


Institute for Language, Logic and Information

**A RANDOMIZED ALGORITHM FOR TWO-PROCESS
WAIT-FREE TEST-AND-SET**

John Tromp
Paul M.B. Vitányi

ITLI Prepublication Series
for Computation and Complexity Theory CT-91-10



University of Amsterdam

The ITLI Prepublication Series

1986

86-01 Peter van Emde Boas
 86-02 Johan van Benthem
 86-03 Reinhard Muskens
 86-04 Kenneth A. Bowen, Dick de Jongh
 86-05 Johan van Benthem

The Institute of Language, Logic and Information
 A Semantical Model for Integration and Modularization of Rules
 Categorical Grammar and Lambda Calculus
 A Relational Formulation of the Theory of Types
 Some Complete Logics for Branched Time, Part I Well-founded Time,
 Logical Syntax Forward looking Operators

1987

87-01 Jeroen Groenendijk, Martin Stokhof
 87-02 Renate Bartsch
 87-03 Jan Willem Klop, Roel de Vrijer
 87-04 Johan van Benthem
 87-05 Víctor Sánchez Valencia
 87-06 Eleonore Oversteegen
 87-07 Johan van Benthem
 87-08 Renate Bartsch
 87-09 Herman Hendriks

Type shifting Rules and the Semantics of Interrogatives
 Frame Representations and Discourse Representations
 Unique Normal Forms for Lambda Calculus with Surjective Pairing
 Polyadic quantifiers
 Traditional Logicians and de Morgan's Example
 Temporal Adverbials in the Two Track Theory of Time
 Categorical Grammar and Type Theory
 The Construction of Properties under Perspectives
 Type Change in Semantics: The Scope of Quantification and Coordination

1988

LP-88-01 Michiel van Lambalgen
 LP-88-02 Yde Venema
 LP-88-03 Reinhard Muskens
 LP-88-04 Johan van Benthem
 LP-88-05 Johan van Benthem
 LP-88-06 Johan van Benthem
 LP-88-07 Renate Bartsch
 LP-88-08 Jeroen Groenendijk, Martin Stokhof
 LP-88-09 Theo M.V. Janssen
 LP-88-10 Anneke Kleppe

Logic, Semantics and Philosophy of Language: Algorithmic Information Theory
 Expressiveness and Completeness of an Interval Tense Logic
 Year Report 1987
 Going partial in Montague Grammar
 Logical Constants across Varying Types
 Semantic Parallels in Natural Language and Computation
 Tenses, Aspects, and their Scopes in Discourse
 Context and Information in Dynamic Semantics
 A mathematical model for the CAT framework of Eurotra
 A Blissymbolics Translation Program

ML-88-01 Jaap van Oosten
 ML-88-02 M.D.G. Swaen
 ML-88-03 Dick de Jongh, Frank Veltman
 ML-88-04 A.S. Troelstra
 ML-88-05 A.S. Troelstra

Mathematical Logic and Foundations: Lifschitz' Realizability
 The Arithmetical Fragment of Martin Löf's Type Theories with weak Σ -elimination
 Provability Logics for Relative Interpretability
 On the Early History of Intuitionistic Logic
 Remarks on Intuitionism and the Philosophy of Mathematics

CT-88-01 Ming Li, Paul M.B. Vitanyi
 CT-88-02 Michiel H.M. Smid
 CT-88-03 Michiel H.M. Smid, Mark H. Overmars
 Leen Torenvliet, Peter van Emde Boas
 CT-88-04 Dick de Jongh, Lex Hendriks
 Gerard R. Renardel de Lavalette

Computation and Complexity Theory: Two Decades of Applied Kolmogorov Complexity
 General Lower Bounds for the Partitioning of Range Trees
 Maintaining Multiple Representations of
 Dynamic Data Structures
 Computations in Fragments of Intuitionistic Propositional Logic

CT-88-05 Peter van Emde Boas
 CT-88-06 Michiel H.M. Smid
 CT-88-07 Johan van Benthem
 CT-88-08 Michiel H.M. Smid, Mark H. Overmars
 Leen Torenvliet, Peter van Emde Boas

Machine Models and Simulations (revised version)
 A Data Structure for the Union-find Problem having good Single-Operation Complexity
 Time, Logic and Computation
 Multiple Representations of Dynamic Data Structures

CT-88-09 Theo M.V. Janssen
 CT-88-10 Edith Spaan, Leen Torenvliet, Peter van Emde Boas
 CT-88-11 Sieger van Denneheuvel, Peter van Emde Boas

Towards a Universal Parsing Algorithm for Functional Grammar
 Nondeterminism, Fairness and a Fundamental Analogy
 Towards implementing RL

X-88-01 Marc Jumelet

Other prepublications: On Solovay's Completeness Theorem

1989

LP-89-01 Johan van Benthem
 LP-89-02 Jeroen Groenendijk, Martin Stokhof

Logic, Semantics and Philosophy of Language: The Fine-Structure of Categorical Semantics
 Dynamic Predicate Logic, towards a compositional,
 non-representational semantics of discourse

LP-89-03 Yde Venema
 LP-89-04 Johan van Benthem
 LP-89-05 Johan van Benthem
 LP-89-06 Andreja Prijatelj
 LP-89-07 Heinrich Wansing
 LP-89-08 Víctor Sánchez Valencia
 LP-89-09 Zhisheng Huang

Two-dimensional Modal Logics for Relation Algebras and Temporal Logic of Intervals
 Language in Action
 Modal Logic as a Theory of Information
 Intensional Lambek Calculi: Theory and Application
 The Adequacy Problem for Sequential Propositional Logic
 Peirce's Propositional Logic: From Algebra to Graphs
 Dependency of Belief in Distributed Systems

ML-89-01 Dick de Jongh, Albert Visser
 ML-89-02 Roel de Vrijer
 ML-89-03 Dick de Jongh, Franco Montagna
 ML-89-04 Dick de Jongh, Marc Jumelet, Franco Montagna
 ML-89-05 Rineke Verbrugge
 ML-89-06 Michiel van Lambalgen
 ML-89-07 Dirk Roorda
 ML-89-08 Dirk Roorda
 ML-89-09 Alessandra Carbone

