

PROGRAMMING MULTI-CORE AND MANY-CORE COMPUTING SYSTEMS



PROGRAMMING MULTI-CORE AND MANY-CORE COMPUTING SYSTEMS

**in the Wiley Series on Parallel and
Distributed Computing**

Sabri Pllana et al.

A JOHN WILEY & SONS, INC., PUBLICATION

Copyright ©2010 by John Wiley & Sons, Inc. All rights reserved.

Published by John Wiley & Sons, Inc., Hoboken, New Jersey.
Published simultaneously in Canada.

No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning, or otherwise, except as permitted under Section 107 or 108 of the 1976 United States Copyright Act, without either the prior written permission of the Publisher, or authorization through payment of the appropriate per-copy fee to the Copyright Clearance Center, Inc., 222 Rosewood Drive, Danvers, MA 01923, (978) 750-8400, fax (978) 646-8600, or on the web at www.copyright.com. Requests to the Publisher for permission should be addressed to the Permissions Department, John Wiley & Sons, Inc., 111 River Street, Hoboken, NJ 07030, (201) 748-6011, fax (201) 748-6008.

Limit of Liability/Disclaimer of Warranty: While the publisher and author have used their best efforts in preparing this book, they make no representations or warranties with respect to the accuracy or completeness of the contents of this book and specifically disclaim any implied warranties of merchantability or fitness for a particular purpose. No warranty may be created or extended by sales representatives or written sales materials. The advice and strategies contained herein may not be suitable for your situation. You should consult with a professional where appropriate. Neither the publisher nor author shall be liable for any loss of profit or any other commercial damages, including but not limited to special, incidental, consequential, or other damages.

For general information on our other products and services please contact our Customer Care Department with the U.S. at 877-762-2974, outside the U.S. at 317-572-3993 or fax 317-572-4002.

Wiley also publishes its books in a variety of electronic formats. Some content that appears in print, however, may not be available in electronic format.

Library of Congress Cataloging-in-Publication Data:

Title, etc
Printed in the United States of America.

10 9 8 7 6 5 4 3 2 1





CONTENTS

1	Lock-free Concurrent Data Structures	1
	Daniel Cederman, Anders Gidenstam, Phuong Ha, Håkan Sundell, Marina Papatrantafileou and Philippos Tsigas	
1.1	Introduction	1
1.2	Synchronization Primitives	3
1.2.1	Fundamental synchronization primitives	3
1.2.2	Synchronization power	4
1.2.3	Scalability and Combinability	4
1.2.4	Multi-word Primitives	5
1.3	Lock-Free Data Structures	5
1.3.1	Overview	7
1.3.2	Producer-Consumer Collections	7
1.3.3	Lists	9
1.3.4	Sets and Dictionaries	10
1.4	Memory Management for Concurrent Data-Structures	11
1.4.1	Memory Allocation	11
1.4.2	Memory Reclamation	12
		vii

viii CONTENTS

1.5	Graphics Processors	14
1.5.1	Data Parallel Model	14
1.5.2	New Algorithmic Design	15
	References	16

CHAPTER 1

LOCK-FREE CONCURRENT DATA STRUCTURES

DANIEL CEDERMAN¹, ANDERS GIDENSTAM², PHUONG HA³, HÅKAN SUNDELL²,
MARINA PAPATRIANTAFILOU¹ AND PHILIPPAS TSIGAS¹

¹ Chalmers University of Technology, Sweden

² University of Borås, Sweden

³ University of Tromsø, Norway

Algorithms + Data Structures = Programs

—Niklaus Wirth

1.1 INTRODUCTION

Concurrent data structures are the data sharing side of parallel programming. Data structures give the means to the program to store data but also provide operations to the program to access and manipulate these data. These operations are implemented through algorithms that have to be efficient. In the sequential setting, data structures are crucially important for the performance of the respective computation. In the parallel programming setting, their importance becomes more crucial because of the increased use of data and resource sharing for utilizing parallelism. In parallel programming, computations are split into subtasks in order to introduce parallelization at the control/computation level. To utilize this opportunity of concurrency, subtasks share data and various resources (dictionaries, buffers, and so forth). This makes

it possible for logically independent programs to share various resources and data structures. A subtask that wants to update a data structure, say add an element into a dictionary, that operation may be logically independent of other subtasks that use the same dictionary.

Concurrent data structure designers are striving to maintain consistency of data structures while keeping the use of mutual exclusion and expensive synchronization to a minimum, in order to prevent the data structure from becoming a sequential bottleneck. Maintaining consistency in the presence of many simultaneous updates is a complex task. Standard implementations of data structures are based on locks in order to avoid inconsistency of the shared data due to concurrent modifications. In simple terms, a single lock around the whole data structure may create a bottleneck in the program where all of the tasks serialize, resulting in a loss of parallelism because too few data locations are concurrently in use. Deadlocks, priority inversion, and convoying are also side-effects of locking. The risk for deadlocks makes it hard to compose different blocking data structures since it is not always possible to know how closed source libraries do their locking. It is worth noting that in graphics processors (GPUs) locks are not recommended for designing concurrent data structures. GPUs prior to the NVIDIA Fermi architecture do not have writable caches, so for those GPUs, repeated checks to see if a lock is available or not require expensive repeated accesses to the GPU's main memory. While Fermi GPUs do support writable caches, there is no guarantee that the thread scheduler will be fair, which can make it difficult to write deadlock-free locking code. OpenCL explicitly disallows locks for these and other reasons.

Lock-free implementations of data structures support concurrent access. They do not involve mutual exclusion and make sure that all steps of the supported operations can be executed concurrently. Lock-free implementations employ an optimistic conflict control approach, allowing several processes to access the shared data object at the same time. They suffer delays only when there is an actual conflict between operations that causes some operations to retry. This feature allows lock-free algorithms to scale much better when the number of processes increases.

An implementation of a data structure is called *lock-free* if it allows multiple processes/threads to access the data structure concurrently and also guarantees that at least one operation among those finishes in a finite number of its own steps regardless of the state of the other operations. A consistency (safety) requirement for lock-free data structures is *linearizability* [44], which ensures that each operation on the data appears to take effect instantaneously during its actual duration and the effect of all operations are consistent with the object's sequential specification. Lock-free data structures offer several advantages over their blocking counterparts, such as being immune to deadlocks, priority inversion, and convoying, and have been shown to work well in practice in many different settings [92, 84]. They have been included in Intel's Threading Building Blocks Framework [75] the NOBLE library [84] and the Java concurrency package [56], and will be included in the forthcoming parallel extensions to the Microsoft .NET Framework [69]. They have also been of interest to designers of languages such as C++ [12] and Java [56].

This chapter has two goals. The first and main goal is to provide a sufficient background and intuition to help the interested reader to navigate in the complex research area of lock-free data structures. The second goal is to offer the programmer familiarity to the subject that will allow her to use truly concurrent methods.

The chapter is structured as follows. First we discuss the fundamental and commonly-supported synchronization primitives on which efficient lock-free data structures rely. Then we give an overview of the research results on lock-free data structures that appeared in the literature with a short summary for each of them. The problem of managing dynamically-allocated memory in lock-free concurrent data structures and general concurrent environments is discussed separately. Following this is a discussion on the idiosyncratic architectural features of graphics processors that are important to consider when designing efficient lock-free concurrent data structures for this emerging area.

