

# Enabling Modular Verification with Abstract Interference Specifications for a Concurrent Queue

Alan Weide<sup>1</sup>, Paolo A. G. Sivilotti<sup>1</sup>, and Murali Sitaraman<sup>2</sup>

<sup>1</sup> The Ohio State University, Columbus OH 43221, USA  
weide.3@osu.edu, paolo@ccse.ohio-state.edu

<sup>2</sup> Clemson University, Clemson SC 29634, USA  
murali@clemson.edu

**Abstract.** When concurrent threads of execution do not modify shared data, their parallel execution is trivially equivalent to their sequential execution. For many imperative programming languages, however, the modular verification of this independence is often frustrated by (i) the possibility of aliasing between variables mentioned in different threads, and (ii) the lack of abstraction in the description of read/write effects of operations on shared data structures. We describe a specification and verification framework in which abstract specifications of functional behavior are augmented with abstract interference effects that permit verification of client code with concurrent calls to operations of a data abstraction. To illustrate the approach, we present a classic concurrent data abstraction: the bounded queue. Three different implementations are described, each with different degrees of entanglement and hence different degrees of possible synchronization-free concurrency.

## 1 Introduction

Parallel programming is important for both large-scale high performance systems and, increasingly, small-scale multi-core commodity software. Programming with multiple threads, however, is error-prone. Furthermore, when errors are made, they can be difficult to debug and correct because parallel programs are often nondeterministic. Non-trivial parallel programs designed with software engineering consideration will be invariably composed from reusable components, often ones that encapsulate data abstractions.

Given this context, we propose a specification and verification framework to guarantee entanglement-free execution of concurrent code that invokes operations on data abstractions. Guaranteeing, simultaneously, modularity of the verification process and the independence of concurrent threads is complicated by two key problems. The first of these is the possibility of aliasing between objects involved in different threads. The second problem concerns guaranteeing safe parallel execution of data abstraction operations on an object without violating abstraction.

At the core of a solution to the aliasing problem is a notion of *clean* operation calls whereby effects of calls are restricted to objects that are explicit parameters or to global objects that are explicitly specified as affected. Under this notion, regardless of the level of granularity, syntactically independent operation calls are always safe to parallelize. While both the problem and the solution are of interest, this paper focuses only on a solution to the second problem.

To illustrate the ideas, the paper presents a bounded queue data abstraction and outlines three different implementations that vary in their potential for parallelism among different queue operations. The data abstraction specification is typical, except that it is designed to avoid unintended aliasing. To capture the parallel potential in a class of implementations we augment the data abstraction specification with an interference specification that introduces additional modeling details to facilitate guarantees of safe execution of concurrent client code. The second-level specification is typically still quite abstract and is devoid of concrete implementation details. The novelty of the proposed solution is that it modularizes the verification problem along abstraction boundaries. Specifically, verification of implementation code with respect to both its data abstract and interference specification is done once in the lifetime of the implementation. Verification of client code relies strictly on the specifications.

This paper is strictly work in progress. We outline, for example, the specification and verification framework, but do not include formal proof rules. The rest of the paper is organized as follows. Section 2 summarizes the most related work. Section 3 describes the central example and alternative implementations. Section 4 describes the solution. It begins with a presentation of the interference specification that forms the basis for the subsequent discussion on verification. The last section summarizes and gives directions for further research.

## 2 Related Work

The summary here is meant to be illustrative of the type of related work, not exhaustive.

Classical solutions to the interference problem (e.g., [3]) would involve defining and using locks, but neither the solutions nor the proofs of absence of interference here involve abstraction or specification. Lock-free solutions built using atomic read-write-modify primitives (e.g., compare-and-swap) allow finer granularity of parallelism, but the proofs of serializability in that context are often not modular and do not involve complex properties.

The objective of modular verification is widely shared. The work in [1], for example, involves specifying interference points. For data abstractions, the interference points would be set at the operation level, meaning two operations may not execute concurrently on an object, even if they are disentangled at a “fine-grain” level. The work by Rodriguez, et al [7] to extend JML for concurrent code makes it possible to specify methods to be atomic through locking and other properties. Using JML\* and a notion of dynamic frames, the work in [6] address safe concurrent execution in the context of more general solutions to

address aliasing and sharing for automated verification. The work in [9] makes it possible to specify memory locations that fall within the realm of an object’s lock. Chalice allows specification of various types of permissions and includes a notion of permission transfer [5]. Using them, it is possible to estimate an upper bound on the location sets that may be affected by a thread in Chalice.

