Concurrent

Compact-fit

Concurrency & Scalability
versus
Fragmentation & Compaction

Questions

- Does allocation/deallocation throughput scale with multiple processors?

Questions

- Does allocation/deallocation throughput scale with multiple processors?
- Which aspects influence scalability?

Questions

- Does allocation/deallocation throughput scale with multiple processors?
- Which aspects influence scalability?
- Does compaction of large objects harm system latency?

Questions

- Does allocation/deallocation throughput scale with multiple processors?
- Which aspects influence scalability?
- Does compaction of large objects harm system latency?
- Does concurrency and incrementality affect memory consumption?

Partial Compaction

- Per-size-class partial compaction bound κ bounds size-class fragmentation:
  - κ = 1: fully compacting
  - 1 < κ < ∞: partially compacting
  - κ = ∞: non-compacting

Partial Compaction

- Per-size-class partial compaction bound κ bounds size-class fragmentation:
  - κ = 1: fully compacting
  - 1 < κ < ∞: partially compacting
  - κ = ∞: non-compacting
- Non-compacting CF can be optimized by not using abstract addresses

Fragmentation through Partitioning

- Fragmentation through partitioning is fixed at compile time and is not controlled by partial compaction:
  - Page-block-internal fragmentation
  - Page-internal fragmentation

Fragmentation through Partitioning

- Fragmentation through partitioning is fixed at compile time and is not controlled by partial compaction:
  - Page-block-internal fragmentation
  - Page-internal fragmentation
- May dominate overall fragmentation

Size Class 1

Size Class 2

Size Class 3

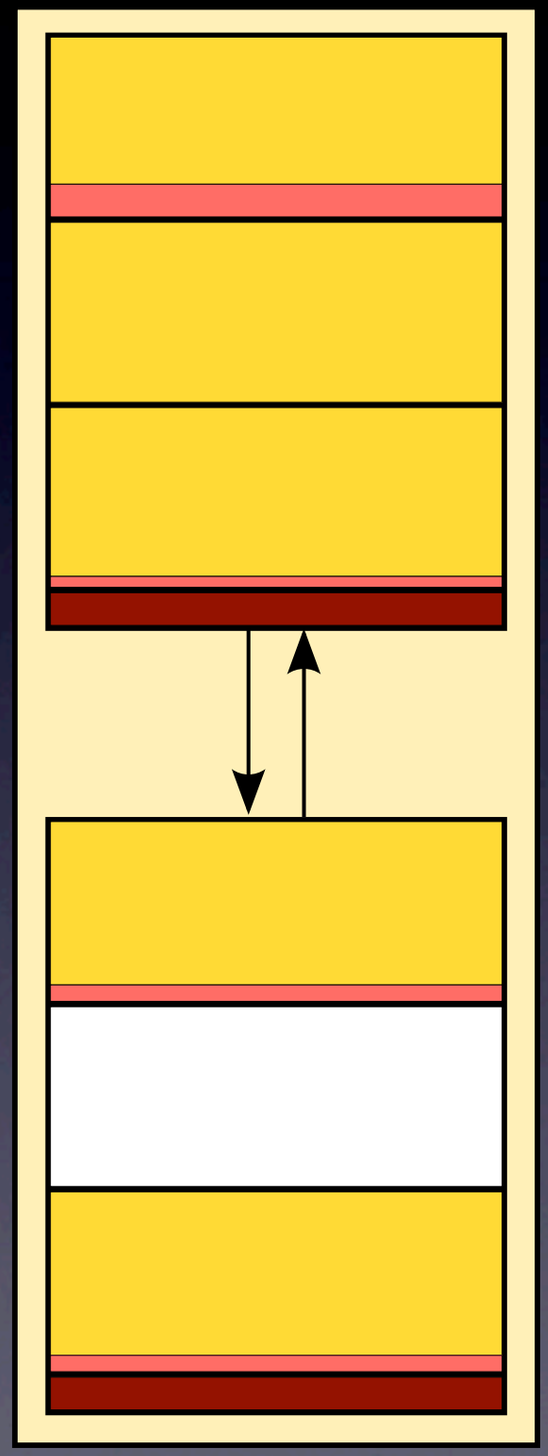

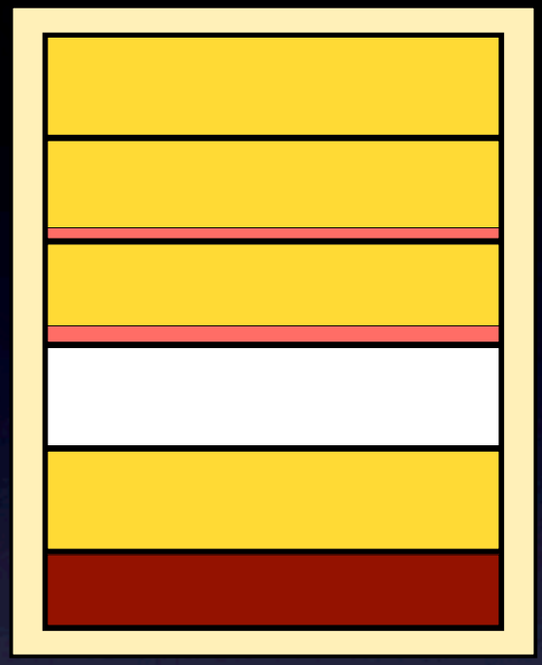

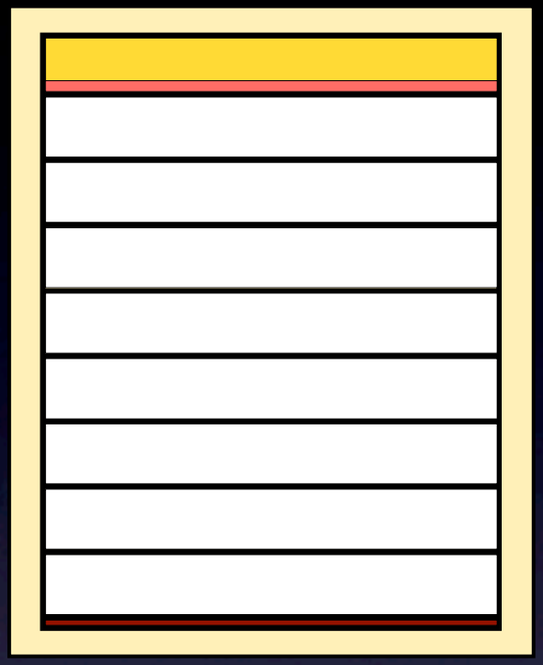

- free range
- used space
- page-block-internal fragmentation
- page-internal fragmentation

