

Operation-Valency and the Cost of Coordination

Danny Hendler¹ and Nir Shavit

Tel-Aviv University

This paper introduces *operation-valency*, a generalization of the valency proof technique originated by Fischer, Lynch, and Paterson. By focusing on critical events that influence the return values of individual operations rather than on critical events that influence a protocol’s single return value, the new technique allows us to derive a collection of realistic lower bounds for lock-free implementations of concurrent objects such as linearizable queues, stacks, sets, hash tables, shared counters, approximate agreement, and more. By realistic we mean that they follow the real-world model introduced by Dwork, Herlihy, and Waarts, counting both memory-references and memory-stalls due to contention, and that they allow the combined use of read, write, and read-modify-write operations available on current machines.

By using the operation-valency technique, we derive an $\Omega(\sqrt{n})$ *non-cached shared memory accesses* lower bound on the worst-case time complexity of lock-free implementations of objects in *Influence(n)*, a wide class of concurrent objects including all of those mentioned above, in which an individual operation can be influenced by all others.

We also prove the existence of a fundamental relationship between the space complexity, latency, contention, and “influence level” of any lock-free object implementation. Our results are broad in that they hold for implementations combining read/write memory and *any* collection of read-modify-write operations, and in that they apply even if shared memory words have unbounded size.

1. INTRODUCTION

In 1993, Dwork et. al [7] introduced a formal model to capture the real world phenomenon of memory contention on today’s shared memory machines, machines that allow read, write, and read-modify-write (RMW) operations. Using FLP-style valency arguments, they proved that there are inherent tradeoffs between contention and latency² in concurrent data structure design. Their work was extended in several directions, most notably in the context of mutual exclusion [1; 2; 6] and counting networks [3; 4].

This paper presents *operation-valency*, a generalization of the valency proof technique of Fischer et. al (FLP) [9], and uses it to continue the above work in deriving real-world time complexity lower bounds for state of the art concurrent objects. As surveyed by Lynch [14] and by Fich and Rupert [8], there are numerous elegant extensions and reformulations of the FLP-style valency technique. The main difference between the operation-valency approach we present here and FLP-style arguments is that we focus on *temporary* changes in the anticipated results of *individual-operation solo executions* rather than on permanent changes in the valency of the protocol as a whole. In doing so, we are able to capture some of the

¹This work was supported in part by a grant from Sun Microsystems. Contact author: Danny Hendler, hendlerd@post.tau.ac.il.

²In [7] Dwork et. al define a protocol’s latency to be the maximal number of shared variable accesses, over all executions, a single high-level operation does.

complexity resulting from the influence among shared object operations that return distinct yet dependent values.

The time metric we use, which we call *memory steps*, counts only first-access shared memory events, and memory stalls due to contention in writing to shared locations. It is stricter than the time metric used by [7], as the later counts all shared memory references and also counts memory stalls due to contention in *reading*; it is similar to the *communication cost* metric used by Cypher [6], and to the *remote memory references* metric used by Anderson and Yang [1], and by Anderson and Kim [2], in that a single unit of both metrics corresponds to a shared memory reference that cannot be resolved by a local cache ³.

We use the new operation-valency technique to derive a collection of tradeoffs and lower bound results. Specifically, we are able to show an $\Omega(\sqrt{n})$ time-complexity lower bound on lock-free implementations of objects in a class we call *Influence(n)*, a wide class of concurrent objects in which an individual operation can be influenced by $\Omega(n)$ other operations. *Influence(n)* includes data structures such as linearizable queues, stacks, shared counters, hash tables, sets, multi-sets, and approximate agreement objects. Our results are the first known time complexity lower bounds for implementing these objects using *any* RMW operation. Before listing our results in detail, let us briefly describe the operation-valency technique.

1.1 The Operation-Valency Technique

Valency arguments, introduced by Fischer et. al [9], have been used extensively [5; 7; 10; 13] to derive impossibility results and lower bounds for consensus and related problems. In problems such as consensus, a protocol is required to eventually output a *single* protocol value. FLP-style valency classifies system states according to whether they are *univalent* or *multivalent*. A state S is *univalent* if, in any two execution extensions starting from S , the protocol outputs the same value, and *multivalent* otherwise. Thus, essentially, a state S is univalent iff the protocol's output value is already determined in S ⁴. The valency technique looks at critical events that atomically change the system state from multivalent to univalent. Valency arguments are then applied w.r.t. to these critical events to derive impossibility results or lower bounds.

The basic idea behind our operation-valency technique is to generalize the approach of [9] by looking at the return values of individual operations instead of the single return value of the protocol as a whole. Similarly to FLP-style valency, we identify critical events that atomically affect return values, and we argue about the order and location of these events to obtain our results. We note, however, the significant differences between operation-valency and FLP-style valency:

—FLP-style valency looks at a single protocol output value; operation-valency looks at the return value of specific operations, in protocols where different operations

³Note, however, that the communication cost metric and the remote memory references metric are stricter than our memory steps metric, since in distributed shared memory (DSM) systems, they do not count references to a process' local segment of shared memory, whereas such references *may* be counted as memory steps by our metric.

⁴The output value may still be unknown to all participating processes, however, even in a univalent state.

are allowed to return different values;

- FLP-style valency looks at critical events that have *permanent* effect on the protocol’s output value: before the event, there exist two different execution extensions that yield two different output values; after the critical event is executed, *all* execution extensions yield the same protocol output value. Operation-valency looks at a different class of critical events, which we call *modifying events*. These events atomically modify the return value of a *solo continuation* of a specific operation R by some process p : before such a modifying event e is executed, R ’s solo continuation must return some value V ; *right after* e is executed, R ’s solo continuation must return some other value. The effect of a modifying event may be *temporary*: the execution may be extended such that if a solo continuation of R takes place at a later stage, it would, once again, have to return V .

