

In this presentation we take a closer look at complexity classes in which the bound is on the amount of memory it takes to compute the problem.

**Goal:**

- Explore space complexity

**Plan:**

- Low space classes: L, NL
- Savitch's Theorem
- Immerman's Theorem
- TQBF

2

In particular, we'll look at low complexity classes, such as LOGSPACE and non-deterministic LOGSPACE. Among others, we prove three fundamental theorems regarding those classes.

# Space-Complexity

## Definition:

- Let  $t: \mathbb{N} \rightarrow \mathbb{N}$  be a complexity function

## Deterministic space:

$$SPACE[t(n)] \equiv \{L \mid L \text{ decided by } O(t(n))\text{-space deterministic TM}\}$$

## Nondeterministic space:

$$NSPACE[t(n)] \equiv \{L \mid L \text{ decided by } O(t(n))\text{-space nondeterministic TM}\}$$



the input takes  $n$  cells; how can a TM use only  $\log n$  space?

## Det. Log space:


$$L \equiv SPACE[\log(n)]$$

## Nondet. Log space:

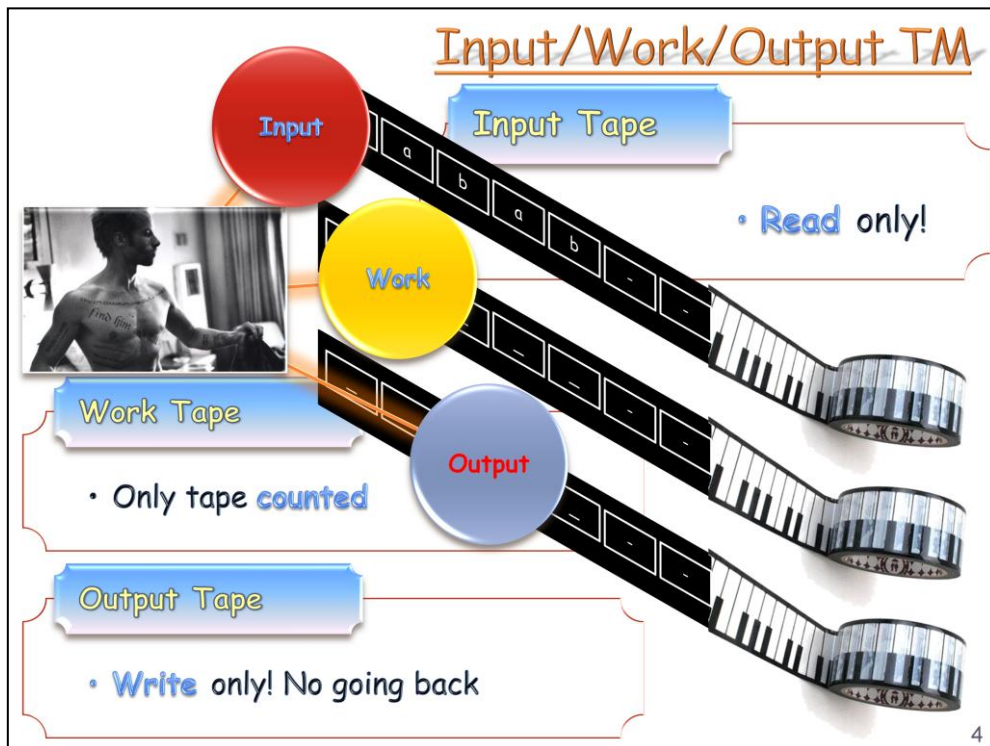
$$NL \equiv NSPACE[\log(n)]$$

## Det polynomial space:

$$PSPACE \equiv \bigcup_k SPACE[n^k]$$



Let us recall our definition of space complexity classes: it is quite straightforward, however, we need to clarify what does it mean for an algorithm to use sub linear space.