Size Class 1

Size Class 2

Size Class 3

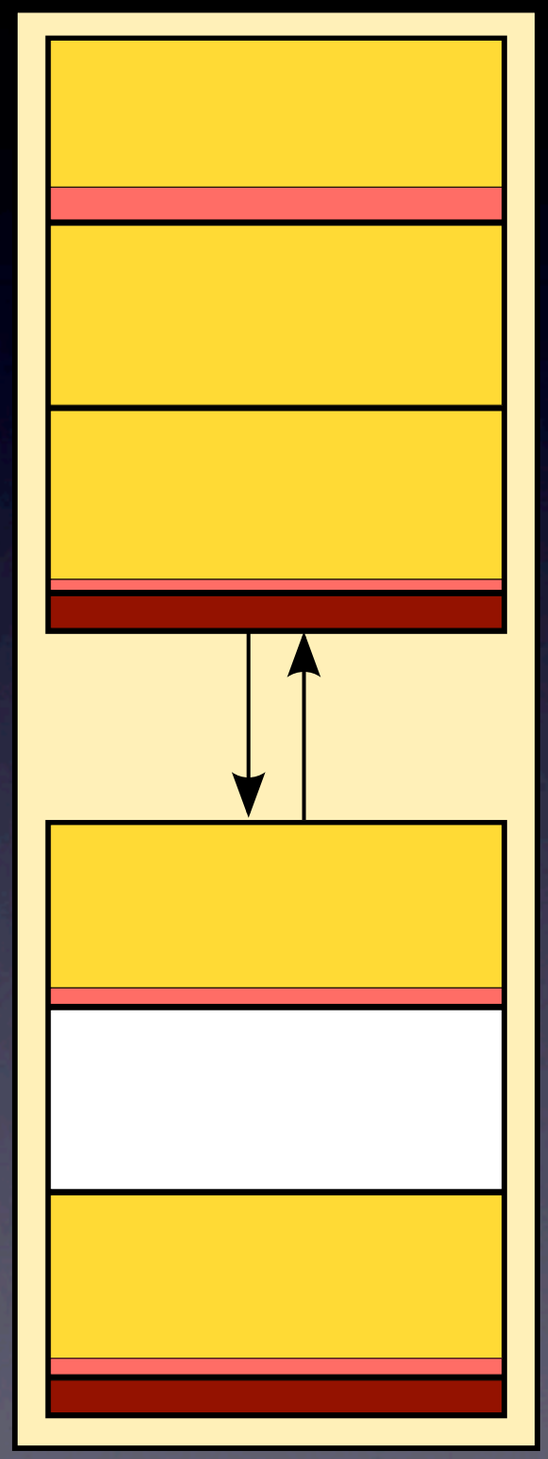

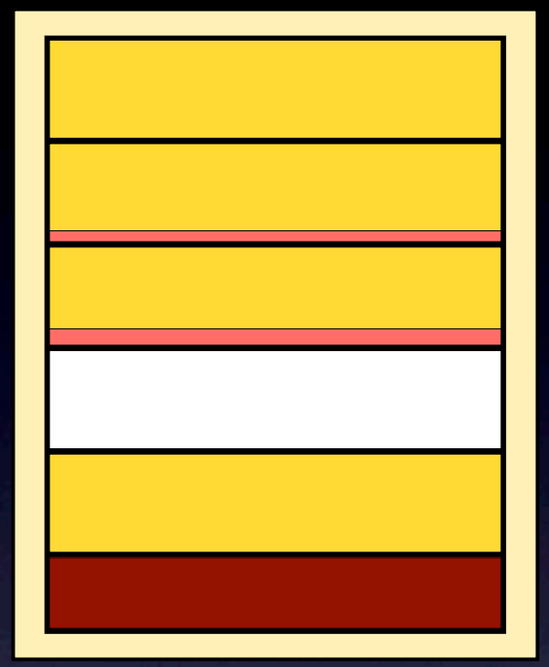

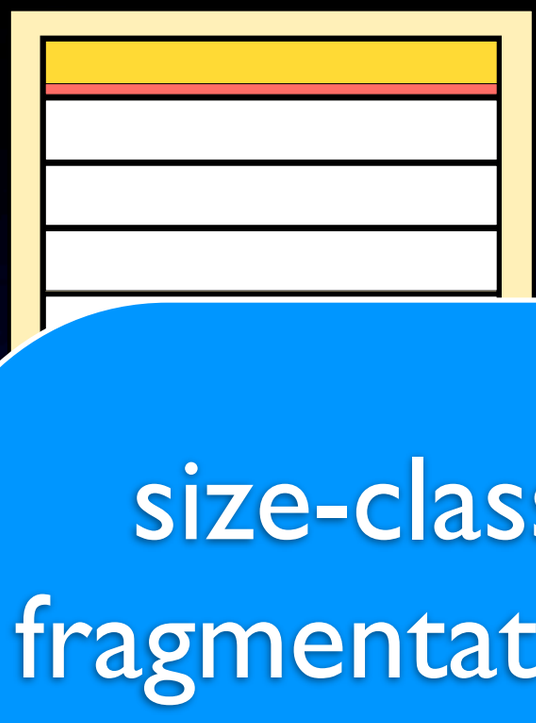

size-class fragmentation

- free range
- used space
- page-block-internal fragmentation
- page-internal fragmentation

Incremental Compaction

- Global compaction increment ι bounds size of memory involved in any atomic compaction operation:
  - $1 < \iota < \infty$: incremental compaction of objects larger than ι
  - $\iota = \infty$: non-incremental compaction

Incremental Compaction

- Global compaction increment ι bounds size of memory involved in any atomic compaction operation:
  - $1 < \iota < \infty$: incremental compaction of objects larger than ι
  - $\iota = \infty$: non-incremental compaction
- Incremental compaction creates transient size-class fragmentation

CF Configurations

- 1-CF(κ, ι)
  - one CF instance for multiple threads
  - partial compaction bound κ
  - compaction increment ι

CF Configurations

- 1-CF(κ, ι)
  - one CF instance for multiple threads
    - partial compaction bound κ
    - compaction increment ι
- n-CF(κ, ι)
  - n CF instances for n threads
  - allows to control degree of sharing

