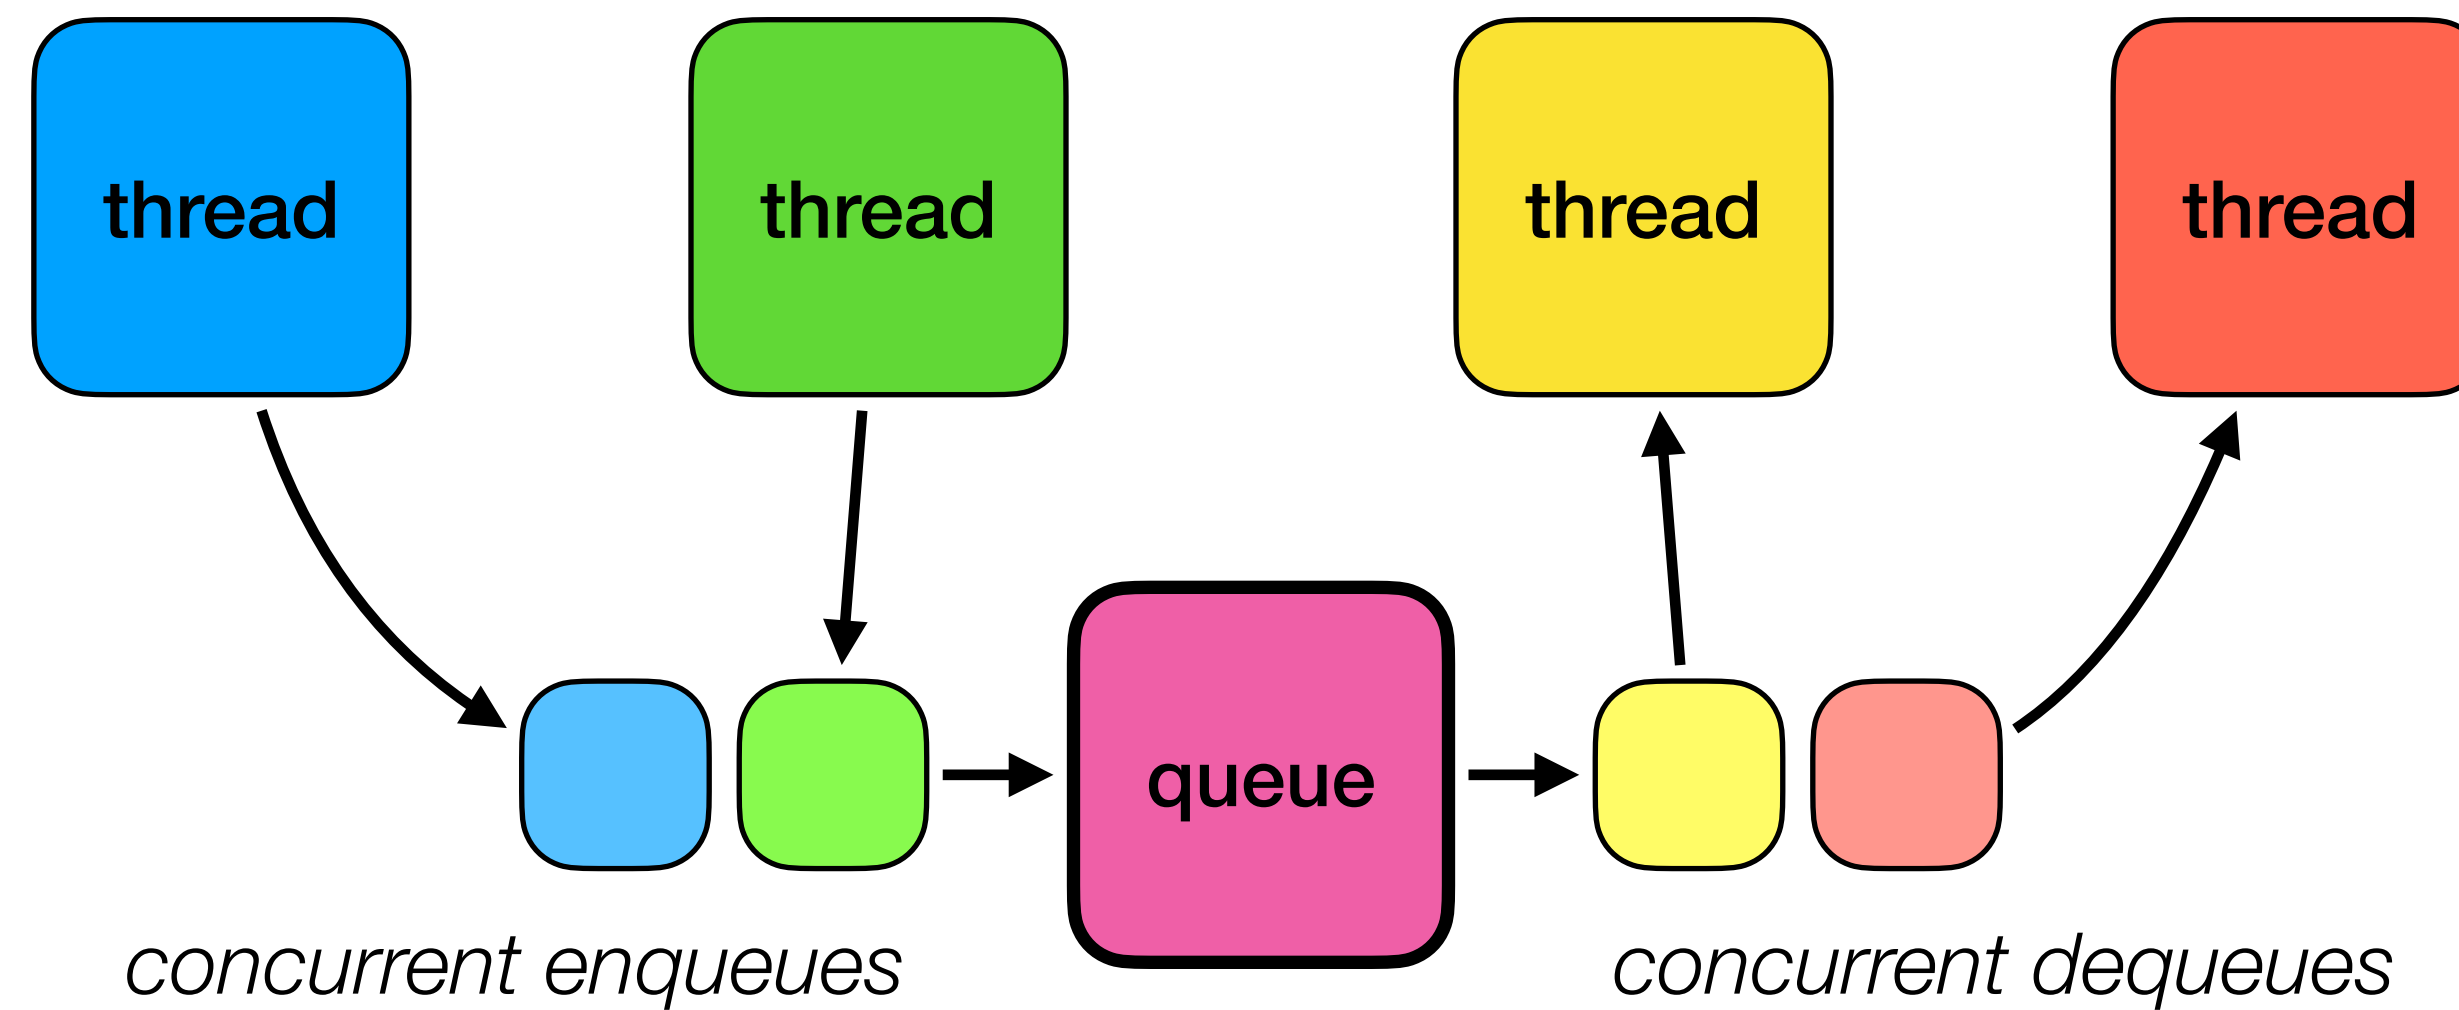


# CHECKING LINEARIZABILITY: THEORETICAL LIMITS

Constantin Enea  
Ecole Polytechnique

# Concurrent Objects

Multi-threaded programming



**e.g. Java Development Kit SE**


dozens of objects, including queues, maps, sets, lists, locks, atomic integers, ...




# Linearizability [Herlihy&Wing 1990]

Effects of each invocation appear to occur instantaneously

Execution history

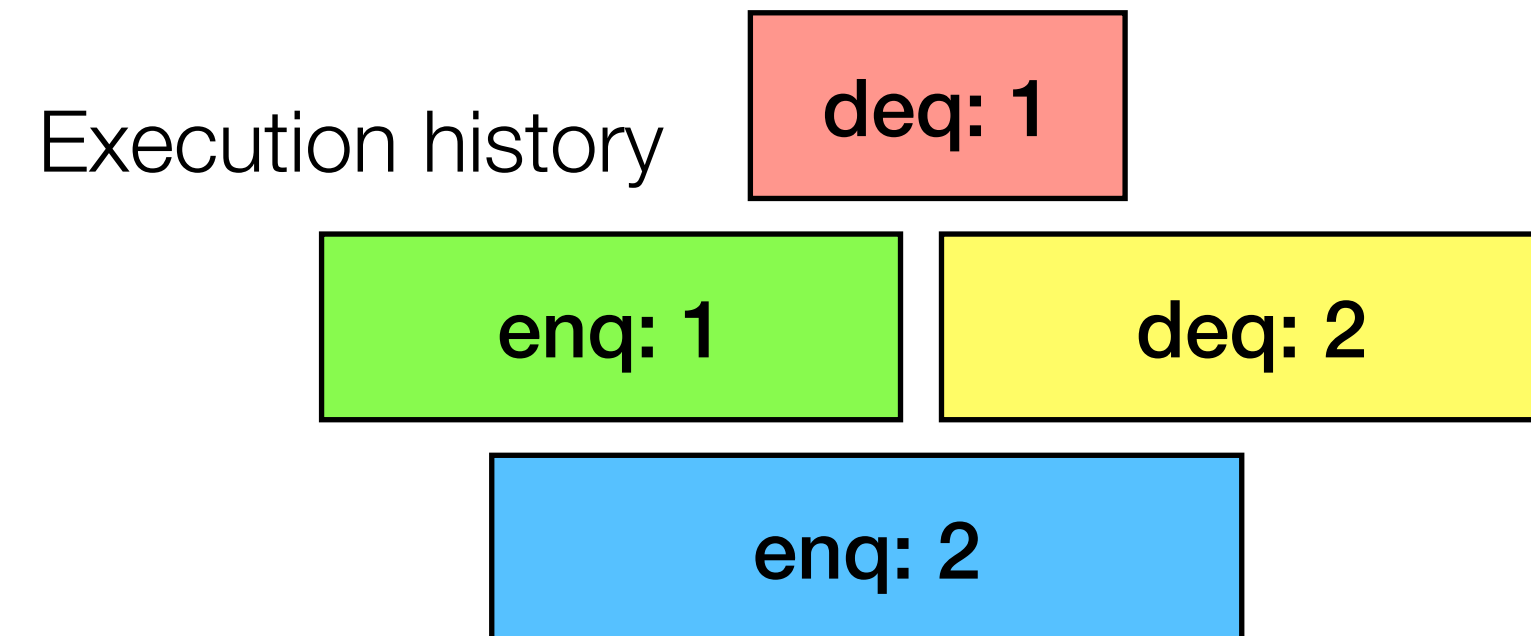


Linearization admitted by Queue ADT

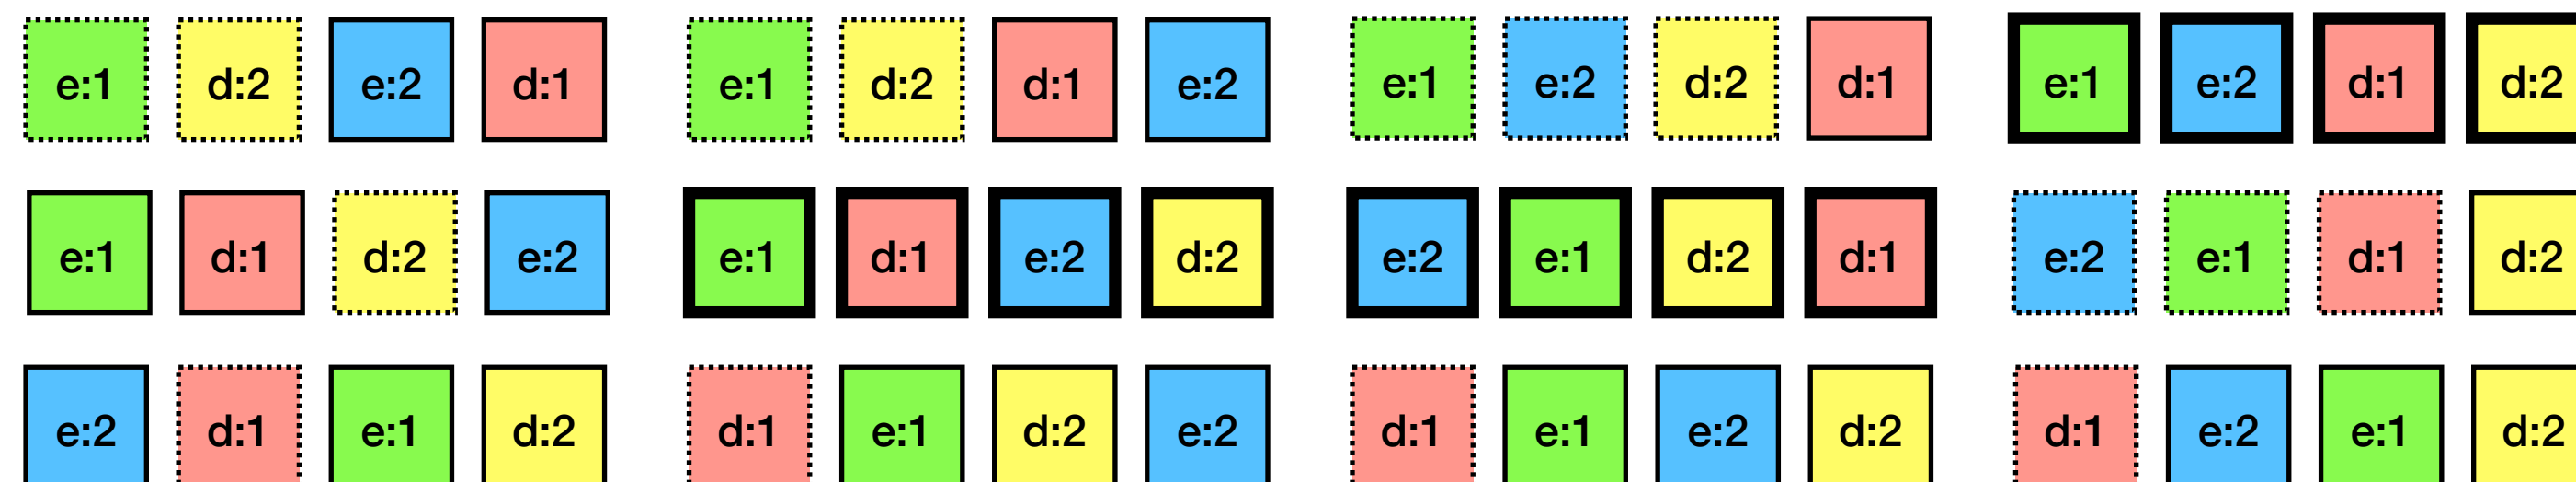


$\exists \text{ lin. } rb \subseteq \text{lin} \wedge \text{lin} \in \text{Queue ADT}$

# Complexity of Testing Linearizability



Exponentially many linearizations to consider



**Theorem** [Gibbons.et.al.'97]

Checking linearizability for a fixed execution is NP-hard

# Checking Linearizability: Complexity (finite-state implementations)

## Bounded Nb. of Threads:

- EXSPACE-complete [Alur et al., 1996, Hamza 2015]

## Unbounded Nb. of Threads:

- Undecidable [Bouajjani et al., 2013]
- Decidable with “fixed linearization points” [Bouajjani et al. 2013]

**Alur et al. 1996:** Rajeev Alur, Kenneth L. McMillan, Doron A. Peled: Model-Checking of Correctness Conditions for Concurrent Objects. LICS 1996

**Bouajjani et al., 2013:** Ahmed Bouajjani, Michael Emmi, Constantin Enea, Jad Hamza: Verifying Concurrent Programs against Sequential Specifications. ESOP 2013

**Hamza 2015:** Jad Hamza: On the Complexity of Linearizability. NETYS 2015

# Checking Linearizability: Complexity (finite-state implementations)

## **Bounded Nb. of Threads:**

- EXSPACE-complete [Alur et al., 1996, Hamza 2015]

## **Unbounded Nb. of Threads:**

- Undecidable [Bouajjani et al., 2013]
- Decidable with “fixed linearization points” [Bouajjani et al. 2013]

# Concurrent Languages

- Concurrent language =  $(\Sigma, D)$ , where  $\Sigma$  is an alphabet,  $D \subseteq \Sigma \times \Sigma$  (Mazurkiewicz traces -  $D$  is symmetric)
- $a$  and  $b$  are called **independent** when  $(a,b) \notin D$
- $\Rightarrow_D$  a relation that permutes independent symbols: for all  $(a,b) \notin D$ ,  
 $\sigma ab \sigma' \Rightarrow_D \sigma' ba \sigma$  (and trans. closure)
- $cl_D(L)$  = all strings  $\sigma'$  such that  $\sigma' \Rightarrow_D \sigma$  for some  $\sigma \in L$
- Ex:  $\Sigma = \{a,b\}$ ,  $L=(ab)^*$ ,  $D=\emptyset$  and  $D=\{(b,a)\}$

# Specifications, Implementations

- **Specification** = a language over an alphabet containing symbols  $p:m(a) \Rightarrow b$ 
  - Example: bounded-value register, bounded size queue
- **Implementation** = a language over an alphabet containing symbols  $p:call\ m(a)$  and  $p:ret\ m(a) \Rightarrow b$  where returns “match” previous calls
- $\Sigma_p = ( \Sigma_{call}(p) \cup \Sigma_{ret}(p) )$  and  $\Sigma = \bigcup_p \Sigma_p$

# Defining Linearizability

- $\text{lin} = \bigcup_p ( \Sigma_p \times \Sigma_p ) \cup ( \Sigma_{\text{ret}} \times \Sigma_{\text{call}} )$
- $\text{Spec}^*$  = replacing  $p:m(a) \Rightarrow b$  with call/ret actions
- an execution  $\sigma$  is **linearizable** iff  $\sigma \in \text{cl}_{\text{lin}}(\text{Spec}^*)$
- Impl is linearizable iff  $\text{Impl} \subseteq \text{cl}_{\text{lin}}(\text{Spec}^*)$ 
  - this inclusion check is undecidable in general (for regular languages)

# Defining Linearizability

- **Linearizability**: an execution  $\sigma$  is linearizable iff there is a sequence  $\tau$  that contains  $\sigma$  and linearization points (symbols  $p:m(a)\Rightarrow b$ ) s.t.:
  - every projection over “actions” of the same process is “sequential”
  - the projection over linearization point actions is included in the specification



# Defining Linearizability

- $\text{lin} = \cup_p ( \Sigma_p \times \Sigma_p ) \cup ( \Sigma_{\text{ret}} \times \Sigma_{\text{call}} )$
- $\text{Spec}^* =$  replacing  $p:m(a) \Rightarrow b$  with call/ret actions
- an execution  $\sigma$  is **linearizable** iff  $\sigma \in \text{cl}_{\text{lin}}(\text{Spec}^*)$
- Impl is linearizable iff  $\text{Impl} \subseteq \text{cl}_{\text{lin}}(\text{Spec}^*)$ 
  - this inclusion check is undecidable in general (for regular languages)
- $\text{cl}_{\text{lin}}(\text{Spec}^*) = ( \parallel_p L_{\text{lin\_points}}(p) \parallel \text{Spec} ) \downarrow ( \Sigma_{\text{call}} \cup \Sigma_{\text{ret}} )$

# EXPSPACE-hardness

*Problem 2 (Letter Insertion).* Input: A set of *insertable* letters  $A = \{a_1, \dots, a_l\}$ .  
An NFA  $N$  over an alphabet  $\Gamma \uplus A$ .

Question: For all words  $w \in \Gamma^*$ , does there exist a decomposition  $w = w_0 \cdots w_l$ , and a permutation  $p$  of  $\{1, \dots, l\}$ , such that  $w_0 a_{p[1]} w_1 \dots a_{p[l]} w_l$  is accepted by  $N$ ?

## Reducing Letter Insertion to Linearizability:


1. there exists a word  $w$  in  $\Gamma^*$ , such that there is no way to insert the letters from  $A$  in order to obtain a word accepted by  $N$
2. there exists an execution of  $Lib$  with  $k$  threads which is not linearizable w.r.t.  $S_N$


$$k = l+2$$


# EXPSPACE-hardness

Define  $k$ , the number of threads, to be  $l + 2$ .  
We will define a library *Lib* composed of


- methods  $M_1, \dots, M_l$ , one for each letter of  $A$
- methods  $M_\gamma$ , one for each letter of  $\Gamma$
- a method  $M_{\text{Tick}}$ .



**Fig. 4.** Description of  $M_\gamma$ ,  $\gamma \in \Gamma$



**Fig. 5.** Description of  $M_1, \dots, M_l$

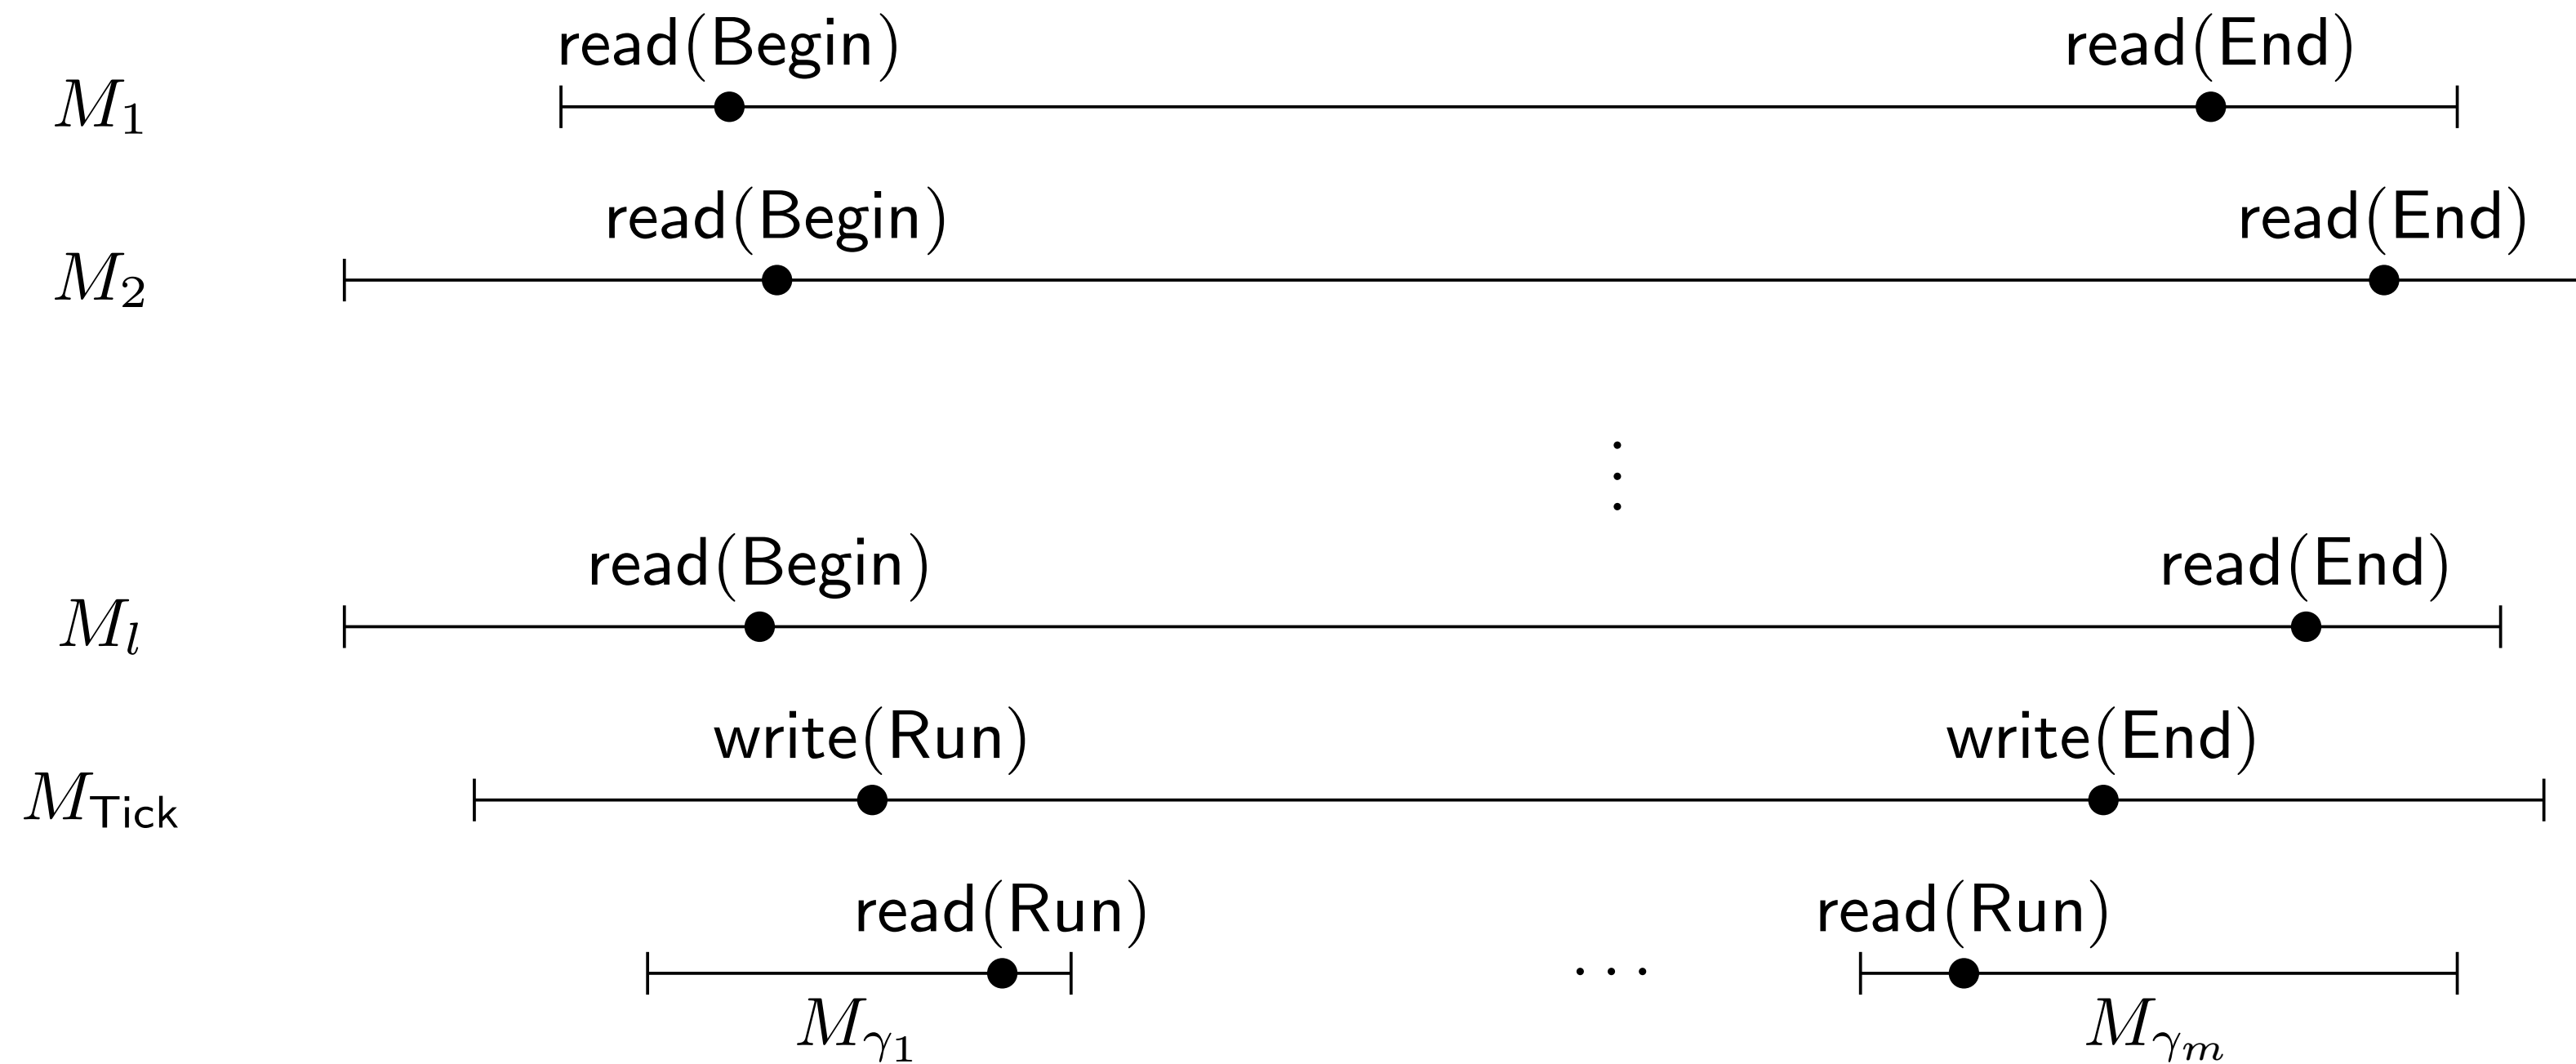


**Fig. 6.** Description of  $M_{\text{Tick}}$

The specification  $S_N$  is defined as the set of words  $w$  over the alphabet  $\{M_1, \dots, M_l\} \cup \{M_{\text{Tick}}\} \cup \{M_\gamma | \gamma \in \Gamma\}$  such that one the following condition holds:

- $w$  contains 0 letter  $M_{\text{Tick}}$ , or more than 1, or
- for a letter  $M_i$ ,  $i \in \{1, \dots, l\}$ ,  $w$  contains 0 such letter, or more than 1, or
- when projecting over the letters  $M_\gamma$ ,  $\gamma \in \Gamma$  and  $M_i$ ,  $i \in \{1, \dots, l\}$ ,  $w$  is in  $N_M$ , where  $N_M$  is  $N$  where each letter  $\gamma$  is replaced by the letter  $M_\gamma$ , and where each letter  $a_i$  is replaced by the letter  $M_i$ .

# EXPSPACE-hardness



**Fig. 7.** Non-linearizable execution corresponding to a word  $\gamma_1 \dots \gamma_m$  in which we cannot insert the letters from  $A = \{a_1, \dots, a_l\}$  to make it accepted by  $N$ . The points represent steps in the automata.

# Checking Linearizability: Complexity (finite-state implementations)

## **Bounded Nb. of Threads:**

- EXSPACE-complete [Alur et al., 1996, Hamza 2015]

## **Unbounded Nb. of Threads:**

- Undecidable [Bouajjani et al., 2013]
- Decidable with “fixed linearization points” [Bouajjani et al. 2013]

# Undecidability

- Reduction from reachability in counter machines
- Given a counter machine  $A$ , we construct a library  $L_A$  and a specification  $S_A$  such that  $L_A$  is **not** linearizable w.r.t.  $S_A$  iff  $A$  reaches the target state
- $L_A$  = transition methods  $T[t]$ , increments  $I[c_i]$ , decrements  $D[c_i]$  and zero-tests  $Z[c_i]$
- $L_A$  allows only valid sequences of transitions
- $S_A$  allows executions which don't reach the target state, or which erroneously pass

some zero-test

- it doesn't contain  $M[q_f]$ ,
- it ends in  $M[q_f]$  and it contains a prefix of the form

$$(M\_inc[i] M\_dec[i])^* (M\_inc[i]^+ + M\_dec[i]^+) M\_zero[i]$$

- it ends in  $M_f$  and it contains a subword of the form

$$M\_zero[i] (M\_inc[i] M\_dec[i])^* (M\_inc[i]^+ + M\_dec[i]^+) M\_zero[i].$$

# Undecidability

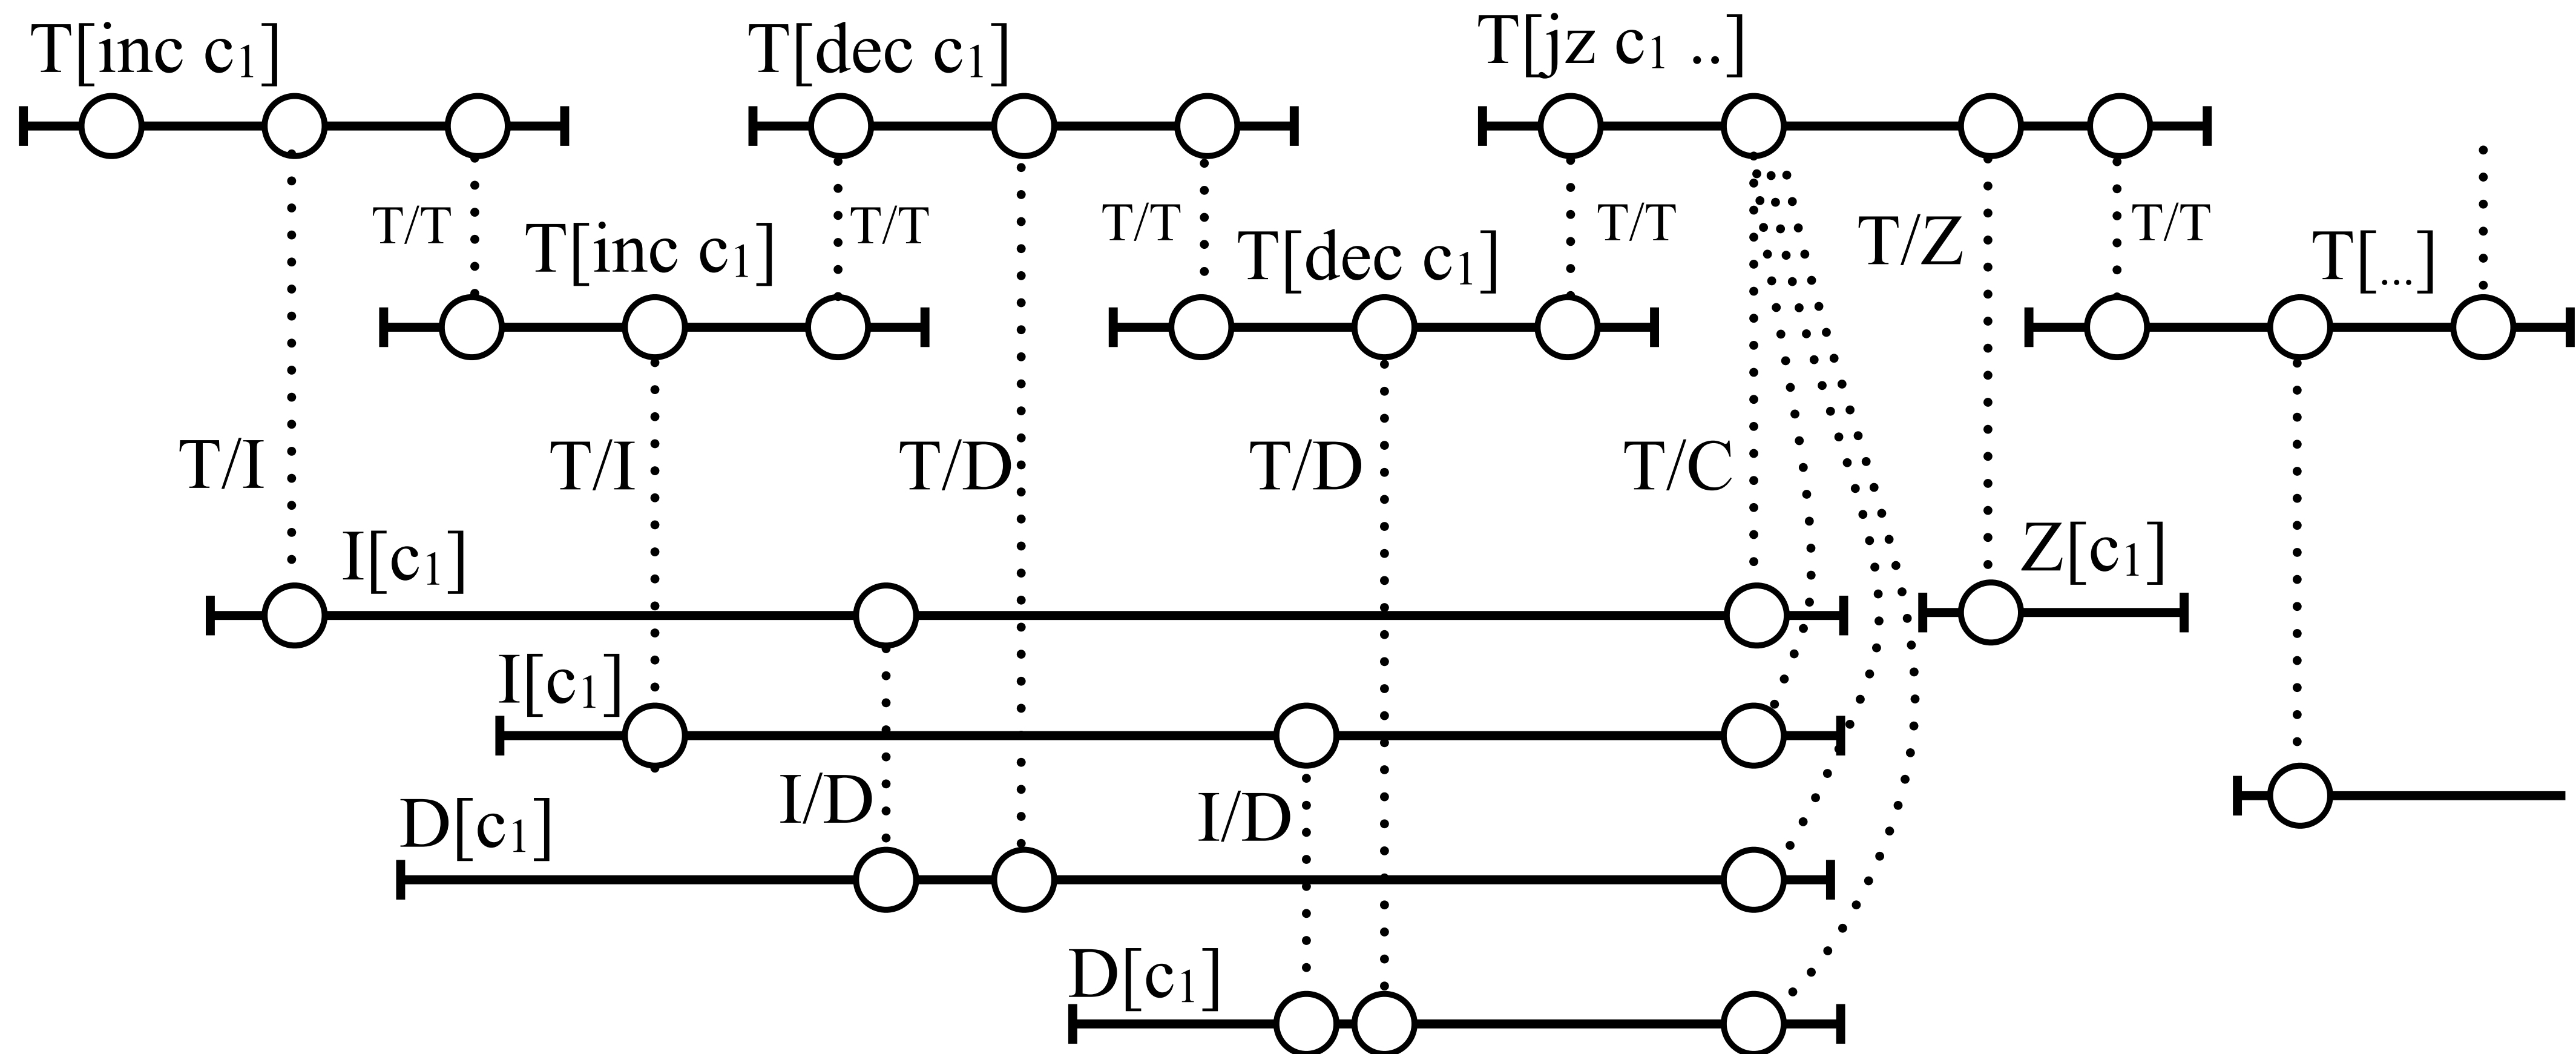
1. A sequence  $t_1 t_2 \dots t_i$  of  $\mathcal{A}$ -transitions is modeled by a pairwise-overlapping sequence of  $T[t_1] \cdot T[t_2] \cdots T[t_i]$  operations.
2. Each  $T[t]$ -operation has a corresponding  $I[c_i]$ ,  $D[c_i]$ , or  $Z[c_i]$  operation, depending on whether  $t$  is, resp., an increment, decrement, or zero-test transition with counter  $c_i$ .
3. Each  $I[c_i]$  operation has a corresponding  $D[c_i]$  operation.
4. For each counter  $c_i$ , all  $I[c_i]$  and  $D[c_i]$  between  $Z[c_i]$  operations overlap.
5. For each counter  $c_i$ , no  $I[c_i]$  nor  $D[c_i]$  operations overlap with a  $Z[c_i]$  operation.
6. The number of  $I[c_i]$  operations between two  $Z[c_i]$  operations matches the number of  $D[c_i]$  operations.

# Undecidability

- a T/T signal between  $T[*]$  operations
- for each counter  $c$ , a T/I, T/D, T/Z between  $T[*]$  operations and, resp.,  $I[c_i]$ ,  $D[c_i]$  and  $Z[c_i]$  operations
- an I/D signal between  $I[c_i]$  and  $D[c_i]$  operations
- a T/C signal between  $T[t]$  operations and  $I[c_i]$ ,  $D[c_i]$  operations, for zero-testing transitions  $t$



# Undecidability



# Undecidability

```
1 var  $q \in Q$ :  $T$ 
2 var req[ $U$ ]:  $T$ 
3 var ack[ $U$ ]:  $T$ 
4 var dec[ $i \in \mathbb{N} : i < d$ ]:  $T$ 
5 var zero[ $i \in \mathbb{N} : i < d$ ]:  $\mathbb{B}$ 
6
7 // for each transition  $\langle q, n, q' \rangle$ 
8 method M[ $q, n, q'$ ] ()
9     atomic
10         wait( $q$ );
11         signal(req[ $n$ ]);
12     atomic
13         wait(ack[ $n$ ]);
14         signal( $q'$ );
15     return ()
16
17 // for each transition  $\langle q, i, q' \rangle$ 
18 method M[ $q, i, q'$ ] ()
19     atomic
20         wait( $q$ );
21         zero[ $i$ ] := true;
22     atomic
23         if !zero[ $i$ ] then
24             signal( $q'$ );
25     return ()
26
27 // for each final state  $q_f$ 
28 method M[ $q_f$ ] ()
29     wait( $q_f$ );
30     return
31
32 method M_inc[ $i$ ] ()
33     atomic
34         if !zero[ $i$ ] then
35             wait(req[ $u_i$ ]);
36             signal(ack[ $u_i$ ]);
37
38             signal(dec[ $i$ ])
39     assume zero[ $i$ ];
40     return ()
41
42 method M_dec[ $i$ ] ()
43     atomic
44         if !zero[ $i$ ] then
45             wait(dec[ $i$ ]);
46     atomic
47         wait(req[ $-u_i$ ]);
48         signal(ack[ $-u_i$ ]);
49     assume zero[ $i$ ];
50     return ()
51
52 method M_zero[ $i$ ] ()
53     atomic
54         if zero[ $i$ ] then
55             zero[ $i$ ] := false;
56
57     return ()
```

# Checking Linearizability: Complexity (finite-state implementations)

## **Bounded Nb. of Threads:**

- EXSPACE-complete [Alur et al., 1996, Hamza 2015]

## **Unbounded Nb. of Threads:**

- Undecidable [Bouajjani et al., 2013]
- Decidable with “fixed linearization points” [Bouajjani et al. 2013]

# Libraries

A *method* is a finite automaton  $M = \langle Q, \Sigma, I, F, \hookrightarrow \rangle$  with labeled transitions  $\langle m_1, v_1 \rangle \xrightarrow{a} \langle m_2, v_2 \rangle$  between method-local states  $m_1, m_2 \in Q$  paired with finite-domain shared-state valuations  $v_1, v_2 \in V$ . The initial and final states  $I, F \subseteq Q$  represent the method-local states passed to, and returned from,  $M$ .

A *client* of a library  $L$  is a finite automaton  $C = \langle Q, \Sigma, \ell_0, \hookrightarrow \rangle$  with initial state  $\ell_0 \in Q$  and transitions  $\hookrightarrow \subseteq Q \times \Sigma \times Q$  labeled by the alphabet  $\Sigma = \{M(m_0, m_f) : M \in L, m_0, m_f \in Q_M\}$  of library method calls

*most general client*  $C^* = \langle Q, \Sigma, \ell_0, \hookrightarrow \rangle$  of a library  $L$  nondeterministically calls  $L$ 's methods in any order:  $Q = \{\ell_0\}$  and  $\hookrightarrow = Q \times \Sigma \times Q$ .

# Example: Coarse-Grain Sets

```
remove(k):  
Node p, c  
Lock l  
1 lock l; LP  
2 p = H;  
3 c = p.next  
4 while (c.key < k)  
5     p = c;  
6     c = c.next;  
7 if (c.key = k)  
8     p.next = c.next  
9     return true  
10 else  
11     return false  
12 unlock l
```

```
ctn(k):  
Node p, c  
Lock l  
1 lock l; LP  
2 p = H;  
3 c = p.next  
4 while (c.key < k)  
5     p = c;  
6     c = c.next;  
7 if (c.key = k)  
8     return true  
9 else  
10     return false  
11 unlock l
```

# Example: Fine-Grain Sets

```
add(k):  
Node p, c  
1 lock(H)  
2 p = H  
3 c = p.next;  
4 lock(c);  
5 while (c.key < k)  
6     unlock(p);  
7     p = c;  
8     c = c.next;  
9     lock(c);  
10 if (c.key = k)  
11     return false  
12 else  
13     n = new Node(k,-)  
14     n.next = c  
15     p.next = n  
16     return true  
17     unlock(c)  
18     unlock(p)
```

LP

**“time point at which I find  
a key  $\geq$  input key”**

**“time point at which I  
lock the cell  
with key  $\geq$  input key”**



# The Treiber Stack Algorithm

## CAS: Compare And Swap

```
CAS (x,expected,new):  
Integer tmp  
1  Atomically  
2    tmp = x  
3    if (tmp == expected)  
4      x = new  
5      return true  
6  else  
7    return false
```

hardware

Intel, AMD

CMPXCHG

SPARC

CAS

software

java

CompareAndSet

C#

CompareExchange

CAS(x,5,8)

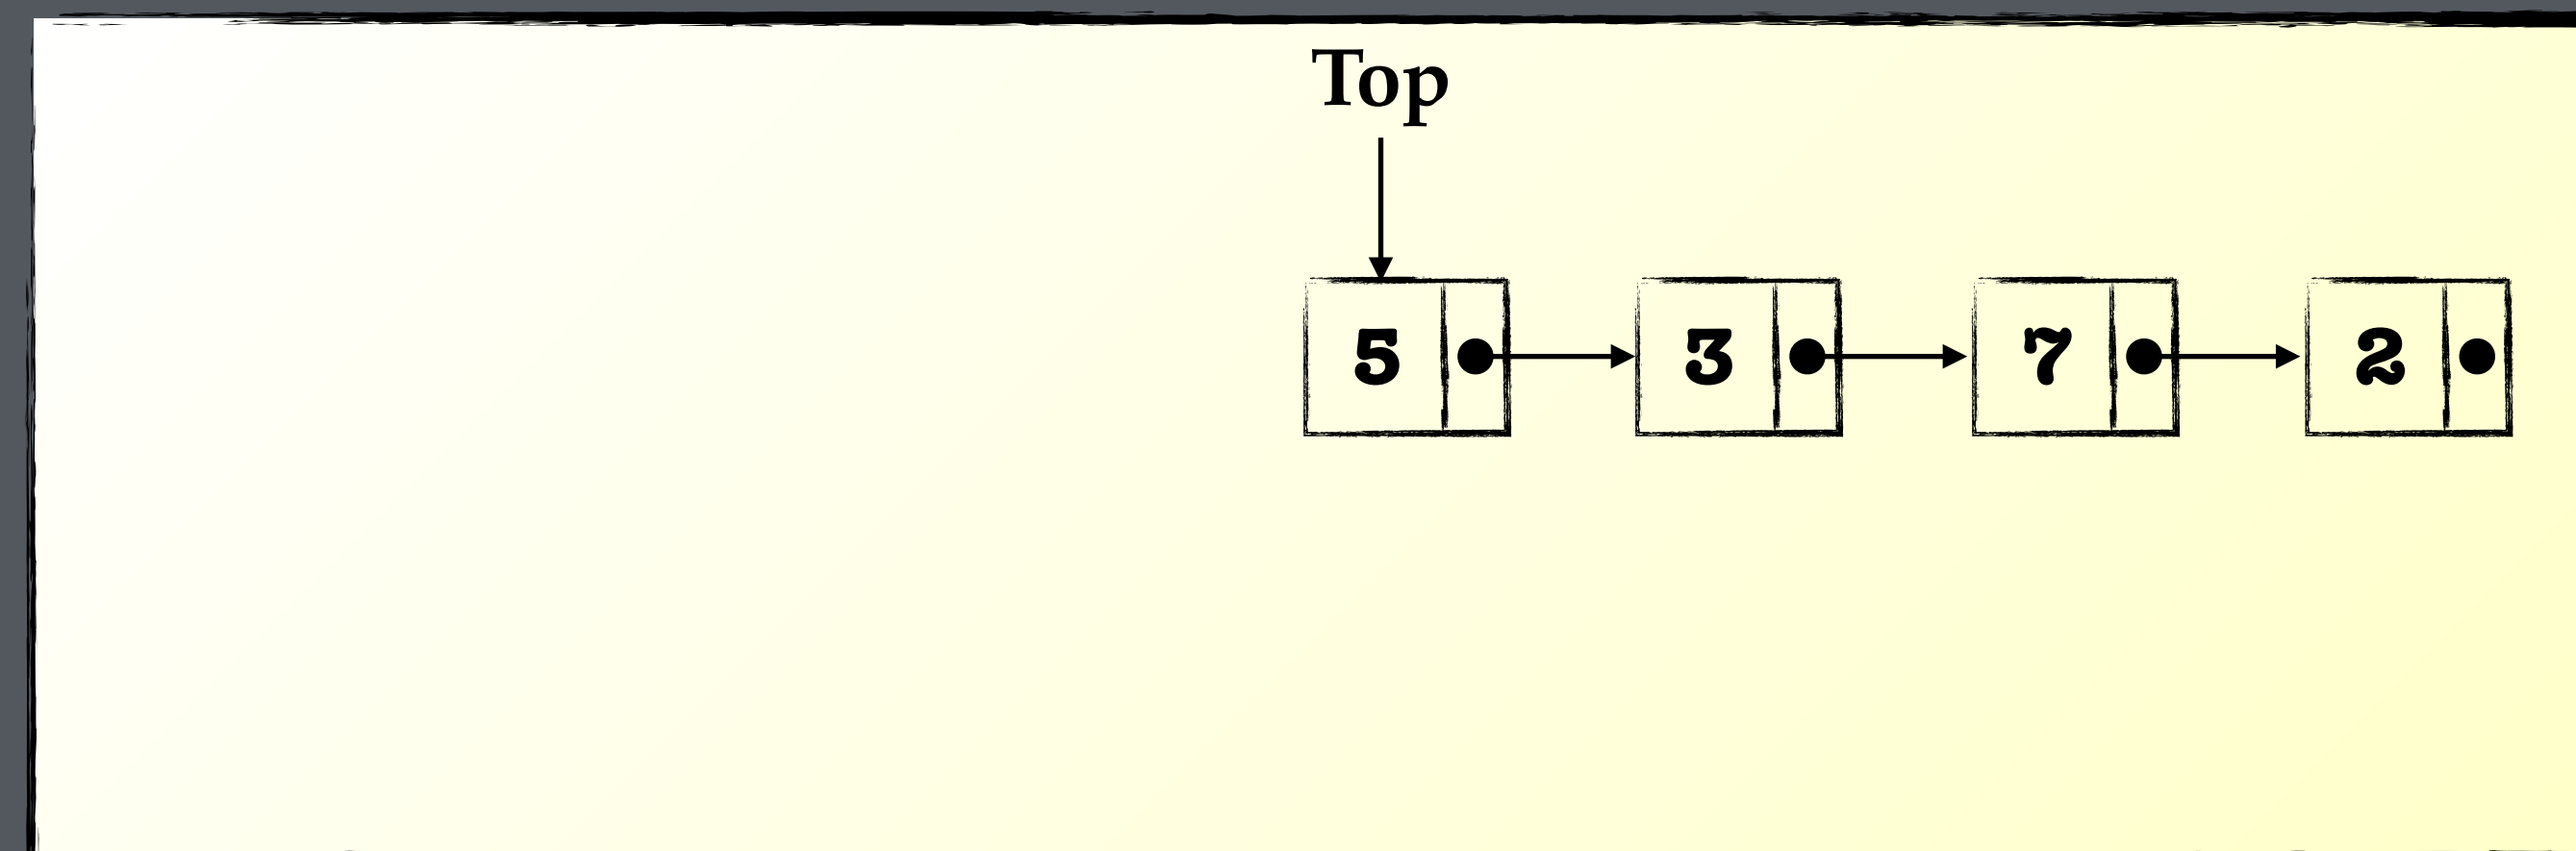

false

memory

7

x


# The Treiber Stack Algorithm





# The Treiber Stack Algorithm

```
push(k):  
Node t  
1 n = new Node(k,-) ←  
2 while (true)  
3   t = Top  
4   n.next = t  
5   if (CAS (Top,t,n))  
6     exit
```







push(4)


stack


5  
3  
7  
2

Top



# The Treiber Stack Algorithm

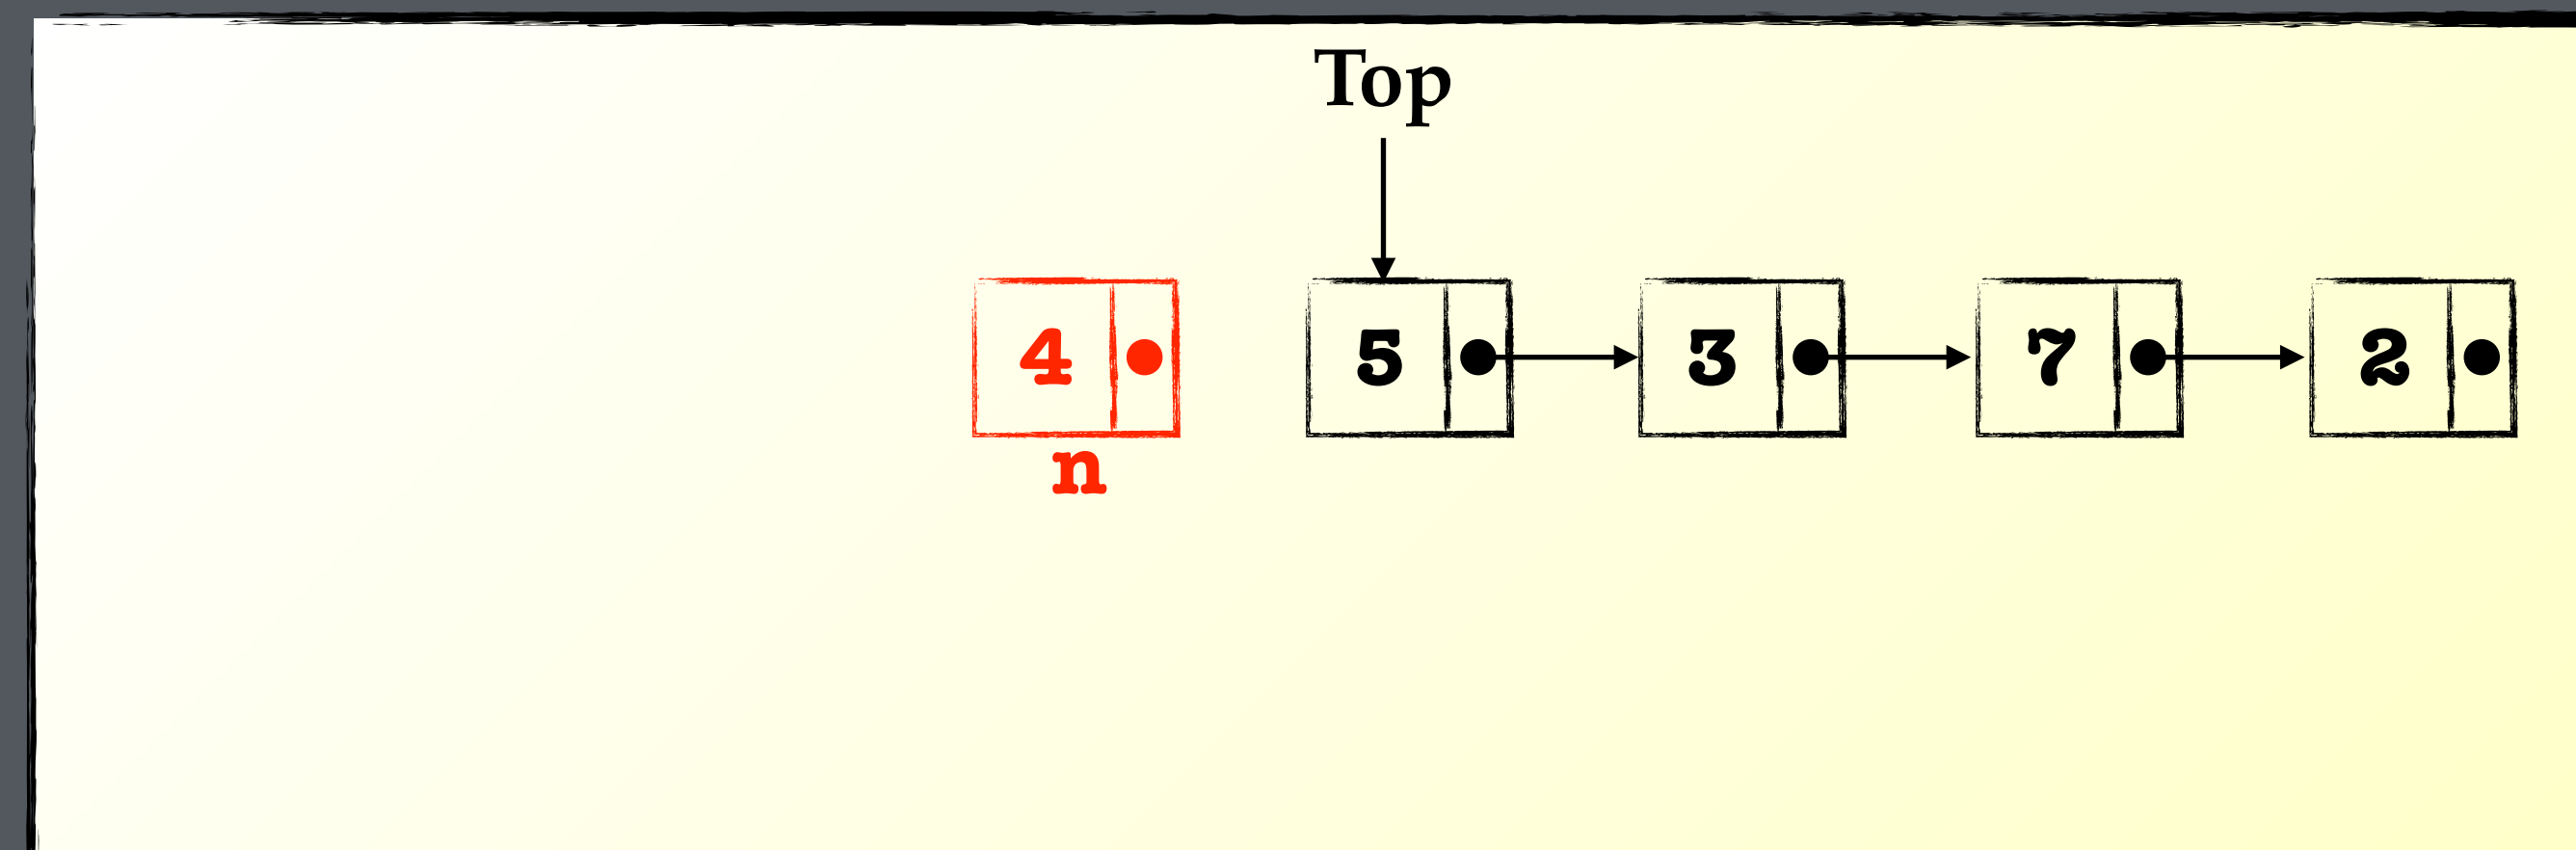
```
push(k):  
Node t  
1 n = new Node(k,-)  
2 while (true)   
3   t = Top  
4   n.next = t  
5   if (CAS (Top,t,n))  
6     exit
```