As an example, consider an implementation of a “one time” n -process linearizable counter object allowing *fetch-and-increment* (FAI) operations. Before execution starts, any process may start a solo execution that returns 1. Let E be an execution where process p is idle, and assume some process q , $q \neq p$, completes its FAI operation in E , then a solo execution by p after E must return a value bigger than 1. We identify the critical modifying memory events, write, or RMW events, following which the return value of p ’s solo execution is modified. Our proof technique constructs executions in which such modifying events are pending, and shows that the operations whose return values are about to be atomically influenced by them must read the memory locations on which they are pending, otherwise we can construct indistinguishable executions that will lead to contradicting outcomes.

1.2 Our Results

To characterize the coordination requirements of shared objects, we introduce the *influence level* metric \mathbb{I} , informally defined as the maximal number of high level operations by other processes that can influence the outcome of another given process’ high level operation. For example, in a linearizable shared counter, the outcome of a given operation can be influenced by $n - 1$ others: if it runs alone it will return one value, but if any of the $n - 1$ other processes precedes it, the value returned will be different.

1.2.1 New Fundamental Tradeoffs. We prove that the following fundamental relationships exist for all lock-free protocols. Let P be a lock-free protocol for a shared object with influence level \mathbb{I} ; let $\mathbb{L}(P)$, $\mathbb{S}(P)$, and $\mathbb{C}(P)$ respectively denote the latency, space complexity, and write-contention of P , then:

$$\mathbb{L}(P) \geq \mathbb{I}/\mathbb{C}(P), \quad \mathbb{S}(P) \geq \mathbb{I}/\mathbb{C}(P) \quad (1)$$

For example, for linearizable counting this tradeoff strengthens a result of Herlihy et. al [11], which try to capture contention via a static measure of *capacity*: the maximal number of processes $c(P)$ that access any particular variable in any execution. They prove the existence of the tradeoff $\mathbb{L}(P) \geq (n-1)/c(P)$ between the capacity and latency of linearizable counters. However, they note that the capacity $c(P)$ is not necessarily correlated with contention. Our tradeoff captures a stronger relationship between latency and the *actual write-contention* for a broad class of problems, which for linearizable shared counting implies the desired relationship

$\mathbb{L}(\mathbf{P}) \geq (n-1)/\mathbf{C}(\mathbf{P})$. The above tradeoff also answers an open question posed by Dwork et. al [7] as to whether there exists a tradeoff between latency and contention for the approximate agreement problem.

1.2.2 New Time Complexity Lower Bounds. We identify *Influence*(n), the class of problems where a single operation has influence level $\mathbb{I} \in \Omega(n)$. This class includes lock-free implementations of key objects such as linearizable queues, stacks, hash-tables, sets, shared counters, approximate agreement, and consensus. We prove a lower bound of $\Omega(\sqrt{n})$ memory steps for any object belonging to *Influence*(n). This bound has an immediate real-world implication: any lock-free implementation for any of the above objects on any of today’s architectures, using any combination of reads, writes, and RMWs, has a worst-case operation time-complexity of at least \sqrt{n} shared memory accesses; moreover, we show that these accesses cannot be resolved by a local cache.

Our lower bounds are the first real-world time bounds for such objects. Though they seem stronger than Jayanti’s interesting $\Theta(\log n)$ time bounds [12] on similar objects, they are in a sense orthogonal. This is because Jayanti’s time metric does not count contention when accessing shared locations. For example, according to Jayanti’s metric, a shared linearizable counter for n processes can be implemented in constant time. Unlike our results, Jayanti’s bounds are restricted to a model with only load-locked/store-conditional, a specific type of RMW operation.

Finally, we show that there exists an object in *Influence*(n), which we name *First-Generation*, for which our bound is tight, that is, it has $\Theta(\sqrt{n})$ memory steps complexity. However, we believe the tight bounds on many interesting problems in *Influence*(n) are higher, and that the operation-valency approach might be the key to deriving them.

2. PRELIMINARIES

2.1 Shared Memory System Model

Our model of an asynchronous shared memory system is based on the model described by Cypher in [6], which is based, in turn, on the model given by Merritt and Taubenfeld [15]. We will assume that shared objects are specified as in [10]. Our model allows three types of shared memory *events*: atomic Read, atomic Write, and Read-Modify-Write events. No bound is assumed on register size. An n -process shared memory protocol $(\mathbf{E}, \mathbf{P}, \mathbf{R}, \mathbf{I})$ consists of a non-empty set \mathbf{E} of executions, a set \mathbf{P} of n processes, a set \mathbf{R} of memory registers, and a function \mathbf{I} that assigns an initial value to each register in \mathbf{R} . An execution is a sequence (either finite or infinite) of events, where an event is an atomic memory access performed by a single process. An event can have one of the following three forms:

- Atomic Read: $read(p, r, v)$ indicates that process p reads the value v from register r ;
- Atomic Write: $write(p, r, w)$ indicates that process p writes the value w to register r ;
- Atomic RMW: $RMW(p, r, v, w)$ indicates that process p atomically does the following: it reads the value v from register r and computes w based on v . p then

proceeds in one of the following two ways: if w is different from the special value $null$, p writes the value w to register r , otherwise p does not write to register r .⁵

Given any event e , $\text{read}(e)$ is true if e is either a Read or a RMW event, and $\text{write}(e)$ is true if e is a Write event, or if e is a $\text{RMW}(p, r, v, w)$ event with $w \neq null$. $\text{mem}(e)$ is the memory register *accessed* (read and/or written) by e ; $\text{proc}(e)$ is the process that executed e . For any $e \in E$, $\text{index}(e, E)$ is the number of events that precede e in E ; when the execution discussed is clear from the context, we just use $\text{index}(e)$.

