair of consecutive steps by the read-only transaction; an *update* is a transaction with a singleton write set and an empty read set. We first define a special (finite) execution of this form, called *flippable*, and show that such a read-only transaction cannot terminate successfully. Then we show how a flippable execution can be repeatedly extended to construct successively longer flippable executions.

An execution is called *flippable* since there are two similar executions in which we flip the position of two update transactions and one of the executions is indistinguishable from the original execution. One type of flipped execution is called a *forward flip* since it moves an update transaction forward in the execution, while other is called a *backward flip* since it defers the execution of an update transaction. Formally:

DEFINITION 2. A flippable execution of length k with t updaters is a finite execution $E_k = U_0 s_1 U_1 \dots s_k U_k$ executed

$$\begin{array}{rcl}
q & : & s_1 \quad \dots \quad s_{l-1} \quad s_l \quad \dots \quad s_k \\
p_0 & : & U_0 \quad \dots \quad U_{l-1} \quad \dots \quad U_k \\
p_1 & : & \quad U_1 \quad \dots \quad \quad U_l \quad \dots
\end{array}$$

(a) E_k .

$$\begin{array}{rcl}
q & : & s_1 \quad \dots \quad s_{l-1} \quad s_l \quad \dots \quad s_k \\
p_0 & : & U_0 \quad \dots \quad U_{l-1} \quad \dots \quad U_k \\
p_1 & : & \quad U_1 \quad \dots \quad \boxed{U_l} \quad \dots
\end{array}$$

(b) Forward flip: U_l is performed before $U_{l-1}s_l$.

$$\begin{array}{rcl}
q & : & s_1 \quad \dots \quad s_{l-1} \quad s_l \quad \dots \quad s_k \\
p_0 & : & U_0 \quad \dots \quad \boxed{U_{l-1}} \quad \dots \quad U_k \\
p_1 & : & \quad U_1 \quad \dots \quad U_l \quad \dots
\end{array}$$

(c) Backward flip: U_{l-1} is performed after $s_l U_l$.

Figure 2: A flippable execution of length k with two updaters: Figure 2(a) shows a flippable execution E_k ; Figure 2(b) shows the *forward flip* execution of E_k , where the update transaction U_l by process p_1 is executed before the update transaction U_{l-1} by process p_0 and before the step s_l of the read-only transaction; Figure 2(c) shows the *backward flip* execution of E_k , where the update U_{l-1} by process p_0 is deferred after the update transaction U_l by process p_1 and after the step s_l of the read-only transaction.

by processes $p_0, \dots, p_{t-1}$ executing update transactions and process q executing a read-only transaction, which reads and returns the value of t data items $v_0 \dots v_{t-1}$. The execution E_k satisfies all the following conditions:

1. for $j = 1, \dots, k$, s_j is a single step by q ,
2. for $j = 0, \dots, k$, U_j is a solo execution of a complete update transaction, in which process $p_h \in \{p_0, \dots, p_{t-1}\}$, writes $j+1$ to the data item v_h
3. consecutive updates are executed by different processes, and
4. for any l , $0 < l \leq k$, the execution

$$E_k = U_0 s_1 U_1 \dots s_{l-1} U_{l-1} s_l U_l \dots s_k U_k$$

is indistinguishable to all processes from one of the following executions:

$$\overleftarrow{F}_l = U_0 s_1 U_1 \dots s_{l-1} U_l U_{l-1} s_l \dots s_k U_k$$

in which the update transaction U_l is executed before $U_{l-1}s_l$ instead of after $U_{l-1}s_l$ (forward flip) or

$$\overrightarrow{F}_l = U_0 s_1 U_1 \dots s_{l-1} s_l U_l U_{l-1} \dots s_k U_k$$

in which the update transaction U_{l-1} is executed after $s_l U_l$ instead of before $s_l U_l$ (backward flip).

Figures 2(b) and 2(c) present the forward and the backward flips of the execution in Figure 2(a).

This definition, and the structure of our proof, is similar to the lower bound of Attiya, Ellen and Fatourou [2] on the step complexity of update operations in implementations of atomic snapshot objects. The main difference is that our definition of a flippable execution has *two* types of flipped executions, and t processes executing update transactions instead of just two.

The next lemma is proved by arguments similar to those applied in [2], extended to handle the possibility of two kinds of flips (forward and backward).

LEMMA 1. *The read-only transaction in a flippable execution does not terminate successfully.*

PROOF. Let $E_k = U_0 s_1 U_1 \dots s_k U_k$ be a flippable execution. Assume, towards a contradiction, that q successfully terminates its read-only transaction in E_k , with a result $\vec{v} = (v_0, \dots, v_{t-1})$. Since the update transactions in the execution E_k do not overlap, they must be serialized in the order $U_0, \dots, U_l$. Since all steps of the read-only transaction by q are after U_0 and before U_k , it has a unique serialization point between U_{l-1} and U_l , for some l , $1 \leq l \leq k$. Let v_h be the item written by U_{l-1} , and recall that U_{l-1} writes l to v_h ; hence $v_h = l$.

