

MIT Open Access Articles

Are lock-free concurrent algorithms practically wait-free?

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Dan Alistarh, Keren Censor-Hillel, and Nir Shavit. 2014. Are lock-free concurrent algorithms practically wait-free?. In Proceedings of the 46th Annual ACM Symposium on Theory of Computing (STOC '14). ACM, New York, NY, USA, 714-723.

As Published: <http://dx.doi.org/10.1145/2591796.2591836>

Publisher: Association for Computing Machinery (ACM)

Persistent URL: <http://hdl.handle.net/1721.1/91155>

Version: Original manuscript: author's manuscript prior to formal peer review

Terms of use: Creative Commons Attribution-Noncommercial-Share Alike



Are Lock-Free Concurrent Algorithms Practically Wait-Free?

Dan Alistarh
MIT*

Keren Censor-Hillel
Technion[†]

Nir Shavit
MIT & Tel-Aviv University[‡]

Abstract

Lock-free concurrent algorithms guarantee that *some* concurrent operation will always make progress in a finite number of steps. Yet programmers prefer to treat concurrent code as if it were *wait-free*, guaranteeing that *all* operations always make progress. Unfortunately, designing wait-free algorithms is generally a very complex task, and the resulting algorithms are not always efficient. While obtaining efficient wait-free algorithms has been a long-time goal for the theory community, most non-blocking commercial code is only lock-free.

This paper suggests a simple solution to this problem. We show that, for a large class of lock-free algorithms, under scheduling conditions which approximate those found in commercial hardware architectures, lock-free algorithms behave as if they are wait-free. In other words, programmers can keep on designing simple lock-free algorithms instead of complex wait-free ones, and in practice, they will get wait-free progress.

Our main contribution is a new way of analyzing a general class of lock-free algorithms under a *stochastic scheduler*. Our analysis relates the individual performance of processes with the global performance of the system using *Markov chain lifting* between a complex per-process chain and a simpler system progress chain. We show that lock-free algorithms are not only wait-free with probability 1, but that in fact a general subset of lock-free algorithms can be closely bounded in terms of the average number of steps required until an operation completes.

To the best of our knowledge, this is the first attempt to analyze progress conditions, typically stated in relation to a worst case adversary, in a stochastic model capturing their expected asymptotic behavior.

1 Introduction

The introduction of multicore architectures as today's main computing platform has brought about a renewed interest in concurrent data structures and algorithms, and a considerable amount of research has focused on their modeling, design and analysis.

The behavior of concurrent algorithms is captured by *safety properties*, which guarantee their correctness, and *progress properties*, which guarantee their termination. Progress properties can be quantified using two main criteria. The first is whether the algorithm is *blocking* or *non-blocking*, that is, whether the delay of a single process will cause others to be blocked, preventing them from terminating. Algorithms that use locks are blocking, while algorithms that do not use locks are non-blocking. Most of the code in the world today is lock-based, though the fraction of code without locks is steadily growing [11].

The second progress criterion, and the one we will focus on in this paper, is whether a concurrent algorithm guarantees *minimal* or *maximal progress* [12]. Intuitively, minimal progress means that *some*

*alistarh@csail.mit.edu.

[†]ckeren@cs.technion.ac.il. Shalon Fellow.

[‡]shanir@csail.mit.edu.

process is always guaranteed to make progress by completing its operations, while maximal progress means that *all* processes always complete all their operations.

Most non-blocking commercial code is *lock-free*, that is, provides minimal progress without using locks [6, 12]. Most blocking commercial code is *deadlock-free*, that is, provides minimal progress when using locks. Over the years, the research community has devised ingenious, technically sophisticated algorithms that provide *maximal progress*: such algorithms are either *wait-free*, i.e. provide maximal progress without using locks [9], or *starvation-free* [15], i.e. provide maximal progress when using locks. Unexpectedly, maximal progress algorithms, and wait-free algorithms in particular, are not being adopted by practitioners, despite the fact that the completion of all method calls in a program is a natural assumption that programmers implicitly make.

Recently, Herlihy and Shavit [12] suggested that perhaps the answer lies in a surprising property of lock-free algorithms: in practice, they often behave as if they were wait-free (and similarly, deadlock-free algorithms behave as if they were starvation-free). Specifically, most operations complete in a timely manner, and the impact of long worst-case executions on performance is negligible. In other words, in real systems, the scheduler that governs the threads' behavior in long executions does not single out any particular thread in order to cause the theoretically possible bad behaviors. This raises the following question: could the choice of *wait-free* versus *lock-free* be based simply on what assumption a programmer is willing to make about the underlying scheduler, and, with the right kind of scheduler, one will not need wait-free algorithms except in very rare cases?

This question is important because the difference between a wait-free and a lock-free algorithm for any given problem typically involves the introduction of specialized “helping” mechanisms [9], which significantly increase the complexity (both the design complexity and time complexity) of the solution. If one could simply rely on the scheduler, adding a helping mechanism to guarantee wait-freedom (or starvation-freedom) would be unnecessary.

Unfortunately, there is currently no analytical framework which would allow answering the above question, since it would require predicting the behavior of a concurrent algorithm over long executions, under a scheduler that is not adversarial.

Contribution. In this paper, we take a first step towards such a framework. Following empirical observations, we introduce a *stochastic scheduler* model, and use this model to predict the long-term behavior of a general class of concurrent algorithms. The stochastic scheduler is similar to an adversary: at each time step, it picks some process to schedule. The main distinction is that, in our model, the scheduler's choices *contain some randomness*. In particular, a stochastic scheduler has a probability threshold $\theta > 0$ such that every (non-faulty) process is scheduled with probability at least θ in each step.

We start from the following observation: under *any stochastic scheduler*, every *bounded lock-free* algorithm is actually *wait-free with probability 1*. (A *bounded lock-free* algorithm guarantees that *some* process always makes progress within a finite progress bound.) In other words, for any such algorithm, the schedules which prevent a process from ever making progress must have probability mass 0. The intuition is that, with probability 1, each specific process eventually takes enough consecutive steps, implying that it completes its operation. This observation generalizes to any bounded minimal/maximal progress condition [12]: we show that under a stochastic scheduler, bounded minimal progress becomes maximal progress, with probability 1. However, this intuition is insufficient for explaining why lock-free data structures are *efficient* in practice: because it works for arbitrary algorithms, the upper bound it yields on the number of steps until an operation completes is unacceptably high.

Our main contribution is analyzing a general class of lock-free algorithms under a specific stochastic scheduler, and showing that not only are they wait-free with probability 1, but that in fact they provide a

pragmatic bound on the number of steps until each operation completes.

We address a refined *uniform* stochastic scheduler, which schedules each non-faulty process with uniform probability in every step. Empirical data suggests that, in the long run, the uniform stochastic scheduler is a reasonable approximation for a real-world scheduler (see Figures 3 and 4). We emphasize that we do not claim real schedulers are uniform stochastic, but only that such a scheduler gives a good approximation of what happens in practice for our complexity measures, over long executions.

We call the algorithmic class we analyze *single compare-and-swap universal* (SCU). An algorithm in this class is divided into a *preamble*, and a *scan-and-validate* phase. The preamble executes auxiliary code, such as local updates and memory allocation. In the second phase, the process first determines the data structure state by scanning the memory. It then locally computes the updated state after its method call would be performed, and attempts to commit this state to memory by performing an atomic *compare-and-swap* (CAS) operation. If the CAS operation succeeds, then the state has been updated, and the method call completes. Otherwise, if some other process changes the state in between the scan and the attempted update, then the CAS operation fails, and the process must restart its operation.

This algorithmic class is widely used to design lock-free data structures. It is known that every sequential object has a lock-free implementation in this class using a lock-free version of Herlihy’s universal construction [9]. Instances of this class are used to obtain efficient data structures such as stacks [21], queues [17], or hash tables [6]. The read-copy-update (RCU) [7] synchronization mechanism employed by the Linux kernel is also an instance of this pattern.

We examine the class *SCU* under a uniform stochastic scheduler, and first observe that, in this setting, every such algorithm behaves as a Markov chain. The computational cost of interest is *system steps*, i.e. shared memory accesses by the processes. The complexity metrics we analyze are *individual latency*, which is the expected number of steps of the system until a specific process completes a method call, and *system latency*, which is the expected number of steps of the system to complete *some* method call. We bound these parameters by studying the stationary distribution of the Markov chain induced by the algorithm.

We prove two main results. The first is that, in this setting, all algorithms in this class have the property that the individual latency of any process is n times the system latency. In other words, the expected number of steps for any two processes to complete an operation is *the same*; moreover, the expected number of steps for the system to complete any operation is the expected number of steps for a specific process to complete an operation, divided by n . The second result is an upper bound of $O(q + s\sqrt{n})$ on the system latency, where q is the number of steps in the preamble, s is the number of steps in the scan-and-validate phase, and n is the number of processes. This bound is asymptotically tight.

The key mathematical tool we use is *Markov chain lifting* [3, 8]. More precisely, for such algorithms, we prove that there exists a function which *lifts* the complex Markov chain induced by the algorithm to a simplified *system* chain. The asymptotics of the system latency can be determined directly from the minimal progress chain. In particular, we bound system latency by characterizing the behavior of a new type of *iterated* balls-into-bins game, consisting of iterations which end when a certain condition on the bins first occurs, after which some of the bins change their state and a new iteration begins. Using the lifting, we prove that the individual latency is always n times the system latency.

In summary, our analysis shows that, under an approximation of the real-world scheduler, a large class of lock-free algorithms provide virtually the same progress guarantees as wait-free ones, and that, roughly, the system completes requests at a rate that is n times that of individual processes. More generally, it provides for the first time an analytical framework for predicting the behavior of a class of concurrent algorithms, over long executions, under a scheduler that is not adversarial.

Related work. To the best of our knowledge, the only prior work which addresses a probabilistic sched-

uler for a shared memory environment is that of Aspnes [2], who gave a fast consensus algorithm under a probabilistic scheduler model different from the one considered in this paper. The observation that many lock-free algorithms behave as wait-free in practice was made by Herlihy and Shavit in the context of formalizing minimal and maximal progress conditions [12], and is well-known among practitioners. For example, reference [1, Figure 6] gives empirical results for the latency distribution of individual operations of a lock-free stack. Recent work by Petrank and Timnat [20] states that most known lock-free algorithms can be written in a canonical form, which is similar to the class *SCU*, but more complex than the pattern we consider. Significant research interest has been dedicated to transforming obstruction-free or lock-free algorithms to wait-free ones, e.g. [14, 20], while minimizing performance overhead. In particular, an efficient strategy has been to divide the algorithm into a lock-free *fast path*, and a wait-free *backup path*, which is invoked if an operation fails repeatedly. Our work does not run contrary to this research direction, since the progress guarantees we prove are only probabilistic. Instead, it could be used to bound the cost of the backup path during the execution.

