



Fully Adaptive Snapshots

Hagit Attiya and Idan Zach
Department of Computer Science
The Technion



The Collect Problem

- Collect up-to-date values of active processes
- A simple solution uses an array with N entries
⇒ $O(N)$ step complexity

But ...

- N is large
- Often only few of the processes take steps

2


Attiya and Zach
Efficient Adaptive Snapshots

DISC, October 2004



Adaptive Algorithms

- Step complexity depends only on the number of participating processes
 - **Point contention** during an operation: Max number of processes taking steps **together**



3

Attiya and Zach
Efficient Adaptive Snapshots

DISC, October 2004



Adaptive Collect Algorithms

- [Attiya, Fouren & Gafni, 1997] $O(K)$ (total contention)
- [AfeK, Stupp & Touitou, 1999] $O(k^4)$
- [Attiya & Fouren, 2003] $O(k^2)$

But local step complexity depends on the total contention!

- [AfeK, Stupp & Touitou, 2000] $O(k^3)$ local + shared
 - This work $O(k) / O(k^2)$ local + shared
- For a generalized problem

4



Attiya and Zach
Efficient Adaptive Snapshots

DISC, October 2004



Collect & f

- Compute a function f on the stored values
 - associative, commutative and idempotent
 - $put&f(val)$, $collect&f()$
- A weaker variant: $gather&f$
 - Earlier gather may be more up-to-date than a later gather




Implementing Gather & f / Collect & f

- Naïve method: Use a gather / collect object
 - Local step complexity is not adaptive since f is applied to **all** stored values
- Better method: Incrementally calculate the result as values are stored, and save it for later use
 - Extend idea from [Afek, Stupp & Touitou, 2000]
 - Fit into [Attiya and Fouren, 2003]



Two Examples

- **Pile** maintains the maximum value stored
 - [Afek et al. 2000] $O(k^4)$
 - using our **gather & max** $O(k^2)$
- **Active set** returns the current set of active processes
 - [Afek et al. 2000] $O(k^3)$
 - using our **gather & union** $O(k^2)$



Repetitive Collect & f

- Many algorithms invoke many put / collect operations within one high-level operation
 - ⇒ Amortize costs over all invocations



- $put&f$ a single step
- $collect&f$ $O(k)$ steps
- $start$ and end $O(k^2)$ steps

Implementing Repetitive Collect&f


```

start {
  sign-in (active set) }

put&f {
  Ri = value }

collect&f {
  get active set
  read regs of active proc.
  read from Pile
  merge the results }

end {
  put last result in pile
  sign-out (active set) }
  
```



9

Attiya and Zach
Efficient Adaptive Snapshots

DISC, October 2004

Implications

- Atomic snapshots
 - [Afek, Stupp & Touitou, 2000] $O(k^k)$ local+shared
 - [Attiya & Fouren, 2003] $O(k^3)$ shared only
 - This work $O(k^2)$ local+shared
- Immediate snapshots
 - [Attiya, Fouren & Gafni, 1997] $O(K^3)$ shared (total)
 - [Afek, Stupp & Touitou, 2000] $O(k^5)$ local+shared
 - This work $O(k^3)$ local+shared
- Long-lived $(2k-1)$ -renaming
 - [Attiya, Fouren & Gafni, 2003] $O(k^3)$ local+shared
 - Using our work

10

Attiya and Zach
Efficient Adaptive Snapshots

DISC, October 2004