### 3 A Bounded Queue Data Abstraction

#### 3.1 RESOLVE Background

RESOLVE[8] is an imperative, object-based programming and specification framework designed to support modular verification of sequential code. *Contracts* contain functional specifications and invariants in terms of abstract state. Abstract state is given in terms of mathematical types, such as sets or strings.<sup>1</sup> *Realizations* provide executable implementations as well as correspondence information connecting concrete and abstract state. The fundamental data movement operation is swap ( $:=$ ), a constant-time operator that avoids introducing aliasing while also avoiding deep or shallow copying [2].

In addition to pre- and post-condition based specifications, operation signatures in contracts include *parameter modes*, whereby the modification frame is defined. For example, the value of a restores-mode parameter is the same at the end of the operation as it was the beginning. In the realm of concurrency, restore-mode alone is not sufficient to ensure noninterference since it does not preclude the temporary modification of a parameter during the execution of an operation. Other parameter modes include clears (changed to be an initial value), replaces (can change, incoming value is irrelevant), and updates (can change, incoming value may be relevant).

#### 3.2 Abstract Specification

The BoundedQueueTemplate concept models a queue as a mathematical string of items. This concept defines queue operations including Enqueue, Dequeue, SwapFirstEntry, Length, and RemCapacity. The operations have been designed and specified to avoid aliasing that arises when queues contain non-trivial objects [2] and to facilitate clean semantics [4].

The operations in the contract are given in the listing below. For Enqueue, the requires clause says that there must be space in the queue for the new element ( $|q| < \text{MAX\_LENGTH}$ ). The ensures clause says that the outgoing value of  $q$  is the string concatenation of the incoming value of  $q$  (i.e.,  $\#q$ ) and the string consisting of a single item, the old value of  $e$ . Less formally, Enqueue puts  $e$  at the end of the queue. The parameter mode for  $e$  defines its outgoing value: an initial value for its type.

<sup>1</sup> A string is a sequence of values such as  $\langle 1, 2, 1, 3 \rangle$ . The string concatenation operator is  $\circ$ .

```

operation Enqueue (clears e: Item, updates q: Queue)
  requires |q| < MAX_LENGTH
  ensures q = #q o <#e>

operation Dequeue (replaces r: Item, updates q: Queue)
  requires q /= empty_string
  ensures #q = <r> o q

operation SwapFirstEntry (updates e: Item, updates q: Queue)
  requires q /= empty_string
  ensures
    <e> = substring(#q, 0, 1) and
    q = <#e> o substring(#q, 1, |#q|)

operation Length (restores q: Queue) : Integer
  ensures Length = |q|

operation RemCapacity (restores q: Queue) : Integer
  ensures RemCapacity = MAX_LENGTH - |q|

```

**Listing 1.1.** Contracts for Queue Operations

The requires clause for Dequeue says that q must not be empty. The ensures clause says that the concatenation of the resulting element r and outgoing value of q is the original value of q.

The SwapFirstEntry operation makes it possible to retrieve or update the first entry, without introducing aliasing.

The functions Length and RemCapacity behave as expected: Length returns an integer equal to the number of elements in the queue, and RemCapacity returns an integer equal to the number of free slots left in the queue before it becomes full. Neither modifies the queue.

### 3.3 Alternative Implementations

We have developed three alternative implementations of the bounded queue specified above, each with different parallelization opportunities. All three are based on a circular array. In the first two implementations, the length of the underlying array is equal to the maximum length of the queue, MAX\_LENGTH, while in the third the length of the array is one greater.

The first implementation has two Integer fields, front and length, where front is the index of the first element of the queue and length is the number of elements in the queue. This implementation cannot handle concurrent calls to Enqueue and Dequeue without synchronization because both of those calls must necessarily write to length. A client can, however, make concurrent calls to SwapFirstEntry and Enqueue when the precondition for both methods is met before