The execution E_k is indistinguishable to process q from F_l , which is either the forward flip

$$\overleftarrow{F}_l = U_0 s_1 U_1 \dots s_{l-1} U_l U_{l-1} s_l s_{l+1} \dots U_k$$

in which update U_l is executed before $U_{l-1}s_l$ instead of after $U_{l-1}s_l$; or the backward flip

$$\overrightarrow{F}_l = U_0 s_1 U_1 \dots s_{l-1} s_l U_l U_{l-1} s_{l+1} \dots U_k$$

in which update U_{l-1} is executed after $s_l U_l$ instead of before $s_l U_l$. Hence, the read-only transaction executed by q in F_l returns the same result, $\vec{v}$, as in E_k .

Since the update transactions do not overlap in F_l , they are serialized in the order $U_0, \dots, U_l, U_{l-1}, \dots, U_k$, that is, the same as for E_k , except that U_{l-1} and U_l are flipped. Since two consecutive update transactions are on different items, the values of $\{v_0, \dots, v_{t-1}\}$ are the same after both update transactions have been executed, no matter which has been executed first. Hence, at all points in the serialization of F_l , except between U_l and U_{l-1} , the value of all items $\{v_0, \dots, v_{t-1}\}$ is the same as its value in the corresponding points in the serialization of E_k . Thus, the read-only transaction of q can only be serialized after U_l and before U_{l-1} in F_l . However, since U_{l-1} is the first write of l to v_h , the value of v_h is not l before U_{l-1} , and hence, the read-only transaction executed by q cannot be serialized between U_l and

Figure 3: Illustration for the proof of Lemma 2

U_{l-1} . This contradicts the assumption that the read-only transaction terminates successfully. $\square$

Lemma 3 (below) proves that when read-only transactions are invisible, we can inductively construct a flippable execution. The crux of this lemma is quite different from [2], as it relies on weakly disjoint-access parallelism. A critical step in the proof is provided by the next lemma, showing that in a weakly disjoint-access parallel STM, two consecutive updates by different processes on different items cannot contend on the same base objects. The proof of the lemma shows that two such consecutive updates can be perturbed to *concurrently* contend on the same base object.

LEMMA 2. *Given a weakly disjoint-access parallel STM implementation and a quiescent configuration C , consider the consecutive execution of two update transactions $U_{j_h} U_{j_{h'}}$, executed by a process p_h on an item i_h and by process $p_{h'}$ on an item $i_{h'}$, $h \neq h'$, respectively, from C . Then p_h and $p_{h'}$ do not contend on the same base object when executing U_{j_h} and $U_{j_{h'}}$.*

PROOF. Assume, towards a contradiction, that p_h and $p_{h'}$ contend on a base object when executing $U_{j_h} U_{j_{h'}}$ from a quiescent configuration C . If in U_{j_h} , p_h applies a non-trivial primitive to a base object on which they contend, let ϕ_h be the last event in U_{j_h} in which p_h applies such a primitive, say, to base object o . Let $\phi_{h'}$ be the first event in $U_{j_{h'}}$ that accesses o . Otherwise, p_h only applies trivial primitives in U_{j_h} to base objects on which it contends with $p_{h'}$ in $U_{j_{h'}}$; let $\phi_{h'}$ be the first event in $U_{j_{h'}}$ in which $p_{h'}$ applies a non-trivial primitive to some base object, say, o , on which they contend. Let ϕ_h be the last event of p_h in U_{j_h} that accesses o . In both cases, denote by $\alpha_h \phi_h$ the prefix of the execution of U_h from C and by $\alpha_{h'} \phi_{h'}$ the prefix of the execution of $U_{h'}$ after U_h (see Figure 3(a)).

We now consider an overlapping execution of the update transactions U_{j_h} and $U_{j_{h'}}$, by processes p_h and $p_{h'}$, from C . We argue that p_h and $p_{h'}$ perform the same steps up to the events ϕ_h and $\phi_{h'}$, and as shown in Figure 3(b), p_h and $p_{h'}$ concurrently contend on base object o .

In more detail, consider the execution $\alpha_h \alpha_{h'}$ from C , in which p_h executes U_{j_h} until it is about to perform ϕ_h , and then $p_{h'}$ executes $U_{j_{h'}}$ until it is about to perform $\phi_{h'}$. Clearly, p_h is about to perform ϕ_h also after $\alpha_h \alpha_{h'}$. By

construction, the execution interval $\alpha_h \alpha_{h'}$ from C is indistinguishable to $p_{h'}$ from the execution interval $U_{j_h} \alpha_{h'}$ from C . Hence, $p_{h'}$ is about to perform the event $\phi_{h'}$ also after $\alpha_h \alpha_{h'}$, that is, $p_{h'}$ and p_h concurrently contend on o . However, the conflict graph of the execution interval $\alpha_h \alpha_{h'} \phi_{h'} \phi_h$ does not contain a path between the data sets of U_{j_h} and $U_{j_{h'}}$, contradicting the assumption that the implementation is weakly disjoint-access parallel. $\square$

Since two consecutive updates do not contend on the same base object, we can construct an execution where either the previous update is deferred or the next update is moved forward in the execution without affecting the single step of the read-only transaction in between them. This allows us to inductively construct a flippable execution, in the proof of the next lemma.