Roadmap. We describe the model, progress guarantees, and complexity metrics in Section 2. In particular, Section 2.3 defines the stochastic scheduler. We show that minimal progress becomes maximal progress with probability 1 in Section 4. Section 5 defines the class *SCU*(q, s), while Section 6.1 analyzes individual and global latency. The Appendix contains empirical justification for the model, and a comparison between the predicted behavior of an algorithm and its practical performance.

2 System Model

2.1 Preliminaries

Processes and Objects. We consider a shared-memory model, in which n processes p_1, \dots, p_n , communicate through registers, on which they perform atomic read, write, and compare-and-swap (CAS) operations. A CAS operation takes three arguments $(R, expVal, newVal)$, where R is the register on which it is applied, $expVal$ is the expected value of the register, and $newVal$ is the new value to be written to the register. If $expVal$ matches the value of R , then we say that the CAS is successful, and the value of R is updated to $newVal$. Otherwise, the CAS fails. The operation returns *true* if it is successful, and *false* otherwise.

We assume that each process has a unique identifier. Processes follow an algorithm, composed of shared-memory steps and local computation. The order of process steps is controlled by the *scheduler*. A set of at most $n - 1$ processes may fail by crashing. A crashed process stops taking steps for the rest of the execution. A process that is not crashed at a certain step is *correct*, and if it never crashes then it takes an infinite number of steps in the execution.

The algorithms we consider are implementations of shared objects. A shared object O is an abstraction providing a set of *methods* M , each given by its sequential specification. In particular, an implementation of a method m for object O is a set of n algorithms, one for each executing process. When process p_i invokes method m of object O , it follows the corresponding algorithm until it receives a response from the algorithm. In the following, we do not distinguish between a method m and its implementation. A method invocation is *pending* if it has not received a response. A method invocation is *active* if it is made by a *correct* process (note that the process may still crash in the future).

Executions, Schedules, and Histories. An execution is a sequence of operations performed by the processes. To represent executions, we assume discrete time, where at every time unit only one process is scheduled. In a time unit, a process can perform any number of local computations or coin flips, after which it issues a *step*, which consists of a single shared memory operation. Whenever a process becomes active,

as decided by the scheduler, it performs its local computation and then executes a step. The *schedule* is a (possibly infinite) sequence of process identifiers. If process p_i is in position $\tau \geq 1$ in the sequence, then p_i is active at time step τ .

Raising the level of abstraction, we define a *history* as a finite sequence of method invocation and response events. Notice that each schedule has a corresponding history, in which individual process steps are mapped to method calls. On the other hand, a history can be the image of several schedules.

2.2 Progress Guarantees

We now define minimal and maximal progress guarantees. We partly follow the unified presentation from [12], except that we do not specify progress guarantees for each method of an object. Rather, for ease of presentation, we adopt the simpler definition which specifies progress provided by an implementation. Consider an execution e , with the corresponding history H_e . An implementation of an object O provides *minimal progress* in the execution e if, in every suffix of H_e , some pending active instance of some method has a matching response. Equivalently, there is no point in the corresponding execution from which all the processes take an infinite number of steps without returning from their invocation.

An implementation provides *maximal progress* in an execution e if, in every suffix of the corresponding history H_e , every pending active invocation of a method has a response. Equivalently, there is no point in the execution from which a process takes infinitely many steps without returning.

Scheduler Assumptions. We say that an execution is *crash-free* if each process is always correct, i.e. if each process takes an infinite number of steps. An execution is *uniformly isolating* if, for every $k > 0$, every correct process has an interval where it takes at least k consecutive steps.

Progress. An implementation is *deadlock-free* if it guarantees minimal progress in every crash-free execution, and maximal progress in some crash-free execution.¹ An implementation is *starvation-free* if it guarantees maximal progress in every crash-free execution. An implementation is *clash-free* if it guarantees minimal progress in every uniformly isolating history, and maximal progress in some such history [12]. An implementation is *obstruction-free* if it guarantees maximal progress in every uniformly isolating execution². An implementation is *lock-free* if it guarantees minimal progress in every execution, and maximal progress in some execution. An implementation is *wait-free* if it guarantees maximal progress in every execution.

Bounded Progress. While the above definitions provide reasonable measures of progress, often in practice more explicit progress guarantees may be desired, which provide an upper bound on the number of steps until some method makes progress. To model this, we say that an implementation guarantees *bounded minimal progress* if there exists a bound $B > 0$ such that, for any time step t in the execution e at which there is an active invocation of some method, some invocation of a method returns within the next B steps by all processes. An implementation guarantees *bounded maximal progress* if there exists a bound $B > 0$ such that every active invocation of a method returns within B steps by all processes. We can specialize the definitions of bounded progress guarantees to the scheduler assumptions considered above to obtain definitions for *bounded deadlock-freedom*, *bounded starvation-freedom*, and so on.

¹According to [12], the algorithm is required to guarantee maximal progress in some execution to rule out pathological cases where a thread locks the object and never releases the lock.

²This is the definition of obstruction freedom from [12]; it is weaker than the one in [10] since it assumes uniformly isolating schedules only, but we use it here as it complies with our requirements of providing maximal progress.

2.3 Stochastic Schedulers

We define a stochastic scheduler as follows.

Definition 1 (Stochastic Scheduler). *For any $n \geq 0$, a scheduler for n processes is defined by a triple $(\Pi_\tau, A_\tau, \theta)$. The parameter $\theta \in [0, 1]$ is the threshold. For each time step $\tau \geq 1$, Π_τ is a probability distribution for scheduling the n processes at τ , and A_τ is the subset of possibly active processes at time step τ . At time step $\tau \geq 1$, the distribution Π_τ gives, for every $i \in \{1, \dots, n\}$ a probability γ_τ^i , with which process p_i is scheduled. The distribution Π_τ may depend on arbitrary outside factors, such as the current state of the algorithm being scheduled. A scheduler $(\Pi_\tau, A_\tau, \theta)$ is stochastic if $\theta > 0$. For every $\tau \geq 1$, the parameters must ensure the following:*

1. (Well-formedness) $\sum_{i=1}^n \gamma_\tau^i = 1$;
2. (Weak Fairness) For every process $p_i \in A_\tau$, $\gamma_\tau^i \geq \theta$;
3. (Crashes) For every process $p_i \notin A_\tau$, $\gamma_\tau^i = 0$;
4. (Crash Containment) $A_{\tau+1} \subseteq A_\tau$.

The well-formedness condition ensures that some process is always scheduled. Weak fairness ensures that, for a stochastic scheduler, possibly active processes do get scheduled with some non-zero probability. The crash condition ensures that failed processes do not get scheduled. The set $\{p_1, p_2, \dots, p_n\} \setminus A_\tau$ can be seen as the set of crashed processes at time step τ , since the probability of scheduling these processes at every subsequent time step is 0.

An Adversarial Scheduler. Any classic asynchronous shared memory adversary can be modeled by “encoding” its adversarial strategy in the probability distribution Π_τ for each step. Specifically, given an algorithm A and a worst-case adversary \mathcal{A}_A for A , let p_i^τ be the process that is scheduled by \mathcal{A}_A at time step τ . Then we give probability 1 in Π_τ to process p_i^τ , and 0 to all other processes. Things are more interesting when the threshold θ is strictly more than 0, i.e., there is some randomness in the scheduler’s choices.

The Uniform Stochastic Scheduler. A natural scheduler is the *uniform* stochastic scheduler, for which, assuming no process crashes, we have that Π_τ has $\gamma_i^\tau = 1/n$, for all i and $\tau \geq 1$, and $A_\tau = \{1, \dots, n\}$ for all time steps $\tau \geq 1$. With crashes, we have that $\gamma_i^\tau = 1/|A_\tau|$ if $i \in A_\tau$, and $\gamma_i^\tau = 0$ otherwise.

2.4 Complexity Measures

Given a concurrent algorithm, standard analysis focuses on two measures: *step complexity*, the worst-case number of steps performed by a single process in order to return from a method invocation, and *total step complexity*, or *work*, which is the worst-case number of system steps required to complete invocations of all correct processes when performing a task together. In this paper, we focus on the analogue of these complexity measures for long executions. Given a stochastic scheduler, we define (*average*) *individual latency* as the maximum over all inputs of the expected number of steps taken by the system between the returns times of two consecutive invocations of the same process. Similarly, we define the (*average*) *system latency* as the maximum over all inputs of the expected number of system steps between consecutive returns times of any two invocations.

3 Background on Markov Chains

We now give a brief overview of Markov chains. Our presentation follows standard texts, e.g. [16, 18]. The definition and properties of Markov chain lifting are adapted from [8].

Given a set S , a sequence of random variables $(X_t)_{t \in \mathbb{N}}$, where $X_t \in S$, is a (discrete-time) *stochastic process* with states in S . A *discrete-time Markov chain* over the state set S is a discrete-time stochastic process with states in S that satisfies the *Markov condition*

$$\Pr[X_t = i_t | X_{t-1} = i_{t-1}, \dots, X_0 = i_0] = \Pr[X_t = i_t | X_{t-1} = i_{t-1}].$$

The above condition is also called the *memoryless property*. A Markov chain is *time-invariant* if the equality $\Pr[X_t = j | X_{t-1} = i] = \Pr[X_{t'} = j | X_{t'-1} = i]$ holds for all times $t, t' \in \mathbb{N}$ and all $i, j \in S$. This allows us to define the *transition matrix* P of a Markov chain as the matrix with entries

$$p_{ij} = \Pr[X_t = j | X_{t-1} = i].$$

The *initial distribution* of a Markov chain is given by the probabilities $\Pr[X_0 = i]$, for all $i \in S$. We denote the time-invariant Markov chain X with initial distribution λ and transition matrix P by $M(P, \lambda)$.

The random variable $T_{ij} = \min\{n \geq 1 | X_n = j, \text{ if } X_0 = i\}$ counts the number of steps needed by the Markov chain to get from i to j , and is called the *hitting time* from i to j . We set $T_{i,j} = \infty$ if state j is unreachable from i . Further, we define $h_{ij} = E[T_{ij}]$, and call $h_{ii} = E[T_{ii}]$ the (expected) return time for state $i \in S$.

Given P , the transition matrix of $M(P, \lambda)$, a *stationary distribution* of the Markov chain is a state vector π with $\pi = \pi P$. (We consider *row* vectors throughout the paper.) The intuition is that if the state vector of the Markov chain is π at time t , then it will remain π for all $t' > t$. Let $P^{(k)}$ be the transition matrix P multiplied by itself k times, and $p_{ij}^{(k)}$ be element (i, j) of $P^{(k)}$. A Markov chain is *irreducible* if for all pairs of states $i, j \in S$ there exists $m \geq 0$ such that $p_{ij}^{(m)} > 0$. (In other words, the underlying graph is strongly connected.) This implies that $T_{ij} < \infty$, and all expectations h_{ij} exist, for all $i, j \in S$. Furthermore, the following is known.