Mathematical Logic and Foundations: Explicit Fixed Points for Interpretability Logic
 Extending the Lambda Calculus with Surjective Pairing is conservative
 Rosser Orderings and Free Variables
 On the Proof of Solovay's Theorem
 Σ -completeness and Bounded Arithmetic
 The Axiomatization of Randomness
 Elementary inductive Definitions in HA: from Strictly Positive towards Monotone
 Investigations into Classical Linear Logic
 Provable Fixed points in $\text{ID}_0 + \Omega_1$

CT-89-01 Michiel H.M. Smid
 CT-89-02 Peter van Emde Boas
 CT-89-03 Ming Li, Herman Neuféglise, Leen Torenvliet, Peter van Emde Boas
 CT-89-04 Harry Buhrman, Leen Torenvliet
 CT-89-05 Pieter H. Hartel, Michiel H.M. Smid
 Leen Torenvliet, Willem G. Vree

Computation and Complexity Theory: Dynamic Deferred Data Structures
 Machine Models and Simulations
 On Space Efficient Simulations
 A Comparison of Reductions on Nondeterministic Space
 A Parallel Functional Implementation of Range Queries

CT-89-06 H.W. Lenstra, Jr.
 CT-89-07 Ming Li, Paul M.B. Vitanyi

Finding Isomorphisms between Finite Fields
 A Theory of Learning Simple Concepts under Simple Distributions and
 Average Case Complexity for the Universal Distribution (Prel. Version)

CT-89-08 Harry Buhrman, Steven Homer
 Leen Torenvliet

Honest Reductions, Completeness and
 Nondeterministic Complexity Classes

CT-89-09 Harry Buhrman, Edith Spaan, Leen Torenvliet
 CT-89-10 Sieger van Denneheuvel

On Adaptive Resource Bounded Computations
 The Rule Language RL/1


CT-89-11 Zhisheng Huang, Sieger van Denneheuvel
 Peter van Emde Boas

Towards Functional Classification of Recursive Query Processing

X-89-01 Marianne Kalsbeek
 X-89-02 G. Wagemakers
 X-89-03 A.S. Troelstra
 X-89-04 Jeroen Groenendijk, Martin Stokhof
 X-89-05 Maarten de Rijke
 X-89-06 Peter van Emde Boas

Other Prepublications: An Orey Sentence for Predicative Arithmetic
 New Foundations: a Survey of Quine's Set Theory
 Index of the Heyting Nachlass
 Dynamic Montague Grammar, a first sketch
 The Modal Theory of Inequality
 Een Relationele Semantiek voor Conceptueel Modelleren: Het RL-project

1990 SEE INSIDE BACK COVER



Instituut voor Taal, Logica en Informatie
Institute for Language, Logic and
Information

Faculteit der Wiskunde en Informatica
(Department of Mathematics and Computer Science)
Plantage Muidergracht 24
1018TV Amsterdam

Faculteit der Wijsbegeerte
(Department of Philosophy)
Nieuwe Doelenstraat 15
1012CP Amsterdam

A RANDOMIZED ALGORITHM FOR TWO-PROCESS
WAIT-FREE TEST-AND-SET

John Tromp
CWI

Paul M.B. Vitányi
Department of Mathematics and Computer Science
University of Amsterdam
& CWI

ITLI Prepublication Series
for Computation and Complexity Theory
ISSN 0924-8374

Received June 1991

A Randomized Algorithm for Two-Process Wait-Free Test-and-Set

John Tromp
CWI*

Paul Vitányi
CWI & UvA†

Abstract

It is known to be impossible to implement wait-free test-and-set deterministically in a concurrent setting using only atomic shared variables. It can be shown that n -process wait-free test-and-set can be deterministically implemented from 2-process wait-free test-and-set. Here we present a simple direct *randomized* algorithm for 2-process wait-free test-and-set, implemented with two 4-valued single writer single reader atomic variables. The worst-case (over all adversary schedulers) expected number of steps to execute a test-and-set between two processes is at most 11, while the reset takes exactly 1 step. Based on a finite-state analysis, the proofs of correctness and expected length are compressed into one table.

1 Introduction

A concurrent system consists of n processes communicating through concurrent data objects R_0, \dots, R_{m-1} . An implementation of a new concurrent data object x (rather, a family of objects, one for each n) is *wait-free* if there is a total function f , such that each process can complete any operation associated with x within $f(n)$ steps, irrespective of the timing and execution speeds of the other processes. A step to is a single access to one of the R_i 's. Local

*Centrum voor Wiskunde en Informatica, Kruislaan 413, 1098 SJ Amsterdam, The Netherlands. email: tromp@cwi.nl

†Faculteit Wiskunde en Informatica, Universiteit van Amsterdam. email: paulv@cwi.nl

events, including coin-flips, are not counted. The paper develops a similar definition of wait-freeness for randomized protocols based on the worst case expected length of an operation.

A concurrent object is *constructible* if it can be implemented deterministically with boundedly many safe bits, the mathematical analogues of electronic hardware ‘flip-flops’, [11]. What concurrent wait-free object is the most powerful constructible one? It has been shown that wait-free atomic multi-user variables, and atomic snapshot objects, are constructible [16, 11, 21, 10, 18, 13, 7, 4, 2]. In contrast, wait-free consensus—viewed as an object on which each of n processes can execute just one operation—is not constructible [6, 1], although randomized implementations are possible [6, 1, 5, 19]. Wait-free concurrent test-and-set can deterministically implement 2-process wait-free consensus, and therefore is not deterministically constructible [12, 8]. This raises the question of whether randomized algorithms for test-and-set exist.

In [8] it is shown that repeated use of ‘consensus’ on *unbounded* hardware can implement ‘test-and-set’. In [17, 19, 9] a solution can be obtained by combining several intermediate constructions.

Wait-free n -process test-and-set can be implemented deterministically from wait-free 2-process test-and-set, [3], showing that the impossibility of a deterministic algorithm for n -process test-and-set is solely due to the 2-process case. We present a simple randomized algorithm which directly implements wait-free test-and-set between 2 processes. Randomization means that the algorithm contains a branch conditioned on the outcome of a fair coin flip (as in [20]). We use a finite-state based proof technique for verifying correctness and worst-case expected execution length.