LEMMA 3. *For every $k \geq 0$, every weakly disjoint-access parallel implementation of an STM with invisible read-only transactions, has a flippable execution $E_k = U_0 s_1 U_1 s_2 \dots U_k$ with two updaters p_0 and p_1 , which is indistinguishable to p_0 and p_1 from the execution $E'_k = U_0 U_1 \dots U_k$ in which only p_0 and p_1 take steps.*

PROOF. The proof is by induction on the length k of the flippable execution E_k executed by a process q and two updaters p_0 and p_1 on two items $\{i_0, i_1\}$. In the base case, $k = 0$, the lemma holds with a solo execution of U_0 , an update transaction by p_0 that writes 1 to i_0 . U_0 successfully terminates since it runs solo from a quiescent configuration.

For the induction step, consider a flippable execution of length k , $E_k = U_0 s_1 U_1 s_2 \dots U_k$. By Lemma 1, the read-only transaction does not terminate successfully in E_k . Let s_{k+1} be the next step by q . Assume U_k is executed by $p_{h'}$ and let $h = 1 - h'$; note that $h \neq h'$. Let $E_{k+1} = E_k s_{k+1} U_{k+1}$, where process p_h writes $k+2$ to i_h in the update transaction U_{k+1} . Note that U_{k+1} terminates successfully, by our progress condition; although the configuration at the end of $E_{k+1} = E_k s_{k+1}$ is not quiescent, it is indistinguishable from the quiescent configuration at the end of E'_k .

Since the read-only transaction by q is invisible, E_{k+1} is indistinguishable to p_0 and p_1 from the execution $E'_k U_{k+1}$.

It remains to prove that for every l , $0 < l \leq k+1$, the execution E_{k+1} is indistinguishable to all processes from either $\overleftarrow{F}_l$ or $\overrightarrow{F}_l$. For every l , $0 < l \leq k$, by the inductive assumption, the execution

$$E_k = U_0 s_1 U_1 \dots s_{l-1} U_{l-1} s_l U_l \dots s_k U_k$$

is indistinguishable to all processes from the flipped execution F_l which is either

$$\overleftarrow{F}_l = U_0 s_1 U_1 \dots s_{l-1} U_l U_{l-1} s_l \dots U_k$$

or

$$\overrightarrow{F}_l = U_0 s_1 U_1 \dots s_{l-1} s_l U_l U_{l-1} \dots s_k U_k.$$

In particular, the configurations at the end of the two executions E_k and F_l are the same. Hence, $E_{k+1} = E_k s_{k+1} U_{k+1}$ and $F_l s_{k+1} U_{k+1}$ are indistinguishable to all processes.

To prove the condition for $l = k+1$, let C'_{k-1} be the configuration at the end of E'_{k-1} ; C'_{k-1} is quiescent, and Lemma 2 implies that $p_{h'}$ and p_h do not contend on the same base object when executing U_k followed by U_{k+1} from C'_{k-1} , namely, in the suffix of E'_{k+1} . By the indistinguishability of E'_{k+1} and E_{k+1} , $p_{h'}$ and p_h do not contend on the same base

object while executing U_k and U_{k+1} also in the execution E_{k+1} . Moreover, if q accesses a base object o in s_{k+1} , then either at least one of the two processes p_h or $p_{h'}$ does not access o in U_{k+1} or U_k , respectively, or they both apply a trivial primitive to o . In the former case, if p_h does not access o in U_{k+1} then

$$\overleftarrow{F}_{k+1} = U_0 s_1 U_1 \dots s_k U_{k+1} U_k s_{k+1}$$

is indistinguishable to all processes from E_{k+1} , while if $p_{h'}$ does not access o in U_k , then

$$\overrightarrow{F}_{k+1} = U_0 s_1 U_1 \dots s_k s_{k+1} U_{k+1} U_k$$

is indistinguishable to all processes from E_{k+1} . If both p_h and $p_{h'}$ apply a trivial primitive to o , then both flipped executions, $\overleftarrow{F}_{k+1}$ and $\overrightarrow{F}_{k+1}$, are indistinguishable to all processes from E_{k+1} . $\square$

The impossibility result follows from Lemmas 1 and 3.

THEOREM 4. *There is no weakly disjoint-access parallel implementation with invisible read-only transactions of a strictly serializable STM, in which read-only transactions always terminate successfully.*

The impossibility result stated in Theorem 4 holds also for opaque STMs [10], since opacity implies strict serializability.

3.2 Lower Bound for Read-Only Transactions

The technique of the previous section can be extended to prove that a read-only transaction of t items in a disjoint-access parallel STM implementation, which successfully terminates in a finite number of steps, must apply non-trivial primitives to $t - 1$ base objects; this assumes that there are at least $t + 1$ processes.

The proof of Lemma 1—showing that the read-only transaction in a flippable execution cannot terminate successfully—does not rely on the fact that the read-only transaction is invisible, and the lemma continues to hold. On the other hand, we must modify the proof showing the existence of the flippable execution.

This result relies on a stronger notion of disjoint-access parallelism, which requires two transactions to be connected (in the conflict graph) even if they both just apply a trivial primitive to the same base object. (This is the definition in [17].) Two processes *concurrently access* a base object o if both have a pending access to o at some configuration.