Theorem 1. *An irreducible finite Markov chain has a unique stationary distribution π , namely*

$$\pi_j = \frac{1}{h_{jj}}, \forall j \in S.$$

The periodicity of a state j is the maximum positive integer α such that $\{n \in \mathbb{N} | p_{jj}^{(n)} > 0\} \subseteq \{i\alpha | i \in \mathbb{N}\}$. A state with periodicity $\alpha = 1$ is called *aperiodic*. A Markov chain is *aperiodic* if all states are aperiodic. If a Markov chain has at least one self-loop, then it is aperiodic. A Markov chain that is irreducible and aperiodic is *ergodic*. Ergodic Markov chains converge to their stationary distribution as $t \rightarrow \infty$ independently of their initial distributions.

Theorem 2. *For every ergodic finite Markov chain $(X_t)_{t \in \mathbb{N}}$ we have independently of the initial distribution that $\lim_{t \rightarrow \infty} q_t = \pi$, where π denotes the chain's unique stationary distribution, and q_t is the distribution on states at time $t \in \mathbb{N}$.*

Ergodic Flow. It is often convenient to describe an ergodic Markov chain in terms of its *ergodic flow*: for each (directed) edge ij , we associate a flow $Q_{ij} = \pi_i p_{ij}$. These values satisfy $\sum_i Q_{ij} = \sum_i Q_{ji}$ and $\sum_{i,j} Q_{ij} = 1$. It also holds that $\pi_j = \sum_i Q_{ij}$.

Lifting Markov Chains. Let M and M' be ergodic Markov chains on finite state spaces S, S' , respectively. Let P, π be the transition matrix and stationary distribution for M , and P', π' denote the corresponding objects for M' . We say that M' is a *lifting* of M [8] if there is a function $f : S' \rightarrow S$ such that

$$Q_{ij} = \sum_{x \in f^{-1}(i), y \in f^{-1}(j)} Q'_{xy}, \forall i, j \in S.$$

Informally, M' is collapsed onto M by clustering several of its states into a single state, as specified by the function f . The above relation specifies a homomorphism on the ergodic flows. An immediate consequence of this relation is the following connection between the stationary distributions of the two chains.

Lemma 1. *For all $v \in S$, we have that*

$$\pi(v) = \sum_{x \in f^{-1}(v)} \pi'(x).$$

4 From Minimal Progress to Maximal Progress

We now formalize the intuition that, under a stochastic scheduler, all algorithms ensuring bounded minimal progress guarantee in fact maximal progress with probability 1. We also show the *bounded* minimal progress assumption is necessary: if minimal progress is not bounded, then maximal progress may not be achieved.

Theorem 3 (Min to Max Progress). *Let \mathcal{S} be a stochastic scheduler with probability threshold $1 \geq \theta > 0$. Let A be an algorithm ensuring bounded minimal progress with a bound T . Then A ensures maximal progress with probability 1. Moreover, the expected maximal progress bound of A is at most $(1/\theta)^T$.*

Proof. Consider an interval of T steps in an execution of algorithm A . Our first observation is that, since A ensures T -bounded minimal progress, any process that performs T consecutive steps in this interval must complete a method invocation. To prove this fact, we consider cases on the minimal progress condition. If the minimal progress condition is T -bounded *deadlock-freedom* or *lock-freedom*, then every sequence of T steps by the algorithm must complete some method invocation. In particular, T steps by a single process must complete a method invocation. Obviously, this completed method invocation must be by the process itself. If the progress condition is T -bounded *clash-freedom*, then the claim follows directly from the definition.

Next, we show that, since \mathcal{S} is a stochastic scheduler with positive probability threshold, each correct process will eventually be scheduled for T consecutive steps, with probability 1. By the weak fairness condition in the definition, for every time step τ , every active process $p_i \in A_\tau$ is scheduled with probability at least $\theta > 0$. A process p_i is *correct* if $p_i \in A_\tau$, for all $\tau \geq 1$. By the definition, at each time step τ , each correct process $p_i \in A_\tau$ is scheduled for T consecutive time units with probability at least $\theta^T > 0$. From the previous argument, it follows that every correct process eventually completes each of its method calls with probability 1. By the same argument, the expected completion time for a process is at most $(1/\theta)^T$. \square

The proof is based on the fact that, for every correct process p_i , eventually, the scheduler will produce a solo a schedule of length T . On the other hand, since the algorithm ensures minimal progress with bound T , we show that p_i must complete its operation during this interval.

We then prove that the finite bound for minimal progress is necessary. For this, we devise an *unbounded* lock-free algorithm which is not wait-free with probability > 0 . The main idea is to have processes that fail to change the value of a CAS repeatedly increase the number of steps they need to take to complete an operation. (See Algorithm 1.)

Lemma 2. *There exists an unbounded lock-free algorithm that is not wait-free with high probability.*

Proof. Consider the initial state of Algorithm 1. With probability at least $1/n$, each process p_i can be the first process to take a step, performing a successful CAS operation. Assume process p_1 takes the first step.

```

1 Shared: CAS object  $C$ , initially 0;
2 Register  $R$ 
3 Local: Integers  $v, val, j$ , initially 0
4 while true do
5      $val \leftarrow CAS(C, v, v + 1)$ 
6     if  $val = v$  then return;
7     else
8          $v \leftarrow val$ 
9     for  $j = 1 \dots n^2 v$  do  $read(R)$ ;

```

Algorithm 1: An unbounded lock-free algorithm.

```

10 Shared: registers  $R, R_1, R_2, \dots, R_{s-1}$ 
11 procedure method-call()
12 Take preamble steps  $O_1, O_2, O_q$  /* Preamble region */
13 while true do
14     /* Scan region: */
15      $v \leftarrow R.read()$ 
16      $v_1 \leftarrow R_1.read(); v_2 \leftarrow R_2.read(); \dots; v_{s-1} \leftarrow R_{s-1}.read()$ 
17      $v' \leftarrow$  new proposed state based on  $v, v_1, v_2, \dots, v_{s-1}$ 
18     /* Validation step: */
19      $flag \leftarrow CAS(R, v, v')$ 
20     if  $flag = true$  then
21         output  $success$ 

```

Algorithm 2: The structure of the lock-free algorithms in $SCU_{q,s}$.

Conditioned on this event, let P be the probability that p_1 is not the next process that performs a successful CAS operation. If p_1 takes a step in any of the next $n^2 \cdot v$ steps, then it is the next process that wins the CAS. The probability that this does not happen is at most $(1 - 1/n)^{n^2}$. Summing over all iterations, the probability that p_1 ever performs an unsuccessful CAS is therefore at most $\sum_{\ell=1}^{\infty} (1 - 1/n)^{n^2 \cdot \ell} \leq 2(1 - 1/n)^{n^2} \leq 2e^{-n}$. Hence, with probability at least $1 - 2e^{-n}$, process p_1 always wins the CAS, while other processes never do. This implies that the algorithm is not wait-free, with high probability. \square

5 The Class of Algorithms $SCU(q, s)$

In this section, we define the class of algorithms $SCU(q, s)$. An algorithm in this class is structured as follows. (See Algorithm 2 for the pseudocode.) The first part is the *preamble*, where the process performs a series of q steps. The algorithm then enters a *loop*, divided into a *scan* region, which reads the values of s registers, and a *validation* step, where the process performs a CAS operation, which attempts to change the value of a register. The of the scan region is to obtain a view of the data structure state. In the validation step, the process checks that this state is still valid, and attempts to change it. If the CAS is successful, then the operation completes. Otherwise, the process restarts the loop. We say that an algorithm with the above structure with parameters q and s is in $SCU(q, s)$.

We assume that steps in the preamble may perform memory updates, including to registers R_1, \dots, R_{s-1} , but do not change the value of the decision register R . Also, two processes never propose the same value for

the register R . (This can be easily enforced by adding a timestamp to each request.) The order of steps in the scan region can be changed without affecting our analysis. Such algorithms are used in several CAS-based concurrent implementations. In particular, the class can be used to implement a concurrent version of every sequential object [9]. It has also been used to obtain efficient implementations of several concurrent objects, such as fetch-and-increment [4], stacks [21], and queues [17].

6 Analysis of the Class $SCU(q, s)$

We analyze the performance of algorithms in $SCU(q, s)$ under the uniform stochastic scheduler. We assume that all threads execute the same method call with preamble of length q , and scan region of length s . Each thread executes an infinite number of such operations. To simplify the presentation, we assume all n threads are correct in the analysis. The claim is similar in the crash-failure case, and will be considered separately.

We examine two parameters: system latency, i.e., how often (in terms of system steps) does a new operation complete, and individual latency, i.e., how often does *a certain thread* complete a new operation. Notice that the worst-case latency for the whole system is $\Theta(q + sn)$ steps, while the worst-case latency for an individual thread is ∞ , as the algorithm is not wait-free. We will prove the following result:

Theorem 4. *Let A be an algorithm in $SCU(q, s)$. Then, under the uniform stochastic scheduler, the system latency of A is $O(q + s\sqrt{n})$, and the individual latency is $O(n(q + s\sqrt{n}))$.*

We prove the upper bound by splitting the class $SCU(q, s)$ into two separate components, and analyzing each under the uniform scheduler. The first part is the loop code, which we call the *scan-validate* component. The second part is the *parallel code*, which we use to characterize the performance of the preamble code. In other words, we first consider $SCU(0, s)$ and then $SCU(q, 0)$.

6.1 The Scan-Validate Component

Notice that, without loss of generality, we can simplify the pseudocode to contain a single read step before the CAS. We obtain the performance bounds for this simplified algorithm, and then multiply them by s , the number of scan steps. That is, we start by analyzing $SCU(0, 1)$ and then generalize to $SCU(0, s)$.

Proof Strategy. We start from the Markov chain representation of the algorithm, which we call the *individual chain*. We then focus on a simplified representation, which only tracks *system-wide progress*, irrespective of which process is exactly in which state. We call this the *system chain*. We first prove the individual chain can be related to the system chain via a lifting function, which allows us to relate the individual latency to the system latency (Lemma 5). We then focus on bounding system latency. We describe the behavior of the system chain via an iterated balls-and-bins game, whose stationary behavior we analyze in Lemmas 8 and 9. Finally, we put together these claims to obtain an $O(\sqrt{n})$ upper bound on the system latency of $SCU(0, 1)$.