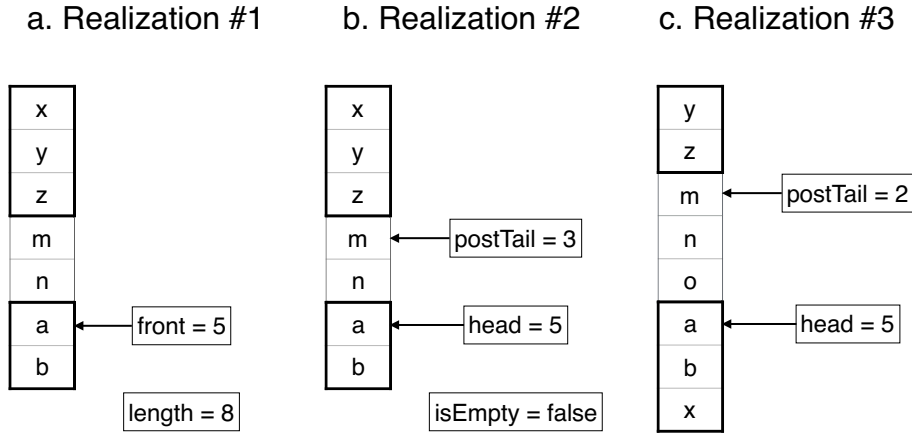


Fig. 1. Three alternatives for implementing a bounded queue on a circular array.

the parallel block (that is, if  $0 < |q|$  and  $|q| < \text{MAX\_LENGTH}$ ). These two methods may be executed in parallel because `SwapFirstEntry` touches only the head of the queue and does not modify `length`, while `Enqueue` will write `length` and touch the end of the queue (which we know is different from the head of the queue because there was already an element in the queue before `Enqueue` was called). An empty queue in this implementation has `length = 0` and  $0 \leq \text{front} < \text{MAX\_LENGTH}$ , and a full queue has `length = MAX\_LENGTH` and  $0 \leq \text{front} < \text{MAX\_LENGTH}$ .

The second implementation also has two Integer fields: `head` and `postTail`, and an additional Boolean field `isEmpty`. While `head` is the index of the array at which the first element of the queue is located, `postTail` is the index of the first element of the array after the last element of the queue. The boolean `isEmpty` is necessary to distinguish between a full queue and an empty queue since in both cases, `head = postTail`. As in implementation 1, a client can concurrently call `Enqueue` and `SwapFirstEntry` as long as both preconditions are satisfied. However, because the length of the queue is computed from the `head` and `postTail` fields (and not another variable written by both `Enqueue` and `Dequeue`), we can also concurrently call `Enqueue` and `Dequeue`, but only in a more limited set of circumstances than is described by their respective preconditions: the queue must have at least 2 entries in it and there must be at least 2 “free” slots in the array. This restriction is important because both `Enqueue` and `Dequeue` must at least read `isEmpty` to determine if the queue is empty when `head = postTail`. By restricting concurrent calls to these two methods to those situations when `isEmpty` will not be changed by either method (that is, when the queue will be made neither full nor empty by either `Enqueue` or `Dequeue`), we can guarantee deterministic behavior when they are executed in parallel.

The third and final implementation is similar to the second in that its two Integer fields are `head` and `postTail` (and they represent the same things), but

in lieu of a Boolean `isEmpty` field, there is a sentinel node added to the array so that when `head = postTail` it can only be the case that the queue is empty (a full queue has `head = (postTail + 1) mod (MAX_LENGTH + 1)`). Because the length of the array is greater than `MAX_LENGTH`, there will always be some element of the array that is not part of the queue. This differentiation between a full and empty queue without the need to have a separate variable ensures that even when the queue might become either full or empty during a call to `Enqueue` or `Dequeue`, it will not write anything that the other method reads or writes.

## 4 Interference Contracts and Modular Verification

Modular reasoning about the safe execution of concurrent threads can be separated into three distinct parts: (i) a description of the conditions under which operations are independent, (ii) a proof that client code ensures these independence conditions, and (iii) a proof that an implementation guarantees non-interference under these conditions.

Our approach to these three tasks is described below and illustrated using the first bounded queue realization from the previous section.

### 4.1 Interference Contract

A functional specification, as given in section 3.2, does not reveal the degree to which different parts of the abstract state are entangled in the implementation. The correspondence relation between concrete state and abstract state is part of the proof of correctness for the implementation, and modular verification precludes its use in reasoning about client code.

Reasoning about the independence of concurrent threads in client code, however, requires exposing more information. Our approach for describing this independence involves creating an intermediate model consisting of orthogonal components, and encapsulating the description of this intermediate model in a distinct specification, an *interference contract*. While this segmentation is all that is necessary for the example in this paper, in general, an augmentation may additionally supplement the abstract model with more elaboration in order to specify absence of interference among operations. In this case, the specification will also need to state the additional guarantees (ensures clauses) on the supplemental model for each operation, not just interference-related specifications as in the present example.

An interference contract for the bounded queue is given below.