DEFINITION 3. *An STM implementation is disjoint-access parallel if two processes p_1, p_2 concurrently access the same base object when executing transactions T_1 and T_2 , respectively, only if T_1 and T_2 are not disjoint-access.*

Since we now put a stronger requirement on disjoint-access parallel STM implementations, Lemma 2, assuming a weaker requirement, still holds.

We first show (in Lemma 5) that, in a disjoint-access parallel STM implementation, two update transactions executed by different processes on different items do not access a common base object when each of them runs solo from a quiescent configuration. This is used in Lemma 6 to prove the existence of a flippable execution, when a read-only transaction of t data items applies non-trivial primitives to at most $t - 2$ base objects.

Figure 4: Illustration for the proof of Lemma 5.

LEMMA 5. *Given a disjoint-access parallel STM implementation and a quiescent configuration C , consider the execution of an update transaction U_{j_h} to the item v_h by process p_h , and an update transaction $U_{j_{h'}}$ to the item $v_{h'}$ by process $p_{h'}$, $h \neq h'$, from C . Then, p_h and $p_{h'}$ do not access a common base object when executing U_{j_h} and $U_{j_{h'}}$, respectively.*

PROOF. Assume, towards a contradiction, that p_h and $p_{h'}$ access the same base object while executing U_{j_h} and $U_{j_{h'}}$, respectively, from C . Let o be the first base object accessed by p_h that is also accessed by $p_{h'}$. Let $\alpha_h \phi_h$ be the prefix of the execution of U_{j_h} from C , where ϕ_h is the first event in which p_h accesses o (see Figure 4(a)). Let $\alpha_{h'} \phi_{h'}$ be the prefix of the execution of $U_{j_{h'}}$ from C , where $\phi_{h'}$ is the first access of $p_{h'}$ to o (see Figure 4(b)).

Consider the execution $\alpha_h \alpha_{h'}$ from C , where p_h executes U_{j_h} until it is about to access o , and then $p_{h'}$ executes $U_{j_{h'}}$ until it is about to access o (see Figure 4(c)). By construction, the execution $\alpha_h \alpha_{h'}$ from C is indistinguishable to p_h and $p_{h'}$ from the corresponding executions α_h and $\alpha_{h'}$ from C . Thus, $p_{h'}$ has the event $\phi_{h'}$ pending and p_h has the event ϕ_h pending after $\alpha_h \alpha_{h'}$; thus, $p_{h'}$ and p_h have concurrently pending accesses to o . However, in the conflict graph of the execution interval $\alpha_h \alpha_{h'} \phi_{h'} \phi_h$ from C , there is no path between the data sets of U_{j_h} and $U_{j_{h'}}$, contradicting the assumption that the implementation is disjoint-access parallel. $\square$

We show that at any point during the execution of the read-only transaction, there is a process that can write to its item without accessing any base object to which q applies non-trivial primitives, thus making the read-only transaction “invisible” to the other processes. Note that, by the definition of a flippable execution, each process always updates the same item. We prove such a process exists by applying a “pigeon hole” argument to show that the process does not access any base object to which the read-only transaction applies non-trivial primitives. Since there are $t - 1$ processes to choose from, each accessing a different item, and since the read-only transaction applies non-trivial primitives to at most $t - 2$ base objects, at least two update transactions by different processes access the same base object, which can be shown to violate disjoint-access parallelism.

LEMMA 6. *For every $k \geq 0$, a disjoint-access parallel implementation of an STM in which a read-only transaction of t data items applies non-trivial primitives to at most $t - 2$ base objects, has a flippable execution $E_k = U_0 s_1 U_1 s_2 \dots U_k$ with t updaters, which is indistinguishable to $p_0, \dots, p_{t-1}$ from the execution $E'_k = U_0 U_1 \dots U_k$ in which only $p_0, \dots, p_{t-1}$ take steps.*

PROOF. The proof is by induction on the length k of the flippable execution E_k . The base case is when $k = 0$. The lemma holds with a solo execution of an update transaction, U_0 , by process p_0 that writes 1 to v_1 . U_0 successfully terminates since it runs solo from a quiescent configuration.

For the induction step, consider a flippable execution of length k , $E_k = U_0 s_1 U_1 s_2 \dots U_k$, which is indistinguishable to $p_0, \dots, p_{t-1}$ from the execution $E'_k = U_0 U_1 \dots U_k$. By Lemma 1, the read-only transaction does not terminate successfully in E_k . Let s_{k+1} be the next step by q and let C_{k+1} denote the configuration at the end of $E_k s_{k+1}$; also, let C'_{k+1} be the configuration at the end of E'_k .

The process p_h to execute U_{k+1} is chosen from $p_0, \dots, p_{t-1}$ such that p_h did not execute U_k and a solo execution of U_{k+1} from C_{k+1} by p_h does not access any base objects to which q applies non-trivial primitives in $E_k s_{k+1}$. Note that this transaction must terminate successfully, by our progress condition; although C_{k+1} is not quiescent, it is indistinguishable from C'_{k+1} , which is quiescent.