**push(4)**


**stack**

**5  
3  
7  
2**



# The Treiber Stack Algorithm

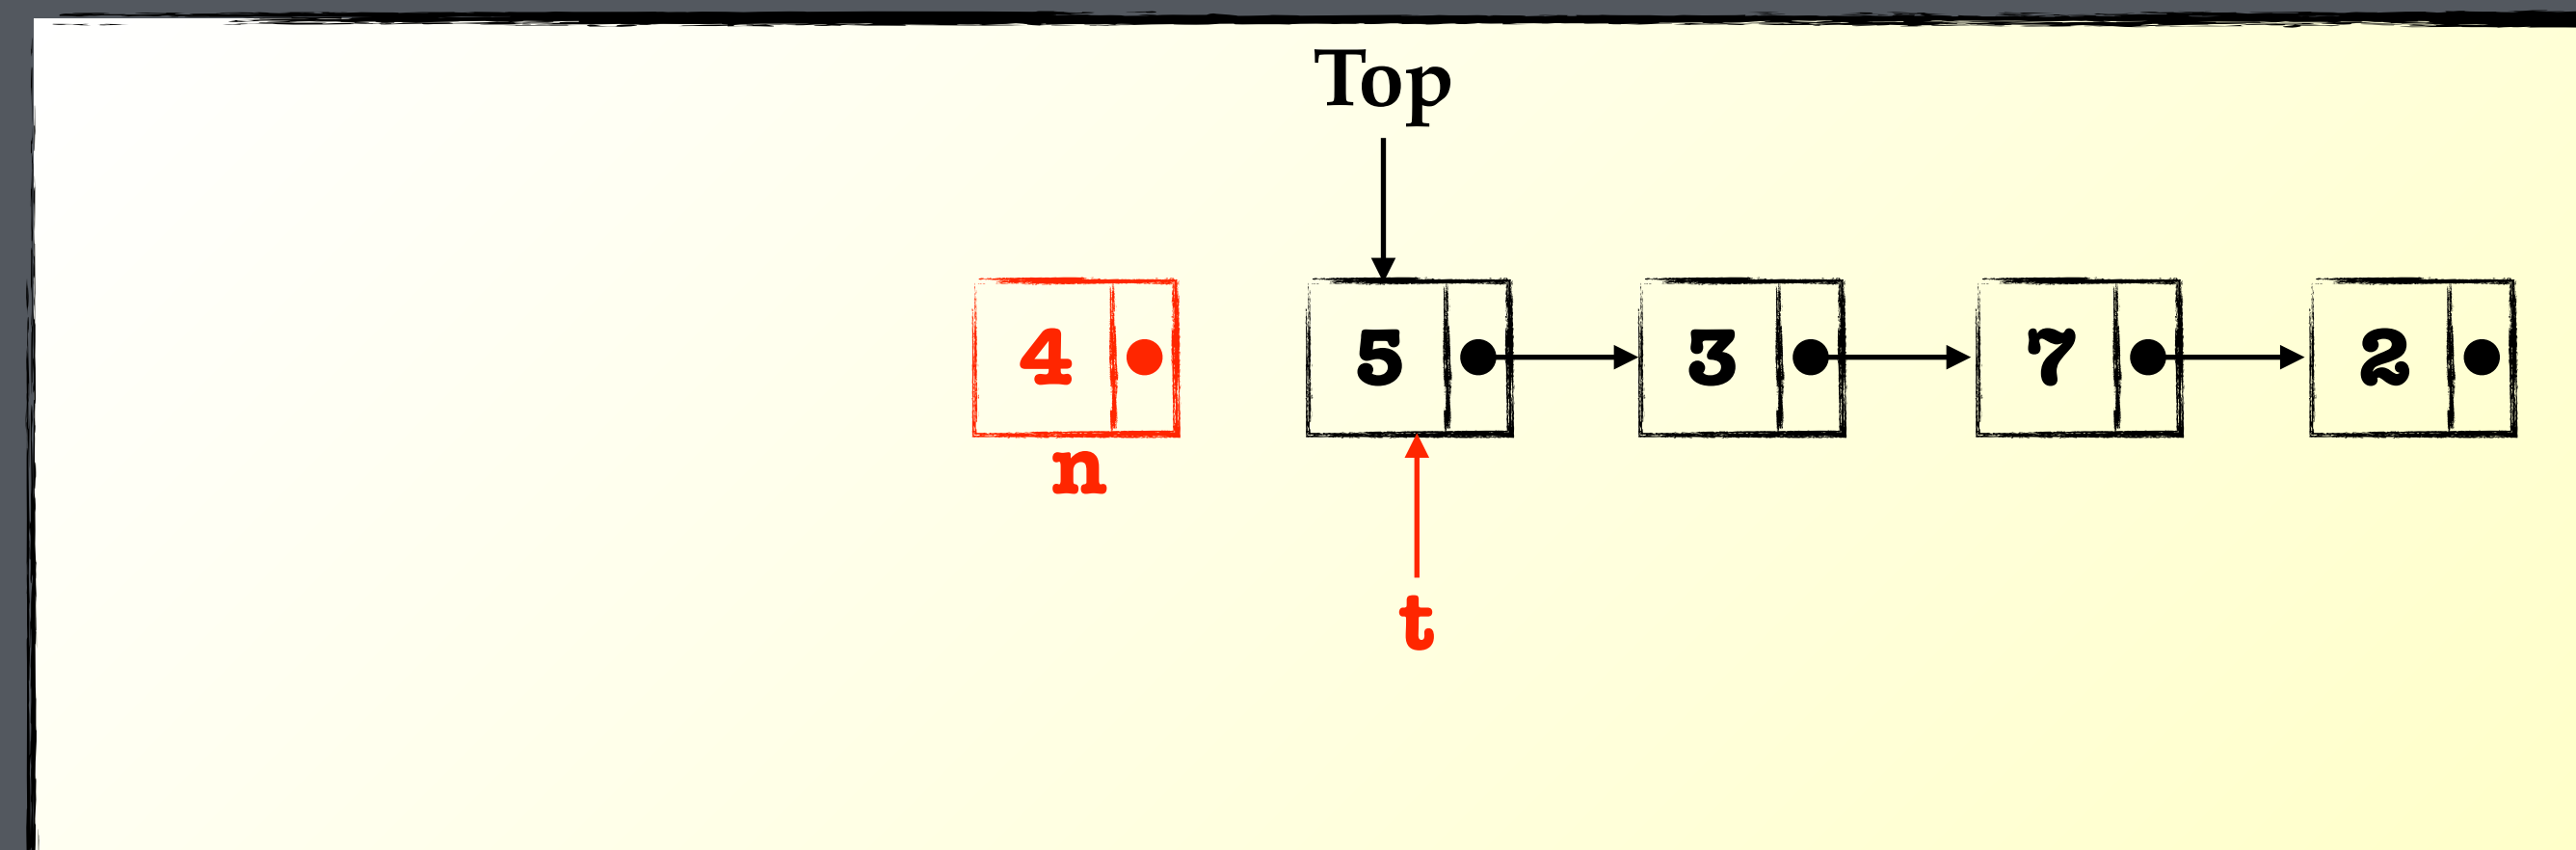
```
push(k):  
Node t  
1 n = new Node(k,-)  
2 while (true)  
3   t = Top ←~~~~~  
4   n.next = t  
5   if (CAS (Top,t,n))  
6     exit
```



push(4)


stack

5  
3  
7  
2



# The Treiber Stack Algorithm


```
push(k):  
Node t  
1 n = new Node(k,-)  
2 while (true)  
3   t = Top  
4   n.next = t  
5   if (CAS (Top,t,n))  
6     exit
```



push(4)


stack

5  
3  
7  
2



# The Treiber Stack Algorithm

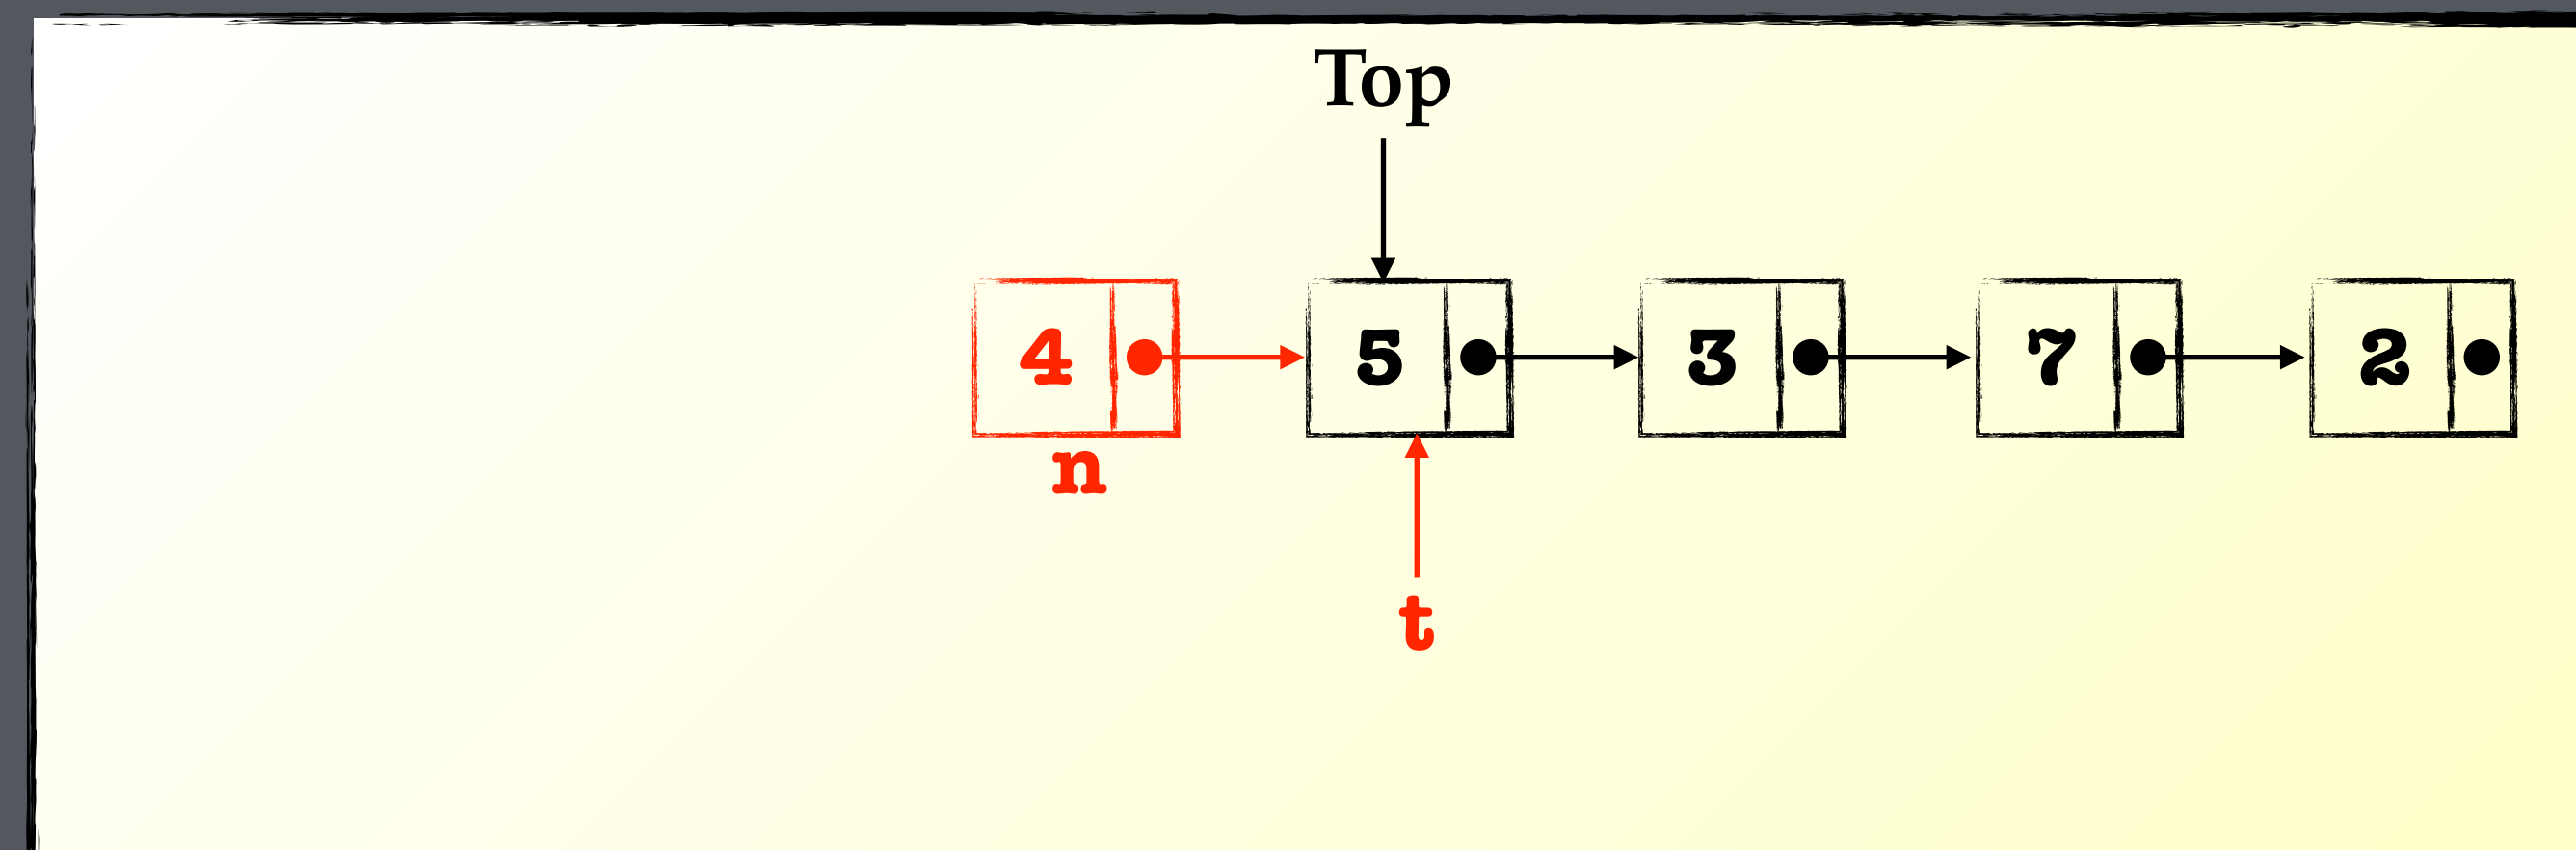
```
push(k):  
Node t  
1 n = new Node(k,-)  
2 while (true)  
3   t = Top  
4   n.next = t  
5   if (CAS (Top,t,n))  
6     exit
```



push(4)

stack


5  
3  
7  
2





# The Treiber Stack Algorithm

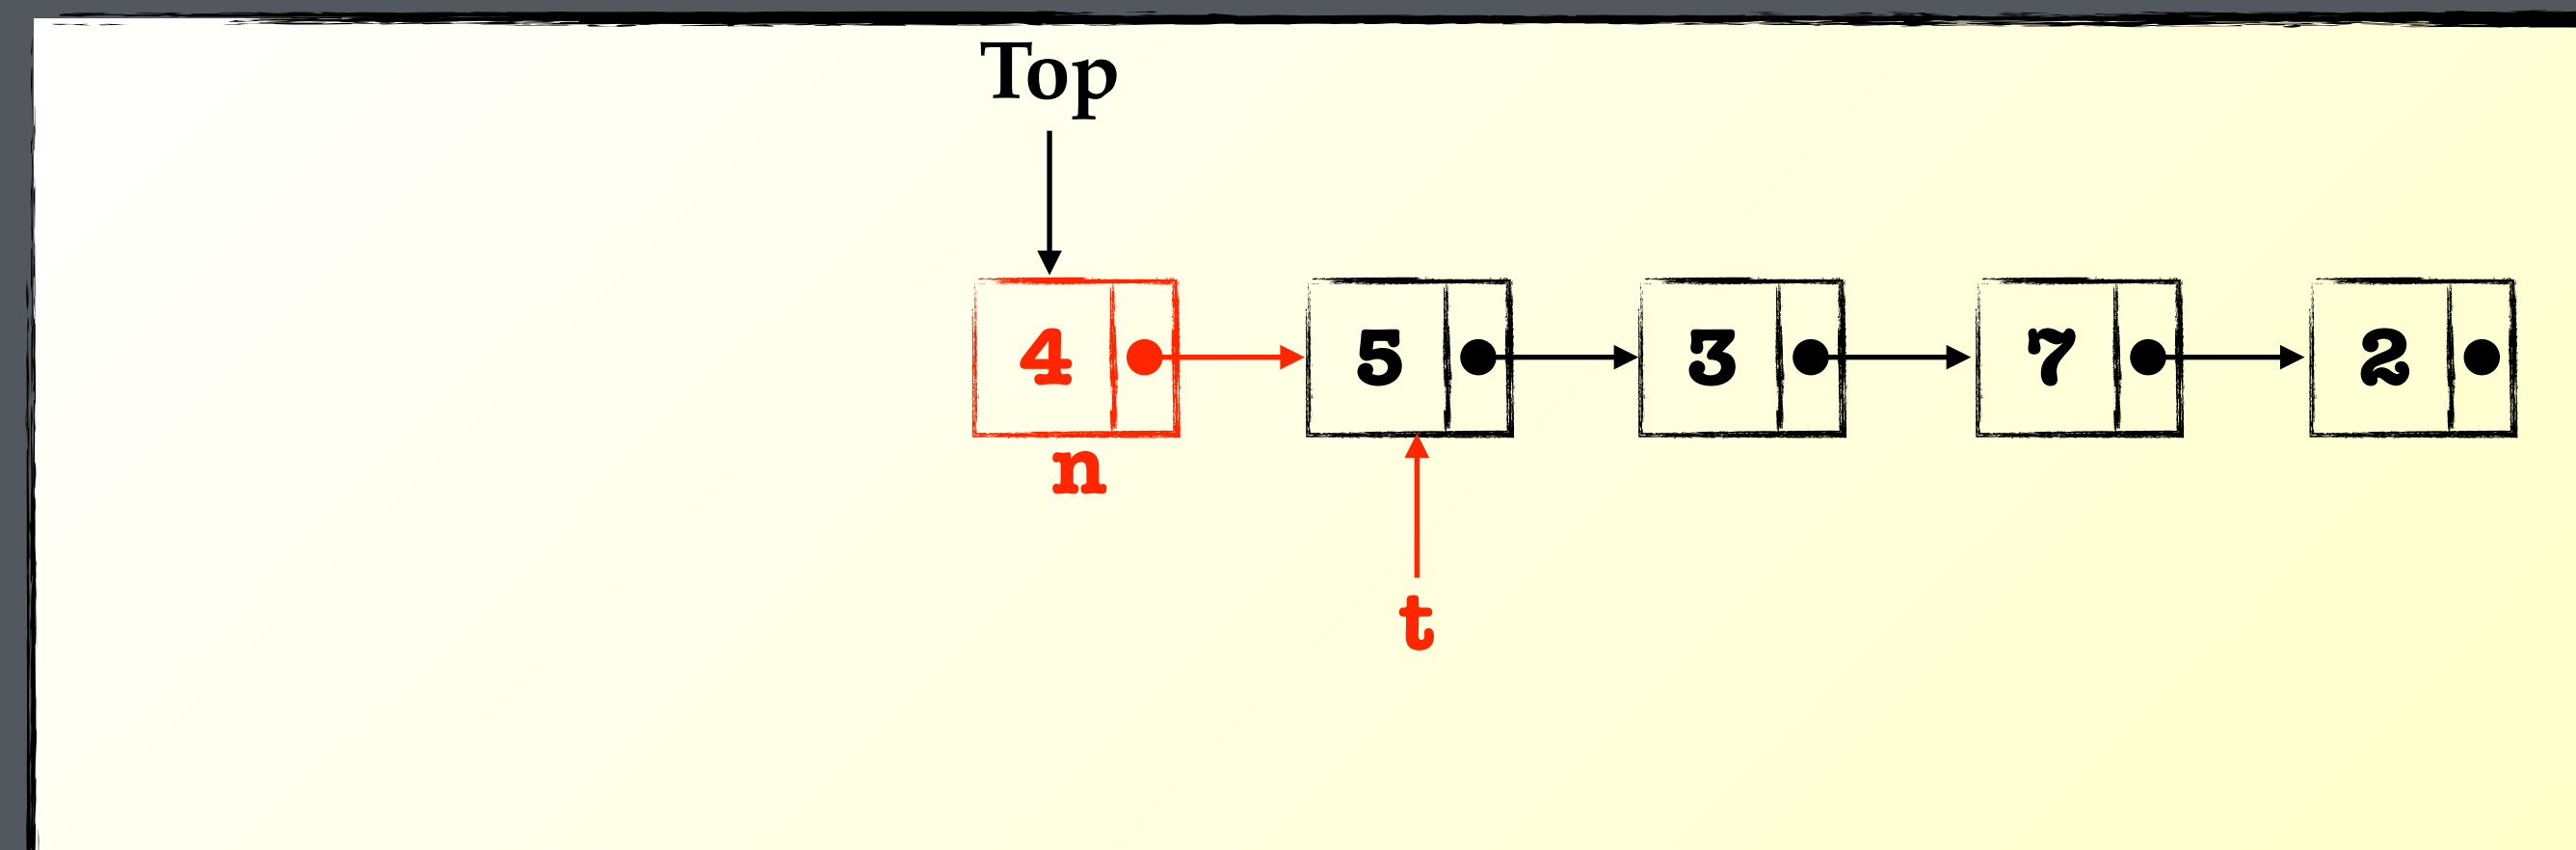
```
push(k):  
Node t  
1 n = new Node(k,-)  
2 while (true)  
3   t = Top  
4   n.next = t  
5   if (CAS (Top,t,n))  
6     exit
```



push(4)


stack

5  
3  
7  
2



# The Treiber Stack Algorithm

```
push(k):  
Node t  
1 n = new Node(k,-)  
2 while (true)  
3   t = Top  
4   n.next = t  
5   if (CAS (Top,t,n))  
6     exit ← wavy arrow
```




**push(4)**

**stack**


5  
3  
7  
2

Top



# The Treiber Stack Algorithm

```
push(k):  
Node t  
1 n = new Node(k,-)  
2 while (true)  
3   t = Top  
4   n.next = t  
5   if (CAS (Top,t,n))  
6     exit ←~~~~~
```



**push(4)**

**stack**

4  
5  
3  
7  
2

Top






# The Treiber Stack Algorithm


Concurrent push operations

```
push(k):  
Node t  
1 n = new Node(k,-) ← wavy red arrow  
2 while (true)  
3   t = Top  
4   n.next = t  
5   if (CAS (Top,t,n))  
6     exit
```

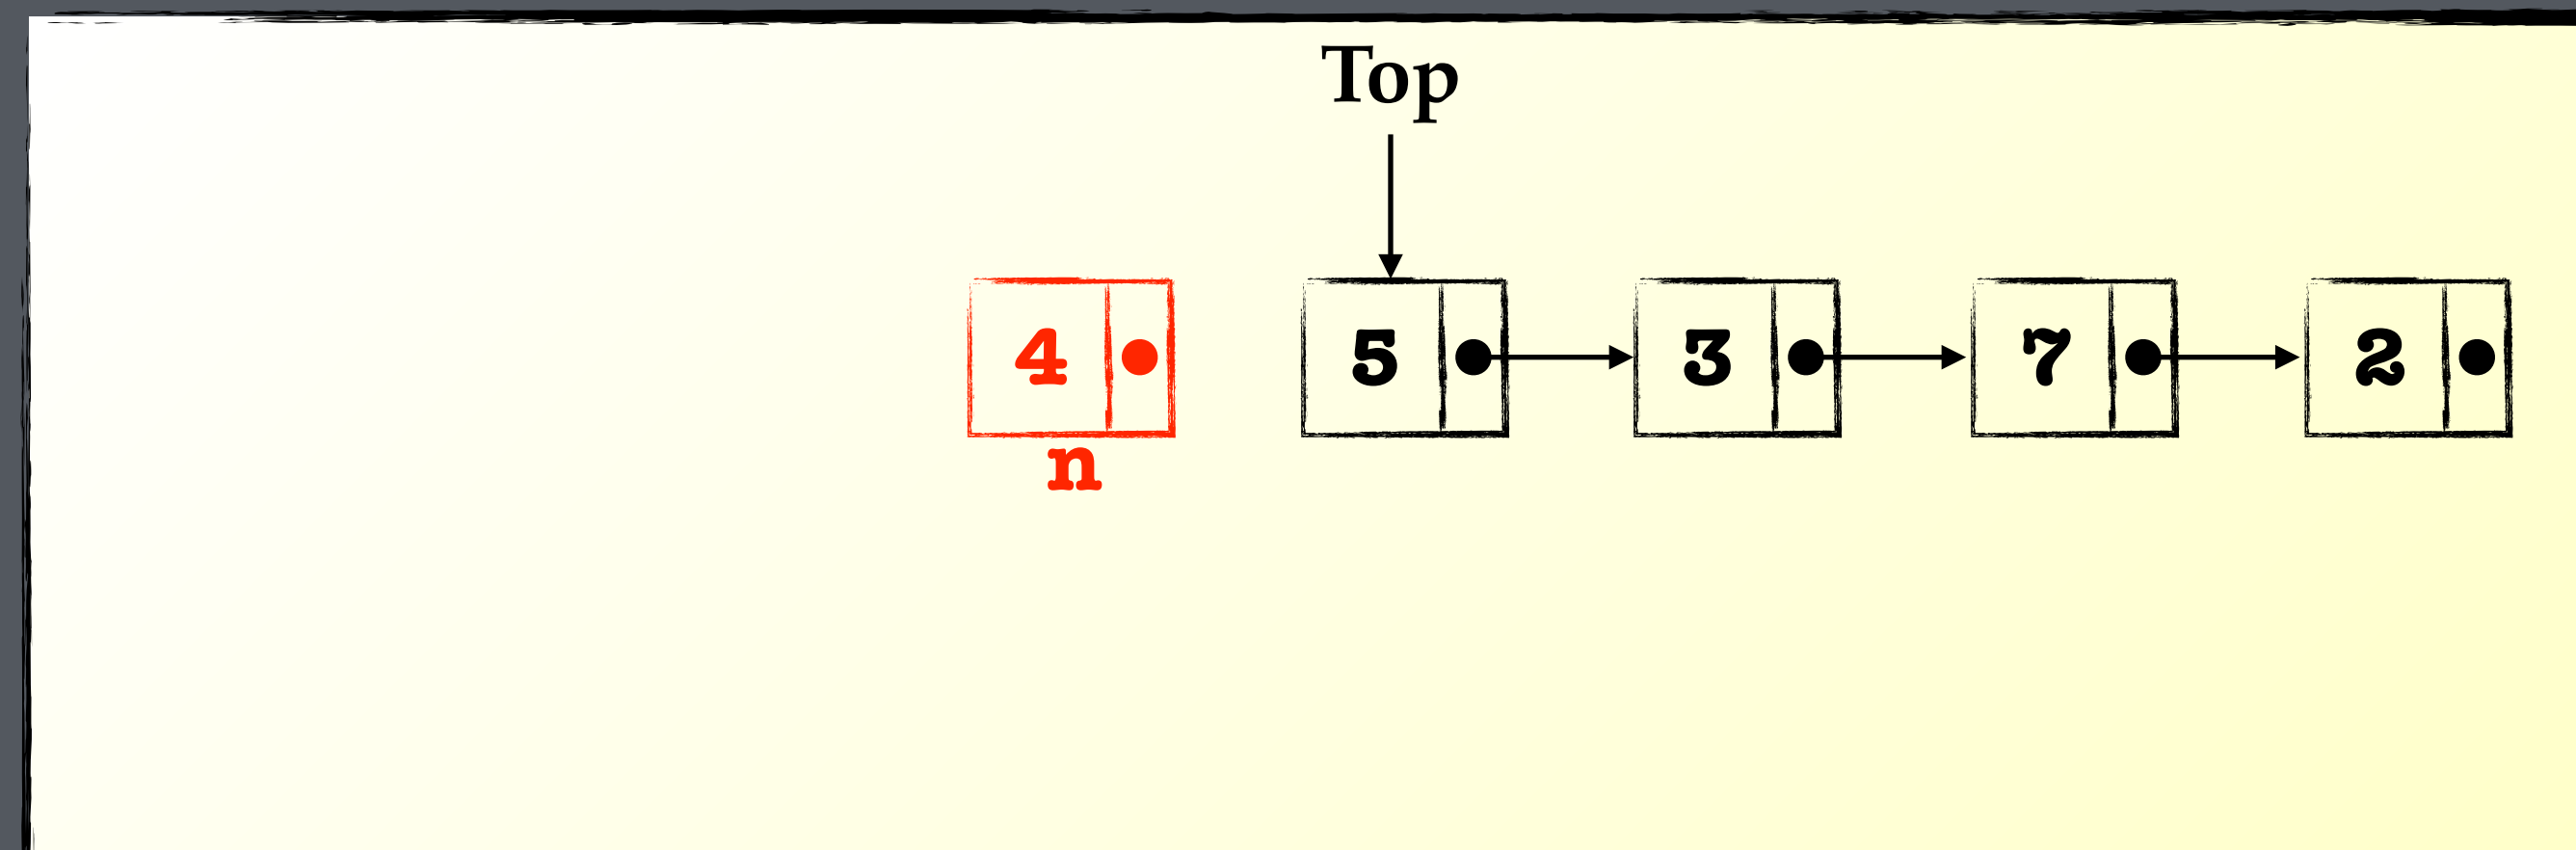


**push(4)**

**push(2)**




```
push(k):  
Node t  
1 n = new Node(k,-) ← wavy red arrow  
2 while (true)  
3   t = Top  
4   n.next = t  
5   if (CAS (Top,t,n))  
6     exit
```



# The Treiber Stack Algorithm


Concurrent push operations

```
push(k):  
Node t  
1 n = new Node(k,-)  
2 while (true) ← wavy red arrow  
3   t = Top  
4   n.next = t  
5   if (CAS (Top,t,n))  
6     exit
```




**push(4)**

**push(2)**




```
push(k):  
Node t  
1 n = new Node(k,-) ← wavy red arrow  
2 while (true)  
3   t = Top  
4   n.next = t  
5   if (CAS (Top,t,n))  
6     exit
```



# The Treiber Stack Algorithm


Concurrent push operations

```
push(k):  
Node t  
1 n = new Node(k,-)  
2 while (true)  
3   t = Top  
4   n.next = t  
5   if (CAS (Top,t,n))  
6     exit
```

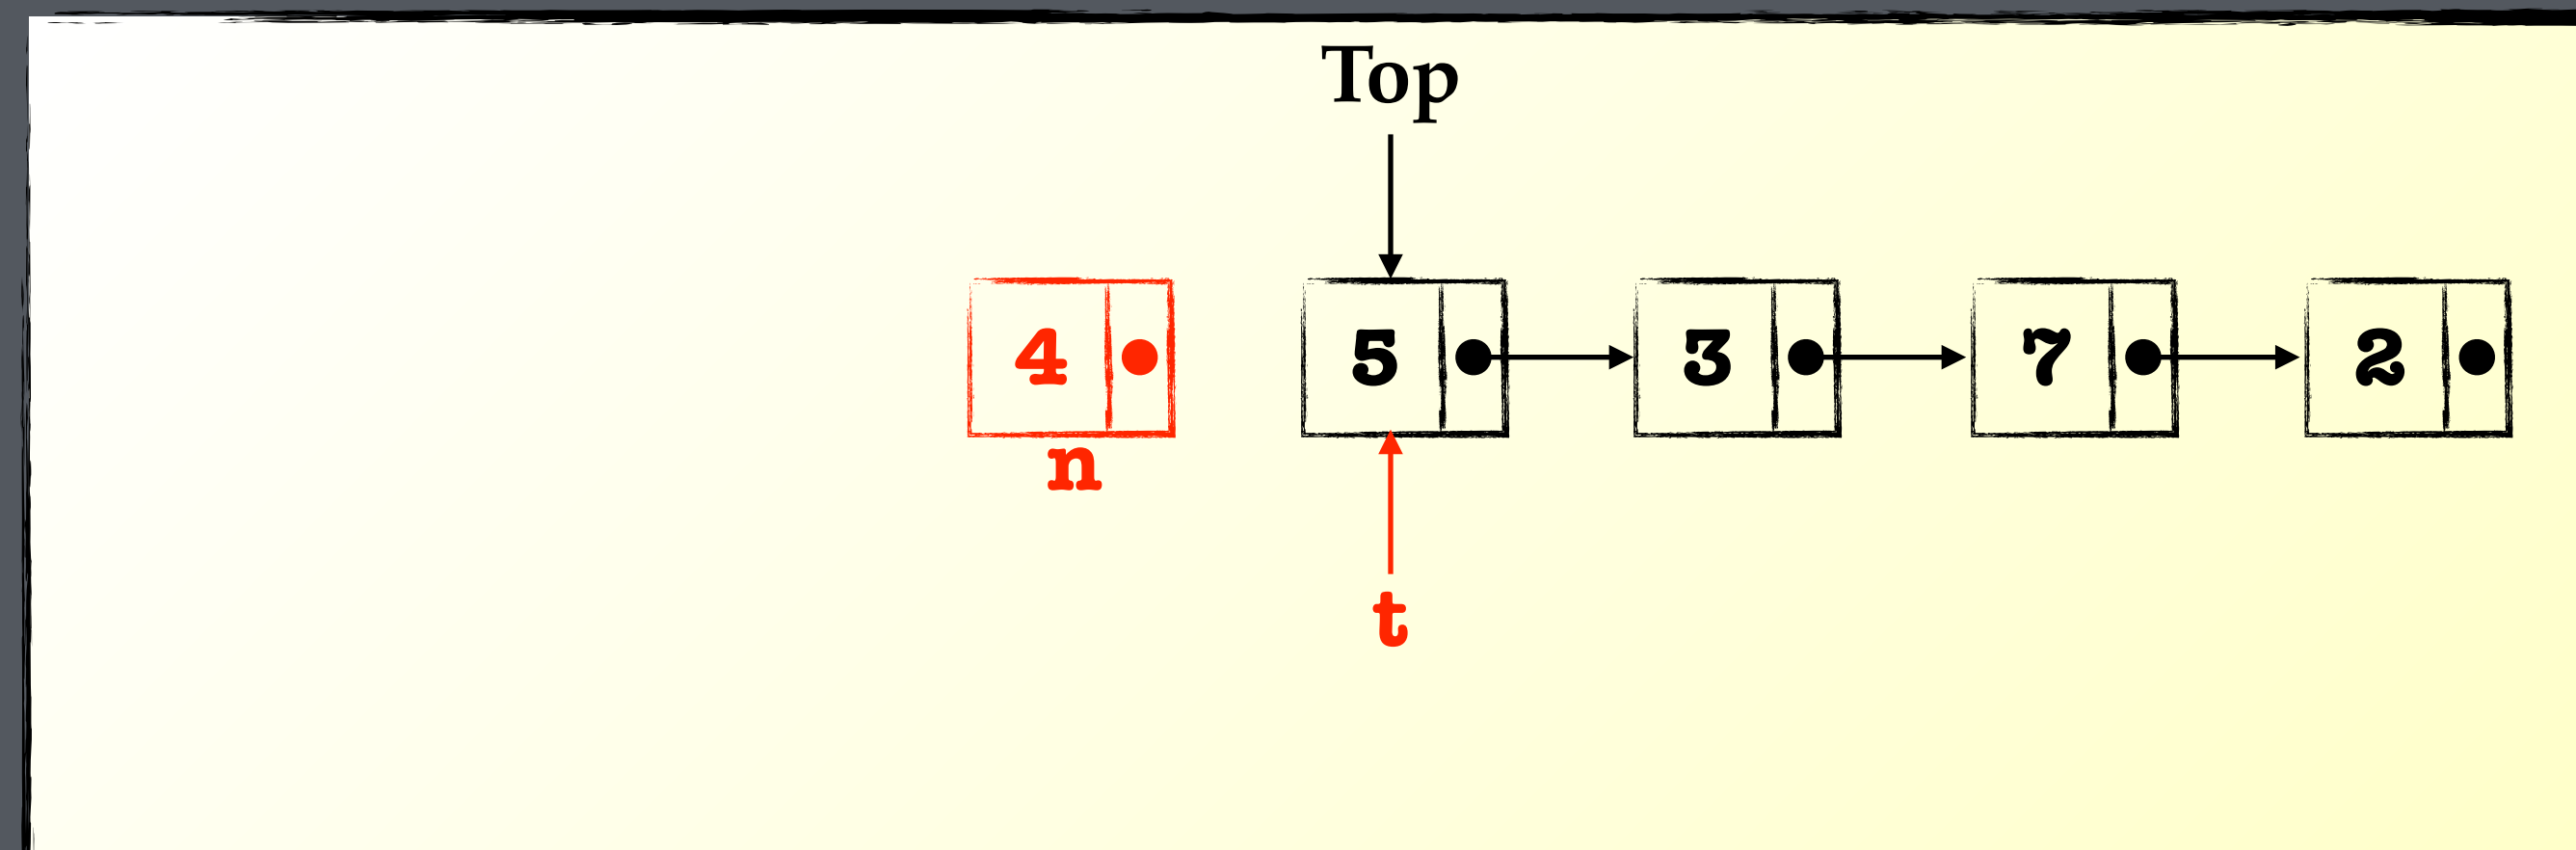


push(4)

push(2)




```
push(k):  
Node t  
1 n = new Node(k,-)  
2 while (true)  
3   t = Top  
4   n.next = t  
5   if (CAS (Top,t,n))  
6     exit
```



# The Treiber Stack Algorithm


Concurrent push operations

```
push(k):  
Node t  
1 n = new Node(k,-)  
2 while (true)  
3   t = Top  
4   n.next = t  
5   if (CAS (Top,t,n))  
6     exit
```

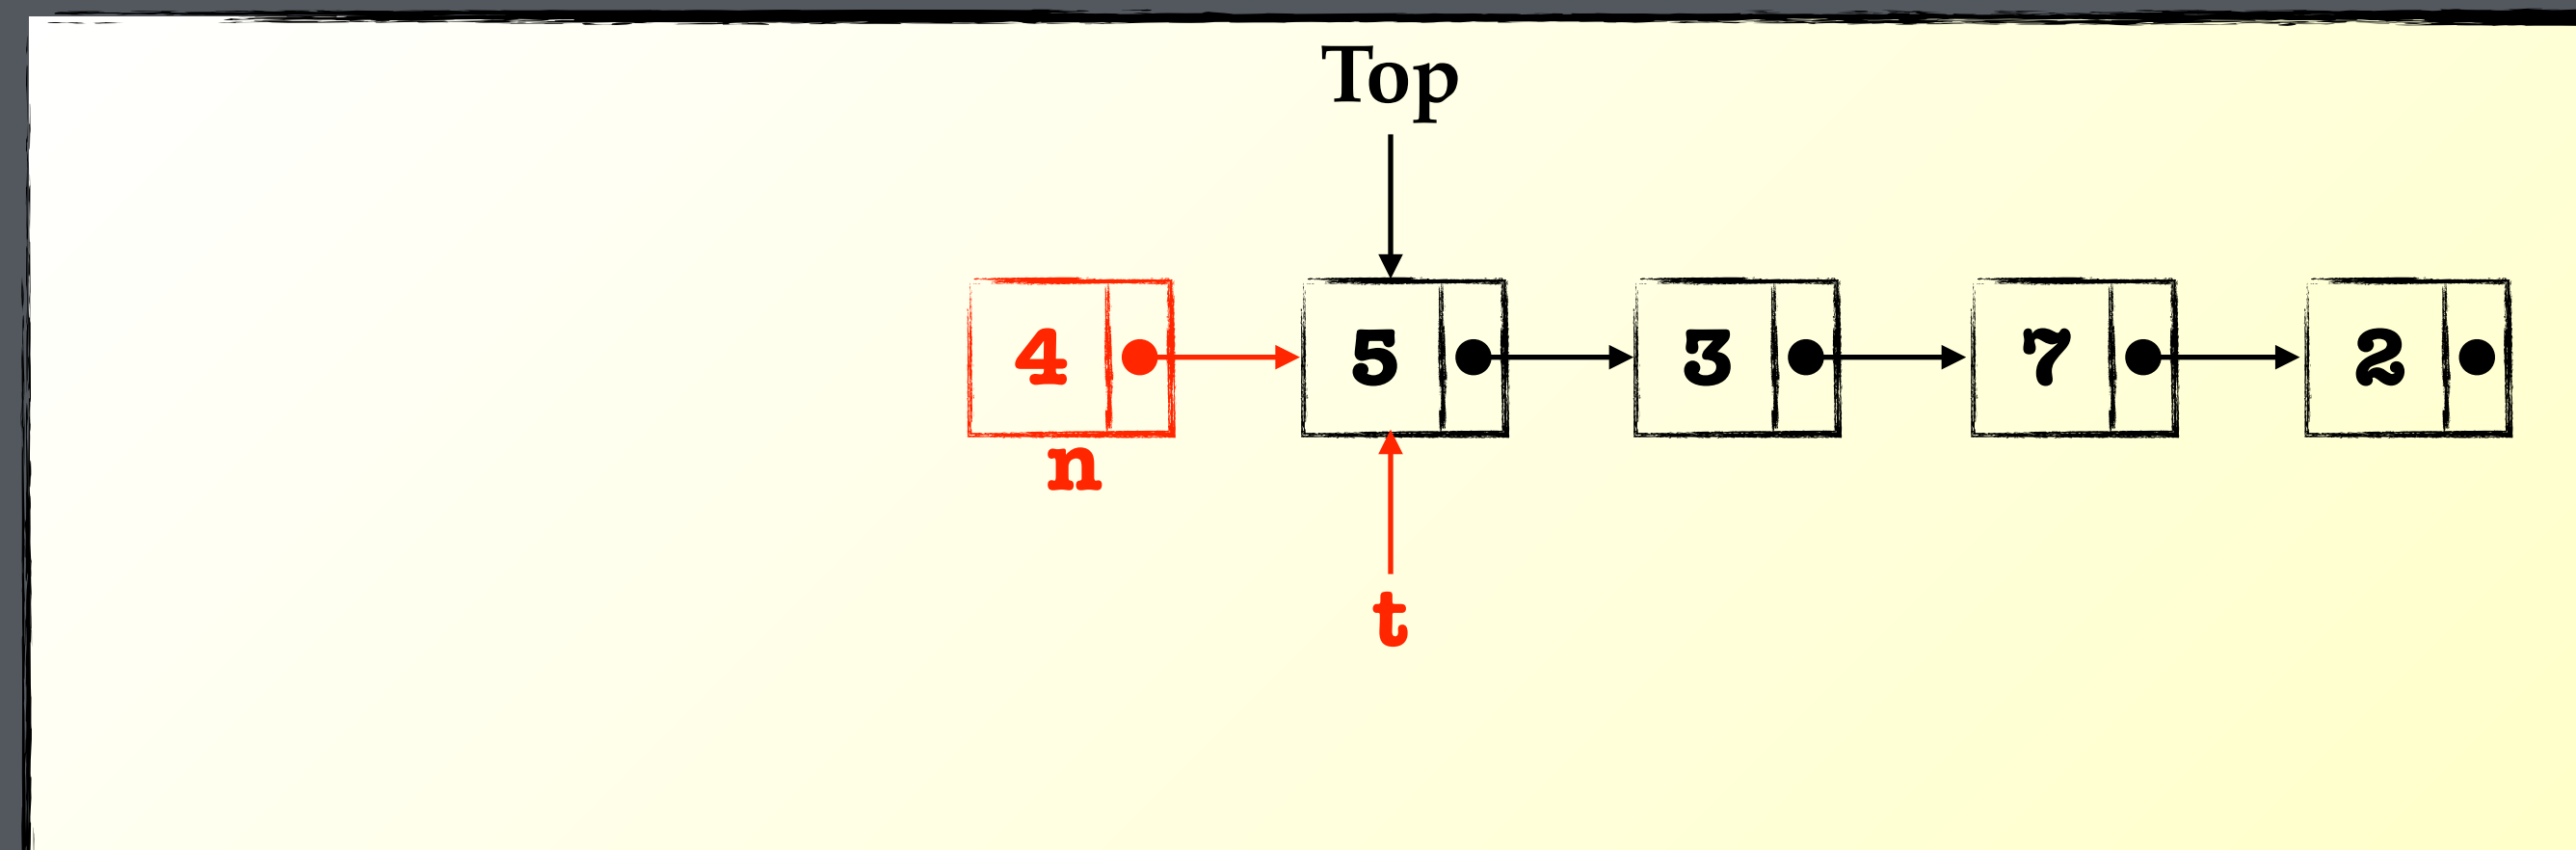


push(4)

push(2)



```
push(k):  
Node t  
1 n = new Node(k,-)  
2 while (true)  
3   t = Top  
4   n.next = t  
5   if (CAS (Top,t,n))  
6     exit
```






# The Treiber Stack Algorithm


Concurrent push operations

```
push(k):  
Node t  
1 n = new Node(k,-)  
2 while (true)  
3   t = Top  
4   n.next = t  
5   if (CAS (Top,t,n))  
6     exit
```




push(4)

push(2)




```
push(k):  
Node t  
1 n = new Node(k,-)  
2 while (true)  
3   t = Top  
4   n.next = t  
5   if (CAS (Top,t,n))  
6     exit
```



# The Treiber Stack Algorithm


Concurrent push operations

```
push(k):  
Node t  
1 n = new Node(k,-)  
2 while (true)  
3   t = Top  
4   n.next = t  
5   if (CAS (Top,t,n))  
6     exit
```

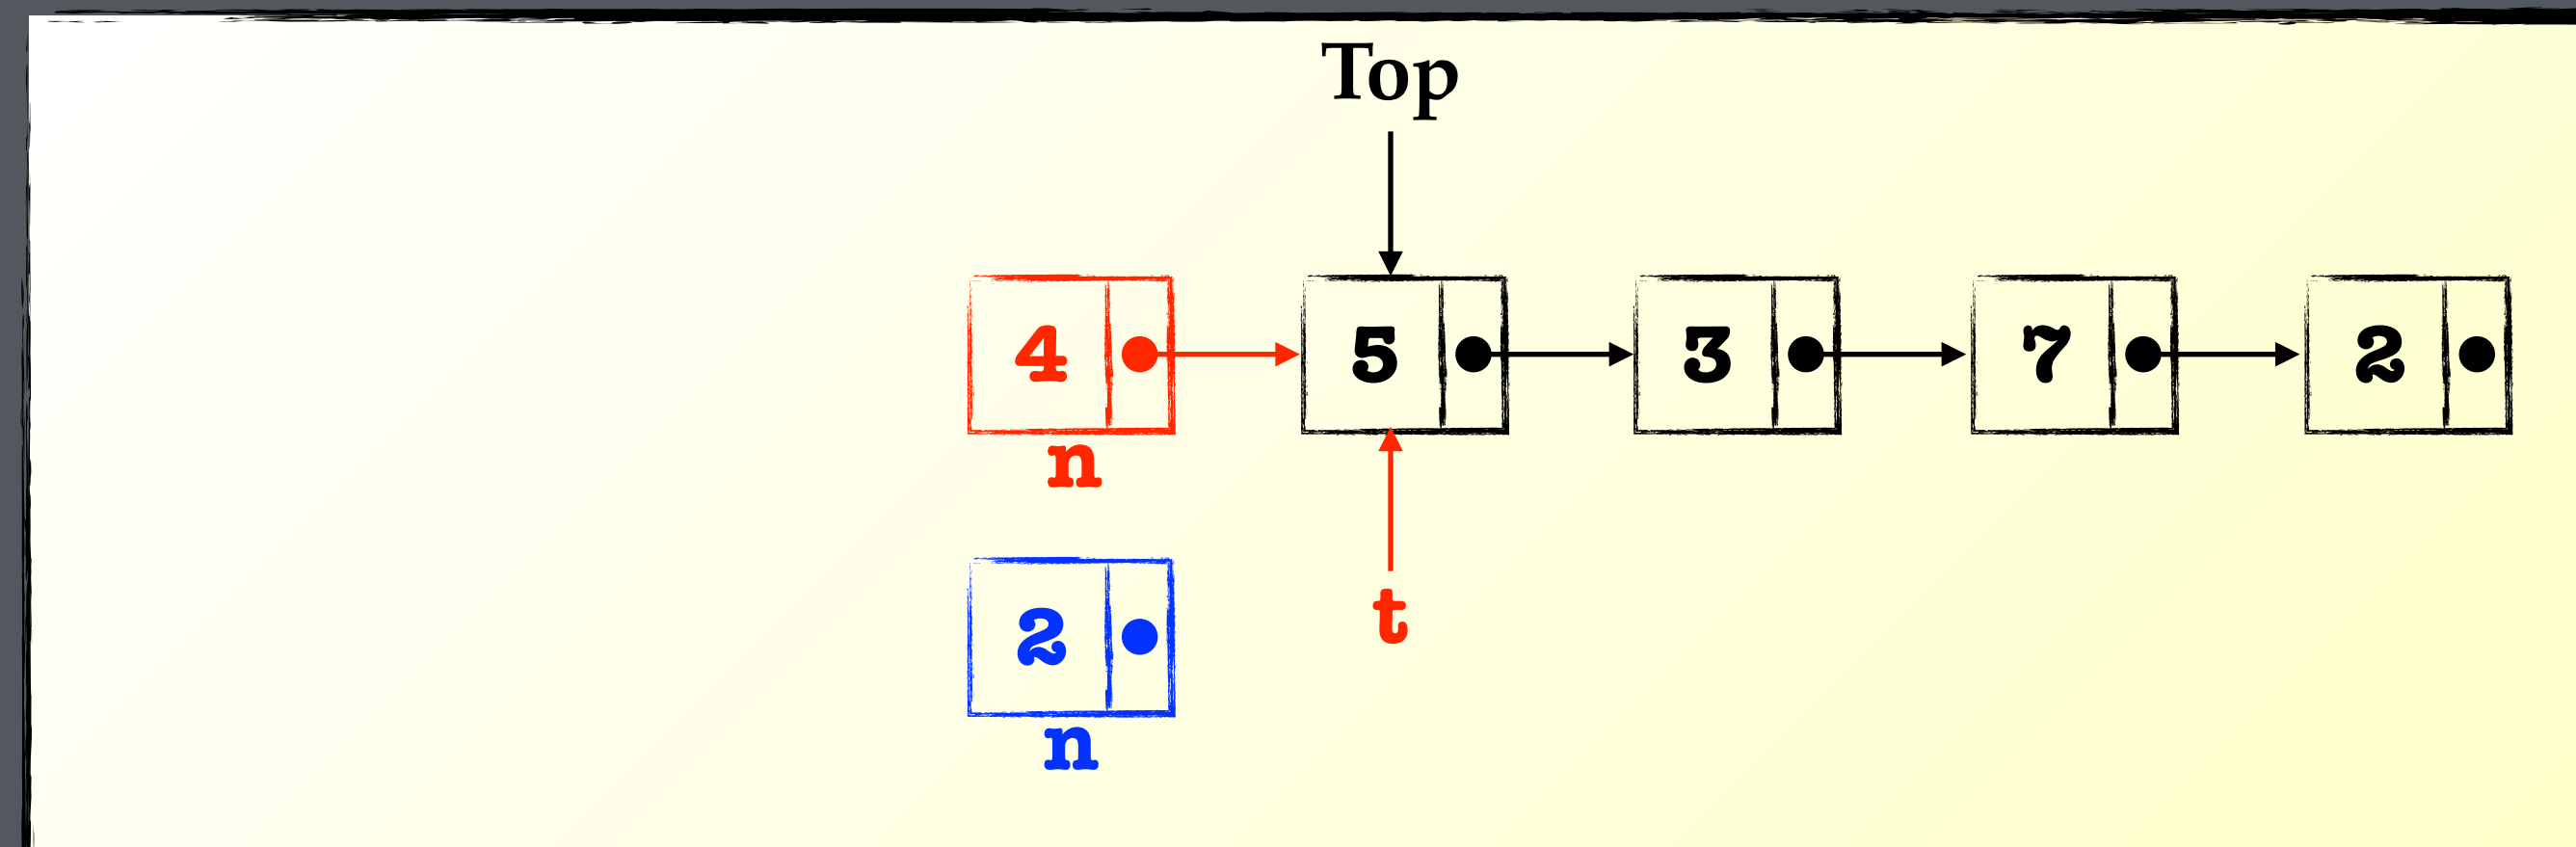


push(4)

push(2)




```
push(k):  
Node t  
1 n = new Node(k,-)  
2 while (true)  
3   t = Top  
4   n.next = t  
5   if (CAS (Top,t,n))  
6     exit
```