1.2 SYNCHRONIZATION PRIMITIVES

To synchronize processes efficiently, multi-/many-core systems usually support certain synchronization primitives. This section discusses the fundamental synchronization primitives, which typically read the value of a *single* memory word, modify the value and write the new value back to the word *atomically*.

1.2.1 Fundamental synchronization primitives

The definitions of the primitives are described in Figure 1.1, where x is a memory word, v , old , new are values and op can be operators *add*, *sub*, *or*, *and* and *xor*. Operations between angle brackets $\langle \rangle$ are executed atomically.

TAS (x) /* test-and-set, init: $x \leftarrow 0$ */ $\langle oldx \leftarrow x; x \leftarrow 1; \mathbf{return} \ oldx; \rangle$	LL (x) /* load-linked */ $\langle \mathit{return} \ \mathit{the\ value\ of\ } x \ \mathit{so\ that\ it\ may\ be\ subsequently\ used\ with\ SC} \rangle$
FAO (x, v) /* fetch-and-op */ $\langle oldx \leftarrow x; x \leftarrow op(x, v); \mathbf{return} \ oldx; \rangle$	SC (x, v) /* store-conditional */ $\langle \mathbf{if} \ (\mathit{no\ process\ has\ written\ to\ } x \ \mathit{since\ the\ last\ LL}(x)) \ \{ x \leftarrow v; \mathbf{return}(true); \}$ $\mathbf{else} \ \mathbf{return}(false); \rangle$
CAS (x, old, new) /* compare-and-swap */ $\langle \mathbf{if}(x = old) \ \{ x \leftarrow new; \mathbf{return}(true); \}$ $\mathbf{else} \ \mathbf{return}(false); \rangle$	

Figure 1.1 Synchronization primitives

Note that there is a problem called the *ABA problem* that may occur with the *CAS* primitive. The reason is that the *CAS* operation can not detect if a variable was read to be A and then later changed to B and then back to A by some concurrent processes. The *CAS* primitive will perform the update even though this might not be intended

by the algorithm’s designer. The *LL/SC* primitives can instead detect any concurrent update on the variable between the time interval of a *LL/SC* pair, independent of the value of the update.

1.2.2 Synchronization power

The primitives are classified according to their synchronization power or *consensus number* [57], which is, roughly speaking, the maximum number of processes for which the primitives can be used to solve a *consensus problem* in a fault tolerant manner. In the consensus problem, a set of n asynchronous processes, each with a given input, communicate to achieve an agreement on one of the inputs. A primitive with a consensus number n can achieve consensus among n processes even if up to $n - 1$ processes stop [93].

According to the consensus classification, read/write registers have consensus number 1, i.e. they cannot tolerate any faulty processes in the consensus setting. There are some primitives with consensus number 2 (e.g. *test-and-set (TAS)* and *fetch-and-op (FAO)*) and some with infinite consensus number (e.g. *compare-and-swap (CAS)* and *load-linked/store-conditional (LL/SC)*). It has been proven that a primitive with consensus number n cannot implement a primitive with a higher consensus number in a system of more than n processes [57]. For example, the *test-and-set* primitive, whose consensus number is two, cannot implement the *compare-and-swap* primitive, whose consensus number is unbounded, in a system of more than two processes.

1.2.3 Scalability and Combinability

As many-core architectures with thousands of cores are expected to be our future chip architectures [5], synchronization primitives that can support scalable thread synchronization for such large-scale architectures are desired. In addition to synchronization power criterion, synchronization primitives can be classified by their scalability or *combinability* [54]. Primitives are combinable if their memory requests to the same memory location (arriving at a switch of the processor-to-memory interconnection network) can be combined into *only one* memory request. Separate replies to the original requests are later created from the reply to the combined request (at the switch). The combining technique has been implemented in the NYU Ultracomputer [30] and the IBM RP3 machine [73], and has been shown to be a scalable technique for large-scale multiprocessors to alleviate the performance degradation due to a synchronization “hot spot”. The set of combinable primitives includes *test-and-set*, *fetch-and-op* (where *op* is an associative operation or boolean operation), blocking full-empty bits [54] and non-blocking full-empty bits [36]. For example, two consecutive requests *fetch-and-add(x, a)* and *fetch-and-add(x, b)* can be combined into a *single* request *fetch-and-add(x, a + b)*. When receiving a reply *oldx* to the combined request *fetch-and-add(x, a + b)*, the switch at which the requests were combined, creates a reply *oldx* to the first request *fetch-and-add(x, a)* and a reply (*oldx + a*) to the successive request *fetch-and-add(x, b)*.

The *CAS* primitives are not combinable since the success of a $CAS(x, a, b)$ primitive depends on the current value of the memory location x . For m -bit locations (e.g. 64-bit words), there are 2^m possible values and therefore, a combined request that represents k $CAS(x, a, b)$ requests, $k < 2^m$, must carry as many as k different checking-values a and k new values b . The *LL/SC* primitives are not combinable either since the success of a *SC* primitive depends on the state of its reservation bit at the memory location that has been set previously by the corresponding *LL* primitive. Therefore, a combined request that represents k *SC* requests (from different processes/processors) must carry as many as k store values.

1.2.4 Multi-word Primitives

Although the *single-word* hardware primitives are conceptually powerful enough to support higher-level synchronization, from the programmer's point of view they are not as convenient as *multi-word* primitives. The multi-word primitives can be built in hardware [52, 16, 11], or in software (in a lock-free manner) using single-word hardware primitives [3, 19, 34, 50, 71, 79]. Sun's third generation chip-multithreaded (CMT) processor called Rock is the first processor supporting transactional memory in hardware [11]. The transactional memory is supported by two new instructions *checkpoint* and *commit*, in which *checkpoint* denotes the beginning of a transaction and *commit* denotes the end of the transaction. If the transaction succeeds, the memory accesses within the transaction take effect atomically. If the transaction fails, the memory accesses have no effect.

Another emerging construct is the Advanced Synchronization Facility (ASF), an experimental AMD64 extension that AMD's Operating System Research Center develops to support lock-free data structures and software transactional memory [16]. ASF is a simplified hardware transactional memory in which all memory objects to be protected should be statically specified before transaction execution. Processors can protect and speculatively modify up to 8 memory objects of cache-line size. There is also research on new primitives aiming at identifying new efficient and powerful primitives, with the non-blocking full/empty bit (NB-FEB) being an example that was shown to be as powerful as *CAS* or *LL/SC* [36].

1.3 LOCK-FREE DATA STRUCTURES

The main characterization on which one can classify the various implementations of lock-free data structures available in the literature, is what *abstract data type* that it intends to implement. For each abstract data type there are usually numerous implementations, each motivated by some specific targeted purposes, where each implementation is characterized by the various properties that it fulfills to different amounts. As many of these properties are orthogonal, for each specific implementation, one or more properties are often strengthened at the cost of some others. Some of the most important properties that differentiate the various lock-free data structure implementations in the literature are:

Semantic fulfillments Due to the complexity of designing lock-free data structures it might not be possible to support all operations normally associated with a certain abstract data type. Hence, some algorithms omit a subset of the normally required operations and/or support operations with a modified semantics.

Time complexity Whether an operation can terminate in a time (without considering concurrency) that is linearly or logarithmically related to e.g. the size of the data structure, can have significant impact on performance. Moreover, whether the maximum execution time can be determined at all or if it can be expected in relation to the number of concurrent threads is of significant importance to time-critical systems (e.g. real-time systems).