We claim such a process exists. Assume, towards a contradiction, that for every process $p_{h_{k+1}}$, $h_{k+1} \neq h_k$, the solo execution by $p_{h_{k+1}}$ from C_{k+1} of the update transaction that writes $k+2$ to $v_{h_{k+1}}$ accesses a base object to which q applies a non-trivial primitive in $E_k s_{k+1}$. We consider $t-1$ possible processes, each writing to a different item. Since the read-only transaction applies non-trivial primitives to at most $t-2$ base objects, at least two update transactions executed by different processes p_h and $p_{h'}$ to different items v_h and $v_{h'}$, starting from configuration C_{k+1} , access the same base object in their first access to a base object to which q applies a non-trivial primitive. Recall that C'_{k+1} is quiescent. Since the execution $E_k s_{k+1}$ is indistinguishable to processes p_h and $p_{h'}$ from the execution E'_k , they access the same base object also when executing the update transactions from C'_{k+1} , which by Lemma 5, violates the assumption that the implementation is disjoint-access parallel.

Pick some process $p_{h_{k+1}}$, $h_{k+1} \neq h_k$, that does not access any base objects to which q applies non-trivial primitives in $E_k s_{k+1}$; let U_{k+1} be an update by $p_{h_{k+1}}$ that writes $k+2$ to $v_{h_{k+1}}$ and denote $E_{k+1} = E_k s_{k+1} U_{k+1}$.

Next, we prove that the execution E_{k+1} is indistinguishable to $p_0, \dots, p_{t-1}$ from the execution E'_{k+1} . This holds for processes other than $p_{h_{k+1}}$ by the inductive assumption and since these processes take no steps in the suffix of this execution. For $p_{h_{k+1}}$, this holds by the inductive assumption and since the solo execution U_{k+1} of an update transaction by $p_{h_{k+1}}$ does not access base objects to which q applies a non-trivial primitive in $E_k s_{k+1}$.

It remains to prove that for every l , $0 < l \leq k+1$, the execution E_{k+1} is indistinguishable to all processes from the flipped execution F_l which is either $\overleftarrow{F}_l$ or $\overrightarrow{F}_l$, as defined in Definition 2. For every l , $0 < l \leq k$, by the inductive assumption, the execution

$$E_k = U_0 s_1 U_1 \dots s_{l-1} U_{l-1} s_l U_l \dots s_k U_k$$

is indistinguishable to all processes from the flipped execu-

tion F_l which is either

$$\overleftarrow{F}_l = U_0 s_1 U_1 \dots s_{l-1} U_l U_{l-1} s_l \dots U_k$$

or

$$\overrightarrow{F}_l = U_0 s_1 U_1 \dots s_{l-1} s_l U_l U_{l-1} \dots s_k U_k.$$

In particular, the configurations at the end of the two executions E_k and F_l are the same. Hence, the executions $E_{k+1} = E_k s_{k+1} U_{k+1}$ and $F_l s_{k+1} U_{k+1}$ are indistinguishable to all processes.

For $l = k+1$, consider the flipped executions $\overleftarrow{F}_{k+1}$ and $\overrightarrow{F}_{k+1}$. The configuration C'_{k-1} at the end of E'_{k-1} is quiescent. Any STM implementation which is disjoint-access parallel is also weakly disjoint-access parallel, hence we can apply Lemma 2 to deduce that p_{h_k} and $p_{h_{k+1}}$ do not contend on, and hence do not access the same base object while executing U_k and U_{k+1} from C'_{k-1} . The indistinguishability property implies that p_{h_k} and $p_{h_{k+1}}$ do not access the same base object while executing U_k and U_{k+1} also in E_{k+1} .

Moreover, if q applies a trivial primitive to some base object o in s_{k+1} , then either at least one of the two processes $p_{h_{k+1}}$ and p_{h_k} does not access o in U_{k+1} and in U_k respectively, or they both apply a trivial primitive to o . In the former case, if $p_{h_{k+1}}$ does not access in U_{k+1} any object that q accesses in s_{k+1} , then

$$\overleftarrow{E}_{k+1} = U_0 s_1 U_1 \dots s_k U_{k+1} U_k s_{k+1}$$

is indistinguishable to all processes from E_{k+1} , while if p_{h_k} does not access in U_k any object that q accesses in s_{k+1} , then

$$\overrightarrow{E}_{k+1} = U_0 s_1 U_1 \dots s_k s_{k+1} U_{k+1} U_k$$

is indistinguishable to all processes from E_{k+1} . If $p_{h_{k+1}}$ and p_{h_k} apply a trivial primitive to o , then both flipped executions are indistinguishable to all processes from E_{k+1} . $\square$

The lower bound follows:

THEOREM 7. *In a strict serializable disjoint-access parallel STM implementation for $t+1$ processes, where all read-only transactions terminate successfully, some read-only transaction of $t \geq 2$ data items applies non-trivial primitives to at least $t-1$ base objects.*

This lower bound holds also for opaque STMs, since opacity implies strict serializability.

4. EXTENDING THE RESULTS TO WEAKER CONSISTENCY CONDITIONS

In this section, we show that both Theorem 4 and Theorem 7 hold for weaker consistency conditions, namely, serializability and snapshot isolation. This uses an additional process.

Recall that an STM is *serializable* if transactions appear to execute sequentially, one after the other; we further require that transactions of the same process preserve their order (*per-process* order).