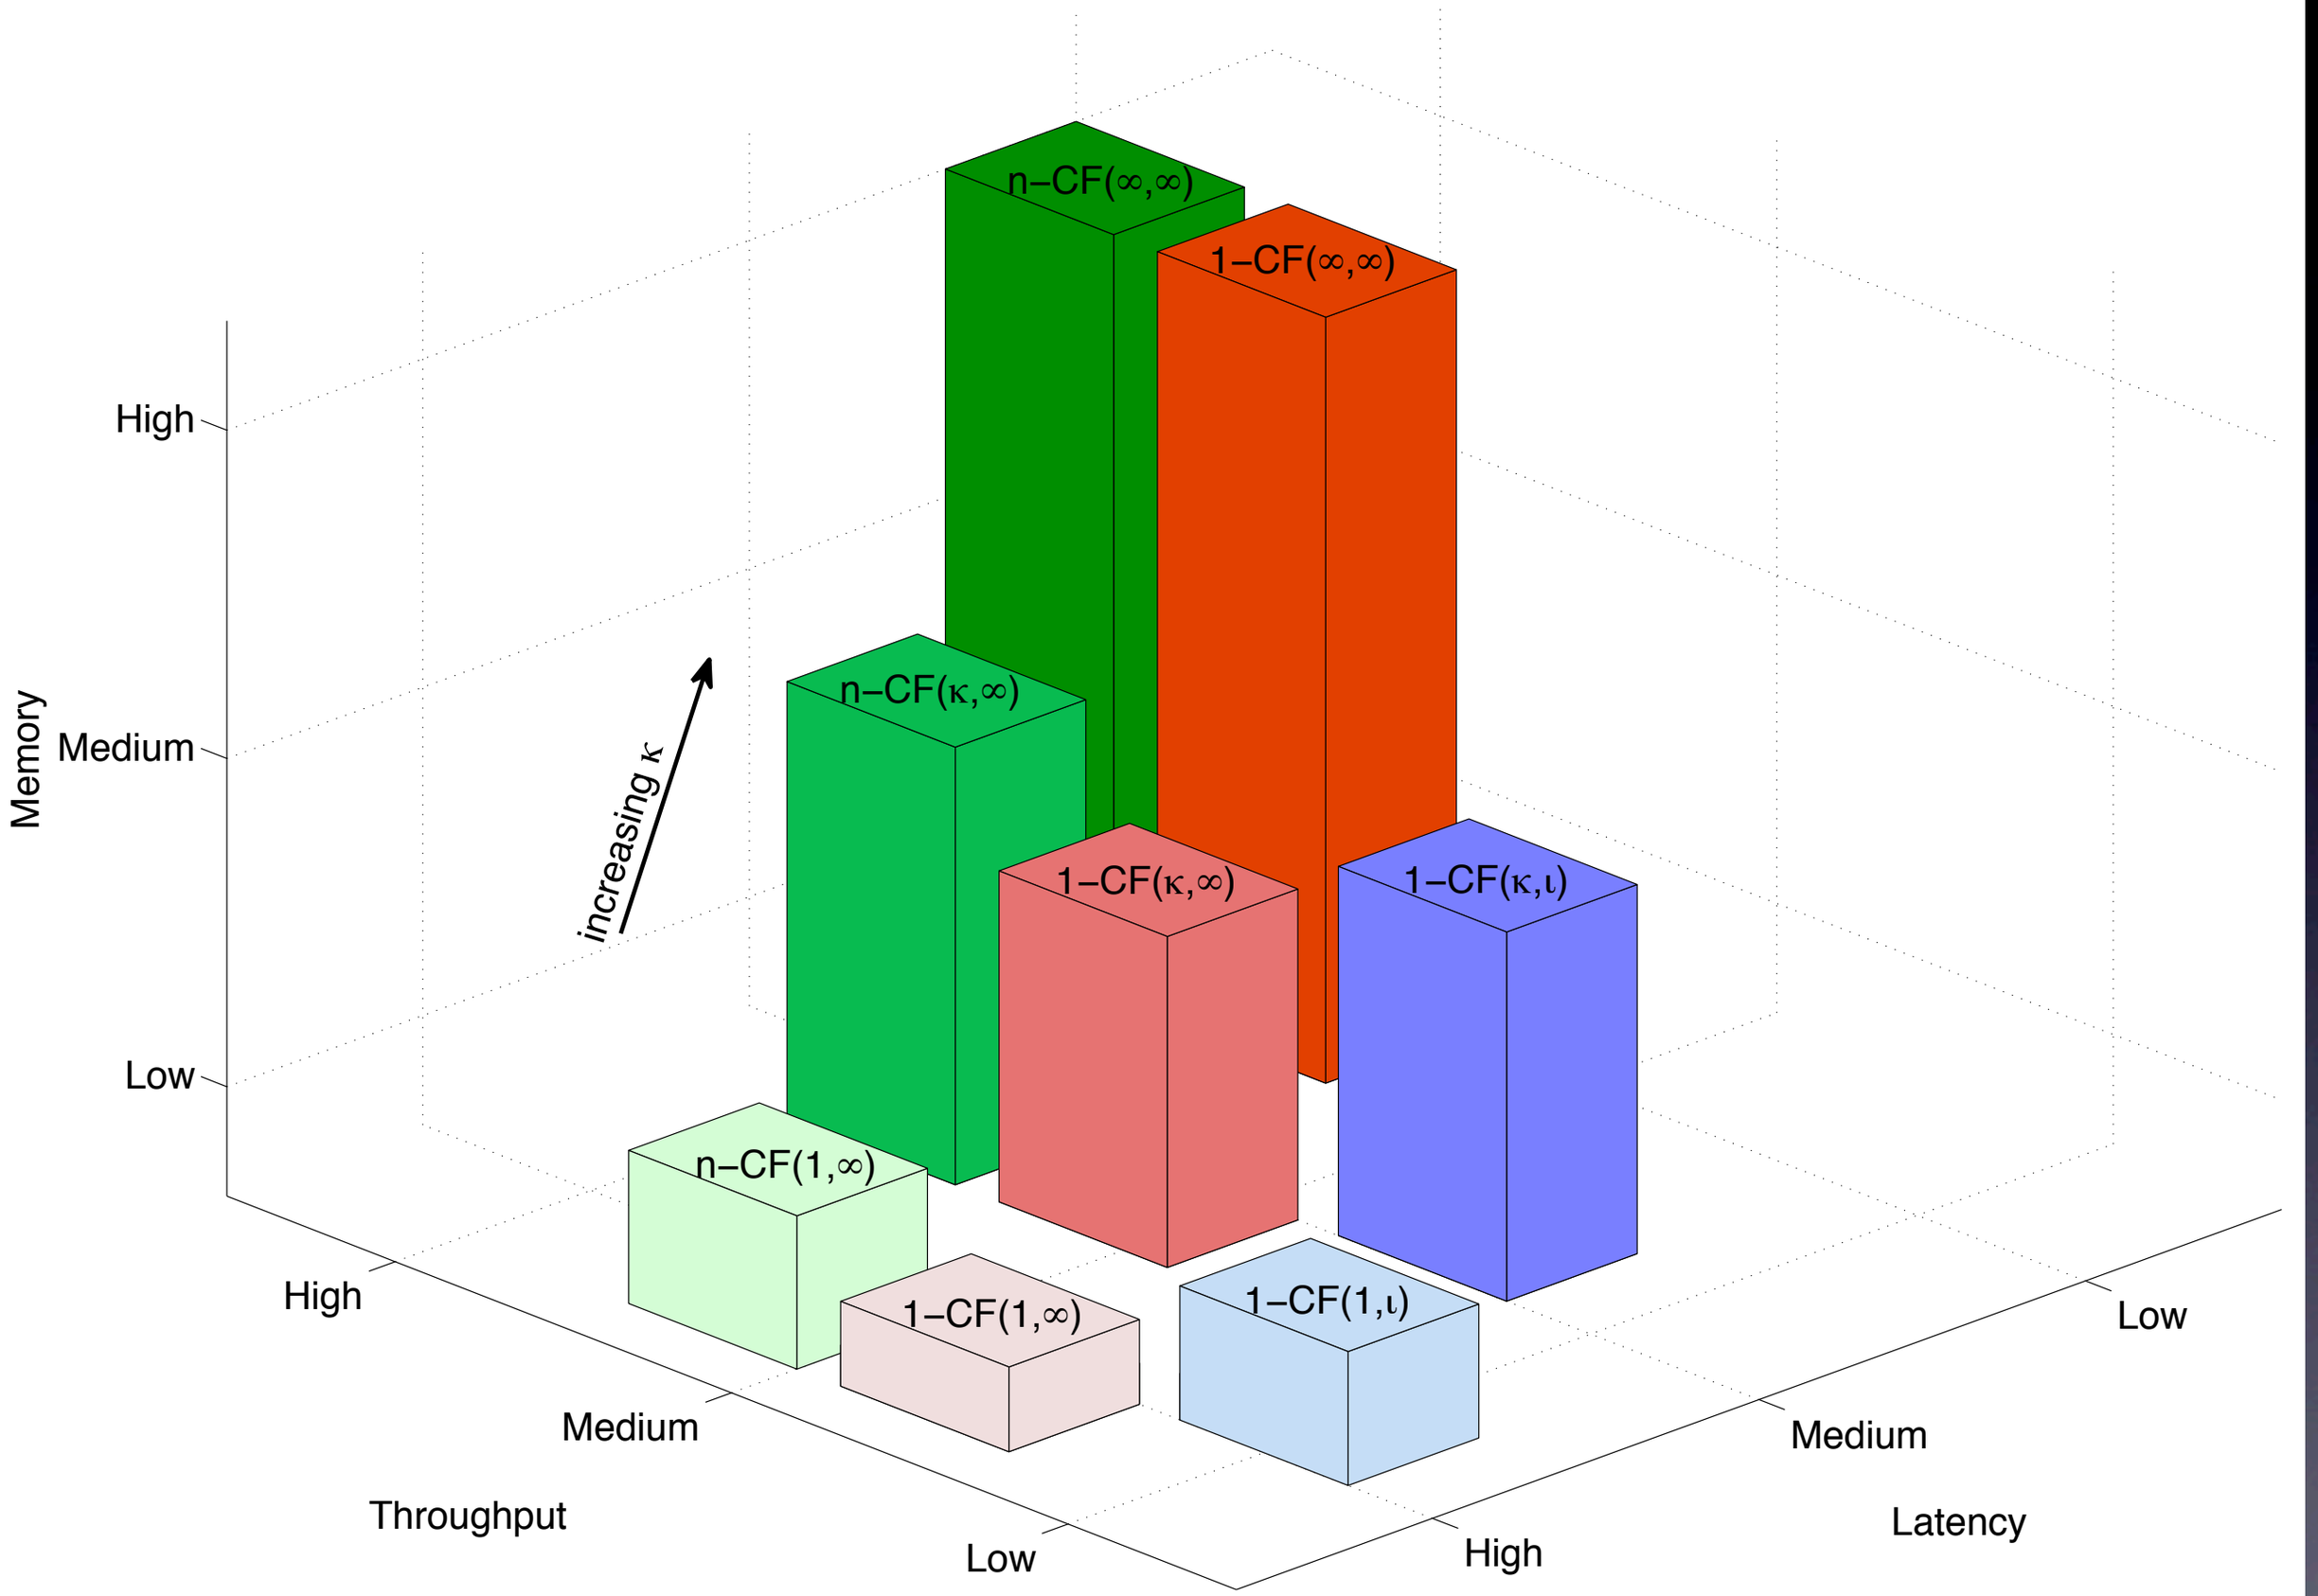

To make CF
concurrent and incremental
we model the algorithm
as a
finite state machine
whose transitions
must be atomic!

