

Between Linearizability and Quiescent Consistency*

Quantitative Quiescent Consistency

Radha Jagadeesan and James Riely

DePaul University

Abstract Linearizability is the de facto correctness criterion for concurrent data structures. Unfortunately, linearizability imposes a performance penalty which scales linearly in the number of contending threads. Quiescent consistency is an alternative criterion which guarantees that a concurrent data structure behaves correctly when accessed sequentially. Yet quiescent consistency says very little about executions that have any contention.

We define quantitative quiescent consistency (QQC), a relaxation of linearizability where the degree of relaxation is proportional to degree of contention. When quiescent, no relaxation is allowed, and therefore QQC refines quiescent consistency, unlike other proposed relaxations of linearizability. We show that high performance counters and stacks designed to satisfy quiescent consistency continue to satisfy QQC. The precise assumptions under which QQC holds provides fresh insight on these structures. To demonstrate the robustness of QQC, we provide three natural characterizations and prove compositionality.

1 Introduction

This paper defines *Quantitative Quiescent Consistency (QQC)* as a criterion that lies between linearizability [9] and quiescent consistency [3, 10]. The following example should give some intuition about these criteria.

Example 1.1. Consider a counter object with a single `getAndIncrement` method. The counter's sequential behavior can be defined as a set of strings such as $[^+]_0^+ \{^+ \}_1^+ (^+)_2^+$ where $[^+]_i^+$ denotes an invocation (or call) of the method and $]_i^+$ denotes the response (or return) with value i . Suppose each invocation is initiated by a different thread.

A concurrent execution may have overlapping method invocations. For example, in $(^+ [^+]_0^+ \{^+ \}_1^+)_2^+$ the execution of $(^+)_2^+$ overlaps with both $[^+]_0^+$ and $\{^+ \}_1^+$, whereas $[^+]_0^+$ finishes executing before $\{^+ \}_1^+$ begins. Consider the following four executions.

$$(^+ [^+]_0^+ \{^+ \}_1^+)_2^+ \quad (^+ \{^+ \}_1^+ [^+]_0^+)_2^+ \quad [^+ (^+)_2^+ \{^+ \}_1^+]_0^+ \quad [^+ (^+)_2^+]_0^+ \{^+ \}_1^+$$

Linearizability states roughly that every response-to-invocation order in a concurrent execution must be consistent with the sequential specification. Thus, the first execution is linearizable, since the response $]_0^+$ precedes the invocation $\{^+ \}_1^+$ in the specification. However, none of the other executions is linearizable. For example, $]_1^+$ precedes $[^+]_0^+$ in the second execution, but $]_1^+$ does not precedes $[^+]_0^+$ in the specification.

* Research supported by NSF 0916741.

Linearizability can also be understood in terms the *linearization point* of a method execution, which must occur between the invocation and response. From this perspective, the first execution above is linearizable because we can find a sequence of linearization points that agrees with the specification; this requires only that the linearization point of $(^+)_2^+$ follow that of $(^+)_1^+$. No such sequence of linearization points exists for the two other executions.

Quiescent consistency is similar to linearizability, except that the response-to-invocation order must be respected only across a quiescent point, that is, a point with no open method calls. The first three executions above are quiescently consistent simply because there are no non-trivial quiescent points. The last execution fails to be quiescently consistent since the order from $)_2^+$ to $(^+)_1^+$ is not preserved in the specification.

We define *Quantitative Quiescent Consistency (QQC)* to require that the number of response-to-invocation pairs that are out-of-order at any point be bounded by the number of open calls at that point. We also give a *counting characterization* of QQC, which requires that if a response matches the i^{th} method call in the specification, then it must be preceded by at least i invocations.

The first two executions above are QQC; however, the last two are not. In the second execution, the open call to $(^+)_1^+$ is “enough” to justify the return of $(^+)_1^+$ before $(^+)_0^+$. However, in the third execution, the return of $(^+)_2^+$ before $(^+)_1^+$ cannot be justified only by the call to $(^+)_1^+$; it is “too far off.” Following the counting characterization sketched above, the third execution fails since $(^+)_2^+$ is the third method call in the specification trace, but $)_2^+$ is only preceded by two invocations: $(^+)_1^+$ and $(^+)_0^+$. \square

Quiescent consistency is too coarse to be of much use in reasoning about concurrent executions. For example, a sequence of interlocking calls never reaches a quiescent point; therefore it is trivially quiescently consistent. This includes obviously correct executions, such as $(^+)_0^+ (^+)_1^+ (^+)_2^+ (^+)_3^+ (^+)_4^+ \dots$, nearly correct executions, such as $(^+)_1^+ (^+)_0^+ (^+)_3^+ (^+)_2^+ (^+)_5^+ \dots$, and also ridiculous executions, such as $(^+)_1^{1074} (^+)_17^+ (^+)_2^{2344} (^+)_3^+ (^+)_5^+ \dots$.

Linearizability has proven quite useful in reasoning about concurrent executions; however, it fundamentally constrains efficiency in a multicore setting: Dwork, Herlihy, and Waarts [5] show that if many threads concurrently access a linearizable counter, there must be either a location with high contention or an execution path that accesses many shared variables.

Shavit [11] argues that the performance penalty of linearizable data structures is increasingly unacceptable in the multi-core age. This observation has led to a recent renewal of interest in nonlinearizable data structures. As a simple example, consider the following counter implementation: a simplified version of the counting networks of Aspnes, Herlihy, and Shavit [3].

```

1 class Counter<N:Int> {
2   field b:[0..N-1] = 0; // 1 balancer
3   field c:Int[] = [0, 1, ..., N-1]; // N counters
4   method getAndIncrement():Int {
5     val i:[0..N-1];
6     atomic { i = b; b++; }
7     atomic { val v = c[i]; c[i] += N; return v; } }

```

The N -Counter has two fields: a *balancer* b and an array c of N integer counters. There are two atomic actions in the code: The first reads and updates the balancer, setting the local index variable i . The second reads and updates the i^{th} counter. Although the balancer has high contention in our simplified implementation, the counters do not; balancers that avoid high contention are described in [3].

Example 1.2. The N -Counter behaves like a sequential counter if calls to `getAndIncrement` are sequentialized. To see this, consider a 2-Counter, with initial state $\langle b = 0, c = [[0], [1]] \rangle$. In a series of sequential calls, the state progresses as follows, where we show the execution of the first atomic with the invocation and the second atomic with the response. The execution $\llbracket^+ \rrbracket_0^+ \llbracket^+ \rrbracket_1^+ \llbracket^+ \rrbracket_2^+$ can be elaborated as follows.

$$\begin{aligned} \langle b = 0, c = [[0], [1]] \rangle &\xrightarrow{\llbracket^+ \rrbracket_0^+} \langle b = 1, c = [[0], [1]] \rangle \xrightarrow{\llbracket^+ \rrbracket_1^+} \langle b = 1, c = [[2], [1]] \rangle \\ &\xrightarrow{\llbracket^+ \rrbracket_1^+} \langle b = 0, c = [[2], [1]] \rangle \xrightarrow{\llbracket^+ \rrbracket_1^+} \langle b = 0, c = [[2], [3]] \rangle \\ &\xrightarrow{\llbracket^+ \rrbracket_2^+} \langle b = 1, c = [[2], [3]] \rangle \xrightarrow{\llbracket^+ \rrbracket_2^+} \langle b = 1, c = [[4], [3]] \rangle \end{aligned}$$

When there is concurrent access, the 2-Counter allows nonlinearizable executions, such as $\llbracket^+ \rrbracket_1^+ \llbracket^+ \rrbracket_0^+ \llbracket^+ \rrbracket_2^+$.

$$\begin{aligned} \langle b = 0, c = [[0], [1]] \rangle &\xrightarrow{\llbracket^+ \rrbracket_1^+} \langle b = 1, c = [[0], [1]] \rangle \\ &\xrightarrow{\llbracket^+ \rrbracket_0^+} \langle b = 0, c = [[0], [1]] \rangle \xrightarrow{\llbracket^+ \rrbracket_1^+} \langle b = 0, c = [[0], [3]] \rangle \\ &\xrightarrow{\llbracket^+ \rrbracket_2^+} \langle b = 1, c = [[0], [3]] \rangle \xrightarrow{\llbracket^+ \rrbracket_0^+} \langle b = 1, c = [[2], [3]] \rangle \\ &\xrightarrow{\llbracket^+ \rrbracket_2^+} \langle b = 1, c = [[4], [3]] \rangle \end{aligned}$$

With a sequence of interlocking calls, it is also possible for the N -Counter to execute as $\llbracket^+ \rrbracket_1^+ \llbracket^+ \rrbracket_0^+ \llbracket^+ \rrbracket_3^+ \llbracket^+ \rrbracket_2^+ \llbracket^+ \rrbracket_5^+ \llbracket^+ \rrbracket_4^+ \dots$, producing an infinite sequence of values that are just slightly out of order. Using the results of this paper, one can conclude that with a maximum of two open calls, the value returned by `getAndIncrement` will be “off” by no more than 2, but this does not follow from quiescent consistency. \square

Our results are related to those of [2–4, 13]. In particular, Aspnes, Herlihy, and Shavit [3] prove that in any *quiescent* state (with no call that has not returned), such a counter has a “step-property”, indicating the shape of c . Between $\llbracket^+ \rrbracket_1^+$ and $\llbracket^+ \rrbracket_0^+$ in the second displayed execution of Example 1.2, the states with $c = [[0], [3]]$ do *not* have the step property, since the two adjacent counters differ by more than 1.

Aspnes, Herlihy, and Shavit imply that the step property is related to quiescent consistency, but they do not formally state this. Indeed, they do not provide a formal definition of quiescent consistency. It appears that they have in mind is something like the following: An execution is *weakly quiescent consistent* if any subsequence of *sequential* calls (single calls separated by quiescent points) is a subtrace of a specification trace.

The situation is delicate: Although the increment-only counters of [3] are quiescently consistent in the sense we defined in Example 1.1 (indeed, they are QQC), the increment-decrement counters of [2, 4, 13] are only *weakly* quiescent consistent. Indeed, the theorems proven in [13] state only that, at a quiescent point, a variant of the step property holds. It states nothing about the actual values read from the individual counters. Instead, we expect that a quiescently consistent execution should be a permutation of *some* specification trace, even if it has no nontrivial quiescent points.

Example 1.3. Consider an extension of the 2-Counter with `decrementAndGet`.

```
method decrementAndGet():Int {
  val i:[0..N-1];
  atomic { i = b-1; b--; }
  atomic { c[i] -= N; return c[i]; } }
```

The execution $[^+ \{^+ \langle^- \langle^- \rangle_{-2} \rangle_{-2} \}^+]_1^+$ is possible, although this is not a permutation of any specification trace. The execution proceeds as follows.

$$\begin{aligned} \langle b = 0, c = [[0], [1]] \rangle &\xrightarrow{I^+} \langle b = 1, c = [[0], [1]] \rangle \xrightarrow{I^+} \langle b = 0, c = [[0], [1]] \rangle \\ &\xrightarrow{C^-} \langle b = 1, c = [[0], [1]] \rangle \xrightarrow{C^-} \langle b = 0, c = [[0], [1]] \rangle \\ &\xrightarrow{Z_{-2}^+} \langle b = 0, c = [[-2], [1]] \rangle \xrightarrow{J_{-2}^+} \langle b = 0, c = [[0], [1]] \rangle \\ &\xrightarrow{J_1^+} \langle b = 0, c = [[0], [3]] \rangle \xrightarrow{J_1^+} \langle b = 0, c = [[0], [1]] \rangle \quad \square \end{aligned}$$

It is important to emphasize that this increment-decrement counter is not even quiescently consistent. There is no hope that it could satisfy any stronger criterion.