For any finite execution E and any sequence of events E' , $E \circ E'$ denotes the concatenation of E and E' . For any sequence of events E' , we define $\text{procs}(E')$ to be the set of processes that perform some event in E' . Let $r \in R$ be a memory register, and $E \in \mathbf{E}$ a finite execution, then $\text{value}(r, E)$ (the value of r after E) is the value written by the last event in E that wrote to r , or $\mathbf{I}(r)$ if there was no such event. Given an execution E and any subset $P \subseteq \mathbf{P}$, we let $\text{proj}(E, P)$ denote the subsequence of E containing only the events in E that were issued by processes in P . If $P = \{p\}$, we also use the notation $\text{proj}(E, p)$ instead of $\text{proj}(E, \{p\})$. If $\text{proj}(E_1, p) = \text{proj}(E_2, p)$, we say that the executions E_1, E_2 are *indistinguishable* by p . Let $E \in \mathbf{E}$ be a finite execution, and e be an event; if $E \circ e \in \mathbf{E}$ holds, we say that e is *enabled after* E .

Definition 2.1 *A shared memory protocol satisfies the following shared memory axioms A1 - A3:*

- **A1:** *If $E \circ e \in \mathbf{E}$, then $E \in \mathbf{E}$.*
- **A2:** *Let $E \circ e \in \mathbf{E}$ be an execution, and let e be a Read or RMW event, then the value read by e is $\text{value}(\text{mem}(e), E)$.*
- **A3:** *Let $E \circ e \in \mathbf{E}$ be an execution, and assume $\text{proj}(E, p) = \text{proj}(E', p)$, then $E' \circ e \in \mathbf{E}$ holds.*

Intuitively, Axiom **A1** states that a prefix of any possible execution is also a possible execution; Axiom **A2** states that the value read by any Read or RMW event returns the value of the most recent write to the accessed register (or the register's initial value if there were no earlier writes to that register); finally, Axiom **A3** states that if there are two possible executions that are indistinguishable by process p , and it is possible for p to perform a memory-operation (either Read, Write or RMW) following one of the executions, then it is also possible for p to perform the same operation following the other execution.

2.2 High Level Operations

Shared memory protocols support high-level operations that processes can execute. We consider protocols that support at least one operation-type that returns a value. High-level operations involve, in general, both private- and shared-memory events; in this paper we only deal with shared-memory events, and so we view a high-level operation Op as consisting of a sequence of one or more atomic shared memory

⁵This definition of RMW events captures both conditional (such as: Compare-and-swap, Test-and-set, Load-linked/Store-conditional) and non conditional RMW operations.

events, each of which can be either *Read*, *Write*, or *RMW*. We only consider executions where each process performs at most a single high-level operation.⁶ Let Op be an operation performed by some process in some execution E . We denote by $proc(Op, E)$ the process that executes Op in E ; we denote by $events(Op, E)$ the sequence of memory-events performed by $proc(Op)$ while executing Op in E . Whenever E is clear from the context, we simply write $proc(Op)$ and $events(Op)$. We denote by $first(Op, E)$ and $last(Op, E)$ the first and last events, respectively, in $events(Op, E)$. (Note that $first(Op, E)$ and $last(Op, E)$ may be the same event.) As before, we omit E when it is clear from the context, and simply write $first(Op)$ and $last(Op)$. We say that an operation Op is contained within an execution E ($Op \subset E$) if all the events of $events(Op)$ appear in E . If $Op \subset E$, we denote by $result(Op, E)$ the value returned by Op in E . Let Op_1, Op_2 be two operations contained in E ; if $index(last(Op_1), E) < index(first(Op_2), E)$, we write $Op_1 \xrightarrow{E} Op_2$. Note that for any execution E , \xrightarrow{E} is a partial order on all the operations contained in E . We say that an execution E is *quiescent*, if there is no operation that started in E but did not terminate in E . We say that a high-level operation Op is enabled after an execution E , and write $enabled(E, Op)$, if Op has a enabled event e after E . Assume $enabled(E, Op)$ holds, then we denote by $solo(E, Op)$ the sequence of events which constitute the solo execution of operation Op after E .

Slightly abusing notation, and for presentation simplicity, we sometimes refer to an underlying execution as a *state*. Thus, e.g., instead of saying that a high-level operation Op is enabled after execution E , we say that after E the system is in a state S where Op is enabled; instead of referring to a solo execution of Op after E by $solo(E, Op)$, we refer to it by $solo(S, Op)$, where S is the system state after execution E .

2.3 The Time Complexity Metric

Our time complexity metric counts the worst-case number of *memory steps* that a single high-level operation may incur. Our metric counts both first-access shared memory events and *stalls* that are incurred when a few processes concurrently attempt to perform a Write or RMW operation to the same memory register. Formal definitions follow.