Scalability Scalability means showing some performance gain with increasing number of threads. Synchronization primitives are normally not scalable in themselves; therefore it is important to avoid unnecessary synchronization. Israeli and Rappoport [50] have defined the term *disjoint-access-parallelism* to identify algorithms that do not synchronize on data that is not logically involved simultaneously in two or more concurrent operations.

Dynamic capacity In situations where it can be difficult to determine the maximum number of items that will be stored in a data structure, it is necessary that the data structure can dynamically allocate more memory when the current capacity is about to be exceeded. If the data structure is based on statically allocated storage, capacity is fixed throughout the lifetime of the data structure.

Space complexity Some algorithms can guarantee an upper bound of memory required, while some others can transiently need an indefinite amount depending on the concurrent operations' invocation order, and can thus not be deterministically determined.

Concurrency limitations Due to the limitations (e.g. consensus number) of the chosen synchronization primitives, some or all operations might not allow more than a certain number of concurrent invocations.

Synchronization primitives Contemporary multi-core and many-core systems typically only support single-word *CAS* or weak and non-nestable variants of *LL/SC* (cf. section 1.2). However, many algorithms for lock-free data structure depend on more advanced primitives as e.g. double-word *CAS* (called e.g. *DCAS* or *CAS2*), ideal *LL/VL/SC* or even more complex primitives. These algorithms then need (at least) one additional abstraction layer for actual implementation, where these more advanced primitives are implemented in software using another specific algorithm. The *LL/VL/SC* primitives can be implemented e.g. by *CAS* [50, 3, 70, 51, 65]. Multi-word *CAS* (called e.g. *MWCAS* or *CASN*) can be implemented e.g. by *CAS* [38, 82] or by *LL/VL/SC* [50, 3, 79, 71, 34].

Reliability Some algorithms try to avoid the ABA problem by the means of e.g. version counters. As these counters are bounded and can overflow, there is a potential risk of the algorithm to actually perform incorrectly and possibly cause inconsistencies. Normally, by design this risk can be kept low enough that it fits for practical purposes, although the risk increases as the computational speed increases. Often, version counters can be removed altogether by the means of proper memory management.

Compatibility and Dependencies Some algorithms only work together with certain memory allocators and reclamation schemes, specific types (e.g. real-time) of system-level process scheduler, or require software layers or semantic constructions only found in certain programming languages (e.g. Java).

1.3.1 Overview

The following sections include a systematic overview of the research result in the literature. For a more in-depth look and a case-study in the design of a lock-free data structure and how it can be used in practice, we would like to refer the reader to our chapter in “GPU Computing Gems” [10], which describes in detail how to implement a lock-free work-stealing deque and the reasoning behind the design decisions.

1.3.2 Producer-Consumer Collections

A common approach to parallelizing applications is to divide the problem into separate threads that act as either producers or consumers. The problem of synchronizing these threads and streaming of data items between them, can be alleviated by utilizing a shared collection data structure.

Bag The Bag abstract data type is a collection of items in which items can be stored and retrieved in any order. Basic operations are *Add* (add an item) and *TryRemoveAny* (remove an arbitrary chosen item). *TryRemoveAny* returns the item removed. Data structures with similar semantics are also called *buffer*, unordered collection, unordered queue, pool, and pile in the literature.

All lock-free stacks, queues and deques implicitly implements the selected bag semantics. Afek et al. [1] presented an explicit pool data structure. It is lock-free, although not linearizable, utilizes distributed storage and is based on randomization to establish a probabilistic level of disjoint-access-parallelism.

In [26, 27] a data structure called *flat-sets*, was introduced and used as a building block in the concurrent memory allocation service. This is a bag-like structure that supports lock-free insertion and removal of items as well as an “inter-object” operation, for moving an item from one flat-set to another in a lock-free and linearizable manner, thus offering the possibility of combining data structures.

In [83] a lock-free bag implementation is presented; the algorithm supports multiple producers and multiple consumers, as well as dynamic collection sizes. To handle concurrency efficiently, the algorithm was designed to optimize for disjoint-access-parallelism for the supported semantics.

Stack The Stack abstract data type is a collection of items in which only the most recently added item may be removed. The latest added item is at the top. Basic operations are *Push* (add to the top) and *Pop* (remove from the top). *Pop* returns the item removed. The data structure is also known as a “last-in, first-out” or LIFO buffer.

Treiber presented a lock-free stack (a.k.a. IBM Freelist) based on linked lists, which was later efficiently fixed from the ABA problem by Michael [64]. Also

Valois [95] presented a lock-free implementation that uses the *CAS* atomic primitive. Hendler et al. [41] presented an extension where randomization and elimination are used for increasing scalability when contention is detected on the *CAS* attempts.

Queue The Queue abstract data type is a collection of items in which only the earliest added item may be accessed. Basic operations are *Enqueue* (add to the tail) and *Dequeue* (remove from the head). Dequeue returns the item removed. The data structure is also known as a “first-in, first-out” or FIFO buffer.

Lamport [55] presented a lock-free (actually wait-free) implementation of a queue based on a static array, with a limited concurrency supporting only one producer and one consumer. Giacomoni et al. [23] presented a cache-aware modification which, instead of using shared head and tail indices, synchronize directly on the array elements. Herman and Damian-Iordache [46] outlined a wait-free implementation of a shared queue for any number of threads, although non-practical due to its high time complexity and limited capacity.

Gong and Wing [29] and later Shann et al. [78] presented a lock-free shared queue based on a cyclic array and the *CAS* primitive, though with the drawback of using version counters, thus requiring double-width *CAS* for storing actual items. Tsigas and Zhang [90] presented a lock-free extension of [55] for any number of threads where synchronization is done both on the array elements and the shared head and tail indices using *CAS*, and the ABA problem is avoided by exploiting two (or more) null values.

Valois [94, 95] makes use of linked lists in his lock-free implementation which is based on the *CAS* primitive. Prakash et al. [74] also presented an implementation using linked lists and the *CAS* primitive, although with the drawback of using version counters and having low scalability. Michael and Scott [68] presented a lock-free queue that is more efficient, synchronizing via the shared head and tail pointers as well via the next pointer of the last node. Moir et al. [72] presented an extension where elimination is used as a back-off strategy when contention on *CAS* is noticed, although elimination is only possible when the queue contains very few items. Hoffman et al. [47] takes another approach for a back-off strategy by allowing concurrent *Enqueue* operations to insert the new node at adjacent positions in the linked list if contention is noticed. Gidenstam et al. [28] combines the efficiency of using arrays and the dynamic capacity of using linked lists, by providing a lock-free queue based on linked lists of arrays, all updated using *CAS* in a cache-aware manner.

Deque The Deque (or doubly-ended queue) abstract data type is a combination of the stack and the queue abstract data types. The data structure is a collection of items in which the earliest as well as the latest added item may be accessed. Basic operations are *PushLeft* (add to the head), *PopLeft* (remove from the head), *PushRight* (add to the tail), and *PopRight* (remove from the tail). *PopLeft* and *PopRight* return the item removed.

Large efforts have been put on the work on so called work-stealing dequeues. These data structures only support three operations and with a limited level of concurrency, and are specifically aimed for scheduling purposes. Arora et al. [4] presented a

lock-free work-stealing deque implementation based on the *CAS* atomic primitive. Hendler et al. [40] improved this algorithm to also handle dynamic sizes.