# The Treiber Stack Algorithm


Concurrent push operations

```
push(k):  
Node t  
1 n = new Node(k,-)  
2 while (true)  
3   t = Top  
4   n.next = t  
5   if (CAS (Top,t,n))  
6     exit
```

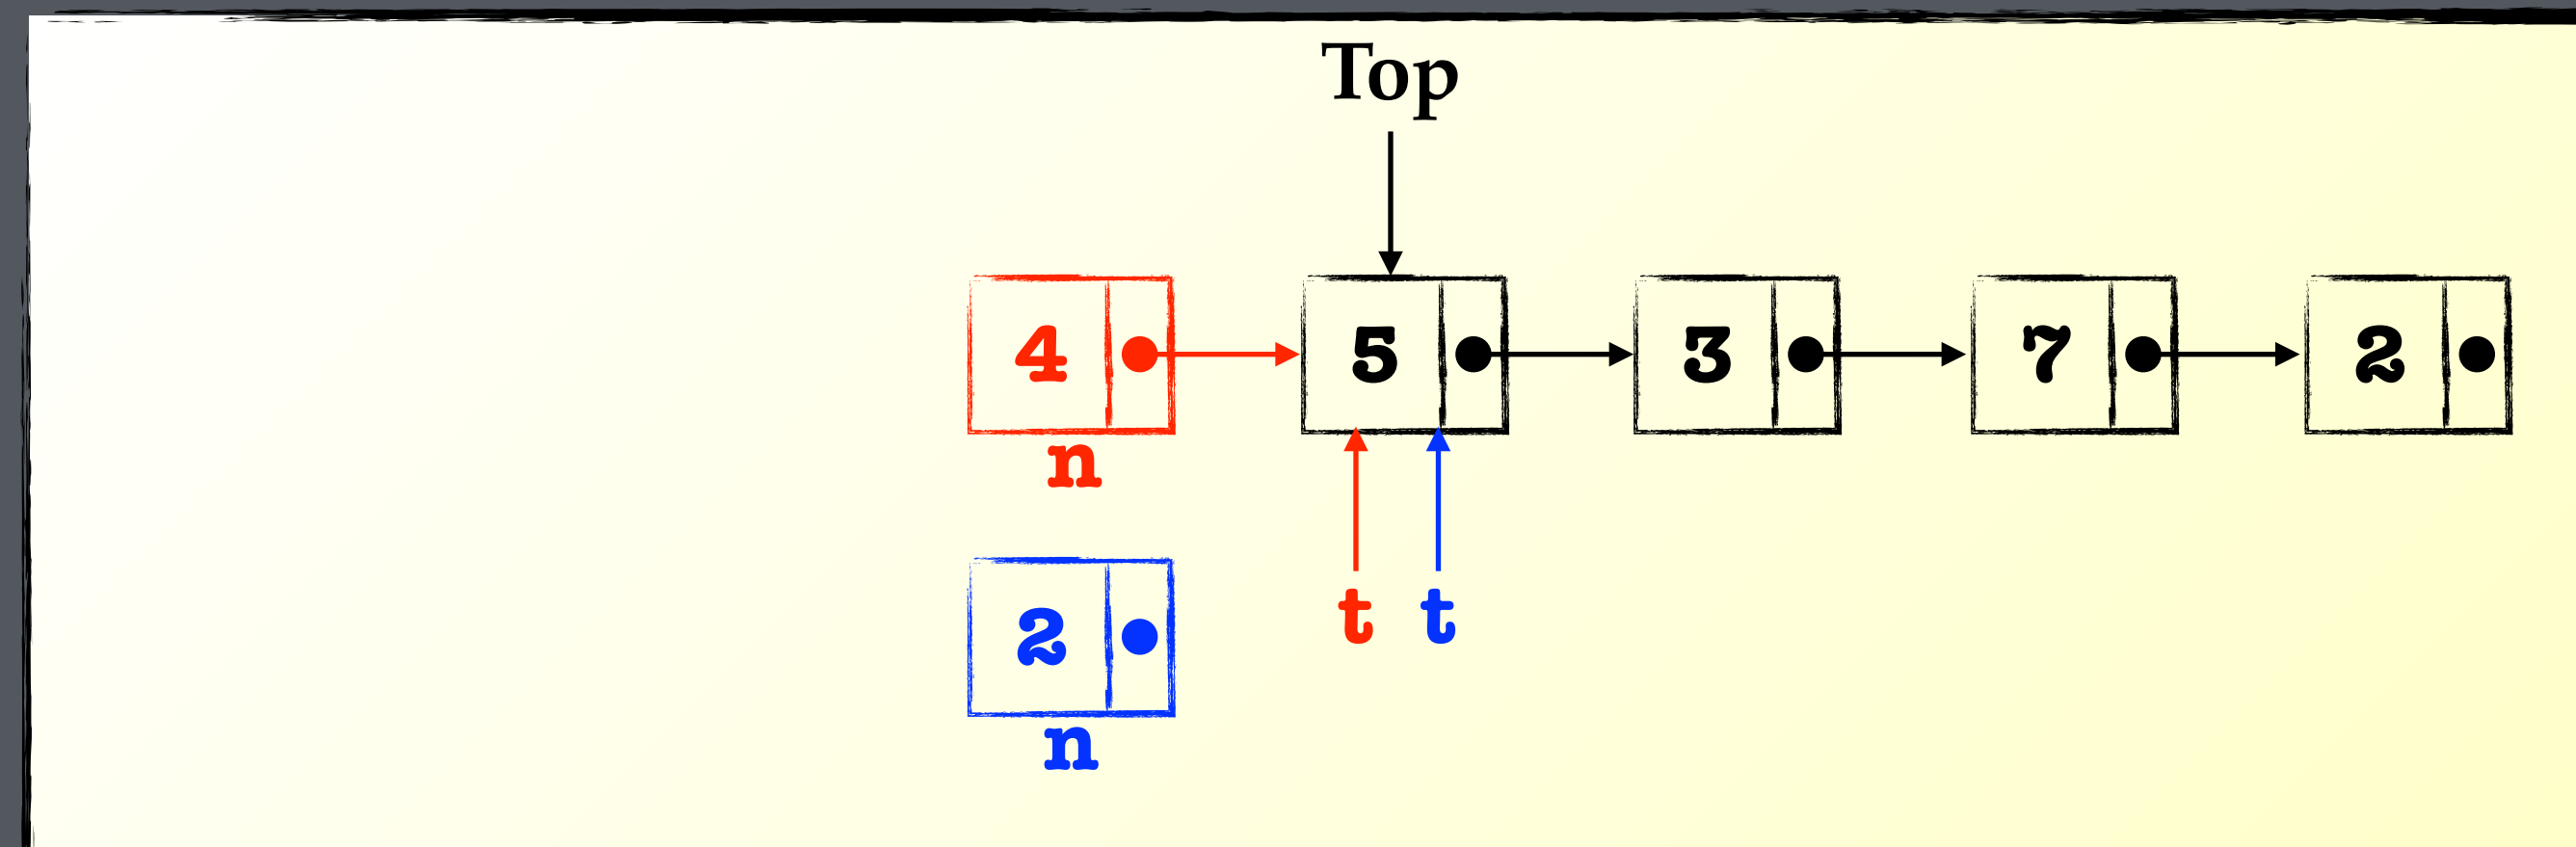


push(4)

push(2)



```
push(k):  
Node t  
1 n = new Node(k,-)  
2 while (true)  
3   t = Top  
4   n.next = t  
5   if (CAS (Top,t,n))  
6     exit
```






# The Treiber Stack Algorithm


Concurrent push operations

```
push(k):  
Node t  
1 n = new Node(k,-)  
2 while (true)  
3   t = Top  
4   n.next = t  
5   if (CAS (Top,t,n))  
6     exit
```

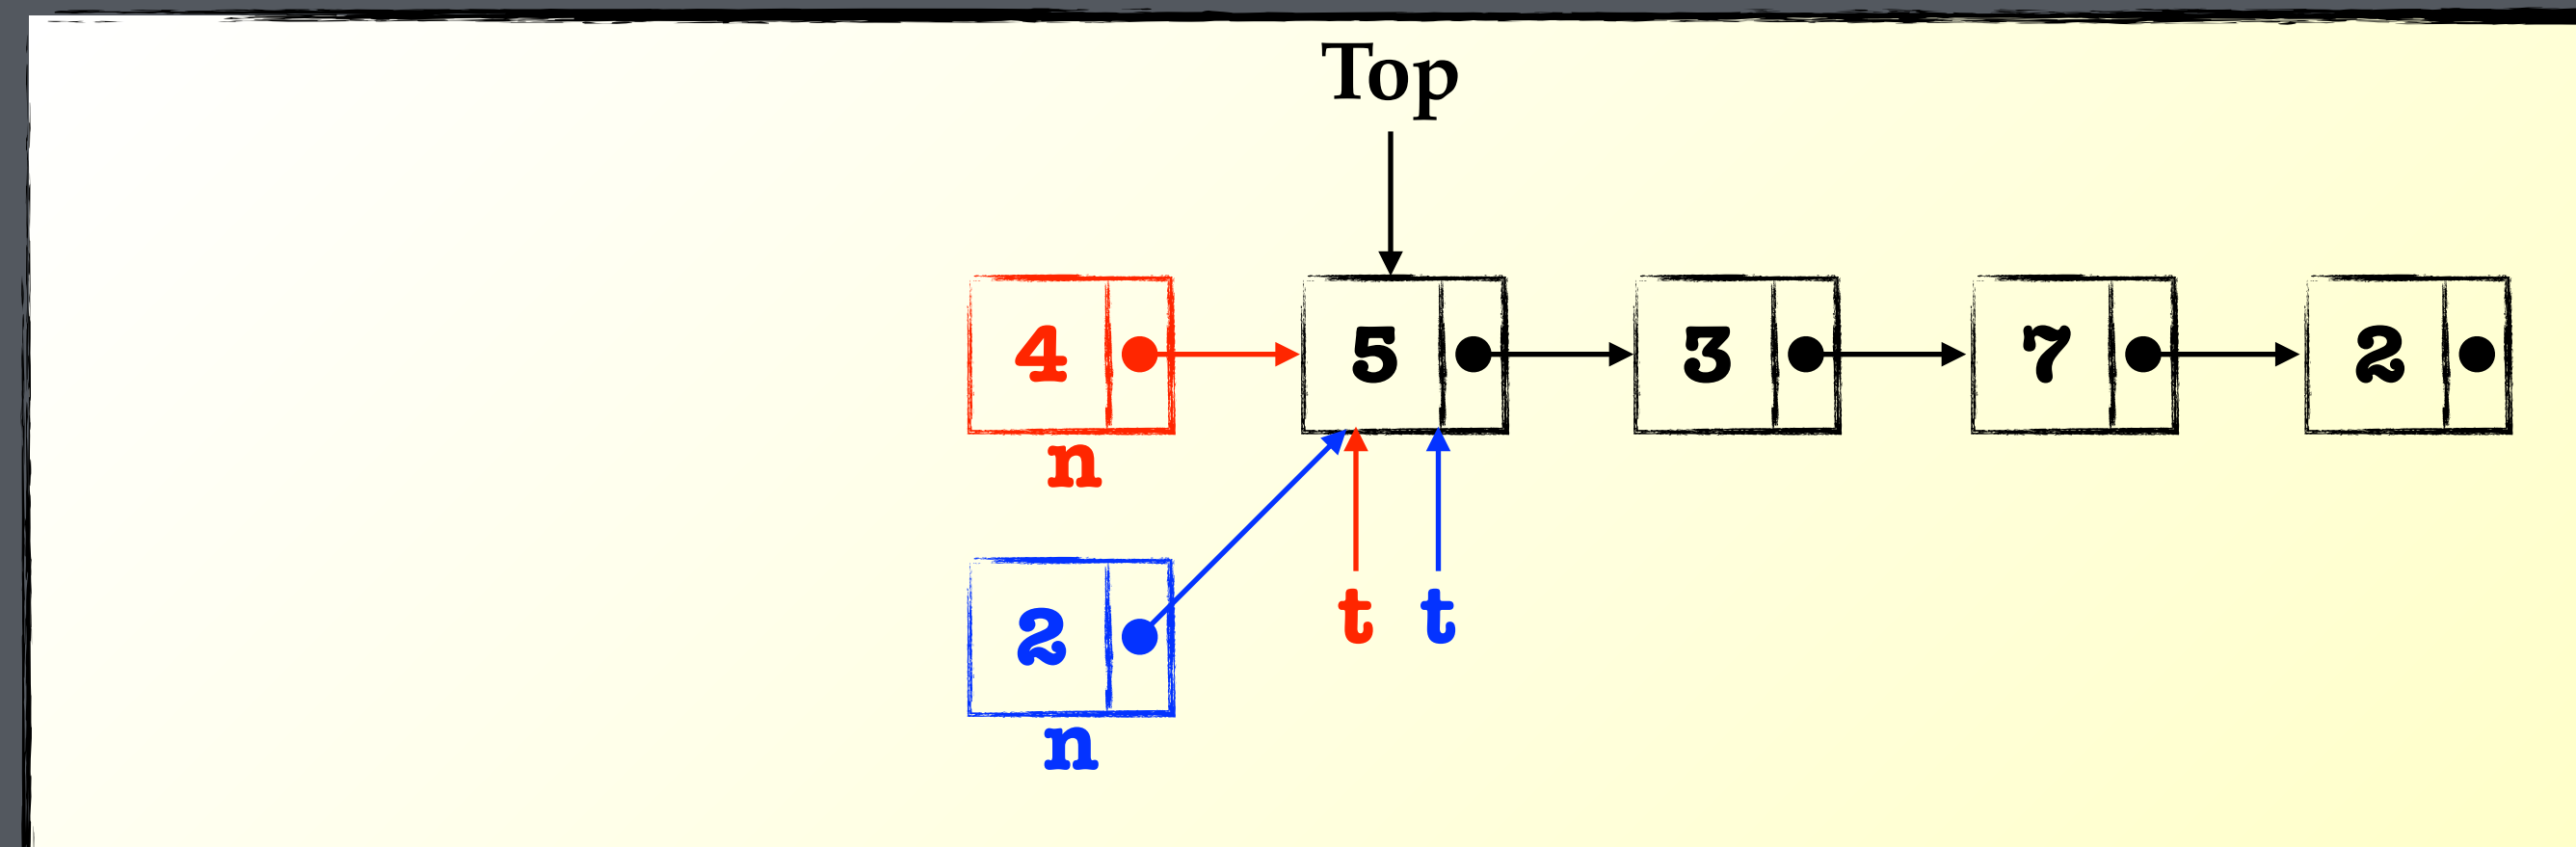


push(4)

push(2)



```
push(k):  
Node t  
1 n = new Node(k,-)  
2 while (true)  
3   t = Top  
4   n.next = t  
5   if (CAS (Top,t,n))  
6     exit
```






# The Treiber Stack Algorithm


Concurrent push operations

```
push(k):  
Node t  
1 n = new Node(k,-)  
2 while (true)  
3   t = Top  
4   n.next = t  
5   if (CAS (Top,t,n))  
6     exit
```

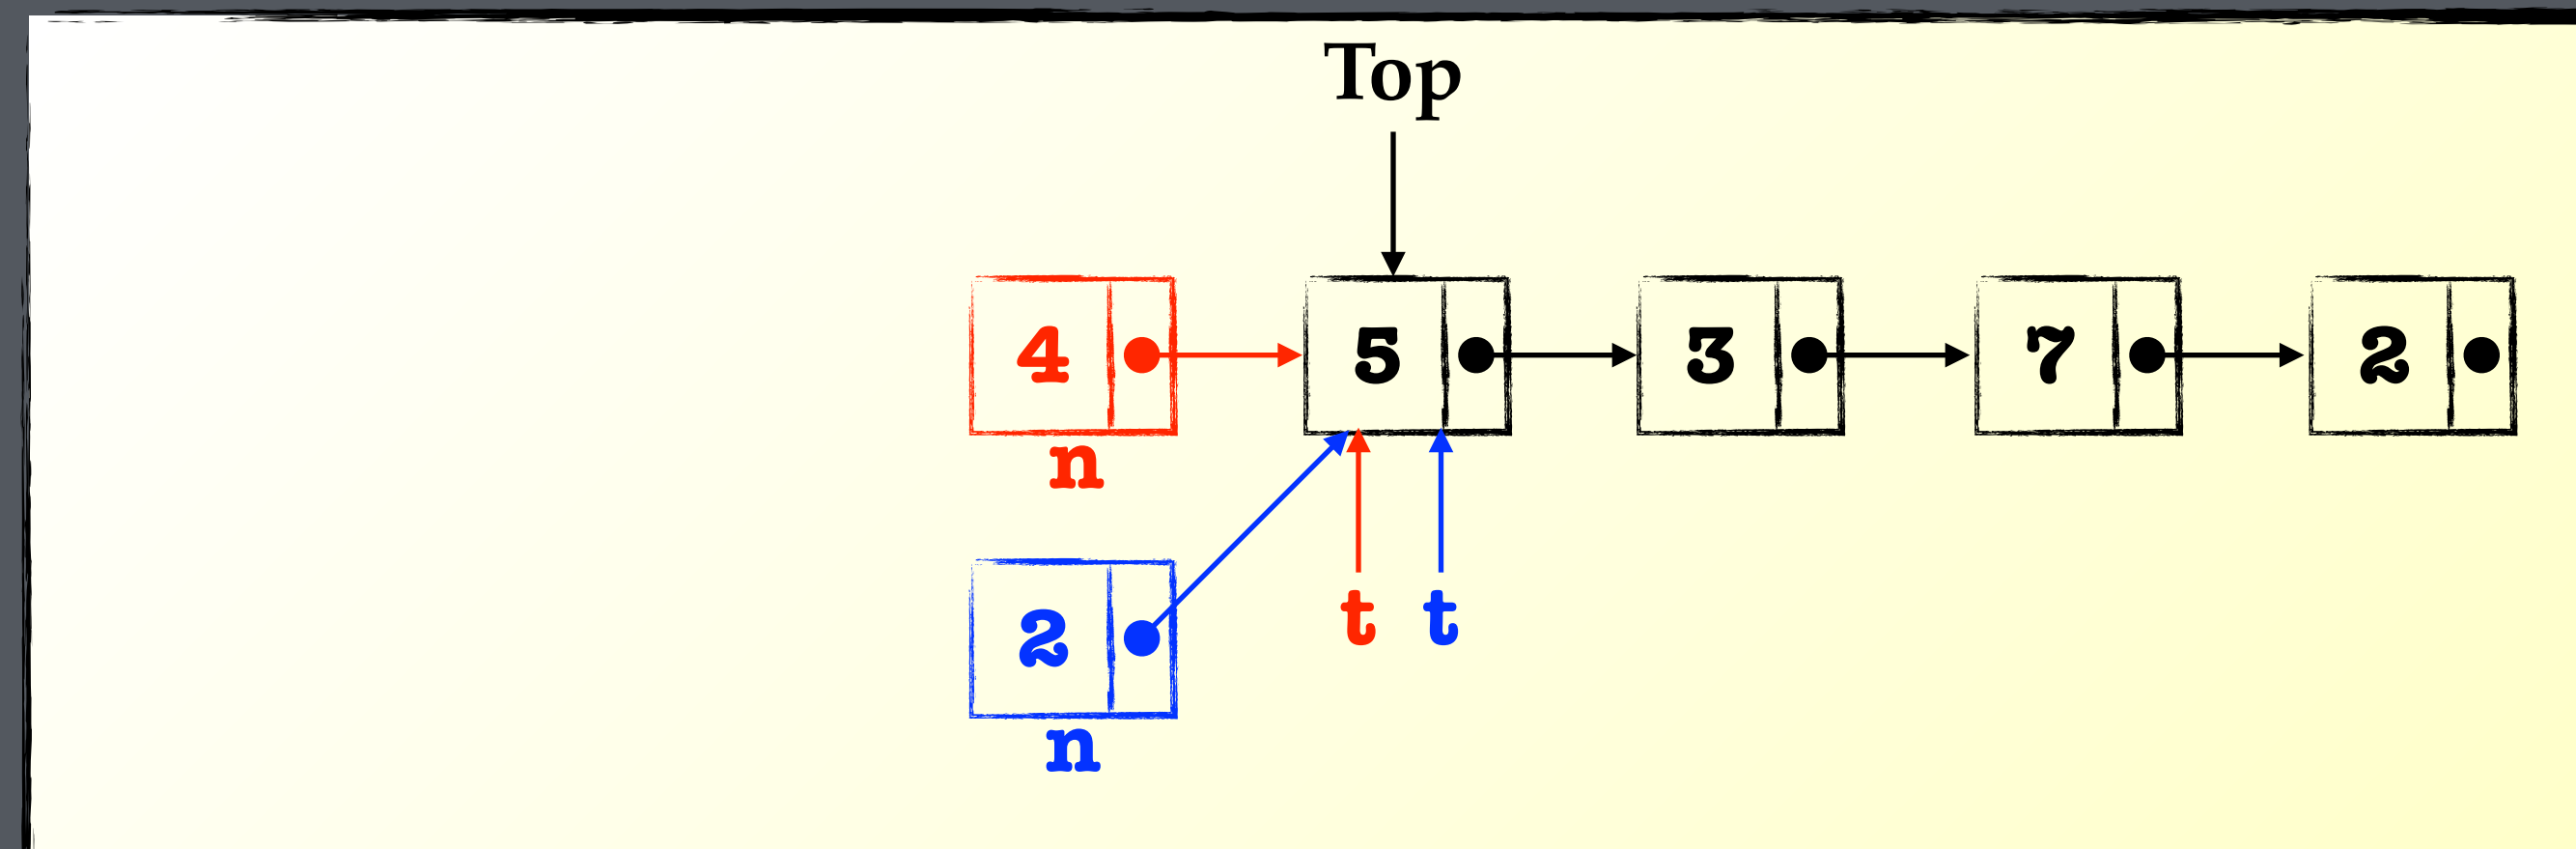


push(4)

push(2)




```
push(k):  
Node t  
1 n = new Node(k,-)  
2 while (true)  
3   t = Top  
4   n.next = t  
5   if (CAS (Top,t,n))  
6     exit
```



# The Treiber Stack Algorithm


Concurrent push operations

```
push(k):  
Node t  
1 n = new Node(k,-)  
2 while (true)  
3   t = Top  
4   n.next = t  
5   if (CAS (Top,t,n))  
6     exit
```




push(4)

push(2)




```
push(k):  
Node t  
1 n = new Node(k,-)  
2 while (true)  
3   t = Top  
4   n.next = t  
5   if (CAS (Top,t,n))  
6     exit
```



# The Treiber Stack Algorithm


Concurrent push operations

```
push(k):  
Node t  
1 n = new Node(k,-)  
2 while (true)  
3   t = Top  
4   n.next = t  
5   if (CAS (Top,t,n))  
6     exit
```




push(4)

push(2)



```
push(k):  
Node t  
1 n = new Node(k,-)  
2 while (true)  
3   t = Top  
4   n.next = t  
5   if (CAS (Top,t,n))  
6     exit
```





# The Treiber Stack Algorithm

Concurrent push operations

push(k):

Node t

1 n = new Node(k,-)


2 while (true)

3 t = Top

4 n.next = t


5 if (CAS (Top,t,n)) ←

6 exit



push(4)

push(2)



push(k):

Node t

1 n = new Node(k,-)


2 while (true)

3 t = Top

4 n.next = t

5 if (CAS (Top,t,n))


6 exit ←



# The Treiber Stack Algorithm


Concurrent push operations

```
push(k):  
Node t  
1 n = new Node(k,-)  
2 while (true) ←~~~~~  
3   t = Top  
4   n.next = t  
5   if (CAS (Top,t,n))  
6     exit
```




push(4)

push(2)




```
push(k):  
Node t  
1 n = new Node(k,-)  
2 while (true)  
3   t = Top  
4   n.next = t  
5   if (CAS (Top,t,n))  
6     exit ←~~~~~
```



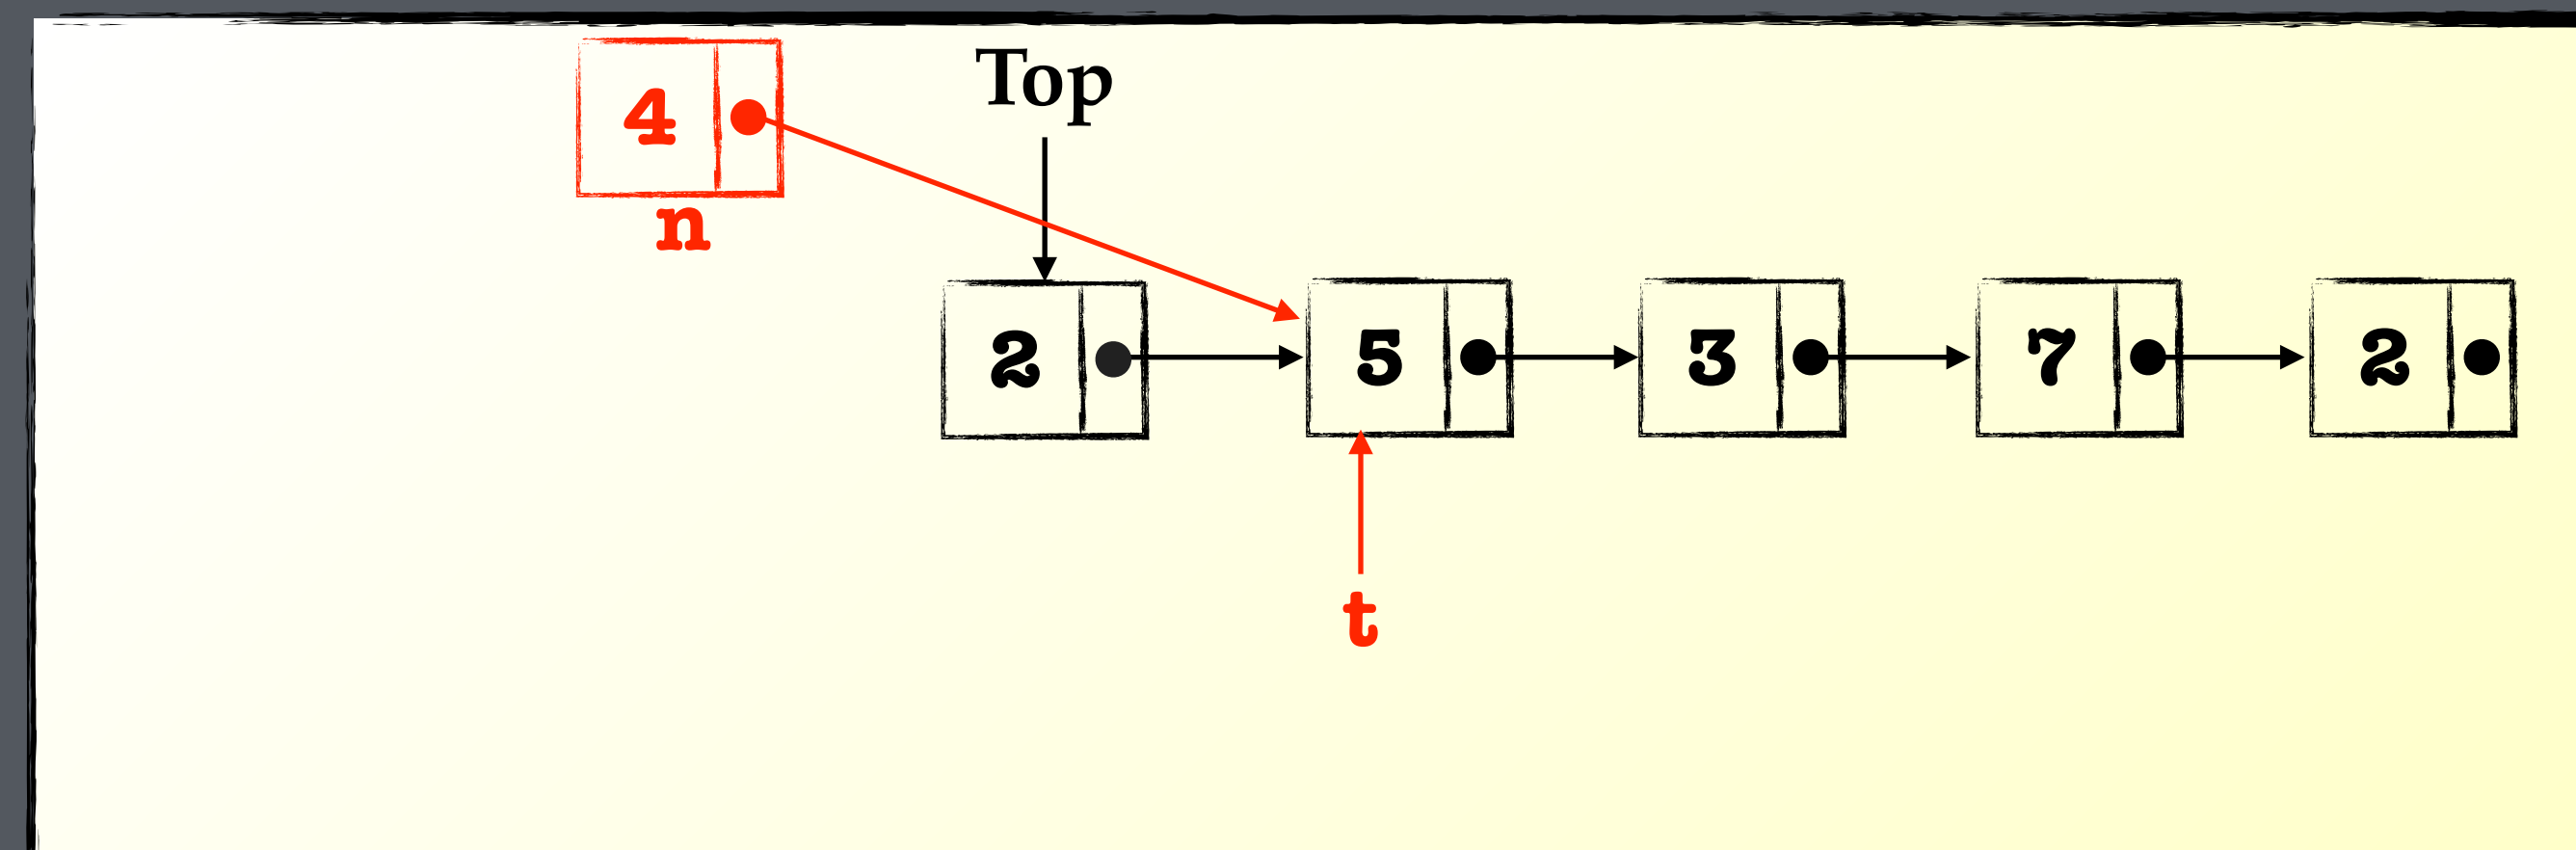
# The Treiber Stack Algorithm

Concurrent push operations

```
push(k):  
Node t  
1  n = new Node(k,-)  
2  while (true) ← wavy  
3    t = Top  
4    n.next = t  
5    if (CAS (Top,t,n))  
6      exit
```




push(4)



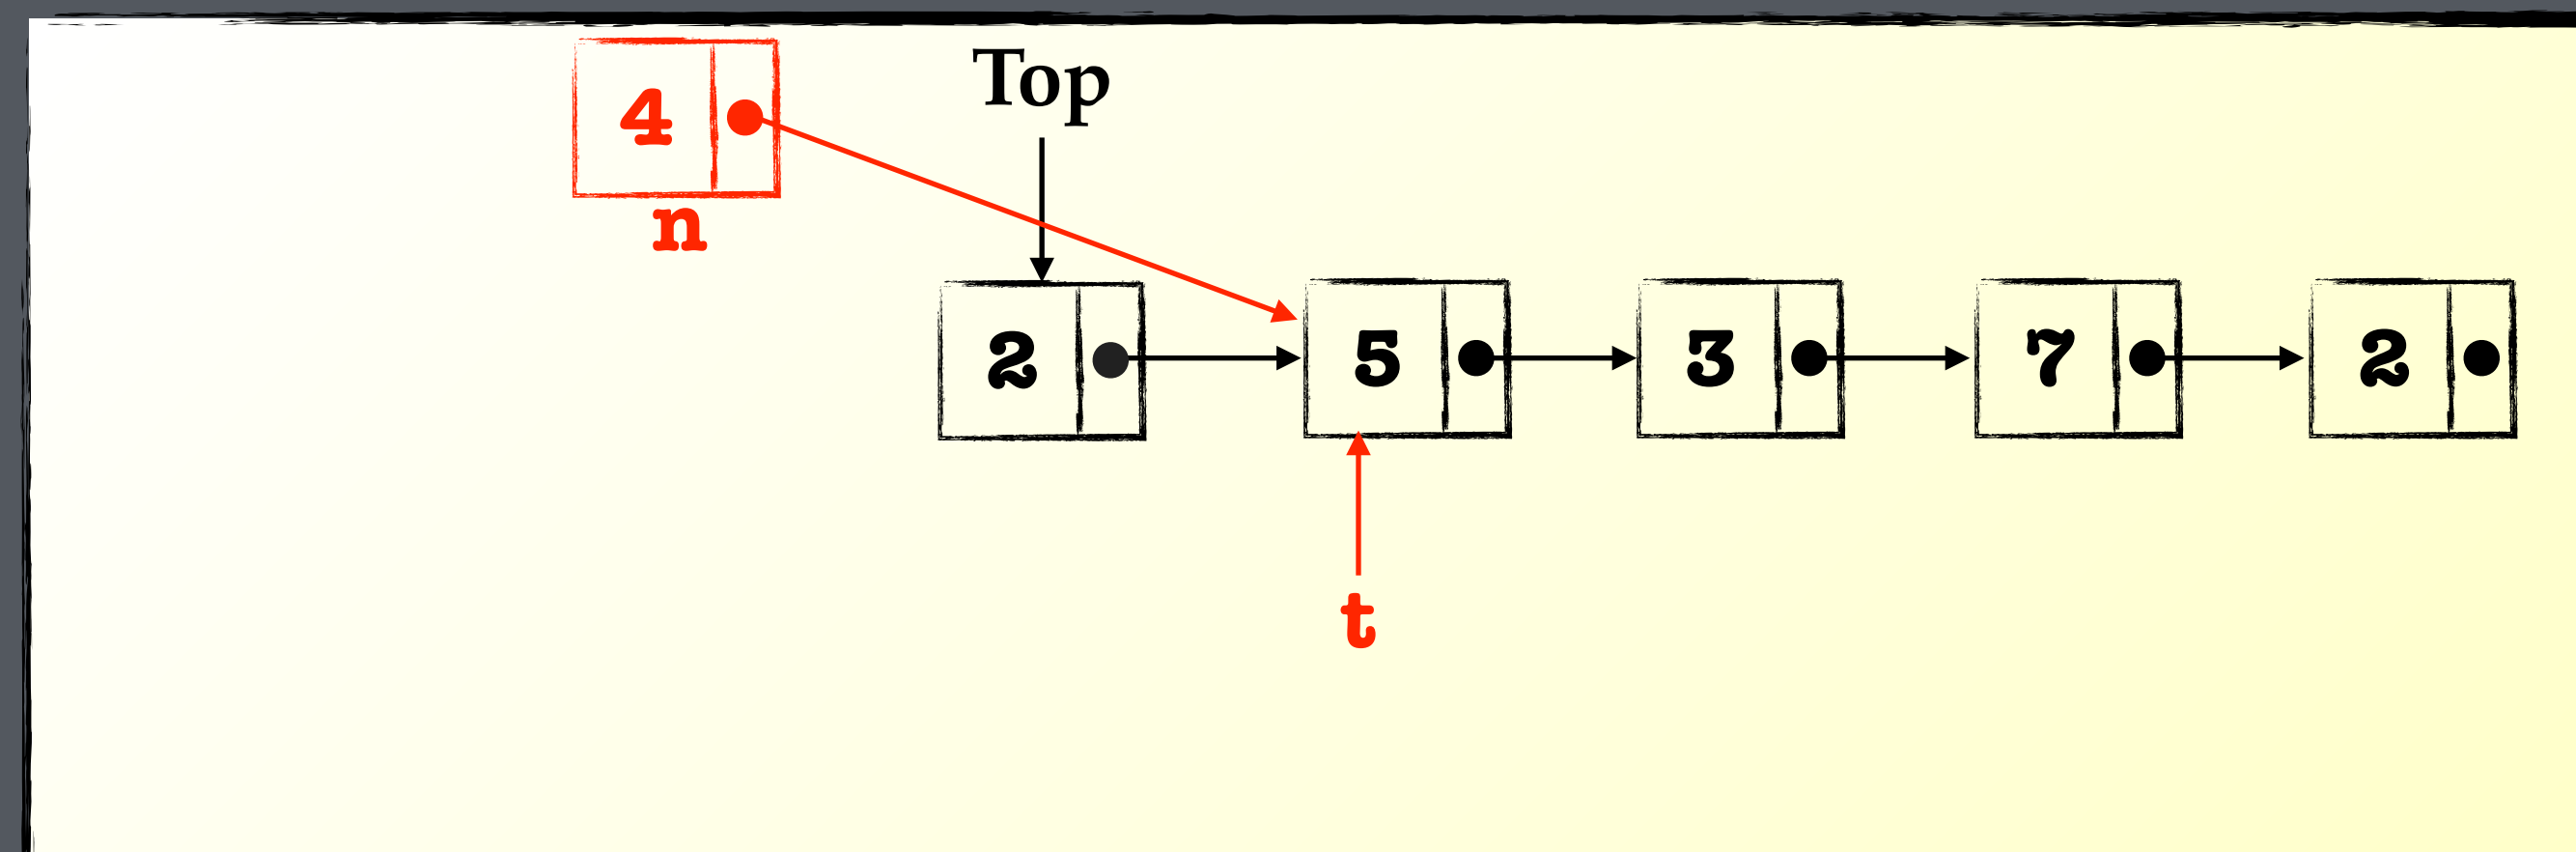
# The Treiber Stack Algorithm

Concurrent push operations

```
push(k):  
Node t  
1  n = new Node(k,-)  
2  while (true)  
3    t = Top ←~~~~~  
4    n.next = t  
5    if (CAS (Top,t,n))  
6      exit
```



push(4)






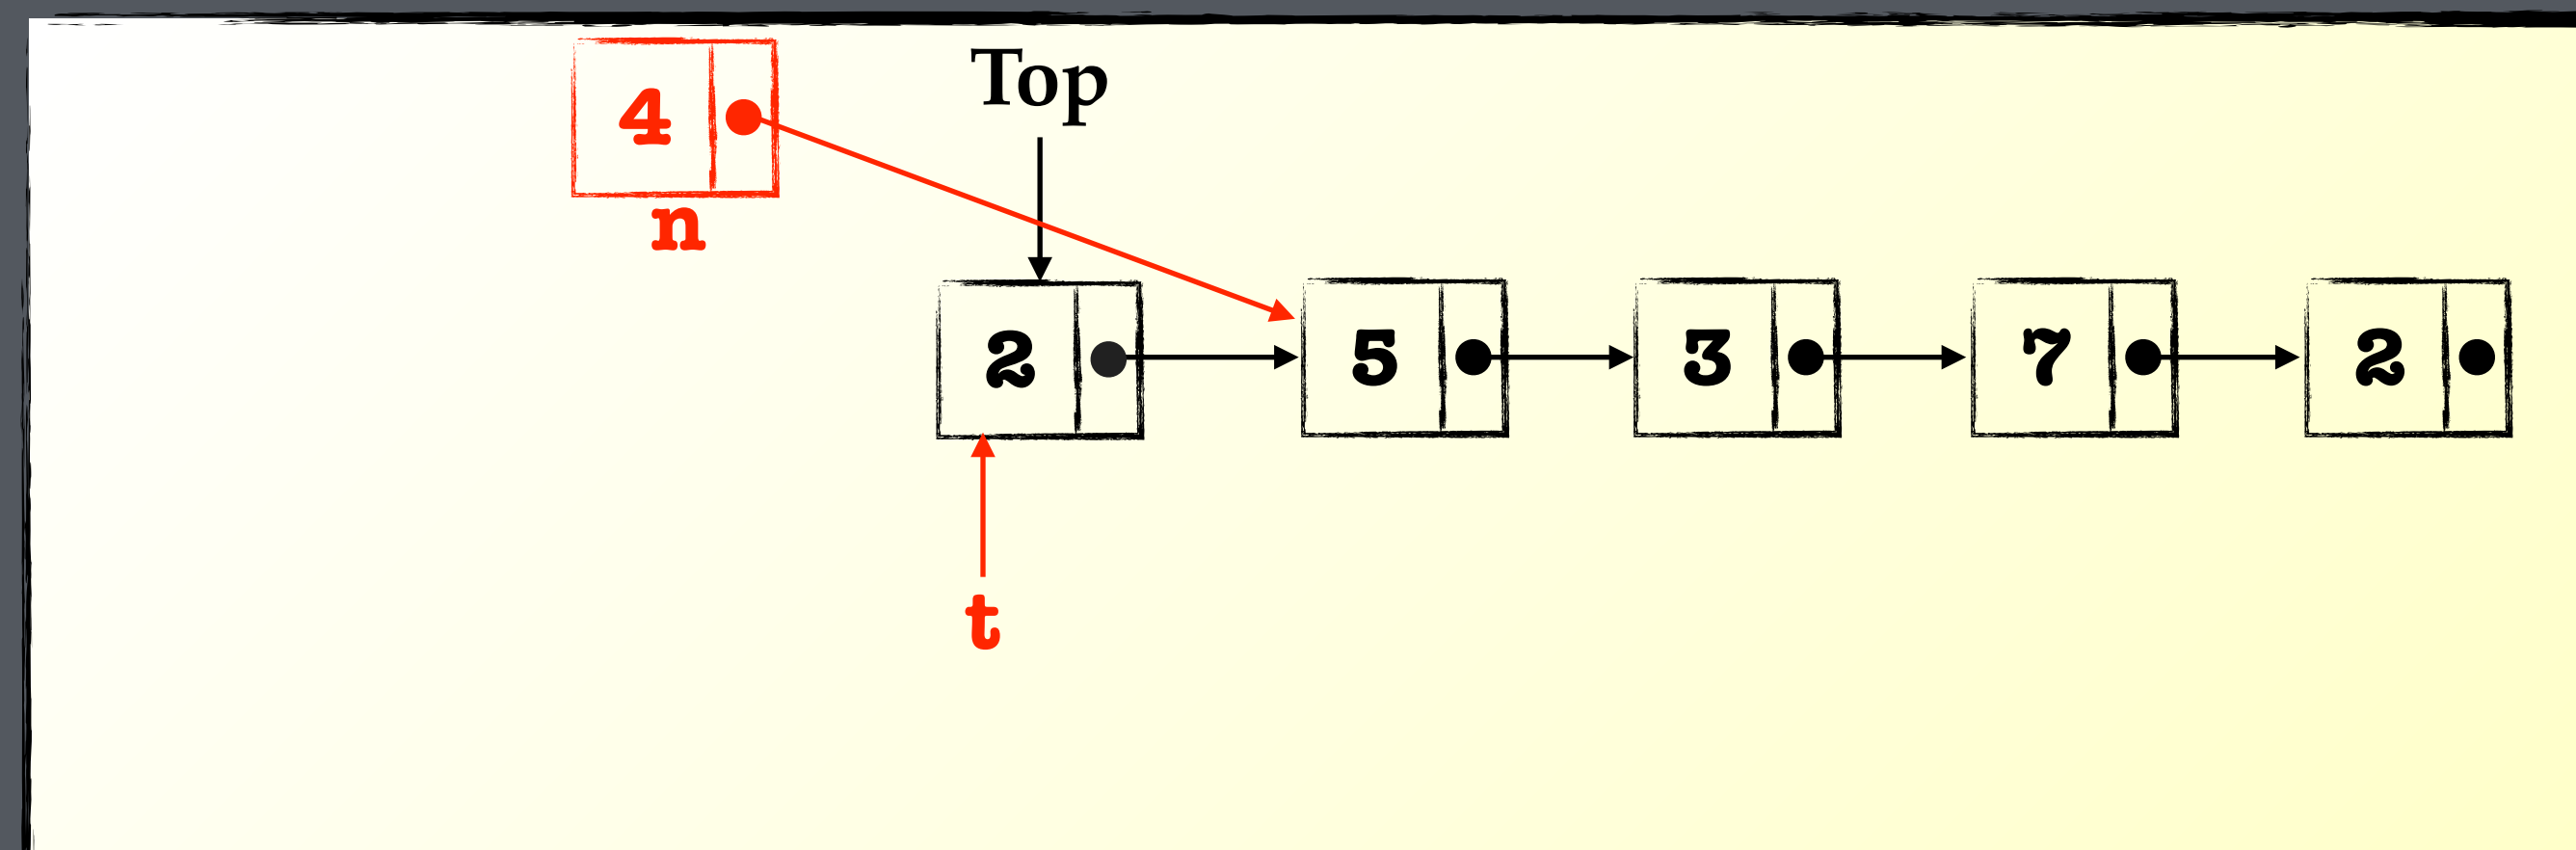
# The Treiber Stack Algorithm

Concurrent push operations

```
push(k):  
Node t  
1  n = new Node(k,-)  
2  while (true)  
3    t = Top ←~~~~~  
4    n.next = t  
5    if (CAS (Top,t,n))  
6      exit
```



push(4)






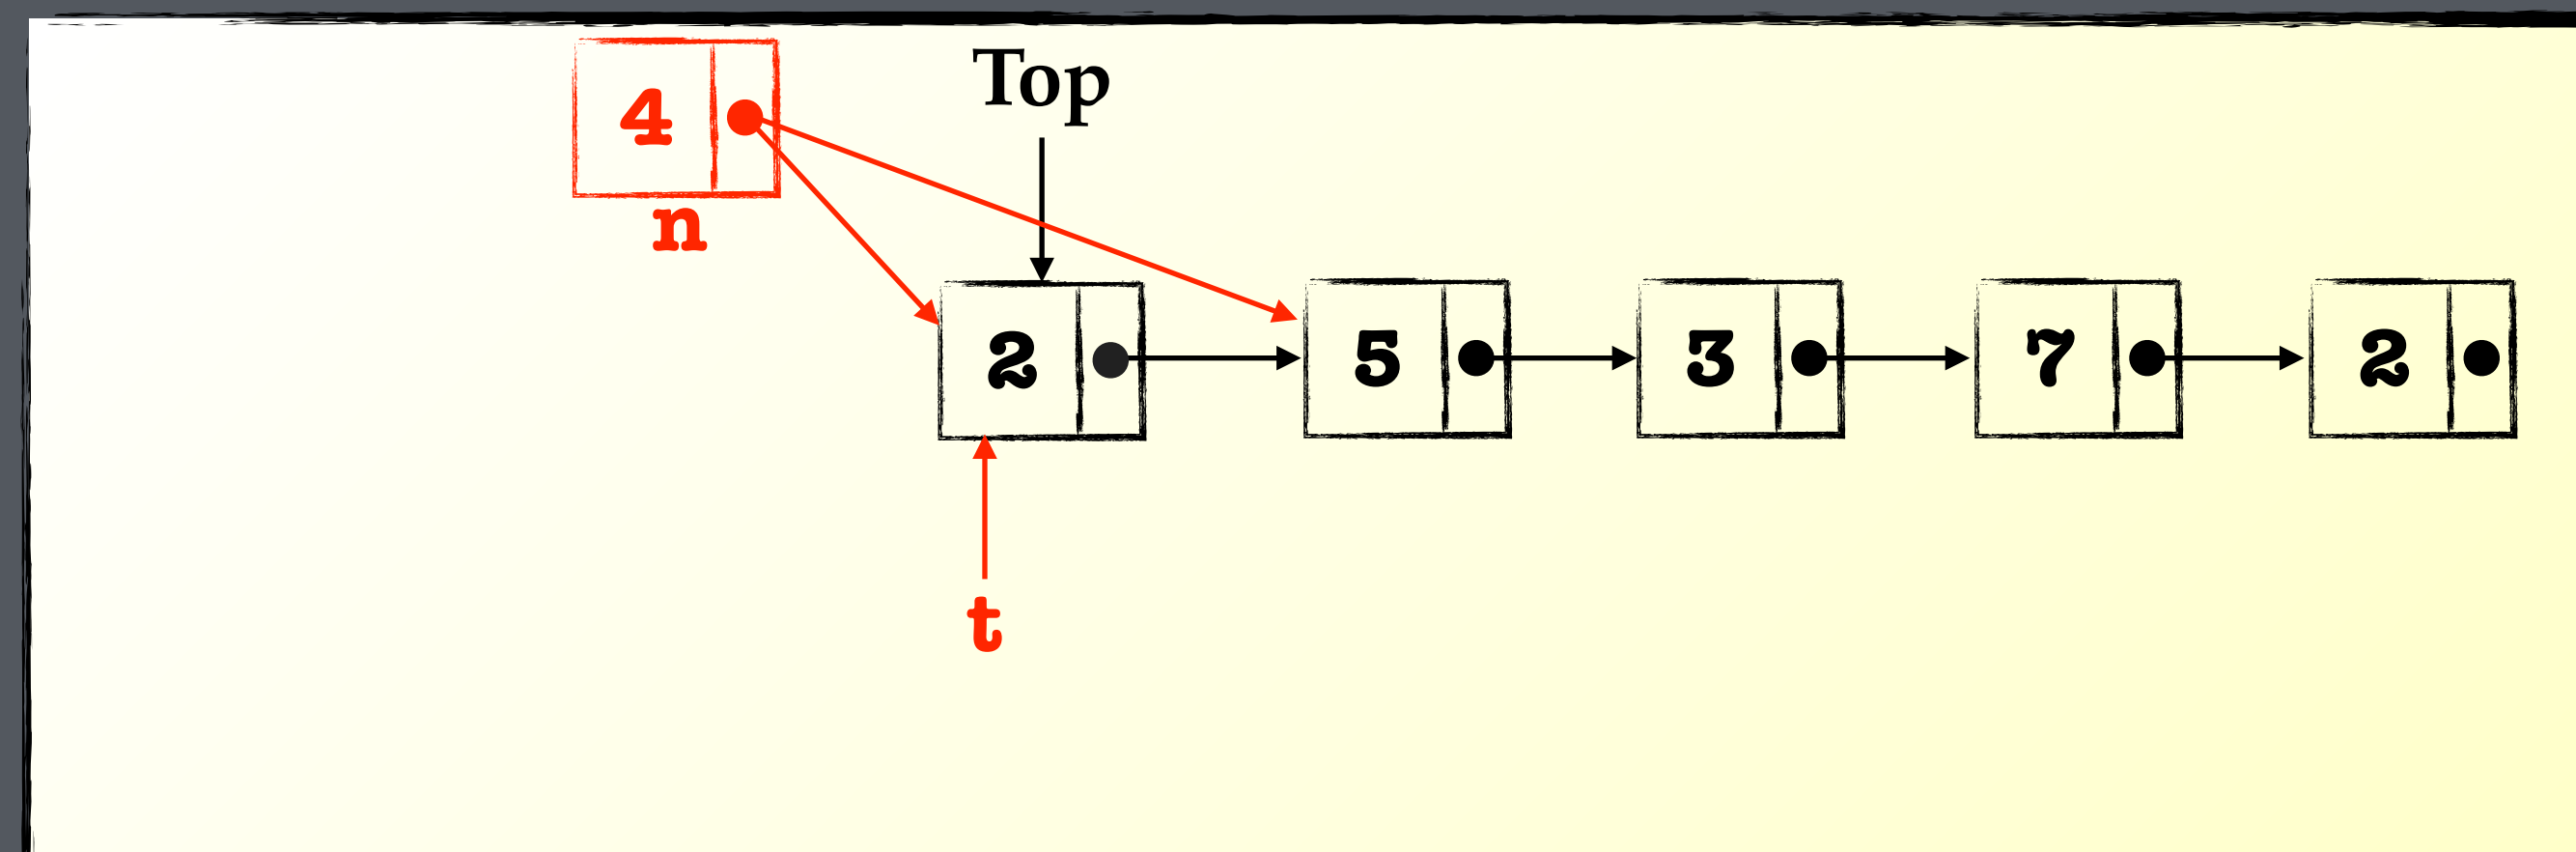
# The Treiber Stack Algorithm

Concurrent push operations

```
push(k):  
Node t  
1  n = new Node(k,-)  
2  while (true)  
3    t = Top  
4    n.next = t  
5    if (CAS (Top,t,n))  
6      exit
```




push(4)



# The Treiber Stack Algorithm


Concurrent push operations

```
push(k):  
Node t  
1 n = new Node(k,-)  
2 while (true)  
3   t = Top  
4   n.next = t  
5   if (CAS (Top,t,n))  
6     exit
```