Size-Class Automaton for π = 1

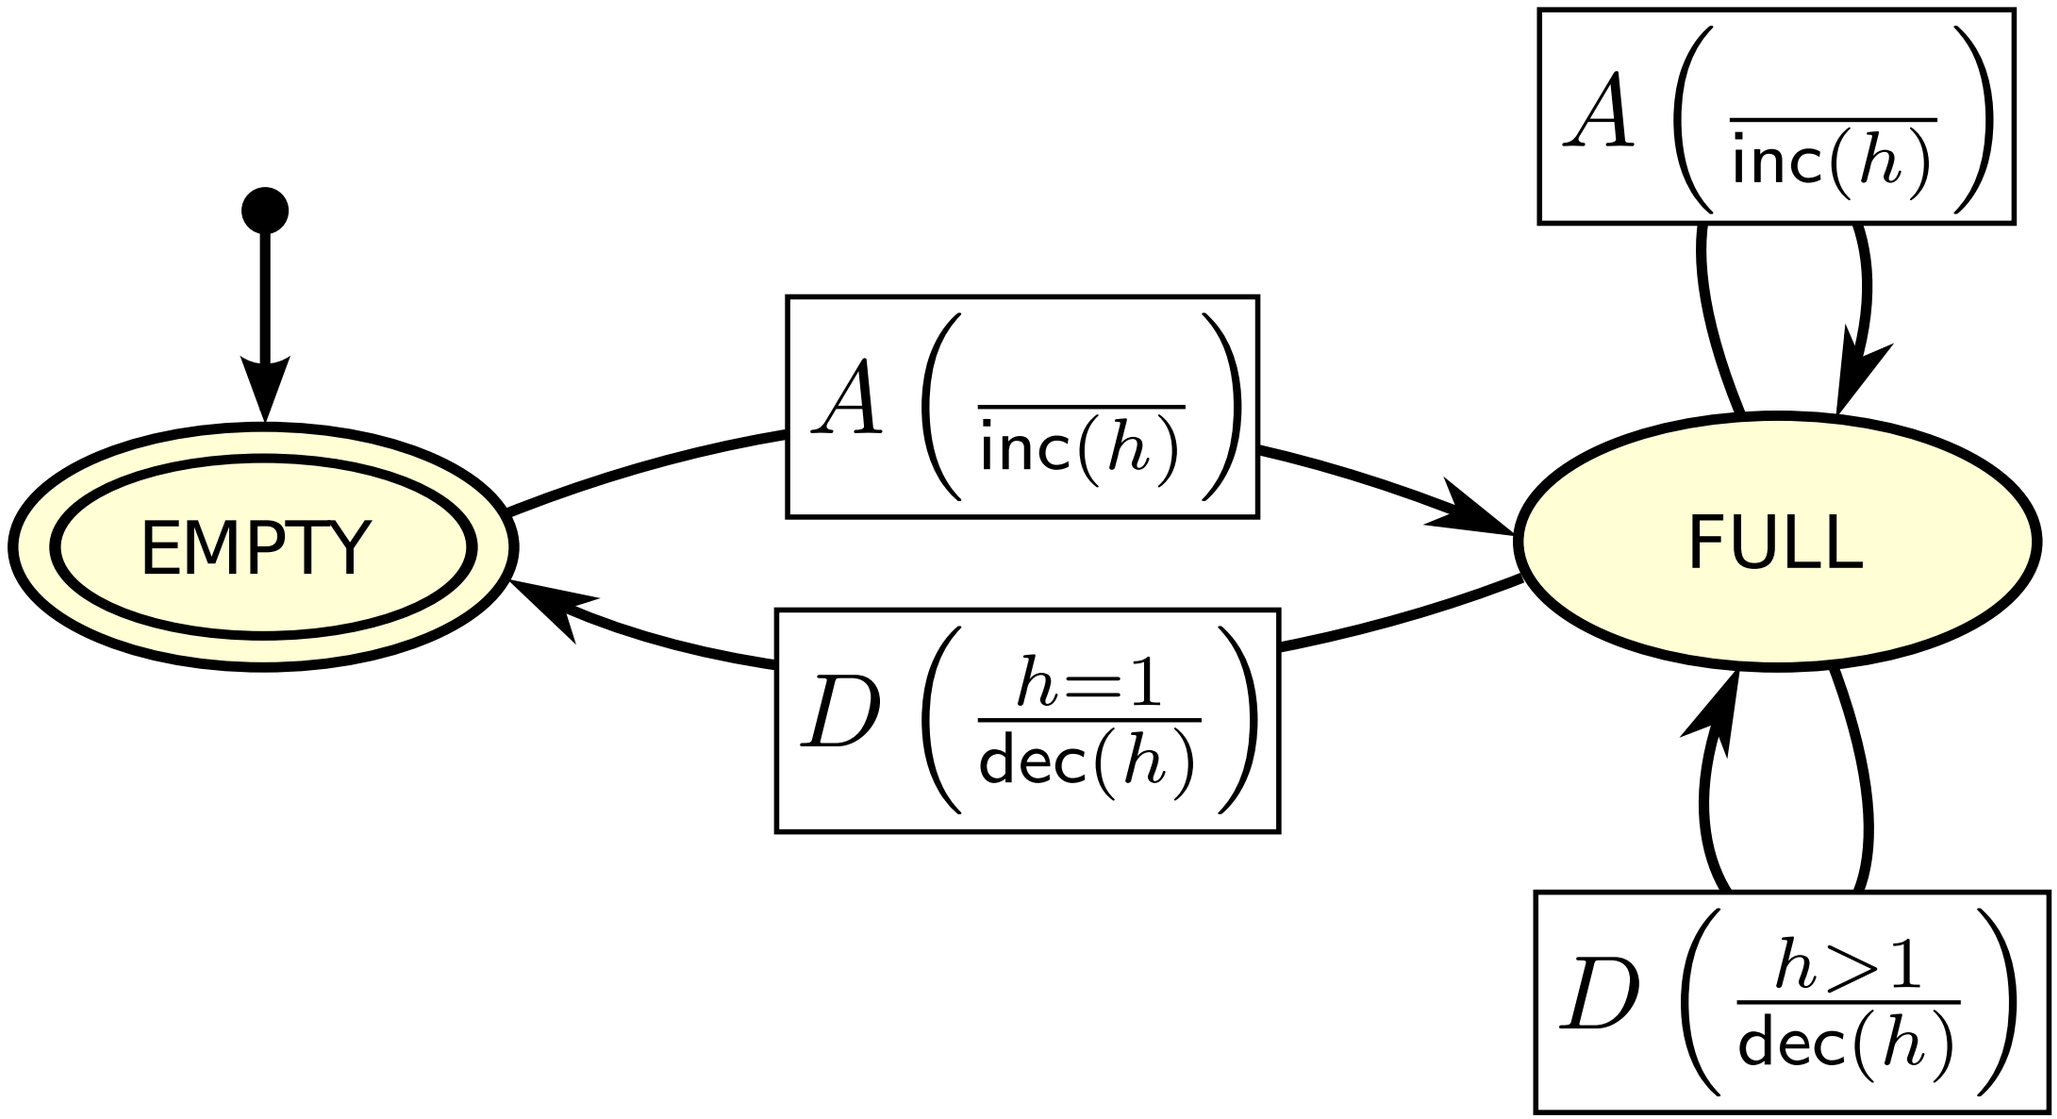

Size-Class Automaton for π_malloc