The solution is very simple: it uses two 4-valued 1-writer 1-reader atomic variables. The worst-case expected number of steps in a test-and-set operation is 11, whereas a reset always takes 1 step.

2 Preliminaries

A *test-and-set* bit is a concurrent data object X shared between n processes $0, \dots, n - 1$. The value of X is 0 or 1. Each process i has a local binary variable x_i . At any time exactly one of X, x_0, \dots, x_{n-1} has value 0, all others have value 1. A process i with $x_i = 1$ can atomically execute a test-and-set

operation

read $x_i := X$; write $X := 1$; return x_i .

A process i with $x_i = 0$ can execute a reset operation

$x_i := 1$; write $X := 0$.

This naturally leads to the definition of the *state* of the test-and-set bit, or *0-owner* as a member of $\{\perp, 0, \dots, n - 1\}$ according to which of X and the x_i 's is 0.

2.1 Test-and-Set Definition

Instead of assuming an atomic test-and-set, we want to implement it with actions that are sequences of accesses to atomic shared variables, R_0, \dots, R_{m-1} , executed by processes $0, \dots, n - 1$, according to some *protocol*. Since the executions by the different processes happen concurrently and asynchronously, the implementation should guarantee that each system execution of the implementation is equivalent to a system execution of the above defined construct. This leads to the following set of definitions.

Definition. For a given execution of the system, denote the set of *actions* that have been started as $A = R \cup T, T = T0 \cup T1$, where R is the set of *resets*, and Tx is the set of *test-and-sets* returning $x, x = 0, 1$. By $r, t, t0, t1$ we denote elements from $R, T, T0, T1$, respectively. We partition the set of actions according to the processes executing them: define A_i to be the actions by process i , and similarly define $R_i = A_i \cap R$, and $Tx_i = A_i \cap Tx, x = \epsilon, 0, 1$. Define an *event* as an execution of a statement in a protocol, that is, a write or a read on a R_i or a coin-flip. A read event is qualified by the value obtained and a coin-flip by its outcome. Every new execution of a statement represents a unique event. Number the events of a test-and-set or reset action $a \in A$ as $a.1, a.2, \dots$. Let $l = l(a)$, the *length* of a , be its number of events in the execution (possibly infinite). Denote $a.1$, the *start* of a , as $s(a)$. If a finishes during the execution, then denote $a.l$, the *finish* of a , as $f(a)$. Each event is assumed to execute *atomically*. The sequence of the events of all actions in A is called the *history*. A history induces a partial ordering of the actions: $a \rightarrow b$ iff $f(a) < s(b)$ (the last event of a precedes the first of b in the history). The number of b such that $b \rightarrow a$ is assumed to be finite for each a .

The pair (A, \rightarrow) is called a *run*. An implementation of a concurrent object shared between processes $0, \dots, n-1$, such that each run (A, \rightarrow) satisfies the following atomicity axiom, is an atomic test-and-set.

Atomicity: We can extend \rightarrow on A to a total order \Rightarrow on A in which the sequence of actions satisfies the test-and-set semantics:

1. the system is initially in state \perp .
2. from state \perp , an action $t0 \in T0_i$ moves the system to state i .
3. from state i , an action $r \in R_i$ moves the system to state \perp .
4. from state i , an action $t1 \in T1 - T1_i$ leaves the system in state i .
5. no other state transitions than the above are allowed.

2.2 Randomization, Adversaries and Wait-Freeness

In the above definition of atomicity we did not use the notion of adversary. The reason is that atomicity must hold for all possible histories, and hence for all possible outcomes of coin-flips. The adversary is introduced to enable a quantification of the *wait-freeness*. While it is inevitable that for some histories a test-and-set action may last arbitrarily many steps, the probability of such histories occurring should be minimized. This leads us to define the probability of a certain history occurring.

Fix a protocol P . Let H (H^∞) be the set of finite (infinite) histories that can arise from this protocol. I.e. the set of h such that

1. for all $i < n$, $h|A_i$, the restriction of h to events by process i , satisfies the protocol for process i , and
2. for all $j < m$, $h|R_j$, the restriction of h to events that access R_j , satisfies the usual semantics of such an atomic variable (a read event returns the value written by the last write event preceding it).

For $h \in H$, let the *cylinder* Γ_h be the set of all histories in H^∞ that start with h . Write he to denote history h followed by event e .

An *adversary* is then defined as a probability measure μ on H^∞ satisfying:

1. $\mu(\Gamma_\epsilon) = 1$, where ϵ is the empty history;

2. $\mu(\Gamma_h) = \sum_{h e \in H} \mu(\Gamma_{h e})$, for $h \in H$ and e is a single event; and
3. $\mu(\Gamma_{h c(\text{heads})}) = \mu(\Gamma_{h c(\text{tails})})$, for each coin-flip event $c()$ with $h c() \in H$.

The first two conditions—already implied by the notion of probability measure—are included for completeness. The third condition ensures that the adversary has no control over the outcome of a fair coin flip: both outcomes are equally likely. This definition is readily generalized to biased coins and multi-branch decisions.

Note that this notion of adversary is the strongest possible short of allowing it to predict the future. For example, it includes nonrecursive adversaries using omniscient oracles and randomization.

Now that adversaries have been defined, we can define the expected length $E(h, i)$ of process i 's current (next if idle) action following a finite initial history segment h . Let $\omega \in \Gamma_h$ be an infinite history starting with h . Let $l_{h,i}(\omega)$ be the length (number of events) of process i 's current action following h in ω . If process i is idle at h , then by 'current' we mean 'next', leaving $l_{h,i}$ undefined for ω in which a doesn't start a new action. Define

$$E(h, i) = \sum_{k=1}^{k=\infty} k \cdot \frac{\mu(\{\omega \in \Gamma_h : l_{h,i}(\omega) = k\})}{\mu(\Gamma_h)}.$$

The summation includes the case $k = \infty$ so that the expected length is infinite if (but not necessarily only if) the set of infinite histories in which an operation execution has infinitely many events, has positive measure. The normalization w.r.t. h gives the adversary a free choice of 'starting' configuration.

Definition. An implementation of a concurrent object shared between n processes is *wait-free*, if there is a constant $f(n)$ bounding the expected length $E(h, i)$, for all h, i , under all adversaries.