Of course counters are not the only data structures of interest. In this paper we treat concurrent stacks in detail. We define a simplified N -Stack below; the full, tree-based data structure is defined in Shavit and Touitou [13] and summarized in section 6.

```
1 class Stack<N:Int> {
2   field b:[0..N-1] = 0; // 1 balancer
3   field s:Stack[] = [[], [], ..., []]; // N stacks of values
4   method push(x:Object):Unit {
5     val i:[0..N-1];
6     atomic { i = b; b++; }
7     atomic { val v = s[i].push(x); return v; } }
8   method pop():Object {
9     val i:[0..N-1];
10    atomic { i = b-1; b--; }
11    atomic { val v = s[i].pop(); return v; } }
```

The trace given in Example 1.3 for the increment-decrement counter is also a trace of the stack, where we interpret $+$ as push, $-$ as pop, and a negative return value as stack underflow. Whereas this is a nonsense execution for a counter, it is a linearizable execution of a stack: simply choose the linearization point of the pops before the pushes. Nonetheless, the N -Stack is only *weakly* quiescent consistent in general.

Example 1.4. The N -Stack generates the execution $[^+ [^+]^+ \langle^+ \langle^- \rangle_a \rangle_b \rangle^+$ as follows.

$$\begin{aligned} \langle b = 0, s = [[], []] \rangle &\xrightarrow{I^+} \langle b = 1, s = [[], []] \rangle \xrightarrow{J^+} \langle b = 1, s = [[a], []] \rangle \\ &\xrightarrow{C_a^+} \langle b = 0, s = [[a], []] \rangle \xrightarrow{J^+} \langle b = 0, s = [[a], [b]] \rangle \\ &\xrightarrow{C^+} \langle b = 1, s = [[a], [b]] \rangle \\ &\xrightarrow{C^-} \langle b = 0, s = [[a], [b]] \rangle \xrightarrow{Z_a^+} \langle b = 0, s = [[], [b]] \rangle \\ &\xrightarrow{J^+} \langle b = 0, s = [[c], [b]] \rangle \end{aligned}$$

However, this specification is not quiescently consistent with any stack execution: There is a quiescent point after each of the first two pushes; therefore it is impossible to pop a before b . This execution is possible even when there are several pushes beforehand. \square

In the case of the *N-Stack*, a simple *local* constraint can be imposed in order to establish quiescent consistency: intuitively, we require that no pop *overtakes* a push on the same stack $s[i]$. In section 6 we show that the stack is actually *QQC* under this constraint, and therefore quiescently consistent. (We have not found a local constraint under which the increment-decrement counter is quiescently consistent; we believe that it may be achievable with a global toggle that determines how to resolve the races at each point, but this, of course, defeats the point.)

The correctness result that we prove for elimination-tree stacks in section 6 is much stronger than that of Shavit and Touitou [13], who only prove *weak* quiescent consistency. The same is holds for increment-only counters [3], although in this case, we have elided the proofs: The proofs for the tree-based increment-only counter follow the structure proofs for the elimination-tree stacks.

The preliminary version of Shavit and Touitou’s paper [12] suggests an upcoming definition *ϵ -linearizability*, “a variant of linearizability that captures the notion of ‘almostness’ by allowing a certain fraction of concurrent operations to be out-of-order.” Since the details did not make it into the final version of the paper [13], it is unclear whether the “fraction of concurrent operations” is meant to vary depending on the amount of concurrency available in execution at any given moment, or if the “fraction” is fixed at the outset. If it is meant to vary, then it would be very similar to *QQC*.

This thread was picked up by Afek, Korland, and Yanovsky [1] and improved by Henzinger, Kirsch, Payer, Sezgin, and Sokolova [8]. As defined in [8], the idea is to define a cost metric on relaxations of strings and to bound the relaxation cost for the specification trace that matches an execution. This approach has been used to validate several novel concurrent data structures [1, 6].

Unlike *QQC*, the metrics in these papers do not depend on the available concurrency in the execution. In fact, with the exception of the increment-only counter validated in [1], all of the concurrent data structures of [1, 6] intentionally violate quiescent consistency. As such, their work is orthogonal to the approach we pursue here, which specifically refines quiescent consistency.

The primary difference between *QQC* and the relaxation-based approach of [1, 8] can be stated as follows: with *QQC*, the maximal degradation depends upon the amount of concurrent access, whereas in the relaxation-based approach it does not. Thus, *QQC* “degrades gracefully” as concurrency increases. In particular, a *QQC* data structure that is accessed sequentially will exactly obey the sequential specification, whereas a data structure validated against the relaxation-based approach may not.

In the rest of the paper, we formalize *QQC* and study its properties. The heart of the paper is section 5, which defines *QQC* and establishes its properties. The impatient reader can safely skim up to that section, referring back as necessary.

Our contributions are as follows.

- We define linearizability (section 3), quiescent consistency (section 4) and *QQC* (section 5) in terms of partial orders over events with duration. The formalities of the model are described in section 2. As in Example 1.1, the definitions are given in terms of the order between a response and a subsequent invocation.

- We provide an alternative characterization of QQC in terms of the number of invocations that precede a response, as well as a similar characterization of linearizability.
- We provide an alternative characterization of QQC in terms of a proxy that controls access to the underlying sequential data structure. The proxy adds a form of *speculation* to the flat combining technique of Hendler, Incze, Shavit, and Tzafrir [7]. We show that the operational semantics is sound and complete for QQC; that is, it generates *all and only* traces that are QQC.
- Like linearizability and quiescent consistency [10], QQC is non-blocking and compositional. Like quiescent consistency and unlike linearizability, a QQC execution may not respect program order. We provide a proof for compositionality.
- We show that QQC is useful for reasoning about data structures in the literature. In section 6, we prove that the elimination tree stacks of Shavit and Touitou [13] are QQC, as long as no pop overtakes a push on the same stack.

2 Model

The semantics of a concurrent program is given as a process. A *process* is a set of traces. A *trace* is a finite, polarized LPO (labelled partial order). Formally, we define traces to be finite sets of named *events*. The event names are the carrier set for the LPO, and the order is embedded in the events themselves using name sets.

2.1 Events

An event is a quadruple, consisting of a polarity, a label, a name (identifying a node the partial order) and a set of names (identifying the preceding nodes in the partial order). As a standard example, the reader may want to consider labels generated by the grammar $\ell ::= \text{call } \tau of w \mid \text{ret } \tau of w$ where τ is a thread identifier, o is an object name, f is a function name, and w is the actual parameter or return value.

Let $a, b \in \text{Name}$ range over names and $A, B \subseteq \text{Name}$ range over finite sets of names. And let $\ell \in \text{Label}$ range over labels (with some interpretation in the application domain). Then events are defined as follows¹.

$$u, v ::= \langle ?\ell \rangle_A^a \mid \langle b\ell \rangle_A^a$$

Under our standard example, we would expect events to come in pairs of the form $\langle ?\text{call } \tau of w \rangle_A^a$ and $\langle a\text{ret } \tau of w' \rangle_B^b$, where w is the actual parameter and w' is the returned value. The appearance of a in the return event indicates that this event closes the open call named a .

Three of the components in an event can be retrieved simply. We use the following functions: $\text{label}(\langle ?\ell \rangle_A^a) \triangleq \ell$, $\text{id}(\langle ?\ell \rangle_A^a) \triangleq a$ and $\text{before}(\langle ?\ell \rangle_A^a) \triangleq A$. For the remaining component, we define both the functions pol and brak . Let $p \in \{?, !\}$ range over the polarities for input (?) and output (!) and let none be a reserved name.

$$\text{pol}(u) \triangleq \begin{cases} ? & \text{if } u = \langle ?\ell \rangle_A^a \\ ! & \text{if } u = \langle b\ell \rangle_A^a \end{cases} \quad \text{brak}(u) \triangleq \begin{cases} \text{none} & \text{if } u = \langle ?\ell \rangle_A^a \\ b & \text{if } u = \langle b\ell \rangle_A^a \end{cases}$$

¹ In this paper, we consider the simple case of non-interacting composition. This allows us to ignore the internal polarity which arise from the interaction of input and output.

Because the standard example is so familiar, we will consider invocation/call/input/? to be synonymous, and likewise response/return/output/!.

We sometimes use superscripts on name metavariables, such as $a^!$ and $a^?$. Any name bound to $a^!$ must be associated with an output event, and likewise for input events. The superscript makes these distinct metavariables. Thus we have $a^! \neq a^?$.

Turning to the order between events, we write $u \Rightarrow v$ to indicate that u precedes v : $(u \Rightarrow v) \triangleq \text{id}(u) \in \text{before}(v)$.

2.2 Traces

We use p - t to range over *event sets* (finite sets of events). Define $\text{ids}(s) \triangleq \{\text{id}(u) \mid u \in s\}$ and let $a \in s$ be shorthand for $a \in \text{ids}(s)$.

Given an event set s and name set A , define *indexing* as $s[A] \triangleq \{u \in s \mid \text{id}(u) \in A\}$. Thus $s[\text{ids}(s)] = s$. If event names are unique, this generates the partial function $s[a]$ for single names: if $s[\{a\}] = \emptyset$ then $s[a]$ is undefined; if $s[\{a\}] = \{u\}$ then $s[a] \triangleq u$. Indexing provides a natural way to lift ordering relations from events to names: $(a \Rightarrow_s b) \triangleq (s[a] \Rightarrow s[b])$. Let \Rightarrow_s be the reflexive closure of \Rightarrow_s .

An event set s is a *trace* if it satisfies the following, $\forall u, v \in s$.


- (1) event names are unique: if $\text{id}(u) = \text{id}(v)$ then $u = v$
- (2) before okay: $\forall a \in \text{before}(u). \exists w \in s. a = \text{id}(w)$
- (3) brak okay: if $\text{pol}(u) = !$ then $\text{brak}(u) \in \text{before}(u)$ and $\text{pol}(s[\text{brak}(u)]) = ?$
- (4) input acquires control: if $a \Rightarrow_s b^?$ then $\exists c^!. a \Rightarrow_s c^! \Rightarrow_s b^?$
- (5) output releases control: if $a^! \Rightarrow_s b$ then $\exists c^?. a^! \Rightarrow_s c^! \Rightarrow_s b$
- (6) \Rightarrow_s defines a strict partial order (irreflexive, antisymmetric and transitive)

A trace s is *operational* if $\forall a^?, b^! \in s$. either $a^? \Rightarrow_s b^!$ or $b^! \Rightarrow_s a^?$.

A trace s is *sequential* if $\forall a, b \in s$. either $a \Rightarrow_s b$ or $b \Rightarrow_s a$.

Our model can be viewed as a labelled partial order enriched with polarity and bracketing. Most significant here are requirements (4) and (5) in the definition of a trace. One immediate consequence is that input events cannot be related to other input events unless there is an intervening output event, and similarly for the dual case.

Consider two bracketed event sequences $[]$ and $()$. As indicated by condition (3) in the definition of traces, the open brackets must be ? events. There are six possible relations among the events. Three of these are familiar: it could be that $[]$ precedes $()$, or that $()$ precedes $[]$ or that they are concurrent.




All of these traces are fully specified in the sense that every ? is ordered with respect to every !, and dually every ! is ordered with respect to every ?. We call such traces *operational* in that they correspond to traces generated by an interleaving semantics. In addition, the first two traces are *sequential*, since there is a total order on the events.

Note that in any sequential trace, the initial event must be an input; this follows from properties (2) and (3) in the definition of traces.