Definition 2.2 *Let $E \in \mathbf{E}$ be an execution; let e be an event in E ; let $r = mem(e)$ be the memory register accessed by e , and let Op be a high-level operation such that $e \in events(Op, E)$, then we say that e is a first-access event of Op in E , if e is the first event in $events(Op, E)$ that accesses r . We denote the number of first-access events performed by Op in E by $first_access(E, Op)$.*

In words, an event is a first-access event of a high-level operation Op , if it is the first event of Op to access some memory register.

Definition 2.3 *Let $E \in \mathbf{E}$ be an execution, and let $e_j, 0 \leq j \leq l, l > 0$, be a maximal sequence of consecutive events in E such that the following holds:*

- (1) $\forall j, 0 \leq j \leq l : e_j$ is either a Write or a RMW event.

⁶Obviously, this just strengthens our lower bounds.

$$(2) \forall j_1, j_2, 0 \leq j_1 \neq j_2 \leq l : \\ (\text{proc}(e_{j_1}) \neq \text{proc}(e_{j_2})) \wedge (\text{mem}(e_{j_1}) = \text{mem}(e_{j_2}))$$

Let Op be the high-level operation whose execution issued e_j , then we say that Op incurs j stalls in E on account of e_j , and write $\text{stalls}(E, e_j) = j$.

The above definition of stalls captures the fact that in shared memory systems, when a group of processes have pending Write or RMW events to the same memory location, then a scheduling adversary can release all these events simultaneously, thus causing the operation that issued the second event to incur a single stall, the operation that issued the third event to incur two stalls, and so on.⁷

Definition 2.4 Let E be an execution and let Op be a high level operation such that $Op \subset E$; the memory steps complexity of Op in E , denoted $\text{mem_steps}(E, Op)$, is defined as follows:

$$\text{mem_steps}(E, Op) = \text{first_access}(E, Op) + \sum_{e \in \text{events}(E, Op)} \text{stalls}(e)$$

3. THE INFLUENCE METRIC FOR COORDINATION LEVEL

In this section we define a quantitative metric which is a measure of the coordination level of distributed protocols. More specifically, the *influence level* metric is a measure of the extent to which concurrently executing operations (which we call *influencing* operations or just *influencers*) can influence the result of another operation (which we call the *influenced* operation). To get a feel for this metric, consider an n -process protocol that implements a linearizable stack, with *push* and *pop* high-level operations. Consider a quiescent state S , where the stack contains a single item - the number 1. Assume that process p_1 has an enabled *pop* operation, and each of the processes $p_i, 2 \leq i \leq n$ has a enabled *push*(i), for $2 \leq i \leq n$, respectively. Clearly, the value returned by the *pop* operation can be influenced by the *push* operations: if the *pop* operation is allowed a solo-execution while the *push* operations have not yet begun, then, from linearizability, it has to return 1; on the other hand, if we allow any interleaved execution of the *push* operations where at least one of them terminates, and only then start a solo-execution of the *pop* - then (again from linearizability) the *pop* must return a different value. Dependencies of this type are what we capture in the following definitions and lemmata.

Definition 3.1 Let R be a high-level operation which is enabled in state S ; we say that S has influence level K w.r.t. R (and write $\mathbb{I}(S, R) = K$) if the following holds:

- $\text{solo}(S, R)$ returns some value V ;
- There are exactly K enabled high-level operations $W_i, 1 \leq i \leq K$, that meet the following condition: each operation W_i has an enabled event e_i , such that $\text{solo}(S \circ e_i, R)$ returns a value other than V .

We say that R is K -influenced in S and that the operations W_i are the influencers of R in S . We call V the distinguished value of R in S . We call the events e_i modifying

⁷This definition of stalls does *not* assume that concurrent write events to the same memory-register are serviced in FIFO order, or in any other order.

events. We define the influence level of S as the maximum influence level over all enabled high-level operations: $\mathbb{I}(S) = \max\{\mathbb{I}(S, Op) \mid \text{enabled}(S, Op)\}$

We next extend the above definition of influence level to executions, protocols, and concurrent objects.

Definition 3.2 *The influence level of an execution E , denoted by $\mathbb{I}(E)$, is the maximum influence level over all the states E reaches; the influence level of a protocol P , denoted by $\mathbb{I}(P)$, is the maximum influence level over all its executions.*

Slightly abusing notation, we now define the influence level of objects.

Definition 3.3 *A concurrent object O has influence level \mathbb{I} , if the influence level of every lock-free protocol implementing it is at least \mathbb{I} .*

Based on the above definitions, the next two lemmas prove that we can determine a lower bound on the influence level of linearizable implementations of an object, based on the object's sequential specification.

Lemma 3.1 *Let P be lock-free protocol that can be brought to a state S where the following holds:*

- *There is a process that has an enabled operation R so that $\text{solo}(S, R)$ returns some value V ;*
- *There are K other processes, each having an enabled operation $W_i, 1 \leq i \leq K$, such that: after any execution E' , consisting of events issued by the operations W_i , where at least one operation terminates - $\text{solo}(S \circ E', R)$ returns a value other than V .⁸*

Then P has influence level at least K .

Proof We construct an execution, E , that leads to a state with influence level at least K . E is constructed iteratively. Initially E is the empty execution. In each iteration, we pick some W_i whose next event e is not a modifying event; we extend E with e , and we let W_i execute e . From the lemma assumptions, no operation can terminate before a modifying event is performed by some W_i operation; also, as P is lock free, some W_i operation must terminate after a finite number of events are executed, and so the construction is guaranteed to bring us to the desired state after some finite number of iterations. Q.E.D.