Several lock-free implementations of the deque abstract data type for general purposes, although based on the non-available *CAS2* atomic primitive, have been published in the literature [31, 2, 13, 58, 6]. Michael [63] presented a lock-free deque implementation based on the *CAS* primitive, although not supporting any level of disjoint-access-parallelism. Sundell and Tsigas [88] presented a lock-free implementation that allows both disjoint-access-parallelism as well as dynamic sizes using the standard *CAS* atomic primitive.

Priority Queue The Priority Queue abstract data type is a collection of items which can efficiently support finding the item with the highest priority. Basic operations are *Insert* (add an item), *FindMin* (finds the item with minimum (or maximum) priority), and *DeleteMin* (removes the item with minimum (or maximum) priority). *DeleteMin* returns the item removed.

Israeli and Rappoport [49] have presented a wait-free algorithm for a shared priority queue, that requires the non-available multi-word *LL/VL/SC* atomic primitives. Greenwald [31] has presented an outline for a lock-free priority queue based on the non-available *CAS2* atomic primitive. Barnes [7] presented an incomplete attempt for a lock-free implementation that uses atomic primitives available on contemporary systems. Sundell and Tsigas [87] presented the first lock-free implementation of a priority queue based on skip lists and the *CAS* atomic primitive.

1.3.3 Lists

The List abstract data type is a collection of items where two items are related only with respect to their relative position to each other. The data structure should efficiently support traversals among the items. Depending on what type of the underlying data structure, e.g. *arrays* or *linked lists*, different strengths of traversal functionality are supported.

Array List implementations based on the fundamental array data structure can support traversals to absolute index positions. Higher level abstractions as extendable arrays are in addition supporting stack semantics. Consequently, the Array abstract data type would support the operations *ReadAt* (read the element at index), *WriteAt* (write the element at index), *Push* (add to the top) and *Pop* (remove from the top). *Pop* returns the item removed.

A lock-free extendable array for practical purposes has been presented by Dechev et al. [12].

Linked List In a concurrent environment with List implementations based on linked lists, traversals to absolute index positions are not feasible. Consequently, traversals are only supported relatively to a current position. The current position is maintained by the cursor concept, where each handle (i.e. thread or process) maintains one independent cursor position. The first and last cursor positions do not

refer to real items, but are instead used as end markers, i.e. before the first item or after the last item. Basic operations are *InsertAfter* (add a new item after the current), *Delete* (remove the current item), *Read* (inspect the current item), *Next* (traverse to the item after the current), *First* (traverse to the position before the first item). Additional operations are *InsertBefore* (add a new item before the current), *Previous* (traverse to the item before the current), and *Last* (traverse to the position after the last item).

Lock-free implementations of the singly-linked list based on the *CAS* atomic primitive and with semantics suitable for the Dictionary abstract type rather than the List has been presented by Harris [39], Michael [61], and Fomitchev and Ruppert [18]. Greenwald [32] presented a doubly-linked list implementation of a dictionary based on the non-available *CAS2* atomic primitive. Attiya and Hillel [6] presented a *CAS2*-based implementation that also supports disjoint-access-parallelism. Valois [95] outlined a lock-free doubly-linked list implementation with all list semantics except delete operations. A more general doubly-linked list implementation supporting general list semantics was presented by Sundell and Tsigas [88].

1.3.4 Sets and Dictionaries

The Set abstract data type is a collection of special items called *keys*, where each key is unique and can have at most one occurrence in the set. Basic operations are *Add* (adds the key), *ElementOf* (checks if key is present), and *Delete* (removes the key).

The Dictionary abstract data type is a collection of items where each item is associated with a unique key. The data structure should efficiently support finding the item associated with the specific key. Basic operations are *Insert* (add an item associated with a key), *Find* (finds the item associated with a certain key), and *Delete* (removes the item associated with a certain key). *Delete* returns the item removed. In concurrent environments, an additional basic operation is the *Update* (re-assign the association of a key with another item) operation.

Implementations of Sets and Dictionaries are often closely related in a way that most implementations of a set can be extended to also support dictionary semantics in a straight forward manner. However, the *Update* operation mostly needs specific care in the fundamental part of the algorithmic design to be linearizable. Non-blocking implementations of sets and dictionaries are mostly based on hash-tables or linked lists as done by Valois [95]. The path using concurrent linked lists was improved by Harris [39]. Other means to implement sets and dictionaries are the skip-list and tree data structures.

Skip-List Valois [95] outlined an incomplete idea of how to design a concurrent skip list. Sundell and Tsigas presented a lock-free implementation of a skip list in the scope of priority queues [85, 87] as well as dictionaries [86, 81] using the *CAS* primitive. Similar constructions have appeared in the literature by Fraser [20], and Fomitchev and Ruppert [18].

Hash-Table Michael [61] presented a lock-free implementation of the set abstract data type based on a hash-table with its chaining handled by an improved linked

list compared to [39]. To a large part, its high efficiency is thanks to the memory management scheme applied. The algorithm was improved by Shalev and Shavit[77] to also handle dynamic sizes of the hash-table's underlying array data structure. Greenwald [32] have presented a dictionary implementation based on chained hash-tables and the non-available *CAS2* atomic primitive.

Gao et al. [21] presented a lock-free implementation of the dictionary abstract data type based on a hash-table data structure using open addressing. The hash-table is fully dynamic in size, although its efficiency is limited by its relatively complex memory management.

Tree Tsay and Li [89] presents an approach for designing lock-free implementations of a tree data structure using the *LL/VL/SC* atomic primitives and extensive copying of data. However, the algorithm is not provided with sufficient evidence for showing linearizability. Ellen et al. [17] presented a lock-free implementation of the set abstract data type based on a binary tree data structure using the *CAS* atomic primitive. Spiegel and Reynolds [80] presents a lock-free implementation of the set abstract data type based on a skip-tree and the *CAS* atomic primitive.

1.4 MEMORY MANAGEMENT FOR CONCURRENT DATA-STRUCTURES

The problem of managing dynamically allocated memory in a concurrent environment has two parts, keeping track of the free memory available for allocation and safely reclaim allocated memory when it is no longer in use, i.e. memory allocation and memory reclamation.

1.4.1 Memory Allocation

A memory allocator manages a pool of memory (heap), e.g. a contiguous range of addresses or a set of such ranges, keeping track of which parts of that memory are currently given to the application and which parts are unused and can be used to meet future allocation requests from the application. A traditional (such as the "libc" malloc) general purpose memory allocator is not allowed to move or otherwise disturb memory blocks that are currently owned by the application.

Some of the most important properties that distinguish memory allocators for concurrent applications in the literature are:

Fragmentation To minimize fragmentation is to minimize the amount of free memory that cannot be used (allocated) by the application due to the size of the memory blocks.

False-sharing False sharing is when different parts of the same cache-line are allocated to separate objects that end up being used by threads running on different processors.

Efficiency and scalability The concurrent memory allocator should be as fast as a good sequential one when executed on a single processor and its performance should scale with the load in the system.

Here we focus on lock-free memory allocators but there is also a considerable number of lock-based concurrent memory allocators in the literature.

Early work on lock-free memory allocation is the work on non-blocking operating systems by Massalin and Pu [60, 59] and Greenwald and Cheriton [33, 31]. Dice and Garthwaite [15] presented LFMalloc, a memory allocator based on the architecture of the Hoard lock-based concurrent memory allocator [8] but with reduced use of locks. Michael [66] presented a fully lock-free allocator, also loosely based on the Hoard architecture. Gidenstam et al. [26] presented NBmalloc, another lock-free memory allocator loosely based on the Hoard architecture. NBmalloc is designed from the requirement that the first-remove-then-insert approach to moving references to large internal blocks of memory (superblocks) around should be avoided and therefore introduces and uses a move operation that can move a reference between different internal data-structures atomically. Schneider et al. [76] presented Streamflow, a lock-free memory allocator that has improved performance over previous solutions due to allowing thread local allocations and deallocations without synchronization.