Given a flippable execution $E_k = U_0 s_1 U_1 \dots s_k U_k$, we construct an *augmented flippable execution*

$$\widehat{E}_k = U_0 s_1 S_1^* U_1 \dots s_k S_k^* U_k,$$

where an additional process q' performs invisible read-only transactions. For every $j \in \{1, \dots, k\}$, q' performs solo a

Figure 5: An augmented flippable execution $\hat{E}_k$ derived from the flippable execution E_k of Figure 2.

sequence S_j^* of read-only transactions after the event s_j by process q and before the update U_j . Each read-only transaction in S_j^* accesses the items $\iota_{f_{j-1}}$ and ι_{f_j} updated by U_{j-1} and U_j . The result of the last read-only transaction in the sequence S_j^* , denoted S_j , is the value written by U_{j-1} to $\iota_{f_{j-1}}$ and the last value of ι_{f_j} before U_j updates it.

Figure 5 shows the augmented flippable execution obtained by augmenting the flippable execution E_k of Figure 2 with sequences of read-only transactions performed by an additional process q' .

We apply the per-process ordering of transactions to prove that the read-only transactions of q' must eventually read the latest value written in U_{j-1} , and thus, S_j^* is finite.

LEMMA 8. *Consider an augmented flippable execution of length $k \geq 0$, $\hat{E}_k = U_0 s_1 S_1^* U_1 \dots s_k S_k^* U_k$. In any serialization of $\hat{E}_k$ that preserves the per-process order, $U_0, U_1, \dots, U_k$ appear in their order of execution.*

PROOF. We show, by induction on ℓ , that $U_0, U_1, \dots, U_\ell$ appear in their order of execution. The base case is trivial.

For the induction step, consider $U_{\ell+1}$. By the induction assumption, the updates $U_0, U_1, \dots, U_\ell$ are serialized by their execution order in $\hat{E}_k$. By construction $S_{\ell+1}^*$ accesses the items ι_{f_ℓ} and $\iota_{f_{\ell+1}}$ repeatedly up to a read-only transaction $S_{\ell+1}$, which returns the value written by U_ℓ and the last value of $\iota_{f_{\ell+1}}$ before the one written by $U_{\ell+1}$.

$S_{\ell+1}^*$ is finite since the STM is serializable and so, eventually, some transaction must return the latest values written to ι_{f_ℓ} and $\iota_{f_{\ell+1}}$, and by the induction assumption, U_ℓ is the last to write to ι_{f_ℓ} . Moreover, $S_{\ell+1}$ completes before $U_{\ell+1}$ starts, so it cannot return the value written by $U_{\ell+1}$, since due to serializability, a read operation can not return a value not written.

Since each data item is written by a different process, and due to per-process order, $U_{\ell+1}$ can not be serialized before the last update of $\iota_{f_{\ell+1}}$ preceding $U_{\ell+1}$.

Moreover, $U_{\ell+1}$ can not be serialized after this update and before $S_{\ell+1}$, since $S_{\ell+1}$ does not return the value written by $U_{\ell+1}$. Hence, $U_{\ell+1}$ is serialized after $S_{\ell+1}$. $\square$

We use Lemma 8 to prove the analogue of Lemma 1.

LEMMA 9. *Consider an augmented flippable execution of length $k \geq 0$ with t updaters, $\hat{E}_k = U_0 s_1 S_1^* U_1 \dots s_k S_k^* U_k$. If the read-only transactions by process q' are invisible, then the read-only transaction by process q does not terminate successfully.*

PROOF. Assume, towards a contradiction, that the read-only transaction of process q in $\hat{E}_k$ terminates successfully and returns a value $\vec{v} = (v_0, \dots, v_{t-1})$, which does not violate serializability. Let the augmented flippable execution $\hat{E}_k = U_0 s_1 S_1^* U_1 \dots s_k S_k^* U_k$ correspond to a flippable execution $E_k = U_0 s_1 U_1 \dots s_k U_k$.

By Lemma 8, the updates in $\hat{E}_k$ are serialized in the order $U_0, U_1, \dots, U_k$. The value $\vec{v}$ determines where q 's read-only transaction is serialized. In particular, for some l , $0 < l \leq k$, the read-only transaction of q is serialized after U_{l-1} and before U_l , and for each item ι_f in $\{\iota_0 \dots \iota_{t-1}\}$, either v_f is zero and no update wrote to ι_f before U_l , or the last update to ι_f before U_l wrote v_f to ι_f . Let S be the serialization of execution $\hat{E}_k$.

Since the read-only transactions executed by process q' are invisible, $\hat{E}_k$ and E_k are indistinguishable to $p_0, \dots, p_{t-1}$ and q . Thus, they will execute the same steps in both executions. Note that S is a serialization also for E_k . Since S preserves the real-time order among transactions, E_k is a flippable execution where the read-only transaction terminates and strict serializability is preserved, contradicting Lemma 1. $\square$

As discussed before the lemma, the existence of a flippable execution (guaranteed by Lemma 3) implies there is an augmented flippable execution, and hence, Lemma 9 implies the following impossibility result:

THEOREM 10. *There is no weakly disjoint-access parallel STM implementation with invisible read-only transactions of a serializable STM, in which read-only transactions always terminate successfully.*

When a read-only transaction of $t \geq 2$ data items applies non-trivial primitives to at most $t-2$ base objects, the read-only transactions of q' in the augmented flippable execution are, in fact, invisible since their read set contains only two data items. As discussed before Lemma 9, the existence of a flippable execution (guaranteed by Lemma 6) implies there is an augmented flippable execution, and hence, Lemma 9 implies the following lower bound:

THEOREM 11. *In a serializable disjoint-access parallel STM implementation for $t+2$ processes, where all read-only transactions terminate successfully, some read-only transaction of $t \geq 2$ data items applies non-trivial primitives to at least $t-1$ base objects.*

Snapshot isolation [19, 23, 25] decouples the consistency of the reads and the writes, and guarantees a snapshot of the read set not older than the start of the transaction. The proof can be adapted to hold also when the consistency condition of the STM is snapshot isolation.