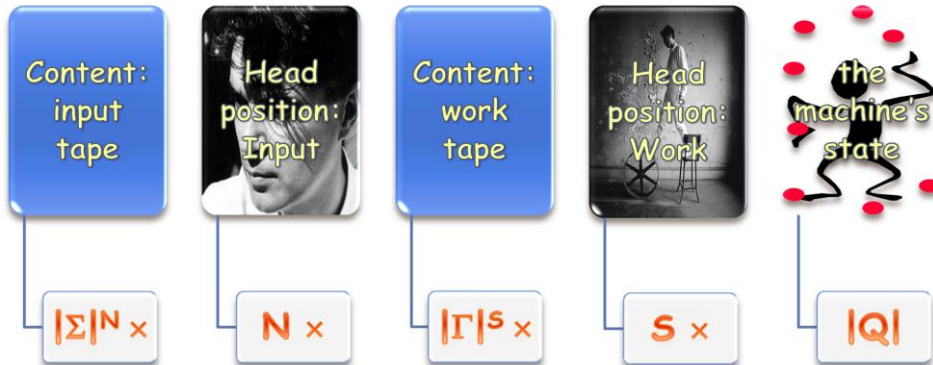


For that purpose, we change a little our model of computation to consist of an input tape, which is read only, an output tape, which is write only, and a work tape, which is the only one counted for purposes of complexity bounds.

## Configurations


How many distinct configurations may a TM with input-size  $N$  and work-tape of size  $S$  have?

What about output?



5

Let us now figure out how many configurations such a machine has. The location of the heads on the input tape and on the work tape are counted. Both the content of the output tape and the location of the head on it are not considered in counting the configurations. The content of only the work tape is included.



Try to put the following computational problems in as small a class as you can.  
 Try also to come up with a problem that is in non-deterministic LOGSPACE,  
 however is not known to be in LOGSPACE.



# Log-space Reductions

$A$  is log-space reducible to  $B$   
(denoted  $A \leq_L B$ )

If  
there  
exists a

log-space-computable  
function  $f: \Sigma^* \rightarrow \Sigma^*$

s.t. for  
every  $w$

$w \in A \Leftrightarrow f(w) \in B$

i.e.,  $\exists$  log-space  
TM that outputs  
 $f(w)$  on input  $w$


$f$  is a log-space  
reduction of  $A$  to  $B$

**Theorem:**

- L, NL, P, NP, PSPACE and EXPTIME are closed  
under log-space reductions.

We can now define LOGSPACE reductions: they're the same as Karp reductions, with the added restriction that the reduction-function must be computed using only logarithmic memory.





$L$  Closed under  $\leq_L$

**WRONG!!** Why not simply apply  $f$  then solve  $A_2$  on the outcome?

Claim:

- $f$  is a LOGSPACE reduction from  $A_1$  to  $A_2$  and  $A_2 \in L \Rightarrow A_1$  is in  $L$

Proof:

- on input  $x$ : Simulate  $M$  for  $A_2$ ; whenever  $M$  reads the  $i^{\text{th}}$  symbol of its input, run  $f$  on  $x$  and wait for the  $i^{\text{th}}$  bit to be outputted

8

Let us now see that these reductions can indeed be applied appropriately.

Think of the following scenario: you have a little chip that can play a DVD in a given format. You have a DVD encoded with a different format. You have another little chip that can convert the format the DVD is written in to the format the other chip can read. Is it possible to combine the two and build a machine that can play the DVD?

The wrong solution would be to store the output of the first chip and apply the second chip to that -there is simply not enough memory for that solution to work.

The correct solution is to run the second chip and give it the appropriate bits of the output of the first chip; if necessary, restart the first chip, and let it read the DVD from start.



# Graph Connectivity (CONN)


## Instance:

- a directed graph  $G=(V,E)$  and two vertices  $s, t \in V$

## Decision Problem:

- Is there a path from  $s$  to  $t$  in  $G$ ?