Lemma 3.2 *Let O be an object whose sequential specification S includes a history H , such that:*

- *H can be extended by a value-returning operation R , and $\text{result}(R, H \circ R) = V$ for some V ;*
- *There are K operations, $W_i, 1 \leq i \leq K$, such that the following holds: for any non-empty subset of indices $T \subseteq \{1, \dots, K\}$, and for every permutation σ_T of T , the following extension of H exists: $H_{\sigma_T} = H \circ W_{\sigma_T(1)} \circ \dots \circ W_{\sigma_T(|T|)} \circ R$ and $\text{result}(R, H_{\sigma_T}) \neq V$.*

Then any linearizable implementation of O has influence level at least K .

⁸It is not assumed that the operations W_i return a value.

Proof Immediate from Lemma 3.1 and from the linearizability of the implementation. Q.E.D.

4. TRADEOFFS AND LOWER BOUNDS

We consider a lock-free n -process protocol, P , for $n \geq 2$, that has influence level K . We define P 's *latency* as the maximal number of shared memory events issued by a single high-level operation, over all executions, and denote it by $\mathbb{L}(P)$; we define P 's *read latency* as the maximal number of read/RMW shared memory events issued by a single high-level operation, over all executions, and denote it by $\mathbb{L}_{\mathbb{R}}(P)$; we define P 's *space complexity* as the total number of shared memory registers read/written by P , over all executions, and denote it by $\mathbb{S}(P)$; we define P 's *write-contention* as the maximal number of consecutive write/RMW events by different processes to the same memory register, over all executions, and denote it by $\mathbb{C}(P)$. Our proofs need only consider executions where each process executes at most a single operation⁹.

The following lemma proves that modifying events are either write or RMW events.

Lemma 4.1 *Let e be a modifying event, then $\text{write}(e)$ holds.*

Proof Let e be a modifying event w.r.t. some operation R in a state S . Assume by way of contradiction that the claim does not hold, then e does not write any value. From the definition of modifying events, $\text{solo}(S, R)$ returns value V , but $\text{solo}(S \circ e, R)$ returns a different value. Since e does not involve a write, and since e is not an event of p , the states S and $S \circ e$ are indistinguishable w.r.t p and therefore, by Axiom A3 of the shared memory model, $\text{solo}(S, R)$ and $\text{solo}(S \circ e, R)$ are identical, which implies that a solo execution of R returns V also right after e is performed. This obviously implies that e is not a modifying event, a contradiction. Q.E.D.

We now prove a tradeoff between the space-complexity and write-contention of any lock-free object implementation.

Theorem 4.2 *Let O be an object with influence level \mathbb{I} and let P be a lock-free implementation of O , then the following holds:*

$$\mathbb{S}(P) \geq \lceil \mathbb{I}/\mathbb{C}(P) \rceil$$

Proof Since P implements O , $\mathbb{I}(P) \geq \mathbb{I}$. Consequently, from Lemma 4.1 and from the definition of influence level, P can be brought to a state where at least \mathbb{I} write or RMW events are enabled. Since at most $\mathbb{C}(P)$ such events can be outstanding on any single register, the result follows. Q.E.D.

We next prove a similar tradeoff between the latency and write-contention of lock-free implementations. We actually prove a stronger result, by showing that the tradeoff holds even if we exclude write events from the latency count.

Theorem 4.3 *Let O be an object with influence level \mathbb{I} and let P be a lock-free implementation of O , then the following holds:*

$$\mathbb{L}(P) \geq \mathbb{L}_{\mathbb{R}}(P) \geq \lceil \mathbb{I}/\mathbb{C}(P) \rceil$$

⁹Obviously this just strengthens our lower bounds.

Proof As all the events counted by $\mathbb{L}_{\mathbb{R}}(P)$ are also counted by $\mathbb{L}(P)$, the left-hand inequality is obvious. As for the right-hand inequality, since P implements O , we have $\mathbb{I}(P) \geq \mathbb{I}$. From the definition of influence level, there is some execution E of P that leads to a state S with influence level at least I . Let R be an I -influenced operation by process p in S ; let V be R 's distinguished value in S ; let q_1, \dots, q_I be influencing processes, each having an enabled modifying event w.r.t R ; also, let B be the set of registers to which these outstanding modifying events are about to write. Note that $|B| \geq \lceil \mathbb{I}/\mathbb{C}(P) \rceil$. Since no modifying event has been performed yet, then from the definition of modifying events $result(Solo(S, R)) = V$. We prove the theorem by showing that $Solo(S, R)$ must read all the registers in B . Assume to the contrary, then $Solo(S, R)$ returns a value without reading some register $r \in B$. Let e be some modifying event which is outstanding on register r , then clearly $proj(E \circ e, p) = proj(E, p)$. Consequently, by inductive application of Axiom A3 of the shared memory model, we get that $result(Solo(S, R)) = result(Solo(S \circ e, R))$, which is a contradiction to our assumption that e is a modifying event. Q.E.D.

Based on Theorem 4.3, we can now establish a lower bound on the memory steps complexity of any lock-free object implementation.