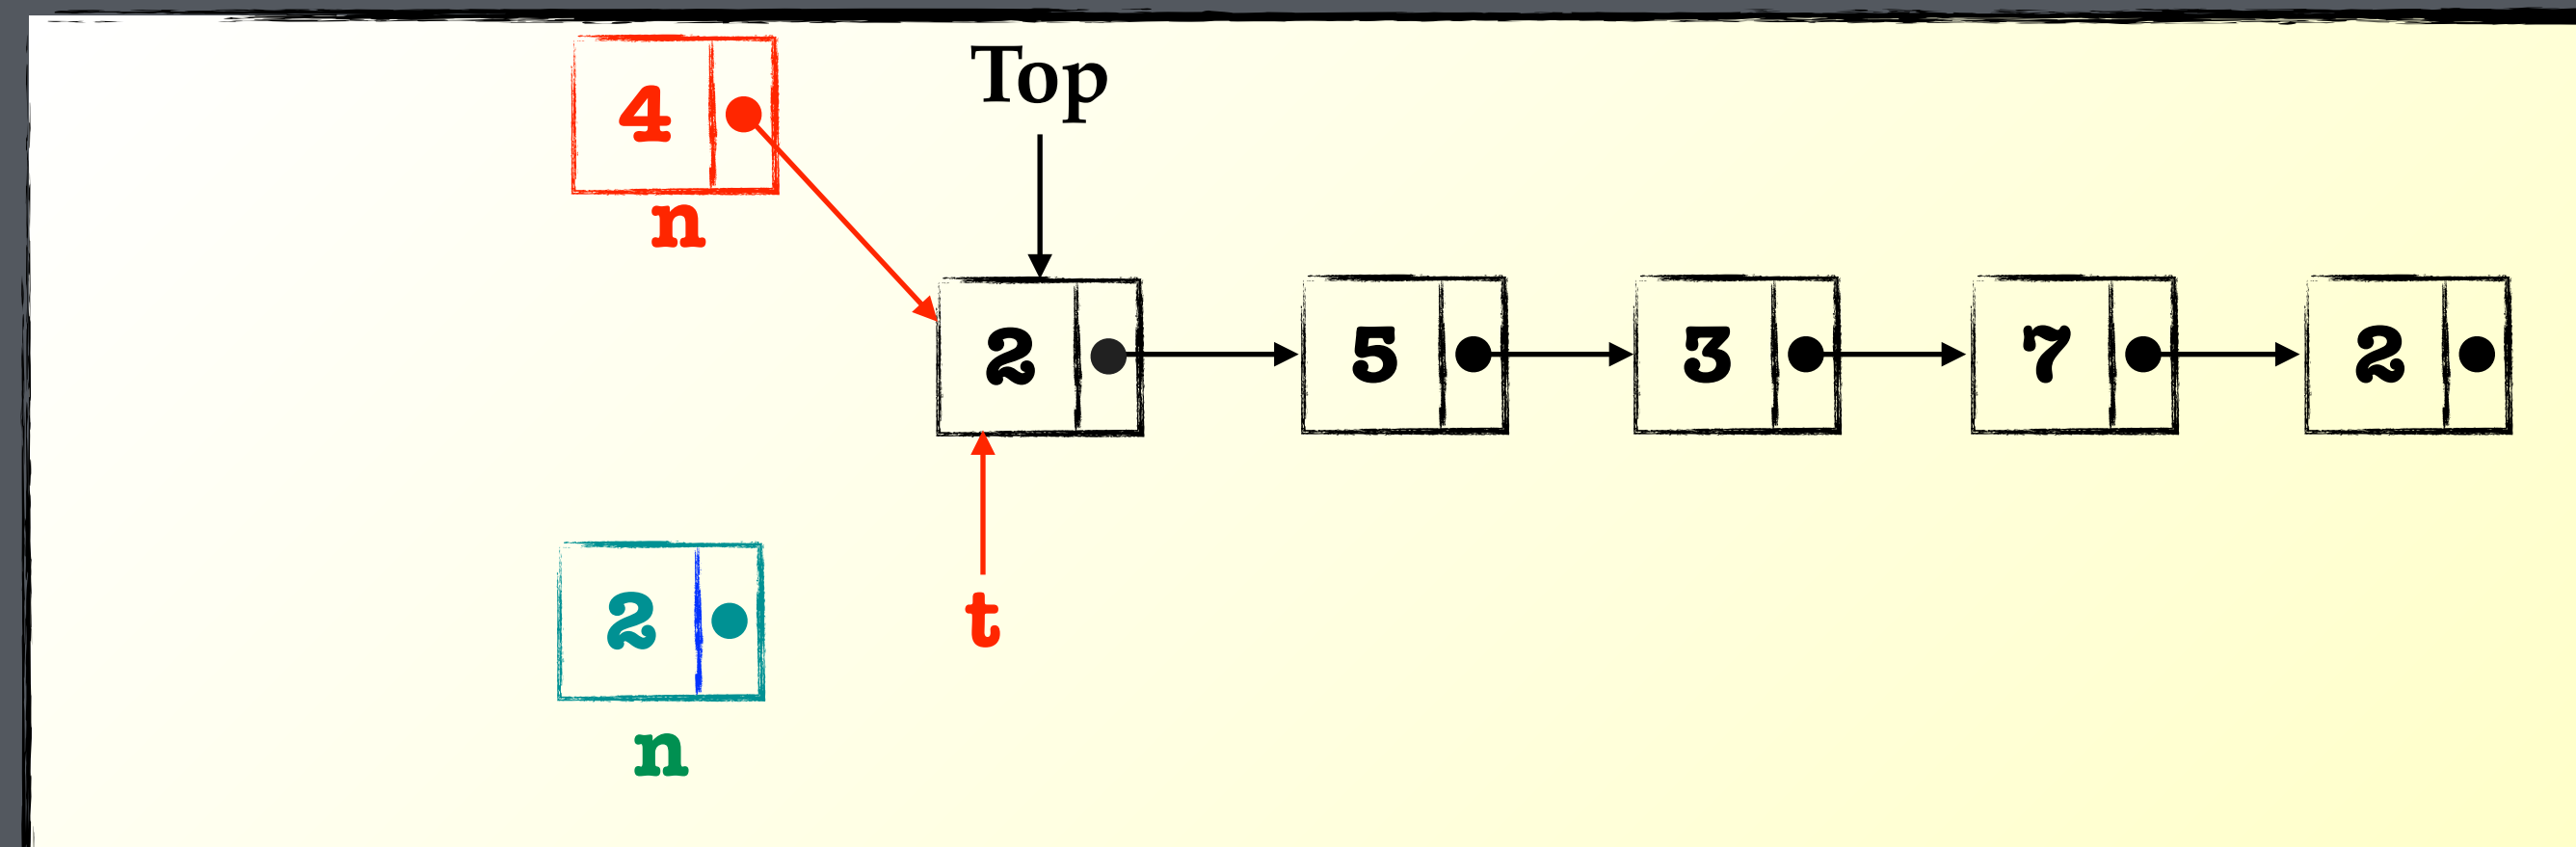


push(4)

push(2)




```
push(k):  
Node t  
1 n = new Node(k,-)  
2 while (true)  
3   t = Top  
4   n.next = t  
5   if (CAS (Top,t,n))  
6     exit
```



# The Treiber Stack Algorithm


Concurrent push operations

```
push(k):  
Node t  
1 n = new Node(k,-)  
2 while (true)  
3   t = Top  
4   n.next = t  
5   if (CAS (Top,t,n))  
6     exit
```

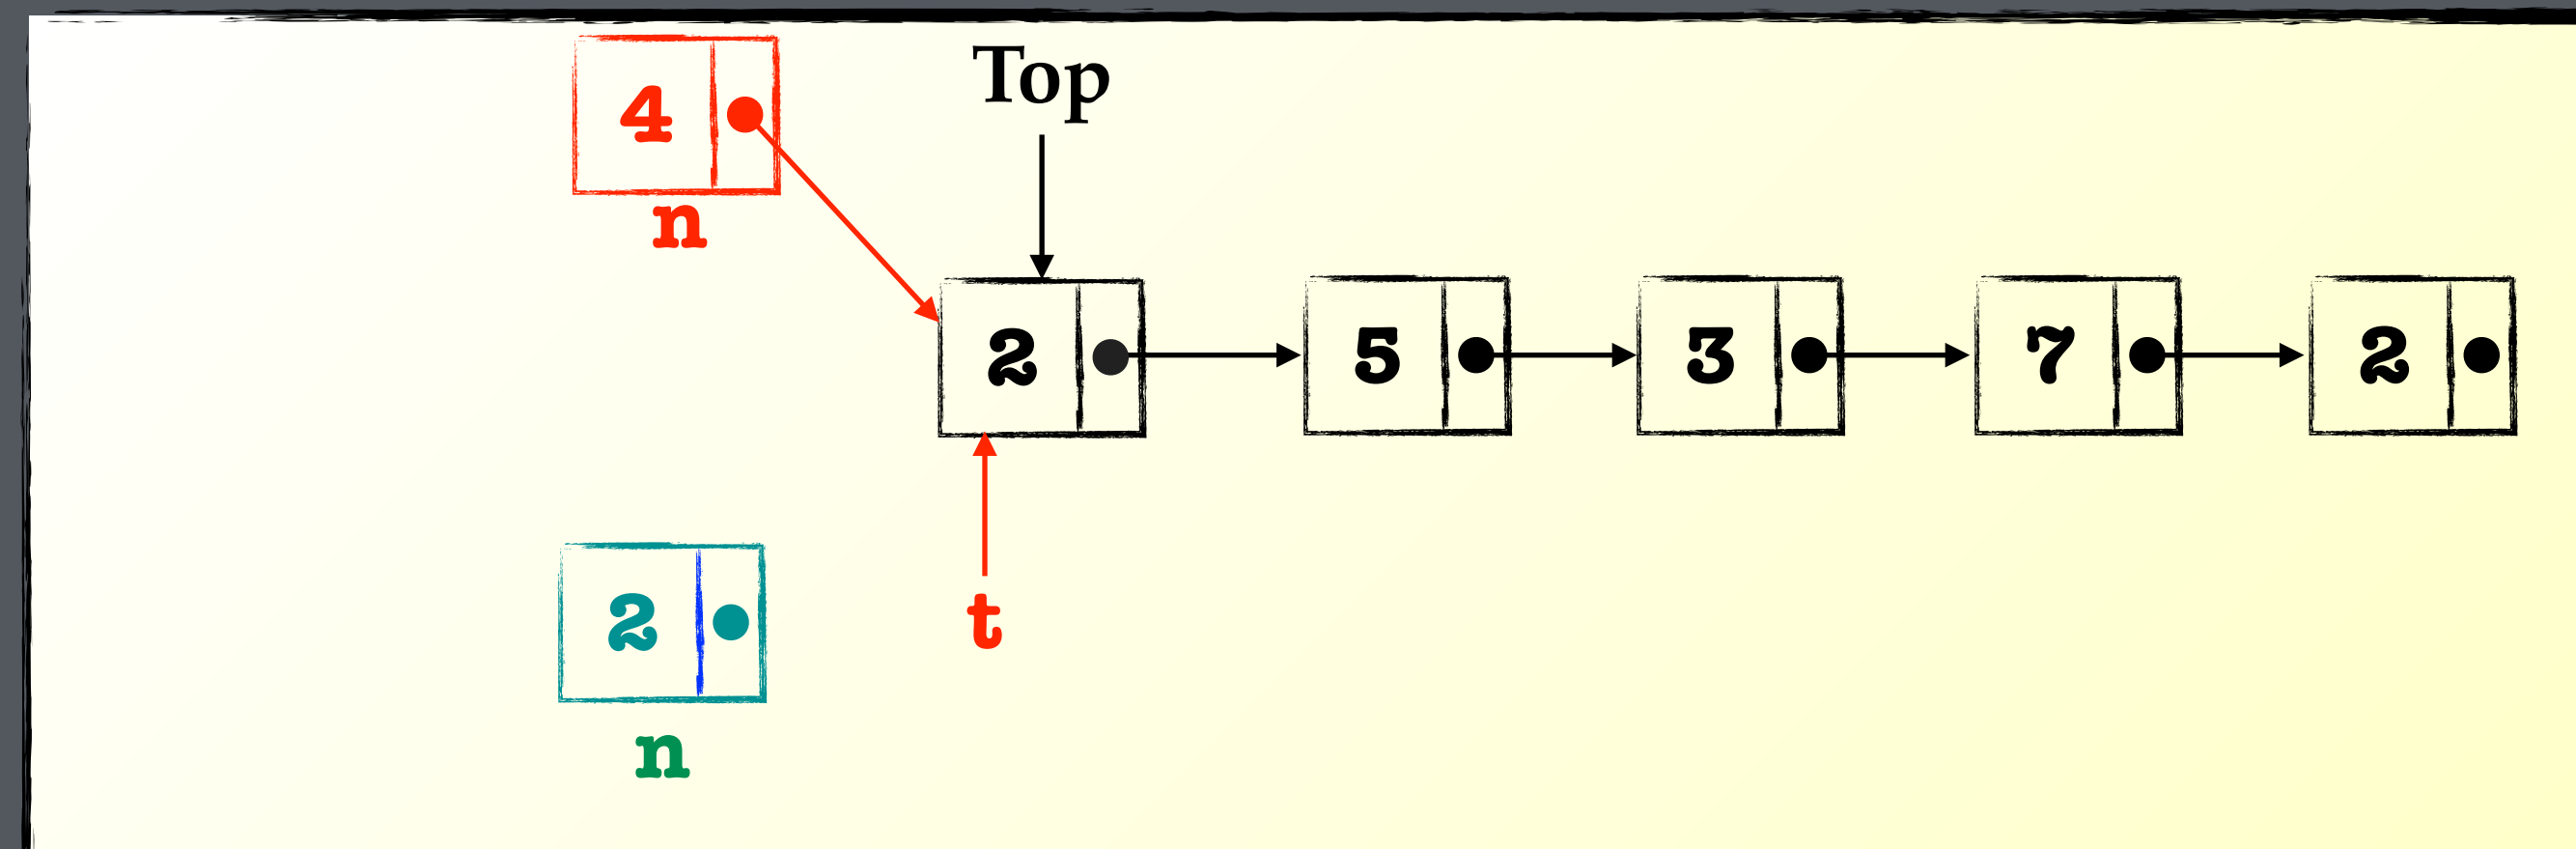


push(4)

push(2)




```
push(k):  
Node t  
1 n = new Node(k,-)  
2 while (true)  
3   t = Top  
4   n.next = t  
5   if (CAS (Top,t,n))  
6     exit
```



# The Treiber Stack Algorithm


Concurrent push operations

```
push(k):  
Node t  
1 n = new Node(k,-)  
2 while (true)  
3   t = Top  
4   n.next = t  
5   if (CAS (Top,t,n))  
6     exit
```

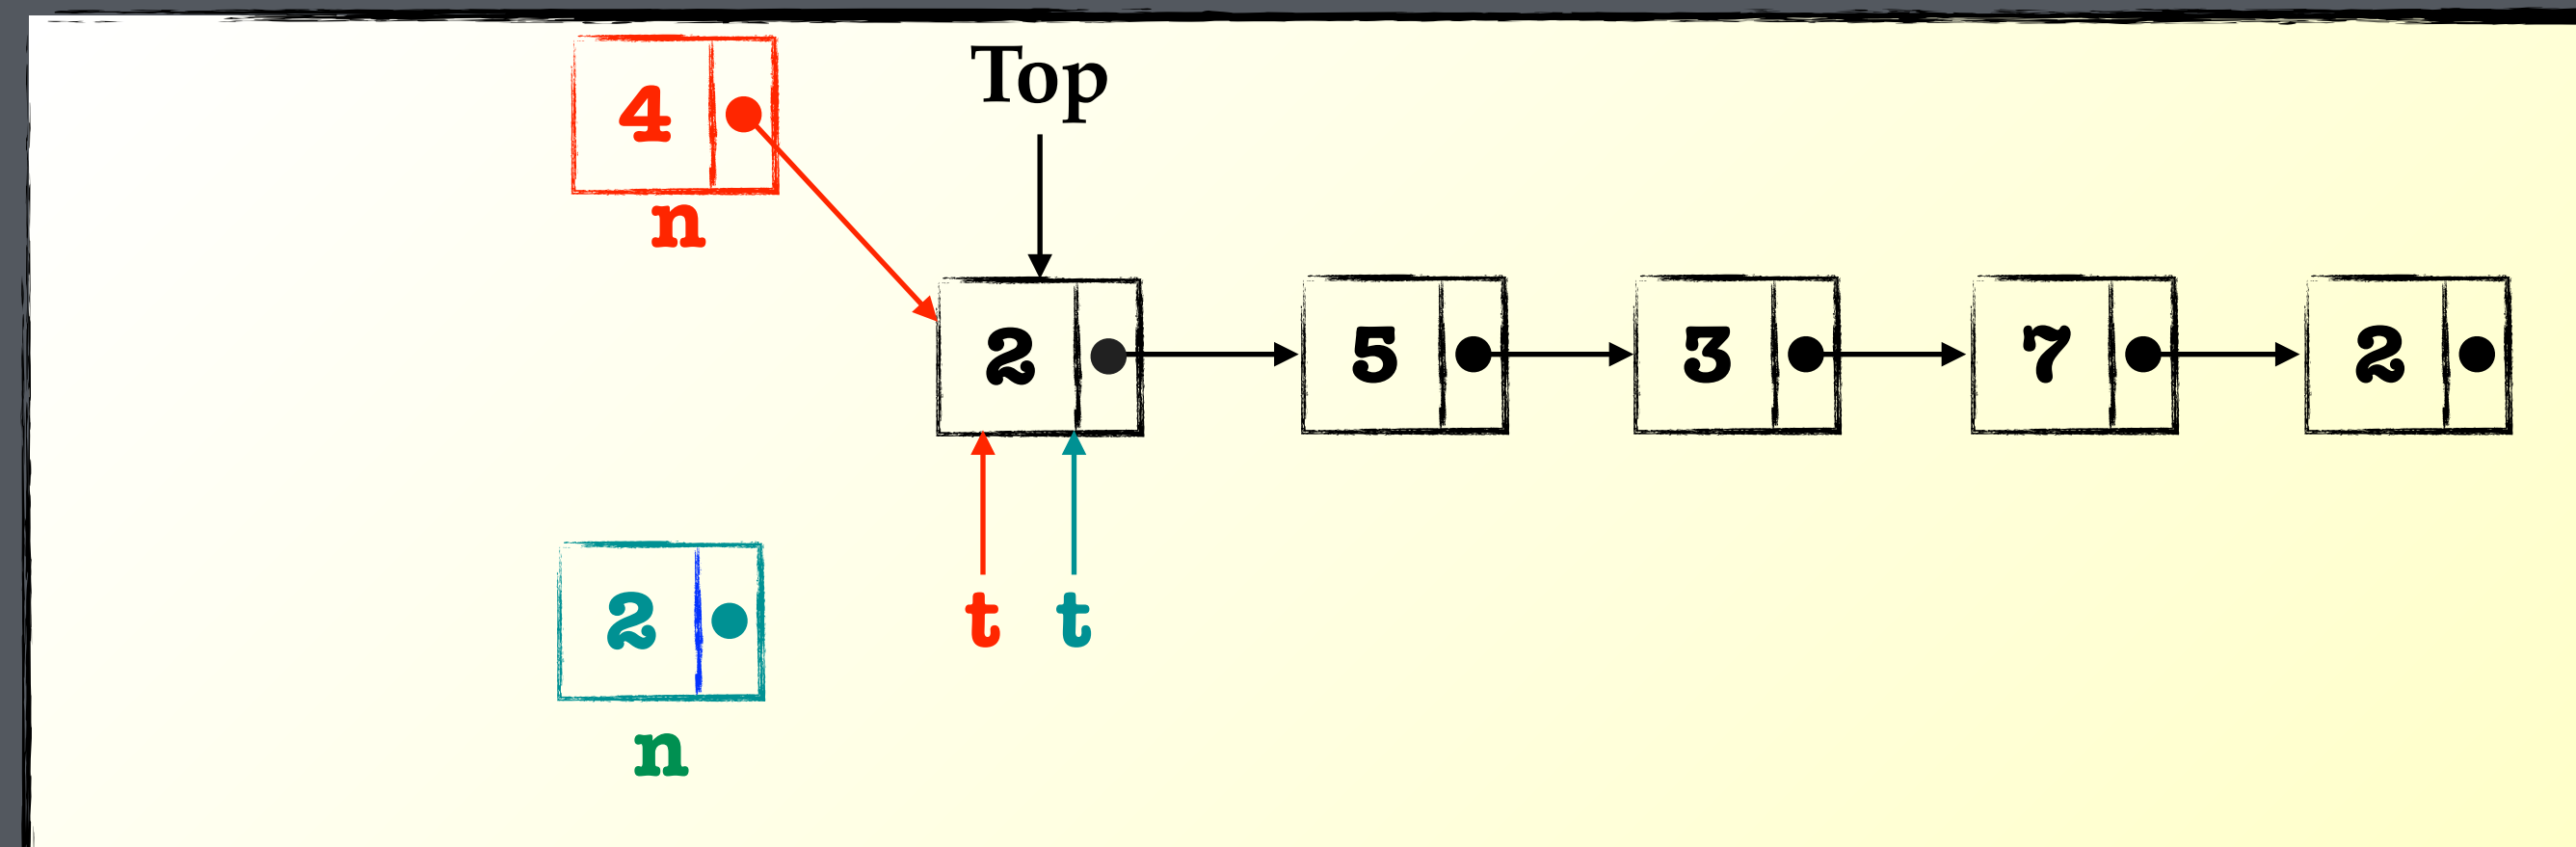


push(4)

push(2)



```
push(k):  
Node t  
1 n = new Node(k,-)  
2 while (true)  
3   t = Top  
4   n.next = t  
5   if (CAS (Top,t,n))  
6     exit
```






# The Treiber Stack Algorithm


Concurrent push operations

```
push(k):  
Node t  
1 n = new Node(k,-)  
2 while (true) ←~~~~~  
3   t = Top  
4   n.next = t ←~~~~~  
5   if (CAS (Top,t,n))  
6     exit
```

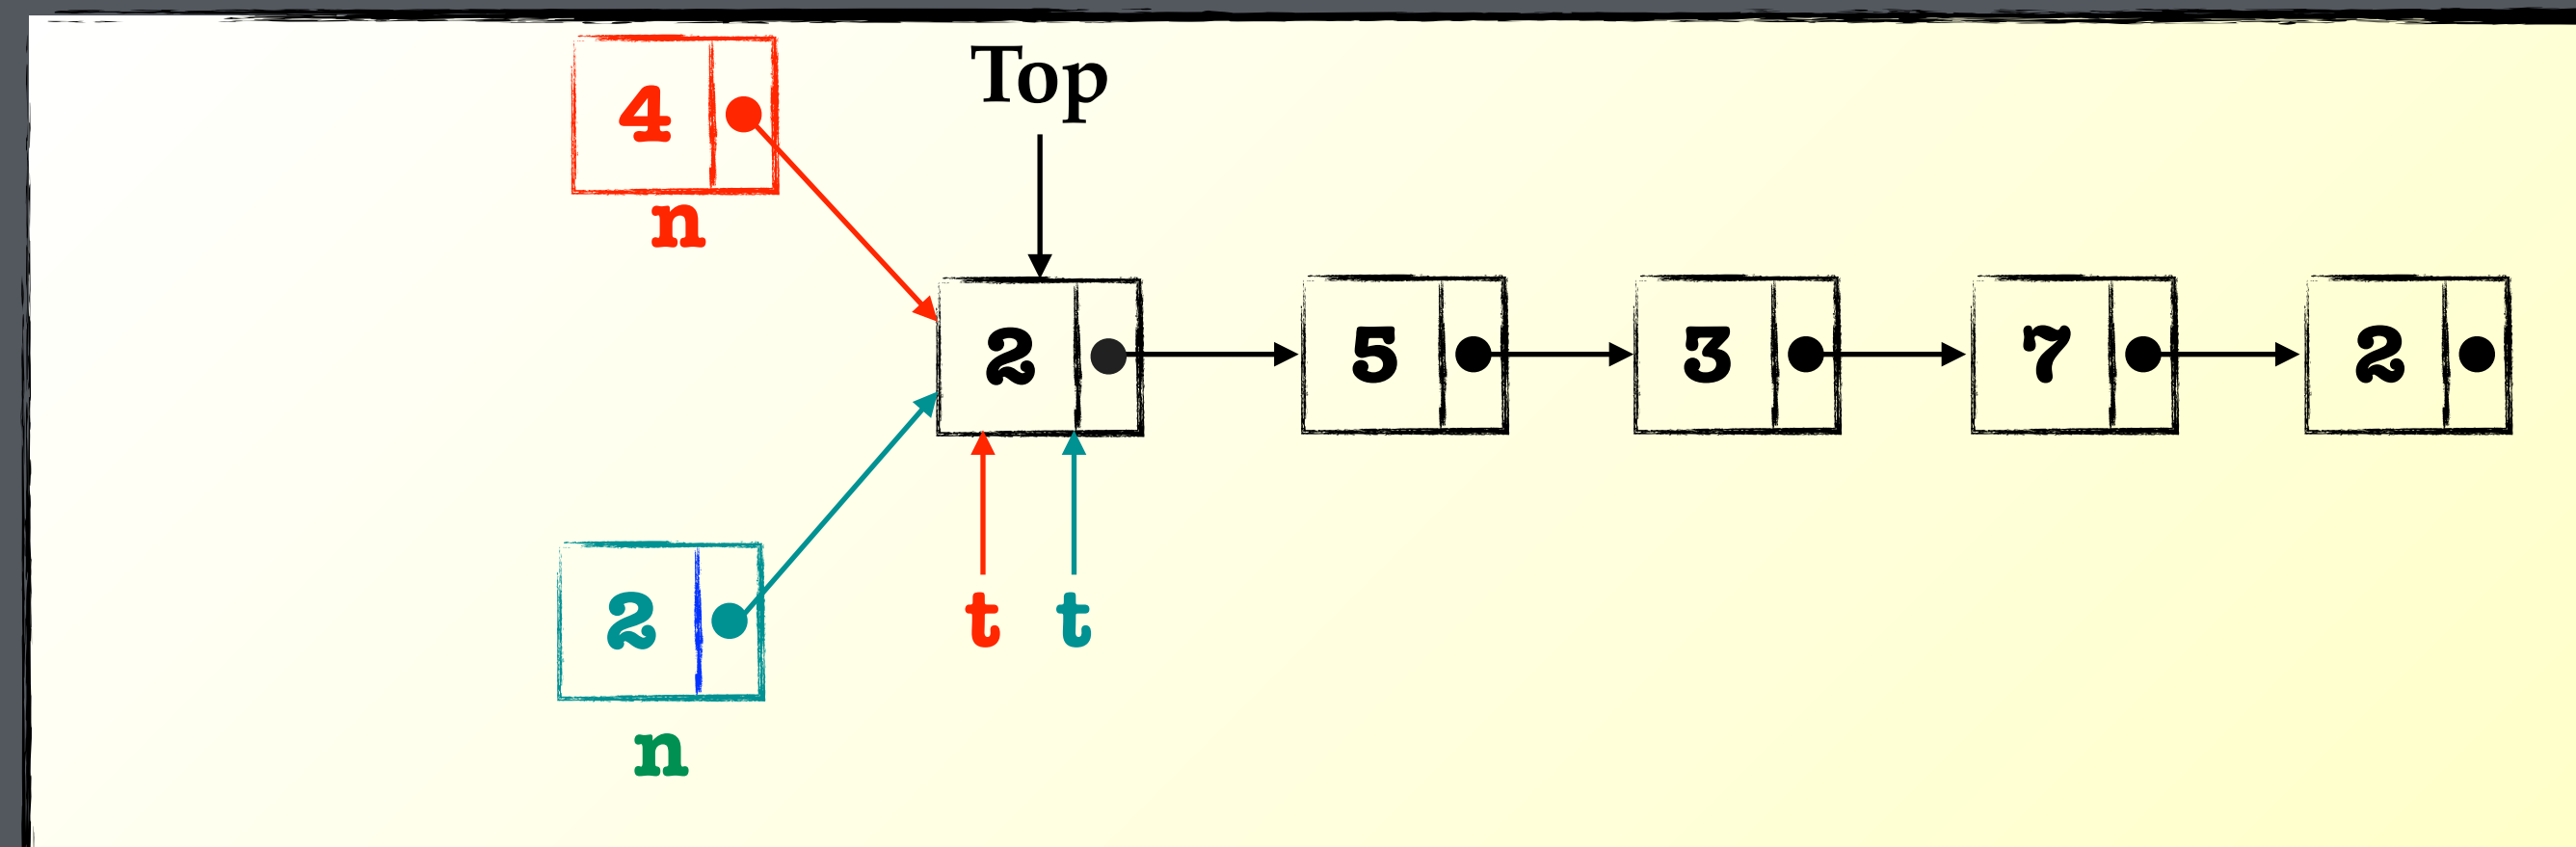


push(4)

push(2)




```
push(k):  
Node t  
1 n = new Node(k,-)  
2 while (true)  
3   t = Top  
4   n.next = t ←~~~~~  
5   if (CAS (Top,t,n))  
6     exit
```



# The Treiber Stack Algorithm


Concurrent push operations

```
push(k):  
Node t  
1 n = new Node(k,-)  
2 while (true) ← wavy red arrow  
3   t = Top  
4   n.next = t ← wavy red arrow  
5   if (CAS (Top,t,n))  
6     exit
```




push(4)

push(2)




```
push(k):  
Node t  
1 n = new Node(k,-)  
2 while (true)  
3   t = Top  
4   n.next = t  
5   if (CAS (Top,t,n)) ← wavy red arrow  
6     exit
```



# The Treiber Stack Algorithm


Concurrent push operations

```
push(k):  
Node t  
1 n = new Node(k,-)  
2 while (true)  
3   t = Top  
4   n.next = t  
5   if (CAS (Top,t,n))  
6     exit
```

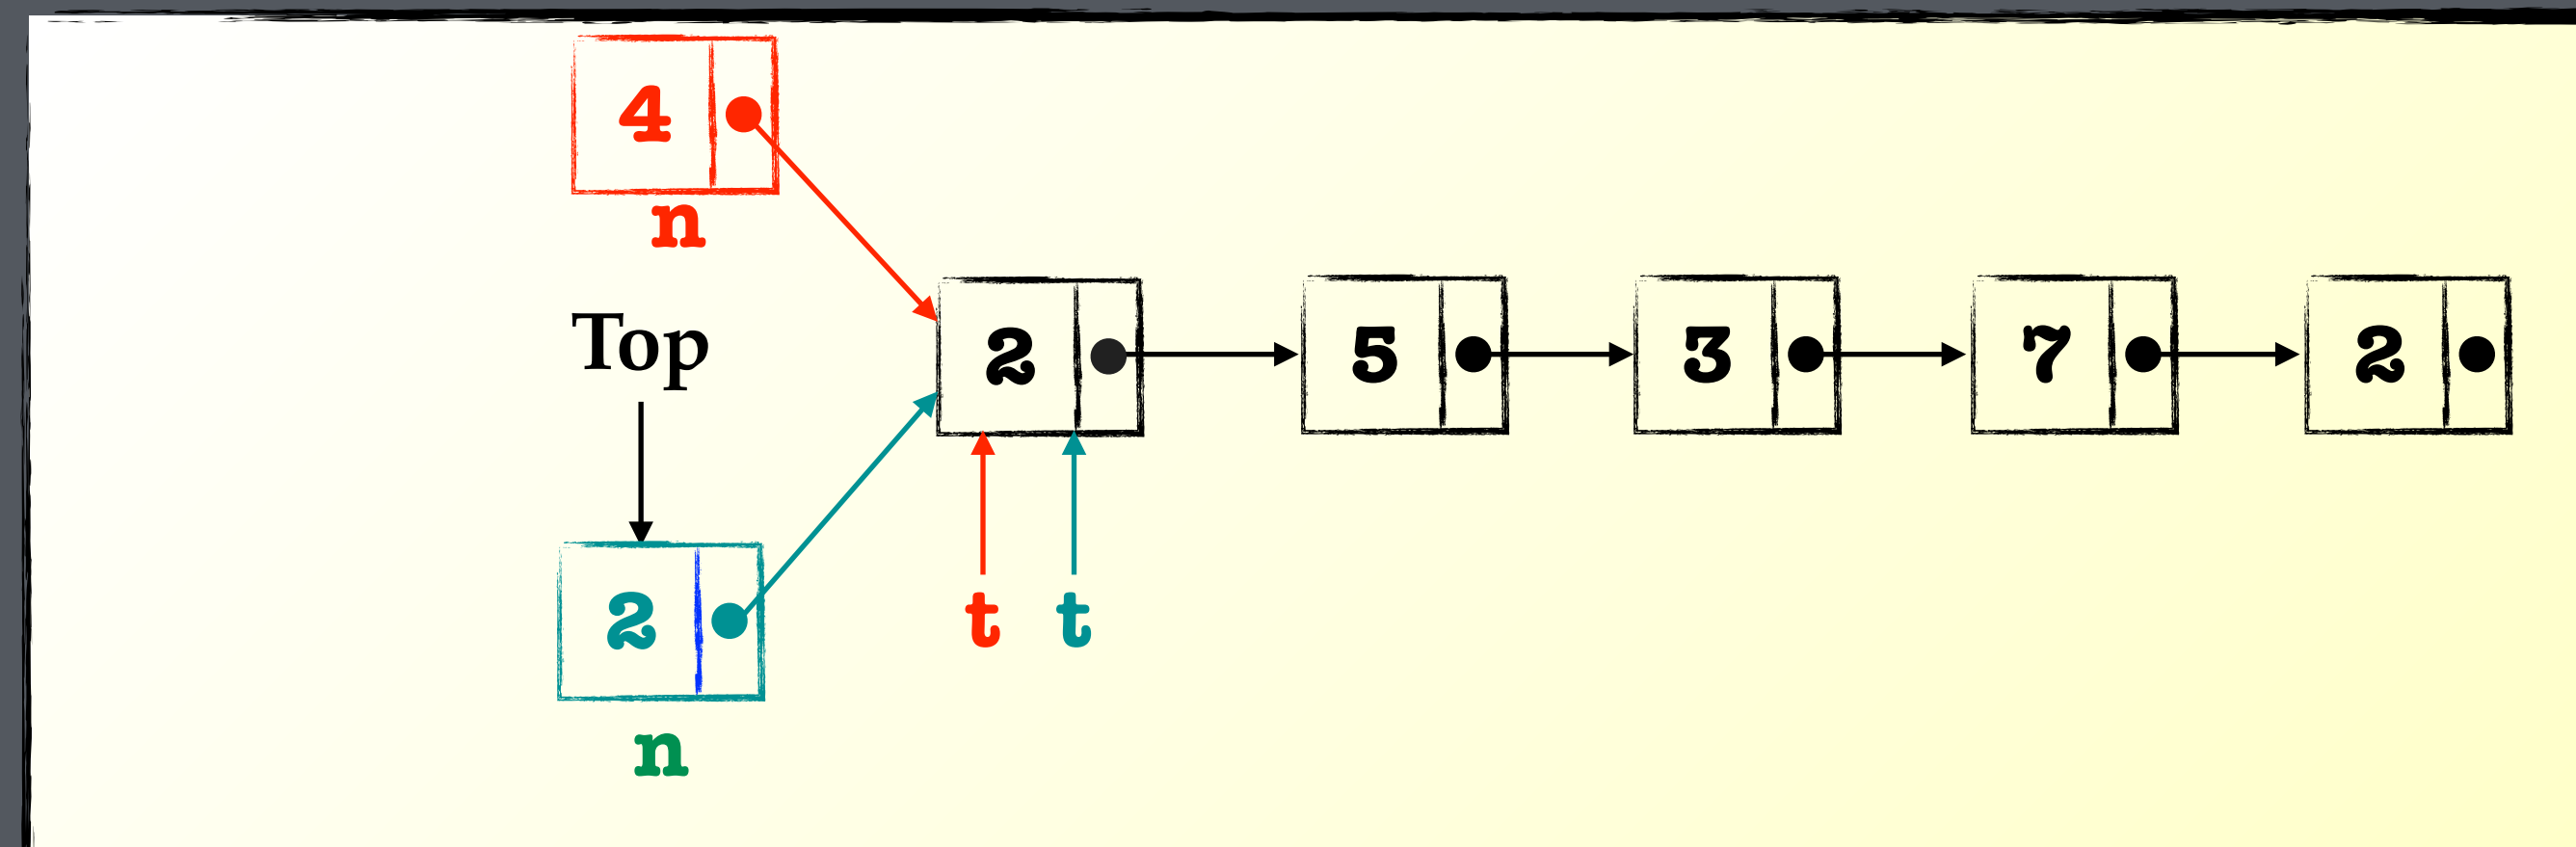


push(4)

push(2)



```
push(k):  
Node t  
1 n = new Node(k,-)  
2 while (true)  
3   t = Top  
4   n.next = t  
5   if (CAS (Top,t,n))  
6     exit
```






# The Treiber Stack Algorithm


Concurrent push operations

```
push(k):  
Node t  
1 n = new Node(k,-)  
2 while (true)  
3   t = Top  
4   n.next = t  
5   if (CAS (Top,t,n))  
6     exit
```




push(4)

push(2)



```
push(k):  
Node t  
1 n = new Node(k,-)  
2 while (true)  
3   t = Top  
4   n.next = t  
5   if (CAS (Top,t,n))  
6     exit
```






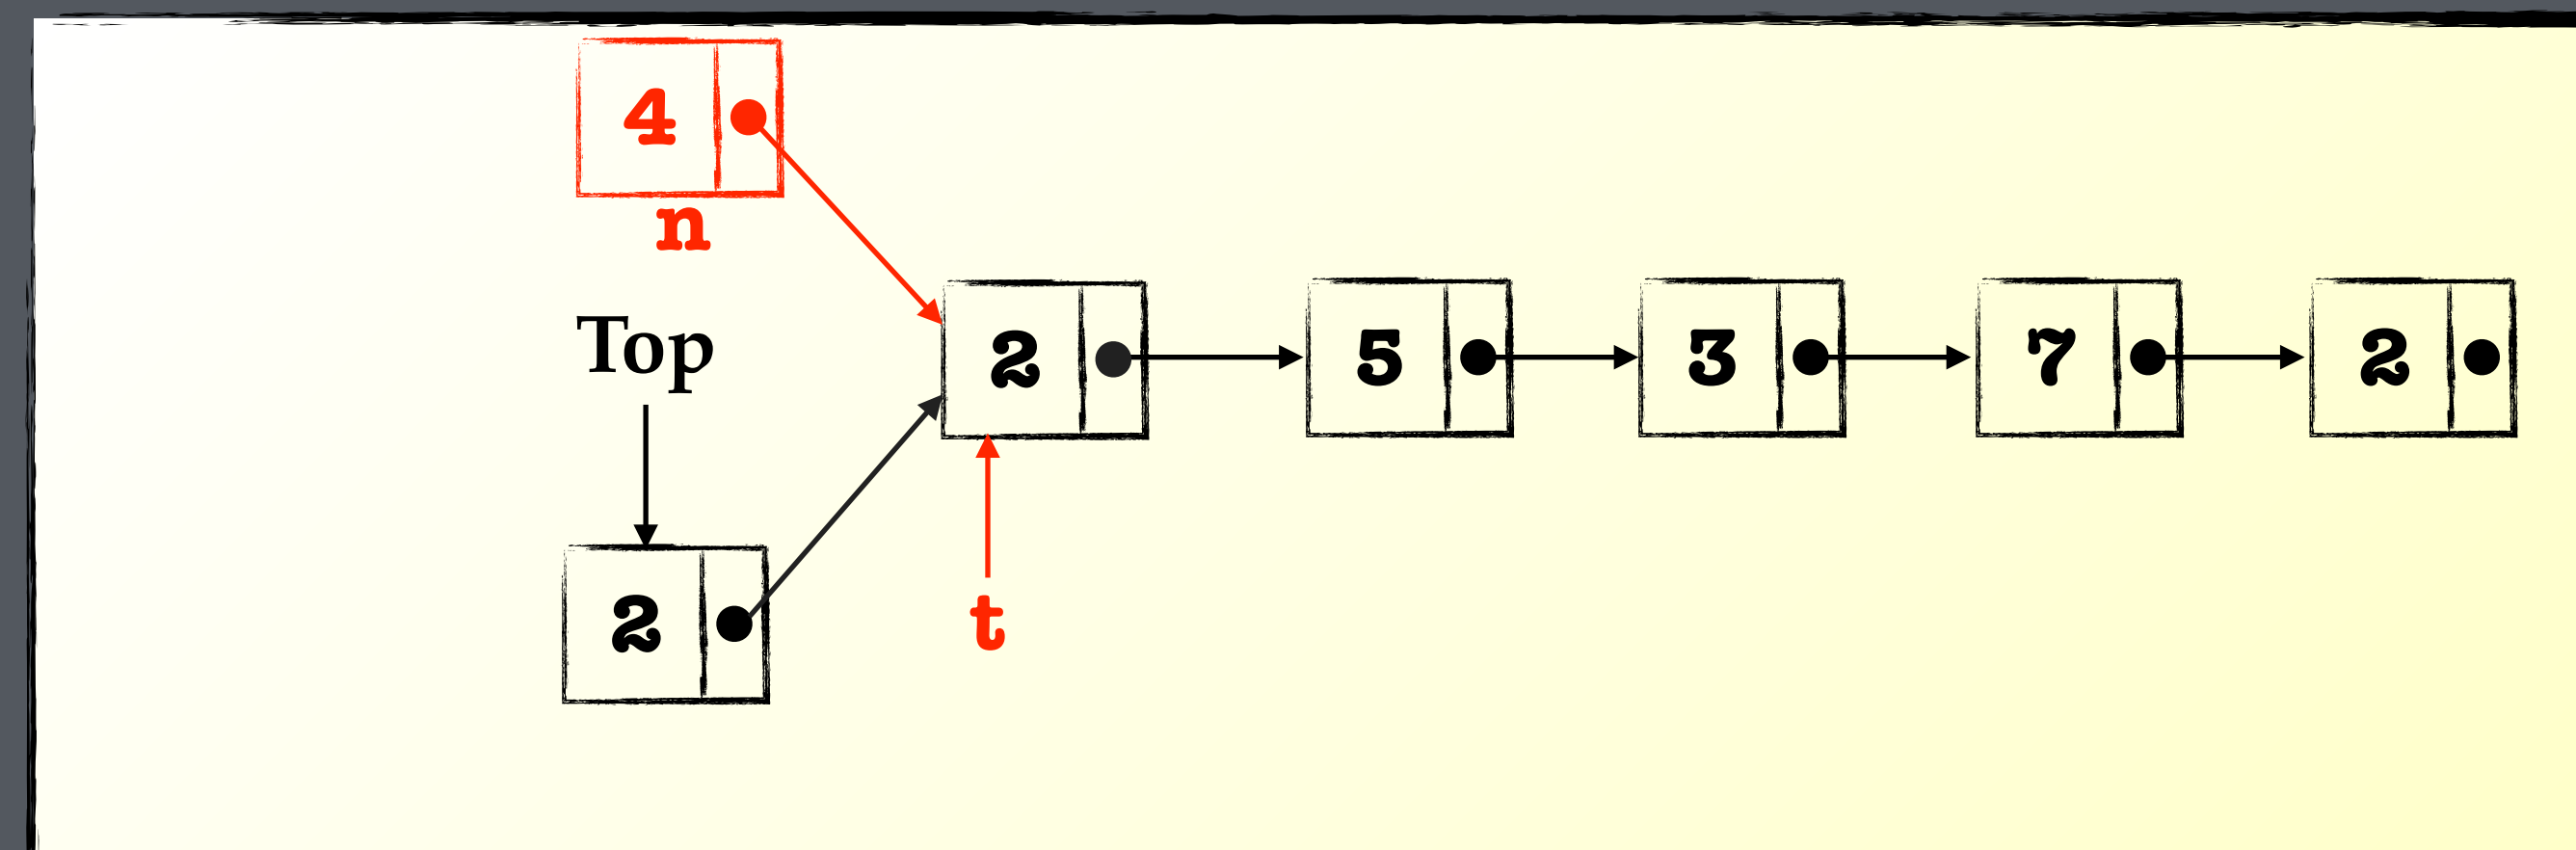
# The Treiber Stack Algorithm

Concurrent push operations

```
push(k):  
Node t  
1  n = new Node(k,-)  
2  while (true)  
3    t = Top  
4    n.next = t  
5    if (CAS (Top,t,n))  
6      exit
```




push(4)




# The Treiber Stack Algorithm

Concurrent push operations


```
push(k):  
Node t  
1  n = new Node(k,-)  
2  while (true)  
3    t = Top  
4    n.next = t  
5    if (CAS (Top,t,n))  
6      exit
```




push(4)



# The Treiber Stack Algorithm


```
push(k):  
Node t  
1 n = new Node(k,-)  
2 while (true)   
3   t = Top  
4   n.next = t  
5   if (CAS (Top,t,n))  
6     exit
```




**push(4)**


**not wait-free (starvation)**

**lock-free**



# The Treiber Stack Algorithm

```
push(k):  
Node t  
1 n = new Node(k,-)  
2 while (true)   
3   t = Top  
4   n.next = t  
5   if (CAS (Top,t,n))  
6     exit
```



**push(4)**


**not wait-free (starvation)**

**lock-free**



**essentially sequential**


**one at a time**






# The Treiber Stack Algorithm

```
pop:
Node t
1 while (true) ←~~~~~
2   t = Top
3   if (t = NULL)
4     return *
5     exit
6   if (CAS (Top,t,t.next))
7     return t.val
8   exit
```




pop

Top




# The Treiber Stack Algorithm

```
pop:
Node t
1 while (true)
2   t = Top ←~~~~~
3   if (t = NULL)
4     return *
5     exit
6   if (CAS (Top,t,t.next))
7     return t.val
8   exit
```




pop

Top




# The Treiber Stack Algorithm

```
pop:
Node t
1 while (true)
2   t = Top
3   if (t = NULL) ←~~~~~
4     return *
5     exit
6   if (CAS (Top,t,t.next))
7     return t.val
8   exit
```




**pop**






# The Treiber Stack Algorithm

```
pop:
Node t
1 while (true)
2   t = Top
3   if (t = NULL)
4     return *
5     exit
6   if (CAS (Top,t,t.next)) ←
7     return t.val
8     exit
```




pop

Top

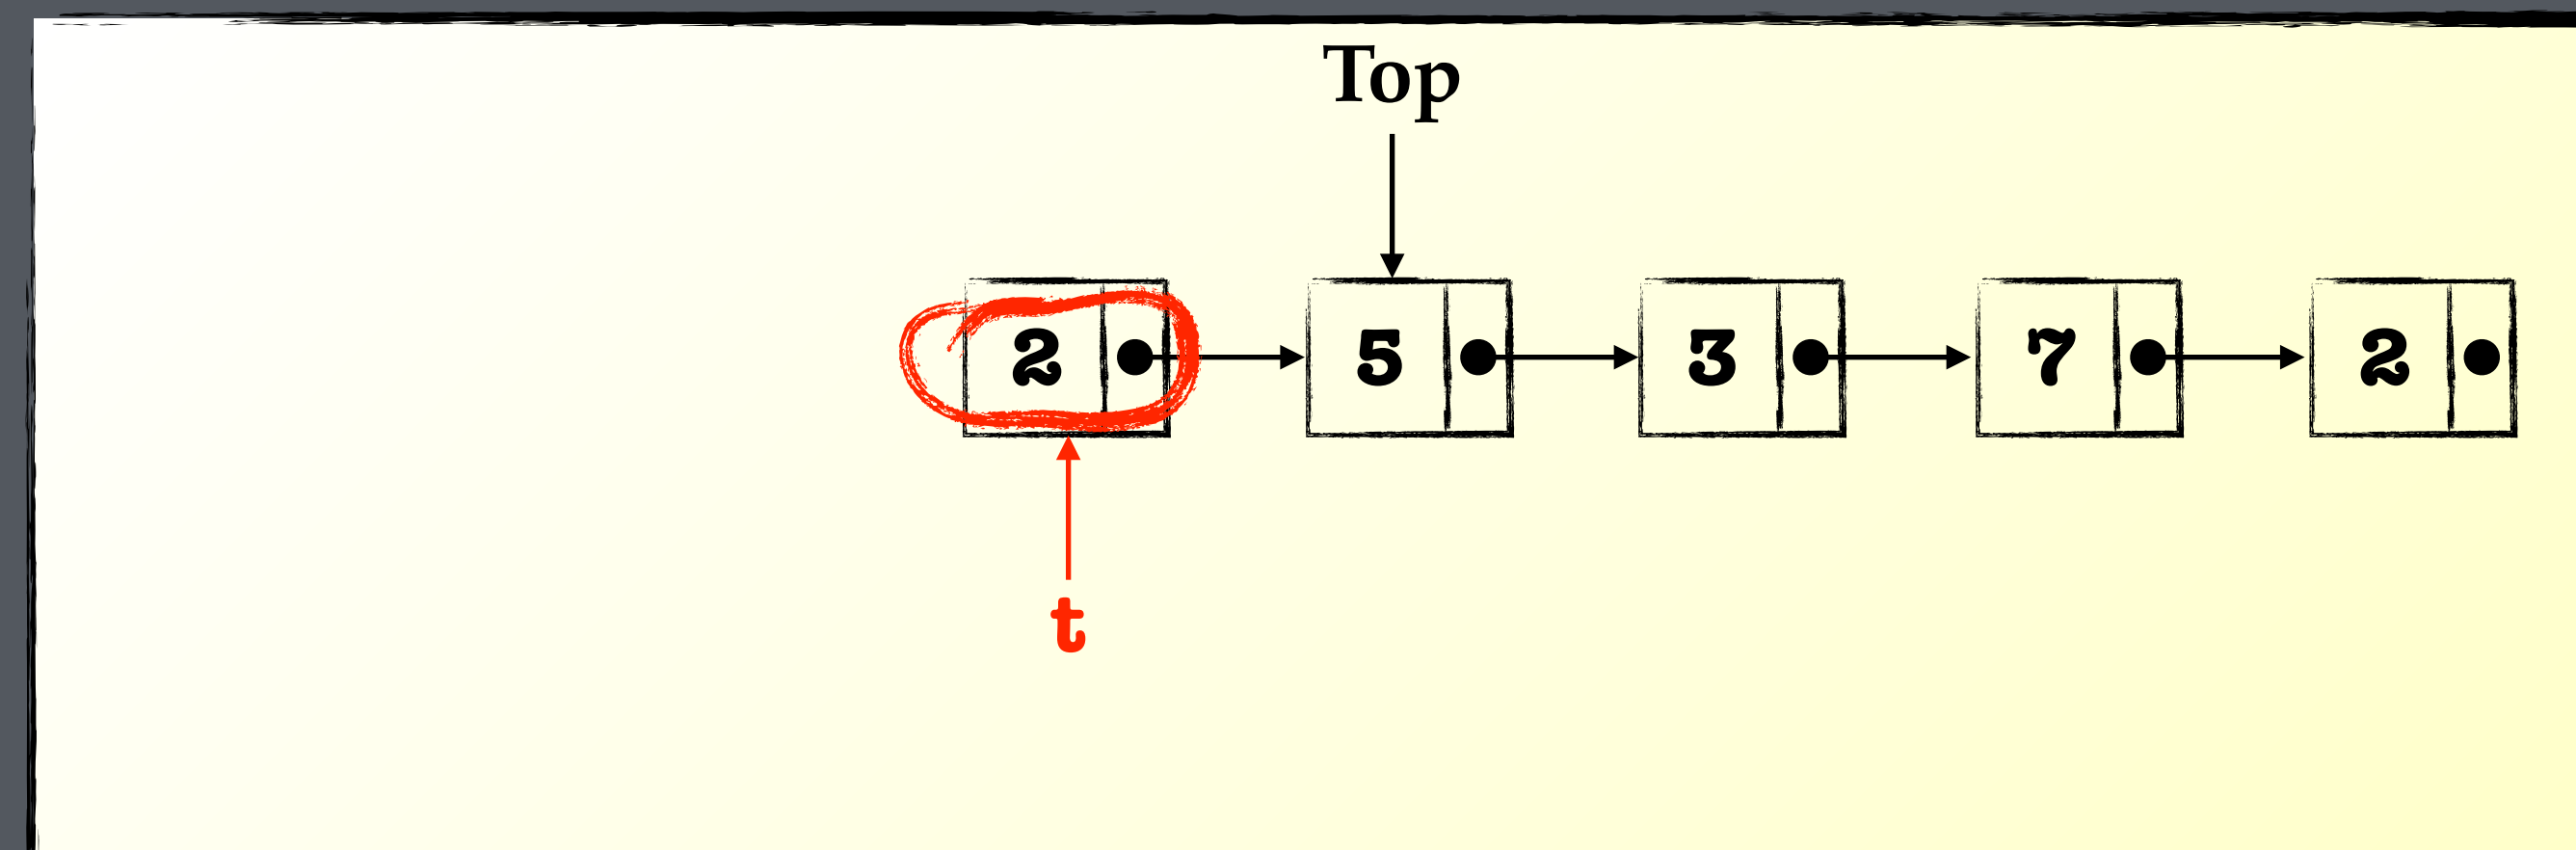


# The Treiber Stack Algorithm

```
pop:
Node t
1 while (true)
2   t = Top
3   if (t = NULL)
4     return *
5     exit
6   if (CAS (Top,t.t.next))
7     return t.val
8   exit
```