There is a homomorphism from strings of bracketed labels to operational traces: each input in the string is ordered with respect to each output that follows it in the string, and dually. If we narrow attention to sequential traces, this is an isomorphism. For example, we can write the first two traces above as the strings $\$ and $(\)[\]$, respectively. The last trace above can be written as any interleaving in $[\] \parallel (\)$ that orders the inputs before both outputs; these are $[\ (\)$, $(\ [\])$, $[\ (\)$, and $(\ [\])$. We use this notation when giving examples of operational traces, as in the introduction.

As a consequence of the homomorphism, we can use string notation on operational traces without ambiguity. Specifically, let st represent the concatenation of s and t and $s \parallel t$ represent the set of their interleavings, with renaming as necessary to avoid collisions between the names of s and t .

Our model also allows underspecification of the relationship.



The leftmost of these says only that $(\)$ cannot precede $[\]$. Said positively, either $[\]$ precedes $(\)$, or they are concurrent. The rightmost of these places no constraints on the relative order of $[\]$ and $(\)$.

Operational traces can be seen as having a global notion of time: everyone agrees what happened before what. The constraints between pairs of inputs and pairs of output simply indicate the limits of observability: it is impossible to tell which of two calls happened first. In this light, one may view an underspecified trace as a representative for the set of operational traces that can be derived by augmenting the partial order. We take this viewpoint in our compositionality result, which is stated only for operational traces.

We define several notations for event sets and traces.

As noted above, for operational traces s and t we use string notation for concatenation (st) and the set of interleavings ($s \parallel t$).

A *renaming* of a trace is identical to the original trace up to a bijection on names². We write $=_\alpha$ for equivalence up to renaming.

A *permutation* of a trace contains events with the same names, labels and polarities, but may differ in ordering. Permutation does *not* allow renaming; the names to pick out the witnessing bijection. We write $=_\pi$ for equivalence up to permutation³.

A *prefix* is a down-closed subtrace⁴. We write $t \leq_{\text{pre}} s$ or $s \geq_{\text{pre}} t$ to indicate that t is a prefix of s , and $\downarrow_s a$ for the smallest down-closed subset of s that includes a .

² ($s =_\alpha t$) is defined to mean that there exists a bijection α on names such that (1) $\text{ids}(s) = \alpha(\text{ids}(t))$, and (2) $\forall a \in \text{ids}(s). s[a] = t[\alpha(a)]$. (In the first condition, we have used the obvious homomorphic extension of α to sets of names.)

³ Let $(u =_{\text{brak}}^{\text{label}} v) \triangleq (\text{label}(u) = \text{label}(v)) \wedge (\text{brak}(u) = \text{brak}(v))$.

Then define $(s =_\pi t) \triangleq (\text{ids}(s) = \text{ids}(t)) \wedge (\forall a \in s. s[a] =_{\text{brak}}^{\text{label}} t[a])$.

⁴ Trace t is a *prefix* of trace s if $\forall a, b \in s. \text{if } a \in t \text{ and } b \Rightarrow_s a \text{ then } b \in t$.

We treat traces both as sets of events and as partial orders. We use $(-)$ for set difference and (\div) for partial order difference⁵.

3 Linearizability

We give two characterizations of linearizability and prove compositionality.

In subsection 3.1, we give a characterization that looks at every way to *cut* a trace into prefix and suffix; linearizability requires that response-to-invocation order be respected across all cuts. This corresponds to characterization of QQC given in subsection 5.1. In the case of QQC, a certain number of invocations may be ignored, proportional to the number of calls that are both open across the cut and out of specification-order with respect to the response.

In subsection 3.2 we give a subset-based characterization, which requires that if a response matches the i^{th} method call in the specification, then it must be preceded by at the first i invocations of the specification. This corresponds to characterization of QQC given in subsection 5.2. In the case of QQC, it is sufficient that a response by the i^{th} method be preceded by any i invocations, not necessarily the first i invocations of the specification.

The proof of compositionality in subsection 3.3 is provided as a warmup for the proof compositionality for QQC in subsection 5.4.

3.1 First characterization: response to invocation

Intuitively, linearizability requires that the response-to-invocation order in an execution be respected by a specification trace. To show that s'' is linearizable, it suffices to do the following

- Choose a specification trace t .
- Choose an extension s' of s'' that closes the open calls in s'' . Let $\text{extensions}(s'')$ be the set of *extensions* of s'' , such that (1) if $s' \in \text{extensions}(s'')$ then s'' is a prefix of s' ($s' \geq_{\text{pre}} s''$), and (2) all of the new events in $s' - s''$ are ordered after all events of opposite polarity in s'' ⁶ (that is, calls after returns and returns after calls).
- Choose a renaming $s =_{\alpha} s'$ such that $s =_{\pi} t$. Informally, this establishes that s' is a permutation of t . Rather than carrying the permutation around in the definition, as usual in definitions of linearizability, we perform a renaming up front, once and

⁵ An event set t is *bracketed* if every output in t has a matching input in t ; that is $\forall u \in t$. if $\text{pol}(u) = !$ then $\text{brak}(u) \in \text{ids}(t)$. A bracketed set may contain unmatched inputs, but not unmatched outputs.

For arbitrary event sets, we write $s - t$ for set difference. For trace s and bracketed event set t , we write $s \div t$ for partial order difference. For example, consider the trace the sequential trace $s = \langle ?\ell_1 \rangle_{\emptyset}^a \langle a\ell_2 \rangle_{\{a\}}^{a'} \langle ?\ell_3 \rangle_{\{a,a'\}}^b \langle b\ell_4 \rangle_{\{a,a',b\}}^{b'} \langle ?\ell_5 \rangle_{\{a,a',b,b'\}}^c$ and let t be the bracketed set $\{s[b], s[b']\}$. Then we have $s - t = \langle ?\ell_1 \rangle_{\emptyset}^a \langle a\ell_2 \rangle_{\{a\}}^{a'} \langle ?\ell_5 \rangle_{\{a,a',b,b'\}}^c$ and $s \div t = \langle ?\ell_1 \rangle_{\emptyset}^a \langle a\ell_2 \rangle_{\{a\}}^{a'} \langle ?\ell_5 \rangle_{\{a,a'\}}^c$.

⁶ $\text{extensions}(p) \triangleq \{s \mid p \leq_{\text{pre}} s \wedge \forall a^! \in p. \forall b^? \in s - p. a^! \Rightarrow_p b^? \wedge \forall a^? \in p. \forall b^! \in s - p. a^? \Rightarrow_p b^!\}$

for all. The names are then witness to the permutation. This works nicely, since our traces are indexed by names. Typically, linearizability is defined over strings, indexed by integers, so this technique is not available.

- Show that for every response a^1 in s and invocation $b^?$ in s , if a^1 precedes $b^?$ in s ($a^1 \Rightarrow_s b^?$), then the same must be true in t ($a^1 \Rightarrow_t b^?$).

Stated compactly, we have the following definition.

Definition 3.1. Trace s'' linearizes to t if $\exists s' \in \text{extensions}(s'')$. $\exists s =_\alpha s'$. $s =_\pi t$ and

$$\forall a^1 \in s. \forall b^? \in s. (a^1 \Rightarrow_s b^?) \text{ implies } (a^1 \Rightarrow_t b^?).$$

Trace set S linearizes to T if $\forall s'' \in S. \exists t \in T. s''$ linearizes to t . □

This definition differs from the traditional one in several small details, but is equivalent under reasonable assumptions. The differences are as follows.

- We do not require that specifications be sequential.
- We do not make requirements specific to threads. A thread is simply a totally ordered sequence of actions, with the result that every pair of invocations must be separated by a response, and similarly for pairs of responses. The fact that thread order is respected by linearizability follows from the general requirement that order from response to invocation must be respected.
- In addition to *returns*, we allow $s \in \text{extensions}(s')$ to include *calls* that are not in s' . Assuming that specifications are prefix-closed, this permissiveness is harmless. For every spec t that includes the extra calls in a suffix, there is a corresponding spec t' such that $t \in \text{extensions}(t')$ that does not include them; if s' linearizes to t , then it also linearizes to t' .
- We require that all incomplete calls remain in s' . Assuming that specifications are input-enabled, this restriction is harmless. For every spec t that does not include the extra calls, there is a corresponding spec $t' \in \text{extensions}(t)$ that does include them; if s' linearizes to t with some incomplete calls removed, then it also linearizes to t' .

We can refactor the definition slightly to pull it into the shape used to define quiescent consistency and QQC.

Definition 3.2. For traces s, t , we write $s \sqsubseteq_{\text{lin}} t$ if $s =_\pi t$ and for every prefix $p \leq_{\text{pre}} s$

$$\forall a^1 \in p. \forall b^? \in s - p. (a^1 \Rightarrow_s b^?) \text{ implies } (a^1 \Rightarrow_t b^?).$$

Then $(s'' \sqsubset_{\text{lin}} t) \triangleq (\exists s' \in \text{extensions}(s''). \exists s =_\alpha s'. s \sqsubseteq_{\text{lin}} t)$. □

Lemma 3.3. s is linearizes to t iff $s \sqsubset_{\text{lin}} t$.

PROOF. This is an immediate consequence of the definition of prefix. □

This characterization of linearizability requires that we look at every way to *cut* the trace s into a prefix p and suffix $s - p$. We then look at the return events in p and the call events in $s - p$ and ensure that the order of events *crossing the cut* is respected in t . The definitions are equivalent since we quantify over all possible cuts.

As an example, consider the incrementing counter specification from Example 1.1: $[^+]_0^+ \{^+ \} _1^+ (^+)_2^+$. For a completely concurrent trace, such as $[^+ \{^+ (^+)_2^+ \} _1^+]_0^+$ linearizability is trivially satisfied since there is no cut that has a return on the left and call on the right. The trace $\{^+ [^+ \} _1^+ (^+)_2^+ \}$ is also linearizable. The interesting cut is $\{^+ [^+ \} _1^+$ which requires only that $\} _1^+$ precede $(^+)_2^+$ in the specification. By the same reasoning, $\{^+ (^+ \} _1^+ [^+]_0^+ \}$, is not linearizable, since it requires that $\} _1^+$ precede $[^+]_0^+$.

3.2 Second characterization: invocation to response

Theorem 3.4. Let t be a sequential trace with name order $(a_1^?, a_1^!, a_2^?, a_2^!, \dots, a_n^?, a_n^!)$. Let s be an operational trace such that $s =_{\pi} t$. Then

$$s \sqsubseteq_{\text{lin}} t \quad \text{iff} \quad \forall j. \{1, \dots, j\} \subseteq \{i \mid a_i^? \Rightarrow_s a_j^!\}$$

PROOF. Using the definition of linearizability and calculating, we have the following proof obligation.

$$(\forall i, j. a_i^! \Rightarrow_s a_j^? \text{ implies } i < j) \quad \Leftrightarrow \quad (\forall i, j. i \leq j \text{ implies } a_i^? \Rightarrow_s a_j^!)$$

(\Rightarrow) Fix $i \leq j$. If $i = j$ the right implication holds by the definition of traces. Suppose $i < j$. By operability, either $a_i^? \Rightarrow_s a_j^!$ or $a_j^! \Rightarrow_s a_i^?$. In the first case, the right implication holds. In the second case, the left implication requires $j < i$, a contradiction.