1.4.2 Memory Reclamation

To manage dynamically allocated memory in non-blocking algorithms is difficult due to overlapping operations that might read, change or *dereference* (i.e. follow) references to dynamically allocated blocks of memory concurrently. One of the most problematic cases is when a slow process dereferences a pointer value that it previously read from a shared variable. This dereference of the pointer value could occur an arbitrarily long time after the shared pointer holding that value was overwritten and the memory designated by the pointer removed from the shared data structure. Consequently it is impossible to safely free or reuse the block of memory designated by this pointer value until we are sure that there are no such slow processes with pointers to that block.

There are several reclamation schemes in the literature with a wide and varying range of properties:

I. Safety of local references For local references, which are stored in private variables accessible only by one thread, to be safe the memory reclamation scheme must guarantee that a dynamically allocated node is never reclaimed while there still are local references pointing to it.

II. Safety of shared references Additionally, a memory reclamation scheme could also guarantee that it is always safe for a thread to dereference any shared references located within a dynamic node the thread has a local reference to. Property I alone does not guarantee this, since for a node that has been deleted but cannot be reclaimed yet any shared references within it could reference nodes that have been deleted and reclaimed since the node was removed from the data structure.

III. Automatic or explicit deletion A dynamically allocated node could either be reclaimed automatically when it is no longer accessible through any local or shared reference, that is, the scheme provides *automatic garbage collection*, or the user algorithm or data structure could be required to explicitly tell the memory reclamation scheme when a node is removed from the active data structure and should be reclaimed

	Property II	Property III	Property IV	Property V
Michael [62, 64]	No	Explicit	Yes	Yes
Herlihy et al. [43]	No	Explicit	Yes	No
Valois et al. [95, 67]	Yes	Automatic	No	Yes
Detlefs et al. [14]	Yes	Automatic	Yes	No
Herlihy et al. [42]	Yes	Automatic	Yes	No
Gidenstam et al. [24, 25]	Yes	Explicit	Yes	Yes
Fraser [20]	Yes	Explicit	Yes	Yes
Herlihy et al. [45]	Yes	Automatic	Integrated	Yes?
Gao et al. [22]	Yes	Automatic	Integrated	Yes?

Table 1.1 Properties of different approaches to non-blocking memory reclamation.

as soon as it has become safe. While automatic garbage collection is convenient for the user, explicit deletion by the user gives the reclamation scheme more information to work with and can help to provide stronger guarantees, e.g. bounds on the amount of deleted but yet unreclaimed memory.

IV. Requirements on the memory allocator Some memory reclamation schemes require special properties from the memory allocator, like, for example, that each allocable node has a permanent (i.e. for the rest of the system’s lifetime) reference counter associated with it. Other schemes are compatible with the well-known and simple *allocate/free* allocator interface where the node has ceased to exist after the call to *free*.

V. Required synchronization primitives Some memory reclamation schemes are defined using synchronization primitives that few if any current processor architectures provide in hardware, such as for example *double word CAS*, which then have to be implemented in software often adding considerable overhead. Other schemes make do with *single word CAS*, *single word LL/VL/SC* or even just reads and writes alone.

The properties of the memory reclamation schemes discussed here are summarized in Table 1.1. One of the most important is Property II, which many lock-free algorithms and data structures need. Among the memory reclamation schemes that guarantee Property II we have the following ones, all based on reference counting: Valois et al. [95, 67], Detlefs et al. [14], Herlihy et al. [42] and Gidenstam et al. [24, 25] and the potentially blocking epoch-based scheme by Fraser [20].

On the other hand, for data structures that do not need Property II, for example stacks, the use of a reclamation scheme that does not provide this property has significant potential to offer reduced overhead compared to the stronger schemes. Among these memory reclamation schemes we have the non-blocking ones by Michael [62, 64] and Herlihy et al. [43].

Fully Automatic Garbage Collection. A fully automatic garbage collector provides property I, II and III with automatic deletion.

There are some lock-free garbage collectors in the literature. Herlihy and Moss presented a lock-free copying garbage collector in [45]. Gao et al. [22] presented a lock-free Mark & Sweep garbage collector and Kliot et al. [53] presented a lock-free stack scanning mechanism for concurrent garbage collectors.

1.5 GRAPHICS PROCESSORS

Currently the two most popular programming environments for general purpose computing for graphics processors are CUDA and OpenCL. Neither provides any direct support for locks, and it is unlikely that this will change in the future. Concurrent data structures that are used on graphics processors will therefore have to be lock-free.

While graphics processors share many features with conventional processors, and many lock-free algorithms can be ported directly, there are some differences that are important to consider, if one also wants to maintain or improve the scalability and throughput of the algorithms.

1.5.1 Data Parallel Model

A graphics processor consists of a number of multiprocessors that can execute the same instruction on multiple data, known as SIMD computing. Concurrent data structures are, as the name implies, designed to support multiple concurrent operations, but when used on a multiprocessor they also need to support concurrent instructions within an operation. This is not straightforward, as most have been designed for scalar processors. Considering that SIMD instructions play an instrumental role in the parallel performance offered by the graphics processor, it is imperative that this issue be addressed.

Graphics processor have a wide memory bus and a high memory bandwidth, which makes it possible to quickly transfer data from the memory to the processor and back. The hardware is also capable of coalescing multiple small memory operations into a single, large, atomic memory operation. As a single large memory operation can be performed faster than many small, this should be taken advantage of in the algorithmic design of the data structure.

The cache in graphics processors is smaller than on conventional SMP processors and in many cases non-existent. The memory latency is instead masked by utilizing thousands of threads and by storing data temporarily in a high-speed multiprocessor local memory area. The high number of threads reinforces the importance of the data structure being highly scalable.

The scheduling of threads on a graphics processor is commonly being performed by the hardware. Unfortunately, the scheme used is often undocumented, thus there is no guarantee that it will be fair. This makes the use of algorithms with blocking behavior risky. For example, a thread holding a lock could be indefinitely swapped out in favor of another thread waiting for the same lock, resulting in a livelock situation. Lock-freeness is thus a must.

Of a more practical concern is the fact that a graphics processor often lacks stacks, making recursive operations more difficult. The lack of a joint address space between the GPU and the CPU also complicates the move of data from the CPU to the graphics processor, as all pointers in the data structure have to be rebased when moved to a new address.

1.5.2 New Algorithmic Design

The use of SIMD instructions means that if multiple threads write to the same memory location, only one (arbitrary) thread can succeed. Thus, allowing threads that will be combined to a SIMD unit by the hardware to concurrently try to enqueue an item to the same position in a queue, will with all likelihood be unnecessarily expensive, as only one thread can succeed in enqueueing its item. Instead, by first combining the operations locally, and then trying to insert all elements in one step, this problem can be avoided. This is a technique used by XMalloc, a lock-free memory allocator for graphics processors [48]. On data structures with more disjoint memory access than a queue, the problem is less pronounced, as multiple operations can succeed concurrently if they access different parts of the memory.