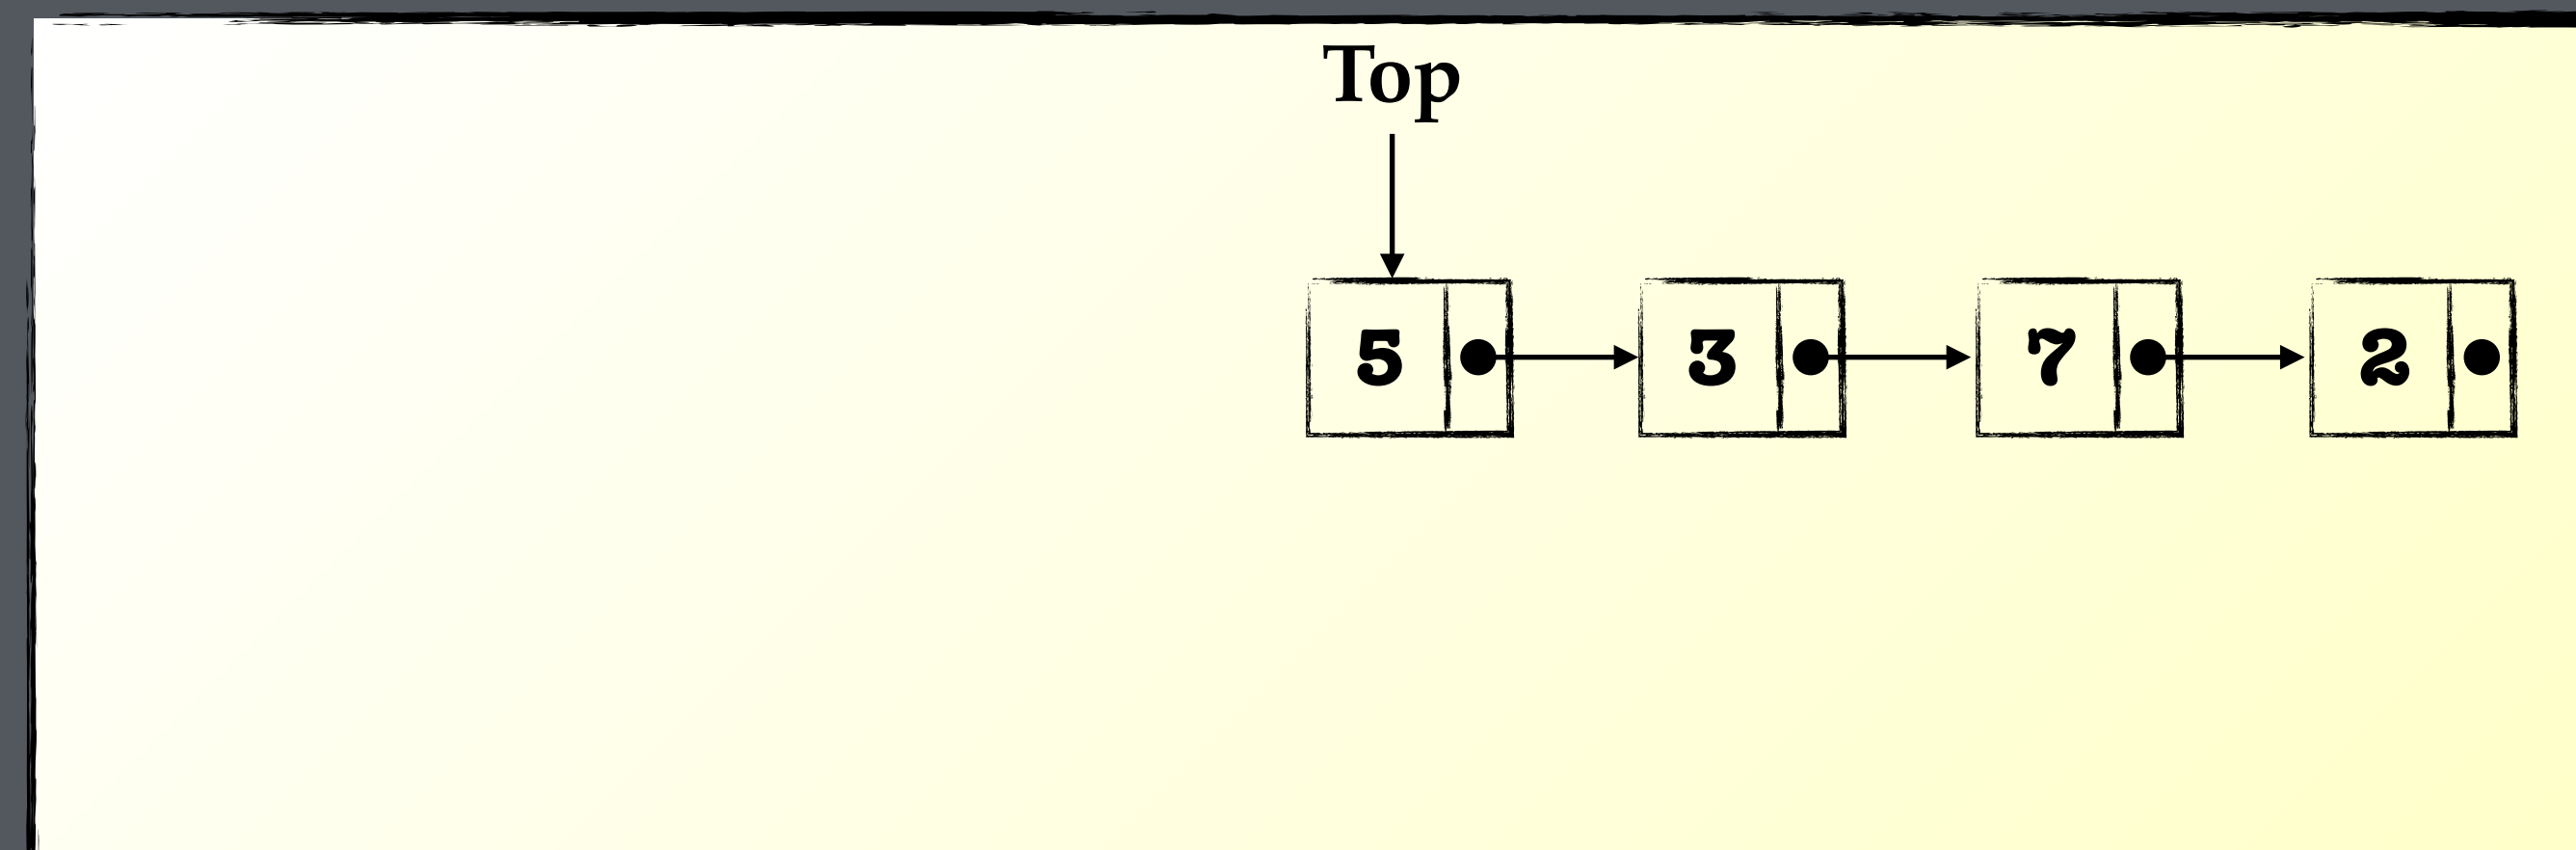


pop



# The Treiber Stack Algorithm


Concurrent pop operations



# The Treiber Stack Algorithm


Concurrent pop operations

```
pop:
Node t
1 while (true) ←~~~~~
2   t = Top
3   if (t = NULL)
4     return *
5     exit
6   if (CAS (Top,t,t.next))
7     return t.val
8   exit
```



pop


Top



# The Treiber Stack Algorithm


Concurrent pop operations

```
pop:
Node t
1 while (true)
2   t = Top ←~~~~~
3   if (t = NULL)
4     return *
5     exit
6   if (CAS (Top,t,t.next))
7     return t.val
8   exit
```



pop

Top






# The Treiber Stack Algorithm


Concurrent pop operations

```
pop:
Node t
1 while (true)
2   t = Top
3   if (t = NULL)
4     return *
5     exit
6   if (CAS (Top,t,t.next))
7     return t.val
8   exit
```




pop

pop



```
pop:
Node t
1 while (true)
2   t = Top
3   if (t = NULL)
4     return *
5     exit
6   if (CAS (Top,t,t.next))
7     return t.val
8   exit
```

Top






# The Treiber Stack Algorithm


Concurrent pop operations

```
pop:
Node t
1 while (true)
2   t = Top
3   if (t = NULL)
4     return *
5     exit
6   if (CAS (Top,t,t.next))
7     return t.val
8   exit
```




pop

pop



```
pop:
Node t
1 while (true)
2   t = Top
3   if (t = NULL)
4     return *
5     exit
6   if (CAS (Top,t,t.next))
7     return t.val
8   exit
```


Top



# The Treiber Stack Algorithm


Concurrent pop operations

```
pop:
Node t
1 while (true)
2   t = Top
3   if (t = NULL)
4     return *
5     exit
6   if (CAS (Top,t,t.next))
7     return t.val
8   exit
```




pop

pop



```
pop:
Node t
1 while (true)
2   t = Top
3   if (t = NULL)
4     return *
5     exit
6   if (CAS (Top,t,t.next))
7     return t.val
8   exit
```


Top



# The Treiber Stack Algorithm


Concurrent pop operations

```
pop:
Node t
1 while (true)
2   t = Top
3   if (t = NULL)
4     return *
5   exit
6   if (CAS (Top,t,t.next))
7     return t.val
8   exit
```

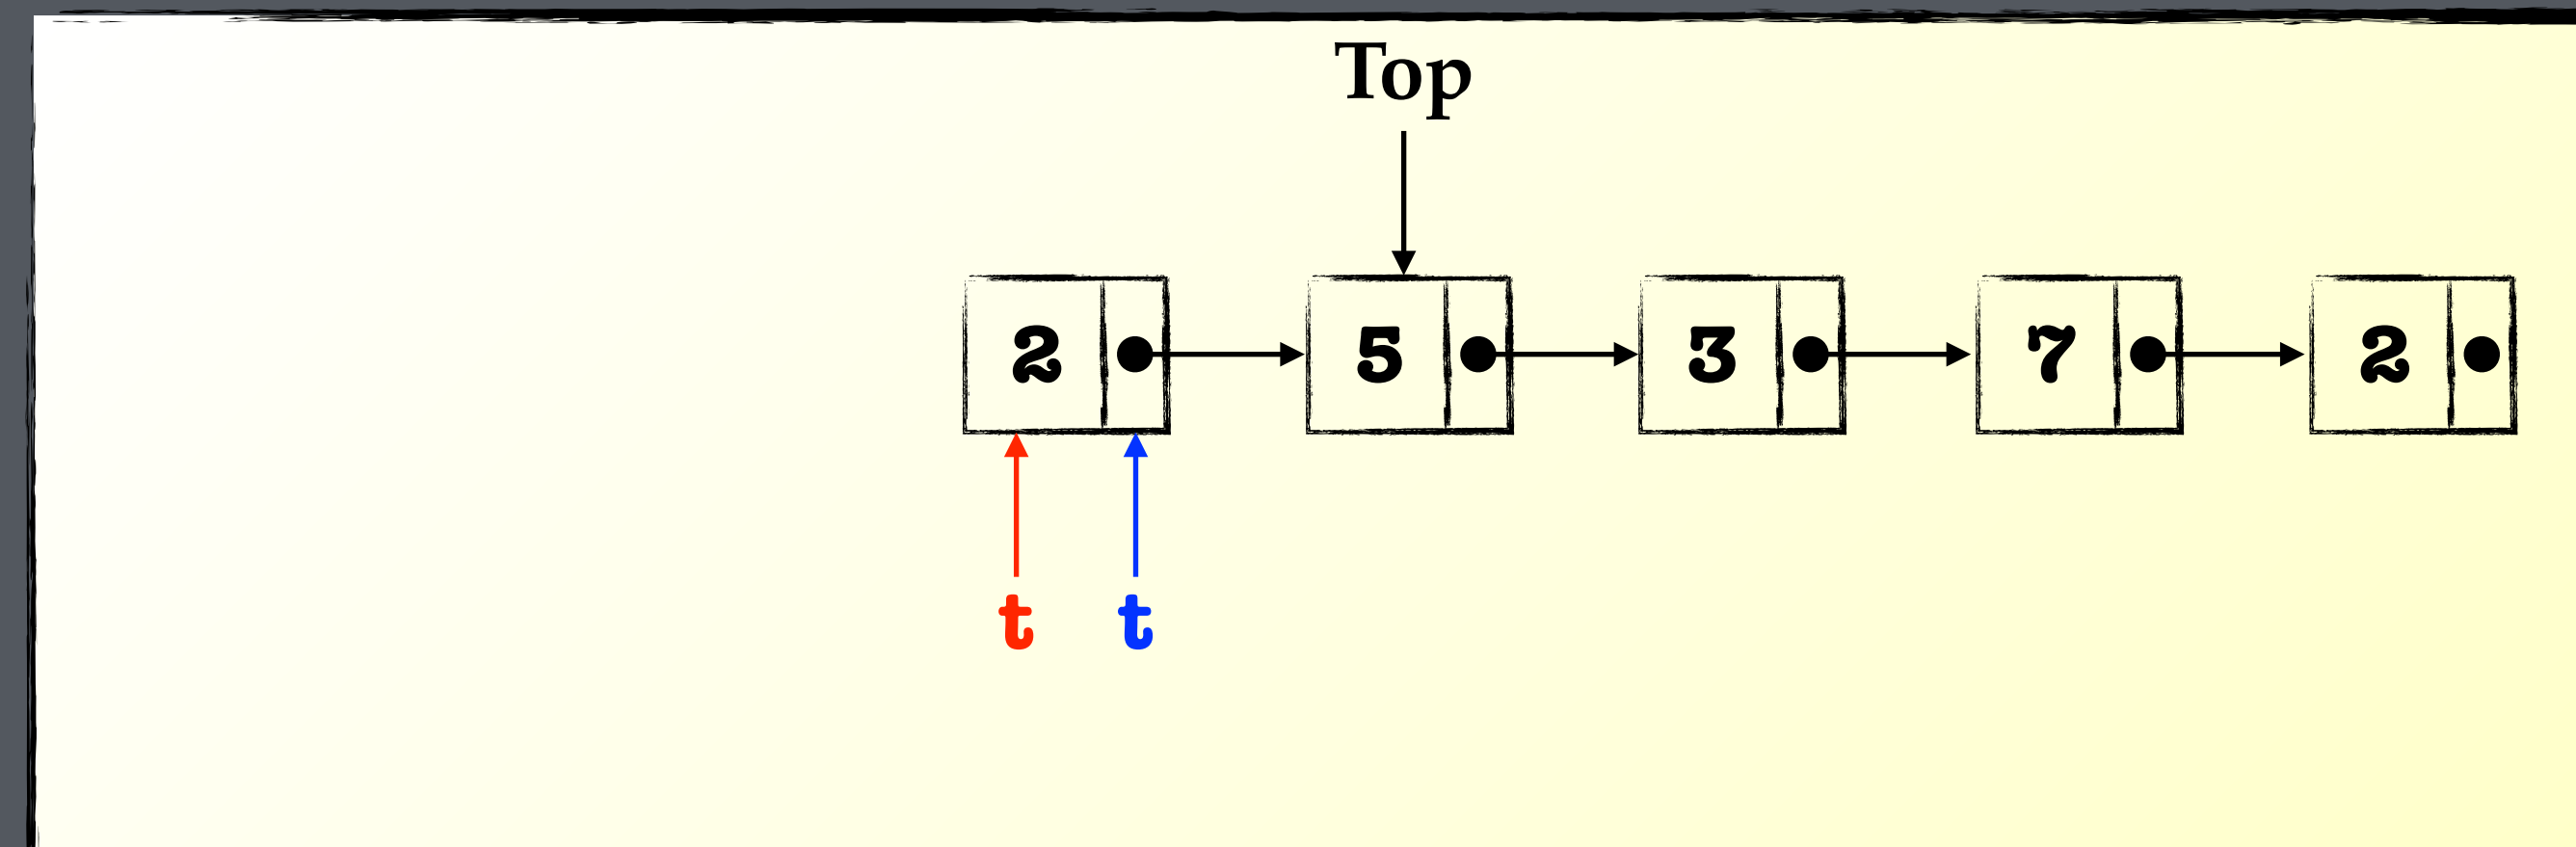


pop

pop



```
pop:
Node t
1 while (true)
2   t = Top
3   if (t = NULL)
4     return *
5   exit
6   if (CAS (Top,t,t.next))
7     return t.val
8   exit
```






# The Treiber Stack Algorithm


Concurrent pop operations

```
pop:
Node t
1 while (true)
2   t = Top
3   if (t = NULL)
4     return *
5     exit
6   if (CAS (Top,t,t.next))
7     return t.val
8   exit
```

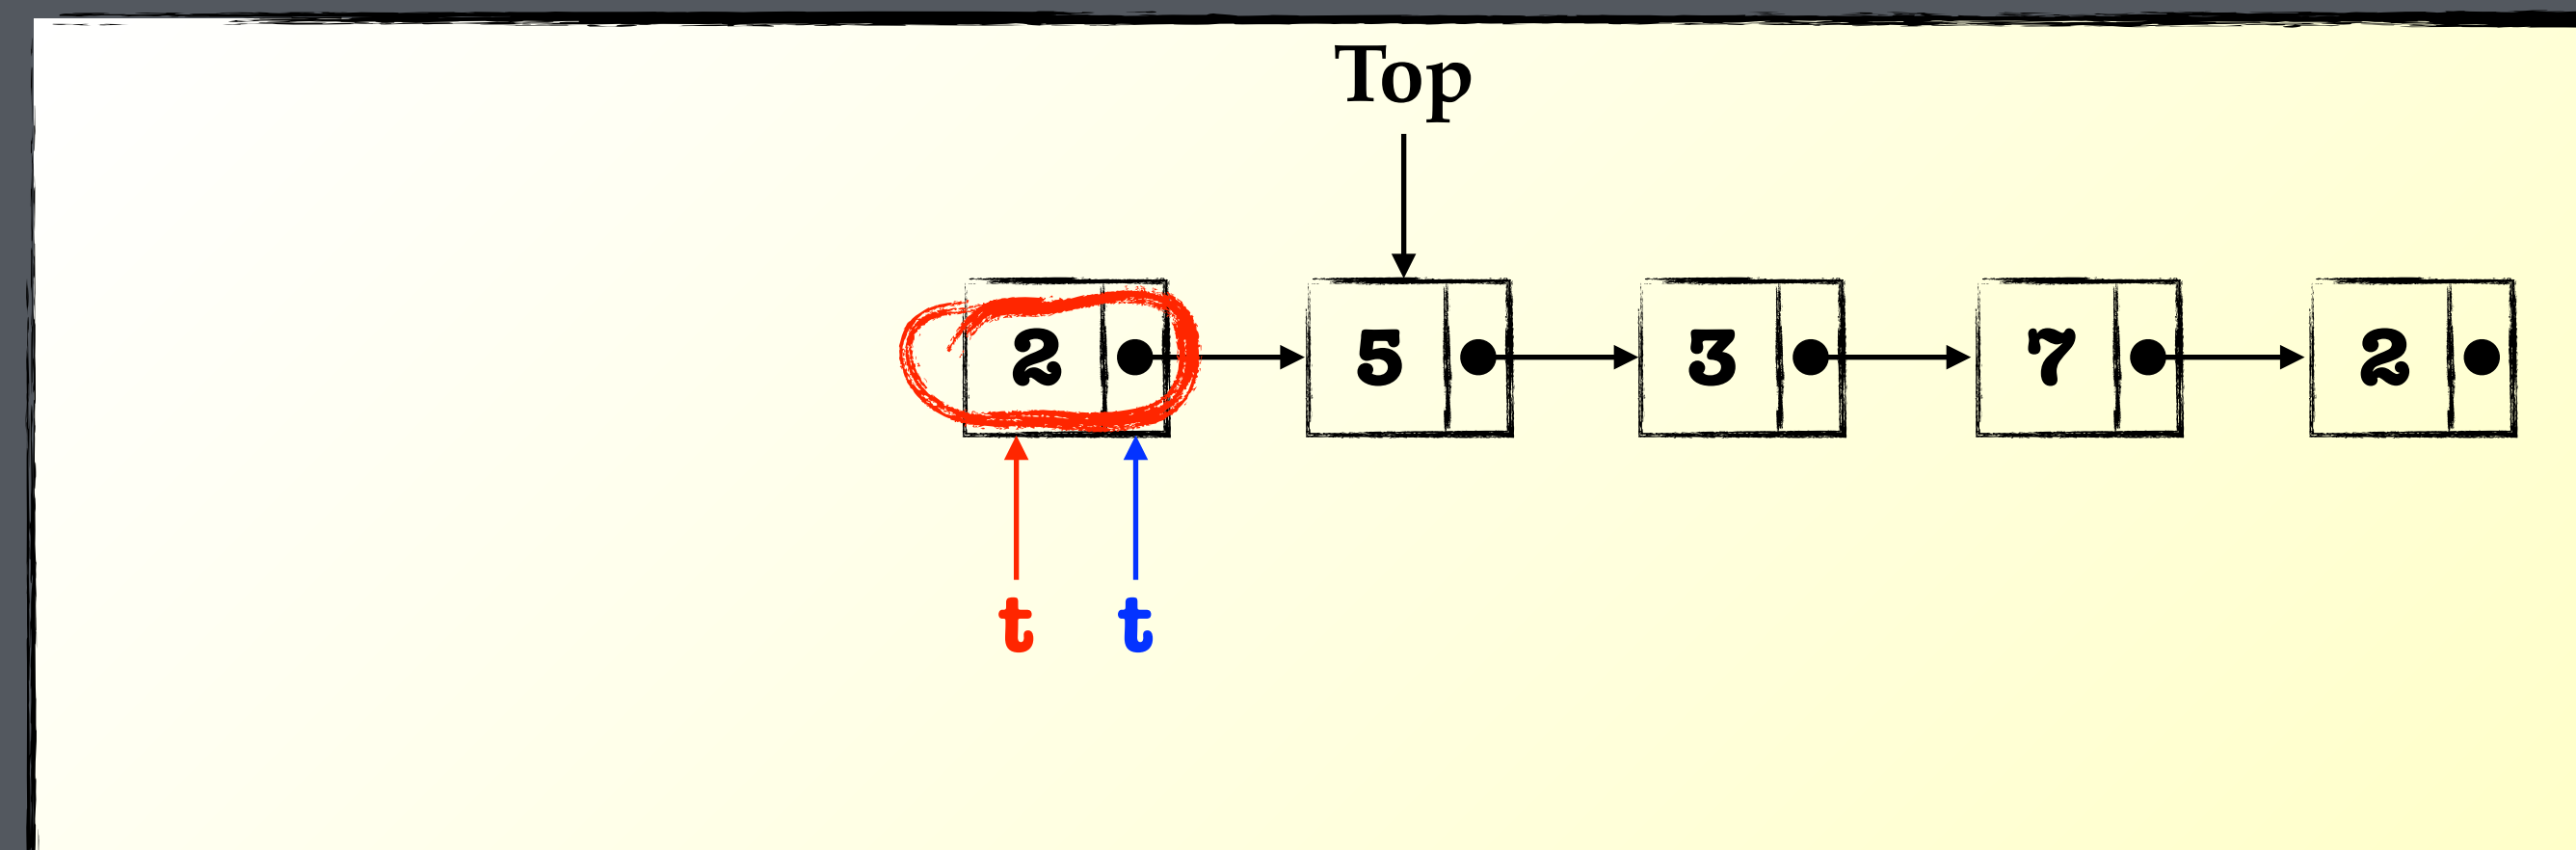


pop

pop




```
pop:
Node t
1 while (true)
2   t = Top
3   if (t = NULL)
4     return *
5     exit
6   if (CAS (Top,t,t.next))
7     return t.val
8   exit
```



# The Treiber Stack Algorithm


Concurrent pop operations

```
pop:
Node t
1 while (true)
2   t = Top
3   if (t = NULL)
4     return *
5     exit
6   if (CAS (Top,t,t.next))
7     return t.val
8   exit
```

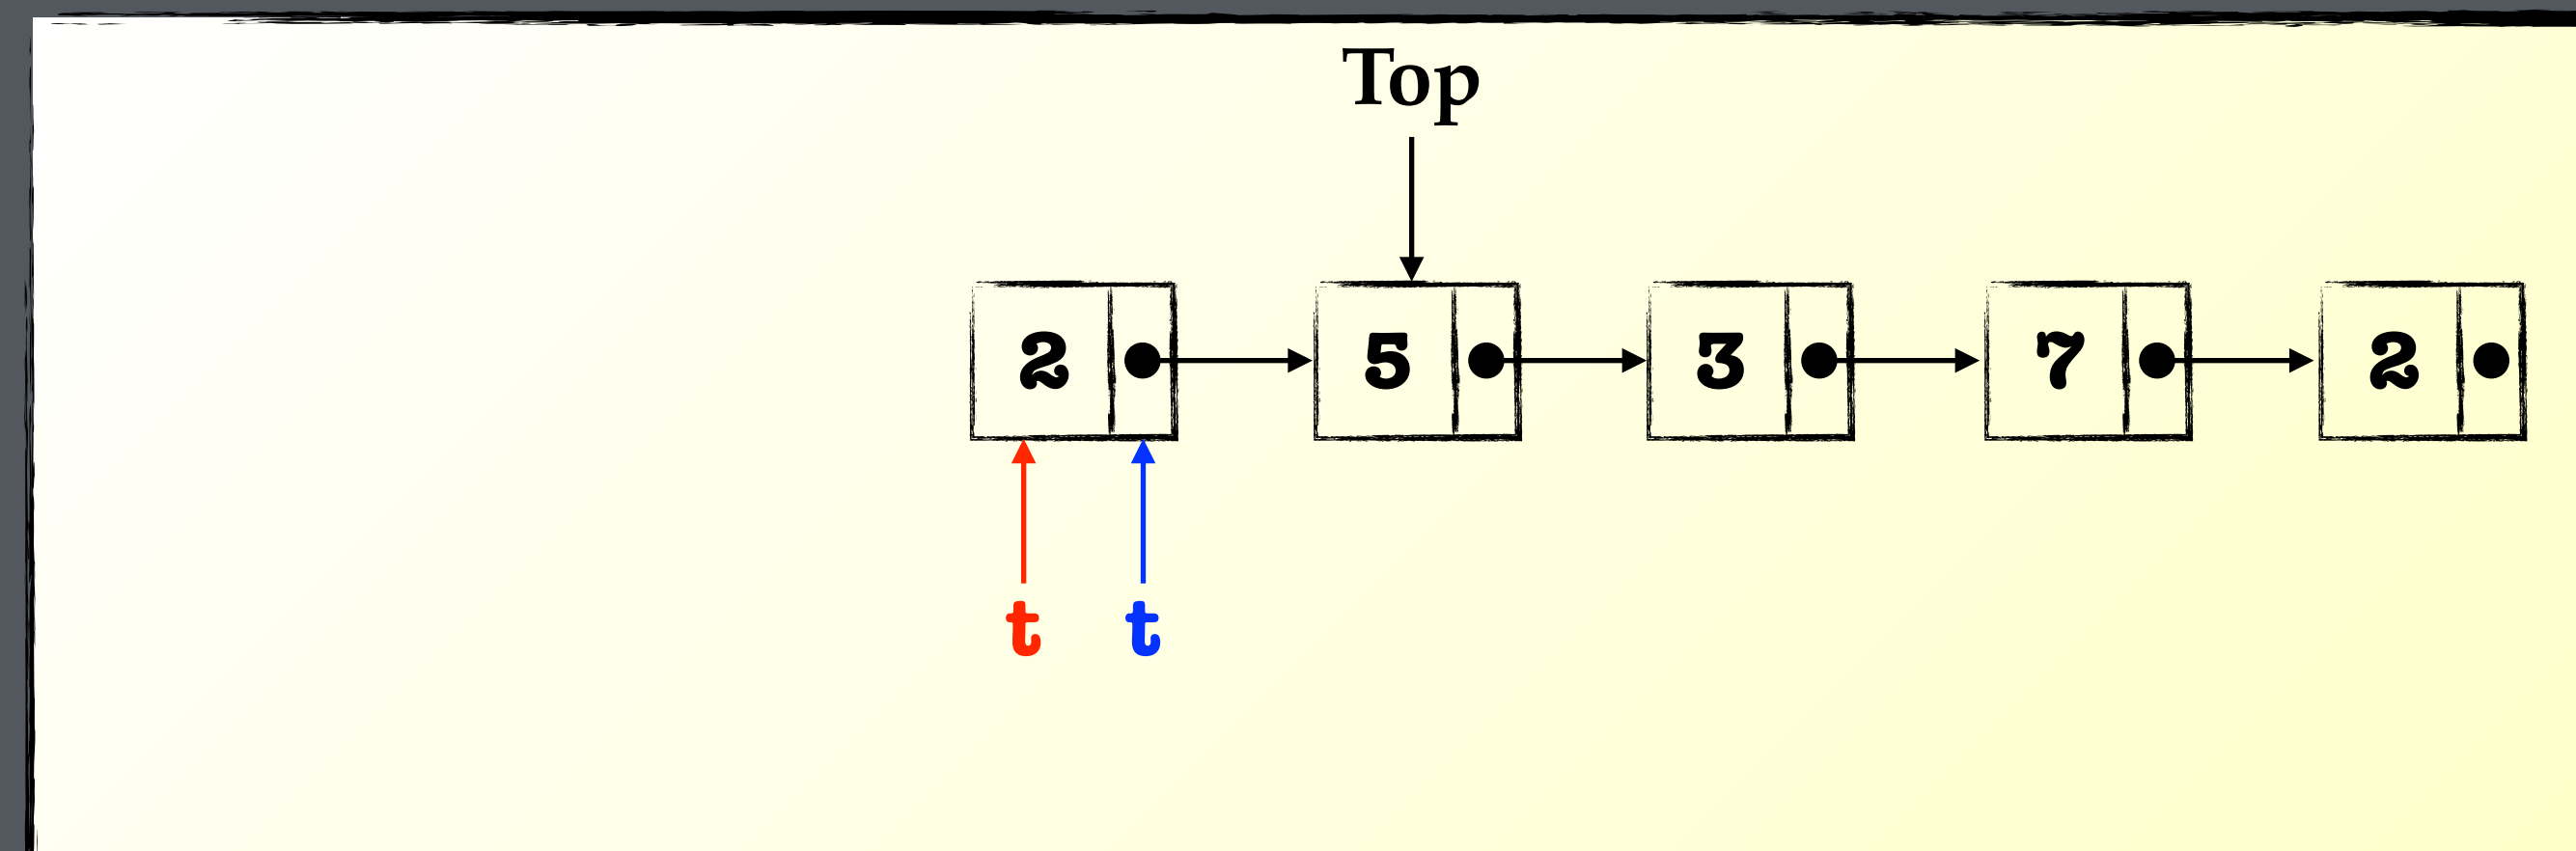


pop

pop




```
pop:
Node t
1 while (true)
2   t = Top
3   if (t = NULL)
4     return *
5     exit
6   if (CAS (Top,t,t.next))
7     return t.val
8   exit
```



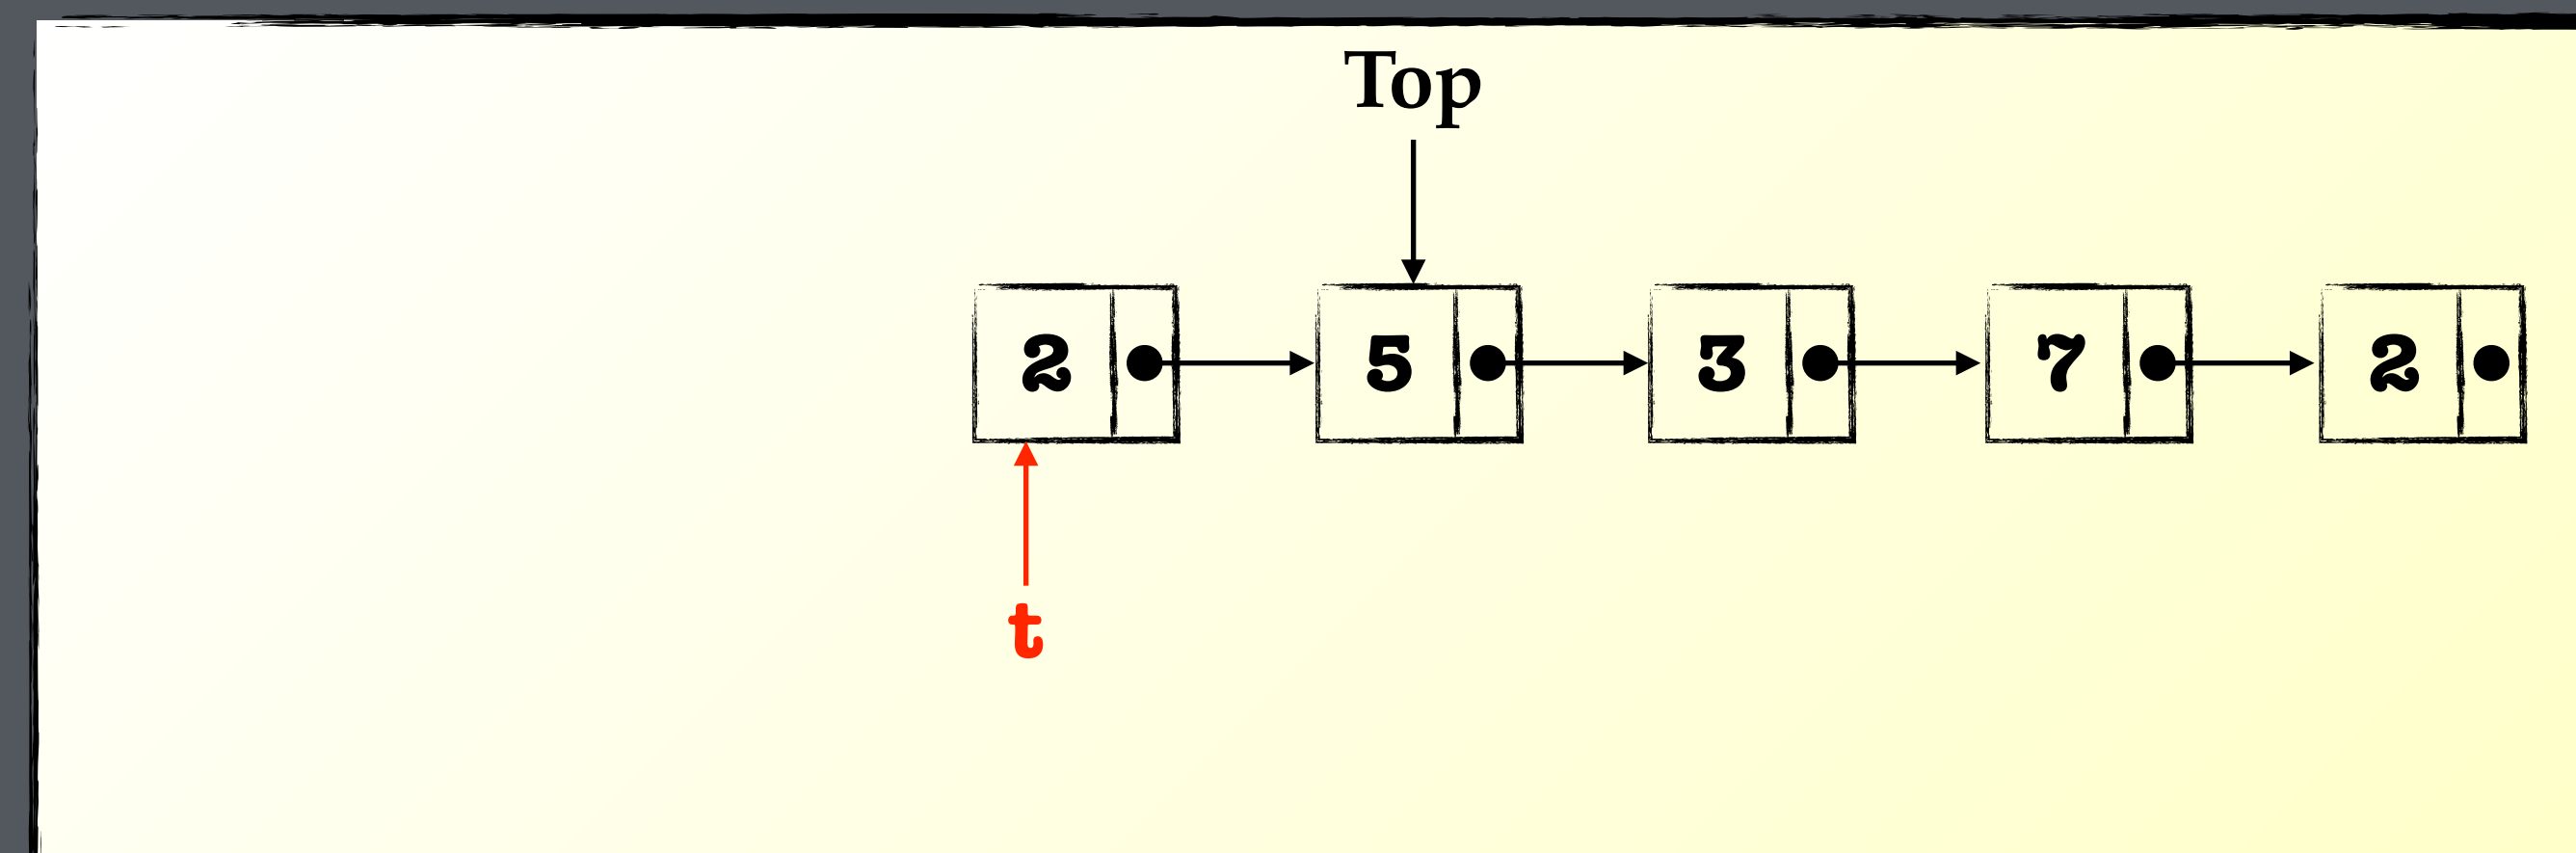
# The Treiber Stack Algorithm

Concurrent pop operations

```
pop:
Node t
1 while (true)
2   t = Top
3   if (t = NULL)
4     return *
5     exit
6   if (CAS (Top, t, t.next))
7     return t.val
8   exit
```



pop






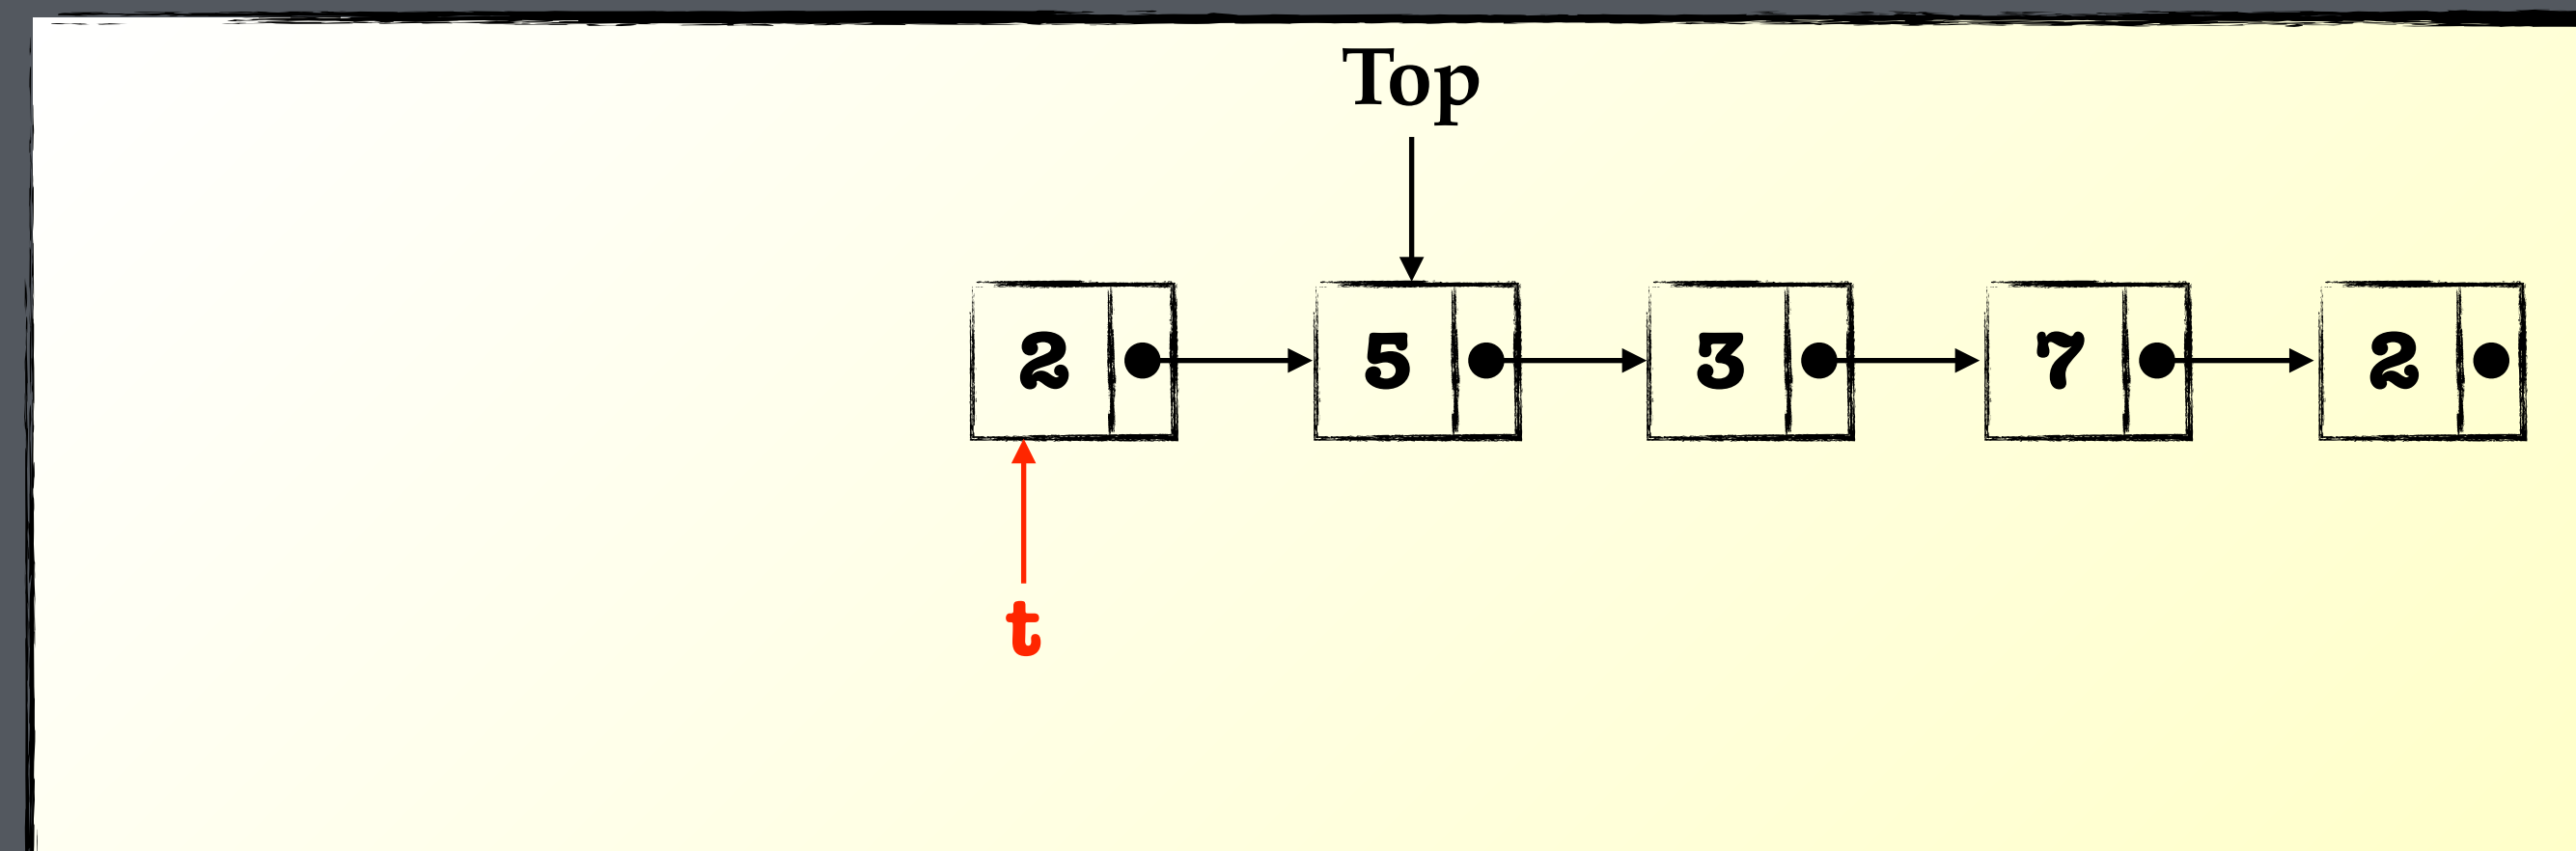
# The Treiber Stack Algorithm

Concurrent pop operations

```
pop:
Node t
1 while (true) ←~~~~~
2   t = Top
3   if (t = NULL)
4     return *
5   exit
6   if (CAS (Top,t,t.next))
7     return t.val
8   exit
```




pop



# The Treiber Stack Algorithm

```
pop:
Node t
1 while (true) ←~~~~~
2   t = Top
3   if (t = NULL)
4     return *
5     exit
6   if (CAS (Top,t,t.next))
7     return t.val
8   exit
```



pop


Top



empty  
stack

# The Treiber Stack Algorithm

```
pop:
Node t
1 while (true)
2   t = Top ←~~~~~
3   if (t = NULL)
4     return *
5     exit
6   if (CAS (Top,t,t.next))
7     return t.val
8   exit
```



pop

Top




t

empty  
stack

# The Treiber Stack Algorithm

```
pop:
Node t
1 while (true)
2   t = Top
3   if (t = NULL) ← wavy line
4     return *
5     exit
6   if (CAS (Top, t, t.next))
7     return t.val
8   exit
```



pop

Top




t

empty  
stack

# The Treiber Stack Algorithm

```
pop:
Node t
1 while (true)
2   t = Top
3   if (t = NULL)
4     return *
5   exit
6   if (CAS (Top,t,t.next))
7     return t.val
8   exit
```



pop

Top




t

empty  
stack



# The Treiber Stack Algorithm

```
pop:
Node t
1 while (true)
2   t = Top
3   if (t = NULL)
4     return *
5     exit ← ~~~~~
6   if (CAS (Top,t,t.next))
7     return t.val
8   exit
```



pop

Top



t

empty  
stack

# The Treiber Stack Algorithm

## Linearization Policy

```
push(k):  
Node t  
1 n = new Node(k,-)  
2 while (true)  
3   t = Top  
4   n.next = t  
5   if (CAS (Top,t,n))  
6     exit
```

when successful

```
pop:  
Node t  
1 while (true)  
2   t = Top  
3   if (t = NULL)  
4     return *  
5   exit  
6   if (CAS (Top,t,t.next))  
7     return t.val  
8   exit
```

when empty

when successful

# The Treiber Stack Algorithm

## Linearization Policy

```
push(k):  
Node t  
1 n = new Node(k,-)  
2 while (true)  
3   t = Top  
4   n.next = t  
5   if (CAS (Top,t,n))  
6     exit
```

when successful

```
pop:  
Node t  
1 while (true)  
2   t = Top  
3   if (t = NULL)  
4     return *  
5   exit  
6   if (CAS (Top,t,t.next))  
7     return t.val  
8   exit
```


when empty

when successful

# The Treiber Stack Algorithm

Linearization Policy




```
push(k):  
Node t  
1 n = new Node(k,-) ←  
2 while (true)  
3   t = Top  
4   n.next = t  
5   if (CAS (Top,t,n))  
6     exit
```



push(4)

push(4)


Top



# The Treiber Stack Algorithm

Linearization Policy






```
push(k):  
Node t  
1 n = new Node(k,-)  
2 while (true)  
3   t = Top ←~~~~~  
4   n.next = t  
5   if (CAS (Top,t,n))  
6     exit
```



push(4)

push(4)

Top






# The Treiber Stack Algorithm

Linearization Policy


```
push(k):  
Node t  
1 n = new Node(k,-)  
2 while (true)  
3   t = Top  
4   n.next = t  
5   if (CAS (Top,t,n))  
6     exit
```



push(4)

push(4)


Top



# The Treiber Stack Algorithm


## Linearization Policy

```
push(k):  
Node t  
1 n = new Node(k,-)  
2 while (true)  
3   t = Top  
4   n.next = t  
5   if (CAS (Top,t,n))  
6     exit
```

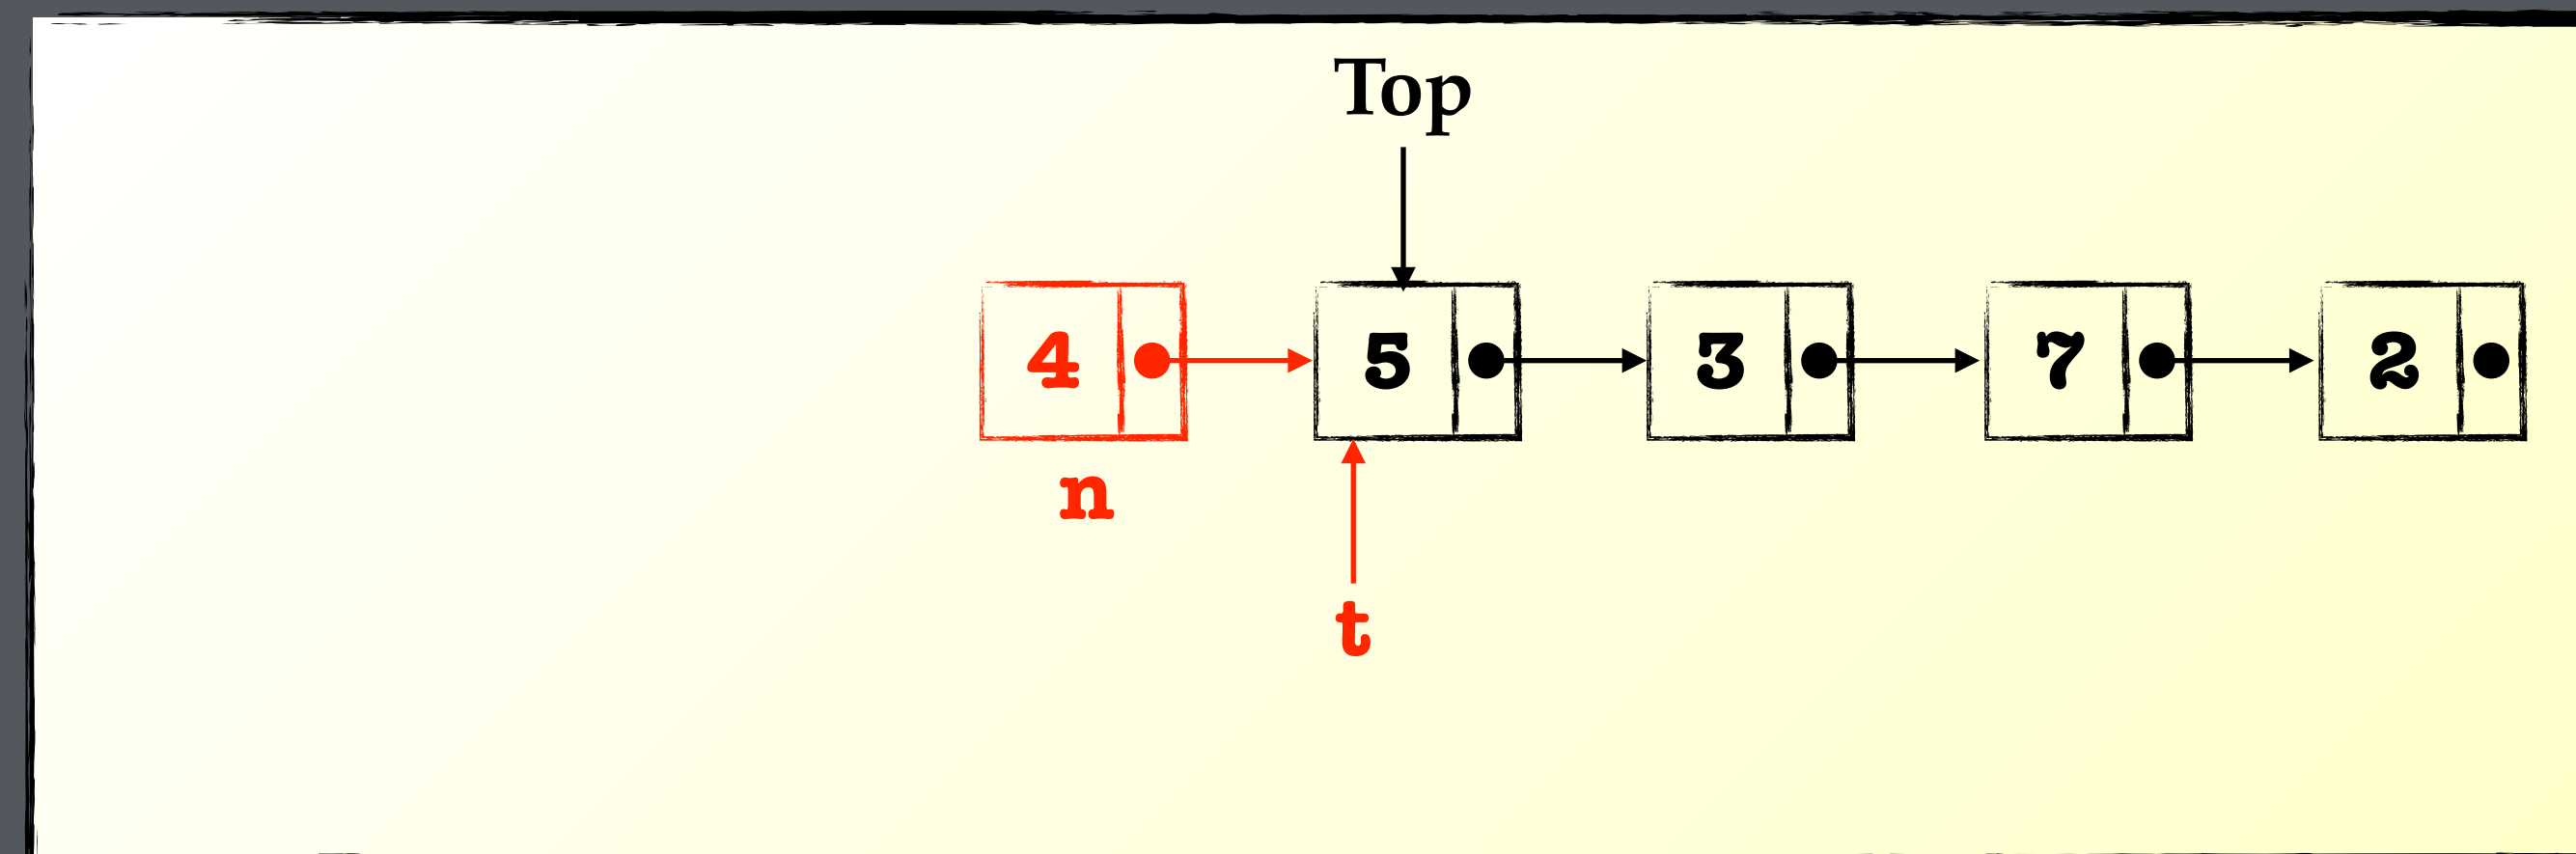



push(4)

pop




```
pop:  
Node t  
1 while (true)  
2   t = Top  
3   if (t = NULL)  
4     return *  
5   exit  
6   if (CAS (Top,t,t.next))  
7     return t.val  
8   exit
```



# The Treiber Stack Algorithm


## Linearization Policy

```
push(k):  
Node t  
1 n = new Node(k,-)  
2 while (true)  
3   t = Top  
4   n.next = t  
5   if (CAS (Top,t,n))  
6     exit
```