(\Leftarrow) Fix $a_i^! \Rightarrow_s a_j^?$. By way of contradiction, suppose $i \leq j$. From the right implication we deduce that $a_i^? \Rightarrow_s a_j^!$. The resulting cycle, $a_i^! \Rightarrow_s a_j^? \Rightarrow_s a_i^!$ contradicts the supposition that s is a trace. Therefore it must be that $i < j$ as required. \square

Let us revisit the incrementing counter specification $[^+]_0^+ \{^+ \} _1^+ (^+)_2^+$. In the completely concurrent trace $[^+ \{^+ (^+)_2^+ \} _1^+]_0^+$ all invocations precede all responses, and therefore linearizability is trivially satisfied. The linearizability of $\{^+ [^+ \} _1^+ (^+)_2^+$ follows from the fact that $\} _1^+$ is preceded by both $[^+]_0^+$ and $\{^+ \}$, and the non-linearizability of $\{^+ (^+ \} _1^+ [^+]_0^+ \}$, follows from the fact that $[^+]_0^+$ does not precede $\} _1^+$.

3.3 Compositionality

We re-prove one of the fundamental properties of linearizability: compositionality [9]. The proof we give here is similar to the proof given for QQC in subsection 5.4, in a simpler setting.

Lemma 3.5 (Operational traces). Suppose that s is an operational trace that imposes the following order:

$$\begin{array}{ccc} a_1^? & a_0^? & b_0^? & b_1^? \\ & \searrow & \swarrow & \downarrow \\ & a_0^! & & b_0^! \end{array}$$

Then either $a_1^? \Rightarrow_s b_0^!$ or $b_1^? \Rightarrow_s a_0^!$.

PROOF. If neither holds, then, by operationality we must have both $b_0^! \Rightarrow_s a_1^?$ and $a_0^! \Rightarrow_s b_1^?$, which results in the cycle $b_0^! \Rightarrow_s a_1^? \Rightarrow_s a_0^! \Rightarrow_s b_1^? \Rightarrow_s b_0^!$. \square

Recall from subsection 2.2 that (\parallel) denotes interleaving and (\div) denotes partial order difference. To split trace s in “half,” it suffices to postulate the existence of s_1 and s_2 such that $s_1 = s \div s_2$ and $s_2 = s \div s_1$.

Theorem 3.6. *Let t_1 and t_2 be sequential traces.*

Let s , s_1 and s_2 be operational traces such that $s_1 = s \div s_2$ and $s_2 = s \div s_1$.

For $i \in \{1, 2\}$, suppose that each $s_i \sqsubseteq_{\text{lin}} t_i$.

Then there exists a sequential trace $t \in (t_1 \parallel t_2)$ such that $s \sqsubseteq_{\text{lin}} t$.

PROOF. Without loss of generality, assume that $\text{ids}(t_1)$ and $\text{ids}(t_2)$ are disjoint. Let the sequence of names in t_1 be $(a_1^?, a_1^!, \dots, a_m^?, a_m^!)$ and sequence of name in t_2 be $(b_1^?, b_1^!, \dots, b_n^?, b_n^!)$. Applying Theorem 3.4 to the supposition $s_1 \sqsubseteq_{\text{lin}} t_1$, we have that $i \leq j$ implies $a_i^? \Rightarrow_s a_j^!$, and similarly for the b s.

Our aim is to construct a sequential interleaving of t_1 and t_2 . To do this, we construct a partial order over event pairs. Any interleaving consistent with the partial order will satisfy the conclusion of the theorem by construction. For the elements of the partial order, let a_i represent the pair $a_i^? a_i^!$ and let b_k represent the pair $b_k^? b_k^!$. Let the a s be totally ordered by subscript, corresponding to the fact that $a_i^? \Rightarrow_s a_j^!$ whenever $i \leq j$, and similarly the b s. Let there be a *cross edge* from a_i to b_ℓ if $a_i^! \Rightarrow_s b_\ell^?$, and symmetrically from b s to a s. Visually, we have an order such as the following.

$$\begin{array}{ccccccc} a_1 & \rightarrow & a_2 & \rightarrow & \dots & \rightarrow & a_i & \rightarrow & \dots & \rightarrow & a_j & \rightarrow & \dots & \rightarrow & a_m \\ & & & & & & \swarrow & & & & \searrow & & & & \\ b_1 & \rightarrow & b_2 & \rightarrow & \dots & \rightarrow & b_k & \rightarrow & \dots & \rightarrow & b_\ell & \rightarrow & \dots & \rightarrow & b_n \end{array}$$

The a - a and b - b edges go from ? to ! in s , whereas the cross edges go from ! to ?.

The proof obligation is to show that this order is acyclic, in which case it induces at least one interleaving. We show that any cycle in the defined order corresponds to a cycle in s , contradicting the supposition that s is a trace. For there to be a cycle in the defined order, there must be $i < j$ and $k < \ell$, such that $a_i^? \Rightarrow_s a_j^! \Rightarrow_s b_k^? \Rightarrow_s b_\ell^! \Rightarrow_s a_i^?$. This contradicts the supposition that s is a trace. \square

4 Quiescent Consistency

Let $\text{open}(s)$ be the set of calls in s that have no matching return⁷. We say that trace s is *quiescent* if $\text{open}(s) = \emptyset$. This notion of quiescence does not require that there be no active thread, but only that there be no open calls. Thus, this notion of quiescence is compatible with libraries that maintain their own thread pools.

The definition of quiescent consistency is similar to Definition 3.2 of linearizability. The difference lies in the quantifier for the prefix $p \leq_{\text{pre}} s$: Whereas linearizability quantifies over *every* prefix, quiescent consistency only quantifies over *quiescent* prefixes.

⁷ $\text{open}(s) \triangleq \{u \in s \mid \text{pol}(u) = ? \wedge \nexists v \in s. \text{brak}(v) = \text{id}(u)\}$

Definition 4.1. We write $s \sqsubseteq_{\text{qc}} t$ if $s =_{\pi} t$ and for any *quiescent* prefix $p \leq_{\text{pre}} s$

$$\forall a^! \in p. \forall b^? \in s - p. (a^! \Rightarrow_s b^?) \text{ implies } (a^! \Rightarrow_t b^?).$$

Then $(s'' \sqsubseteq_{\text{qc}} t) \triangleq (\exists s' \in \text{extensions}(s''). \exists s =_{\alpha} s'. s \sqsubseteq_{\text{qc}} t)$. \square

Again let us revisit the counter specification from Example 1.1: $[^+]_0^+ \{^+ \}_1^+ (^+)_2^+$. This notion quiescent consistency places some constraints on the system even when it has no nontrivial quiescent points. For example, the execution $[^+ \{^+ (^+)_3^+ \}_1^+]_0^+$ is not quiescently consistent with the given specification, since it is not a permutation. If one extends the execution to $[^+ \{^+ (^+)_3^+ \}_1^+]_0^+ <^+ >_2^+$ and attempts to match it against the specification $[^+]_0^+ \{^+ \}_1^+ <^+ >_2^+ (^+)_3^+$, quiescent consistency continues to fail: In the quiescent prefix $[^+ \{^+ (^+)_3^+ \}_1^+]_0^+$, the order across the cut from $]_3^+$ to $<^+$ is not preserved in the specification.

For linearizability, we argued that because specifications are prefix-closed, only responses need be included in the extensions of a trace. The same does not hold for quiescent consistency. For example, since $(^+ \{^+ \}_1^+ [^+]_0^+)_2^+$ is quiescently consistent, its prefix $(^+ \{^+ \}_1^+)$ should also be quiescently consistent. However, there is no specification trace that can be matched that does not include $[^+]_0^+$. Therefore, it does not suffice merely to close the open call by adding $]_2^+$; We must include $[^+$ and $]_0^+$.

Compositionality (as expressed in Theorem 3.6) also holds for quiescent consistency. The proof is straightforward: any quiescent point of $s_1 \cup s_2$ is also a quiescent point for each s_i ; the two specifications may be interleaved arbitrarily between these quiescent points.

As noted in the introduction, if the sequence of interlocking calls $[^+ (^+]_i^+ [^+]_j^+ (^+)_k^+]_k^+ \dots$, never reaches quiescence, then the counter may return any natural number for i, j and k . QQC reduces this permissiveness by looking at every cut. It remains less strict than linearizability by loosening the requirement that *every* response-to-invocation across the cut be respected in the specification.

5 Quantitative Quiescent Consistency

We provide three characterizations of QQC and prove their equivalence.

- In subsection 5.1, we define QQC in the style that we have defined linearizability and quiescent consistency, from response to invocation.
- In subsection 5.2, we give a *counting characterization* of QQC, which requires that if a response matches the i^{th} method call in the specification, then it must be preceded by at least i invocations.
- In subsection 5.3, we give an operational characterization of QQC as a proxy between the concurrent world and an underlying sequential data structure. This can be seen a mix of flat combining Hendler, Ince, Shavit, and Tzafrir [7] with speculation.

Finally, in subsection 5.4, we demonstrate that QQC is compositional, as in [9].

To give some intuition for the what is allowed by QQC, we first give some examples using the 2-Counter from the introduction. First we note that the capability given by an open call can be used repeatedly, as in the execution $(^+ \llbracket^+ \rrbracket_1 \{^+ \rrbracket_0^+ \llbracket^+ \rrbracket_3 \{^+ \rrbracket_2^+ \llbracket^+ \rrbracket_5 \{^+ \rrbracket_4^+ \rrbracket_6^+}$. Alternatively, multiple open calls may be accumulated to create a trace with events that are arbitrarily far off, as in $(^+ \llbracket^+ \rrbracket_1^+ (^+ \llbracket^+ \rrbracket_3^+ (^+ \llbracket^+ \rrbracket_5^+ (^+ \llbracket^+ \rrbracket_7^+ \llbracket^+ \rrbracket_0^+ \rrbracket_2^+ \rrbracket_4^+ \rrbracket_6^+ \rrbracket_8^+)$. Note that $\llbracket^+ \rrbracket_0^+$ follows $\llbracket^+ \rrbracket_7^+$ in this execution! It is worth emphasizing that the order between these actions is observable to the outside: a single thread can call `getAndIncrement` and get 7, then subsequently call `getAndIncrement` and get 0.

In general, an N -Counter can give results that are $k \times N$ off of the expected value, where k is the maximum number of open calls and N is the width of the counter. There is no way to bound the behavior of this counter, as in [8], without also bounding the amount of concurrency, as in [1].

It is also possible for open calls to overlap in nontrivial ways. The trace $(^+ \llbracket^+ \rrbracket_1^+ \{^+ \llbracket^+ \rrbracket_0^+ \rrbracket_3^+ (^+ \rrbracket_2^+ \rrbracket_4^+)$ is QQC. Here, the first $(^+$ justifies the out-of-order execution of $\llbracket^+ \rrbracket_1^+$ and $\llbracket^+ \rrbracket_0^+$. The subsequent $\{^+$ justifies an inversion of the previous justifier, namely $(^+ \rrbracket_3^+$ and $(^+ \rrbracket_2^+$. A similar example is $\{^+ (^+ \rrbracket_1^+ (^+ \llbracket^+ \rrbracket_0^+ \rrbracket_3^+ \llbracket^+ \rrbracket_2^+ \rrbracket_4^+)$.

Finally, we note that the stack execution $\{^+ \llbracket^+ \rrbracket_a^- (^+ \rrbracket_a^+ \rrbracket^+$ is QQC with respect to the specification $(^+ \rrbracket_a^+ \llbracket^+ \rrbracket_a^- \{^+ \rrbracket^+$. This follows from exactly the kind of reasoning that we have done for the counter. This lack of causality may be troubling, but we note that it is typical of weak correctness criteria such as quiescent consistency. We revisit this issue in the conclusions.

5.1 First characterization: response to invocation

Linearizability requires that for *every* cut, *all* response-to-invocation order crossing the cut must be respected in the specification. Quiescent consistency limits this attention to *quiescent* cuts. QQC restores the quantification over every cut, but relaxes the requirement to match all response-to-invocation order crossing the cut. When checking response-to-invocation pairs across the cut, QQC allows some invocations to be ignored. How many?