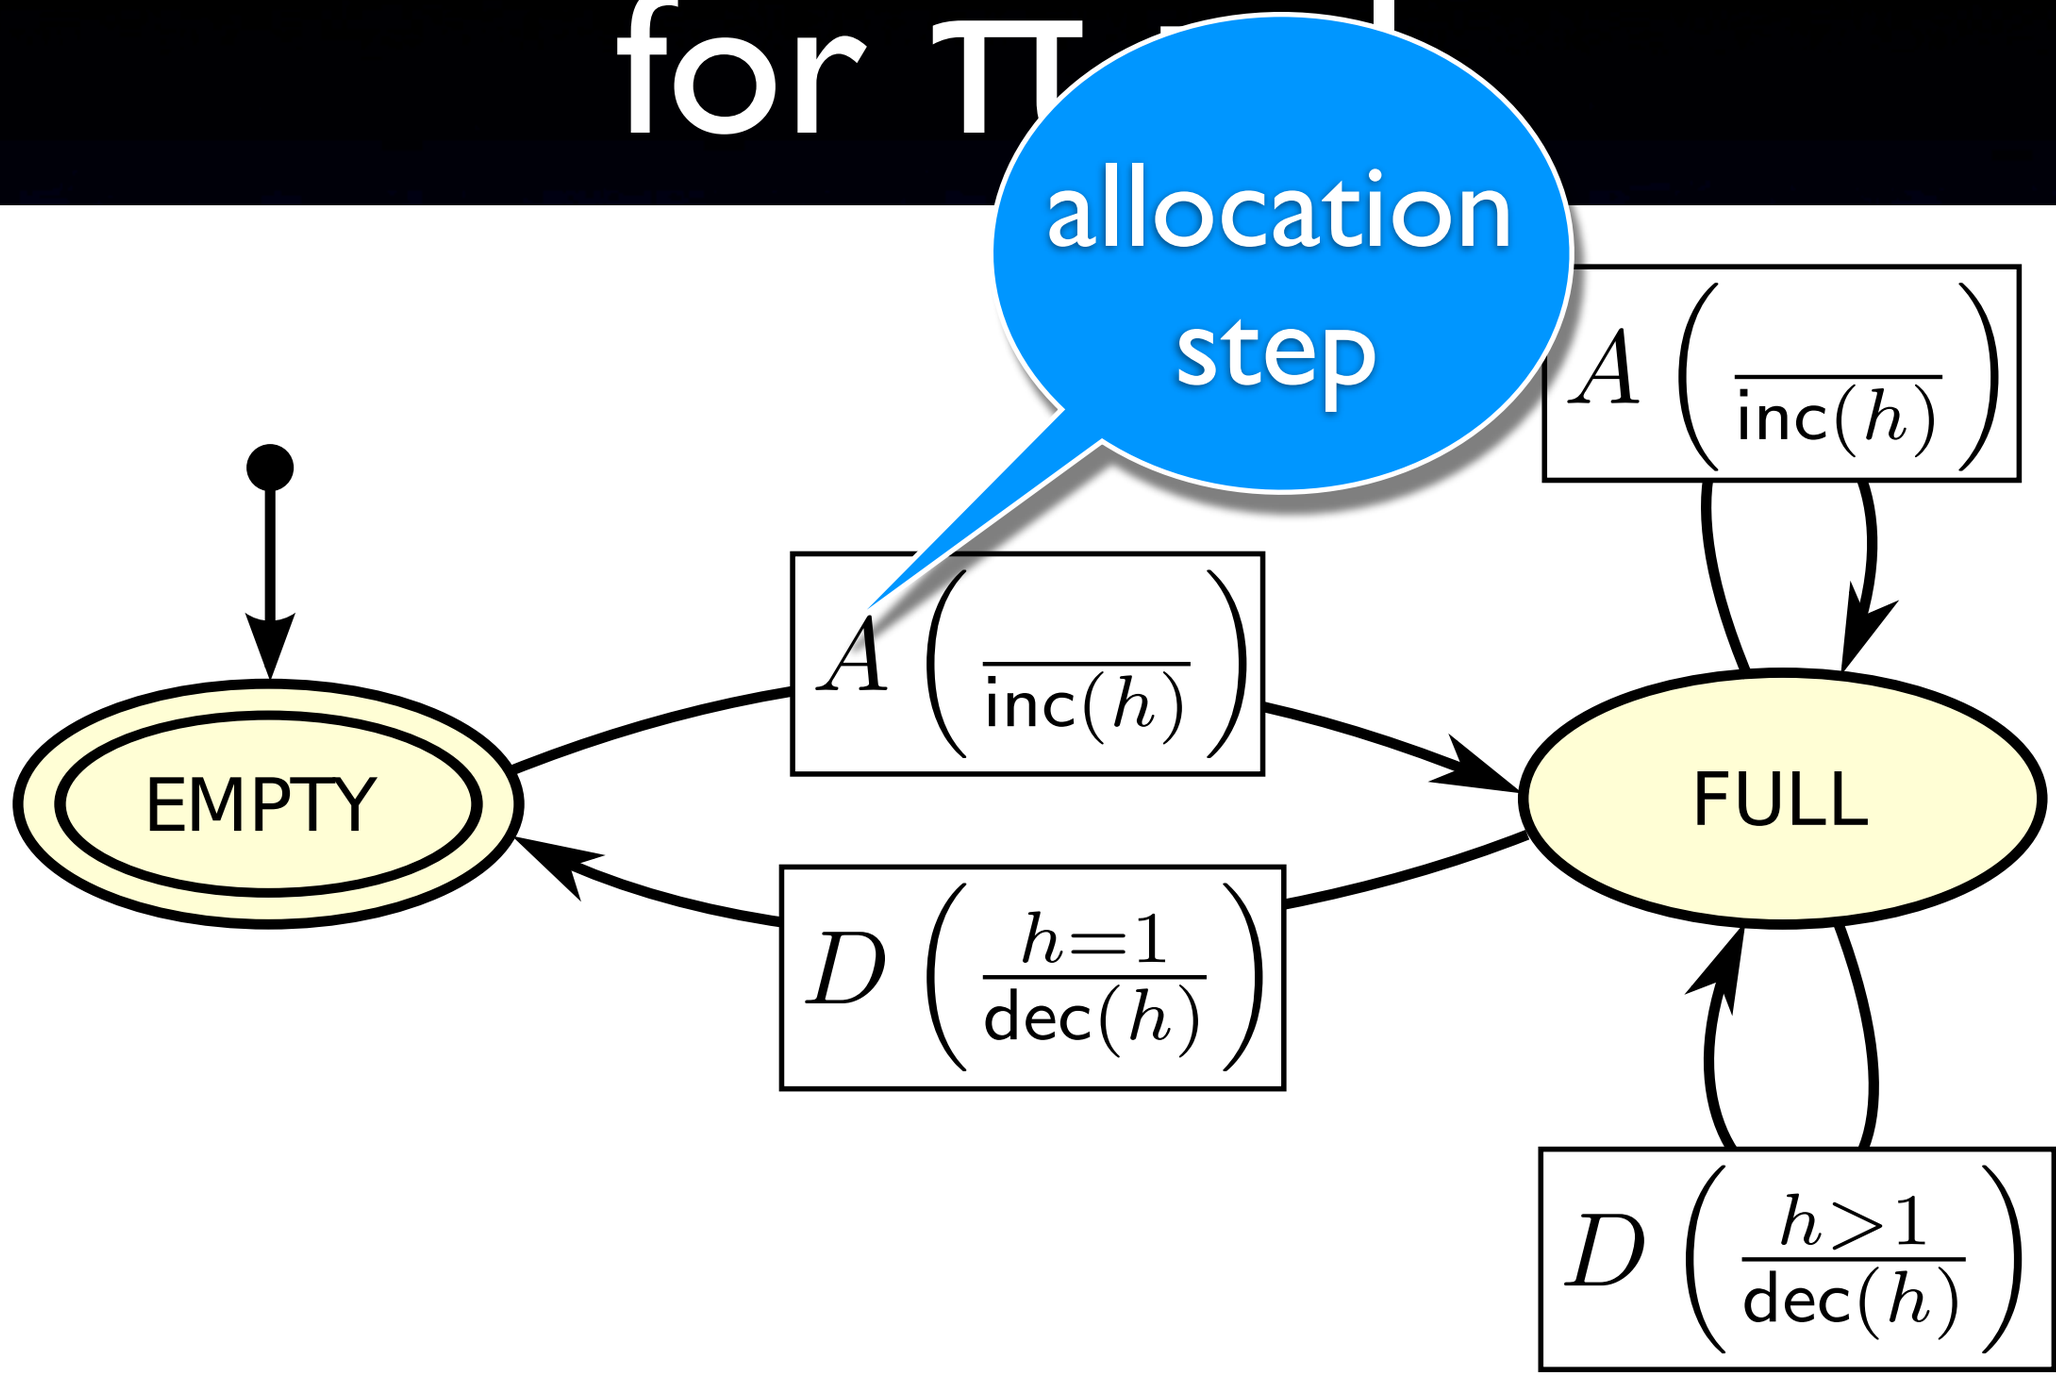

Size-Class Automaton for π = 1

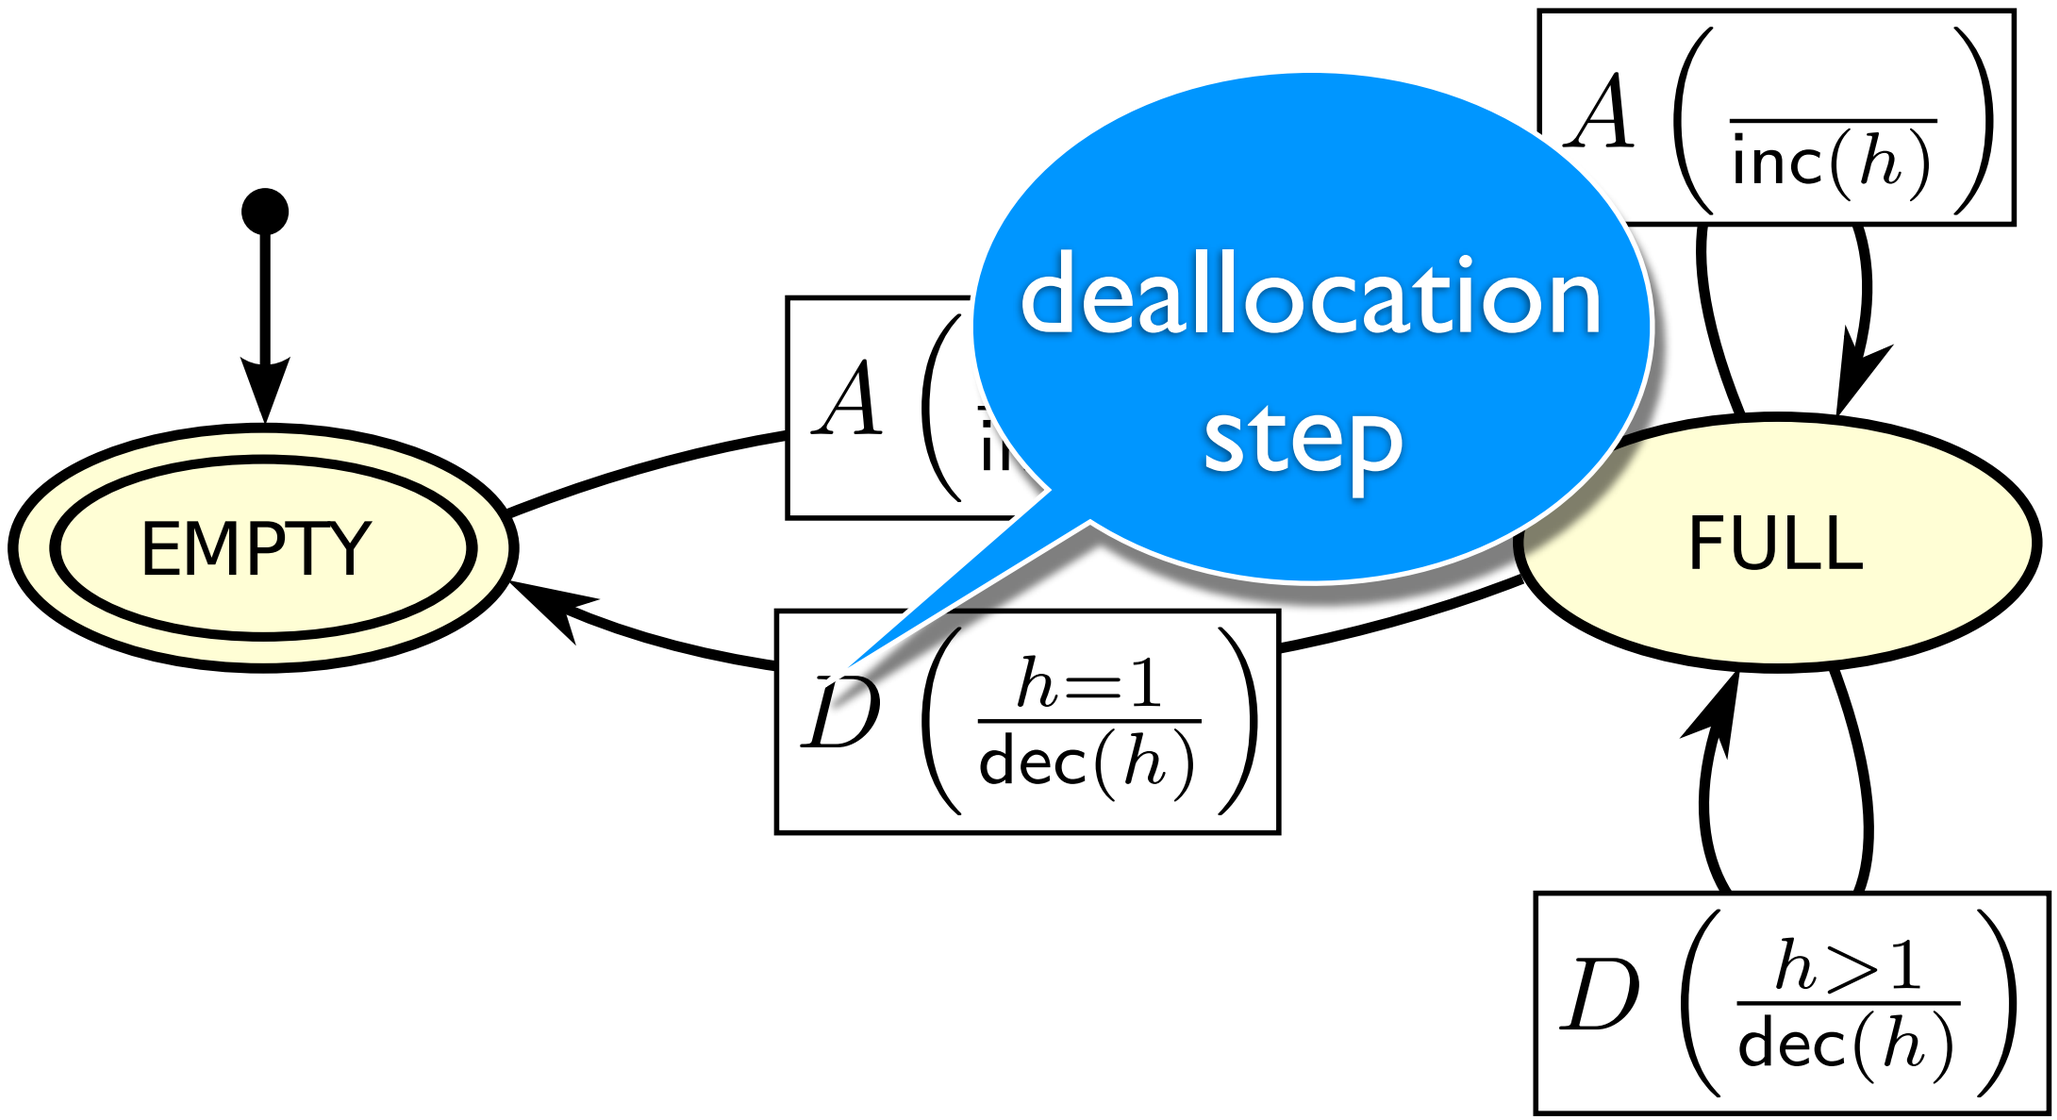