3 Solution for Two Processes

We give a test-and-set implementation between two processes, process 0 and process 1. The construction uses two 4-valued shared variable objects, R_0 and R_1 . The four values are 'me', 'he', 'choose', 'rst'. Process i solely writes

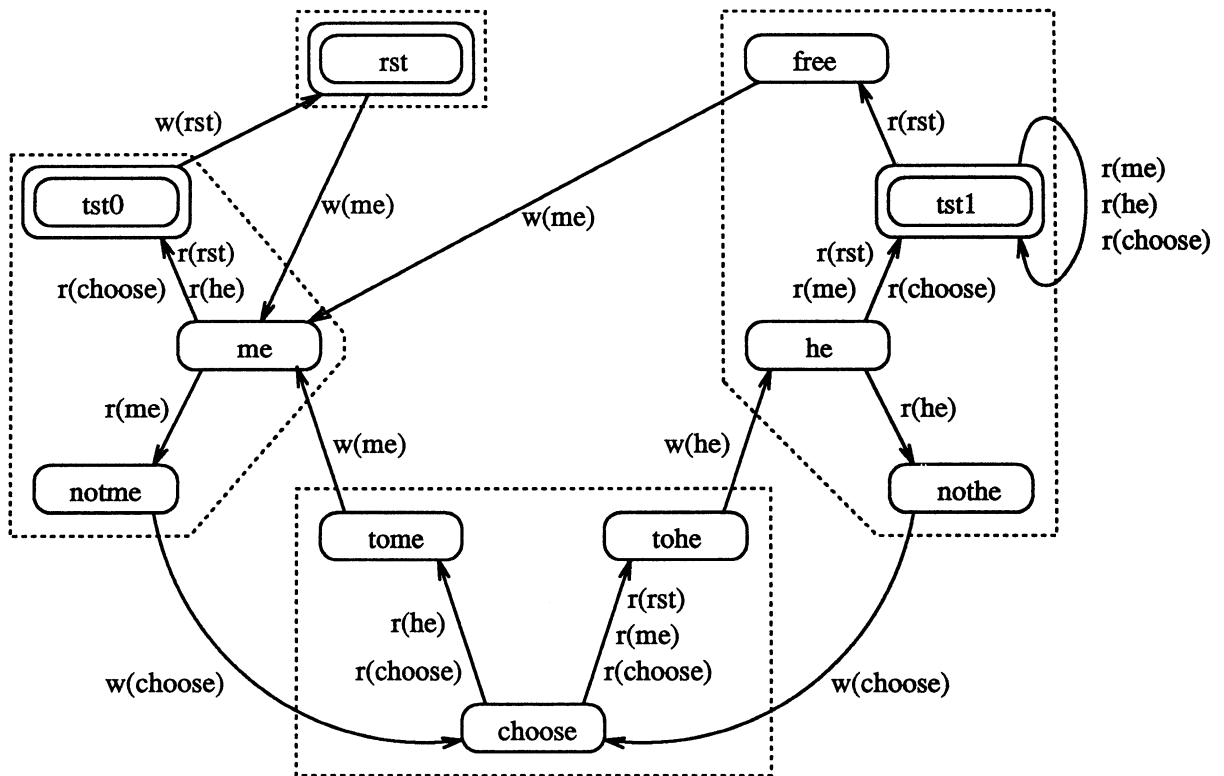


Figure 1: test-and-set protocol

variable R_i , its *own* variable, and solely reads R_{1-i} . For this reason the reads and writes in the protocol don't need to be qualified by the shared variables they access. The protocol, for process i , is first presented as a finite state transition diagram, figure 1. The transitions are labeled with reads $r(value)$ and writes $w(value)$ of the shared variables, where $value$ denotes the value read or written. The 11 states of the protocol are split into 4 groups enclosed by dotted lines. Each group is an equivalence class consisting of the set of states in which that process's variable has the same value. That is, the states in a group are equivalent in the sense that process $1 - i$ cannot distinguish between them by reading R_i . Accordingly, the inter-group events are writes to R_i , whereas the intra-group events are reads of R_{1-i} . Each group is named after the corresponding value of the shared variable. The diagram is deterministic, but for a coin flip which is modeled by the two $r(choose)$ transitions from the choose state.

A more conventional representation of the protocol, for process i , is given below. An occurrence of R_i not preceded by 'write' (resp. R_{1-i} not preceded by 'read') refers to the last value written to it (resp. read from it), stored in correspondingly named local variables. The conditional ' $\text{rnd}(\text{true}, \text{false})$ ' represents the boolean outcome 'true' or 'false' of a fair coin flip. The system is initialized with all local and global variables in state rst .

`test_and_set:`

```

if  $R_i = \text{he}$  AND read  $R_{1-i} \neq \text{rst}$ 
then return 1
write  $R_i := \text{me}$ 
while read  $R_{1-i} = R_i$  do
  write  $R_i := \text{choose}$ 
  if read  $R_{1-i} = \text{he}$  OR ( $R_{1-i} = \text{choose}$  AND  $\text{rnd}(\text{true}, \text{false})$ )
  then write  $R_i := \text{me}$ 
  else write  $R_i := \text{he}$ 
if  $R_i = \text{me}$ 
then return 0
else return 1

```

`reset:`

write $R_i := \text{rst}$

It can be verified in the usual way that the transition diagram represents the operation of the program. The intuition behind the protocol is as follows. The default situation is where both processes are idle in the `rst` state. If process i starts a test-and-set then it writes $R_i := \text{me}$ (indicating its desire to take the 0), and checks whether process $1 - i$ agrees (by *not* having $R_{1-i} := \text{me}$). If so, then it has successfully completed a test-and-set of 0. It is easy to see that in this case process $1 - i$ can not get 0 until process i does a reset by writing $R_i := \text{rst}$. While $R_i = \text{me}$, process $1 - i$ can only move from state 'me' to state 'notme' and on via states 'choose', 'tohe' and 'he' to 'tst1', where it completes a test-and-set of 1.

Problems arise only if both processes see each other's variable equal to 'me'. In this case they are said to *disagree* or *in conflict*. They then proceed to the choose state from where they decide between going for 0 or 1, according to what the other process is seen to be doing. (It is essential that this decision be made in a neutral state, i.e. without a claim of preference for either 0 or 1. If, for example, on seeing a conflict, a process would change preference at random, then a process cannot know for sure whether the other one agrees or is about to write a changed preference.)