One constraint comes from our desire to refine quiescent consistency. For quiescent cuts, we cannot drop any invocations, since this quiescent consistency does not. As a first attempt at a definition, we may take the number of dropped invocations at any cut to be bounded by $|\text{open}(p)|$. However, this is too permissive. For example, this simple criterion cannot distinguish the following executions from Example 1.1.

$$(^+ \{^+ \rrbracket_1^+ \llbracket^+ \rrbracket_0^+ \rrbracket_2^+ \qquad \llbracket^+ (^+ \rrbracket_2^+ \{^+ \rrbracket_1^+ \rrbracket_0^+$$

The interesting cut splits the traces in half at the midpoint. In each case there is one open call. Therefore in the first trace, we can ignore \llbracket^+ in the suffix, and in the second trace, we can ignore $\{^+$ in the suffix. However, we believe that these traces should be distinguished. The second trace is “more off” than the first.

The difference can be seen by looking not only at the number of open calls, but also at *which* calls are open. In the first trace we have $(^+$ before \rrbracket_1^+ , and in the second, we have \llbracket^+ before \rrbracket_2^+ . We say that $(^+$ is *early* for \rrbracket_1^+ , since it does not precede \rrbracket_1^+ in the specification, whereas \llbracket^+ is not early for \rrbracket_2^+ , since it *does* precede \rrbracket_2^+ . We restrict our attention to calls that are both open and early with respect to the response of interest.

Given a specification t and $a^! \in t$ none of the actions in the t -downclosure of $a^!$ could possibly be early for $a^!$; any other action could be. Thus, the actions in $\text{open}(p) - (\downarrow_t a^!)$ are both open and early for $a^!$. This leads us to the following definition. (In subsection 5.2, we show that for sequential specifications, we can swap the quantifiers $(\exists r)$ and $(\forall a^!)$, pulling out the existential.)

Definition 5.1. We write $s \sqsubseteq_{\text{qqc}} t$ if $s =_{\pi} t$ and for any prefix $p \leq_{\text{pre}} s$

$$\begin{aligned} \forall a^! \in p. \exists r \subseteq s. |r| \leq |\text{open}(p) - (\downarrow_t a^!)|. \\ \forall b^? \in ((s - p) - r). (a^! \Rightarrow_s b^?) \text{ implies } (a^! \Rightarrow_t b^?). \end{aligned}$$

Then $(s'' \sqsubseteq_{\text{qqc}} t) \triangleq (\exists s' \in \text{extensions}(s''). \exists s =_{\alpha} s'. s \sqsubseteq_{\text{qqc}} t)$. \square

In this definition, it is safe to restrict attention to sets r consisting only of input events that are concurrent with the open calls. We do not impose these restrictions explicitly because they are not necessary. Choosing outputs does not add any flexibility, effectively wasting an open call. Non-concurrent calls will be revealed by the prefix in which the call is closed.

Theorem 5.2. $(\sqsubseteq_{\text{lin}}) \subset (\sqsubseteq_{\text{qqc}}) \subset (\sqsubseteq_{\text{qc}})$

PROOF. Containment is immediate from the definitions, always taking $r = \varepsilon$ for QQC. To see that the containment is proper, consider the incrementing counter specification from Example 1.1, $[^+]_0^+ (^+)_1^+ \{^+ \}_2^+$. With respect to this specification, $\{^+ (^+)_1^+ [^+]_0^+ \}_2^+$ is QQC but not linearizable $[^+ \{^+ \}_2^+ (^+)_1^+]_0^+$ is quiescently consistent but not QQC. \square

5.2 Second characterization: counting invocations

Given the subtlety of Definition 5.1, it may be surprising that QQC has the following simple characterization for sequential specifications.

Theorem 5.3. Let t be a sequential trace with name order $(a_1^?, a_1^!, \dots, a_n^?, a_n^!)$. Let s be an operational trace such that $s =_{\pi} t$. Then

$$s \sqsubseteq_{\text{qqc}} t \quad \text{iff} \quad \forall j. j \leq |\{a_i^? \mid a_i^? \Rightarrow_s a_j^!\}|$$

PROOF. (\Rightarrow) Fix j , let $p = \downarrow_s a_j^!$, and let q, r', o be the following disjoint sets.

$$\begin{aligned} q &= \{a_i^? \mid i \leq j \wedge a_i^? \Rightarrow_s a_j^!\} \\ r' &= \{a_i^? \mid i \leq j \wedge a_i^? \not\Rightarrow_s a_j^!\} = \{a_i^? \mid i \leq j \wedge a_j^! \Rightarrow_s a_i^?\} \quad (\text{by operationality}) \\ o &= \{a_i^? \mid i > j \wedge a_i^? \Rightarrow_s a_j^!\} \supseteq \text{open}(p) - (\downarrow_t a_j^!) \quad (\text{by calculation}) \end{aligned}$$

Note that $q \cup o = \{a_i^? \mid a_i^? \Rightarrow_s a_j^!\}$; therefore it suffices to show that $|q \cup o| \geq j$.

For every event in $a_i^? \in r'$ we have that $i \leq j$ and therefore $a_j^! \Rightarrow_s a_i^?$ and $a_j^! \not\Rightarrow_t a_i^?$. Hence the set r chosen in Definition 5.1 must include r' . From Definition 5.1, we have that $|r| \leq |\text{open}(p) - (\downarrow_t a_j^!)r|$. Since $r' \subseteq r$ and $\text{open}(p) - (\downarrow_t a_j^!) \subseteq o$, we have $|r'| \leq |o|$. Since $|q \cup r'| = j$, we have $|q \cup o| \geq j$, as required.

(\Leftarrow) Fix p . Following the argument given in the proof of Lemma 5.4, in order to show that the requirements of Definition 5.1 hold for every $a^1 \in p$, it suffices to show that they hold for a_j^1 , where let $j = \max\{k \mid a_k^1 \in p\}$.

Fix $j = \max\{k \mid a_k^1 \in p\}$. We now show that the requirements of Definition 5.1 hold for a_j^1 . We choose q , r and o as before.

$$\begin{aligned} q &= \{a_i^? \mid i \leq j \wedge a_i^? \Rightarrow_s a_j^1\} \\ r &= \{a_i^? \mid i \leq j \wedge a_i^? \not\Rightarrow_s a_j^1\} = \{a_i^? \mid i \leq j \wedge a_j^1 \Rightarrow_s a_i^?\} \\ o &= \{a_i^? \mid i > j \wedge a_i^? \Rightarrow_s a_j^1\} \subseteq \text{open}(p) - (\downarrow a_j^1) \end{aligned}$$

To see that $o \subseteq \text{open}(p)$, consider that if $a_i^? \in o$ then $a_i^1 \notin p$; otherwise $j \neq \max\{k \mid a_k^1 \in p\}$. By the second characterization of r above (which follows from operability), $\forall a_i^? \notin r$. ($a_j^1 \Rightarrow_s a_i^?$) implies $j < i$. Thus, to establish the result it suffices to show that $|r| \leq |\text{open}(p) - (\downarrow a_j^1)|$. By assumption, $|q \cup o| \geq j$. Since $|q \cup r| = j$, we have $|r| \leq |o|$ and therefore $|r| \leq |\text{open}(p) - (\downarrow a_j^1)|$ as required. \square

This characterization provides a simple method for calculating whether a trace is QQC. For example the trace $\{^+ (^+)_1^+ (^+ [^+]_0^+)_3^+ [^+]_2^+]_4^+\}$ is QQC since $]_1^+$ is preceded by two calls, $]_0^+$, $]_3^+$ by four, and $]_2^+$, $]_4^+$ by five. The trace $\{^+ (^+)_1^+ (^+)_3^+ [^+]_0^+ [^+]_2^+]_4^+\}$ is not QQC since $]_3^+$ is only preceded by three calls, yet it is the fourth call in the specification.

For sequential specifications, we can also simplify Definition 5.1 by exchanging the quantifiers ($\exists r$) and ($\forall a^1$), pulling out the existential.

Lemma 5.4. *Let t be a sequential trace with name order $(a_1^?, a_1^1, \dots, a_n^?, a_n^1)$. Let s be an operational trace such that $s =_\pi t$. Fix $p \leq_{\text{pre}} s$. Then the displayed requirement of Definition 5.1 is equivalent to*

$$\begin{aligned} \exists r \subseteq s. |r| \leq |\text{openEarly}_t(p)|. \\ \forall a^1 \in p. \forall b^? \in ((s - p) - r). (a^1 \Rightarrow_s b^?) \text{ implies } (a^1 \Rightarrow_t b^?), \end{aligned}$$

where $\text{openEarly}_t(p) \triangleq \{b^? \in \text{open}(p) \mid \nexists a^1 \in p. b^? \Rightarrow_t a^1\}$.

PROOF. (5.4 \Rightarrow 5.1) Immediate.

(5.1 \Rightarrow 5.4) Consider the proof of the reverse direction (\Leftarrow) in the Theorem 5.3. An examination of the proof shows that the open calls constructed satisfy the more stringent requirements of 5.4. In fact, the proof of 5.3 shows that (5.3 \Rightarrow 5.4). The result follows since the forward direction of 5.3 shows that (5.1 \Rightarrow 5.3). \square

For full concurrent specifications and implementations, we suspect that Lemma 5.4 fails. (To get a sense of the issues, consider a specification that orders $a \rightarrow c$ and $b \rightarrow d$, and an implementation that executes $a \rightarrow d$ and $b \rightarrow c$.) In this paper, however, all of our results concern sequential specifications and operational implementations.

5.3 Third characterization: speculative flat combining

Our third characterization of QQC describes how QQC affects an arbitrary sequential data structure, using a *proxy* that generates QQC traces from an underlying sequential

implementation. The proxy is *sound*, in that every trace that it accepts is QQC, and *complete*, in that it generates every operational trace that is QQC with respect to the sequential data structure.

This characterization of QQC incorporates *speculation* into flat combining [7]. *Flat combining* is a technique for implementing concurrent data structures using sequential ones by introducing a mediator between the concurrent world and the sequential data structure. As for speculation, we push the obligation to predict the future into the underlying sequential object, which must conform to the following interface.

```
interface Object {
    method run(i:Invocation):Response;
    method predict():Invocation; }
```

The `run` method passes invocations to the underlying sequential structure and returns the appropriate response. The `predict` method is an oracle that guesses the invocations that are to come in the future. It is the use of `predict` that makes our code speculative.

Given an `Object o`, the proxy is defined as follows.

```
class QQCProxy<o:Object> {
    field called:ThreadSafeMultiMap<Invocation,Semaphore> = [];
    field returned:ThreadSafeMap <Semaphore, Response> = [];
    method run(i:Invocation):Response { //proxy for external access to o
        val m:Semaphore = [];
        called.add(i, m);
        m.wait();
        return returned.remove(m); }
    thread { //single thread to interact with o
        val received:MultiMap<Invocation,Semaphore> = [];
        val executed:MultiMap<Invocation,Response> = [];
        repeatedly choose {
            choice if called.nonEmpty() {
                received.add(called.removeAny());
                val i:Invocation = o.predict();
                val r:Response = o.run(i);
                executed.add(i, r); }
            choice if exists i in received.keys() intersect executed.keys() {
                val m:Semaphore = received.remove(i);
                val r:Response = executed.remove(i);
                returned.add(m, r);
                m.signal(); } } }
```

Communication between the implementation threads and the underlying `Object` is mediated by two maps. When a thread would like to interact with the `Object`, it creates a semaphore, registers the semaphore in `called` and waits on the semaphore. Upon awakening, the thread removes the result from `returned` and returns.