6.1.1 Markov Chain Representations

We define the *extended local state* of a process in terms of the state of the system, and of the type of step it is about to take. Thus, a process can be in one of three states: either it performs a read, or it CAS-es with the current value of R , or it CAS-es with an invalid value of R . The state of the system after each step is completely described by the n extended local states of processes. We emphasize that this is different than what is typically referred to as the “local” state of a process, in that the extended local state is described from the viewpoint of the entire system. That is, a process that has a pending CAS operation can be in either

```

1 Shared: register  $R$ 
2 Local:  $v$ , initially  $\perp$ 
3 procedure scan-validate()
4 while true do
5    $v \leftarrow R.read()$ ;  $v' \leftarrow$  new value based on  $v$ 
6    $flag \leftarrow CAS(R, v, v')$ 
7   if  $flag = \text{true}$  then
8     output  $success$ 

```

Algorithm 3: The scan-validate pattern.

of two different extended local states, depending on whether its CAS will succeed or not. This is determined by the state of the entire system. A key observation is that, although the “local” state of a process can only change when it takes a step, its extended local state can change also when another process takes a step.

The individual chain. Since the scheduler is uniform, the system can be described as a Markov chain, where each state specifies the extended local state of each process. Specifically, a process is in state *OldCAS* if it is about to CAS with an old (invalid) value of R , it is in state *Read* if it is about to read, and is in state *CCAS* if it about to CAS with the current value of R . (Once CAS-ing the process returns to state *Read*.)

A state S of the individual chain is given by a combination of n states $S = (P_1, P_2, \dots, P_n)$, describing the extended local state of each process, where, for each $i \in \{1, \dots, n\}$, $P_i \in \{OldCAS, Read, CCAS\}$ is the extended local state of process p_i . There are $3^n - 1$ possible states, since the state where each process CAS-es with an old value cannot occur. In each transition, each process takes a step, and the state changes correspondingly. Recall that every process p_i takes a step with probability $1/n$. Transitions are as follows. If the process p_i taking a step is in state *Read* or *OldCAS*, then all other processes remain in the same extended local state, and p_i moves to state *CCAS* or *Read*, respectively. If the process p_i taking a step is in state *CCAS*, then all processes in state *CCAS* move to state *OldCAS*, and p_i moves to state *Read*.

The system chain. To reduce the complexity of the individual Markov chain, we introduce a simplified representation, which tracks system-wide progress. More precisely, each state of the system chain tracks the number of processes in each state, irrespective of their identifiers: for any $a, b \in \{0, \dots, n\}$, a state x is defined by the tuple (a, b) , where a is the number of processes that are in state *Read*, and b is the number of processes that are in state *OldCAS*. Notice that the remaining $n - a - b$ processes must be in state *CCAS*. The initial state is $(n, 0)$, i.e. all processes are about to read. The state $(0, n)$ does not exist. The transitions in the system chain are as follows. $\Pr[(a+1, b-1)|(a, b)] = b/n$, where $0 \leq a \leq n$ and $b > 0$. $\Pr[(a+1, b)|(a, b)] = 1 - (a+b)/n$, where $0 \leq a < n$. $\Pr[(a-1, b)|(a, b)] = 1 - a/n$, where $0 < a \leq n$. (See Figure 1 for an illustration of the two chains in the two-process case.)

6.1.2 Analysis Preliminaries

First, we notice that both the individual chain and the system chain are ergodic.

Lemma 3. *For any $n \geq 1$, the individual chain and the system chain are ergodic.*

Let π be the stationary distribution of the system chain, and let π' be the stationary distribution for the individual chain. For any state $k = (a, b)$ in the system chain, let π_k be its probability in the stationary distribution. Similarly, for state x in the individual chain, let π'_x be its probability in the stationary distribution.

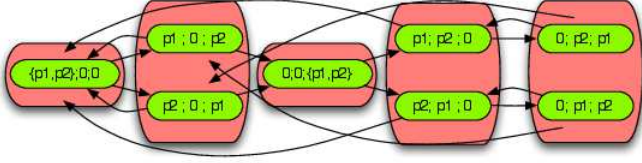


Figure 1: The individual chain and the global chain for two processes. Each transition has probability $1/2$. The red clusters are the states in the system chain. The notation $X; Y; Z$ means that processes in X are in state *Read*, processes in Y are in state *OldCAS*, and processes in Z are in state *CCAS*.

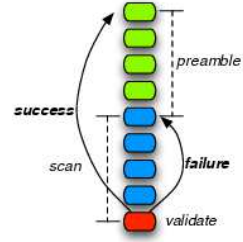


Figure 2: Structure of an algorithm in $SCU(q, s)$.

We now prove that there exists a *lifting* from the individual chain to the system chain. Intuitively, the lifting from the individual chain to the system chain collapses all states in which a processes are about to read and b processes are about to CAS with an old value (the identifiers of these processes are different for distinct states), into to state (a, b) from the system chain.

Definition 2. Let \mathcal{S} be the set of states of the individual chain, and \mathcal{M} be the set of states of the system chain. We define the function $f : \mathcal{S} \rightarrow \mathcal{M}$ such that each state $S = (P_1, \dots, P_n)$, where a processes are in state *Read* and b processes are in state *OldCAS*, is taken into state (a, b) of the system chain.

We then obtain the following relation between the stationary distributions of the two chains.

Lemma 4. For every state k in the system chain, we have $\pi_k = \sum_{x \in f^{-1}(k)} \pi'_x$.

Proof. We obtain this relation algebraically, starting from the formula for the stationary distribution of the individual chain. We have that $\pi' A = \pi'$, where π' is a row vector, and A is the transition matrix of the individual chain. We partition the states of the individual chain into sets, where $G_{a,b}$ is the set of system states S such that $f(S) = (a, b)$. Fix an arbitrary ordering $(G_k)_{k \geq 1}$ of the sets, and assume without loss of generality that the system states are ordered according to their set in the vector π and in the matrix A , so that states mapping to the same set are consecutive.

Let now A' be the transition matrix across the sets $(G_k)_{k \geq 1}$. In particular, a'_{kj} is the probability of moving from a state in the set G_k to some state in the set G_j . Note that this transition matrix is the same as that of the system chain. Pick an arbitrary state x in the individual chain, and let $f(x) = (a, b)$. In other words, state x maps to set G_k , where $k = (a, b)$. We claim that for every set G_j , $\sum_{y \in G_j} \Pr[y|x] = \Pr[G_j|G_i]$.

To see this, fix $x = (P_0, P_1, \dots, P_n)$. Since $f(x) = (a, b)$, there are exactly b distinct states y reachable from x such that $f(y) = (a + 1, b - 1)$: the states where a process in extended local state *OldCAS* takes a step. Therefore, the probability of moving to such a state y is b/n . Similarly, the probability of moving to a state y with $f(y) = (a + 1, b - 1)$ is $1 - (a + b)/n$, and the probability of moving to a state y with $f(y) = (a - 1, b)$ is a/n . All other transition probabilities are 0.

To complete the proof, notice that we can collapse the stationary distribution π' onto the row vector $\bar{\pi}$, where the k th element of $\bar{\pi}$ is $\sum_{x \in G_k} \pi'_x$. Using the above claim and the fact that $\pi' A = \pi'$, we obtain by calculation that $\bar{\pi} A' = \bar{\pi}$. Therefore, $\bar{\pi}$ is a stationary distribution for the system chain. Since the stationary distribution is unique, $\bar{\pi} = \pi$, which concludes the proof. \square

In fact, we can prove that the function $f : \mathcal{S} \rightarrow \mathcal{M}$ defined above induces a lifting from the individual chain to the system chain.

Lemma 5. *The system Markov chain is a lifting of the individual Markov chain.*

Proof. Consider a state k in \mathcal{M} . Let j be a neighboring state of k in the system chain. The ergodic flow from k to j is $p_{kj}\pi_k$. In particular, if k is given by the tuple (a, b) , j can be either $(a + 1, b - 1)$ or $(a + 1, b)$, or $(a - 1, b)$. Consider now a state $x \in \mathcal{M}$, $x = (P_0, \dots, P_n)$, such that $f(x) = k$. By the definition of f , x has a processes in state *Read*, and b processes in state *OldCAS*.

If j is the state $(a + 1, b - 1)$, then the flow from k to j , Q_{kj} , is $b\pi_k/n$. The state x from the individual chain has exactly b neighboring states y which map to the state $(a + 1, b - 1)$, one for each of the b processes in state *OldCAS* which might take a step. Fix y to be such a state. The probability of moving from x to y is $1/n$. Therefore, using Lemma 4, we obtain that

$$\sum_{x \in f^{-1}(k), y \in f^{-1}(j)} Q'_{xy} = \sum_{x \in f^{-1}(k)} \sum_{y \in f^{-1}(j)} \frac{1}{n} \pi'_x = \frac{b}{n} \sum_{x \in f^{-1}(k)} \pi'_x = \frac{b}{n} \pi_k = Q_{kj}.$$

The other cases for state j follow similarly. Therefore, the lifting condition holds. \square

Next, we notice that, since states from the individual chain which map to the same system chain state are symmetric, their probabilities in the stationary distribution must be the same.

Lemma 6. *Let x and x' be two states in \mathcal{S} such that $f(x) = f(x')$. Then $\pi'_x = \pi'_{x'}$.*

Proof (Sketch). The proof follows by noticing that, for any $i, j \in \{1, 2, \dots, n\}$, switching indices i and j in the Markov chain representation maintains the same transition matrix. Therefore, the stationary probabilities for symmetric states (under the swapping of process ids) must be the same. \square

We then use the fact that the code is symmetric and the previous Lemma to obtain an upper bound on the expected time between two successes for a specific process.

Lemma 7. *Let W be the expected system steps between two successes in the stationary distribution of the system chain. Let W_i be the expected system steps between two successes of process p_i in the stationary distribution of the individual chain. For every process p_i , $W = nW_i$.*

Proof. Let μ be the probability that a step is a success by *some* process. Expressed in the system chain, we have that $\mu = \sum_{j=(a,b)} (1 - (a + b)/n) \pi_j$. Let X_i be the set of states in the individual chain in which $P_i = \text{CCAS}$. Consider the event that a system step is a step in which p_i succeeds. This must be a step by p_i from a state in X_i . The probability of this event in the stationary distribution of the individual chain is $\eta_i = \sum_{x \in X_i} \pi'_x/n$.

Recall that the lifting function f maps all states x with a processes in state *Read* and b processes in state *OldCAS* to state $j = (a, b)$. Therefore, $\eta_i = (1/n) \sum_{j=(a,b)} \sum_{x \in f^{-1}(j) \cap X_i} \pi'_x$. By symmetry, we have that $\pi'_x = \pi'_y$, for every states $x, y \in f^{-1}(j)$. The fraction of states in $f^{-1}(j)$ that have p_i in state *CCAS* (and are therefore also in X_i) is $(1 - (a + b)/n)$. Therefore, $\sum_{x \in f^{-1}(j) \cap X_i} \pi'_x = (1 - (a + b)/n) \pi_j$.