Theorem 4.4 *Let O be an object with influence level \mathbb{I} , and let P be a lock-free implementation of O , then the memory steps complexity of P is at least $\lfloor \sqrt{\mathbb{I}} \rfloor$.*

Proof As P implements O , we have $\mathbb{I}(P) \geq \mathbb{I}$. From the definition of influence level, there is an execution E , such that following E there are processes q_1, \dots, q_I , each of which having an enabled modifying event w.r.t some operation R . According to Lemma 4.1 all of the modifying events are either Write or RMW events. Let B be the set of registers to which these outstanding events are about to write. Assume that $|B| < \lfloor \sqrt{\mathbb{I}} \rfloor$, then there is at least one register $r \in B$ that has at least $\lfloor \sqrt{\mathbb{I}} \rfloor$ outstanding events about to write to it. Let q_j be the process whose outstanding event on r is executed last, then the high-level operation of q_j is charged by $\lfloor \sqrt{\mathbb{I}} \rfloor - 1$ memory steps because of the stalls it incurs when accessing r , plus an additional memory step on account of the first access of r , which proves the theorem. Otherwise, $|B| \geq \lfloor \sqrt{\mathbb{I}} \rfloor$, and based on Theorem 4.3 R can be made to access all the registers in B . Consequently we can charge R by $|B|$ memory steps for the first-access events of all the registers in B . Q.E.D.

Figure 1 illustrates the rationale of the lower bound derived in Theorem 4.4. The lower bound is derived by proving that any lock-free implementation, P , of an object O with influence level \mathbb{I} , has at least one of two types of executions. P either has an execution in which a large number of processes have simultaneous enabled writes to the same register (as in Figure 1, (a)) or it has an execution in which some process has to read a large number of registers in the course of a single high-level operation (as in Figure 1, (b)).

4.1 Memory Steps and Time

We now discuss how the memory steps lower bound translates to a time lower bound. For this, we need the following definition:

Definition 4.1 *Let M be a shared memory multiprocessor; we denote by $n\text{lc-time}(M)$ the minimal time a non-local-cache-reference takes in M , i.e. the minimal*

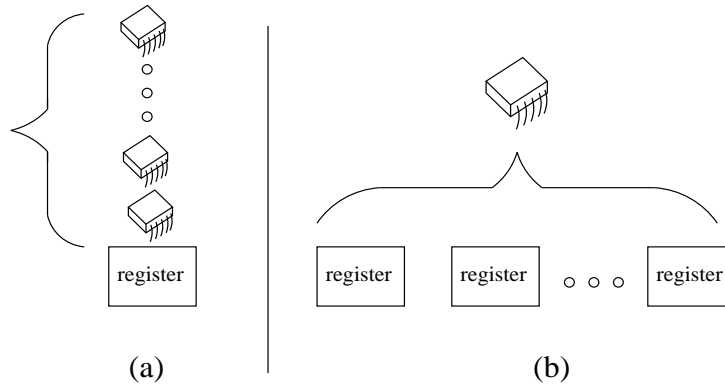


Fig. 1. Every lock-free implementation can be brought to at least one of two states: either a large number of processes are about to write to the same register (a), or some process has to access a large number of registers in the course of its high-level operation (b).

time in M of a memory reference that is not resolved by the local cache (if any) of the processor that issued it.

If M is a distributed shared memory system without caches, $n\text{lcr-time}(M)$ is simply the minimal time it takes a processor in M to access its local segment of shared memory (the minimum taken over all processors); if M is a cache-coherent multiprocessor, then the minimum in the above definition is taken over all memory references that cannot be resolved by the local cache and generate interconnect traffic, such as: references to a memory location that is not in the local cache; writes that generate cache-invalidate transactions; writes that generate cache-update transactions, or any other shared memory references not resolved in the local cache.

In the proof of Theorem 4.4 we have shown that either some operation incurs at least $\lfloor \sqrt{\mathbb{I}} \rfloor$ consecutive stalls, or some other operation performs at least $\lfloor \sqrt{\mathbb{I}} \rfloor$ first-access read events. We now analyze both cases.

- In any shared-memory multiprocessor M , when multiple processors attempt to write to the same memory register simultaneously, the writes are being serialized and are serviced one after the other. Moreover, even if M is a cache-coherent system, x consecutive stalls take at least $x \cdot n\text{lcr-time}(M)$ time: if the cache scheme is *write-through*, then every write generates a cache miss; and even if the cache scheme is *write-back*, then since the writes are by different processors, none of them (except, maybe, the first) can be accomplished by just updating the local cache: they have to either invalidate or update other caches.
- Clearly in no shared-memory multiprocessor can a first-time shared memory read be resolved from the local cache; consequently, if we assume M does not support non-blocking reads, then x first-time access events by an operation take at least $x \cdot n\text{lcr-time}(M)$ time. If M *does* support multiple outstanding reads per processor, then, theoretically, x first-access reads may be resolved in a time equivalent to x cache references.

4.2 The $Influence(n)$ Objects Class

We now define the $Influence(n)$ class of concurrent objects, that contains objects for which every lock-free n -process implementation has influence level in $\Omega(n)$. We then show that many key distributed objects belong to this class, and thus have an inherent operation complexity of $\Omega(\sqrt{n})$ memory steps. We conclude this section by describing the *First Generation* object; we prove that this object belongs to the $Influence(n)$ class, and that it has a $\Theta(\sqrt{n})$ memory steps lock-free implementation. Thus we prove, that the $\Omega(\sqrt{n})$ bound for the $Influence(n)$ class is tight.¹⁰

Definition 4.2 *A generic object O is an object that is specified for any number of processes n . The influence-function of O , denoted \mathbb{I}_O , is defined as follows: $\mathbb{I}_O(n) = K$, if the influence level of every lock-free n -process implementation of O is at least K .*

Definition 4.3 *Influence(n) is the objects class that contains all generic objects O such that \mathbb{I}_O is in $\Omega(n)$.*

It is easily shown that the following objects are in $Influence(n)$: linearizable counters, stacks, queues, hash-tables, sets, approximate agreement. As two examples, we show that approximate agreement and linearizable counting belong to $Influence(n)$.