The `Object` is serviced by a single *proxy* thread which loops forever making one of two nondeterministic choices. The proxy keeps two private maps. Upon receiving an `called` invocation, the proxy moves the invocation from `called` to `received`. Rather than executing the `received` invocation, the proxy asks the oracle to predict an arbitrary

invocation `i` and executes that instead, placing the result in `executed`. Once a invocation is both `received` and `executed`, it may become `returned`.

At the beginning of this section, we noted that the stack execution $\{c^+ [^-]_a^+ (c^+)^+ \}^+$ is QQC with respect to the specification $(c^+)^+ [^-]_a^- \{c^+ \}^+$. How can such a trace possibly be generated? The execution of the proxy proceeds as follows. Upon receipt of $\{c^+\}$, the proxy executes (c^+) , storing the response $\}^+$. Upon receipt of $[^-]$, the proxy executes $[^-]$, storing the response $]_a^-$. At this point $[^-]_a^-$ can be returned. Upon receipt of (c^+) , the proxy executes $\{c^+\}$, storing the response $\}^+$. At this point both $(c^+)^+$ and $\{c^+\}^+$ can be returned.

Theorem 5.5. The concurrent proxy is sound for QQC with respect to the underlying Object. It is also complete for operational traces.

PROOF. For soundness, note that proxy maintains the invariant that the sizes of `received` and `executed` are equal, and therefore the number of returned calls can never exceed the number that has been received. In addition, the number of things added to `received` always exceeds the number added to `returned`.

For completeness, suppose that trace $s \sqsubseteq_{\text{qqc}} t$ and let the sequence of names in t be $(a_1^?, a_1^!, \dots, a_m^?, a_m^!)$. Consider any total order on the events of s that is consistent with the order already present in t . Let $(b_1^?, \dots, b_m^?)$ be the order on the call actions in this total order. When $b_i^?$ arrives, add $b_i^?$ to `received` and execute $a_i^?$, placing $a_i^!$ into `executed`. From Theorem 5.3 we know that whenever a response is required, there will be enough prior invocations so that the required response will be found in `executed`. \square

5.4 Compositionality

We now prove compositionality for QQC, following the proof for linearizability in Theorem 3.6. Below, we give some examples of the construction given in the proof, which is more complex than the one required for linearizability. Recall that (\div) denotes partial order difference.

Theorem 5.6. Let t_1 and t_2 be sequential traces.

Let s , s_1 and s_2 be operational traces such that $s_1 = s \div s_2$ and $s_2 = s \div s_1$.

For $i \in \{1, 2\}$, suppose that each $s_i \sqsubseteq_{\text{qqc}} t_i$.

Then there exists a sequential trace $t \in (t_1 \parallel t_2)$ such that $s \sqsubseteq_{\text{qqc}} t$.

PROOF. As in the proof of Theorem 3.6, assume $\text{ids}(t_1)$ and $\text{ids}(t_2)$ are disjoint, and let the sequence of names in t_1 be $(a_1^?, a_1^!, \dots, a_m^?, a_m^!)$ and sequence of name in t_2 be $(b_1^?, b_1^!, \dots, b_n^?, b_n^!)$. Applying Theorem 5.3 to the supposition $s_1 \sqsubseteq_{\text{in}} t_1$, we have that $j \leq |\{a_i^? \mid a_i^? \Rightarrow_s a_j^!\}|$, and similarly $\ell \leq |\{b_k^? \mid b_k^? \Rightarrow_s b_\ell^!\}|$. It suffices to construct an interleaving $t \in (t_1 \parallel t_2)$ such that whenever t contains a subsequence with names

$$a_j^?, a_j^!, b_k^?, b_k^!, b_{k+1}^?, b_{k+1}^!, \dots, b_{k+x}^?, b_{k+x}^!$$

then for every $k \leq \ell \leq k+x$, we have

$$\{a_i^? \mid a_i^? \Rightarrow_s a_j^!\} \subseteq \{a_i^? \mid a_i^? \Rightarrow_s b_\ell^!\}$$

and symmetrically for subsequences $b_k^?, b_k^!, a_j^?, a_j^!, a_{j+1}^?, a_{j+1}^!, \dots, a_{j+y}^?, a_{j+y}^!$. Given such a t , we know that $j + \ell \leq |\{a_i^? \mid a_i^? \Rightarrow_s b_\ell^!\} \cup \{b_k^? \mid b_k^? \Rightarrow_s b_\ell^!\}|$, as required.

We now demonstrate the existence of such a t . Define the set $\text{merge}(\vec{a}, \vec{b})$ as follows.


$$\begin{aligned} \text{merge}(\vec{a}, \varepsilon) &= \{\vec{a}\} & \text{merge}(\varepsilon, \vec{b}) &= \{\vec{b}\} \\ \text{merge}(\vec{a}a_j^? a_j^!, \vec{b}b_\ell^? b_\ell^!) &\ni \vec{c}b_\ell^? b_\ell^! & \text{if } \vec{c} \in \text{merge}(\vec{a}a_j^? a_j^!, \vec{b}) \\ && \text{and } \{a_i^? \mid a_i^? \Rightarrow_s a_j^!\} \subseteq \{a_i^? \mid a_i^? \Rightarrow_s b_\ell^!\} \\ \text{merge}(\vec{a}a_j^? a_j^!, \vec{b}b_\ell^? b_\ell^!) &\ni \vec{c}a_j^? a_j^! & \text{if } \vec{c} \in \text{merge}(\vec{a}, \vec{b}b_\ell^? b_\ell^!) \\ && \text{and } \{b_k^? \mid b_k^? \Rightarrow_s b_\ell^!\} \subseteq \{b_k^? \mid b_k^? \Rightarrow_s a_j^!\} \end{aligned}$$

To demonstrate the existence of an appropriate t , it suffices to show that $\text{merge}(a_1^? a_1^! \dots a_m^? a_m^!, b_1^? b_1^! \dots b_n^? b_n^!)$ is nonempty. By operationality, it must be the case that either (1) $a_j^? \Rightarrow_s b_\ell^!$, in which case $\{a_i^? \mid a_i^? \Rightarrow_s a_j^!\} \subseteq \{a_i^? \mid a_i^? \Rightarrow_s b_\ell^!\}$, (2) $b_\ell^? \Rightarrow_s a_j^!$, in which case $\{b_k^? \mid b_k^? \Rightarrow_s b_\ell^!\} \subseteq \{b_k^? \mid b_k^? \Rightarrow_s a_j^!\}$, or (3) $a_j^?$ and $b_\ell^!$ are unordered, in which case both conclusions hold. Therefore an appropriate t exists. \square

Example 5.7. We demonstrate the merge function defined in the proof above using the following traces.

$$\begin{aligned} t_1 &= \lceil^+ \rfloor_0^+ (\color{red}{\lceil^+})_1^+ \lceil^+ \rfloor_2^+ & t_2 &= \lceil^+ \rfloor_0^+ \lceil^+ \rfloor_1^+ \lceil^+ \rfloor_2^+ \\ s_1 &= \lceil^+ (\color{red}{\lceil^+})_1^+ \lceil^+ \rfloor_0^+ \rfloor_2^+ & s_2 &= \lceil^+ \lceil^+ \rfloor_1^+ \lceil^+ \rfloor_0^+ \rfloor_2^+ \\ s &= \lceil^+ \lceil^+ \rfloor_1^+ \lceil^+ \rfloor_0^+ (\color{red}{\lceil^+})_1^+ \lceil^+ \rfloor_0^+ \rfloor_2^+ \end{aligned}$$

In the graph below, we draw an edge from a_j to b_ℓ if $\{a_i^? \mid a_i^? \Rightarrow_s a_j^!\} \subseteq \{a_i^? \mid a_i^? \Rightarrow_s b_\ell^!\}$, indicating that b_ℓ may come after a_j . Edges from b_ℓ to a_j are similar. When an edge is bidirectional, we use a dashed line.




The following traces are derived from the merge algorithm.

$$\begin{aligned} &\lceil^+ \rfloor_0^+ \lceil^+ \rfloor_1^+ \lceil^+ \rfloor_2^+ \lceil^+ \rfloor_0^+ (\color{red}{\lceil^+})_1^+ \lceil^+ \rfloor_2^+ \\ &\lceil^+ \rfloor_0^+ \lceil^+ \rfloor_1^+ \lceil^+ \rfloor_0^+ \lceil^+ \rfloor_2^+ (\color{red}{\lceil^+})_1^+ \lceil^+ \rfloor_2^+ \\ &\lceil^+ \rfloor_0^+ \lceil^+ \rfloor_1^+ \lceil^+ \rfloor_0^+ (\color{red}{\lceil^+})_1^+ \lceil^+ \rfloor_2^+ \lceil^+ \rfloor_2^+ \\ &\lceil^+ \rfloor_0^+ \lceil^+ \rfloor_1^+ \lceil^+ \rfloor_0^+ (\color{red}{\lceil^+})_1^+ \lceil^+ \rfloor_2^+ \lceil^+ \rfloor_2^+ \end{aligned}$$

Suppose instead that we have the following s .

$$s = \lceil^+ \lceil^+ \rfloor_1^+ \lceil^+ \rfloor_0^+ (\color{red}{\lceil^+})_1^+ \lceil^+ \rfloor_2^+ \lceil^+ \rfloor_0^+ \rfloor_2^+$$

Then the graph and resulting traces are as follows.



In general, if one were to include the linear order from the specification (eg, from $[^+]_0^+$ to $(^+)_1^+$), the resulting graph might be cyclic, even if the dotted edges were removed. \square

6 Stack example

We show that, under reasonable assumptions, our N -Stack is QQC. We extend this argument to the elimination-tree stacks of [13].

In proving that executions of our N -Stack are QQC, the key step is to generate the corresponding specification trace. To do so, we consider the following instrumentation.

```

1 class Stack<N:Int> {
2   field b:[0..N-1] = 0;           // 1 balancer
3   field s:Stack[] = [[], [], ..., []]; // N stacks of values
4   field e:[0..N-1] = 0;           // 1 emitter
5   field q:Queue[] = [[], [], ..., []]; // N queues of actions
6   method push(x:Object):Unit {
7     val i:[0..N-1];
8     atomic {i=b; b++;}
9     atomic {val v=s[i].push(x); q[i].add("push" x); emit(); return v;} }
10  method pop():Object {
11    val i:[0..N-1];
12    atomic {i=b-1; b--;}
13    atomic {val v=s[i].pop(); q[i].add("pop" v); emit(); return v;} }
14  method emit():Unit {
15    while (q[e].first()=~"push" || q[e-1].first()=~"pop") {
16      if (q[e].first()=~"push") {print (q[e].remove()); e++;}
17      if (q[e-1].first()=~"pop") {print (q[e-1].remove()); e--;} } }

```

The state of the machine includes the values of the balancer b and stacks s . It also includes queues q to store the actions that have been executed on each stack and a *emitter* e , with the same range as b , which indicates the queue that should produce the next specification action. The emitter prints any completed pushes from $s[e]$ and any completed pops from $s[e-1]$. When the emitter prints a push, it removes it from the queue and increments e ; when it prints a pop, it removes it from the queue and decrements e . Emitter actions take place as soon as possible, and the emitter continues until it has nothing left to do.

Atomic blocks can only execute concurrently if they do not touch the same shared state. For the code in the introduction, this imposes an order between all executions of the first atomic (lines 8 and 12), since they touch the shared variable b ; order is only imposed between executions of the second atomic that update the same stack. The presence of `emit` indicates also imposes an order between all executions of the second atomic (lines 9 and 13), since `emit` touches the shared variable e . This total order on calls to `emit` ensures that the printed trace is indeed a stack trace, as we argue below.