5. RELATED WORK

Many STM implementations are centralized; in particular, to determine a unique commit timestamp for transactions, the *Lazy Snapshot Algorithm* (LSA) [22] relies on a single shared monotonically increasing counter, while *Transactional Locking II* (TL2) [6] relies on a global clock. Both

approaches introduce a single hot-spot accessed by all transactions, regardless of their data sets, and are therefore not disjoint-access parallel.

More recently, two STM implementations without a centralized hot-spot have been proposed. Avni and Shavit [4] present a *thread-local clock* mechanism that provides a decentralized solution for maintaining a consistent view. The key idea is using Lamport clock (scalar causal timestamps) instead of the real-time global clock. Integrated with TL2, this mechanism provides an STM supporting invisible read-only transactions, without a centralized contention point. A drawback of this algorithm is that transactions that terminated long before the current one may cause it to fail since the timestamp recorded for them is not current enough. Thus, read-only transactions are not wait-free. Imbs and Raynal [16] propose an opaque lock-based STM with no centralized hot-spot but their solution has visible reads.

Guerraoui and Kapalka [9] prove that obstruction-free implementations of software transactional memory cannot ensure *strict* disjoint-access parallelism. This property requires transactions with disjoint data sets not to access a common base object. This notion is stronger than the one originally proposed by Israeli and Rappoport [17], and commonly used in the literature [14], where two transactions with disjoint data sets are allowed to access the same base objects, provided they are connected via other transactions. All other transactions have to progress in parallel, even if they are concurrent. Their definition of strict disjoint-access parallelism, like our first definition (Definition 1), allows concurrent reads to the same base objects even by transactions that are not connected in the conflict graph.

Our lower bound applies to the notion of disjoint-access parallelism as originally defined in [17]. In contrast, the result of [9] does not hold for this weaker notion. Indeed, Herlihy et al. [13] present an obstruction-free and disjoint-access parallel STM. Obstruction-freedom does not prevent interfering concurrent processes from starving each other and thus, the implementation presented in [13] does not guarantee that a read-only transaction eventually terminates successfully.

Elsewhere, Guerraoui and Kapalka [10] prove a lower bound on the number of steps a process takes to successfully terminate a transaction, for every implementation that uses invisible reads, is single-version, and never aborts a transaction unless it conflicts with another live transaction. Our lower bound allows multi-version implementations, but requires read-only transactions to terminate successfully, regardless of overlapping transactions.

Serializability provides a weaker guarantee on the ordering of transactions (it does not have to respect the real-time order of non-overlapping ones). Nevertheless, our impossibility results hold also for serializable STMs that preserve the per-process order. Indeed, none of the serializable STM implementations presented in the literature, e.g. [5, 7, 20, 24], provides disjoint-access parallelism and wait-free, invisible read-only transactions. In fact, the impossibility results hold also for STMs that satisfy the even weaker condition of *snapshot isolation* known from the database literature [19, 25] and suggested as an efficient alternative to serializability for STMs [23].

Riegel et al. [24] proposed an STM implementation that supports invisible reads and is disjoint-access parallel, but it provides only *causal serializability*; moreover, read-only

transactions may abort infinitely many times. Causal serializability is weaker than serializability since it allows different processes to have a different view of the system. This leaves open the question of whether our results holds for causally serializable STMs, or whether the algorithm of [24] can be extended to have wait-free read-only transactions.

A read-only transaction can be considered as a *partial* SCAN operation [3]: a *partial snapshot objects* is an atomic snapshot object [1], where processes can scan any subset of the components. In the wait-free algorithm for partial snapshot objects [3], scanners announce which components they are currently attempting to scan, i.e., read-only transactions are visible.

Our proof techniques draw ideas from the lower bounds on the step complexity of UPDATE operations in snapshot objects. Israeli and Shirazi [18] prove an $\Omega(m)$ lower bound on the number of steps to update a component in an m -component single-writer snapshot objects, implemented from single-writer registers. Attiya, Ellen and Fatourou [2] extend this lower bound to implementations of m -component multi-writer objects from base objects of any type.

6. DISCUSSION

This paper shows that no transactional memory implementation can be disjoint-access parallel and have invisible, wait-free read-only transactions. There are implementations that are disjoint-access parallel and have invisible but not wait-free read-only transactions [4, 13], while others have invisible, wait-free read-only transactions but are not disjoint-access parallel [22]. In principle, the invisibility of read-only transactions can also be sacrificed in order to keep them wait-free, and the implementation disjoint-access parallel. This can be done by treating the read set together with the write set and adapting a dynamic disjoint-access parallel implementation of multi-location synchronization operator, e.g., [12]. (This algorithm is not wait-free, but it seems that it can be made wait-free without sacrificing the other properties.) Thus, each of the assumptions made in our impossibility result is necessary, since removing either of them admits an implementation with the two remaining properties.