An example of a way to take advantage of the SIMD instructions and memory coalescing, is to allow each node in a tree to have more children. Allowing a node in a tree to have more children will have the effect of making the tree shallower and lower the number of nodes that needs to be checked when searching for an item. As a consequence, the time spent in each node will increase, but with coalesced memory access and SIMD instructions, this increase in time spent can be limited by selecting the number of children to suit the SIMD instruction size. The node can then be read in a single memory operation and the correct child can be found using just two SIMD compare instructions.

Another suggestion is to use memory coalescing to implement lazy operations, where larger read and write operations replace a percentage of expensive CAS operations. An array-based queue for example does not need to update its tail pointer using CAS every time an item is inserted. Instead it could be updated every x :th operation, and the correct tail could be found by quickly traversing the array using large memory reads and SIMD instructions, reducing the traversal time to a low static cost. This type of lazy updating was used in the queue by Tsigas and Zhang [91].

The coalescing memory access mechanism also directly influences the synchronization capabilities of the graphics processor. It has for example been shown that it can be used to facilitate wait-free synchronization between threads, without the need of synchronization primitives other than reads and writes [35, 37].

When it comes to software-controlled load balancing, there have been experiments made comparing the built-in hardware scheduler with a software managed work-stealing approach [9]. It was shown that lock-free implementations of data structures worked better than lock-based, and that lock-free work-stealing could outperform the built-in scheduler.

The lack of a stack can be a significant problem for data structures that require recursive helping for lock-freeness. While it is often possible to rewrite recursive

code to work iteratively instead, it requires that recursive depth can be bounded to lower the amount of memory that needs to be allocated.

REFERENCES

1. Y. Afek, G. Korland, M. Natanzon, and N. Shavit. Scalable Producer-Consumer Pools Based on Elimination-Diffraction Trees. In *Euro-Par 2010*, volume 6272 of *Lecture Notes in Computer Science*, pages 151–162. Springer, 2010.
2. O. Agesen, D. Detlefs, C. H. Flood, A. Garthwaite, P. Martin, N. Shavit, and G. L. Steele Jr. DCAS-based concurrent dequeues. In *ACM Symposium on Parallel Algorithms and Architectures*, pages 137–146, 2000.
3. J. H. Anderson and M. Moir. Universal Constructions for Multi-Object Operations. In *Proceedings of the 14th Annual ACM Symposium on the Principles of Distributed Computing*, August 1995.
4. N. S. Arora, R. D. Blumofe, and C. G. Plaxton. Thread Scheduling for Multiprogrammed Multiprocessors. In *ACM Symposium on Parallel Algorithms and Architectures*, pages 119–129, 1998.
5. K. Asanovic, R. Bodik, B. C. Catanzaro, J. J. Gebis, P. Husbands, K. Keutzer, D. A. Patterson, W. L. Plishker, J. Shalf, S. W. Williams, and K. A. Yelick. The Landscape of Parallel Computing Research: A View from Berkeley. *TR No. UCB/EECS-2006-183*, University of California, Berkeley, 2006.
6. H. Attiya and E. Hillel. Built-In Coloring for Highly-Concurrent Doubly-Linked Lists. In *Proceedings of the 20th International Symposium of Distributed Computing*, pages 31–45, 2006.
7. G. Barnes. Wait-Free Algorithms for Heaps. Technical report, Computer Science and Engineering, University of Washington, February 1992.
8. E. Berger, K. McKinley, R. Blumofe, and P. Wilson. Hoard: A Scalable Memory Allocator for Multithreaded Applications. In *9th International Conference on Architectural Support for Programming Languages and Operating Systems*, pages 117–128, November 2000.
9. D. Cederman and P. Tsigas. On Dynamic Load Balancing on Graphics Processors. In *Proceedings of the 23rd ACM SIGGRAPH/EUROGRAPHICS symposium on Graphics hardware*, pages 57–64, 2008.
10. D. Cederman and P. Tsigas. Dynamic Load Balancing Using Work-Stealing. In Wen-Mei Hwu, editor, *GPU Computing Gems Jade Edition*. Morgan Kaufmann, 2011.
11. S. Chaudhry, R. Cypher, M. Ekman, M. Karlsson, A. Landin, S. Yip, H. Zeffner, and M. Tremblay. Rock: A High-Performance Sparc CMT Processor. *Micro, IEEE*, 29(2):6–16, 2009.
12. D. Dechev, P. Pirkelbauer, and B. Stroustrup. Lock-Free Dynamically Resizable Arrays. In *Proceedings of the 10th International Conference on Principles of Distributed Systems*, Lecture Notes in Computer Science, pages 142–156. Springer Verlag, 2006.
13. D. Detlefs, C. H. Flood, A. Garthwaite, P. Martin, N. Shavit, and G. L. Steele Jr. Even Better DCAS-Based Concurrent Deques. In *International Symposium on Distributed Computing*, pages 59–73, 2000.

14. D. L. Detlefs, P. A. Martin, M. Moir, and G. L. Steele, Jr. Lock-free Reference Counting. In *Proceedings of the 20th annual ACM symposium on Principles of Distributed Computing*, pages 190–199. ACM, 2001.
15. D. Dice and A. Garthwaite. Mostly Lock-Free Malloc. In *Proceedings of the 3rd International Symposium on Memory Management*, pages 163–174. ACM Press, 2002.
16. S. Diestelhorst and M. Hohmuth. Hardware Acceleration for Lock-Free Data Structures and Software Transactional Memory. In *Proceedings of the Workshop on Exploiting Parallelism with Transactional Memory and other Hardware Assisted Methods*, pages 1–8, 2008.
17. F. Ellen, P. Fatourou, E. Ruppert, and F. van Breugel. Non-Blocking Binary Search Trees. In *Proceeding of the 29th ACM SIGACT-SIGOPS symposium on Principles of Distributed Computing*, pages 131–140. ACM, 2010.
18. M. Fomitchev and E. Ruppert. Lock-Free Linked Lists and Skip Lists. In *Proceedings of the 23rd annual symposium on Principles of Distributed Computing*, pages 50–59, 2004.
19. K. Fraser and T. Harris. Concurrent Programming Without Locks. *ACM Transactions on Computer Systems*, 25, May 2007.
20. K. A. Fraser. *Practical Lock-Freedom*. PhD thesis, University of Cambridge, 2003.
21. H. Gao, J. F. Groote, and W. H. Hesselink. Lock-free dynamic hash tables with open addressing. *Distributed Computing*, 18(1):21–42, 2005.
22. H. Gao, J. F. Groote, and W. H. Hesselink. Lock-free Parallel and Concurrent Garbage Collection by Mark&Sweep. *Science of Computer Programming*, 64(3):341–374, 2007.
23. J. Giacomoni, T. Moseley, and M. Vachharajani. FastForward for Efficient Pipeline Parallelism: a Cache-Optimized Concurrent Lock-Free Queue. In *Proceedings of the 13th ACM SIGPLAN Symposium on Principles and practice of parallel programming*, pages 43–52. ACM, 2008.
24. A. Gidenstam, M. Papatriantafidou, H. Sundell, and P. Tsigas. Efficient and Reliable Lock-Free Memory Reclamation Based on Reference Counting. In *Proceedings of the 8th International Symposium on Parallel Architectures, Algorithms and Networks*, pages 202–207. IEEE, 2005.
25. A. Gidenstam, M. Papatriantafidou, H. Sundell, and P. Tsigas. Efficient and Reliable Lock-Free Memory Reclamation Based on Reference Counting. *IEEE Transactions on Parallel and Distributed Systems*, 20(8):1173–1187, 2009.
26. A. Gidenstam, M. Papatriantafidou, and P. Tsigas. Allocating Memory in a Lock-Free Manner. In *Proceedings of the 13th Annual European Symposium on Algorithms*, pages 329–242. LNCS vol. 3669, Springer Verlag, 2005.
27. A. Gidenstam, M. Papatriantafidou, and P. Tsigas. NBmalloc: Allocating Memory in a Lock-Free Manner. *Algorithmica*, 58:304–338, 2010.
28. A. Gidenstam, H. Sundell, and P. Tsigas. Cache-Aware Lock-Free Queues for Multiple Producers/Consumers and Weak Memory Consistency. In *Proceedings of the 14th International Conference on Principles Of Distributed Systems*, pages 302–317, 2010.
29. C. Gong and J.M. Wing. A Library of Concurrent Objects and Their Proofs of Correctness. Technical Report CMU-CS-90-151, Computer Science Department, Carnegie Mellon University, 1990.