Definition 6.1. Let a be a call to push or pop. Then $\text{time}_1(a)$ is the time of the execution of the first atomic statement in the N -Stack, and $\text{time}_2(a)$ is the time of the execution of the second atomic. A *linearized trace* of an N -Stack is one in which the invocations are ordered consistently with time_1 and the responses are ordered consistently with time_2 . \square

For example, from the linearization $(\overset{+}{c} \overset{+}{} \overset{+}{>} \overset{+}{[a]} \overset{+}{J})^+$ we know $\text{time}_1(\overset{+}{c}) < \text{time}_1(\overset{+}{[a]})$ and $\text{time}_2(\overset{+}{J}) < \text{time}_2(\overset{+}{>})$. Such a linearized trace is distinct from other linearizations of the same trace, such as $(\overset{+}{b} \overset{+}{[a]} \overset{+}{J})^+$, $(\overset{+}{[a]} \overset{+}{} \overset{+}{>} \overset{+}{J})^+$ and $(\overset{+}{[a]} \overset{+}{} \overset{+}{>})^+ \overset{+}{J}$.

The response order in the linearized trace is particularly significant. For example, the linearization $(\overset{+}{b} \overset{+}{[a]} \overset{+}{J})^+ \overset{-}{[a]} \overset{-}{} \overset{-}{>}$ cannot result from the execution of a 1-Stack. In this case a is pushed before b and therefore the pop of a cannot be ordered before the pop of b .

Example 6.2. Consider the following linearized trace of a 2-Stack.

$$(\overset{+}{c} \overset{+}{} \overset{+}{>} \overset{+}{[a]} \overset{+}{J})^+ \overset{-}{<c>} \overset{-}{>} \overset{-}{[a]}$$

Execution proceeds as follows. We show the atomic that is being executed above the arrow. Arrows without labels are executed within `emit`, atomically with the prior label. On the right-hand side, we show any emitted actions, followed by the resulting state. The initial state of the machine is $\langle b = 0, e = 0, s = [[], []], q = [[], []] \rangle$.

$$\begin{array}{ll}
 & \langle b = 0, e = 0, s = [[], []], \quad q = [[], []] \rangle \\
 \overset{+}{c} \rightarrow & \langle b = 1, e = 0, s = [[], []], \quad q = [[], []] \rangle \\
 \overset{+}{} \rightarrow & \langle b = 0, e = 0, s = [[], []], \quad q = [[], []] \rangle \\
 \overset{+}{>} \rightarrow & \langle b = 0, e = 0, s = [[], [b]], \quad q = [[], [\overset{+}{} \overset{+}{>}]] \rangle \\
 \overset{+}{[a]} \rightarrow & \langle b = 1, e = 0, s = [[], [b]], \quad q = [[], [\overset{+}{} \overset{+}{>}]] \rangle \\
 \overset{+}{J} \rightarrow & \langle b = 1, e = 0, s = [[a], [b]], \quad q = [[[\overset{+}{[a]} \overset{+}{J}], [\overset{+}{} \overset{+}{>}]] \rangle \\
 \rightarrow \overset{+}{[a]} \overset{+}{J} & \langle b = 1, e = 1, s = [[a], [b]], \quad q = [[], [\overset{+}{} \overset{+}{>}]] \rangle \\
 \rightarrow \overset{+}{} \overset{+}{>} & \langle b = 1, e = 0, s = [[a], [b]], \quad q = [[], []] \rangle \\
 \overset{+}{>} \rightarrow & \langle b = 1, e = 0, s = [[ca], [b]], \quad q = [[(\overset{+}{c})^+], []] \rangle \\
 \rightarrow \overset{+}{<c>} \overset{+}{>} & \langle b = 1, e = 1, s = [[ca], [b]], \quad q = [[], []] \rangle \\
 \overset{-}{<c>} \rightarrow & \langle b = 0, e = 1, s = [[ca], [b]], \quad q = [[], []] \rangle \\
 \overset{-}{>} \rightarrow & \langle b = 0, e = 1, s = [[a], [b]], \quad q = [[(\overset{-}{c})^-], []] \rangle \\
 \rightarrow \overset{-}{<c>} \overset{-}{>} & \langle b = 0, e = 0, s = [[a], [b]], \quad q = [[], []] \rangle \\
 \overset{-}{} \rightarrow & \langle b = 1, e = 0, s = [[a], [b]], \quad q = [[], []] \rangle \\
 \overset{-}{>} \rightarrow & \langle b = 1, e = 0, s = [[a], []], \quad q = [[], [\overset{-}{} \overset{-}{>}]] \rangle \\
 \rightarrow \overset{-}{} \overset{-}{>} & \langle b = 1, e = 1, s = [[a], []], \quad q = [[], []] \rangle \\
 \overset{-}{[a]} \rightarrow & \langle b = 0, e = 1, s = [[a], []], \quad q = [[], []] \rangle \\
 \overset{-}{J} \rightarrow & \langle b = 0, e = 1, s = [[], []], \quad q = [[[\overset{-}{[a]} \overset{-}{J}], []] \rangle \\
 \rightarrow \overset{-}{<[a]>} \overset{-}{>} & \langle b = 0, e = 0, s = [[], []], \quad q = [[], []] \rangle
 \end{array}$$

□

Example 6.3. Consider the following execution of the instrumented counter.

$$\begin{array}{ll}
 & \langle b = 0, e = 0, s = [[], []], \quad q = [[], []] \rangle \\
 \overset{+}{[a]} \rightarrow & \langle b = 1, e = 0, s = [[], []], \quad q = [[], []] \rangle \\
 \overset{+}{<a>} \rightarrow & \langle b = 0, e = 0, s = [[], []], \quad q = [[], [(\overset{+}{a})^+]] \rangle
 \end{array}$$

$\xrightarrow{y^+}$	$\langle \mathbf{b} = 0, \mathbf{e} = 0, \mathbf{s} = [()], [a], \mathbf{q} = [(), [(c_a^+)^+]] \rangle$
$\xrightarrow{c^-}$	$\langle \mathbf{b} = 1, \mathbf{e} = 0, \mathbf{s} = [()], [a], \mathbf{q} = [(), [(c_a^+)^+]] \rangle$
$\xrightarrow{y_a^-}$	$\langle \mathbf{b} = 1, \mathbf{e} = 0, \mathbf{s} = [()], [()], \mathbf{q} = [(), [(c_a^+)^+ (c^-)_a^-]] \rangle$
$\xrightarrow{c_a^+}$	$\langle \mathbf{b} = 0, \mathbf{e} = 0, \mathbf{s} = [()], [()], \mathbf{q} = [(), [(c_a^+)^+ (c^-)_a^-]] \rangle$
$\xrightarrow{y^+}$	$\langle \mathbf{b} = 0, \mathbf{e} = 0, \mathbf{s} = [()], [b], \mathbf{q} = [(), [(c_a^+)^+ (c^-)_a^- (c_b^+)^+]] \rangle$
$\xrightarrow{c_b^+}$	$\langle \mathbf{b} = 1, \mathbf{e} = 0, \mathbf{s} = [()], [b], \mathbf{q} = [(), [(c_a^+)^+ (c^-)_a^- (c_b^+)^+]] \rangle$
$\xrightarrow{c_a^+}$	$\langle \mathbf{b} = 0, \mathbf{e} = 0, \mathbf{s} = [()], [b], \mathbf{q} = [(), [(c_a^+)^+ (c^-)_a^- (c_b^+)^+]] \rangle$
$\xrightarrow{y^+}$	$\langle \mathbf{b} = 0, \mathbf{e} = 0, \mathbf{s} = [()], [bc], \mathbf{q} = [(), [(c_a^+)^+ (c^-)_a^- (c_b^+)^+ (c_c^+)^+]] \rangle$
$\xrightarrow{J^+}$	$\langle \mathbf{b} = 0, \mathbf{e} = 0, \mathbf{s} = [0], [bc], \mathbf{q} = [[c_0^+ J^+], [(c_a^+)^+ (c^-)_a^- (c_b^+)^+ (c_c^+)^+]] \rangle$
$\rightarrow [c_0^+ J^+]$	$\langle \mathbf{b} = 0, \mathbf{e} = 1, \mathbf{s} = [0], [bc], \mathbf{q} = [(), [(c_a^+)^+ (c^-)_a^- (c_b^+)^+ (c_c^+)^+]] \rangle$
$\rightarrow (c_a^+)^+$	$\langle \mathbf{b} = 0, \mathbf{e} = 0, \mathbf{s} = [0], [bc], \mathbf{q} = [(), [(c^-)_a^- (c_b^+)^+ (c_c^+)^+]] \rangle$
$\rightarrow (c^-)_a^-$	$\langle \mathbf{b} = 0, \mathbf{e} = 1, \mathbf{s} = [0], [bc], \mathbf{q} = [(), [(c_b^+)^+ (c_c^+)^+]] \rangle$
$\rightarrow (c_b^+)^+$	$\langle \mathbf{b} = 0, \mathbf{e} = 0, \mathbf{s} = [0], [bc], \mathbf{q} = [(), [(c_c^+)^+]] \rangle$
$\xrightarrow{J^+}$	$\langle \mathbf{b} = 0, \mathbf{e} = 0, \mathbf{s} = [01], [bc], \mathbf{q} = [[c_1^+ J^+], [(c_c^+)^+]] \rangle$
$\rightarrow [c_1^+ J^+]$	$\langle \mathbf{b} = 0, \mathbf{e} = 1, \mathbf{s} = [01], [bc], \mathbf{q} = [(), [(c_c^+)^+]] \rangle$
$\rightarrow (c_c^+)^+$	$\langle \mathbf{b} = 0, \mathbf{e} = 0, \mathbf{s} = [01], [bc], \mathbf{q} = [(), [()]] \rangle$

This produces the following linearized trace s and specification t .

$$s = [c_0^+ (c_a^+)^+ (c^-)_a^- (c_b^+)^+ [c_1^+ (c_c^+)^+ J^+ J^+]$$

$$t = [c_0^+ J^+ (c_a^+)^+ (c^-)_a^- (c_b^+)^+ [c_1^+ J^+ (c_c^+)^+]$$

After the push of c returns, we have $\mathbf{q}[1] = [(c_a^+)^+ (c^-)_a^- (c_b^+)^+ (c_c^+)^+]$. When the first J^+ occurs, the first three actions in the $\mathbf{q}[1]$ must be emitted. \square

Lemma 6.4. Given an instrumented execution of an N -Stack, the linearized trace of the execution is QQC with the emitted specification.

PROOF SKETCH. Let us refer to a sequence like $(c_a^+)^+ (c^-)_a^- (c_b^+)^+$ as a *chain*. A chain is a sequence of calls that can be emitted from a single queue without any intervening change to \mathbf{e} . By Theorem 5.3 suffices to show that after the execution of each atomic, the number of chains is bounded by the number of open calls. This follows by induction on the length of the instrumented execution. \square

In light of Lemma 6.4, to show that the N -Stack is QQC, it suffices to show that the emitted specification is indeed a stack specification. Unfortunately, as observed in [13], this fails to hold.

Example 6.5. As discussed in Example 1.4, the linearized trace $[c_a^+ J^+ (c_b^+)^+ [c_c^- [c_a^-]^-] J^+]$ generates the specification $[c_a^+ J^+ (c_b^+)^+ [c_a^- [c_c^+ J^+]]^+]$. However, this specification is not a stack trace. With some number of initial pushes, this execution is still possible: The linearized trace $[c_x^+ J^+ (c_y^+)^+ [c_a^+ J^+ (c_b^+)^+ [c_c^- [c_a^-]^-] J^+]$ generates the specification $[c_x^+ J^+ (c_y^+)^+ [c_a^+ J^+ (c_b^+)^+ [c_a^- [c_c^+ J^+]]^+]$. \square

This problematic execution occurs because a push and pop are racing at the first stack, yet the pop retrieves a prior value: the pop has *overtaken* the push. We must disallow such executions. It is not sufficient to require only that pop operations block on an empty stack.