An approximate agreement object supports a single *decide* operation. Each participating process calls *decide* with the process' real-number input-value. The values returned by the *decide* operation to different processes are required to be within a given distance ϵ of each other, and are also required to be within the range of the inputs.

Proof We prove that any lock-free approximate-agreement protocol for n processes has influence level $n - 1$. Consider the problem instance where process $p_i, 1 \leq i \leq n$, starts with value $2(i - 1) \cdot \epsilon$. We denote the *decide* operation executed by process i as *decide* _{i} . We now show that the initial state, S , meets the conditions of Lemma 3.1. clearly, if *decide*₁ runs alone, it must return 0. On the other hand, in any execution E which does not involve p_1 , in which some other process decides - the decision value must be in the range $[2\epsilon \cdots 2(n - 1)\epsilon]$, and so if *decide*₁ starts to execute after E , *decide*₁ must return a value no less than ϵ . Consequently, from Lemma 3.1, the influence level of any n -process lock-free implementation of approximate-agreement is at least $n - 1$. From the definition of influence level, the influence level of an n -process protocol cannot be more than $n - 1$, hence the result follows. Q.E.D.

A shared counter object supports a single fetch-and-increment (*FAI*) operation. The counter-values returned by *FAI* operations are required to be unique natural numbers. It is also required that in quiescent states the values distributed by the counter constitute a contiguous range of natural numbers. *Linearizable shared counters* are also required to be linearizable, i.e. if FAI_i and FAI_j are two activations of the *FAI* operation, and $FAI_i \xrightarrow{E} FAI_j$ in an execution E , then $result(FAI_i, E) < result(FAI_j, E)$ must hold.

¹⁰This does not imply, however, that the bound is tight for *all* the problems in $Influence(n)$.

Theorem 4.5 *Linearizable counting is in Influence(n).*

Proof We prove that any lock-free linearizable counting protocol for n processes has influence level $n-1$. The sequential specification, SEQ , of a linearizable counter contains a history $H_1 = FAI_1$, where process p_1 returns 1; additionally, for any $T \subseteq \{2 \cdots n\}$, and for any permutation σ_T of T , SEQ contains the following history: $H_{\sigma_T} = FAI_{\sigma_1} \circ \cdots \circ FAI_{\sigma_{|T|}} \circ FAI_1$, and $result(FAI_1, H_{\sigma_T}) > 1$. Consequently, by using Lemma 3.2, the result follows. Q.E.D.

The proofs that linearizable stacks, queues, sets and hash-tables are in *Influence(n)* are very similar to the proof of Theorem 4.5, and are consequently omitted.

We next present the *First-Generation* problem. We show that it belongs to the *Influence(n)* class and that it can be implemented in $\Theta(\sqrt{n})$ memory steps; thus we prove, that there are problems in *Influence(n)* for which our bound is tight. Let E be an execution; we say that an operation Op belongs to the first-generation of E , and write $Op \in FG(E)$, if it has no predecessor in the partial-order induced by \xrightarrow{E} .

Definition 4.4 *A First-Generation object supports a single operation - First, which every process can call at most once. The operation returns a boolean value. Any correct implementation must meet the following requirements for every non-empty execution E :*

- An operation which is not in $FG(E)$ cannot return true in E ;
- If all the operations in $FG(E)$ terminate, then at least one of them returns true.

Lemma 4.6 *First-Generation is in Influence(n).*

Proof We denote by $First_i$ the *First* operation performed by process i , $1 \leq i \leq n$. Let S be the protocol initial state; we obtain the result by proving that the conditions of Lemma 3.1 hold for S . clearly, $value(Solo(S, First_1)) = true$. Let E be an execution which does not involve p_1 , in which some process p_i , $1 < i \leq n$ completes its $First_i$ operation, and let S' be the resulting state, then, from the definition of the *First-Generation* object, $value(Solo(S', First_1)) = false$. Consequently, from Lemma 3.1, the influence level of any n -process lock-free implementation of *First-Generation* is at least $n-1$. From the definition of influence level, the influence level of an n -process protocol cannot be more than $n-1$, hence the result follows. Q.E.D.

We now present a simple $\Theta(\sqrt{n})$ memory steps lock-free n -process protocol implementing a *First-Generation* object. The protocol uses an array of multi reader multi writer atomic registers, *mark*, of size $\lceil \sqrt{n} \rceil$. The entries of the *mark* array are initialized to *false*. The code implementing the *First* operation is shown in Figure 2. The unique id of each process is stored in a local register called *myId*.

In the following, we prove that the code shown in Figure 2 is a correct implementation of the *First-Generation* object.

Lemma 4.7 *Let E be a non-empty execution such that all the operations in $FG(E)$ has terminated, then at least one of these operations returns true on line 5.*

Proof Let Op be the first operation in $FG(E)$ to exit the loop of lines 1-3; assume Op exits the loop at time t . As no other operation could have set its *mark*

```

boolean First()
{
1:   for (k=0; k< (sizeof mark); k++)
2:       if (mark(k) == true)
3:           return false;
4:   mark[sqrt(myId)] = true;
5:   return true;
}

```

Fig. 2. First Operation Code

flag before time t , Op must have read all $mark$ flags as *false*, and so must have set its $mark$ flag in line 4, and then returned *true* on line 5. Q.E.D.