30. A. Gottlieb, R. Grishman, C. P. Kruskal, K. P. McAuliffe, L. Rudolph, and M. Snir. The NYU Ultracomputer—designing a MIMD, shared-memory parallel machine (Extended Abstract). In *Proceedings of the 9th annual symposium on Computer Architecture, ISCA '82*, pages 27–42. IEEE Computer Society Press, 1982.
31. M. Greenwald. *Non-Blocking Synchronization and System Design*. PhD thesis, Stanford University, 1999.
32. M. Greenwald. Two-handed Emulation: How to Build Non-blocking Implementations of Complex Data-structures using DCAS. In *Proceedings of the 21st annual symposium on Principles of Distributed Computing*, pages 260–269. ACM Press, 2002.
33. M. Greenwald and D. Cheriton. The Synergy Between Non-Blocking Synchronization and Operating System Structure. In *Proceedings of the 2nd Symposium on Operating System Design and Implementation*, pages 123–136, 1996.
34. P. H. Ha and P. Tsigas. Reactive Multi-word Synchronization for Multiprocessors. *Journal of Instruction-Level Parallelism*, 6, April 2004.
35. P. H. Ha, P. Tsigas, and O. J. Anshus. Wait-free Programming for General Purpose Computations on Graphics Processors. In *Proceedings of the IEEE International Parallel and Distributed Processing Symposium*, pages 1–12, 2008.
36. P. H. Ha, P. Tsigas, and O. J. Anshus. NB-FEB: A Universal Scalable Easy-to-Use Synchronization Primitive for Manycore Architectures. In *Proceedings of the International Conference on Principles of Distributed Systems*, pages 189–203, 2009.
37. P. H. Ha, P. Tsigas, and O. J. Anshus. The Synchronization Power of Coalesced Memory Accesses. *IEEE Transactions on Parallel and Distributed Systems*, 21(7):939–953, 2010.
38. T. Harris, K. Fraser, and I. Pratt. A Practical Multi-Word Compare-and-Swap Operation. In *Proceedings of the 16th International Symposium on Distributed Computing*, 2002.
39. T. L. Harris. A Pragmatic Implementation of Non-Blocking Linked Lists. In *Proceedings of the 15th International Symposium of Distributed Computing*, pages 300–314, 2001.
40. D. Hendler, Y. Lev, M. Moir, and N. Shavit. A Dynamic-Sized Nonblocking Work Stealing Deque. *Distributed Computing*, 18(3):189–207, 2006.
41. D. Hendler, N. Shavit, and L. Yerushalmi. A Scalable Lock-Free Stack Algorithm. *Journal of Parallel and Distributed Computing*, 70(1):1–12, 2010.
42. M. Herlihy, V. Luchangco, P. Martin, and M. Moir. Nonblocking Memory Management Support for Dynamic-sized Data Structures. *ACM Transactions on Computer Systems*, 23:146–196, May 2005.
43. M. Herlihy, V. Luchangco, and M. Moir. The Repeat Offender Problem: A Mechanism for Supporting Dynamic-Sized, Lock-Free Data Structure. In *Proceedings of 16th International Symposium on Distributed Computing*, October 2002.
44. M. Herlihy and J. Wing. Linearizability: a Correctness Condition for Concurrent Objects. *ACM Transactions on Programming Languages and Systems*, 12(3):463–492, 1990.
45. M. P. Herlihy and J. E. B. Moss. Lock-Free Garbage Collection for Multiprocessors. *IEEE Transactions on Parallel and Distributed Systems*, 3:304–311, May 1992.
46. T. Herman and V. Damian-Iordache. Space-Optimal Wait-Free Queues. In *Proceedings of the 16th annual ACM symposium on Principles of Distributed Computing*, page 280. ACM Press, 1997.

47. M. Hoffman, O. Shalev, and N. Shavit. The Baskets Queue. In *Proceedings of the 11th International Conference on Principles of Distributed Systems*, pages 401–414, 2007.
48. X. Huang, C. I. Rodrigues, S. Jones, I. Buck, and W. Hwu. XMalloc: A Scalable Lock-free Dynamic Memory Allocator for Many-core Machines. *International Conference on Computer and Information Technology*, 0:1134–1139, 2010.
49. A. Israeli and L. Rappoport. Efficient Wait-Free Implementation of a Concurrent Priority Queue. In *Proceedings of the 7th International Workshop on Distributed Algorithms*, volume 725 of *Lecture Notes in Computer Science*, pages 1–17. Springer Verlag, September 1993.
50. A. Israeli and L. Rappoport. Disjoint-access-parallel Implementations of Strong Shared Memory Primitives. In *Proceedings of the 13th annual ACM symposium on Principles of Distributed Computing*, 1994.
51. P. Jayanti and S. Petrovic. Efficient and Practical Constructions of LL/SC Variables. In *Proceedings of the 22nd Annual Symposium on Principles of Distributed Computing*, pages 285–294. ACM Press, 2003.
52. S. Kelly-Bootle and B. Fowler. *68000, 68010, 68020 Primer*. Howard W. Sams & Co., 1985.
53. G. Kliot, E. Petrank, and B. Steensgaard. A Lock-Free, Concurrent, and Incremental Stack Scanning for Garbage Collectors. In *Proceedings of the 2009 ACM SIGPLAN/SIGOPS international conference on Virtual Execution Environments*, pages 11–20. ACM, 2009.
54. C. P. Kruskal, L. Rudolph, and M. Snir. Efficient Synchronization of Multiprocessors with Shared Memory. *ACM Transactions on Programming Languages and Systems*, 10:579–601, October 1988.
55. L. Lamport. Specifying Concurrent Program Modules. *ACM Transactions on Programming Languages and Systems*, 5(2):190–222, 1983.
56. Doug Lea. The Java Concurrency Package (JSR-166), 2009.
57. M-Herlihy. Wait-Free Synchronization. *ACM Transactions on Programming Languages and Systems*, 11(1):124–149, January 1991.
58. P. Martin, M. Moir, and G. Steele. DCAS-based Concurrent Deques Supporting Bulk Allocation. Technical Report TR-2002-111, Sun Microsystems, 2002.
59. H. Massalin. *Synthesis: An Efficient Implementation of Fundamental Operating System Services*. PhD thesis, Columbia University, 1992.
60. H. Massalin and C. Pu. A Lock-Free Multiprocessor OS Kernel. Technical Report CUCS-005-91, Computer Science Department, Columbia University, June 1991.
61. M. M. Michael. High Performance Dynamic Lock-Free Hash Tables and List-Based Sets. In *Proceedings of the 14th ACM Symposium on Parallel Algorithms and Architectures*, pages 73–82, 2002.
62. M. M. Michael. Safe Memory Reclamation for Dynamic Lock-Free Objects Using Atomic Reads and Writes. In *Proceedings of the 21st ACM Symposium on Principles of Distributed Computing*, pages 21–30, 2002.
63. M. M. Michael. CAS-based lock-free algorithm for shared deques. In *Proceedings of the 9th International Euro-Par Conference*, LNCS. Springer Verlag, August 2003.
64. M. M. Michael. Hazard Pointers: Safe Memory Reclamation for Lock-Free Objects. *IEEE Transactions on Parallel and Distributed Systems*, 15(8), August 2004.