We finally get that, for every process p_i , $\eta_i = (1/n) \sum_{j=(a,b)} (1 - (a + b)/n) \pi_j = (1/n) \mu$. On the other hand, since we consider the stationary distribution, from a straightforward extension of Theorem 1, we have that $W_i = 1/\eta_i$, and $W = 1/\mu$. Therefore, $W_i = nW$, as claimed. \square

6.1.3 System Latency Bound

In this section we provide an upper bound on the quantity W , the expected number of system steps between two successes in stationary distribution of the system chain. We prove the following.

Theorem 5. *The expected number of steps between two successes in the system chain is $O(\sqrt{n})$.*

An iterated balls-into-bins game. To bound W , we model the evolution of the system as a balls-into-bins game. We will associate each process with a bin. At the beginning of the execution, each bin already contains one ball. At each time step, we throw a new ball into a uniformly chosen random bin. Essentially, whenever the process takes a step, its bin receives an additional ball. We continue to distribute balls until the first time a bin acquires *three* balls. We call this event a *reset*. When a reset occurs, we set the number of balls in the bin containing three balls to one, and all the bins containing two balls become empty. The game then continues until the next reset.

This game models the fact that initially, each process is about to read the shared state, and must take two steps in order to update its value. Whenever a process changes the shared state by CAS-ing successfully, all other processes which were CAS-ing with the correct value are going to fail their operations; in particular, they now need to take three steps in order to change the shared state. We therefore reset the number of balls in the corresponding bins to 0.

More precisely, we define the game in terms of *phases*. A phase is the interval between two resets. For phase i , we denote by a_i the number of bins with one ball at the beginning of the phase, and by b_i the number of bins with 0 balls at the beginning of the phase. Since there are no bins with two or more balls at the start of a phase, we have that $a_i + b_i = n$.

It is straightforward to see that this random process evolves in the same way as the system Markov chain. In particular, notice that the bound W is the expected length of a phase. To prove Theorem 5, we first obtain a bound on the length of a phase.

Lemma 8. *Let $\alpha \geq 4$ be a constant. The expected length of phase i is at most $\min(2\alpha n/\sqrt{a_i}, 3\alpha n/b_i^{1/3})$. The phase length is $2\alpha \min(n\sqrt{\log n}/\sqrt{a_i}, n(\log n)^{1/3}/b_i^{1/3})$, with probability at least $1 - 1/n^\alpha$. The probability that the length of a phase is less than $\min(n/\sqrt{a_i}, n/(b_i)^{1/3})/\alpha$ is at most $1/(4\alpha^2)$.*

Proof. Let A_i be the set of bins with one ball, and let B_i be the set of bins with zero balls, at the beginning of the phase. We have $a_i = |A_i|$ and $b_i = |B_i|$. Practically, the phase ends either when a bin in A_i or a bin in B_i first contains three balls.

For the first event to occur, some bin in A_i must receive two additional balls. Let $c \geq 1$ be a large constant, and assume for now that $a_i \geq \log n$ and $b_i \geq \log n$ (the other cases will be treated separately). The number of bins in A_i which need to receive a ball before some bin receives two new balls is concentrated around $\sqrt{a_i}$, by the birthday paradox. More precisely, the following holds.

Claim 1. *Let X_i be random variable counting the number of bins in A_i chosen to get a ball before some bin in A_i contains three balls, and fix $\alpha \geq 4$ to be a constant. Then the expectation of X_i is less than $2\alpha\sqrt{a_i}$. The value of X_i is at most $\alpha\sqrt{a_i \log n}$, with probability at least $1 - 1/n^{\alpha^2}$.*

Proof. We employ the Poisson approximation for balls-into-bins processes. In essence, we want to bound the number of balls to be thrown uniformly into a_i bins until two balls collide in the same bin, in expectation and with high probability. Assume we throw m balls into the $a_i \geq \log n$ bins. It is well-known that the number of balls a bin receives during this process can be approximated as a Poisson random variable with

mean m/a_i (see, e.g., [18]). In particular, the probability that no bin receives two extra balls during this process is at most

$$2 \left(1 - \frac{e^{-m/a_i} \left(\frac{m}{a_i}\right)^2}{2} \right)^{a_i} \leq 2 \left(\frac{1}{e} \right)^{\frac{m^2}{2a_i}} e^{-m/a_i}.$$

If we take $m = \alpha\sqrt{a_i}$ for $\alpha \geq 4$ constant, we obtain that this probability is at most

$$2 \left(\frac{1}{e} \right)^{\alpha^2 e^{-\alpha/\sqrt{a_i}}/2} \leq \left(\frac{1}{e} \right)^{\alpha^2/4},$$

where we have used the fact that $a_i \geq \log n \geq \alpha^2$. Therefore, the expected number of throws until some bin receives two balls is at most $2\alpha\sqrt{a_i}$. Taking $m = \alpha\sqrt{a_i \log n}$, we obtain that some bin receives two new balls within $\alpha\sqrt{a_i \log n}$ throws with probability at least $1 - 1/n^{\alpha^2}$. \square

We now prove a similar upper bound for the number of bins in B_i which need to receive a ball before some such bin receives three new balls, as required to end the phase.

Claim 2. *Let Y_i be random variable counting the number of bins in B_i chosen to get a ball before some bin in B_i contains three balls, and fix $\alpha \geq 4$ to be a constant. Then the expectation of Y_i is at most $3\alpha b_i^{2/3}$, and Y_i is at most $\alpha(\log n)^{1/3} b_i^{2/3}$, with probability at least $1 - (1/n)^{\alpha^3/54}$.*

Proof. We need to bound the number of balls to be thrown uniformly into b_i bins (each of which is initially empty), until some bin gets three balls. Again, we use a Poisson approximation. We throw m balls into the $b_i \geq \log n$ bins. The probability that no bin receives three or more balls during this process is at most

$$2 \left(1 - \frac{e^{-m/b_i} \left(m/b_i\right)^3}{6} \right)^{b_i} = 2 \left(\frac{1}{e} \right)^{\frac{m^3}{6b_i^2}} e^{-m/b_i}.$$

Taking $m = \alpha b_i^{2/3}$ for $\alpha \geq 4$, we obtain that this probability is at most

$$2 \left(\frac{1}{e} \right)^{\frac{\alpha^3}{6} e^{-\alpha/b_i^{1/3}}} \leq \left(\frac{1}{e} \right)^{\alpha^3/54}.$$

Therefore, the expected number of ball thrown into bins from B_i until some such bin contains three balls is at most $3\alpha b_i^{2/3}$. Taking $m = \alpha(\log n)^{1/3} b_i^{2/3}$, we obtain that the probability that no bin receives three balls within the first m ball throws in B_i is at most $(1/n)^{\alpha^3/54}$. \square

The above claims bound the number of steps inside the sets A_i and B_i necessary to finish the phase. On the other hand, notice that a step throws a new ball into a bin from A_i with probability a_i/n , and throws it into a bin in B_i with probability b_i/n . It therefore follows that the expected number of steps for a bin in A_i to reach three balls (starting from one ball in each bin) is at most $2\alpha\sqrt{a_i}n/a_i = 2\alpha n/\sqrt{a_i}$. The expected number of steps for a bin in B_i to reach three balls is at most $3\alpha b_i^{2/3}n/b_i = 3\alpha n/b_i^{1/3}$. The next claim provides concentration bounds for these inequalities, and completes the proof of the Lemma.

Claim 3. *The probability that the system takes more than $2\alpha\frac{n}{\sqrt{a_i}}\sqrt{\log n}$ steps in a phase is at most $1/n^\alpha$. The probability that the system takes more than $2\alpha\frac{n}{b_i^{1/3}}(\log n)^{1/3}$ steps in a phase is at most $1/n^\alpha$.*

Proof. Fix a parameter $\beta > 0$. By a Chernoff bound, the probability that the system takes more than $2\beta n/a_i$ steps without throwing at least β balls into the bins in A_i is at most $(1/e)^\beta$. At the same time, by Claim 1, the probability that $\alpha\sqrt{a_i \log n}$ balls thrown into bins in A_i do not generate a collision (finishing the phase) is at most $1/n^{\alpha^2}$.

Therefore, throwing $2\alpha\frac{n}{\sqrt{a_i}}\sqrt{\log n}$ balls fail to finish the phase with probability at most $1/n^{\alpha^2} + 1/e^{\alpha\sqrt{a_i \log n}}$. Since $a_i \geq \log n$ by the case assumption, the claim follows.

Similarly, using Claim 2, the probability that the system takes more than $2\alpha(\log n)^{1/3}b_i^{2/3}n/b_i = 2\alpha(\log n)^{1/3}n/b_i^{1/3}$ steps without a bin in B_i reaching three balls (in the absence of a reset) is at most $(1/e)^{1+(\log n)^{1/3}b_i^{2/3}} + (1/n)^{\alpha^3/54} \leq (1/n)^\alpha$, since $b_i \geq \log n$. \square

We put these results together to obtain that, if $a_i \geq \log n$ and $b_i \geq \log n$, then the expected length of a phase is $\min(2\alpha n/\sqrt{a_i}, 3\alpha n/b_i^{1/3})$. The phase length is $2\alpha \min(\frac{n}{\sqrt{a_i}}\sqrt{\log n}, \frac{n}{b_i^{1/3}}(\log n)^{1/3})$, with high probability.

It remains to consider the case where either a_i or b_i are less than $\log n$. Assume $a_i \geq \log n$. Then $b_i \geq n - \log n$. We can therefore apply the above argument for b_i , and we obtain that with high probability the phase finishes in $2\alpha n(\log n/b_i)^{1/3}$ steps. This is less than $2\alpha\frac{n}{\sqrt{a_i}}\sqrt{\log n}$, since $a_i \leq \log n$, which concludes the claim. The converse case is similar. \square

Returning to the proof, we characterize the dynamics of the phases $i \geq 1$ based on the value of a_i at the beginning of the phase. We say that a phase i is in *the first range* if $a_i \in [n/3, n]$. Phase i is in *the second range* if $n/c \leq a_i < n/3$, where c is a large constant. Finally, phase i is in *the third range* if $0 \leq a_i < n/c$. Next, we characterize the probability of moving between phases.

Lemma 9. *For $i \geq 1$, if phase i is in the first two ranges, then the probability that phase $i + 1$ is in the third range is at most $1/n^\alpha$. Let $\beta > 2c^2$ be a constant. The probability that $\beta\sqrt{n}$ consecutive phases are in the third range is at most $1/n^\alpha$.*

Proof. We first bound the probability that a phase moves to the third range from one of the first two ranges.