Lemma 4.8 *Let E be a non-empty execution, and let Op be an operation such that $Op \notin FG(E)$, then Op cannot return true in E .*

Proof According to the lemma assumption, there is some operation Op_1 such that: $Op_1 \xrightarrow{E} Op$. There are two possibilities to consider.

- If Op_1 returned *true*, then it must have set its $mark$ flag on line 4 before Op started, and so Op reads *true* from that flag and returns *false* in line 3;
- If Op_1 returns *false*, then Op_1 reads *true* from some $mark$ flags, and so would Op .

Q.E.D.

Lemma 4.9 *The code shown in Figure 2 correctly implements a First-Generation object, and has memory steps complexity of $\Theta(\sqrt{n})$.*

Proof Correctness stems from Lemmas 4.7, 4.8. As for the memory steps complexity, we need to consider the worst case numbers of first-access shared memory events and stalls incurred by the *First* operation.

- The write contention of the protocol in Figure 2 is \sqrt{n} ; as the code performs at most a *single* write, in line 4, it incurs no more than $\sqrt{n} - 1$ stalls;
- A high level operation Op performs at most \sqrt{n} iterations of the loop in lines 1-3, and so performs at most \sqrt{n} first-access reads.¹¹

Q.E.D.

5. DISCUSSION AND FURTHER RESEARCH

This paper introduces the operation-valency technique and the influence metric for reasoning about multi-valued protocols, and uses them to obtain \sqrt{n} time lower bounds for a broad class of objects. The time metric we use - *memory steps* - is similar to the communication-cost metric and the remote-memory-references metric used by [1; 2; 6], in that it counts only memory references that cannot be resolved by a local cache.

¹¹If Op returns *true*, then it also performs a single write on line 5, which is not a first-access event.

We have proven an $\Omega(\sqrt{n})$ time lower bound for all objects in the *Influence*(n) class, and we have also shown that the bound is tight for some objects in it. For most of the interesting objects in *Influence*(n), however, including linearizable counters, stacks and queues, all known lock-free implementations require $\Omega(n)$ time. Note that for linearizable objects such as these, differently from the *First-Generation* object, there's generally a requirement of *distinctness* - i.e. there are scenarios in which all n high-level operations are required to return distinct values. Finding the tight time complexity for this class of objects remains an interesting open problem.

It would also be interesting to see whether our $\Omega(\sqrt{n})$ lower bound for a single operation holds also for the protocol's amortized complexity, possibly by showing that it holds for $\Omega(n)$ different operations.

6. ACKNOWLEDGEMENTS

We would like to thank Faith Fich and Ori Shalev for helpful comments. Comments of the anonymous PODC referees were also very helpful.

REFERENCES

- [1] ANDERSON, AND YANG. Time/contention trade-offs for multiprocessor synchronization. *IN-FCRTL: Information and Computation (formerly Information and Control) 124* (1996).
- [2] ANDERSON, J., AND KIM, Y. An improved lower bound for the time complexity of mutual exclusion, 2001.
- [3] BUSCH, C., HARDAVELLAS, N., AND MAVRONICOLAS, M. Contention in counting networks. In *Symposium on Principles of Distributed Computing* (1994), p. 404.
- [4] BUSCH, C., AND MAVRONICOLAS, M. An efficient counting network. In *Proceedings of the 1st Merged International Parallel Processing Symposium and Symposium on Parallel and Distributed Processing (IPPS/SPDP'98)* (1998), pp. 380–385.
- [5] CHOR, B., ISRAELI, A., AND LI, M. On processor coordination using asynchronous hardware. In *Proceedings of the 6th ACM Symposium on Principles of Distributed Computing (PODC)* (New York, NY, 1987), ACM Press, pp. 86–97.
- [6] CYPHER, R. The communication requirements of mutual exclusion. In *ACM Proceedings of the Seventh Annual Symposium on Parallel Algorithms and Architectures* (1995), pp. 147–156.
- [7] DWORK, C., HERLIHY, M., AND WAARTS, O. Contention in shared memory algorithms. *Journal of the ACM (JACM) 44*, 6 (1997), 779–805.
- [8] FICH, F. E., AND RUPPERT, E. Lower bounds in distributed computing. In *International Symposium on Distributed Computing* (2000), pp. 1–28.
- [9] FISCHER, M. J., LYNCH, N. A., AND PATERSON, M. S. Impossibility of distributed consensus with one faulty process. *Journal of the ACM 32*, 2 (April 1985), 374–382.
- [10] HERLIHY, M. Wait-free synchronization. *ACM Transactions On Programming Languages and Systems 13*, 1 (Jan. 1991), 123–149.
- [11] HERLIHY, M., SHAVIT, N., AND WAARTS., O. Linearizable counting networks. *Distributed Computing 9*, 4 (April 1996), 193–203.
- [12] JAYANTI, P. A time complexity lower bound for randomized implementations of some shared objects. In *Symposium on Principles of Distributed Computing* (1998), pp. 201–210.
- [13] LOUI, M. C., AND ABU-AMARA., H. H. Memory requirements for agreement among unreliable asynchronous processes. *Advances in Computing Research 4* (1987), 163–183.
- [14] LYNCH, N. A hundred impossibility proofs for distributed computing. In *Proceedings of the 8th ACM Symposium on Principles of Distributed Computing (PODC)* (New York, NY, 1989), ACM Press, pp. 1–28.
- [15] MERRIT, M., AND TAUBENFELD, G. Knowledge in shared memory systems. In *ACM Symp. on Principles of Distributed Computing* (1991), pp. 189–200.