Our work joins recent efforts to explore the boundaries of STM implementations, so as to guide algorithm designers in their attempt to find better and more efficient implementations. Such boundaries demonstrate which directions are futile and which might lead to performance gains. It would be interesting to derive additional quantitative results on the complexity of transactions, and in particular, read-only transactions.

Our proof shows that the read-only transaction cannot terminate successfully, but it is possible to terminate it unsuccessfully, by *aborting* it; however, this abort is not justified by data conflicts. Moreover, when the read-only transaction is retried, it is possible to continue the construction and force it to abort again. An implementation is *permissive* with respect to a safety property [8] if it never aborts a transaction unless necessary for ensuring correctness. Our proof shows that a disjoint-access parallel implementation with invisible read-only transactions that always terminate—however, not always successfully—is not permissive with respect to opacity, strict serializability, serializability or snapshot isolation. We would like to further investigate the connections between our results and the study of

unnecessary aborts [7, 8] or *wasted work* in STM implementations.

Acknowledgements. We would like to thank Rachid Guerraoui, Michal Kapalka and Martin Vechev for helpful comments.

7. REFERENCES

- [1] Y. Afek, H. Attiya, D. Dolev, E. Gafni, M. Merritt, and N. Shavit. Atomic snapshots of shared memory. *J. ACM*, 40(4):873–890, 1993.
- [2] H. Attiya, F. Ellen, and P. Fatourou. The complexity of updating multi-writer snapshot objects. In *ICDCN '06*, pages 319–330.
- [3] H. Attiya, R. Guerraoui, and E. Ruppert. Partial snapshot objects. In *SPAA '08*, pages 336–343.
- [4] H. Avni and N. Shavit. Maintaining consistent transactional states without a global clock. In *SIROCCO '08*, pages 131–140.
- [5] U. Aydonat and T. Abdelrahman. Serializability of transactions in software transactional memory. In *TRANSACT '08*.
- [6] D. Dice, O. Shalev, and N. Shavit. Transactional locking II. In *DISC '06*, pages 194–208.
- [7] V. Gramoli, D. Harmanici, and P. Felber. Towards a theory of input acceptance for transactional memories. In *OPODIS '08*, pages 527–533.
- [8] R. Guerraoui, T. A. Henzinger, and V. Singh. Permissiveness in transactional memories. In *DISC '08*.
- [9] R. Guerraoui and M. Kapalka. On obstruction-free transactions. In *SPAA '08*, pages 304–313.
- [10] R. Guerraoui and M. Kapalka. On the correctness of transactional memory. In *PPoPP '08*, pages 175–184.
- [11] R. Guerraoui and M. Kapalka. The semantics of progress in lock-based transactional memory. In *POPL '09*, pages 404–415.
- [12] T. L. Harris, K. Fraser, and I. A. Pratt. A practical multi-word compare-and-swap operation. In *DISC '02*, pages 265–279.
- [13] M. Herlihy, V. Luchangco, M. Moir, and W. N. Scherer III. Software transactional memory for dynamic-sized data structures. In *PODC '03*, pages 92–101.
- [14] M. Herlihy and N. Shavit. *The Art of Multiprocessor Programming*. Morgan Kaufmann, 2008.
- [15] M. P. Herlihy and J. M. Wing. Linearizability: a correctness condition for concurrent objects. *ACM Trans. Program. Lang. Syst.*, 12(3):463–492, 1990.
- [16] D. Imbs and M. Raynal. A lock-based protocol for software transactional memory. In *OPODIS '08*, pages 226–245.
- [17] A. Israeli and L. Rappoport. Disjoint-access-parallel implementations of strong shared memory primitives. *PODC '94*, pages 151–160.
- [18] A. Israeli and A. Shirazi. The time complexity of updating snapshot memories. *Inf. Process. Lett.*, 65(1):33–40, 1998.
- [19] S. Lu, A. Bernstein, and P. Lewis. Correct execution of transactions at different isolation levels. *IEEE Transactions on Knowledge and Data Engineering*, 16(9):1070–1081, 2004.
- [20] J. Napper and L. Alvisi. Lock-free serializable transactions. Technical Report TR-05-04, The University of Texas at Austin, 2005.
- [21] C. H. Papadimitriou. The serializability of concurrent database updates. *J. ACM*, 26(4):631–653, 1979.
- [22] T. Riegel, P. Felber, and C. Fetzer. A lazy snapshot algorithm with eager validation. In *DISC '06*, pages 284–298.
- [23] T. Riegel, C. Fetzer, and P. Felber. Snapshot isolation for software transactional memory. In *TRANSACT '06*.
- [24] T. Riegel, C. Fetzer, H. Sturzrehm, and P. Felber. From causal to z-linearizable transactional memory. In *PODC '07*, pages 340–341.
- [25] G. Weikum and G. Vossen. *Transactional Information Systems: Theory, Algorithms, and the Practice of Concurrency Control and Recovery*. Morgan Kaufmann, 2001.