Size-Class Automaton

π = I

h is the total # of allocated page-blocks in the size-class

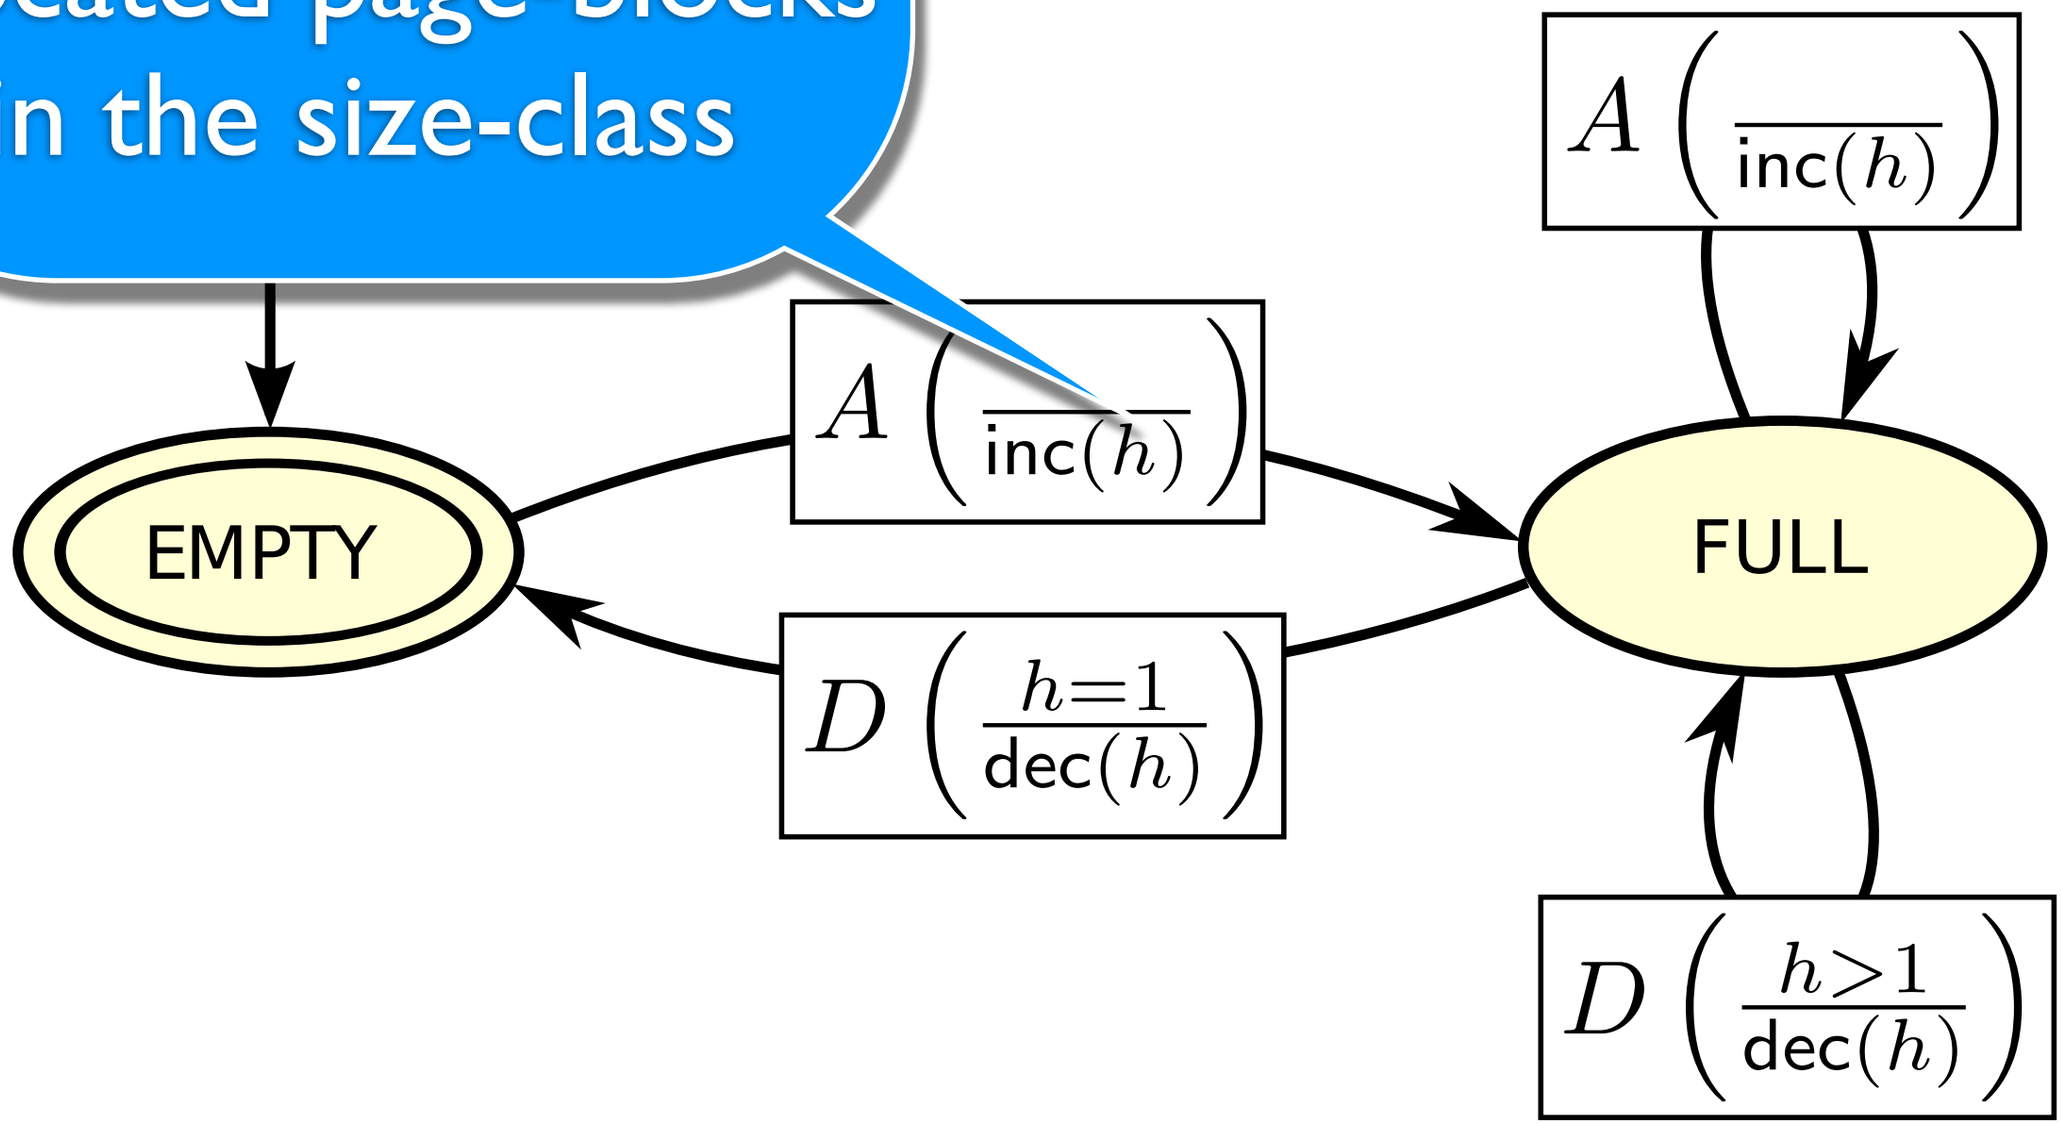