**interference contract** `LookupOffset` **for** `BoundedQueueTemplate`

partition **for** `Queue` **is** (`head`, `tail`, `offset`)

**operation** `Enqueue` (**clears** `e`: `Item`, **updates** `q`: `Queue`)

**affects** `q.tail`

```

preserves q.offset
when q = empty_string affects q.head

operation SwapFirstEntry (updates e: Item, updates q: Queue)
  affects q.head
  preserves q.offset
end LookupOffset

```

Names for different segments of the intermediate model are introduced with the partition keyword. These segments are independent of implementation particulars.

An interference contract includes the *effects* of each operation in terms of this partition. There are two kinds of possible effects: *affects* and *preserves*. The former reflects a possible perturbation (i.e., a write) while the latter reflects non-modifying access (i.e., a read). Standard RESOLVE parameter modes map to these two categories of effects. The partition, however, allows for a finer-granularity description of effects, which is particularly important when concurrent threads use the same variable as a parameter, for example the access of a shared data structure.

RESOLVE's clean semantics ensure that an operation is *oblivious* to (i.e. neither reads nor writes) any variable not explicitly included as a parameter. Similarly, an operation is oblivious any segment of a partition not explicitly mentioned in its effects. For example, SwapFirstEntry is oblivious to q.tail.

A *when* clause gives a condition that restricts the scope of effects. That is, the when predicate must hold initially for the stated effect to occur. For example, in order for Enqueue to affect q.head, the queue must be empty.

Notice that this partitioning of this state space is not the same as establishing the *independence* of these segments from the point of view of the correspondence relation. In this example, the independence of the Enqueue operation on front is conditioned by the queue being non-empty. These independence conditions are in addition to the usual preconditions of the corresponding operations from the template specification, so SwapFirstEntry must be oblivious to q.tail only when the queue is non-empty.

## 4.2 Modular Verification of Client Code

In order for a set of statements to be safely executed in parallel, each variable—or each segment in a variable's intermediate model—can be affected by at most one statement. Furthermore, if any segment is affected by some statement, all of the other statements must be oblivious to this segment.

For example, with the interference contract given above, SwapFirstEntry and Enqueue affect non-overlapping segments (q.head and q.tail, respectively). Furthermore, each is oblivious to the segment affected by the other, assuming the queue is non empty. Finally, the segment used by both (q.offset) is preserved by both. The following client code illustrates the parallel composition of these operations.

```

assume 0 < |q| < MAX_LENGTH
cobegin
  SwapFirstEntry(x, q)
  Enqueue(y, q)
end

```

First we note that the client code above can be executed concurrently only if there is no aliasing between objects  $x$  and  $y$ . This isolation is implied if the programming language is defined to have a clean semantics like RESOLVE or through disciplined programming in a language to avoid unintended aliasing. Under clean semantics, the effects of operations are restricted to their explicit parameters (or explicitly specified global variables) [4].

In addition to satisfying the usual preconditions for functional correctness, the verification of the client code includes establishing the independence conditions of the two operations. This verification is carried out entirely in the context of the client code, using only the abstract functional specification and interference contract of the bounded queue template.

The independence of the constituent statements of a cobegin block means that the statements can be executed in any concurrent or arbitrarily interleaved manner. The semantics of their execution is identical to that of their sequential composition.

### 4.3 Modular Verification of an Implementation

In order to map from concrete implementation state to abstract specification state, realizations provide a representation invariant (convention) and a correspondence function (or relation, more generally). Our approach for establishing operation independence is to augment this correspondence relation with a partitioning of the constituent concrete state space. That is, an implementation must provide a mapping from the concrete data structure involved in the implementation (e.g., contents, front, and length) to the partitioned model of the queue in the interference contract. Specifically, it must place each implementation structure for the queue realization into one of head, tail, and offset.

The restrictions imposed by the effect statements need to be proven for the implementation code of each operation, under the specified conditions. In order for an operation's implementation to meet the obliviousness requirement, all statements in its code must be oblivious to the corresponding parts of the data structure. When a statement does not mention a part of the data structure (e.g., front), it is trivially oblivious to that variable. (This observation also requires clean semantics.) Otherwise, a statement may use parts of the data structure from their obliviousness requirement only in operations which, themselves, are oblivious on the corresponding parts of the data structure. The underlying data structure itself might be built from other data abstractions. This is not a problem, because the lack of entanglement of one component can be layered on top of appropriately disentangled realization components.



```

realization ArrayWithLength for BoundedQueueTemplate
  respects LookupOffset

type representation for Queue is
  (contents: array 0..MAX_LENGTH - 1 of Item,
   front: Integer,
   length: Integer)
exemplar q
convention
  0 <= q.front < MAX_LENGTH and
  0 <= q.length <= MAX_LENGTH
correspondence
  Conc.q = Iterated_Concatenation(i = q.front.. q.front + q.length + 1,
    q.contents(i mod MAX_LENGTH))
interference correspondence
  head: q.contents.c[q.front]
  tail: q.length, q.contents.c except on {q.front}
  offset: q.front
end Queue

procedure Enqueue(clears e: Item; updates q: Queue)
  e := q.contents[q.front + q.length mod MAX_LENGTH]
  q.length := q.length + 1
  Clear(e)
end Enqueue

procedure SwapFirstEntry(updates e: Item; updates q: Queue)
  e := q.contents[q.front]
end SwapFirstEntry

end ArrayWithLength