Definition 6.6. An execution is *properly-popped* if for every push a and pop b that are assigned the same stack $s[i]$,

$$\text{time}_1(a) < \text{time}_1(b) \text{ implies } \text{time}_2(a) < \text{time}_2(b). \quad \square$$

Lemma 6.7. If an execution of the instrumented N -Stack is properly-popped, then its trace is a stack trace.

PROOF SKETCH. It is sufficient to note that the execution of the emitter follows the same pattern as the uninstrumented N -Stack on a sequential execution. (This is only true with proper popping.) The result follows since, as shown in [13], the sequential execution of the N -Stack does simulate a stack. \square


Theorem 6.8. Any properly-popped execution of an N -Stack is *QQC*.

PROOF. By Lemmas 6.4 and 6.7. \square

We have shown that for properly-popped executions (where a pop may not ignore a concurrent push on the same stack) the N -Stack is *QQC*. As noted in the introduction, we know of no analogous condition for increment/decrement counters.

In [13], Shavit and Touitou show that in a quiescent state, their elimination-tree stack reaches a state consistent with a stack. We now consider the relation between our N -Stacks and these elimination-tree stacks.

Example 6.9. A depth-2 elimination-tree stack can be implemented using three atomic booleans—top (t), left (l) and right (r)—and 4 linearizable stacks with addresses 00, 01, 10 and 11.



The *address* of a stack in a depth- d elimination tree is a sequence of d booleans, indicating the value of the boolean at each level, going down a branch of the tree. Both push and pop toggle the booleans as they go down the tree, using an atomic read and update. If $t = 0$, then push sets $t = 1$ and goes left. If $t = 0$, then pop sets $t = 1$ and goes right. The methods follow this same pattern down the tree until they reach the

bottom-level stack, at which point they perform the operation. Initially all booleans are set to 0. For example, one uninstrumented execution proceeds as follows.

$$\begin{aligned}
& \langle \mathbf{t} = 0, \langle \mathbf{l} = 0, \mathbf{s}_1 = [[], []] \rangle, \langle \mathbf{r} = 0, \mathbf{s}_r = [[], []] \rangle \rangle \\
& \xrightarrow{c_e^+} \langle \mathbf{t} = 1, \langle \mathbf{l} = 0, \mathbf{s}_1 = [[], []] \rangle, \langle \mathbf{r} = 0, \mathbf{s}_r = [[], []] \rangle \rangle \\
& \xrightarrow{[b]^+} \langle \mathbf{t} = 0, \langle \mathbf{l} = 0, \mathbf{s}_1 = [[], []] \rangle, \langle \mathbf{r} = 1, \mathbf{s}_r = [[b], []] \rangle \rangle \\
& \xrightarrow{[a]^+} \langle \mathbf{t} = 1, \langle \mathbf{l} = 1, \mathbf{s}_1 = [[a], []] \rangle, \langle \mathbf{r} = 1, \mathbf{s}_r = [[b], []] \rangle \rangle \\
& \xrightarrow{[d]^+} \langle \mathbf{t} = 0, \langle \mathbf{l} = 1, \mathbf{s}_1 = [[a], []] \rangle, \langle \mathbf{r} = 0, \mathbf{s}_r = [[b], [d]] \rangle \rangle \\
& \xrightarrow{[c]^+} \langle \mathbf{t} = 1, \langle \mathbf{l} = 0, \mathbf{s}_1 = [[a], [c]] \rangle, \langle \mathbf{r} = 0, \mathbf{s}_r = [[b], [d]] \rangle \rangle \\
& \xrightarrow{e} \langle \mathbf{t} = 1, \langle \mathbf{l} = 1, \mathbf{s}_1 = [[ea], [c]] \rangle, \langle \mathbf{r} = 0, \mathbf{s}_r = [[b], [d]] \rangle \rangle \\
& \xrightarrow{[e]^-} \langle \mathbf{t} = 0, \langle \mathbf{l} = 0, \mathbf{s}_1 = [[a], [c]] \rangle, \langle \mathbf{r} = 0, \mathbf{s}_r = [[b], [d]] \rangle \rangle \\
& \xrightarrow{[d]^-} \langle \mathbf{t} = 1, \langle \mathbf{l} = 0, \mathbf{s}_1 = [[a], [c]] \rangle, \langle \mathbf{r} = 1, \mathbf{s}_r = [[b], []] \rangle \rangle \\
& \xrightarrow{[c]^-} \langle \mathbf{t} = 0, \langle \mathbf{l} = 1, \mathbf{s}_1 = [[a], []] \rangle, \langle \mathbf{r} = 1, \mathbf{s}_r = [[b], []] \rangle \rangle \\
& \xrightarrow{[b]^-} \langle \mathbf{t} = 1, \langle \mathbf{l} = 1, \mathbf{s}_1 = [[a], []] \rangle, \langle \mathbf{r} = 0, \mathbf{s}_r = [[], []] \rangle \rangle \\
& \xrightarrow{[a]^-} \langle \mathbf{t} = 0, \langle \mathbf{l} = 0, \mathbf{s}_1 = [[], []] \rangle, \langle \mathbf{r} = 0, \mathbf{s}_r = [[], []] \rangle \rangle
\end{aligned}$$

This gives the trace $(c_e^+ [b]^+ [a]^+ [d]^+ [c]^+)^+ \{ [e]^- [d]^- [c]^- [b]^- [a]^- \}$ which is QQC with respect to $[a]^+ [b]^+ [c]^+ [d]^+ (c_e^+)^+ \{ [e]^- [d]^- [c]^- [b]^- [a]^- \}$. Our 4-Stack does not generate this execution trace; however, our 2-Stack does. In general, our N^d -Stack has strictly fewer behaviors than the N -branching elimination-tree stack of depth d . We leave open the question of whether a N -branching elimination-tree stack of depth d has behaviors that not possible for an N -Stack. \square

The instrumented execution of a N -branching elimination-tree stack of depth $d > 1$ can be defined using the execution of elimination-tree stacks of depth $d - 1$, using the same strategy as our N -Stack. While the balancer's behavior is more general in the composed system, the emitter's is not: The emitter code is entirely sequentialized, therefore a 2-nested N -branching emitter has the same behavior as a flat N^2 -branching emitter.

Theorem 6.10. Any properly-popped execution of a N -branching elimination-tree stack of depth d is QQC.

PROOF SKETCH. Following the strategy in Theorem 6.8, we need only prove the corresponding lemmas. In each case, the proof proceeds by induction on d . In each case the basis is the same: a depth 1 elimination tree stack is simply an N -Stack.

The analogue of Lemma 6.4 follows, as before, by induction on the length of the instrumented execution. An open call at depth d may initiate a new chain, but only in one stack of depth $d - 1$.

For the analogue of 6.7 it suffices to observe that the emitter's behavior is the same if levels $d > 1$ and $d - 1$ are flattened into a single level of size N^2 . This follows from the atomicity of the emitter. \square

7 Conclusions

Quantitative quiescent consistency (QQC) is a correctness criterion for concurrent data structures that relaxes linearizability and refines quiescent consistency. To the best of our knowledge, it is the first such criterion to be proposed.

To show that QQC is a robust concept, we have provided three alternate characterizations: (1) in the style of linearizability, (2) counting the number of calls before a return, and (3) using speculative flat combining. We have also proven compositionality (in the style of Herlihy and Wing [9]) and the correctness of data structures defined by Aspnes, Herlihy, and Shavit [3] and Shavit and Touitou [13].

In order to establish the correctness of the elimination-tree stack of [13], we had to restrict attention to traces in which no pop *overtakes* a push on the same stack. A related constraint appears in a footnote of [11]: “To keep things simple, pop operations should block until a matching push appears.” This, however, is not strong enough to guarantee quiescent consistency as we have defined it. Our analysis provides a full account: The stack is QQC with the no-overtaking requirement and only weakly quiescently consistent without it.

As witnessed by the stack trace $\{c^+ [^-]_a^- (a^+)^+ \}^+$ in section 5, QQC does not enforce causality. We have chosen not to treat causality in this paper in order to present the basic idea of QQC as clearly as possible. Causality is an orthogonal concept: One can enforce causality in *addition* to QQC. We have not done so here because causality requires a notion of derivation over the underlying values, in which one distinguishes public values (eg, base types) from secret values (eg, pointers or nonces). For stacks, derivation requires that the value returned by a pop must be either public or a previously pushed secret.

Linearizability has proven to be a valuable foundation for program verification techniques. It remains to be seen if QQC can be of use in this regard.

Linearizability is, at its core, *linear*. We have defined QQC in terms of general partial orders, and yet the results reported here are stated in terms of sequential specifications. Partly we have done this so that we can relate the definition of QQC to the vast amount of existing work on linearizability. However, the general case is interesting.

Acknowledgements

Gustavo Petri participated in the early discussions which motivated this work. Alan Jeffrey and Corin Pitcher made useful comments on previous candidate definitions for QQC.

References

- [1] Y. Afek, G. Korland, and E. Yanovsky, “Quasi-linearizability: relaxed consistency for improved concurrency,” in *OPODIS*, ser. LNCS, vol. 6490, Springer, 2010, pp. 395–410.
- [2] W. Aiello et al., “Supporting increment and decrement operations in balancing networks,” *Chicago J. Theor. Comput. Sci.*, 2000.

- [3] J. Aspnes, M. Herlihy, and N. Shavit, “Counting networks,” *J. ACM*, vol. 41, no. 5, pp. 1020–1048, 1994.
- [4] C. Busch and M. Mavronicolas, “The strength of counting networks (abstract),” in *PODC*, J. E. Burns and Y. Moses, Eds., ACM, 1996, p. 311.
- [5] C. Dwork, M. Herlihy, and O. Waarts, “Contention in shared memory algorithms,” *J. ACM*, vol. 44, no. 6, pp. 779–805, 1997.
- [6] A. Haas et al., “Distributed queues in shared memory: multicore performance and scalability through quantitative relaxation,” in *Conf. Computing Frontiers*, ACM, 2013, p. 17.
- [7] D. Hendler, I. Incze, N. Shavit, and M. Tzafrir, “Flat combining and the synchronization-parallelism tradeoff,” in *SPAA*, 2010, pp. 355–364.
- [8] T. A. Henzinger, C. M. Kirsch, H. Payer, A. Sezgin, and A. Sokolova, “Quantitative relaxation of concurrent data structures,” in *POPL*, 2013, pp. 317–328.
- [9] M. Herlihy and J. M. Wing, “Linearizability: a correctness condition for concurrent objects,” *ACM TOPLAS*, vol. 12, no. 3, pp. 463–492, 1990.
- [10] M. Herlihy and N. Shavit, *The Art of Multiprocessor Programming*. Morgan Kaufmann, 2008.
- [11] N. Shavit, “Data structures in the multicore age,” *Commun. ACM*, vol. 54, no. 3, pp. 76–84, Mar. 2011.
- [12] N. Shavit and D. Touitou, “Elimination trees and the construction of pools and stacks (preliminary version),” in *SPAA*, 1995, pp. 54–63.
- [13] —, “Elimination trees and the construction of pools and stacks,” *Theory Comput. Syst.*, vol. 30, no. 6, pp. 645–670, 1997.