The deterministic choices, those made if the other's variable reads different from 'choose', can be seen to lead to a correct resolution of the conflict. A process ending up in the `tst1` state makes sure that its test-and-set of 1 is justified, by remaining in that state until it can be sure that the other process has taken the 0. Only if the other process is seen to be in the `rst` state need it try to take the 0 itself.

Suppose now that process i has read $R_{1-i} = \text{choose}$ and is about to flip a coin. Assume that process $1 - i$ has already moved to one of the states `tome/tohe` (or else reason with the processes interchanged). With 50 percent chance, process i will move to the same state as process $1 - i$ did and thus the conflict will be resolved.

So, intuitively, the probability of each loop through the choose state is at most one half and the expected number of 'choices' (transitions from state choose) at most two. This shows that the worst case expected test-and-set length is 11. Namely, starting from the `tst1` state, it takes 4 steps to get to state choose, another 4 steps to loop back to choose and 3 more steps to reach `tst0/tst1`. The reset operation always takes 1 step.

4 Proof of correctness of the 2-Process Solution

Let h be a history corresponding to a run (A, \rightarrow) of our implementation. Let $B = \{s(t), f(t) : t \in T\} \cup \{s(r) = f(r) : r \in R\}$ be the set of events which start or finish an action. Note that $h|B$, the restriction of h to events in B , completely determines the partial order of actions \rightarrow . Let $C = \{t^* : t \in T\} \cup R$ be the set of atomic occurrences of actions.

The definition of atomic test-and-set for 2 processes, process 0 and process 1, is captured by DFA2, the DFA in figure 2, which accepts all possible sequences of atomic operations (all states final). The states are labeled with the owner of the bit. The arcs representing actions of process 1 are labeled, whereas the non-labeled arcs represent the corresponding actions of process 0.




Figure 2: DFA2: specification of 2-process atomic test-and-set

Figure 3 shows the DFA, DFA3, that accepts the possible sequences of the following events of one process (all states final):

- the start of a test-and-set action, denoted $s(\text{tas})$,
- the atomic occurrence of a test-and-set 0, denoted tas_0 ,
- the atomic occurrence of a test-and-set 1, denoted tas_1 ,
- the finish of a test-and-set 0 action, denoted $f(\text{tas}_0)$,
- the finish of a test-and-set 1 action, denoted $f(\text{tas}_1)$,

- the reset action, denoted `rst`.

These are the events in BUC . The reason for not splitting a reset action into start, atomic occurrence, and finish is that it's implemented in our protocol as a single atomic write where the above three transitions coincide.




Figure 3: DFA3: non-atomic specification of 1-process test-and-set

The proof is based on the finite state diagram DFA4 in figure 4 below (again all states are final).

It is drawn as a cartesian product of the two component processes—transitions of process 0 are drawn vertically and those of process 1 horizontally. For clarity, the transition names are only given once and only for process 1. Identifying the starts and finishes of test-and-set executions with their atomic occurrences by collapsing the $s()$ and $f()$ arcs, the diagram reduces to the atomic test-and-set diagram. Identifying all nodes in the same column (row) reduces the diagram to the diagram of process 0 (process 1).

In the states labeled ‘a’ through ‘h’, neither process owns the 0; the bit is in state \perp . In the states labeled ‘i’ through ‘n’, process 1 owns the 0; the bit is in state 1. In the states labeled ‘o’ through ‘t’, process 0 owns the 0; and the bit is in state 0.

Formally [14], DFA4 is the composition of DFA2 with 2 copies of DFA3, in the I/O Automata framework.

Let NFA4 be the NFA obtained from DFA4 by turning the broken transitions of figure 4 into ϵ -steps.

We claim that acceptance of $h|B$ by NFA4 implies atomicity of (A, \rightarrow) . This is proven as follows. If NFA4 accepts $h|B$, then, corresponding to the ϵ transitions, we can augment $h|B$ with an atomic transition t^* between the start $s(t)$ and finish $f(t)$ of each test-and-set action $t \in T$, to get a history h' accepted by DFA4. Therefore, DFA2, which composes DFA4, accepts $h'|C$, the sequence of atomic events in h' . Furthermore, if $a \rightarrow b$, then $a^* \leq f(a) \rightarrow s(b) \leq b^*$, so \Rightarrow , the total order of actions in $h'|C$, extends \rightarrow . This proves atomicity of (A, \rightarrow) .




Figure 4: DFA4: non-atomic specification of 2-process test-and-set

To show that for all histories $h \in H$ of our implementation, $h|B$ is accepted by NFA4, and thus the correctness of our construction, we assign to each reachable combination of process states (s_0, s_1) a nonempty set S_{s_0, s_1} of NFA4 states, such that: for each history h ending in process states (s_0, s_1) , the set of states in which NFA4 can be after processing $h|B$ contains S_{s_0, s_1} (*). The assignment is given in figure 5. The table entries were chosen so as to minimize the number of ϵ -steps that can be made from each assigned set of NFA4 states. This gives the most insight into the workings of the protocol.

In the table below each row (column) is labeled with a state of process 1 (process 0) as in diagram figure 1. An entry in the table is labeled with roman letters representing a set of atomicity states in figure 4, assigned to that row/column pair of process states. The number ending an entry gives $E(h, 0)$, the expected number of steps to finish the current operation execution of process 0.

We use induction on the length of the history to check (*):

Base: After processing the empty history, NFA4 can be in $\{initialstate\} \supseteq \{d\} = S_{rst, rst}$.

Induction Step: This reduces to checking whether for all transitions (s_0, s_1) to (t_0, t_1) and all NFA4 states $y \in S_{t_0, t_1}$, there is an NFA4 state $x \in S_{s_0, s_1}$, such that NFA4 can move from x to y by processing: either the event corresponding to the transition if it belongs to B , or no event otherwise (there is a sequence of ϵ -steps from x to y).

It is straightforward to check all transitions (state process 0, state process 1) to (newstate process 0, state process 1) or to (state process 0, newstate process 1), corresponding to the atomic transitions in the two copies of protocol figure 1 concerned, to do the induction on the length of the runs to verify correctness, explained above. Simultaneously, the wait-freeness can be checked.