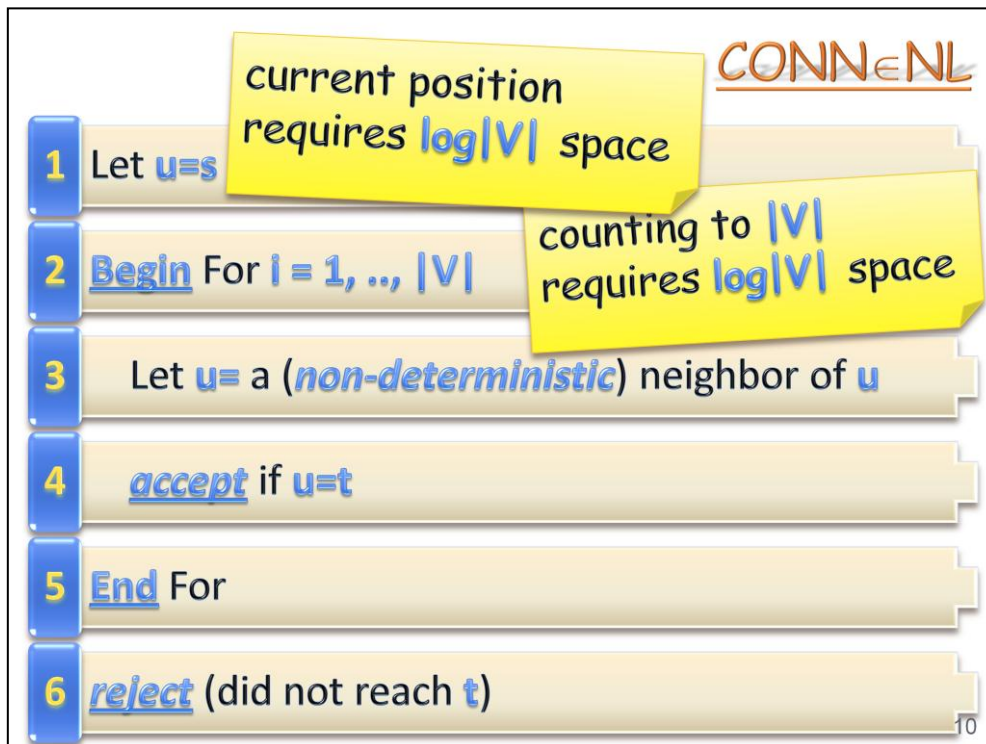
What  
about  
undirected



9

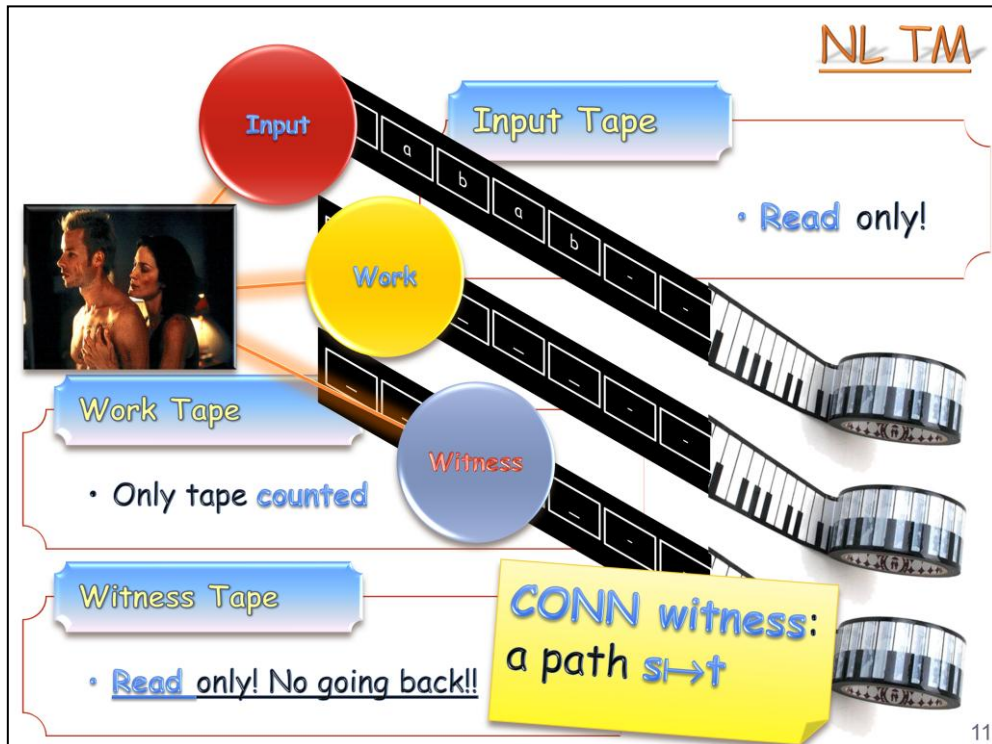
Let us now formally define the connectivity problem: given a graph, a start vertex, and a target vertex, is there a path from start to target?