Size-Class Automaton for π > 1

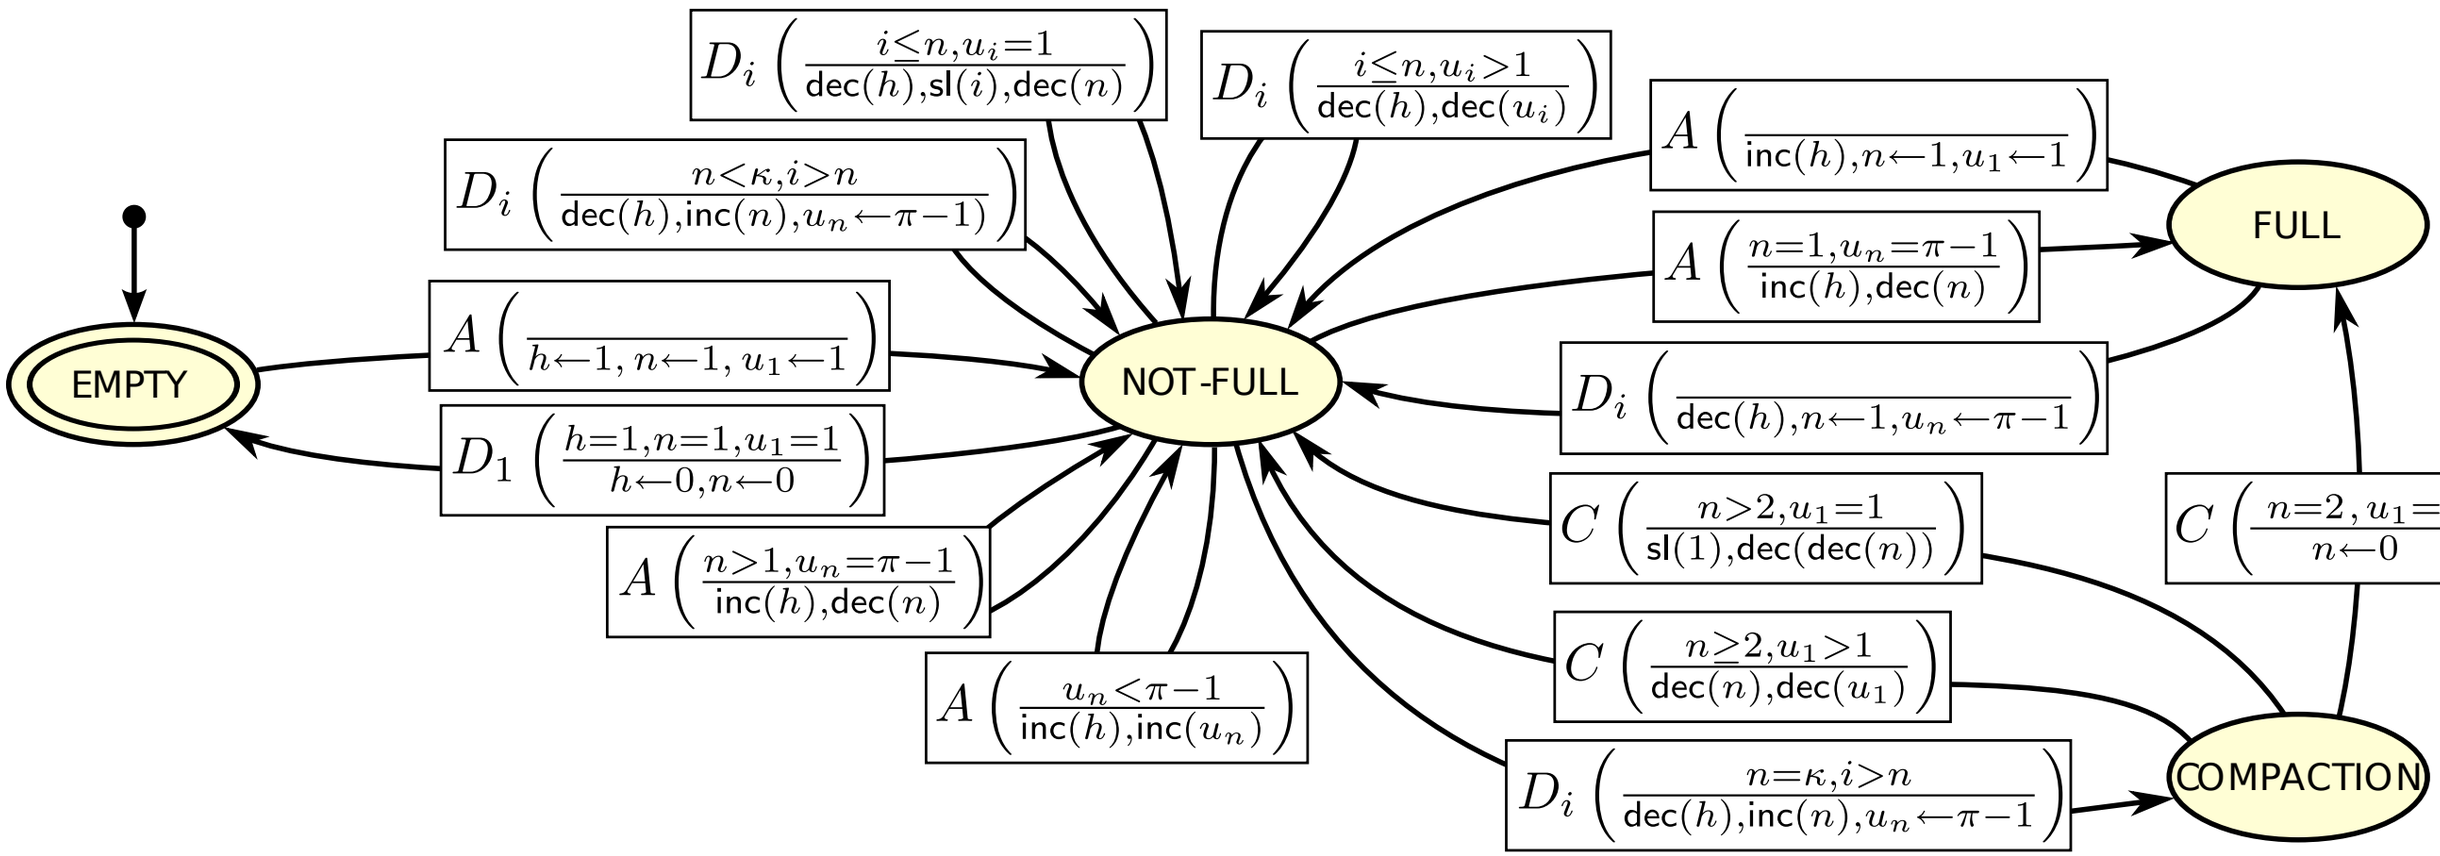

h is the total # of allocated page-blocks in the size-class

n is the # of not-full pages

u_i is the # of used page-blocks in a not-full page *i*