We give an example of checking a few transitions below, and give the interpretation. Verification consists in checking all transitions in the table.

In the default state both processes are in state *rst*. The table entry *d10* gives corresponding state *d*, the start state, in figure 4. The worst-case expected number of steps for a test-and-set by process 0 is 10. Process 0 can start a test-and-set by executing $w(me)$ and entering state *me*. The corresponding table entry *gp9* indicates in figure 4 that the system is now

	rst	tst0	notme	me	tome	choose	tohe	he	nothe	tst1	free
rst	d10	l10	cek10	ek10	ek10	c10	c10	c10	c10	d10	ek10
tst0	s1	*	rt1	rt1	rt1	r1	r1	r1	r1	s1	rt1
notme	agp8	jn8	imoq8	imoq8	*	imoq8	imoq8	o4	*	p4	*
me	gp9	jn9	imoq9	imoq9	imoq9	o1	o1	o1	o1	p1	imoq9
tome	gp10	jn10	*	imoq10	imoq10	imoq6	o2	o2	imoq6	p2	*
choose	a3	j3	imoq7	i3	imoq7	imoq7	imoq7	o3	imoq7	p3	*
tohe	a2	j2	imoq6	i2	i2	imoq6	imoq10	imoq10	*	p6	*
he	a1	j1	i1	i1	i1	i1	imoq9	imoq9	imoq9	p5	*
nothe	a4	j4	*	i4	imoq8	imoq8	*	imoq8	imoq8	p4	*
tst1	d11	l11	k11	k11	k11	k11	k11	k11	k11	*	*
free	gp10	jn10	*	imoq10	*	*	*	*	*	*	*

Figure 5: table to verify correctness and wait-freeness

either in state g meaning that process 0 has executed $s(tas)$, or in state p meaning that process 0 has executed $s(tas)$ and also $tas0$ atomically. The expected number of steps is now $9 \leq 10 - 1$. Suppose process 1 now starts a test-and-set: it executes $w(me)$ and moves to state me . The corresponding table entry $imoq9$ gives the system state as one possibility in $\{i, m, o, q\}$ in figure 4 and the expected number of steps for execution of test-and-set by process 0 is still 9. State m says process 1 has executed $s(tas)$ and $tas0$ atomically, while process 0 has only executed $s(tas)$ —hence the system was previously in state g and not in state p . State i says process 1 has executed $s(tas)$ and $tas0$ atomically, while process 0 has executed $s(tas)$ and $tas1$ atomically—and hence the system was previously in state g and not state p . States o and q imply the same state of affairs with the roles of process 0 and process 1 interchanged, and the previous system state is either p or g .

Note that it is also consistent for the system to be in state h —neither process having executed tas . However, if both processes have started a test-and-set execution, then necessarily, one of them must return 0. We have optimized the table entries by eliminating such spurious states.

Process 0 might now read $R_1 = me$, and move via state $notme$ (table entry $imoq8$) by writing $R_0 := choose$, to state $choose$. Process 1 is idle in the meantime. The table entry is now $i3$. This says that process 1 has atomically executed $tst0$, and process 0 has atomically executed $tst1$. Namely, all subsequent schedules lead in 3 steps of process 0 to state $tst1$ —hence the

expectation 3.

The expected number of remaining steps of process 0's test-and-set has dropped from 8 to 3 by the last step since 8 was the worst-case which could be forced by the adversary. Namely, from the system in state $(notme, me)$, the adversary can schedule process 1 to move to $(notme, notme)$ with table entry $imoq8$, followed by a move of process 1 to state $(notme, choose)$ with table entry $imoq8$, followed by a move of process 0 to state $(choose, choose)$ with table entry $imoq7$. Suppose the adversary now schedules process 0. It now flips a fair coin to obtain the conditional boolean $rnd(true, false)$. If the outcome is *true*, then the system moves to state $(tome, choose)$ with entry $imoq6$. If the outcome is *false*, then the system moves to state $(tohe, choose)$ with table entry $imoq6$. Given a fair coin, this step of process 0 correctly decrements the expected number of steps. Suppose the adversary schedules process 1 in state $(choose, choose)$. Process 1 flips a fair coin. If the outcome is *true* the system moves to state $(choose, tome)$ with table entry $imoq7$; if the outcome is *false* then the system moves to state $(choose, tohe)$ with table entry $imoq7$.

This way the correctness of the implementation can be checked exhaustively by hand. We have done the verification by hand, to optimize the entries, and again by machine.

For the finite-state system as we described, the expected number of remaining steps in a test-and-set execution is *always* bounded by a fixed number. The table shows that, trivially, $1 \leq E(h, 0) \leq 11$ Hence the algorithm is wait-free.

5 On the Difficulty of Multi Process Test And Set

The obvious way to extend the given solution to more than 2 processes would be to arrange them at the leafs of a binary tree. Then, a process wishing to execute an n -process test-and-set, would enter a tournament, as in [15], by executing a separate 2-process test-and-set for each node on the path up to the root. When one of these fails, it would again descend, resetting all the tas-bits on which it succeeded, and return 1. When it succeeds ascending up to the root, it would return 0 and leave the resetting descend to its n -process

reset.

The intuition behind this tree approach is that if a process i fails the test-and-set at some node N , then another process j will get to the root successfully and thus justify the value 1 returned by the former.

The worst case expected length of the n -process operations is only $\log n$ time more than that of the 2-process case.

Unfortunately, this straightforward extension does not work. The problem is that the other process j need not be the one responsible for the failure at node N , and might have started its n -process test-and-set only after process i completes its own. Clearly, the resulting history cannot be linearized.

Nonetheless, it turns out that with a somewhat more complicated construction we can *deterministically* implement n -process test-and-set using 2-process test-and-set as primitives, [3]. This shows that the impossibility of deterministic wait-free atomic n -process test-and-set is completely due to the impossibility of deterministic wait-free atomic 2-process test-and-set. This latter problem we have just solved by a simple direct randomized algorithm.