Do you think the same problem on an undirected graph is easier?



Let us first see that connectivity is in non-deterministic LOGSPACE.

A non-deterministic algorithm for connectivity maintains a pointer to a vertex of the graph. Initially it points to the start vertex. At every stage, the algorithm chooses an edge going out of the vertex it points to, and direct its pointer to the vertex the edge leads to. If it reaches the target, it accepts. If it went too many stages, it rejects.



An alternative formulation of non-deterministic space bounded machines is by introducing the witness tape. The machine can only read that tape and moreover must read it bit by bit and never go back. It is enough that there exists one possible assignment to the content of the tape that causes the machine to accept, for the input to be accepted.

What complexity class do we get if we allow the machine to go back on the witness tape?

# CONN is NL-Complete


## Theorem:

- CONN is NL-hard

Assume a TM  
has 1 accepting  
configuration

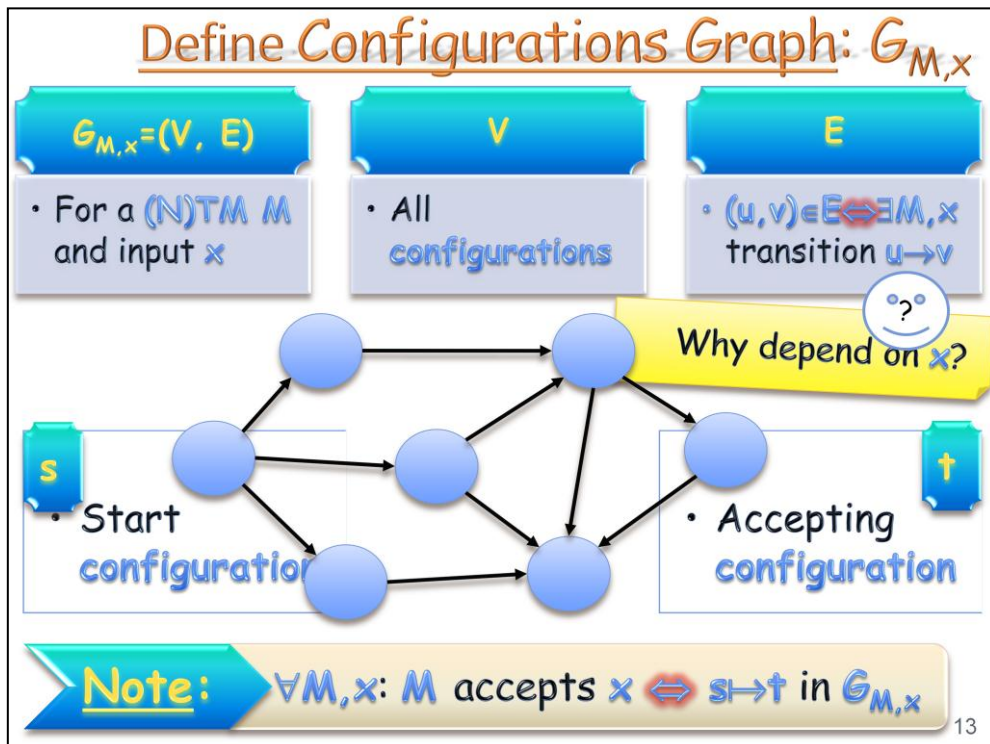
## Proof:

- Given  $M, x$ , construct  
in LOGSPACE a  
CONN instance



12

It turns out that connectivity is in non-deterministic LOGSPACE complete. We will show how to construct the connectivity instance given a machine  $M$  and input  $X$ , so that the machine accepts its input if and only if the instance is in connectivity.