Claim 4. *For $i \geq 1$, if phase i is in the first two ranges, then the probability that phase $i + 1$ is in the third range is at most $1/n^\alpha$.*

Proof. We first consider the case where phase i is in range two, i.e. $n/c \leq a_i < n/3$, and bound the probability that $a_{i+1} < n/c$. By Lemma 8, the total number of system steps taken in phase i is at most $2\alpha \min(n/\sqrt{a_i}\sqrt{\log n}, n/b_i^{1/3}(\log n)^{1/3})$, with probability at least $1 - 1/n^\alpha$. Given the bounds on a_i , it follows by calculation that the first factor is always the minimum in this range.

Let ℓ_i be the number of steps in phase i . Since $a_i \in [n/c, n/3)$, the expected number of balls thrown into bins from A_i is at most $\ell_i/3$, whereas the expected number of balls thrown into bins from B_i is at least $2\ell_i/3$. The parameter a_{i+1} is a_i plus the bins from B_i which acquire a single ball, minus the balls from A_i which acquire an extra ball. On the other hand, the number of bins from B_i which acquire a single ball during ℓ_i steps is tightly concentrated around $2\ell_i/3$, whereas the number of bins in A_i which acquire a single ball during ℓ_i steps is tightly concentrated around $\ell_i/3$. More precisely, using Chernoff bounds, given $a_i \in [n/c, n/3)$, we obtain that $a_i \geq a_{i+1}$, with probability at least $1 - 1/e^{\alpha\sqrt{n}}$.

For the case where phase i is in range one, notice that, in order to move to range three, the value of a_i would have to decrease by at least $n(1/3 - 1/c)$ in this phase. On the other hand, by Lemma 8, the length of the phase is at most $2\alpha\sqrt{3n \log n}$, w.h.p. Therefore the claim follows. A similar argument provides a lower bound on the length of a phase. \square

The second claim suggests that, if the system is in the third range (a low probability event), it gradually returns to one of the first two ranges.

Claim 5. *Let $\beta > 2c^2$ be a constant. The probability that $\beta\sqrt{n}$ phases are in the third range is at most $1/n^\alpha$.*

Proof. Assume the system is in the third range, i.e. $a_i \in [0, n/c)$. Fix a phase i , and let ℓ_i be its length. Let S_b^i be the set of bins in B_i which get a single ball during phase i . Let T_b^i be the set of bins in B_i which get two balls during phase i (and are reset). Let S_a^i be the set of bins in A_i which get a single ball during phase i (and are also reset). Then $b_i - b_{i+1} \geq |S_b^i| - |T_b^i| - |S_a^i|$.

We bound each term on the right-hand side of the inequality. Of all the balls thrown during phase i , in expectation at least $(1 - 1/c)$ are thrown in bins from B_i . By a Chernoff bound, the number of balls thrown in B_i is at least $(1 - 1/c)(1 - \delta)\ell_i$ with probability at least $1 - \exp(-\delta^2\ell_i(1 - 1/c)/4)$, for $\delta \in (0, 1)$. On the other hand, the majority of these balls do not cause collisions in bins from B_i . In particular, from the Poisson approximation, we obtain that $|S_b^i| \geq 2|T_b^i|$ with probability at least $1 - (1/n)^{\alpha+1}$, where we have used $b_i \geq n(1 - 1/c)$.

Considering S_a^i , notice that, w.h.p., at most $(1 + \delta)\ell_i/c$ balls are thrown in bins from A_i . Summing up, given that $\ell_i \geq \sqrt{n}/c$, we obtain that $b_i - b_{i+1} \geq (1 - 1/c)(1 - \delta)\ell_i/2 - (1 + \delta)\ell_i/c$, with probability at least $1 - \max((1/n)^\alpha, \exp(-\delta^2\ell_i(1 - 1/c)/4))$. For small $\delta \in (0, 1)$ and $c \geq 10$, the difference is at least ℓ_i/c^2 . Notice also that the probability depends on the length of the phase.

We say that a phase is *regular* if its length is at least $\min(n/\sqrt{a_i}, n/(b_i)^{1/3})/c$. From Lemma 8, the probability that a phase is regular is at least $1 - 1/(4c^2)$. Also, in this case, $\ell_i \geq \sqrt{n}/c$, by calculation. If the phase is regular, then the size of b_i decreases by $\Omega(\sqrt{n})$, w.h.p.

If the phase is not regular, we simply show that, with high probability, a_i does not decrease. Assume $a_i < a_{i+1}$. Then, either $\ell_i < \log n$, which occurs with probability at most $1/n^{\Omega(\log n)}$ by Lemma 8, or the inequality $b_i - b_{i+1} \geq \ell_i/c^2$ fails, which also occurs with probability at most $1/n^{\Omega(\log n)}$.

To complete the proof, consider a series of $\beta\sqrt{n}$ consecutive phases, and assume that a_i is in the third range for all of them. The probability that such a phase is regular is at least $1 - 1/(4c^2)$, therefore, by Chernoff, a constant fraction of phases are regular, w.h.p. Also w.h.p., in each such phase the size of b_i goes down by $\Omega(\sqrt{n})$ units. On the other hand, by the previous argument, if the phases are not regular, then it is still extremely unlikely that b_i increases for the next phase. Summing up, it follows that the probability that the system stays in the third range for $\beta\sqrt{n}$ consecutive phases is at most $1/n^\alpha$, where $\beta \geq 2c^2$, and $\alpha \geq 4$ was fixed initially. \square

This completes the proof of Lemma 9. \square

Final argument. To complete the proof of Theorem 5, recall that we are interested in the expected length of a phase. To upper bound this quantity, we group the states of the game according to their range as follows: state $S_{1,2}$ contains all states (a_i, b_i) in the first two ranges, i.e. with $a_i \geq n/c$. State S_3 contains all states (a_i, b_i) such that $a_i < n/c$. The expected length of a phase starting from a state in $S_{1,2}$ is $O(\sqrt{n})$, from Lemma 8. However, the phase length could be $\omega(\sqrt{n})$ if the state is in S_3 . We can mitigate this fact given that the probability of moving to range three is low (Claim 4), and the system moves away from range three rapidly (Claim 5): intuitively, the probability of states in S_3 in the stationary distribution has to be very low.

To formalize the argument, we define two Markov chains. The first Markov chain M has two states, $S_{1,2}$ and S_3 . The transition probability from $S_{1,2}$ to S_3 is $1/n^\alpha$, whereas the transition probability from S_3 to $S_{1,2}$ is $x > 0$, fixed but unknown. Each state loops onto itself, with probabilities $1 - 1/n^\alpha$ and $1 - x$, respectively. The second Markov chain M' has two states S and R . State S has a transition to R , with

probability $\beta\sqrt{n}/n^\alpha$, and a transition to itself, with probability $1 - \beta\sqrt{n}/n^\alpha$. State R has a loop with probability $1/n^\alpha$, and a transition to S , with probability $1 - 1/n^\alpha$.

It is easy to see that both Markov chains are ergodic. Let $[s \ r]$ be the stationary distribution of M' . Then, by straightforward calculation, we obtain that $s \geq 1 - \beta\sqrt{n}/n^\alpha$, while $r \leq \beta\sqrt{n}/n^\alpha$.

On the other hand, notice that the probabilities in the transition matrix for M' correspond to the probabilities in the transition matrix for $M^{\beta\sqrt{n}}$, i.e. M applied to itself $\beta\sqrt{n}$ times. This means that the stationary distribution for M is the same as the stationary distribution for M' . In particular, the probability of state $S_{1,2}$ is at least $1 - \beta\sqrt{n}/n^\alpha$, and the probability of state S_3 is at most $\beta\sqrt{n}$.

To conclude, notice that the expected length of a phase is at most the expected length of a phase in the first Markov chain M . Using the above bounds, this is at most $2\alpha\sqrt{n}(1 - \beta\sqrt{n}/n^\alpha) + \beta n^{2/3}\sqrt{n}/n^\alpha = O(\sqrt{n})$, as claimed. This completes the proof of Theorem 5.

6.2 Parallel Code

We now use the same framework to derive a convergence bound for parallel code, i.e. a method call which completes after the process executes q steps, irrespective the concurrent actions of other processes. The pseudocode is given in Algorithm 4.

```

1 Shared: register  $R$ 
2 procedure call()
3 while true do
4   for  $i$  from 1 to  $q$  do
5     Execute  $i$ th step
6   output success

```

Algorithm 4: Pseudocode for parallel code.

Analysis. We now analyze the individual and system latency for this algorithm under the uniform stochastic scheduler. Again, we start from its Markov chain representation. We define the individual Markov chain M_I to have states $S = (C_1, \dots, C_n)$, where $C_i \in \{0, \dots, q-1\}$ is the current step counter for process p_i . At every step, the Markov chain picks i from 1 to n uniformly at random and transitions into the state $(C_1, \dots, (C_i + 1) \bmod q, \dots, C_n)$. A process registers a success every time its counter is reset to 0; the system registers a success every time some process counter is reset to 0. The system latency is the expected number of system steps between two successes, and the individual latency is the expected number of system steps between two successes by a specific process.

We now define the system Markov chain M_S , as follows. A state $g \in M_S$ is given by q values $(v_0, v_1, \dots, v_{q-1})$, where for each $j \in \{0, \dots, q-1\}$ v_j is the number of processes with step counter value j , with the condition that $\sum_{j=0}^{q-1} v_j = n$. Given a state $(v_0, v_1, \dots, v_{q-1})$, let X be the set of indices $i \in \{0, \dots, q-1\}$ such that $v_i > 0$. Then, for each $i \in X$, the system chain transitions into the state $(v_0, \dots, v_i - 1, v_{i+1} + 1, \dots, v_{q-1})$ with probability v_i/n .

It is easy to check that both M_I and M_S are ergodic Markov chains. Let π be the stationary distribution of M_S , and π' be the stationary distribution of M_I . We next define the mapping $f : M_I \rightarrow M_S$ which maps each state $S = (C_1, \dots, C_n)$ to the state $(v_0, v_1, \dots, v_{q-1})$, where v_j is the number of processes with counter value j from S . Checking that this mapping is a lifting between M_I and M_S is straightforward.

Lemma 10. *The function f defined above is a lifting between the ergodic Markov chains M_I and M_S .*

We then obtain bounds on the system and individual latency.