```

The proof of Enqueue’s obliviousness to q.head (when the queue is non-empty) is seen as follows. When the queue is non-empty, q.length  $\geq 1$ . So the part of q.contents that is modified is distinct from q.contents[q.front]. Therefore, Enqueue is oblivious to q.head. SwapFirstEntry, on the other hand, is oblivious to q.tail. Firstly, the operation does not mention q.length. Secondly, only q.contents[q.front] is affected, so it is oblivious to the rest of the contents.

Notice that the partitioning of q.contents involves the interference contract for an array (i.e., q.contents.c). It is the partition at this nested level that is used in the realization’s interference correspondence.

The proof of preserving q.offset amounts to a proof that no statement in the implementation affects q.front. This proof follows from the interference contracts of the operations used by Enqueue and SwapFirstEntry. In particular, the swapping of e and q.contents[q.front] preserves q.front.

## 5 Summary and Future Directions

This paper has presented a novel framework for modular verification of concurrent programs using data abstractions. Specifically, it has explained how multiple operations can be simultaneously invoked on an abstract data object if a set of interference conditions can be specified and verified using an augmentation to the abstract specification of the data abstraction. The proof process is strictly modularized. The paper has presented a concrete example to illustrate the ideas. Future directions include development of a formal proof system and automated verification.

**Acknowledgments.** This research is funded in part by NSF grants CCF-1161916 and DUE-1022941. Any opinions, findings, conclusions, or recommendations expressed here are those of the authors and do not necessarily reflect the views of the NSF.

## References

1. M. Bagherzadeh and H. Rajan. Panini: A concurrent programming model for solving pervasive and oblivious interference. In *Proceedings of the 14th International Conference on Modularity*, MODULARITY 2015, pages 93–108, New York, NY, USA, 2015. ACM.
2. D. E. Harms and B. W. Weide. Copying and swapping: Influences on the design of reusable software components. *IEEE Trans. Software Eng.*, 17:424–435, 1991.
3. M. Herlihy and N. Shavit. *The Art of Multiprocessor Programming*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 2008.
4. G. W. Kulczycki. *Direct Reasoning*. PhD thesis, Clemson University, Clemson, SC, USA, 2004. AAI3125470.
5. K. R. M. Leino, P. Müller, and J. Smans. *Foundations of Security Analysis and Design V: FOSAD 2007/2008/2009 Tutorial Lectures*, chapter Verification of Concurrent Programs with Chalice, pages 195–222. Springer Berlin Heidelberg, Berlin, Heidelberg, 2009.
6. W. Mostowski. *Verified Software: Theories, Tools, and Experiments: 7th International Conference, VSTTE 2015, San Francisco, CA, USA, July 18-19, 2015. Revised Selected Papers*, chapter Dynamic Frames Based Verification Method for Concurrent Java Programs, pages 124–141. Springer International Publishing, 2016.
7. E. Rodriguez, M. Dwyer, C. Flanagan, J. Hatcliff, and G. T. Leavens. Extending jml for modular specification and verification of multi-threaded programs. In *In ECOOP, LNCS 3586*, pages 551–576. Springer, 2005.
8. M. Sitaraman, B. Adcock, J. Avigad, D. Bronish, P. Bucci, D. Frazier, H. M. Friedman, H. Harton, W. Heym, J. Kirschenbaum, J. Krone, H. Smith, and B. W. Weide. Building a push-button resolve verifier: Progress and challenges. *Formal Aspects of Computing*, 23(5):607–626, 2011.
9. J. Smans, B. Jacobs, and F. Piessens. Vericool: An automatic verifier for a concurrent object-oriented language. In *Proceedings of the 10th IFIP WG 6.1 International Conference on Formal Methods for Open Object-Based Distributed Systems, FMOODS '08*, pages 220–239, Berlin, Heidelberg, 2008. Springer-Verlag.