References

- [1] K. Abrahamson, *On achieving consensus using shared memory*, Proc. 7th ACM Symposium on Principles of Distributed Computing, 1988, pp. 291–302.
- [2] Y. Afek, H. Attiya, D. Dolev, E. Gafni, M. Merritt, N. Shavit, *Atomic snapshots of shared memory*, Proc. 9th ACM Symposium on Principles of Distributed Computing, 1990, pp. 1-13.
- [3] Y. Afek, E. Gafni, J. Tromp, P.M.B. Vitanyi, Wait-free test-and-set, submitted to WDAG91.
- [4] J.H. Anderson, *Composite registers*, Proc. 9th ACM Symposium on Principles of Distributed Computing, 1990, pp. 15-29.
- [5] J. Aspnes, *Time- and space-efficient randomized consensus*, Proc. 9th ACM Symposium on Principles of Distributed Computing, 1990, pp. 325-331.

- [6] B. Chor, A. Israeli, M. Li, *On processor coordination using asynchronous hardware*, Proc. 6th ACM Symposium on Principles of Distributed Computing, pp. 86–97, 1987.
- [7] D. Dolev and N. Shavit, *Bounded Concurrent Time-Stamp Systems Are Constructible*, Proc. 21th ACM Symposium on Theory of Computing, pp. 454–466, 1989.
- [8] M.P. Herlihy, *Impossibility and Universality Results for Wait-Free Synchronization*, Proc. 7th ACM Symposium on Principles of Distributed Computing, 1988, pp. 276–290.
- [9] M.P. Herlihy, *Randomized Wait-Free Concurrent Objects*, Proc. 10th ACM Symposium on Principles of Distributed Computing, 1991.
- [10] A. Israeli and M. Li, *Bounded Time-Stamps*, Proc. 28th IEEE Symposium on Foundations of Computer Science, pp. 371–382, 1987.
- [11] L. Lamport, *On Interprocess Communication Parts I and II*, Distributed Computing, vol. 1, pp. 77–101, 1986.
- [12] M. Loui, H.H. Abu-Amara, *Memory requirements for agreement among unreliable asynchronous processes*, pp. 163–183 in: Advances in Computing Research, JAI Press, 1987.
- [13] M. Li, J. Tromp, P.M.B. Vitányi, *How to construct concurrent wait-free variables*, Tech. Rept. CS-8916, CWI, Amsterdam, April 1989. See also: pp. 488-505 in: *Proc. International Colloquium on Automata, Languages, and Programming*, Lecture Notes in Computer Science, Vol. 372, Springer Verlag, 1989.
- [14] N. Lynch and M. Tuttle, *Hierarchical correctness proofs for distributed algorithms*, Proc. 6th ACM Symposium on Principles of Distributed Computing, 1987.
- [15] G.L. Peterson, M. Fischer, *Economical solutions for the critical section problem in a distributed system*, Proc. 9th ACM Symp. on Theory of Computing, pp. 91-97, 1977.

- [16] G.L. Peterson, *Concurrent reading while writing*, ACM Transactions on Programming Languages and Systems, vol. 5, No. 1, pp. 46–55, 1983.
- [17] S. Plotkin, *Sticky bits and universality of consensus*, Proc. 8th ACM Symposium on Principles of Distributed Computing, pp. 159-175.
- [18] R. Schaffer, *On the correctness of atomic multi-writer registers*, Technical Report MIT/LCS/TM-364, MIT lab. for Computer Science, June 1988.
- [19] M. Saks, N. Shavit, and H. Woll. *Optimal time randomized consensus – making resilient algorithms fast in practice*, Proc. of SODA 90, 1990.
- [20] M.O. Rabin, *The choice coordination problem*. Acta Informatica, vol. 17, pp. 121–134, 1982.
- [21] P.M.B. Vitányi, B. Awerbuch, *Atomic Shared Register Access by Asynchronous Hardware*, Proc. 27th IEEE Symposium on Foundations of Computer Science, pp. 233–243, 1986. (Errata, Ibid.,1987)