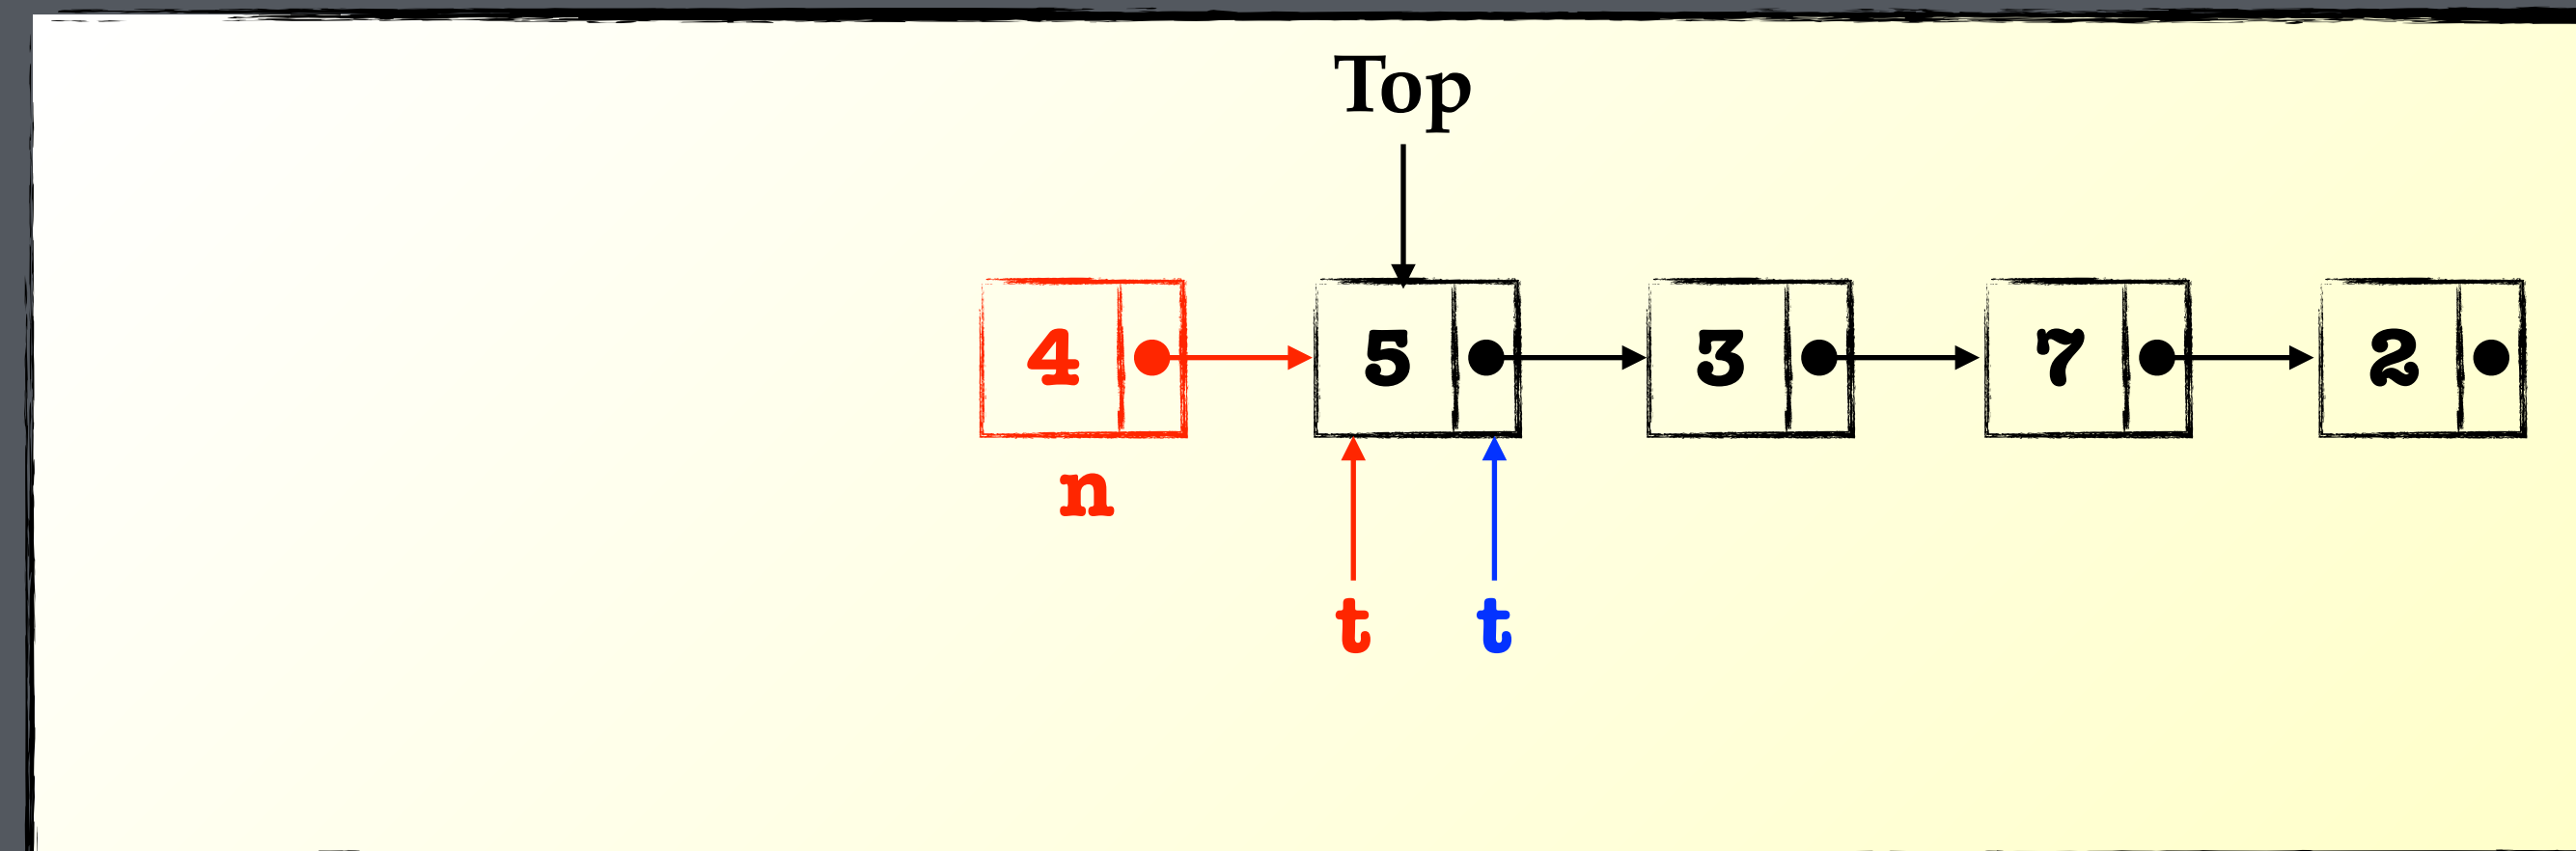



push(4)

pop




```
pop:  
Node t  
1 while (true)  
2   t = Top  
3   if (t = NULL)  
4     return *  
5   exit  
6   if (CAS (Top,t,t.next))  
7     return t.val  
8   exit
```



# The Treiber Stack Algorithm


## Linearization Policy

```
push(k):  
Node t  
1 n = new Node(k,-)  
2 while (true)  
3   t = Top  
4   n.next = t  
5   if (CAS (Top,t,n))  
6     exit
```

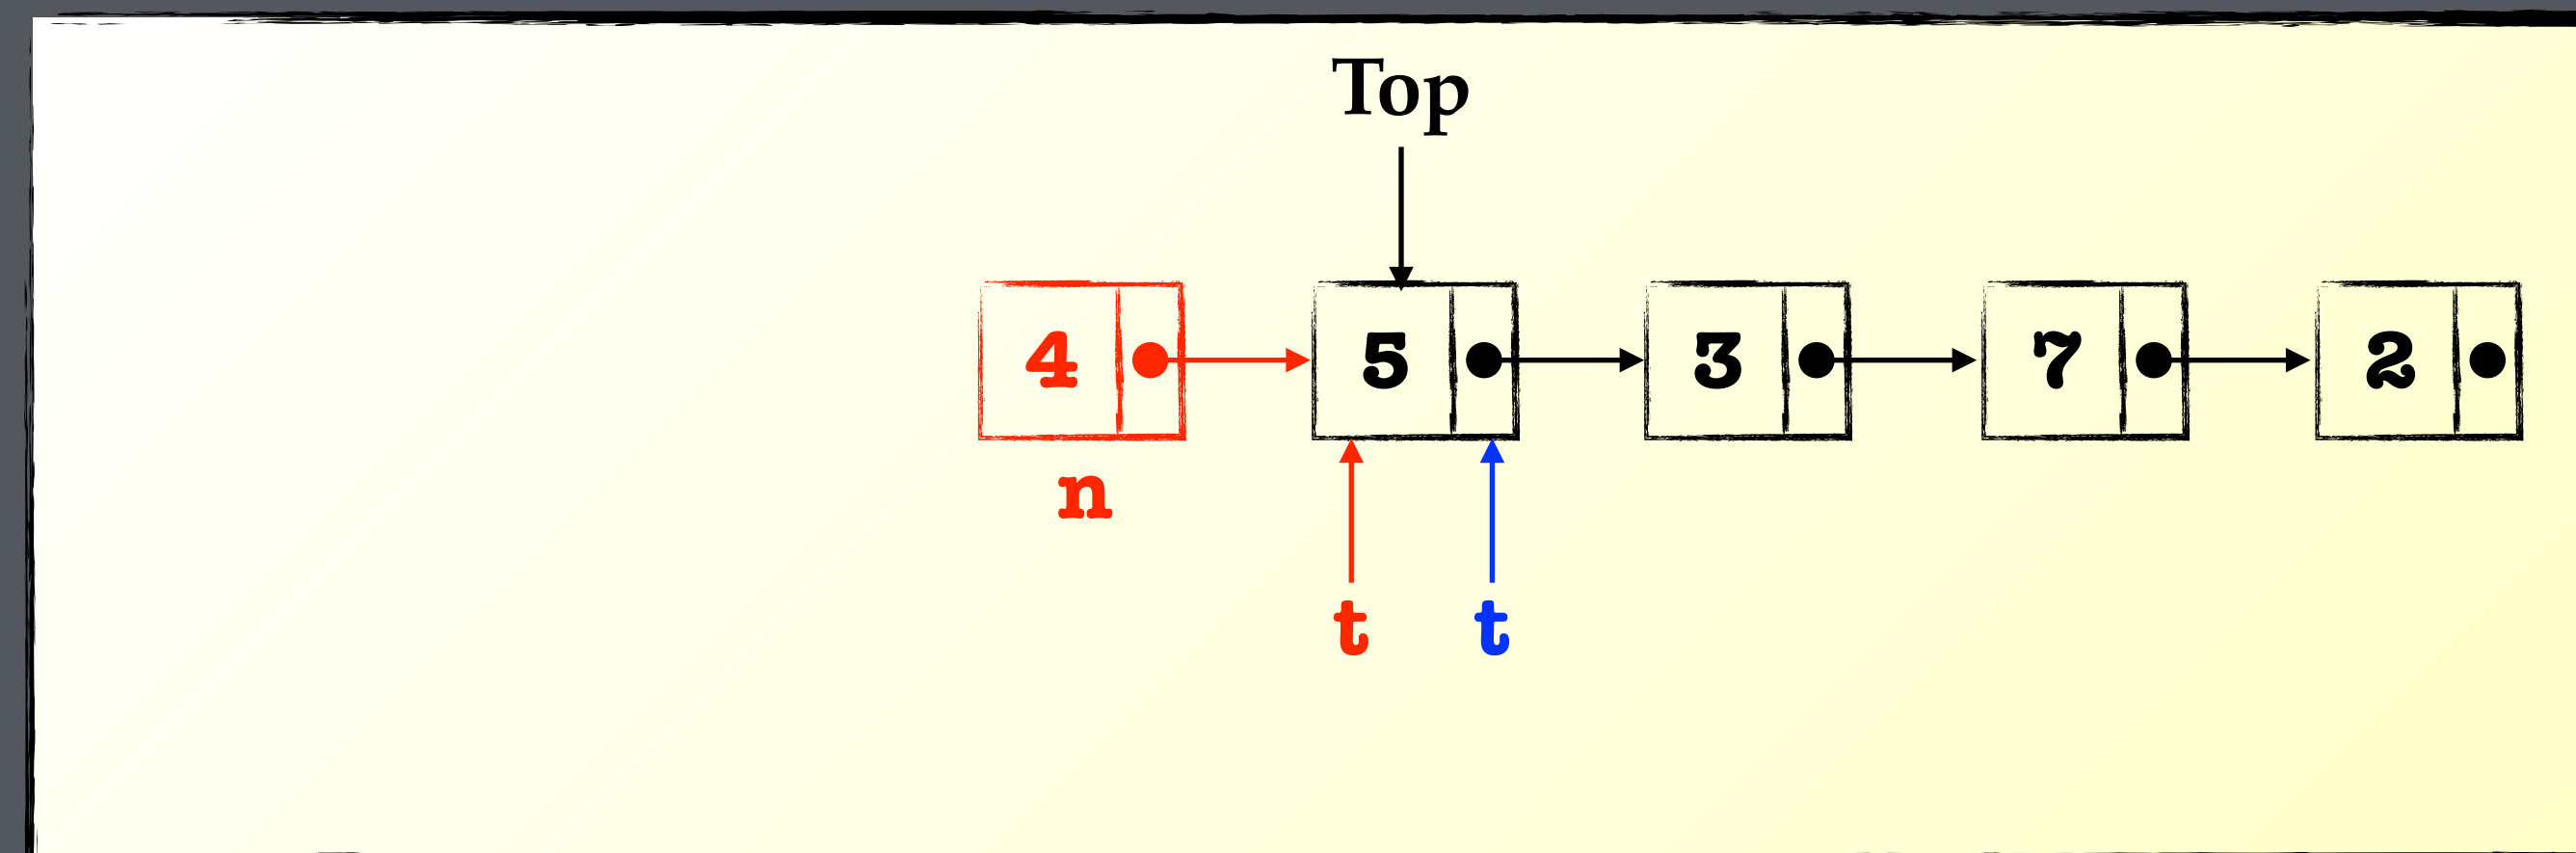



push(4)

pop



```
pop:  
Node t  
1 while (true)  
2   t = Top  
3   if (t = NULL)  
4     return *  
5   exit  
6   if (CAS (Top,t,t.next))  
7     return t.val  
8   exit
```





# The Treiber Stack Algorithm



## Linearization Policy

```
push(k):  
Node t  
1 n = new Node(k,-)  
2 while (true)  
3   t = Top  
4   n.next = t  
5   if (CAS (Top,t,n))  
6     exit
```

**push(4)**

**pop**

```
pop:  
Node t  
1 while (true)  
2   t = Top  
3   if (t = NULL)  
4     return *  
5   exit  
6   if (CAS (Top,t,t.next))  
7     return t.val  
8   exit
```





# The Treiber Stack Algorithm



## Linearization Policy

```
push(k):  
Node t  
1 n = new Node(k,-)  
2 while (true)  
3   t = Top  
4   n.next = t  
5   if (CAS (Top,t,n))  
6     exit
```

**push(4)**

**pop**

```
pop:  
Node t  
1 while (true)  
2   t = Top  
3   if (t = NULL)  
4     return *  
5   exit  
6   if (CAS (Top,t,t.next))  
7     return t.val  
8   exit
```



# The Treiber Stack Algorithm


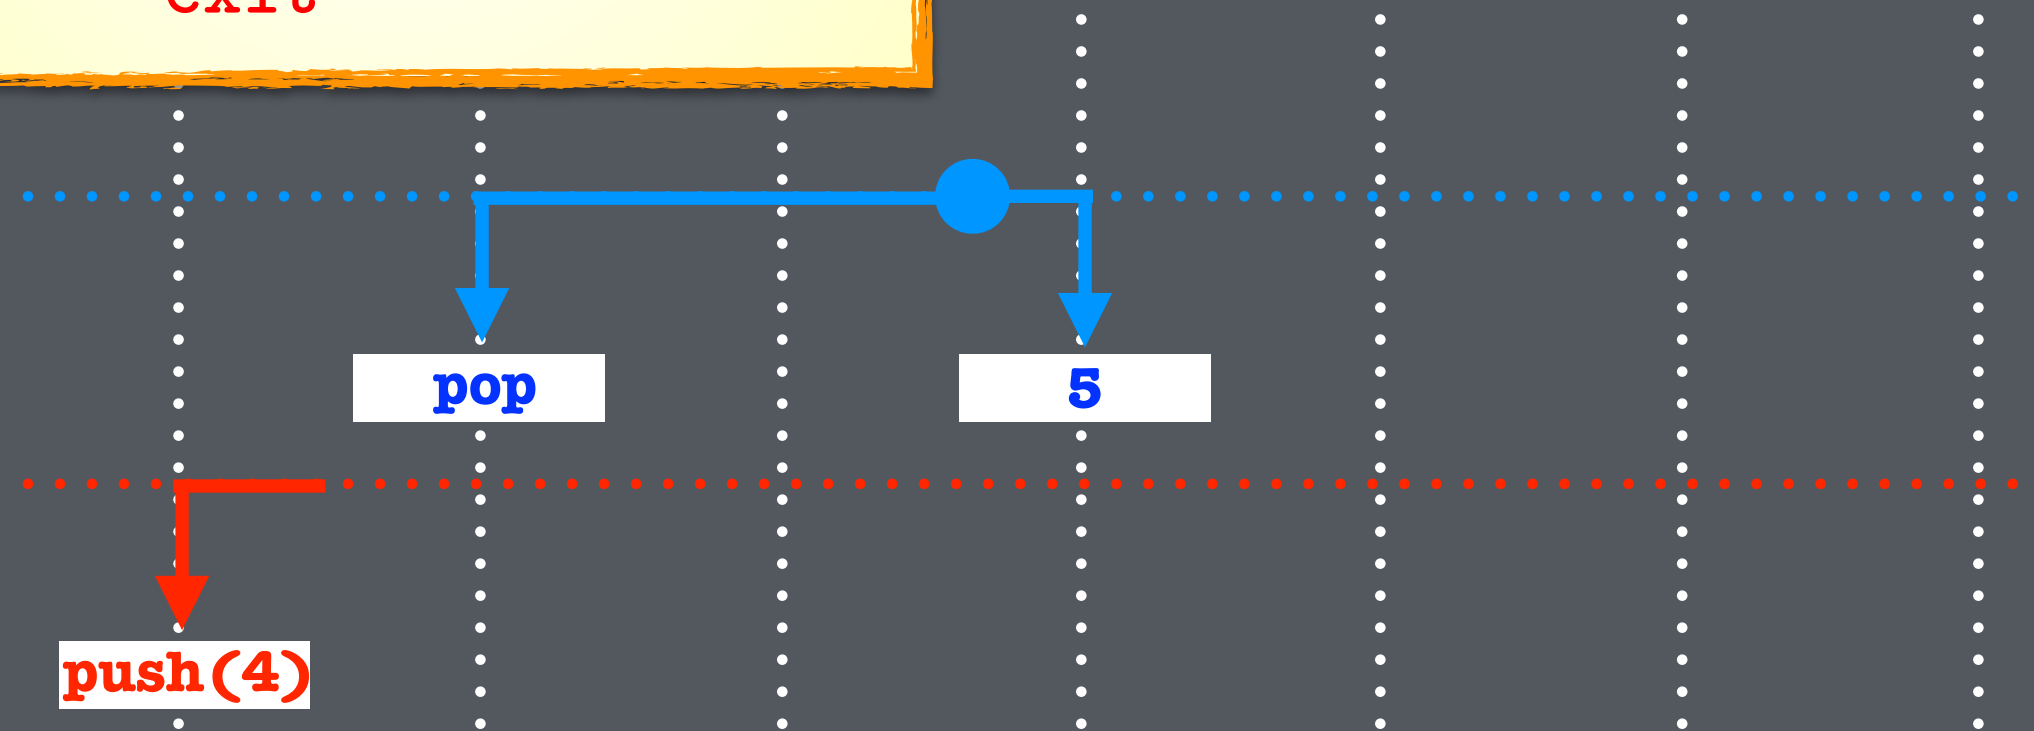
## Linearization Policy

```
push(k):  
Node t  
1 n = new Node(k,-)  
2 while (true)  
3   t = Top  
4   n.next = t  
5   if (CAS (Top,t,n))  
6     exit
```

**push(4)**

**pop**


```
pop:  
Node t  
1 while (true)  
2   t = Top  
3   if (t = NULL)  
4     return *  
5   exit  
6   if (CAS (Top,t,t.next))  
7     return t.val  
8   exit
```



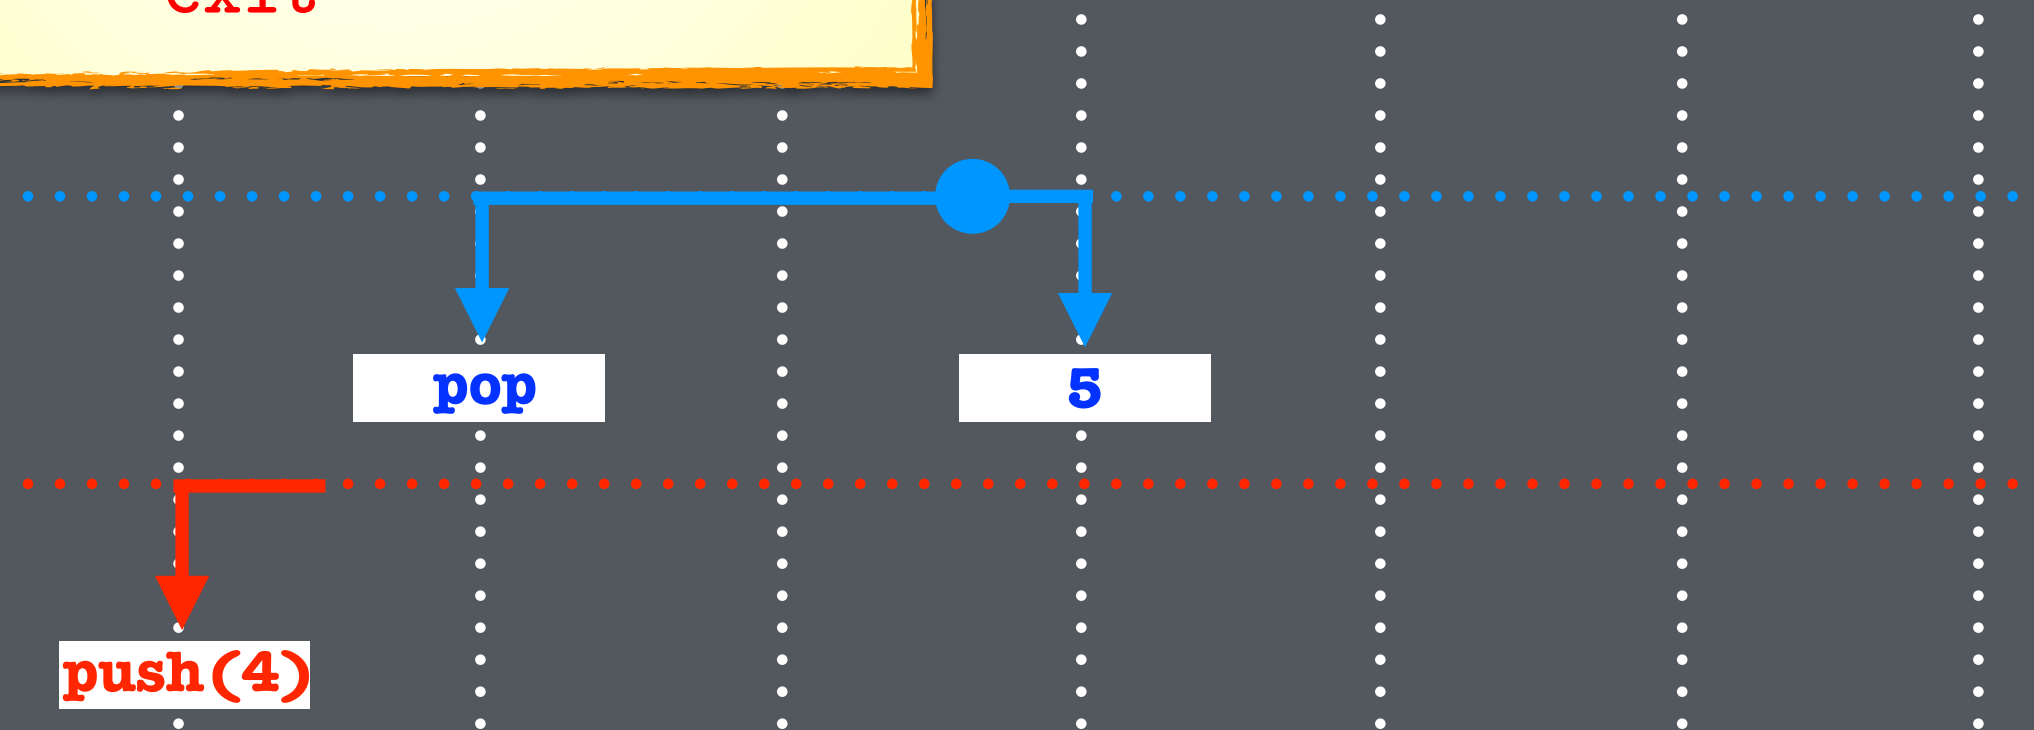
# The Treiber Stack Algorithm

## Linearization Policy


```
push(k):  
Node t  
1 n = new Node(k,-)  
2 while (true)  
3   t = Top  
4   n.next = t  
5   if (CAS (Top,t,n))  
6     exit
```



**push(4)**



```
pop:  
Node t  
1 while (true)  
2   t = Top  
3   if (t = NULL)  
4     return *  
5   exit  
6   if (CAS (Top,t,t.next))  
7     return t.val  
8   exit
```





# The Treiber Stack Algorithm

## Linearization Policy

push(k):

Node t

1 n = new Node(k,-)


2 while (true)

3 t = Top


4 n.next = t

5 if (CAS (Top,t,n)) ← wavy arrow

6 exit



push(4)



pop:

Node t

1 while (true)

2 t = Top

3 if (t = NULL)

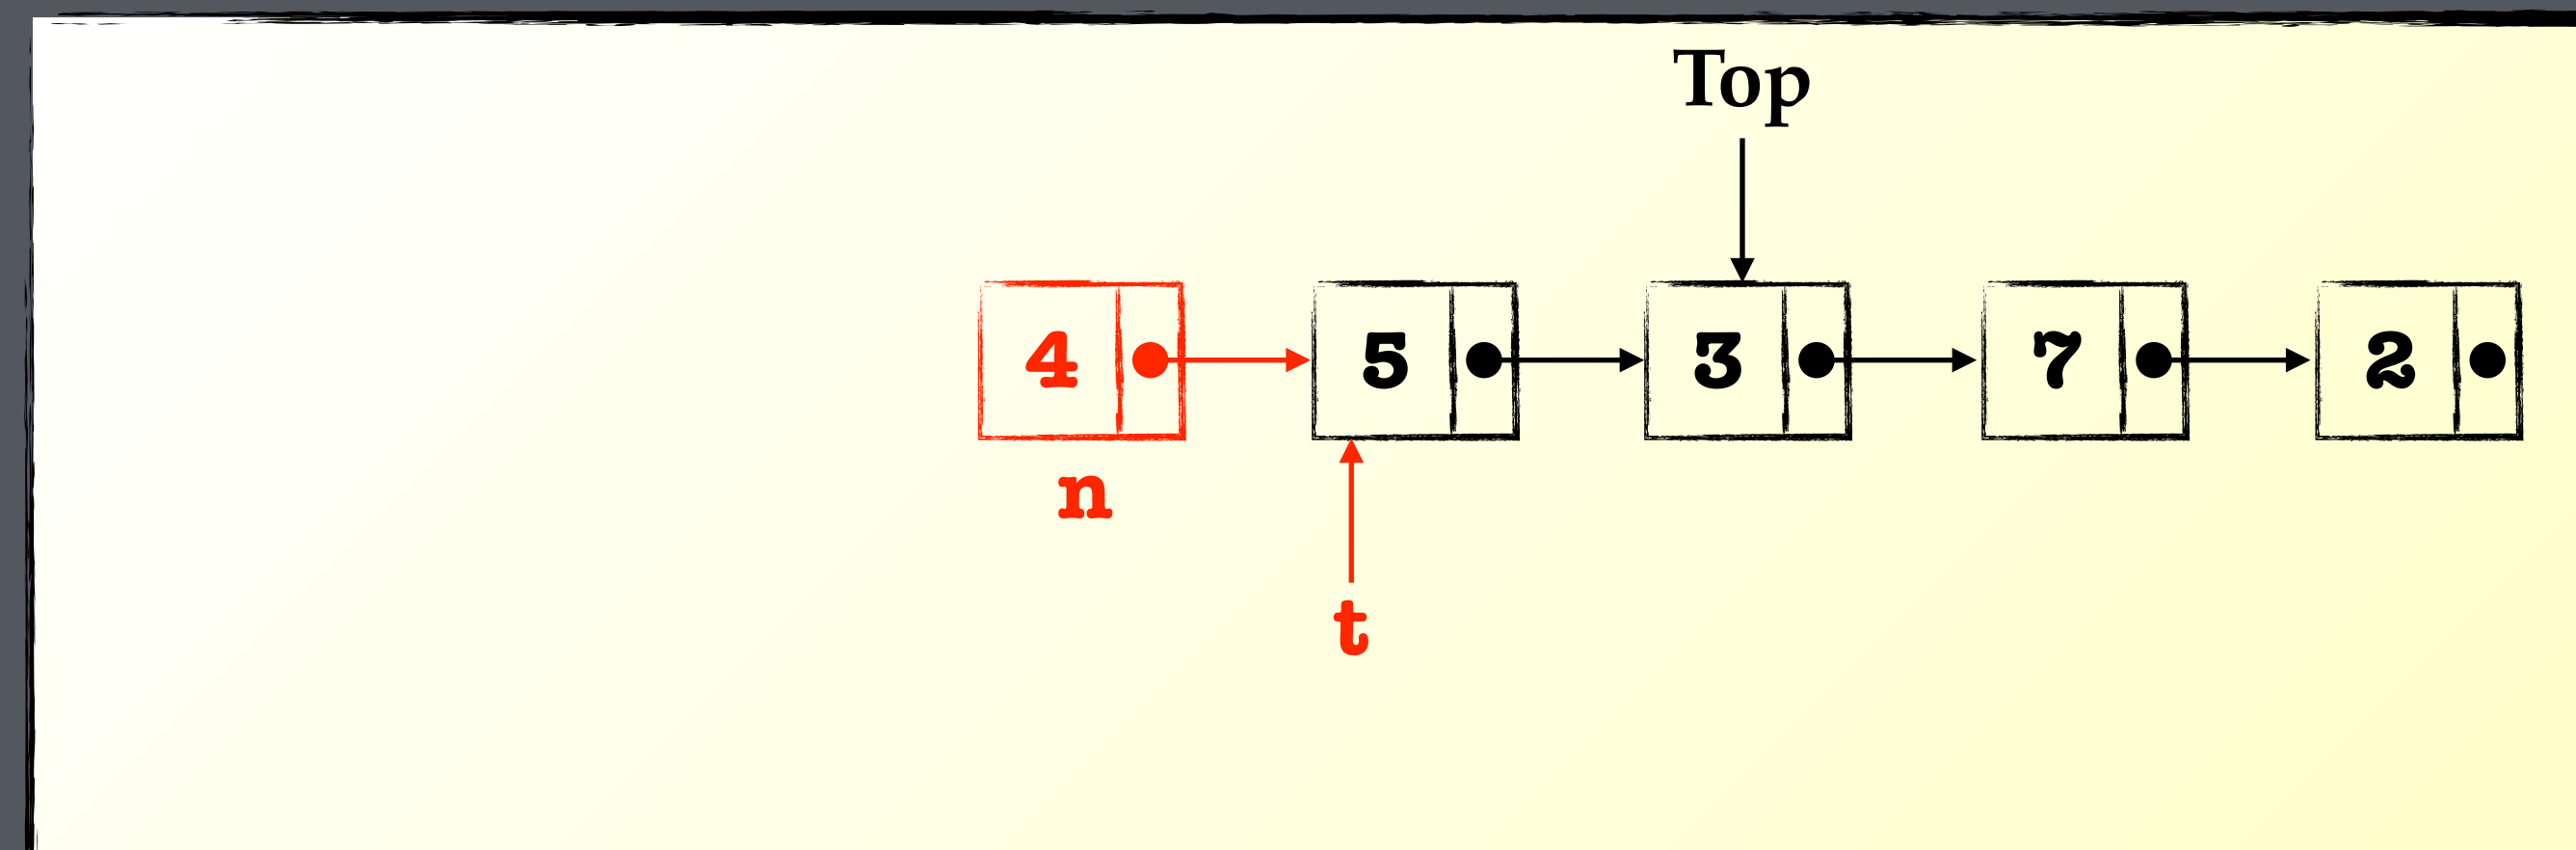
4 return \*

5 exit

6 if (CAS (Top,t,t.next))

7 return t.val

8 exit



# The Treiber Stack Algorithm

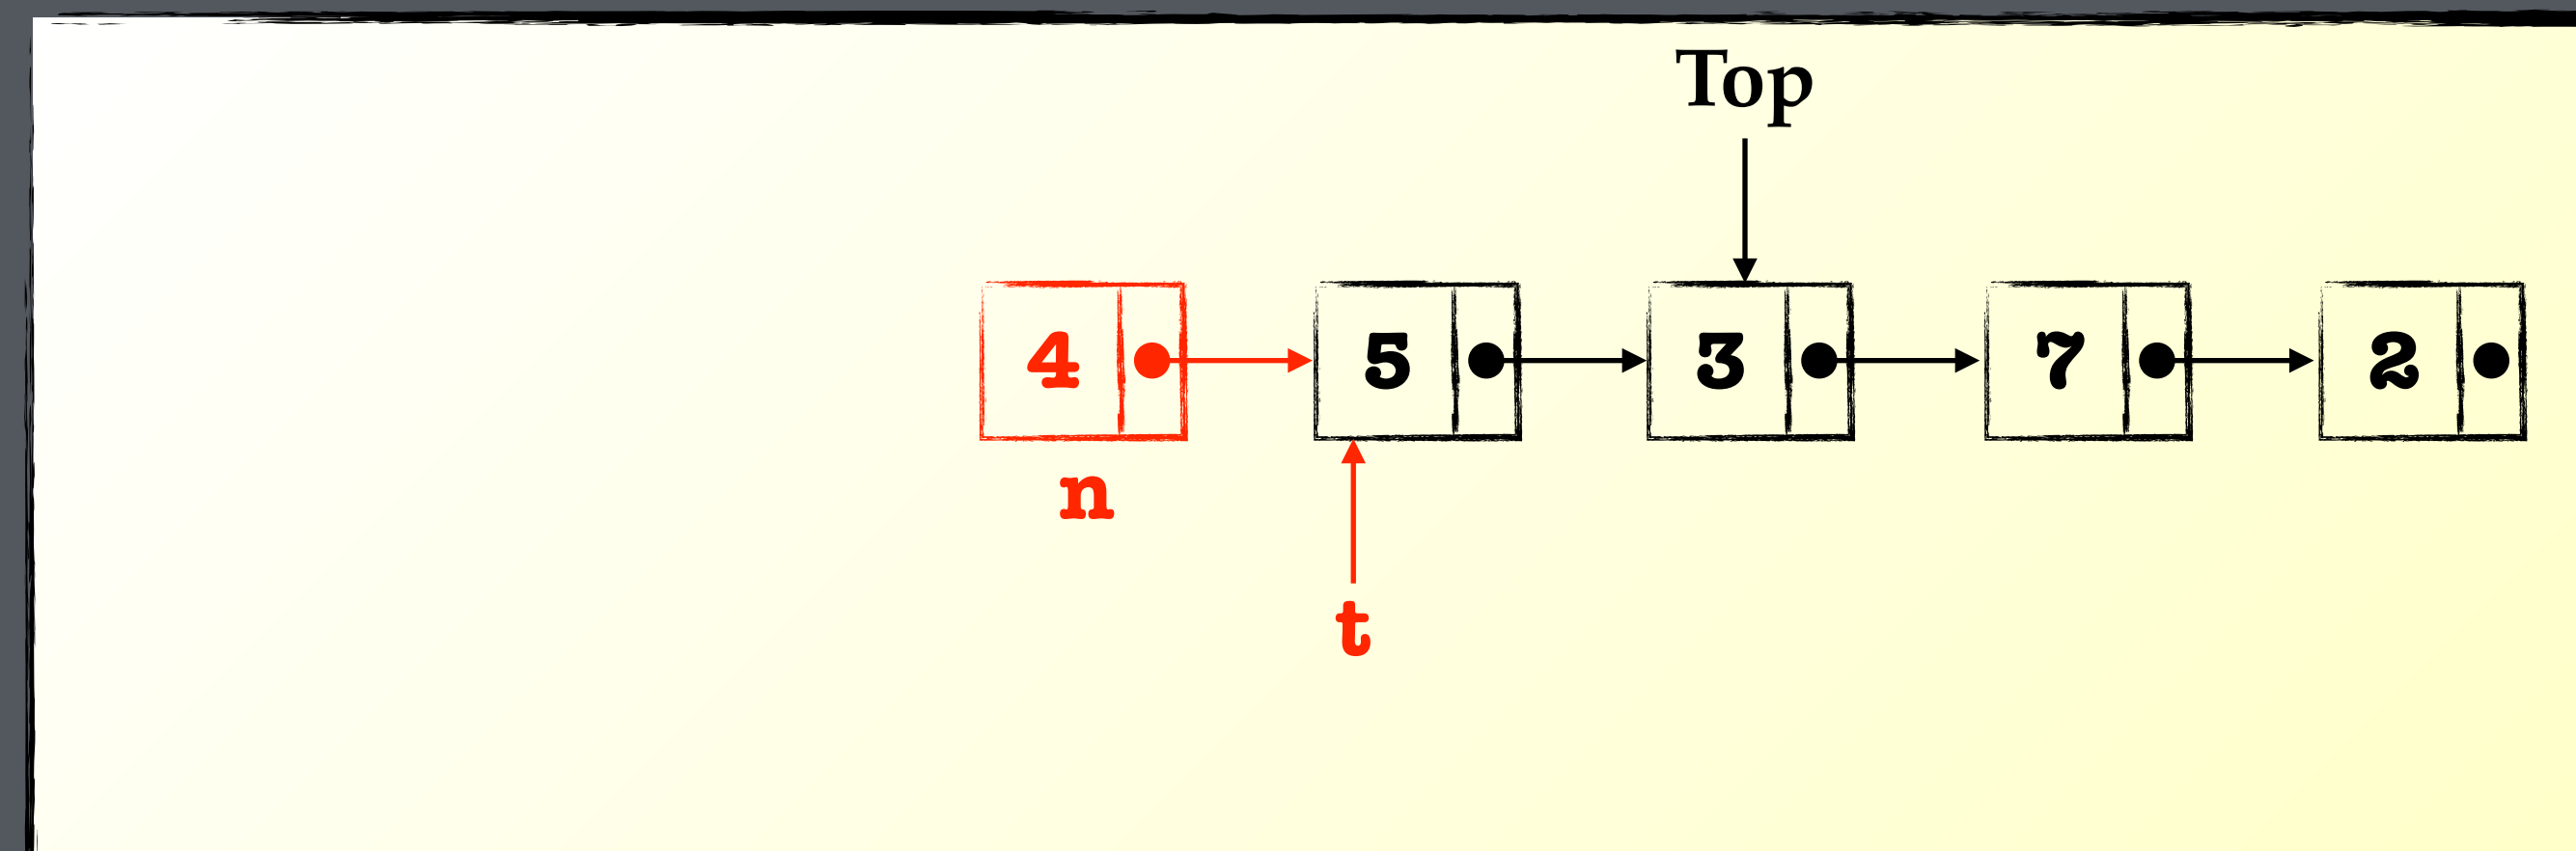
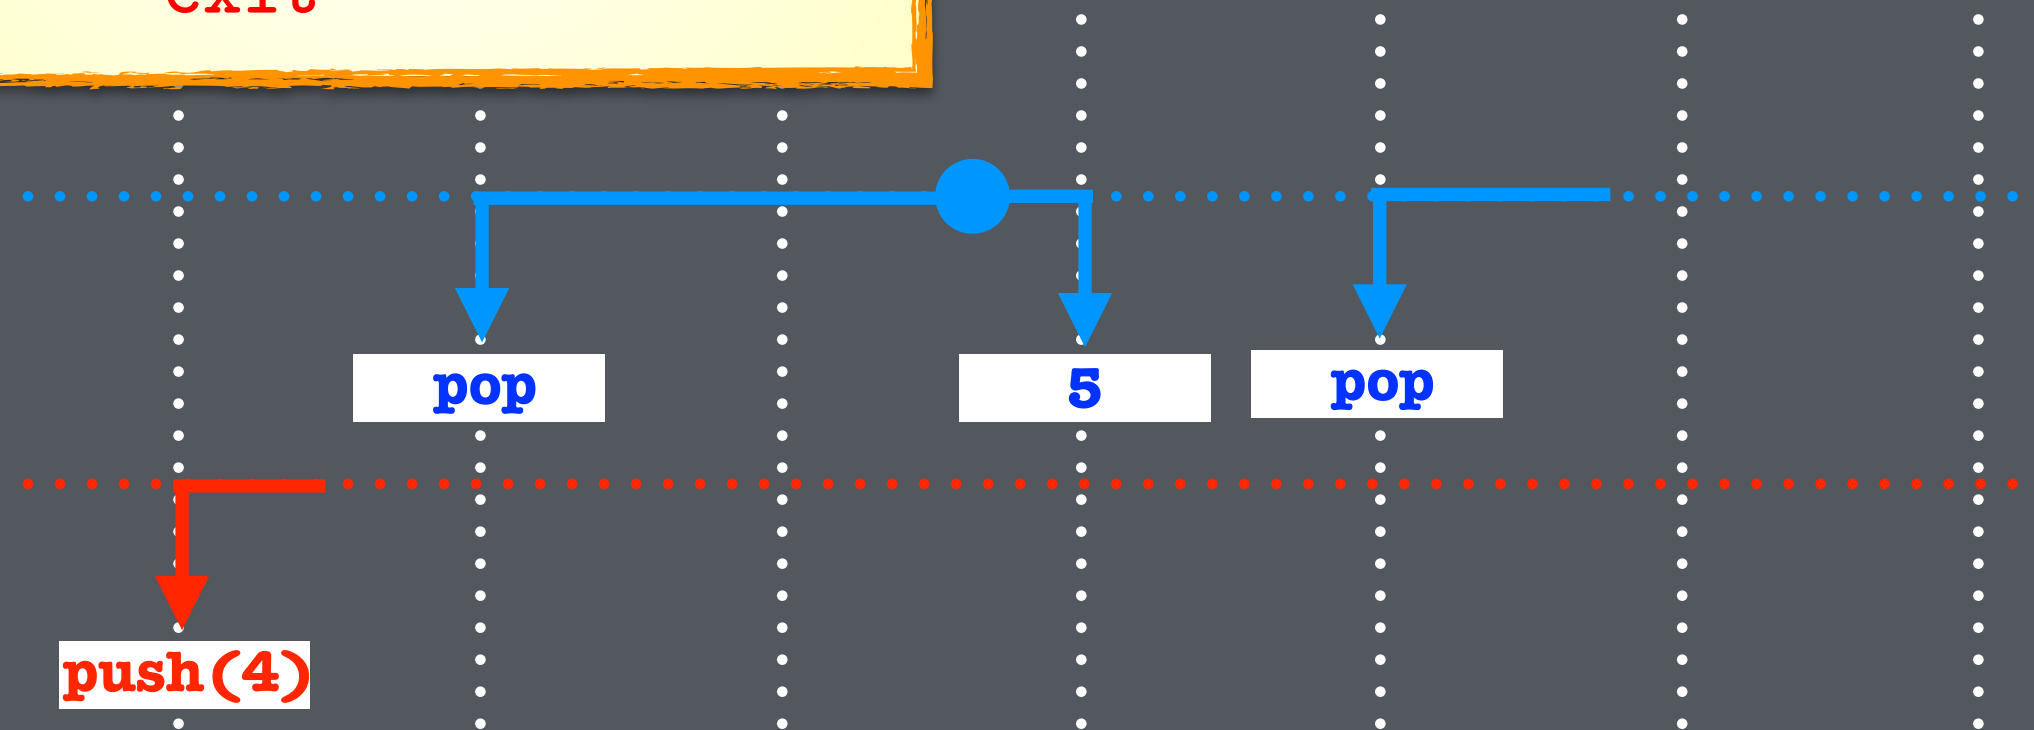
## Linearization Policy

```
push(k):  
Node t  
1 n = new Node(k,-)  
2 while (true) ←~~~~~  
3   t = Top  
4   n.next = t  
5   if (CAS (Top,t,n))  
6     exit
```

**push(4)**

**pop**

```
pop:  
Node t  
1 while (true) ←~~~~~  
2   t = Top  
3   if (t = NULL)  
4     return *  
5   exit  
6   if (CAS (Top,t,t.next))  
7     return t.val  
8   exit
```






# The Treiber Stack Algorithm


## Linearization Policy

```
push(k):  
Node t  
1 n = new Node(k,-)  
2 while (true)  
3   t = Top  
4   n.next = t  
5   if (CAS (Top,t,n))  
6     exit
```

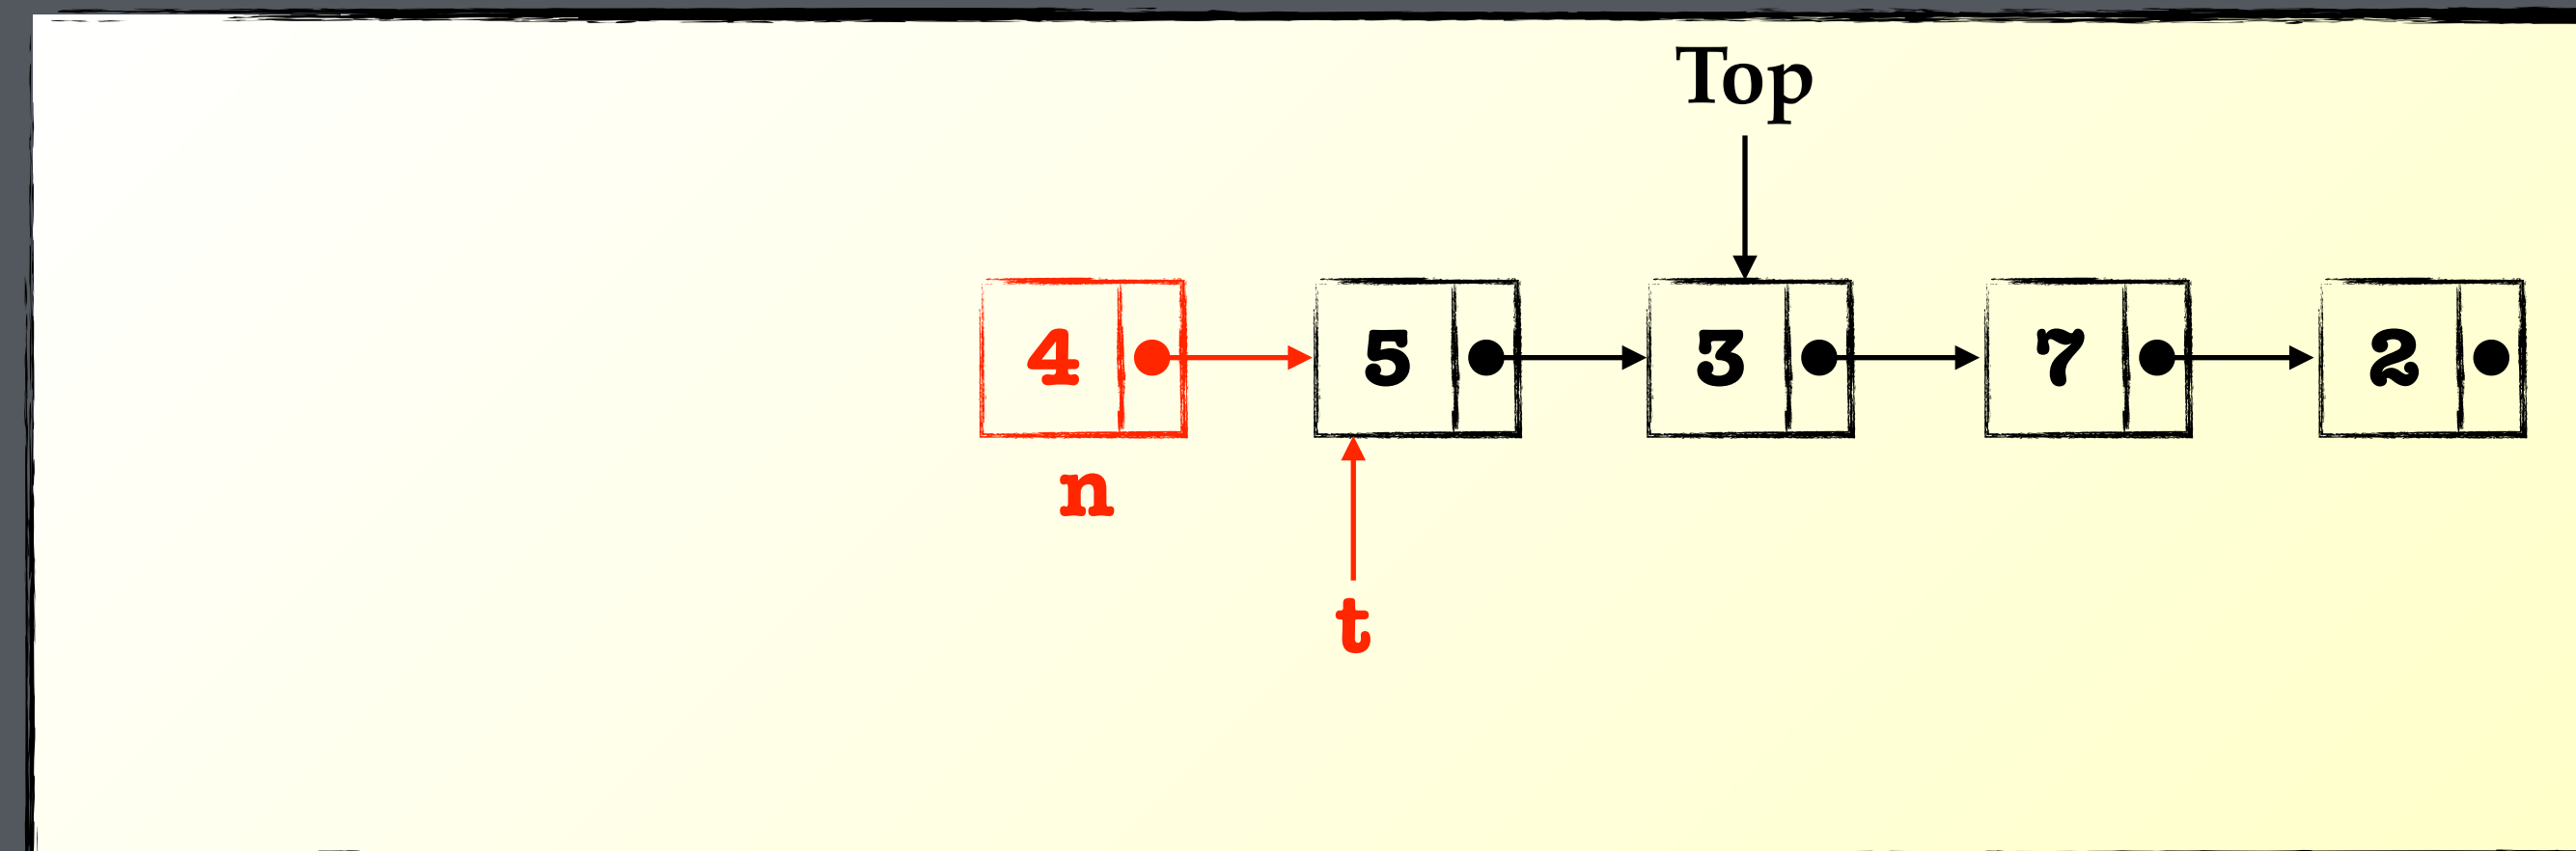
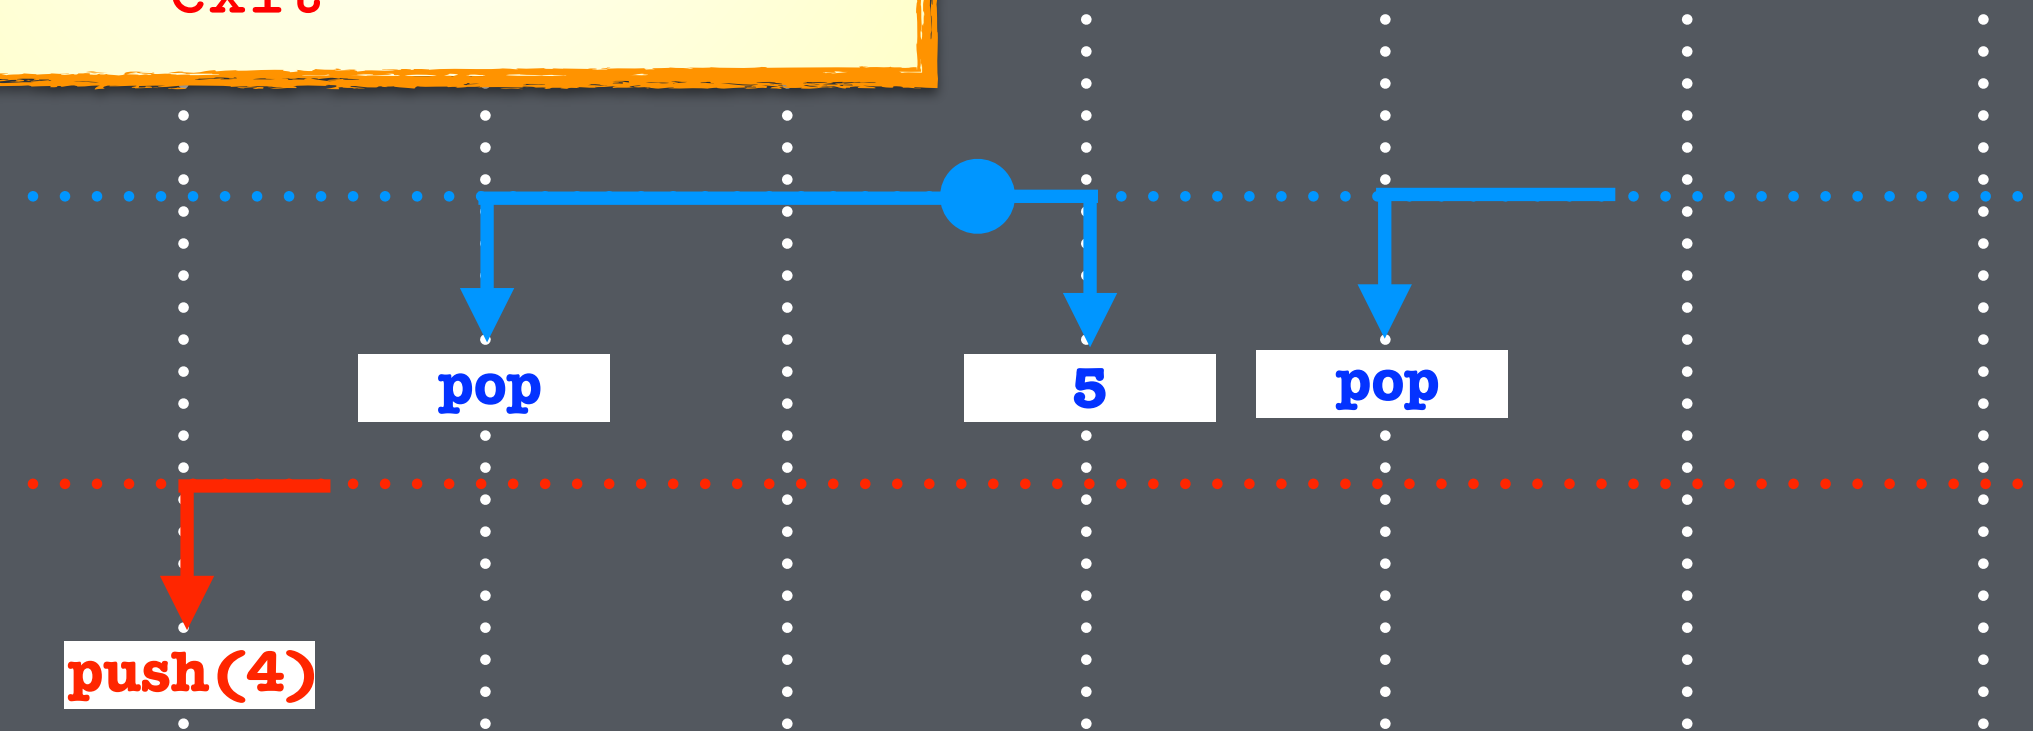