Lemma 11. *For any $1 \leq i \leq n$, the individual latency for process p_i is $W_i = nq$. The system latency is $W = q$.*

Proof. We examine the stationary distributions of the two Markov chains. Contrary to the previous examples, it turns out that in this case it is easier to determine the stationary distribution of the individual Markov chain M_I . Notice that, in this chain, all states have in- and out-degree n , and the transition probabilities are uniform (probability $1/n$). It therefore must hold that the stationary distribution of M_I is *uniform*. Further, notice that a $1/nq$ fraction of the edges corresponds to the counter of a specific process p_i being reset. Therefore, for any i , the probability that a step in M_I is a completed operation by p_i is $1/nq$. Hence, the individual latency for the algorithm is nc . To obtain the system latency, we notice that, from the lifting, the probability that a step in M_S is a completed operation by *some* process is $1/q$. Therefore, the individual latency for the algorithm is q . \square

6.3 General Bound for $SCU(q, s)$

We now put together the results of the previous sections to obtain a bound on individual and system latency. First, we notice that Theorem 5 can be easily extended to the case where the loop contains s scan steps, as the extended local state of a process p can be changed by a step of another process $q \neq p$ only if p is about to perform a CAS operation.

Corollary 1. *For $s \geq 1$, given a scan-validate pattern with s scan steps under the stochastic scheduler, the system latency is $O(s\sqrt{n})$, while the individual latency is $O(ns\sqrt{n})$.*

Obviously, an algorithm in $SCU(q, s)$ is a sequential composition of parallel code followed by s loop steps. Fix a process p_i . By Lemma 11 and Corollary 1, by linearity of expectation, we obtain that the expected individual latency for process p_i to complete an operation is at most $n(q + \alpha s\sqrt{n})$, where $\alpha \geq 4$ is a constant.

Consider now the Markov Chain M_S that corresponds to the sequential composition of the Markov chain for the parallel code M_P , and the Markov chain M_L corresponding to the loop. In particular, a completed operation from M_P does not loop back into the chain, but instead transitions into the corresponding state of M_L . More precisely, if the transition is a step by some processor p_i which completed step number q in the parallel code (and moves to the loop code), then the chain transitions into the state where processor p_i is about to execute the first step of the loop code. Similarly, when a process performs a successful CAS at the end of the loop, the processes' step counter is reset to 0, and its next operation will be the first step of the preamble.

It is straightforward that the chain M_S is ergodic. Let κ_i be the probability of the event that process p_i completes an operation in the stationary distribution of the chain M_S . Since the expected number of steps p_i needs to take to complete an operation is at most $n(q + \alpha\sqrt{n})$, we have that $\kappa_i \geq 1/(n(q + \alpha s\sqrt{n}))$. Let κ be the probability of the event that *some* process completes an operation in the stationary distribution of the chain M_S . It follows that $\kappa = \sum_{i=1}^n \kappa_i \geq 1/(q + \alpha s\sqrt{n})$. Hence, the expected time until the system completes a new operation is at most $q + \alpha s\sqrt{n}$, as claimed.

We note that the above argument also gives an upper bound on the expected number of (individual) steps a process p_i needs to complete an operation (similar to the standard measure of individual *step complexity*). Since the scheduler is uniform, this is also $O(q + s\sqrt{n})$. Finally, we note that, if only $k \leq n$ processes are correct in the execution, we obtain the same latency bounds in terms of k : since we consider the behavior of the algorithm at infinity, the stationary latencies are only influenced by correct processes.

Corollary 2. *Given an algorithm in $SCU(q, s)$ on k correct processes under a uniform stochastic scheduler, the system latency is $O(q + s\sqrt{k})$, and the individual latency is $O(k(q + s\sqrt{k}))$.*

7 Application - A Fetch-and-Increment Counter using Augmented CAS

We now apply the ideas from the previous section to obtain minimal and maximal progress bounds for other lock-free algorithms under the uniform stochastic scheduler.

Some architectures support richer semantics for the CAS operation, which return the *current* value of the register which the operation attempts to modify. We can take advantage of this property to obtain a simpler fetch-and-increment counter implementation based on compare-and-swap. This type of counter implementation is very widely-used [4].

```

7 Shared: register  $R$ 
8 procedure fetch-and-inc()  $v \leftarrow 0$ 
9 while true do
10    $old \leftarrow v$ 
11    $v \leftarrow \text{CAS}(R, v, v + 1)$ 
12   if  $v = old$  then
13     output  $success$ 

```

Algorithm 5: A lock-free fetch-and-increment counter based on compare-and-swap.

7.1 Markov Chain Representations

We again start from the observation the algorithm induces an individual Markov chain and a global one. From the point of view of each process, there are two possible states: *Current*, in which the process has the *current* value (i.e. its local value v is the same as the value of the register R), and the *Stale* state, in which the process has an old value, which will cause its CAS call to fail. (In particular, the *Read* and *OldCAS* states from the universal construction are coalesced.)

The Individual Chain. The per-process chain, which we denote by M_I , results from the composition of the automata representing the algorithm at each process. Each state of M_I can be characterized by the set of processes that have the current value of the register R . The Markov chain has $2^n - 1$ states, since it never happens that *no thread* has the current value.

For each non-empty subset of processes S , let s_S be the corresponding state. The initial state is s_{Π} , the state in which every thread has the current value. We distinguish *winning* states as the states $(s_{\{p_i\}})_i$ in which only *one* thread has the current value: to reach this state, one of the processes must have successfully updated the value of R . There are exactly n winning states, one for each process.

Transitions are defined as follows. From each state s , there are n outgoing edges, one for each process which could be scheduled next. Each transition has probability $1/n$, and moves to state s' corresponding to the set of processes which have the current value at the next time step. Notice that the winning states are the only states with a self-loop, and that from every state s_S the chain either moves to a state s_V with $|V| = |S| + 1$, or to a winning state for one of the threads in S .

The Global Chain. The *global* chain M_G results from clustering the symmetric states from M_I into single states. The chain has n states v_1, \dots, v_n , where state v_i comprises all the states s_S in M_G such that

$|S| = i$. Thus, state v_1 is the state in which *some* process just completed a new operation. In general, v_i is the state in which i processes have the current value of R (and therefore may commit an operation if scheduled next).

The transitions in the global chain are defined as follows. For any $1 \leq i \leq n$, from state v_i the chain moves to state v_1 with probability i/n . If $i < n$, the chain moves to state v_{i+1} with probability $1 - i/n$. Again, the state v_1 is the only state with a self-loop. The intuition is that some process among the i possessing the current value wins if scheduled next (and changes the current value); otherwise, if some other thread is scheduled, then that thread will also have the current value.

7.2 Algorithm Analysis

We analyze the stationary behavior of the algorithm under a uniform stochastic scheduler, assuming each process invokes an infinite number of operations.

Strategy. We are interested in the expected number of steps that some process p_i takes between committing two consecutive operations, in the stationary distribution. This is the *individual latency*, which we denote by W_i . As for the general algorithm, we proceed by first bounding the system latency W , which is easier to analyze, and then show that $W_i = nW$, i.e. the algorithm is *fair*. We will use the two Markov chain representations from the previous section. In particular, notice that W_i is the expected return time of the “win state” v_i of the global chain M_G , and W is the expected return time of the state s_{p_i} in which p_i just completed an operation.

The first claim is an upper bound on the return time for v_1 in M_G .

Lemma 12. *The expected return time for v_1 is $W \leq 2\sqrt{n}$.*

Proof. For $0 \leq i \leq n - 1$, let $Z(i)$ be the hitting time for state v_1 from the state where $n - i$ processes have the current value. In particular, $Z(0)$ is the hitting time from the state where *all* processes have the correct value, and therefore $Z(0) = 1$. Analyzing the transitions, we obtain that $Z(i) = iZ(i - 1)/n + 1$. We prove that $Z(n - 1) \leq 2\sqrt{n}$.

We analyze two intervals: k from 0 to $n - \sqrt{n}$, and then up to $n - 1$. We first claim that, for $0 \leq k \leq n - \sqrt{n}$, it holds that $Z(k) \leq \sqrt{n}$. We prove this by induction. The base case obviously holds. For the induction step, notice that $Z(k) \leq Z(k - 1)(n - \sqrt{n})/n + 1$ in this interval. By the hypothesis, $Z(k - 1) \leq \sqrt{n}$, therefore $Z(k) \leq \sqrt{n}$ for $k \leq n - \sqrt{n}$.

For $k \in \{n - \sqrt{n}, \dots, n\}$, notice that $Z(k)$ can add at most 1 at each iteration, and we are iterating at most \sqrt{n} times. This gives an upper bound of $2\sqrt{n}$, as claimed. \square

Remark. Intuitively, the value $Z(n - 1)$ is related to the birthday paradox, since it counts the number of elements that must be chosen uniformly at random from 1 to n (with replacement) until one of the elements appears twice. In fact, this is the Ramanujan Q function [5], which has been studied previously by Knuth [13] and Flajolet et al. [5] in relation to the performance of linear probing hashing. Its asymptotics are known to be $Z(n - 1) = \sqrt{\pi n/2}(1 + o(1))$ [5].

Markov Chain Lifting. We now analyze W_i , the expected number of total system steps for a specific process p_i to commit a new request. We define a mapping $f : M_I \rightarrow M_G$ between the states of the individual Markov chain. For any non-empty set S of processes, the function maps the state $s_S \in M_I$ to the state v_i of the chain. It is straightforward to prove that this mapping is a correct lifting of the Markov chain, and that both Markov chains are ergodic.

Lemma 13. *The individual chain and the local chain are ergodic. The function f is a lifting between the individual chain and the global chain.*

We then use the lifting and symmetry to obtain the following relation between the stationary distributions of the two Markov chains. The proof is similar to that of Lemma 5. This also implies that every process takes the same number of steps in expectation until completing an operation.

Lemma 14. *Let $\pi = [\pi_1 \dots \pi_n]$ be the stationary distribution of the global chain, and let π' be the stationary distribution of the individual chain. Let π'_i be the probability of $s_{\{p_i\}}$ in π' . Then, for all $i \in \{1, \dots, n\}$, $\pi'_i = \pi/n$. Furthermore, $W_i = nW$.*

This characterizes the asymptotic behavior of the individual latency.

Corollary 3. *For any $i \in \{1, \dots, n\}$, the expected number of system steps between two completed operations by process p_i is $O(n\sqrt{n})$. The expected number of steps by p_i between two completed operations is $O(\sqrt{n})$.*

8 Discussion

This paper is motivated by the fundamental question of relating the theory of concurrent programming to real-world algorithm behavior. We give a framework for analyzing concurrent algorithms which partially explains the wait-free behavior of lock-free algorithms, and their good performance in practice. Our work is a first step in this direction, and opens the door to many additional questions.