The ITLI Prepublication Series

1990 *Logic, Semantics and Philosophy of Language*

- LP-90-01 Jaap van der Does A Generalized Quantifier Logic for Naked Infinitives
 LP-90-02 Jeroen Groenendijk, Martin Stokhof Dynamic Montague Grammar
 LP-90-03 Renate Bartsch Concept Formation and Concept Composition
 LP-90-04 Aarne Ranta Intuitionistic Categorical Grammar
 LP-90-05 Patrick Blackburn Nominal Tense Logic
 LP-90-06 Gennaro Chierchia The Variability of Impersonal Subjects
 LP-90-07 Gennaro Chierchia Anaphora and Dynamic Logic
 LP-90-08 Herman Hendriks Flexible Montague Grammar
 LP-90-09 Paul Dekker The Scope of Negation in Discourse, towards a flexible dynamic Montague grammar
 LP-90-10 Theo M.V. Janssen Models for Discourse Markers
 LP-90-11 Johan van Benthem General Dynamics
 LP-90-12 Serge Lapierre A Functional Partial Semantics for Intensional Logic
 LP-90-13 Zhisheng Huang Logics for Belief Dependence
 LP-90-14 Jeroen Groenendijk, Martin Stokhof Two Theories of Dynamic Semantics
 LP-90-15 Maarten de Rijke The Modal Logic of Inequality
 LP-90-16 Zhisheng Huang, Karen Kwast Awareness, Negation and Logical Omniscience
 LP-90-17 Paul Dekker Existential Disclosure, Implicit Arguments in Dynamic Semantics
- ML-90-01 Harold Schellinx *Mathematical Logic and Foundations* Isomorphisms and Non-Isomorphisms of Graph Models
 ML-90-02 Jaap van Oosten A Semantical Proof of De Jongh's Theorem
 ML-90-03 Yde Venema Relational Games
 ML-90-04 Maarten de Rijke Unary Interpretability Logic
 ML-90-05 Domenico Zambella Sequences with Simple Initial Segments
 ML-90-06 Jaap van Oosten Extension of Lifschitz' Realizability to Higher Order Arithmetic, and a Solution to a Problem of F. Richman
 ML-90-07 Maarten de Rijke A Note on the Interpretability Logic of Finitely Axiomatized Theories
 ML-90-08 Harold Schellinx Some Syntactical Observations on Linear Logic
 ML-90-09 Dick de Jongh, Duccio Pianigiani Solution of a Problem of David Guaspari
 ML-90-10 Michiel van Lambalgen Randomness in Set Theory
 ML-90-11 Paul C. Gilmore The Consistency of an Extended NaDSet
- CT-90-01 John Tromp, Peter van Emde Boas *Computation and Complexity Theory* Associative Storage Modification Machines
 CT-90-02 Sieger van Denneheuvel, Gerard R. Renardel de Lavalette A Normal Form for PCSJ Expressions
 CT-90-03 Ricard Gavaldà, Leen Torenvliet, Osamu Watanabe, José L. Balcázar Generalized Kolmogorov Complexity in Relativized Separations
 CT-90-04 Harry Buhrman, Edith Spaan, Leen Torenvliet Bounded Reductions
 CT-90-05 Sieger van Denneheuvel, Karen Kwast Efficient Normalization of Database and Constraint Expressions
 CT-90-06 Michiel Smid, Peter van Emde Boas Dynamic Data Structures on Multiple Storage Media, a Tutorial
 CT-90-07 Kees Doets Greatest Fixed Points of Logic Programs
 CT-90-08 Fred de Geus, Ernest Rotterdam, Sieger van Denneheuvel, Peter van Emde Boas Physiological Modelling using RL
 CT-90-09 Roel de Vrijer Unique Normal Forms for Combinatory Logic with Parallel
- Other Prepublications*
 X-90-01 A.S. Troelstra Conditional, a case study in conditional rewriting
 X-90-02 Maarten de Rijke Remarks on Intuitionism and the Philosophy of Mathematics, Revised Version
 X-90-03 L.D. Beklemishev Some Chapters on Interpretability Logic
 X-90-04 On the Complexity of Arithmetical Interpretations of Modal Formulae
 X-90-05 Valentin Shehtman Annual Report 1989
 X-90-06 Valentin Goranko, Solomon Passy Derived Sets in Euclidean Spaces and Modal Logic
 X-90-07 V. Yu. Shavrukov Using the Universal Modality: Gains and Questions
 X-90-08 L.D. Beklemishev The Lindenbaum Fixed Point Algebra is Undecidable
 X-90-09 V. Yu. Shavrukov Provability Logics for Natural Turing Progressions of Arithmetical Theories
 X-90-10 Sieger van Denneheuvel, Peter van Emde Boas On Rosser's Provability Predicate
 X-90-11 Alessandra Carbone An Overview of the Rule Language RL/1
 X-90-12 Maarten de Rijke Provable Fixed points in $\text{IA}_0 + \Omega_1$, revised version
 X-90-13 K.N. Ignatiev Bi-Unary Interpretability Logic
 X-90-14 L.A. Chagrova Dzhaparidze's Polymodal Logic: Arithmetical Completeness, Fixed Point Property, Craig's Property
 X-90-15 A.S. Troelstra Undecidable Problems in Correspondence Theory
 Lectures on Linear Logic

1991 *Logic, Semantics and Philosophy of Language*

- LP-91-01 Wiebe van der Hoek, Maarten de Rijke Generalized Quantifiers and Modal Logic
 LP-91-02 Frank Veltman Defaults in Update Semantics
 LP-91-03 Willem Groeneveld Dynamic Semantics and Circular Propositions
- ML-91-01 Yde Venema *Mathematical Logic and Foundations* Cylindric Modal Logic
 ML-91-02 Alessandro Berarducci, Rineke Verbrugge On the Metamathematics of Weak Theories
 ML-91-03 Domenico Zambella On the Proofs of Arithmetical Completeness for Interpretability Logic
 ML-91-04 Raymond Hoofman, Harold Schellinx Collapsing Graph Models by Preorders
 ML-91-05 A.S. Troelstra History of Constructivism in the Twentieth Century
 ML-91-06 Inge Bethke Finite Type Structures within Combinatory Algebras
 ML-91-07 Yde Venema Modal Derivation Rules
- CT-91-01 Ming Li, Paul M.B. Vitányi *Computation and Complexity Theory* Kolmogorov Complexity Arguments in Combinatorics
 CT-91-02 Ming Li, John Tromp, Paul M.B. Vitányi How to Share Concurrent Wait-Free Variables
 CT-91-03 Ming Li, Paul M.B. Vitányi Average Case Complexity under the Universal Distribution Equals Worst Case Complexity
 CT-91-04 Sieger van Denneheuvel, Karen Kwast Weak Equivalence
 CT-91-05 Sieger van Denneheuvel, Karen Kwast Weak Equivalence for Constraint Sets
 CT-91-06 Edith Spaan Census Techniques on Relativized Space Classes
 CT-91-07 Karen L. Kwast The Incomplete Database
 CT-91-08 Kees Doets Levationis Laus
 CT-91-09 Ming Li, Paul M.B. Vitányi Combinatorial Properties of Finite Sequences with high Kolmogorov Complexity
 CT-91-10 John Tromp, Paul Vitányi A Randomized Algorithm for Two-Process Wait-Free Test-and-Set
- Other Prepublications*
 X-91-01 Alexander Chagrov, Michael Zakharyashev The Disjunction Property of Intermediate Propositional Logics
 X-91-02 Alexander Chagrov, Michael Zakharyashev On the Undecidability of the Disjunction Property of Intermediate Propositional Logics
 X-91-03 V. Yu. Shavrukov Subalgebras of Diagonalizable Algebras of Theories containing Arithmetic
 X-91-04 K.N. Ignatiev Partial Conservativity and Modal Logics
 X-91-05 Johan van Benthem Temporal Logic
 X-91-06 Annual Report 1990
 X-91-07 A.S. Troelstra Lectures on Linear Logic, Errata and Supplement
 X-91-08 Giorgie Dzhaparidze Logic of Tolerance
 X-91-09 L.D. Beklemishev On Bimodal Provability Logics for Π_1 -axiomatized Extensions of Arithmetical Theories
 X-91-10 Michiel van Lambalgen Independence, Randomness and the Axiom of Choice
 X-91-11 Michael Zakharyashev Canonical Formulas for K4. Part I: Basic Results
 X-91-12 Herman Hendriks Flexibele Categoriale Syntaxis en Semantiek: de proefschriften van Frans Zwartz en Michael Moortgat
 X-91-13 Max I. Kanovich The Multiplicative Fragment of Linear Logic is NP-Complete