**push(4)**

**pop**




```
pop:  
Node t  
1 while (true)  
2   t = Top  
3   if (t = NULL)  
4     return *  
5   exit  
6   if (CAS (Top,t,t.next))  
7     return t.val  
8   exit
```



# The Treiber Stack Algorithm


## Linearization Policy

```
push(k):  
Node t  
1 n = new Node(k,-)  
2 while (true)  
3   t = Top  
4   n.next = t  
5   if (CAS (Top,t,n))  
6     exit
```

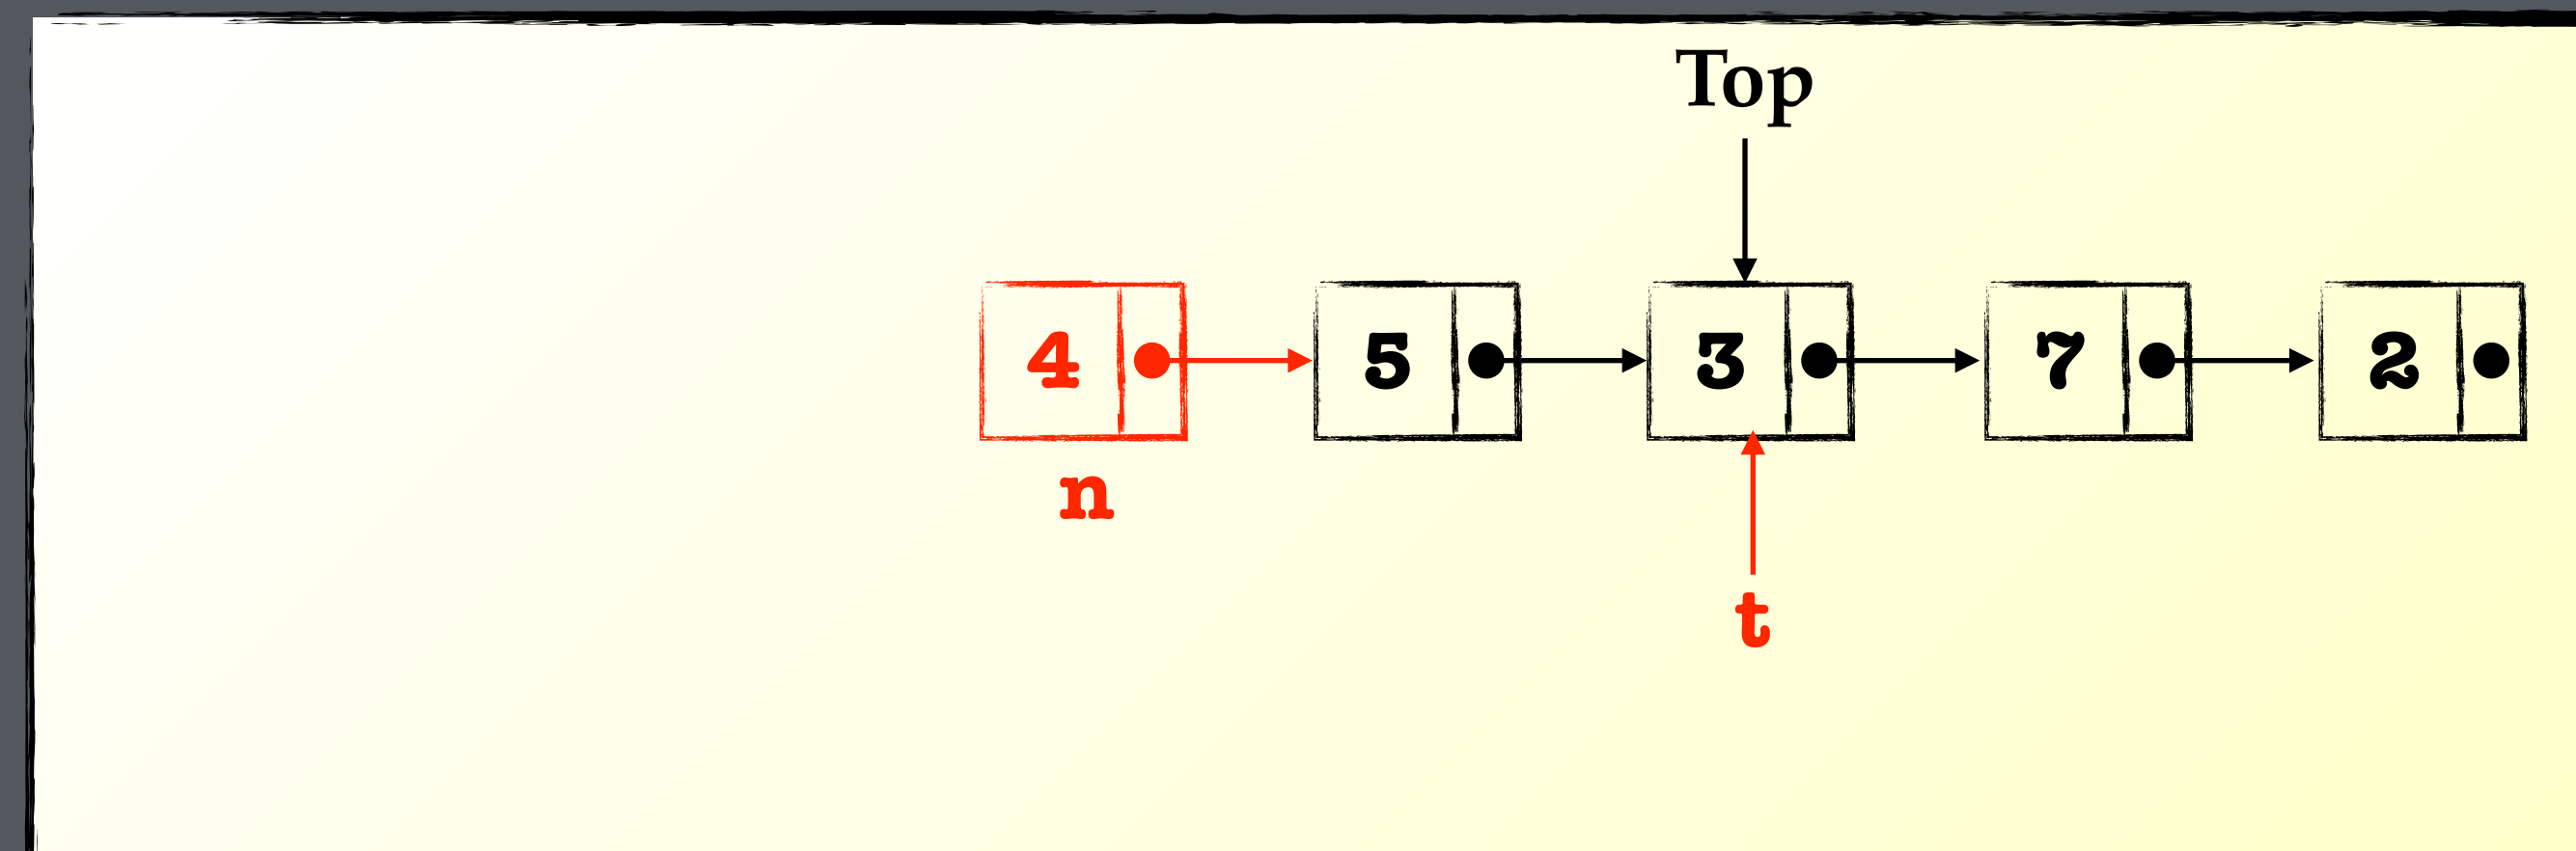



**push(4)**

**pop**




```
pop:  
Node t  
1 while (true)  
2   t = Top  
3   if (t = NULL)  
4     return *  
5   exit  
6   if (CAS (Top,t,t.next))  
7     return t.val  
8   exit
```



# The Treiber Stack Algorithm


## Linearization Policy

```
push(k):  
Node t  
1 n = new Node(k,-)  
2 while (true)  
3   t = Top  
4   n.next = t  
5   if (CAS (Top,t,n))  
6     exit
```

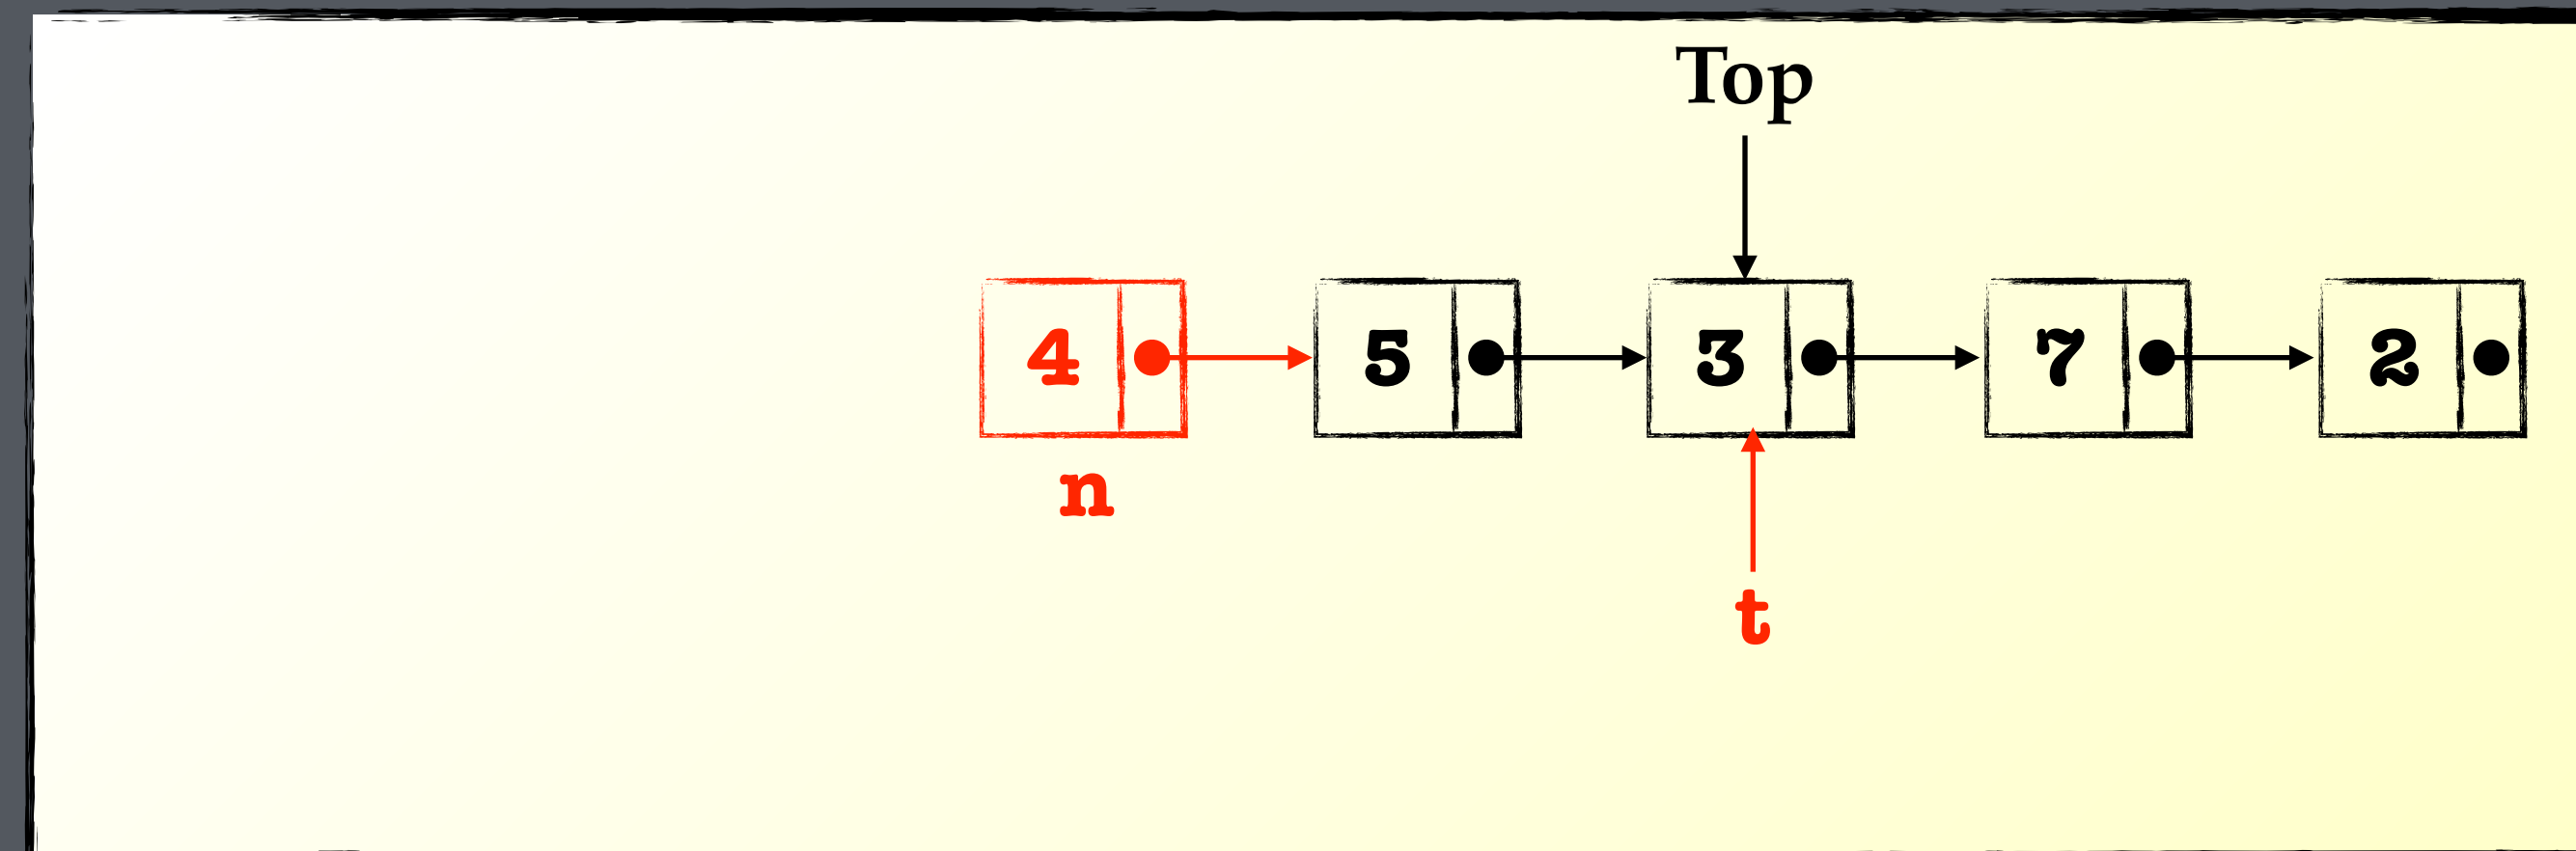
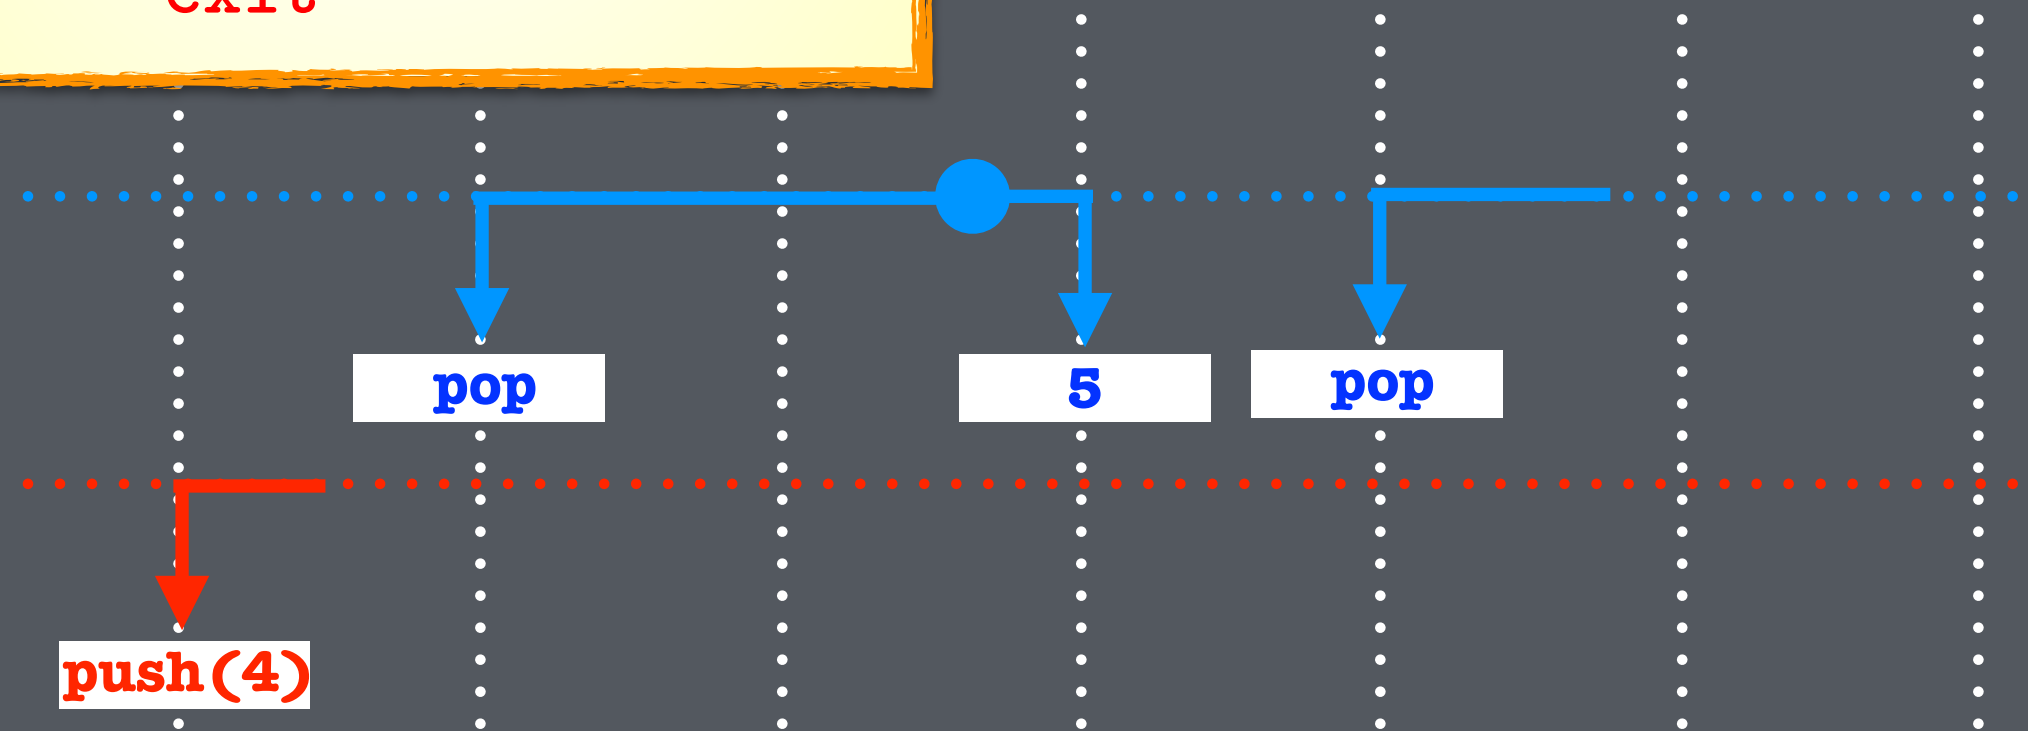


**push(4)**

**pop**



```
pop:  
Node t  
1 while (true)  
2   t = Top  
3   if (t = NULL)  
4     return *  
5   exit  
6   if (CAS (Top,t,t.next))  
7     return t.val  
8   exit
```






# The Treiber Stack Algorithm


## Linearization Policy

```
push(k):  
Node t  
1 n = new Node(k,-)  
2 while (true)  
3   t = Top  
4   n.next = t  
5   if (CAS (Top,t,n))  
6     exit
```


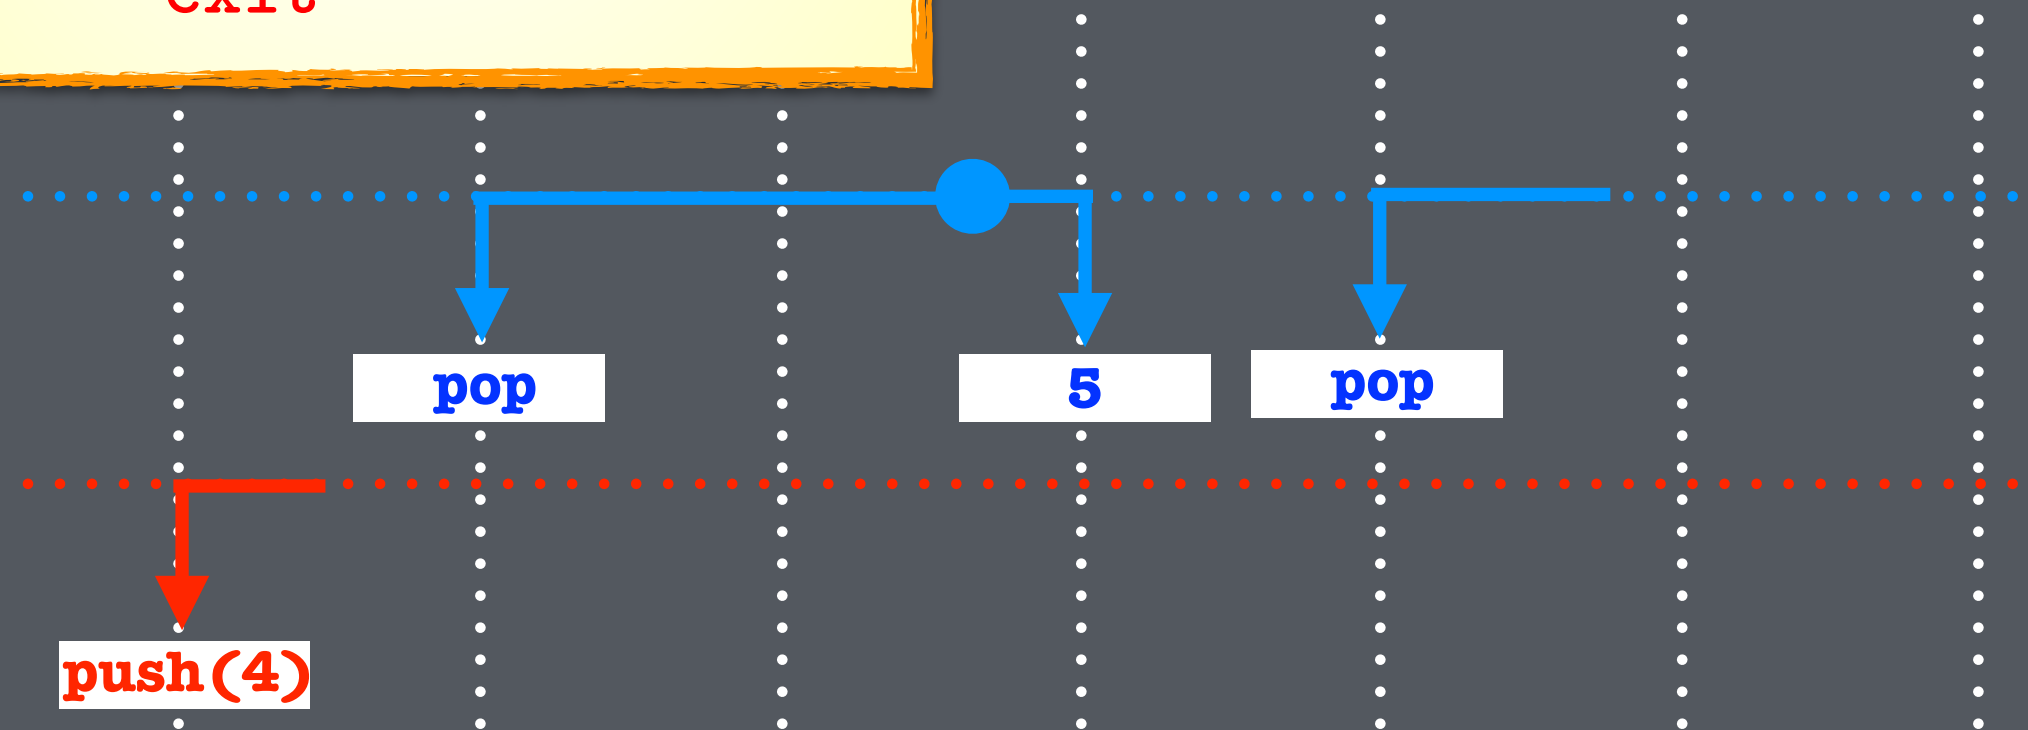


push(4)

pop



```
pop:  
Node t  
1 while (true)  
2   t = Top  
3   if (t = NULL)  
4     return *  
5   exit  
6   if (CAS (Top,t,t.next))  
7     return t.val  
8   exit
```



# The Treiber Stack Algorithm

## Linearization Policy

push(k):

Node t

```
1 n = new Node(k,-)
2 while (true)
3   t = Top
4   n.next = t
5   if (CAS (Top,t,n))
6     exit
```


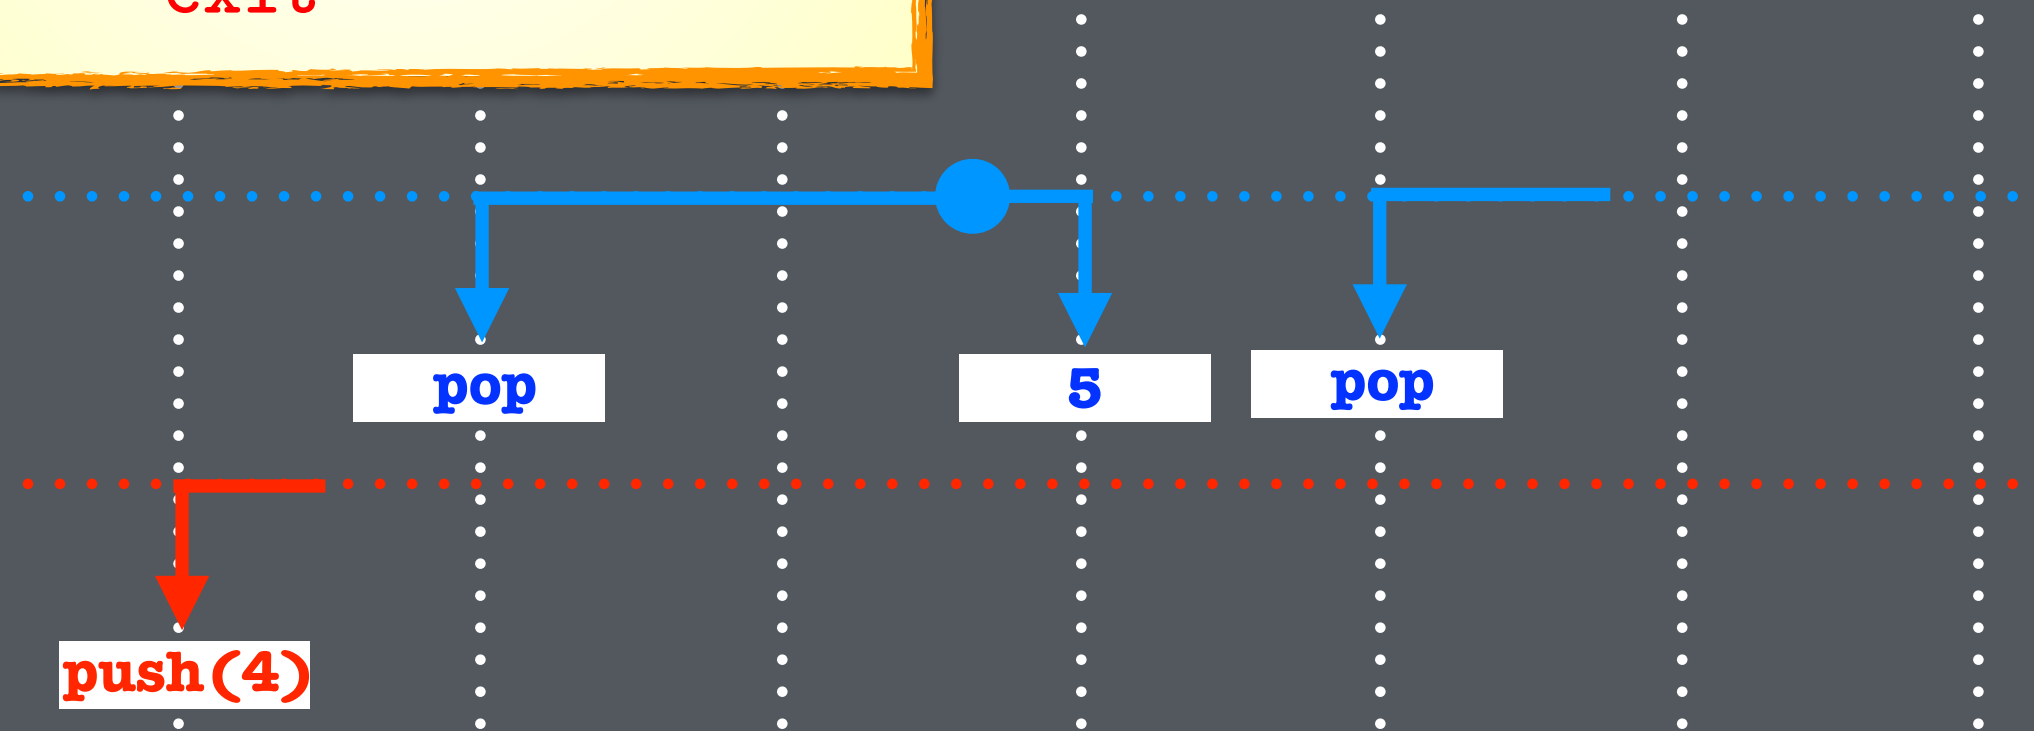
push(4)

pop

pop:

Node t

```
1 while (true)
2   t = Top
3   if (t = NULL)
4     return *
5   exit
6   if (CAS (Top,t,t.next))
7     return t.val
8   exit
```





# The Treiber Stack Algorithm

## Linearization Policy

push(k):

Node t

```
1 n = new Node(k,-)
2 while (true)
3   t = Top
4   n.next = t
5   if (CAS (Top,t,n))
6     exit
```

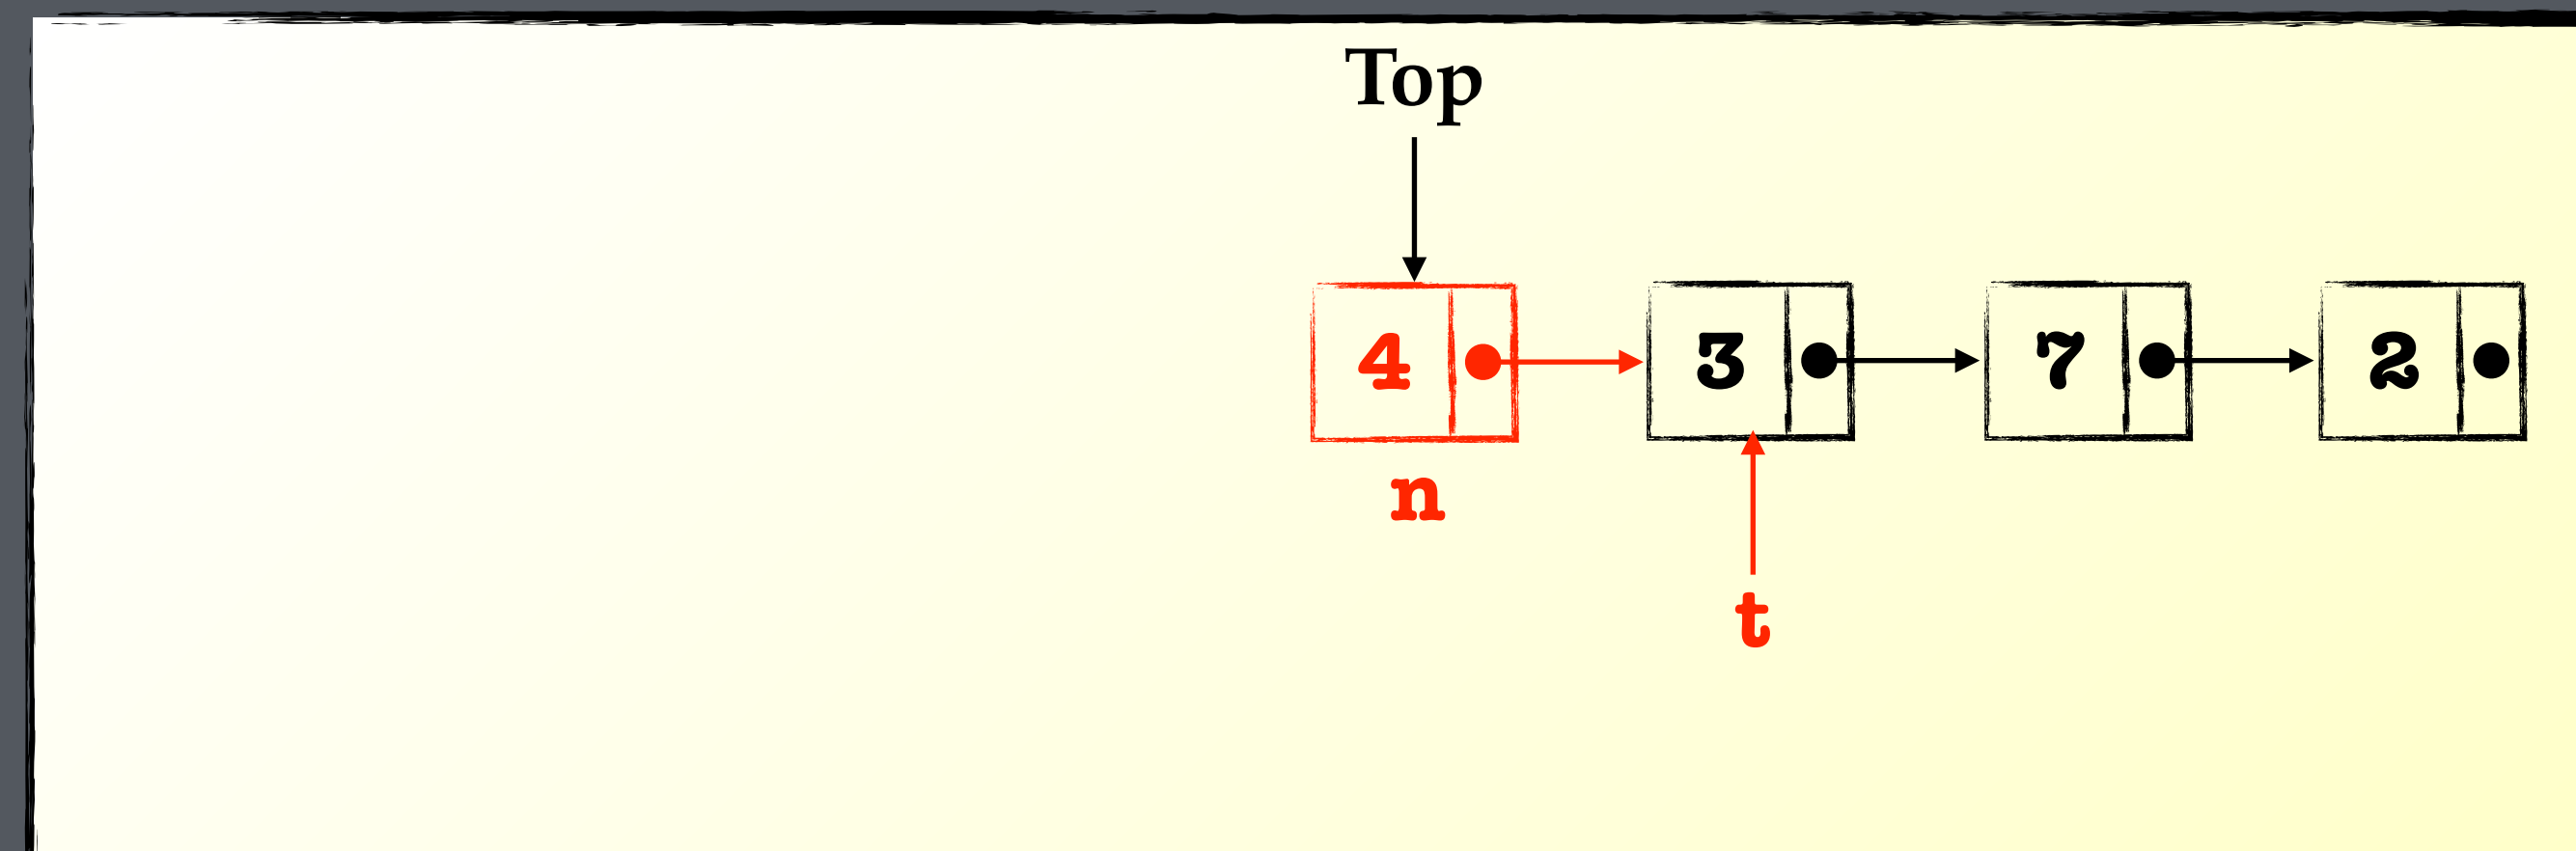

push(4)

pop

pop:

Node t

```
1 while (true)
2   t = Top
3   if (t = NULL)
4     return *
5   exit
6   if (CAS (Top,t,t.next))
7     return t.val
8   exit
```



# The Treiber Stack Algorithm

## Linearization Policy

push(k):

Node t

```
1 n = new Node(k,-)
2 while (true)
3   t = Top
4   n.next = t
5   if (CAS (Top,t,n))
6     exit
```

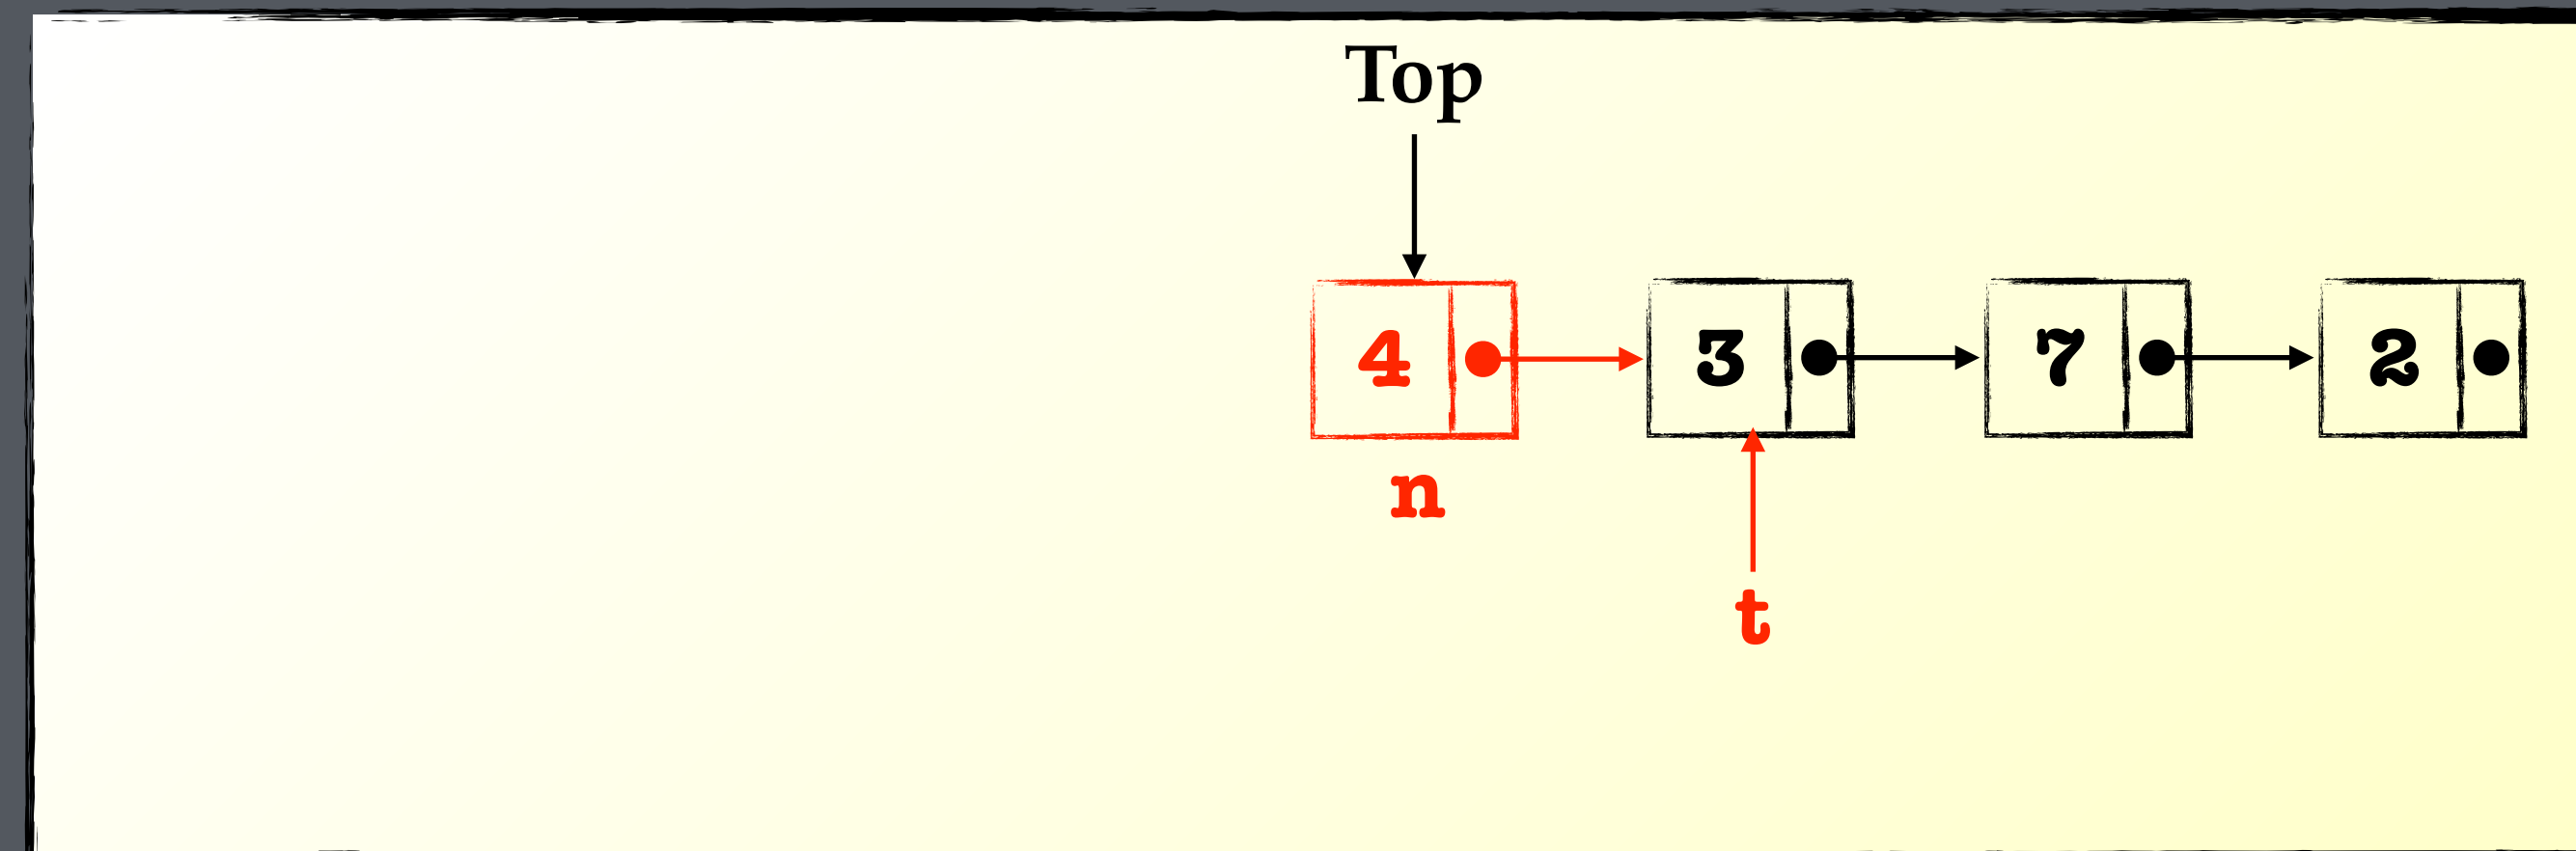
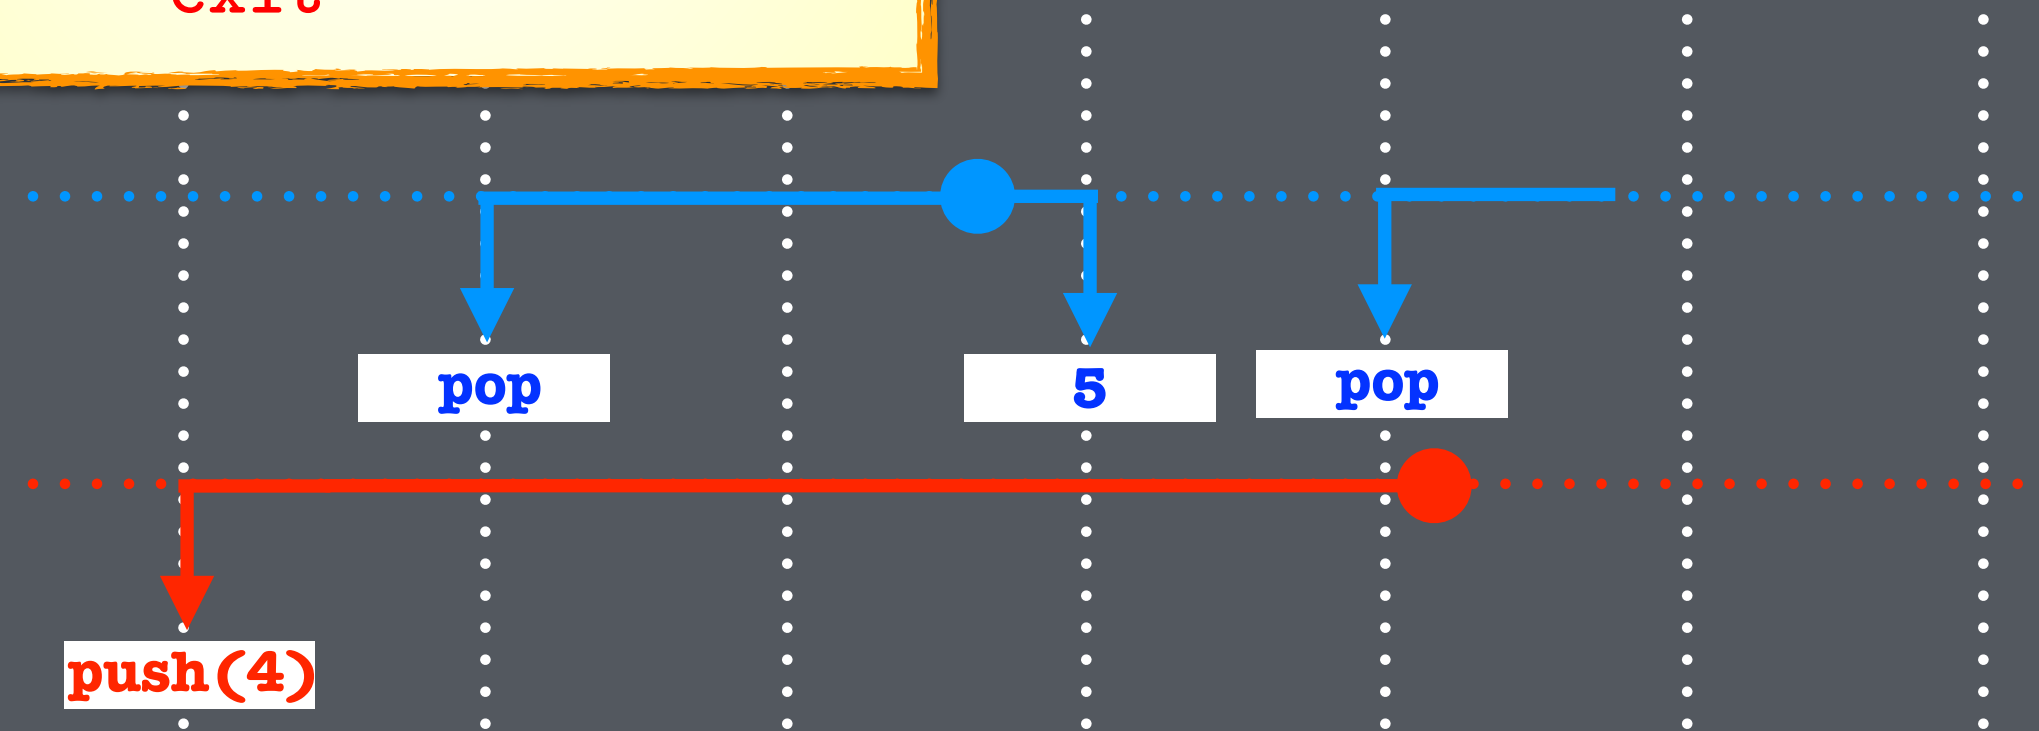
push(4)

pop

pop:

Node t

```
1 while (true)
2   t = Top
3   if (t = NULL)
4     return *
5   exit
6   if (CAS (Top,t,t.next))
7     return t.val
8   exit
```









# The Treiber Stack Algorithm

## Linearization Policy

push(k):

Node t

```
1 n = new Node(k,-)
2 while (true)
3   t = Top
4   n.next = t
5   if (CAS (Top,t,n))
6     exit
```