In particular, we are intrigued by the goal of obtaining a realistic model for the unpredictable behavior of system schedulers. Even though it has some foundation in empirical results, our uniform stochastic model is a rough approximation, and can probably be improved. We believe that some of the elements of our framework (such as the existence of liftings) could still be applied to non-uniform stochastic scheduler models, while others may need to be further developed. A second direction for future work is studying other types of algorithms, and in particular implementations which export several distinct methods. The class of algorithms we consider is *universal*, i.e., covers any sequential object, however there may exist object implementations which do not fall in this class. Finally, it would be interesting to explore whether there exist concurrent algorithms which avoid the $\Theta(\sqrt{n})$ contention factor in the latency, and whether such algorithms are efficient in practice.

Acknowledgements. We thank George Giakkoupis, William Hasenplough, Maurice Herlihy, and Yuval Peres for useful discussions, and Faith Ellen for helpful comments on an earlier version of the paper.

References

- [1] Samy Al-Bahra. Nonblocking algorithms and scalable multicore programming. *Commun. ACM*, 56(7):50–61, 2013.
- [2] James Aspnes. Fast deterministic consensus in a noisy environment. *J. Algorithms*, 45(1):16–39, 2002.
- [3] Fang Chen, László Lovász, and Igor Pak. Lifting markov chains to speed up mixing. In *Proceedings of the thirty-first annual ACM symposium on Theory of computing*, STOC '99, pages 275–281, New York, NY, USA, 1999. ACM.

- [4] Dave Dice, Yossi Lev, and Mark Moir. Scalable statistics counters. In *25th ACM Symposium on Parallelism in Algorithms and Architectures, SPAA '13, Montreal, QC, Canada, 2013*, pages 43–52, 2013.
- [5] Philippe Flajolet, Peter J. Grabner, Peter Kirschenhofer, and Helmut Prodinger. On Ramanujan’s Q-function. *J. Comput. Appl. Math.*, 58(1):103–116, March 1995.
- [6] Keir Fraser. Practical lock-freedom. Technical Report UCAM-CL-TR-579, University of Cambridge, Computer Laboratory, February 2004.
- [7] D. Guniguntala, P.E. McKenney, J. Triplett, and J. Walpole. The read-copy-update mechanism for supporting real-time applications on shared-memory multiprocessor systems with linux. *IBM Systems Journal*, 47(2):221–236, 2008.
- [8] Thomas P. Hayes and Alistair Sinclair. Liftings of tree-structured markov chains. In *Proceedings of the 13th international conference on Approximation, and 14 the International conference on Randomization, and combinatorial optimization: algorithms and techniques, APPROX/RANDOM'10*, pages 602–616, Berlin, Heidelberg, 2010. Springer-Verlag.
- [9] Maurice Herlihy. Wait-free synchronization. *ACM Transactions on Programming Languages and Systems*, 13(1):123–149, January 1991.
- [10] Maurice Herlihy, Victor Luchangco, and Mark Moir. Obstruction-free synchronization: Double-ended queues as an example. In *23rd International Conference on Distributed Computing Systems (ICDCS 2003)*, pages 522–529, 2003.
- [11] Maurice Herlihy and Nir Shavit. *The art of multiprocessor programming*. Morgan Kaufmann, 2008.
- [12] Maurice Herlihy and Nir Shavit. On the nature of progress. In *15th International Conference on Principles of Distributed Systems (OPODIS), Toulouse, France, December 13-16, 2011. Proceedings*, pages 313–328, 2011.
- [13] Donald E. Knuth. *The art of computer programming, volume 3: (2nd ed.) sorting and searching*. Addison Wesley Longman Publishing Co., Inc., Redwood City, CA, USA, 1998.
- [14] Alex Kogan and Erez Petrank. A methodology for creating fast wait-free data structures. In *Proceedings of the 17th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming, PPOPP '12*, pages 141–150, New York, NY, USA, 2012. ACM.
- [15] Leslie Lamport. A new solution of dijkstra’s concurrent programming problem. *Commun. ACM*, 17(8):453–455, 1974.
- [16] David A. Levin, Yuval Peres, and Elizabeth L. Wilmer. *Markov Chains and Mixing Times*. American Mathematical Society, 2008.
- [17] Maged M. Michael and Michael L. Scott. Simple, fast, and practical non-blocking and blocking concurrent queue algorithms. In *PODC*, pages 267–275, 1996.
- [18] Michael Mitzenmacher and Eli Upfal. *Probability and Computing: Randomized Algorithms and Probabilistic Analysis*. Cambridge University Press, New York, NY, USA, 2005.

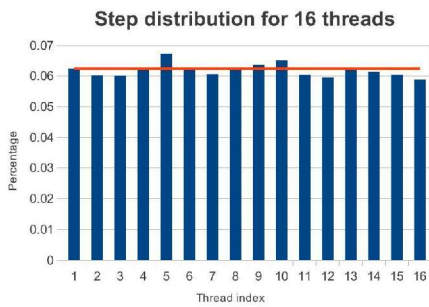


Figure 3: Percentage of steps taken by each process during an execution.

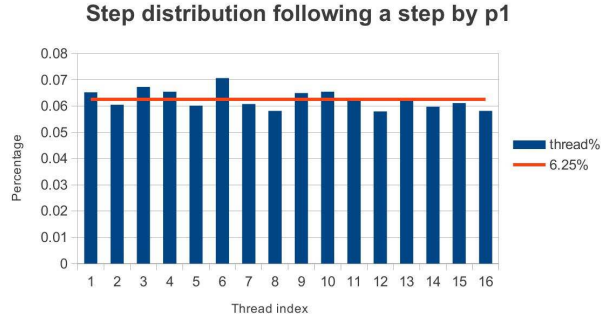


Figure 4: Percentage of steps taken by processes, starting from a step by p_1 . (The results are similar for all threads.)

- [19] David Petrou, John W. Milford, and Garth A. Gibson. Implementing lottery scheduling: matching the specializations in traditional schedulers. In *Proceedings of the annual conference on USENIX Annual Technical Conference, ATEC '99*, pages 1–1, Berkeley, CA, USA, 1999. USENIX Association.
- [20] Shahar Timnat and Erez Petrank. A practical wait-free simulation for lock-free data structures. In *Proceedings of the Symposium on Principles and Practice of Parallel Programming, PPOPP '14*. To Appear., 2014.
- [21] R. K. Treiber. Systems programming: Coping with parallelism. Technical Report RJ 5118, IBM Almaden Research Center, 1986.

Structure of the Appendix. Section A presents empirical data for the stochastic scheduler model, while Section B gives compares the predicted and actual performance of an algorithm in *SCU*.

A The Stochastic Scheduler Model

A.1 Empirical Justification

The real-world behavior of a process scheduler arises as a complex interaction of factors such as the timing of memory requests (influenced by the algorithm), the behavior of the cache coherence protocol (dependent on the architecture), or thread pre-emption (depending on the operating system). Given the extremely complex interactions between these components, the behavior of the scheduler could be seen as *non-deterministic*. However, when recorded for extended periods of time, simple patterns emerge. Figures 3 and 4 present statistics on schedule recordings from a simple concurrent counter algorithm, executed on a system with 16 hardware threads. (The details of the setup and experiments are presented in the next section).

Figure 3 clearly suggests that, in the long run, the scheduler is “fair:” each thread gets to take about the same number of steps. Figure 4 gives an intuition about how the schedule looks like *locally*: assuming process p_i just took a step at time step τ , any process appears to be just as likely to be scheduled in the next step. We note that the structure of the algorithm executed can influence the ratios in Figure 4; also, we only performed tests on an Intel architecture.

Our stochastic scheduler model addresses the non-determinism in the scheduler by associating a distribution with each scheduler time step, which gives the probability of each process being scheduled next.

In particular, we model our empirical observations by considering the uniform stochastic scheduler, which assigns a probability of $1/n$ with which each process is scheduled. We stress that we do not claim that the schedule behaves uniformly random locally; our claim is that the behavior of the schedule over long periods of time can be approximated reasonably in this way, for the algorithms we consider. We note that randomized schedulers attempting to explicitly implement probabilistic fairness have been proposed in practice, in the form of *lottery scheduling* [19].

A.2 Experimental Setup

The machine we use for testing is a Fujitsu PRIMERGY RX600 S6 server with four Intel Xeon E7-4870 (Westmere EX) processors. Each processor has 10 2.40 GHz cores, each of which multiplexes two hardware threads, so in total our system supports 80 hardware threads. Each core has private write-back L1 and L2 caches; an inclusive L3 cache is shared by all cores. We limited experiments to 20 hardware threads, in order to avoid the effects of non-uniform memory access (NUMA), which appear when hardware threads are located on different cores.

We used two methods to record schedules. The first used an atomic fetch-and-increment operation (available in hardware): each process repeatedly calls this operation, and records the values received. We then sort the values of each process to recover the total order of steps. The second method records timestamps during the execution of an algorithm, and sorts the timestamps to recover the total order. We found that the latter method interferes with the schedule: since the timer call causes a delay to the caller, a process is less likely to be scheduled twice in succession. With this exception, the results are similar for both methods. The statistics of the recorded schedule are summarized in Figures 3 and 4. (The graphs are built using 20 millisecond runs, averaged over 10 repetitions; results for longer intervals and for different thread counts are similar.)

B Implementation Results

Let the *completion rate* of the algorithm be the total number of successful operations versus the total number of steps taken during the execution. The completion rate approximates the inverse of the system latency. We consider a fetch-and-increment counter implementation which simply reads the value v of a shared register R , and then attempts to increment the value using a $\text{CAS}(R, v, v + 1)$ call. The predicted completion rate of the algorithm is $\Theta(1/\sqrt{n})$. The actual completion rate of the implementation is shown in Figure 5 for varying thread counts, for a counter implementation based on the lock-free pattern. The $\Theta(1/\sqrt{n})$ rate predicted by the uniform stochastic scheduler model appears to be close to the actual completion rate. Since we do not have precise bounds on the constant in front of $\Theta(1/\sqrt{n})$ for the prediction, we scaled the prediction to the first data point. The worst-case predicted rate $(1/n)$ is also shown.

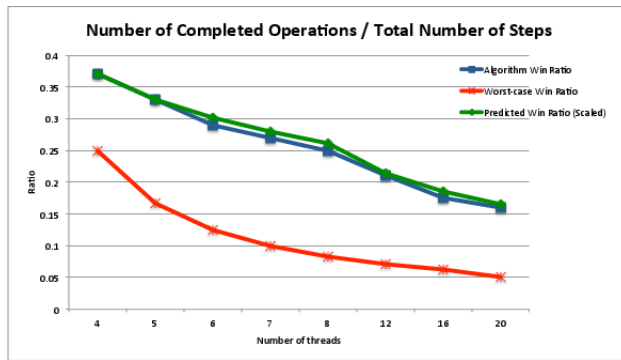


Figure 5: Predicted completion rate of the algorithm vs. completion rate of the implementation vs. worst-case completion rate.