65. M. M. Michael. Practical Lock-Free and Wait-Free LL/SC/VL Implementations Using 64-Bit CAS. In *Proceedings of the 18th International Conference on Distributed Computing*, pages 144–158, 2004.
66. M. M. Michael. Scalable Lock-Free Dynamic Memory Allocation. In *Proceedings of the 2004 ACM SIGPLAN Conference on Programming Language Design and Implementation*, pages 35–46, June 2004.
67. M. M. Michael and M. L. Scott. Correction of a Memory Management Method for Lock-Free Data Structures. Technical report, Computer Science Department, University of Rochester, 1995.
68. M. M. Michael and M. L. Scott. Simple, fast, and practical non-blocking and blocking concurrent queue algorithms. In *Proceedings of the 15th annual ACM Symposium on Principles of Distributed Computing*, pages 267–275, 1996.
69. Microsoft. Parallel Computing Developer Center, 2009.
70. M. Moir. Practical Implementations of Non-Blocking Synchronization Primitives. In *Proceedings of the 15th Annual ACM Symposium on the Principles of Distributed Computing*, August 1997.
71. M. Moir. Transparent Support for Wait-Free Transactions. In *Proceedings of the 11th International Workshop on Distributed Algorithms*, volume 1320, pages 305–319, September 1997.
72. M. Moir, D. Nussbaum, O. Shalev, and N. Shavit. Using Elimination to Implement Scalable and Lock-Free FIFO Queues. In *Proceedings of the 17th annual ACM Symposium on Parallelism in Algorithms and Architectures*, pages 253–262. ACM, 2005.
73. G. F. Pfister, W. C. Brantley, D. A. George, S. L. Harvey, W. J. Kleinfelder, K. P. McAuliffe, E. S. Melton, V. A. Norton, and J. Weiss. The IBM Research Parallel Processor Prototype (RP3): Introduction and Architecture. In *ICPP*, pages 764–771, 1985.
74. S. Prakash, Y. H. Lee, and T. Johnson. A Nonblocking Algorithm for Shared Queues Using Compare-and-Swap. *IEEE Transactions on Computers*, 43(5):548–559, 1994.
75. J. Reinders. *Intel Threading Building Blocks: Outfitting C++ for Multi-core Processor Parallelism*. O’Reilly Media, 2007.
76. S. Schneider, C. D. Antonopoulos, and D. S. Nikolopoulos. Scalable Locality-Conscious Multithreaded Memory Allocation. In *Proceedings of the 5th International Symposium on Memory Management*, pages 84–94. ACM, 2006.
77. O. Shalev and N. Shavit. Split-ordered Lists: Lock-Free Extensible Hash Tables. In *Proceedings of the 22nd Annual Symposium on Principles of Distributed Computing*, pages 102–111. ACM Press, 2003.
78. C. Shann, T. Huang, and C. Chen. A Practical Nonblocking Queue Algorithm using Compare-and-Swap. In *Proceedings of the Seventh International Conference on Parallel and Distributed Systems*, pages 470–475, 2000.
79. N. Shavit and D. Touitou. Software Transactional Memory. In *Proceedings of the fourteenth annual ACM symposium on Principles of distributed computing*, pages 204–213. ACM Press, 1995.
80. M. Spiegel and P. F. Reynolds Jr. Lock-Free Multiway Search Trees. In *Proceedings of the 39th International Conference on Parallel Processing*, pages 604–613, 2010.

81. H. Sundell. *Efficient and Practical Non-Blocking Data Structures*. PhD thesis, Chalmers University of Technology, 2004.
82. H. Sundell. Wait-Free Multi-Word Compare-And-Swap using Greedy Helping and Grabbing. In *Proceedings of the International Conference on Parallel and Distributed Processing Techniques and Applications*, pages 494–500, 2009.
83. H. Sundell, A. Gidenstam, M. Papatriantafilou, and P. Tsigas. A Lock-Free Algorithm for Concurrent Bags. In *Proceedings of the 23rd ACM Symposium on Parallelism in Algorithms and Architectures*. ACM, 2011.
84. H. Sundell and P. Tsigas. NOBLE: A Non-Blocking Inter-Process Communication Library. In *Proceedings of the 6th Workshop on Languages, Compilers and Run-time Systems for Scalable Computers*, 2002.
85. H. Sundell and P. Tsigas. Fast and Lock-Free Concurrent Priority Queues for Multi-Thread Systems. In *Proceedings of the 17th International Parallel and Distributed Processing Symposium*, page 11. IEEE press, 2003.
86. H. Sundell and P. Tsigas. Scalable and Lock-Free Concurrent Dictionaries. In *Proceedings of the 19th ACM Symposium on Applied Computing*, pages 1438–1445. ACM press, 2004.
87. H. Sundell and P. Tsigas. Fast and Lock-Free Concurrent Priority Queues for Multi-Thread Systems. *Journal of Parallel and Distributed Computing*, 65(5):609–627, May 2005.
88. H. Sundell and P. Tsigas. Lock-Free Deques and Doubly Linked Lists. *Journal of Parallel and Distributed Computing*, 68(7):1008–1020, July 2008.
89. J. Tsay and H.-C. Li. Lock-Free Concurrent Tree Structures for Multiprocessor Systems. In *Proceedings of the International Conference on Parallel and Distributed Systems*, pages 544–549, 1994.
90. P. Tsigas and Y. Zhang. A Simple, Fast and Scalable Non-Blocking Concurrent FIFO queue for Shared Memory Multiprocessor Systems. In *Proceedings of the 13th annual ACM Symposium on Parallel Algorithms and Architectures*, pages 134–143, 2001.
91. P. Tsigas and Y. Zhang. Evaluating the Performance of Non-Blocking Synchronization on Shared-Memory Multiprocessors. In *Proceedings of the International Conference on Measurement and Modeling of Computer Systems*, pages 320–321. ACM Press, 2001.
92. P. Tsigas and Y. Zhang. Integrating Non-blocking Synchronisation in Parallel Applications: Performance Advantages and Methodologies. In *Proceedings of the 3rd ACM Workshop on Software and Performance*, pages 55–67. ACM Press, 2002.
93. J. Turek and D. Shasha. The many faces of consensus in distributed systems. *IEEE Computer*, 25(2):8–17, 1992.
94. J. D. Valois. Implementing Lock-Free Queues. In *Proceedings of the 7th International Conference on Parallel and Distributed Computing Systems*, pages 64–69, 1994.
95. J. D. Valois. *Lock-Free Data Structures*. PhD thesis, Rensselaer Polytechnic Institute, Troy, New York, 1995.