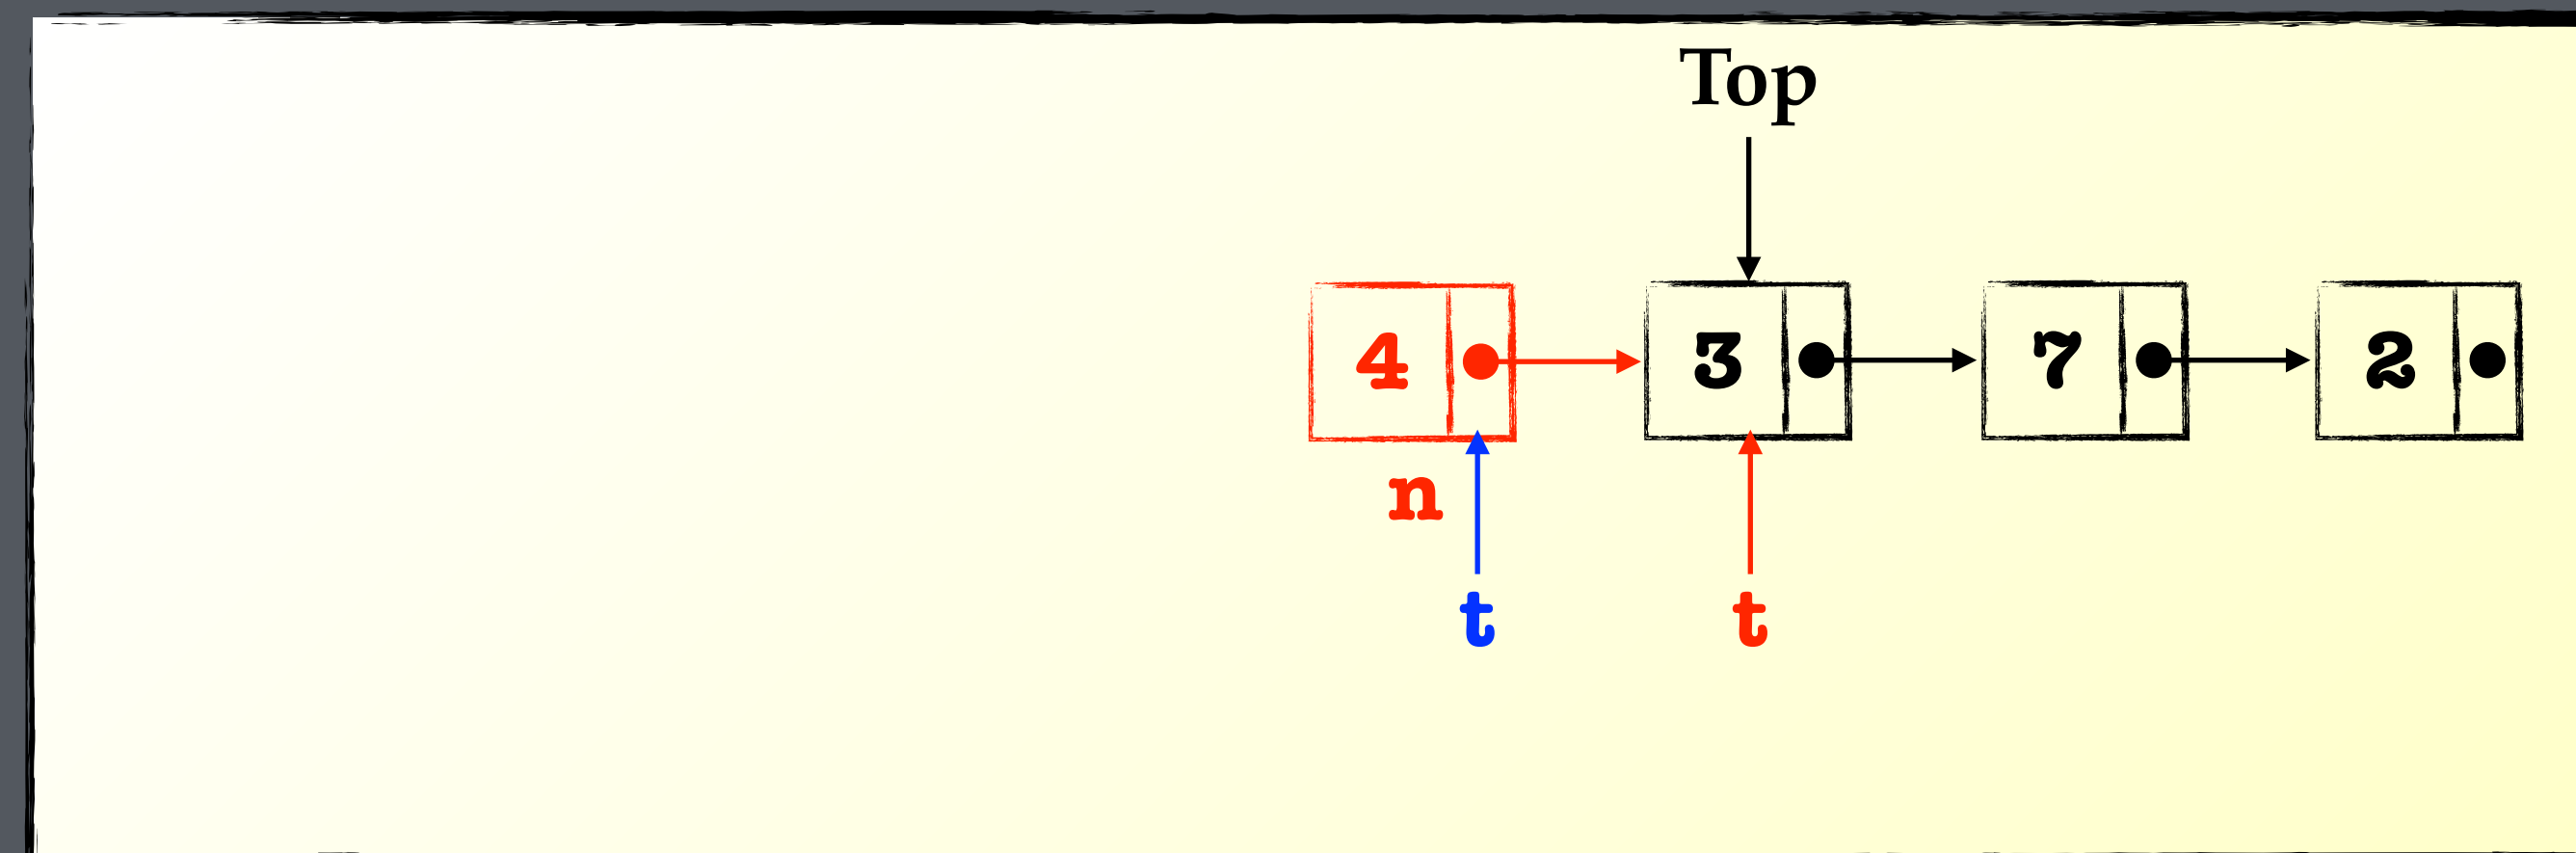
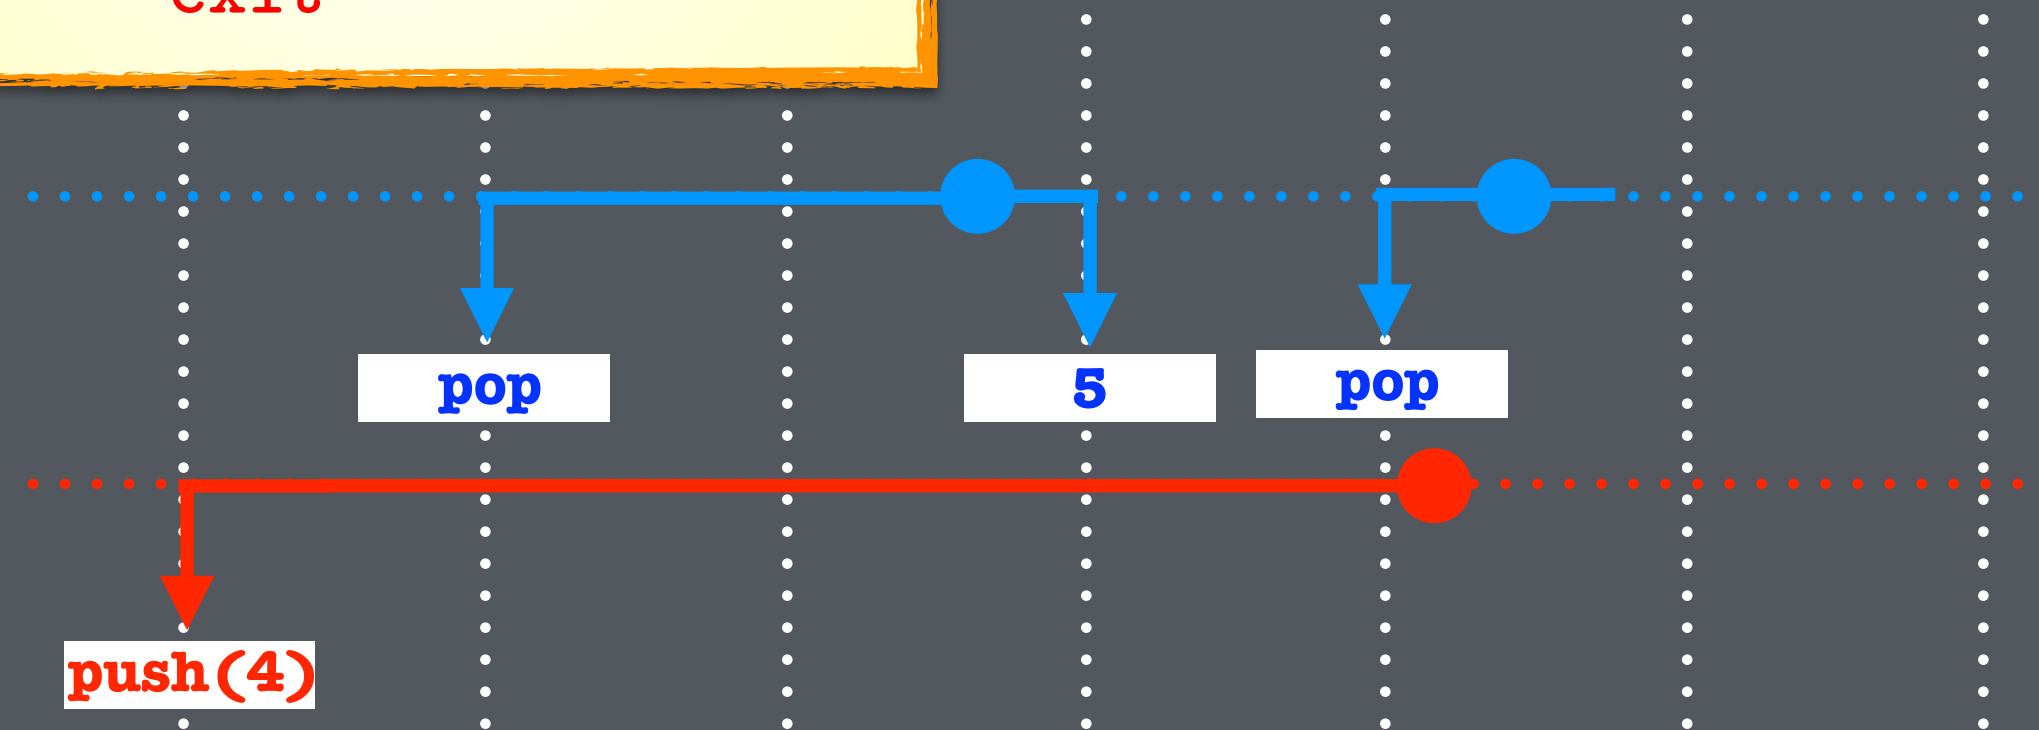
push(4)

pop

pop:

Node t

```
1 while (true)
2   t = Top
3   if (t = NULL)
4     return *
5   exit
6   if (CAS (Top,t,t.next))
7     return t.val
8   exit
```





# The Treiber Stack Algorithm

## Linearization Policy

push(k):

Node t

```
1 n = new Node(k,-)
2 while (true)
3   t = Top
4   n.next = t
5   if (CAS (Top,t,n))
6     exit
```


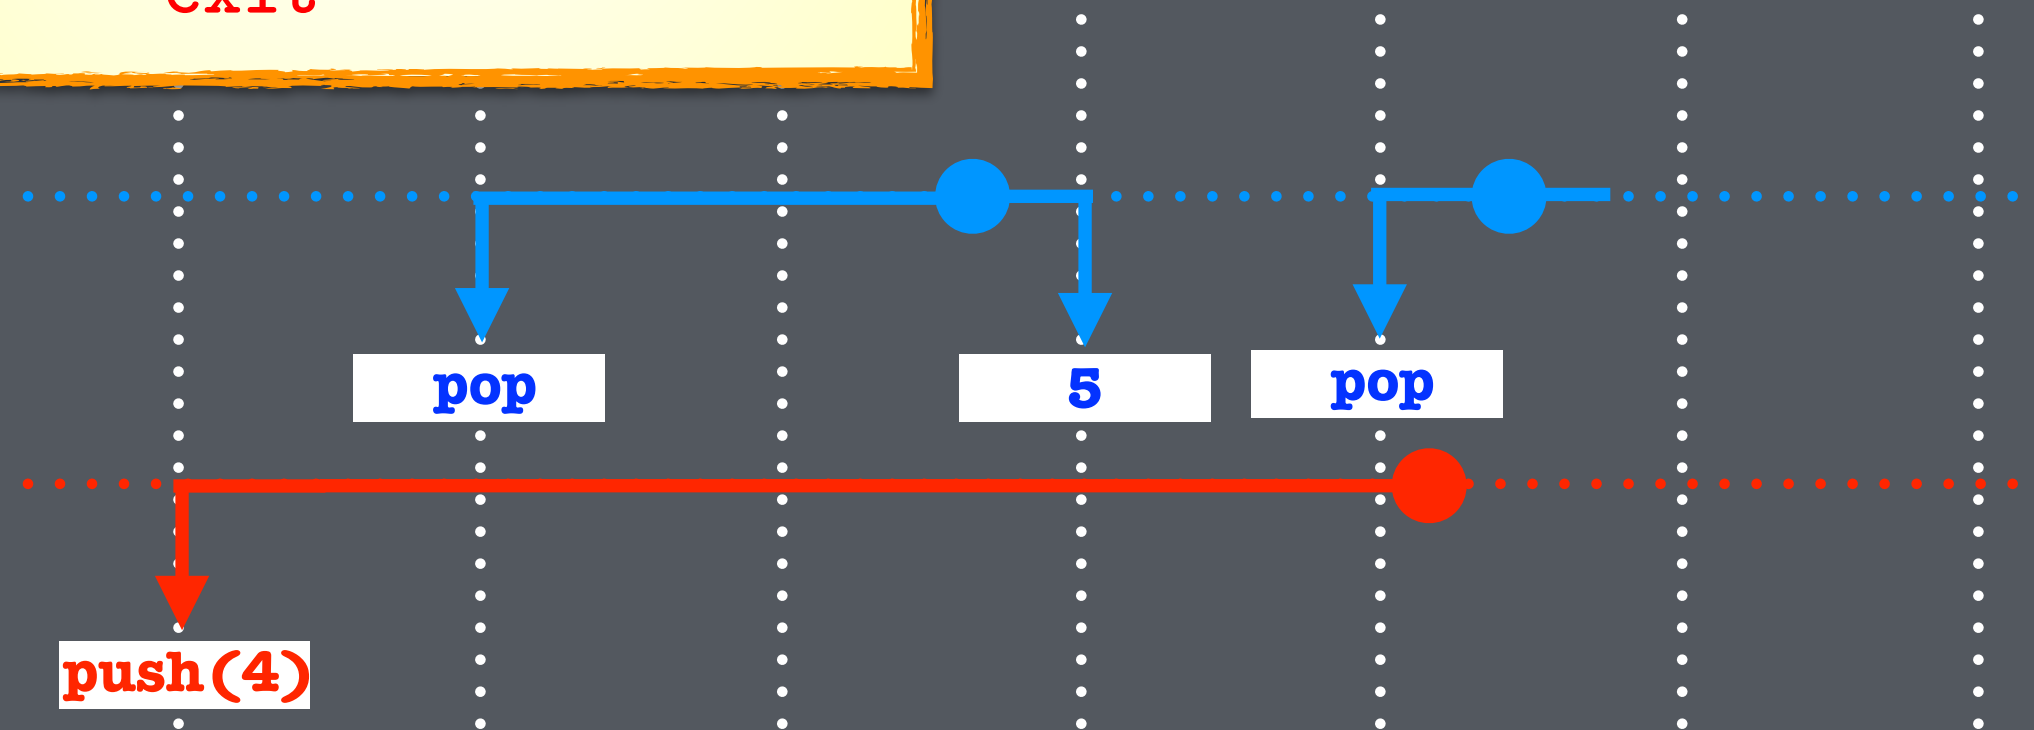
push(4)

pop

pop:

Node t

```
1 while (true)
2   t = Top
3   if (t = NULL)
4     return *
5   exit
6   if (CAS (Top,t,t.next))
7     return t.val
8   exit
```



# The Treiber Stack Algorithm

## Linearization Policy

push(k):

Node t

1 n = new Node(k,-)

2 while (true)

3 t = Top

4 n.next = t

5 if (CAS (Top,t,n))

6 exit

push(4)

pop

pop:

Node t

1 while (true)

2 t = Top

3 if (t = NULL)


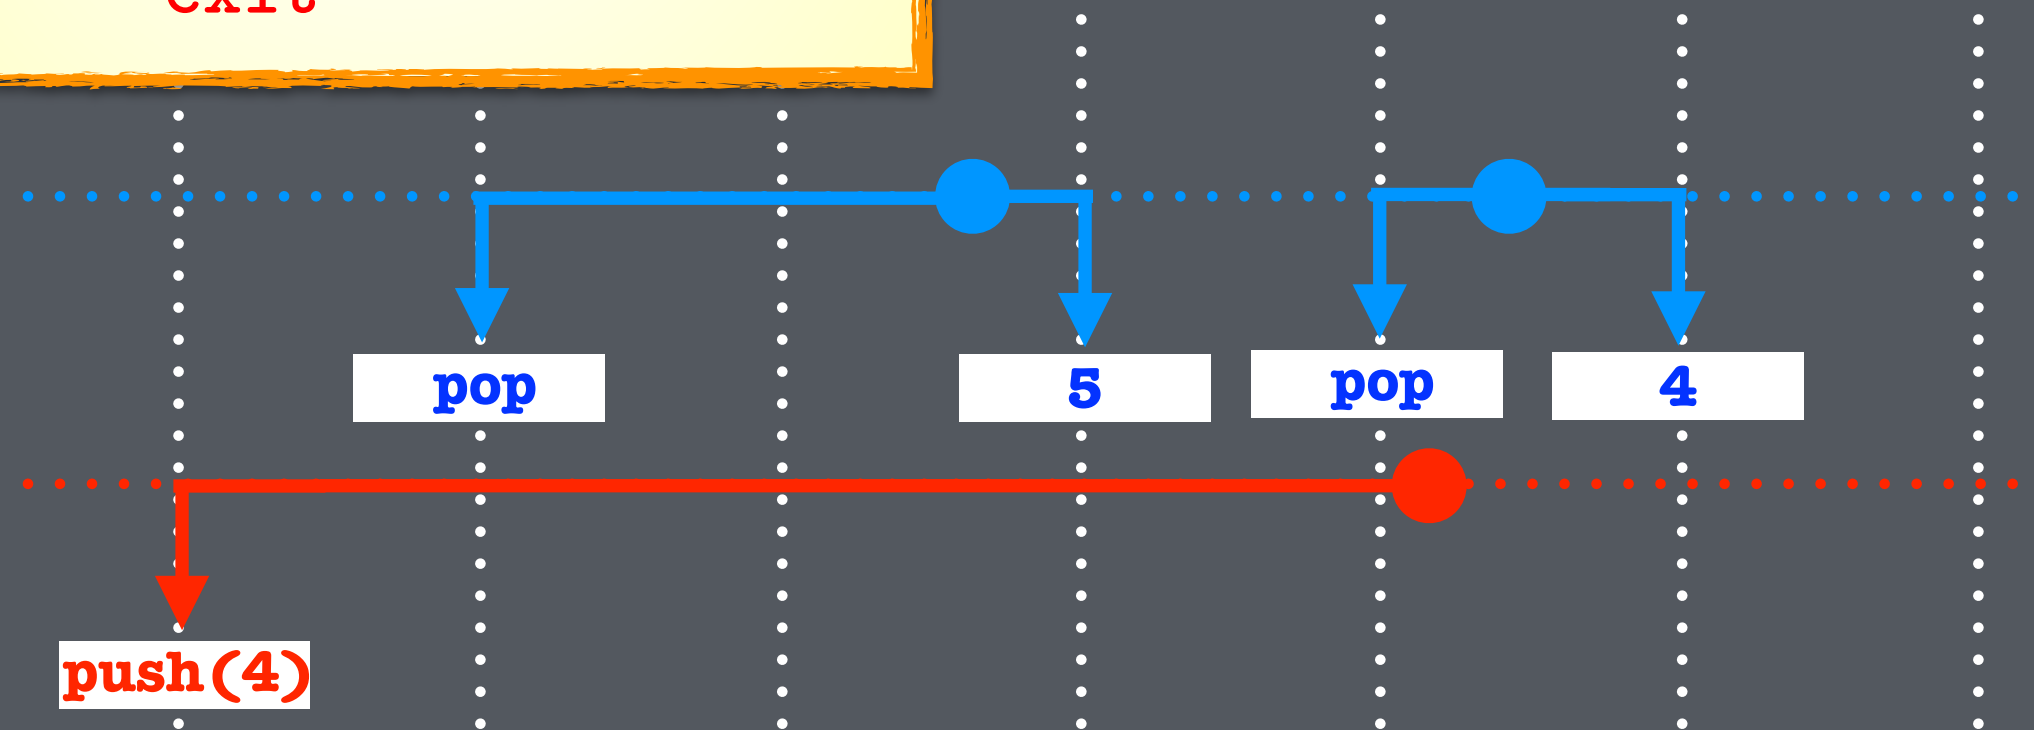
4 return \*

5 exit

6 if (CAS (Top,t,t.next))

7 return t.val

8 exit



# The Treiber Stack Algorithm

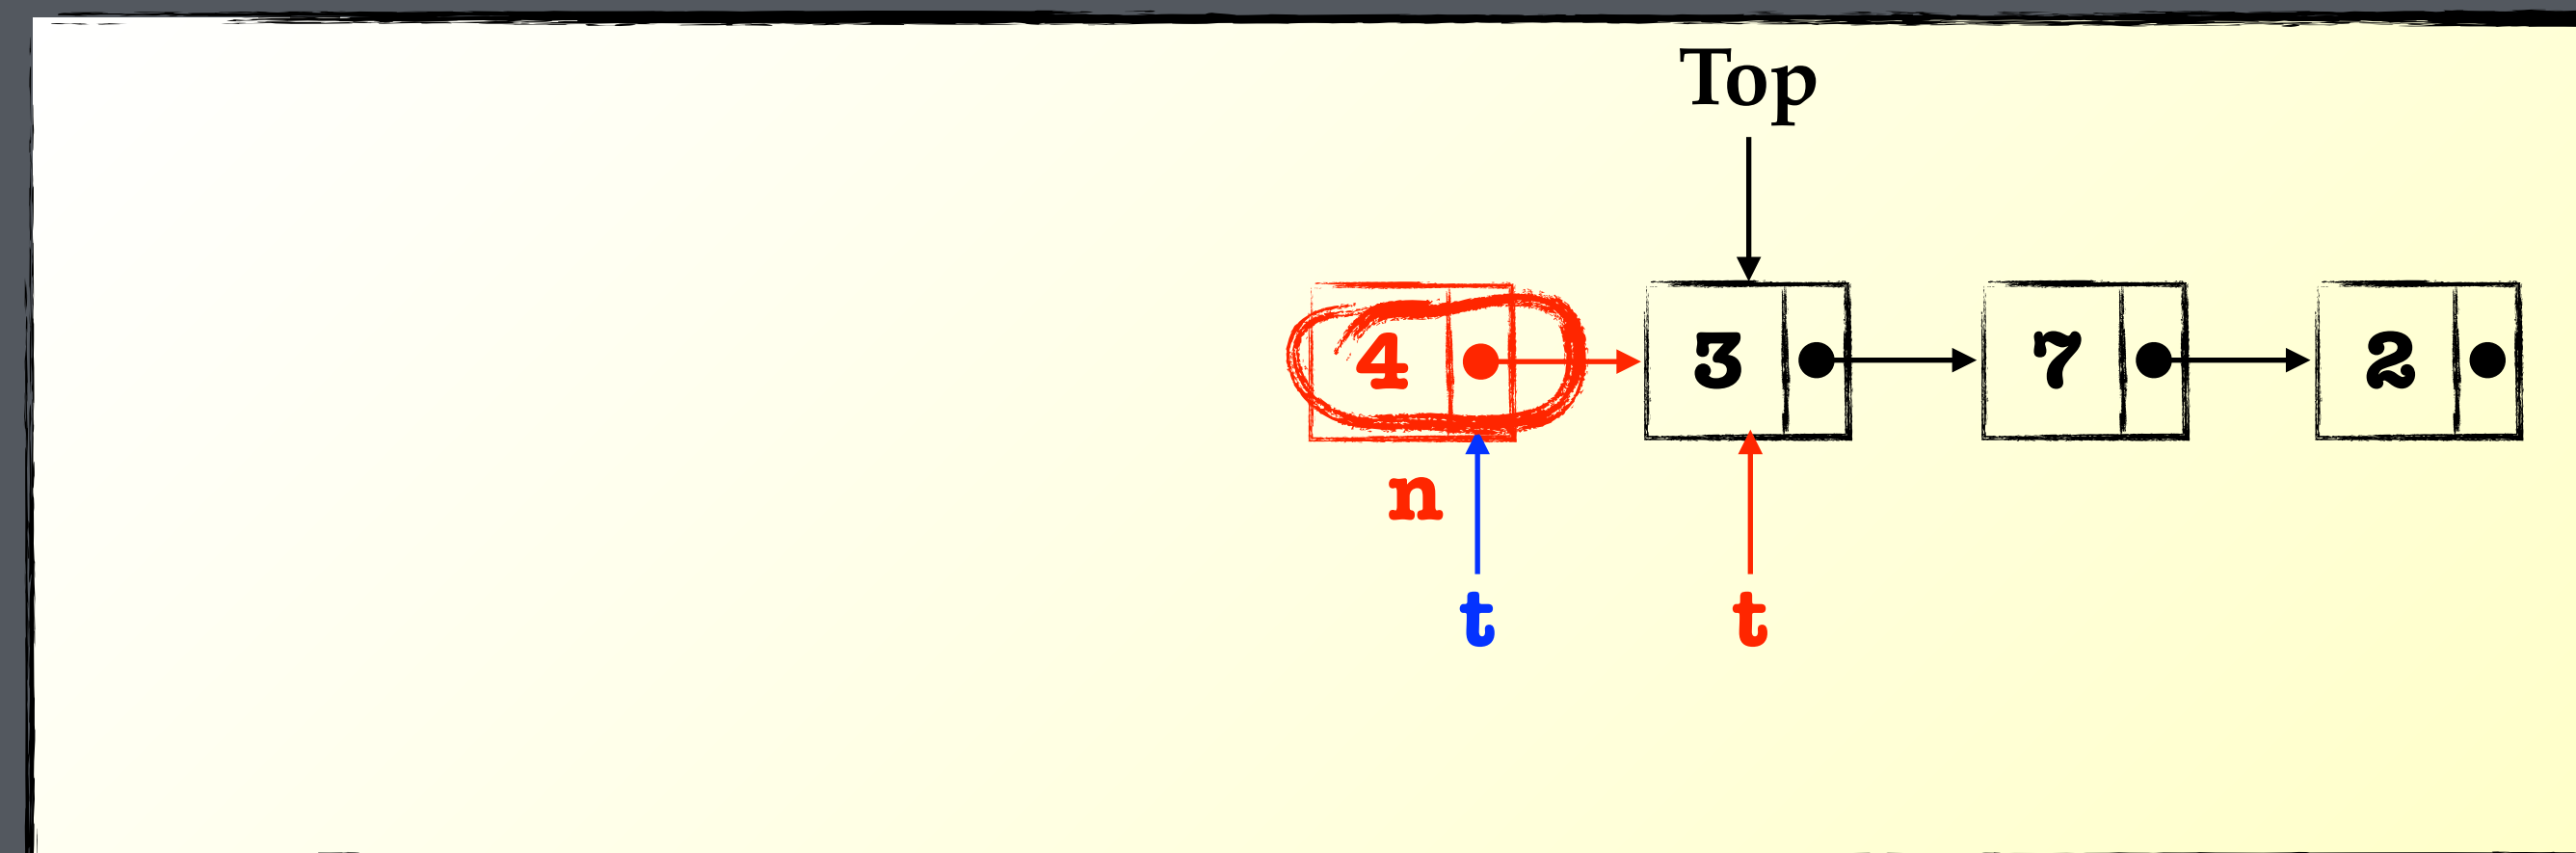
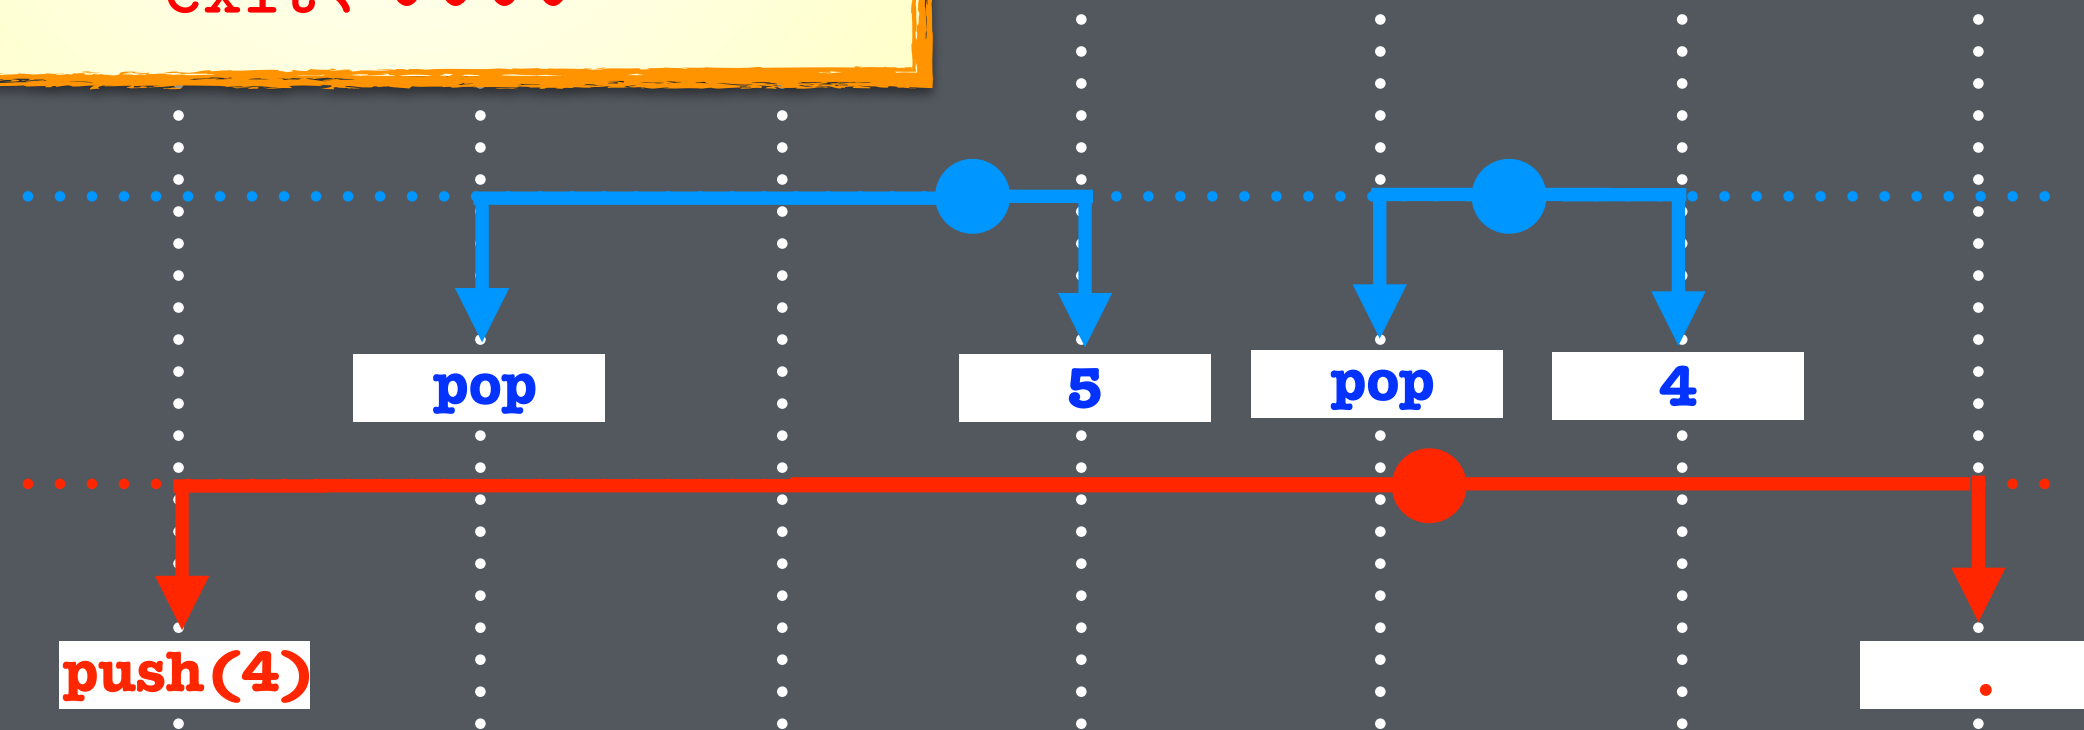
## Linearization Policy

```
push(k):  
Node t  
1 n = new Node(k,-)  
2 while (true)  
3   t = Top  
4   n.next = t  
5   if (CAS (Top,t,n))  
6     exit ← ~~~~~
```

**push(4)**

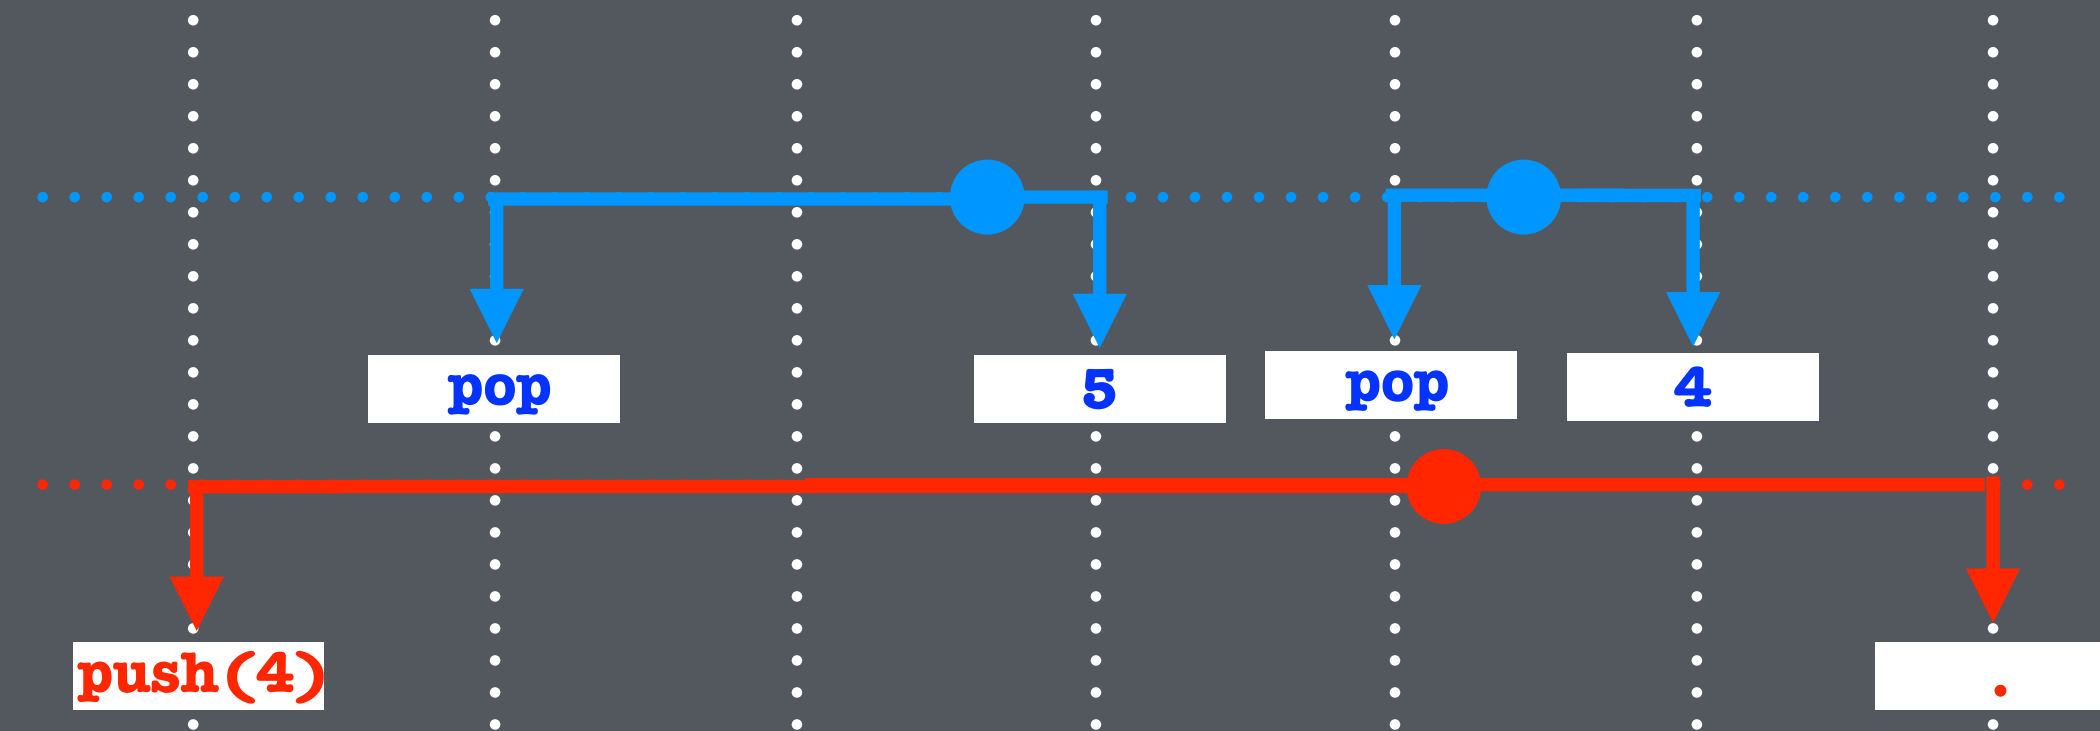
**pop**

```
pop:  
Node t  
1 while (true)  
2   t = Top  
3   if (t = NULL)  
4     return *  
5   exit  
6   if (CAS (Top,t,t.next))  
7     return t.val  
8   exit ← ~~~~~
```



# The Treiber Stack Algorithm

## Linearization Policy





# Libraries

A *configuration*  $c = \langle v, u \rangle$  of  $L[C]$  is a shared memory valuation  $v \in V$ , along with a map  $u$  mapping each thread  $t \in \mathbb{N}$  to a tuple  $u(t) = \langle \ell, m_0, m \rangle$ , composed of a client-local state  $\ell \in Q_C$ , along with initial and current method states  $m_0, m \in Q_L \cup \{\perp\}$ ;  $m_0 = m = \perp$  when thread  $t$  is not executing a library

<p>INTERNAL</p> $\frac{\begin{array}{l} u_1(t) = \langle \ell, m_0, m_1 \rangle \\ \langle m_1, v_1 \rangle \xrightarrow{a} \langle m_2, v_2 \rangle \\ u_2 = u_1 (t \mapsto \langle \ell, m_0, m_2 \rangle) \end{array}}{\langle v_1, u_1 \rangle \xrightarrow[L[C]]{\langle a, t \rangle} \langle v_2, u_2 \rangle}$	<p>CALL</p> $\frac{\begin{array}{l} u_1(t) = \langle \ell_1, \perp, \perp \rangle \\ m_0 \in I_M \quad \ell_1 \xrightarrow[M(m_0, m_f)]{}_C \ell_2 \\ u_2 = u_1 (t \mapsto \langle \ell_1, m_0, m_0 \rangle) \end{array}}{\langle v, u_1 \rangle \xrightarrow[L[C]]{\text{call}(M, m_0, t)} \langle v, u_2 \rangle}$	<p>RETURN</p> $\frac{\begin{array}{l} u_1(t) = \langle \ell_1, m_0, m_f \rangle \\ m_f \in F_M \quad \ell_1 \xrightarrow[M(m_0, m_f)]{}_C \ell_2 \\ u_2 = u_1 (t \mapsto \langle \ell_2, \perp, \perp \rangle) \end{array}}{\langle v, u_1 \rangle \xrightarrow[L[C]]{\text{ret}(M, m_f, t)} \langle v, u_2 \rangle}$
--	--	---

**Fig. 1.** The transition relation  $\rightarrow_{L[C]}$  for the library-client composition  $L[C]$ .

# VASS model

We associate to each concurrent system  $L[C]$  a *canonical VASS*,<sup>2</sup> denoted  $\mathcal{A}_{L[C]}$ , whose states are the set of shared-memory valuations, and whose vector components count the number of threads in each thread-local state; a transition of  $\mathcal{A}_{L[C]}$  from  $\langle v_1, \mathbf{n}_1 \rangle$  to  $\langle v_2, \mathbf{n}_2 \rangle$  updates the shared-memory valuation from  $v_1$  to  $v_2$  and the local state of some thread  $t$  from  $u_1(t)$  to  $u_2(t)$  by decrementing the  $u_1(t)$ -component of  $\mathbf{n}_1$ , and incrementing the  $u_2(t)$ -component, to derive  $\mathbf{n}_2$ .

# Specifications

A *specification*  $S$  of a library  $L$  is a language over the *specification alphabet*

$$\Sigma_S \stackrel{\text{def}}{=} \{M[m_0, m_f] : M \in L, m_0, m_f \in Q_M\}.$$

**Definition 2 (Linearizability [20]).** A trace  $\tau$  is  $S$ -linearizable when there exists a completion<sup>4</sup>  $\pi$  of a strict, serial permutation of  $\tau$  such that  $(\pi \mid S) \in S$ .

# Specifications

The *pending closure* of a specification  $S$ , denoted  $\overline{S}$  is the set of  $S$ -images of serial sequences which have completions whose  $S$ -images are in  $S$ :

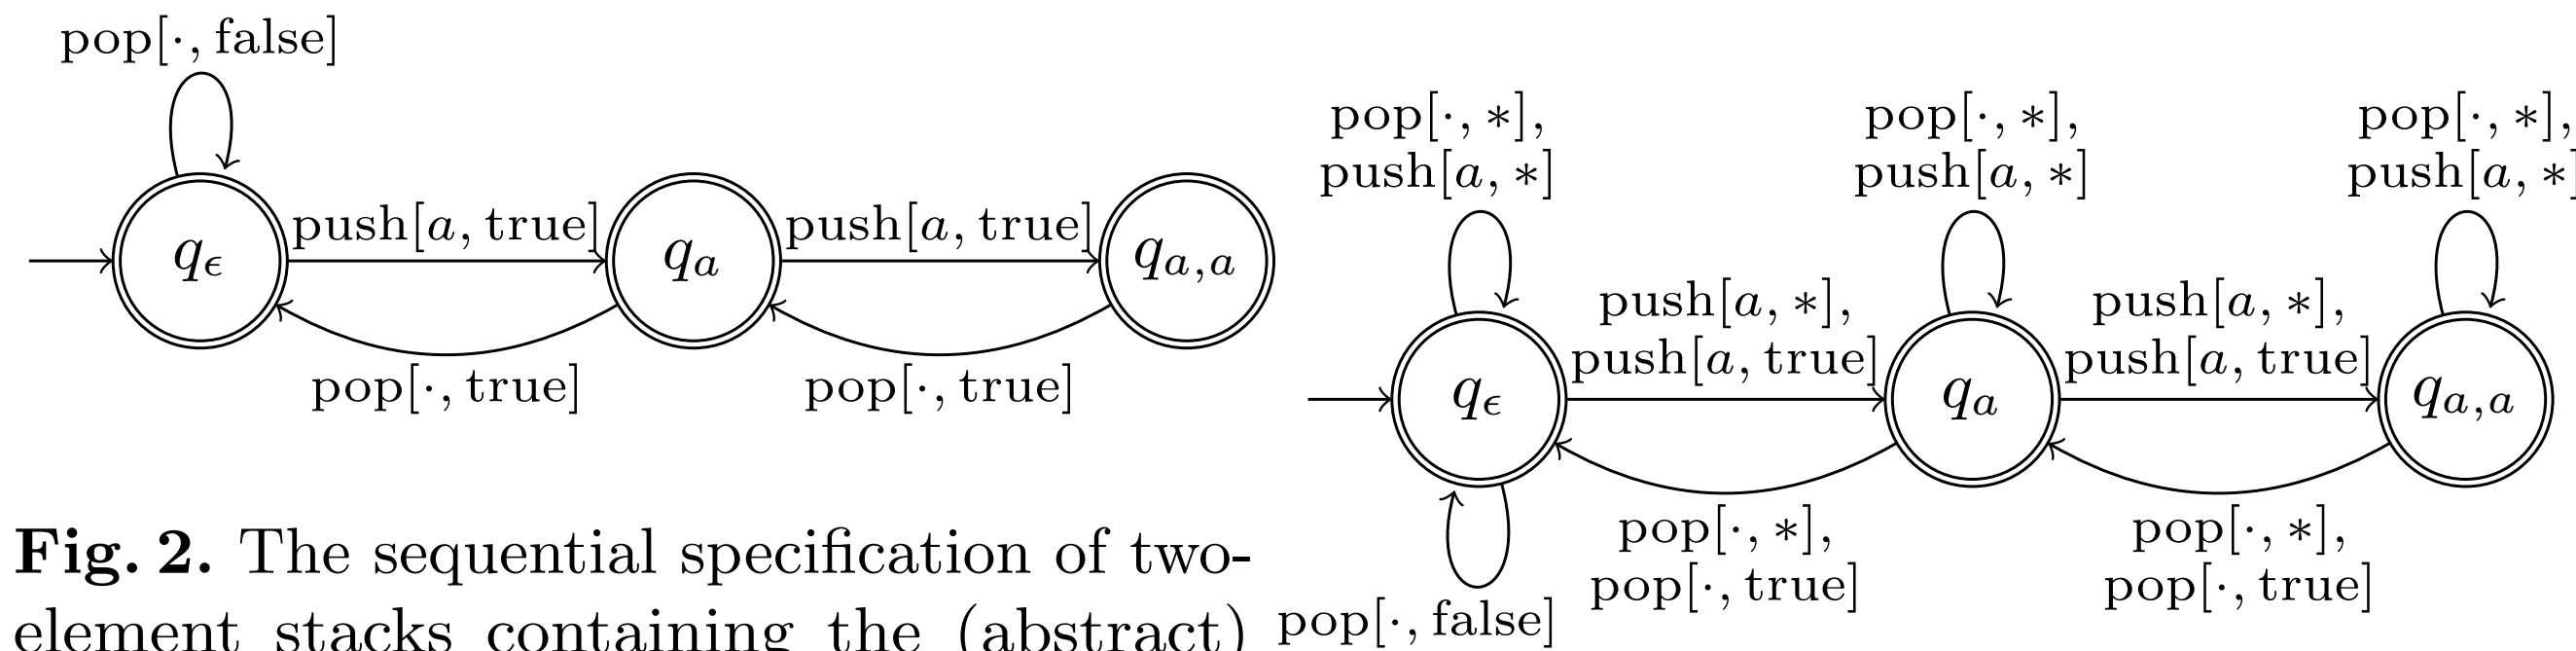
$$\overline{S} \stackrel{\text{def}}{=} \{(\sigma \mid S) \in \overline{\Sigma}_S^* : \exists \sigma' \in \Sigma_S^*. (\sigma' \mid S) \in S \text{ and } \sigma' \text{ is a completion of } \sigma\}.$$



# Specifications

The *pending closure* of a specification  $S$ , denoted  $\overline{S}$  is the set of  $S$ -images of serial sequences which have completions whose  $S$ -images are in  $S$ :

$$\overline{S} \stackrel{\text{def}}{=} \{(\sigma \mid S) \in \overline{\Sigma}_S^* : \exists \sigma' \in \Sigma_S^*. (\sigma' \mid S) \in S \text{ and } \sigma' \text{ is a completion of } \sigma\}.$$



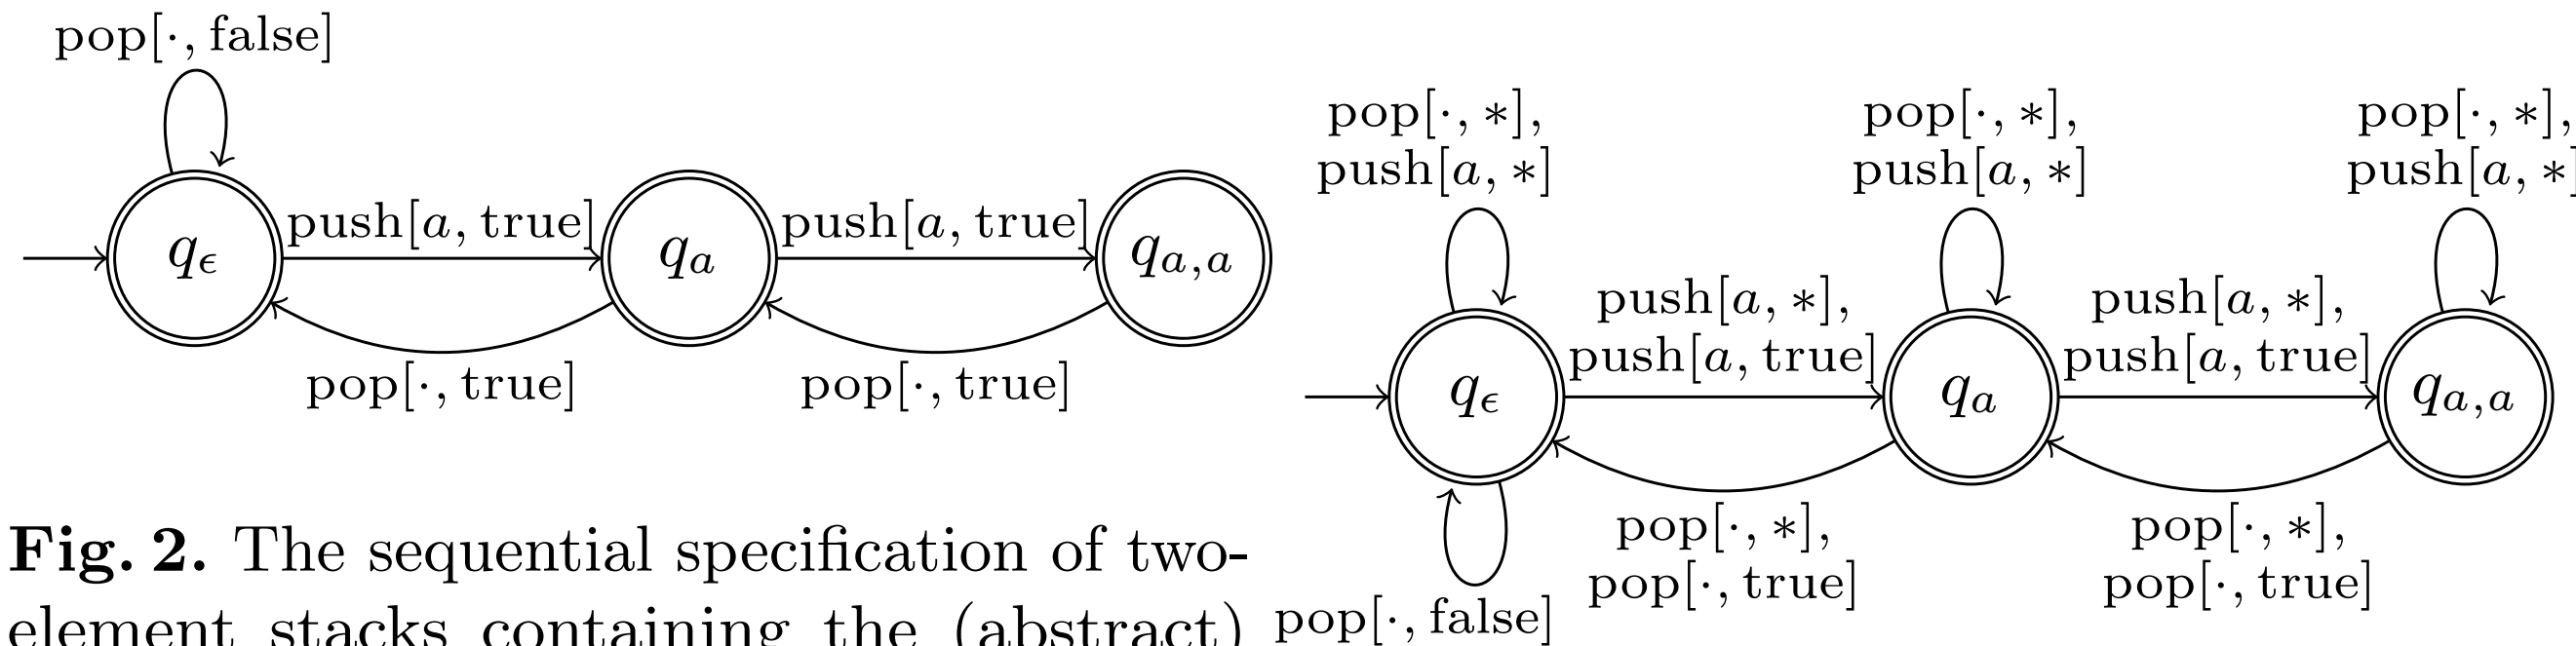
**Fig. 2.** The sequential specification of two-element stacks containing the (abstract) value  $a$ , given as the language of a finite automaton, whose operation alphabet indicates both the argument and return values.

**Fig. 3.** The pending closure of the stack specification from Figure 2.

# Specifications

The *pending closure* of a specification  $S$ , denoted  $\overline{S}$  is the set of  $S$ -images of serial sequences which have completions whose  $S$ -images are in  $S$ :

$$\overline{S} \stackrel{\text{def}}{=} \{(\sigma \mid S) \in \overline{\Sigma}_S^* : \exists \sigma' \in \Sigma_S^*. (\sigma' \mid S) \in S \text{ and } \sigma' \text{ is a completion of } \sigma\}.$$



**Fig. 2.** The sequential specification of two-element stacks containing the (abstract) value  $a$ , given as the language of a finite automaton, whose operation alphabet indicates both the argument and return values.

**Fig. 3.** The pending closure of the stack specification from Figure 2.


**Lemma 1.** *The pending closure  $\overline{S}$  of a regular specification  $S$  is regular.*

**Lemma 2.** *A trace  $\tau$  is  $S$ -linearizable if and only if there exists a strict, serial permutation  $\pi$  of  $\tau$  such that  $(\pi \mid S) \in \overline{S}$ .*

# Read-only operations

Given a method  $M$  of a library  $L$  and  $m_0, m_f \in Q_M$ , an  $M[m_0, m_f]$ -operation  $\theta$  is *read-only* for a specification  $S$  if and only if for all  $w_1, w_2, w_3 \in \Sigma_S^*$ ,

1. If  $w_1 \cdot M[m_0, m_f] \cdot w_2 \in S$  then  $w_1 \cdot M[m_0, m_f]^k \cdot w_2 \in S$  for all  $k \geq 0$ , and
2. If  $w_1 \cdot M[m_0, m_f] \cdot w_2 \in S$  and  $w_1 \cdot w_3 \in S$  then  $w_1 \cdot M[m_0, m_f] \cdot w_3 \in S$ .



# Linearization points

The *control graph*  $G_M = \langle Q_M, E \rangle$  is the quotient of a method  $M$ 's transition system by shared-state valuations  $V$ :  $\langle m_1, a, m_2 \rangle \in E$  iff  $\langle m_1, v_1 \rangle \hookrightarrow_M^a \langle m_2, v_2 \rangle$  for some  $v_1, v_2 \in V$ . A function  $\text{LP} : L \rightarrow \wp(\Sigma_L)$  is called a *linearization-point mapping* when for each  $M \in L$ :

1. each symbol  $a \in \text{LP}(M)$  labels at most one transition of  $M$ ,
2. any directed path in  $G_M$  contains at most one symbol of  $\text{LP}(M)$ , and
3. all directed paths in  $G_M$  containing  $a \in \text{LP}(M)$  reach the same  $m_a \in F_M$ .

An action  $\langle a, i \rangle$  of an  $M$ -operation is called a *linearization point* when  $a \in \text{LP}(M)$ , and operations containing linearization points are said to be *effectuated*;  $\text{LP}(\theta)$  denotes the unique linearization point of an effectuated operation  $\theta$ . A *read-points mapping*  $\text{RP} : \Theta \rightarrow \mathbb{N}$  for an action sequence  $\sigma$  with operations  $\Theta$  maps each read-only operation  $\theta$  to the index  $\text{RP}(\theta)$  of an internal  $\theta$ -action in  $\sigma$ .

# Exercices (1)

- Does the Herlihy & Wing queue admit **fixed** linearization points ?

```
void enq(int x) {
    i = back++; items[i] = x;
}

int deq() {
    while (1) {
        range = back - 1;
        for (int i = 0; i <= range; i++) {
            x = swap(items[i], null);
            if (x != null) return x;
        }
    }
}
```



# Static linearizability

An action sequence  $\sigma$  is called *effectuated* when every completed operation of  $\sigma$  is either effectuated or read-only, and an effectuated completion  $\sigma'$  of  $\sigma$  is *effect preserving* when each effectuated operation of  $\sigma$  also appears in  $\sigma'$ . Given a linearization-point mapping LP, and a read-points mapping RP of an action sequence  $\sigma$ , we say a permutation  $\pi$  of  $\sigma$  is *point preserving* when every two operations of  $\pi$  are ordered by their linearization/read points in  $\sigma$ .

**Definition 4.** A trace  $\tau$  is  $\langle S, \text{LP} \rangle$ -linearizable when  $\tau$  is effectuated, and there exists a read-points mapping RP of  $\tau$ , along with an effect-preserving completion  $\pi$  of a strict, point-preserving, and serial permutation of  $\tau$  such that  $(\pi \mid S) \in S$ .

**Definition 5 (Static Linearizability).** The system  $L[C]$  is  $S$ -static linearizable when  $L[C]$  is  $\langle S, \text{LP} \rangle$ -linearizable for some mapping LP.

# Checking Static Linearizability

- $A_S$  = a deterministic automaton recognizing the Specification
- we define a monitor to be composed with  $L[C]$  that simulates the Specification
  - methods have a new local variable RO which is initially  $\emptyset$  (records return values of read-only operations)
    - if  $m_f \in RO$  in an invocation of  $M$ , then  $M[m_0, m_f]$  is read-only and a state of  $A_S$  in which  $M[m_0, m_f]$  is enabled has been observed
  - $L[C]$  executes a linearization point  $\Rightarrow$  the state of the Specification is advanced to the  $M[m_0, m_f]$  successor ( $m_0$  is the initial state of the current operation and  $m_f$  is the unique final state reachable from this lin. point)
  - $L[C]$  executes an internal action from an  $M[m_0, *]$  operation  $\Rightarrow$  RO is enriched with every  $m_f$  such that  $M[m_0, m_f]$  is read-only and enabled in the current specification state
  - $L[C]$  executes the return of an  $M[m_0, m_f]$  read-only operation  $\Rightarrow$  if  $m_f \notin RO$  then the monitor goes to an error state

# EXPSPACE-hardness

- Reduce control state reachability in VASS (which is EXPSPACE-complete) to static linearizability
- Use the library from the undecidability proof without the zero-test method (the specification excludes only executions not reaching the target state)

# Checking Linearizability: Complexity (finite-state implementations)

## Bounded Nb. of Threads:

- EXSPACE-complete [Alur et al., 1996, Hamza 2015]

## Unbounded Nb. of Threads:

- Undecidable [Bouajjani et al., 2013]
- Decidable with “fixed linearization points” [Bouajjani et al. 2013]

**Alur et al. 1996:** Rajeev Alur, Kenneth L. McMillan, Doron A. Peled: Model-Checking of Correctness Conditions for Concurrent Objects. LICS 1996

**Bouajjani et al., 2013:** Ahmed Bouajjani, Michael Emmi, Constantin Enea, Jad Hamza: Verifying Concurrent Programs against Sequential Specifications. ESOP 2013

**Hamza 2015:** Jad Hamza: On the Complexity of Linearizability. NETYS 2015