For that purpose let us introduce the configurations' graph: vertexes correspond to configurations, edges to transitions, the start vertex correspond to the start configuration, and the target vertex corresponds to the accepting configuration. An accepting computation of the machine corresponds to a path from start to target, while such a path clearly corresponds to accepting computation.

## CONN is NL-Complete

### Proof (end):


- $G_{M,x}$  can be constructed in Log-Space

### Corollary:

- $NL \subseteq P$

### Proof:

- $CONN \in P$  ■

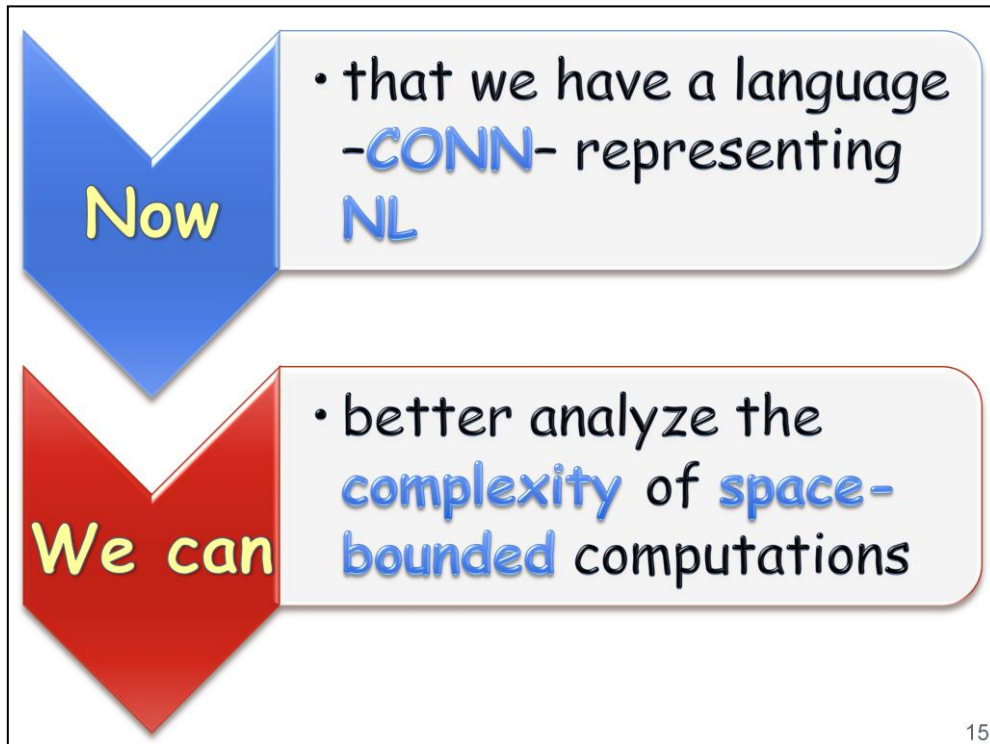


14

Given a non-deterministic LOGSPACE machine, its configuration graph can be computed with logarithmic memory: the algorithm simply needs to compute, given two configurations, whether there is a transition from one to the other.

As a corollary we get that non-deterministic LOGSPACE is contained in P.





- that we have a language  
-**CONN**- representing  
**NL**
- better analyze the  
**complexity** of **space-**  
**bounded** computations

15

The fact that connectivity is NL-complete is fundamental in analyzing space complexity classes: it is crucial in the proof of the following two fundamental theorems we prove.

## Savitch's Theorem

### Theorem:

- $\forall S(n) \geq \log(n): \text{NSPACE}[S(n)] \subseteq \text{SPACE}[S(n)^2]$



NPSPACE=PSPACE

### Proof:

- First  $\text{NL} \subseteq \text{SPACE}[\log^2 n]$  then generalize

### Lemma:

- $\text{NL} \subseteq \text{DSPACE}[\log^2 n]$

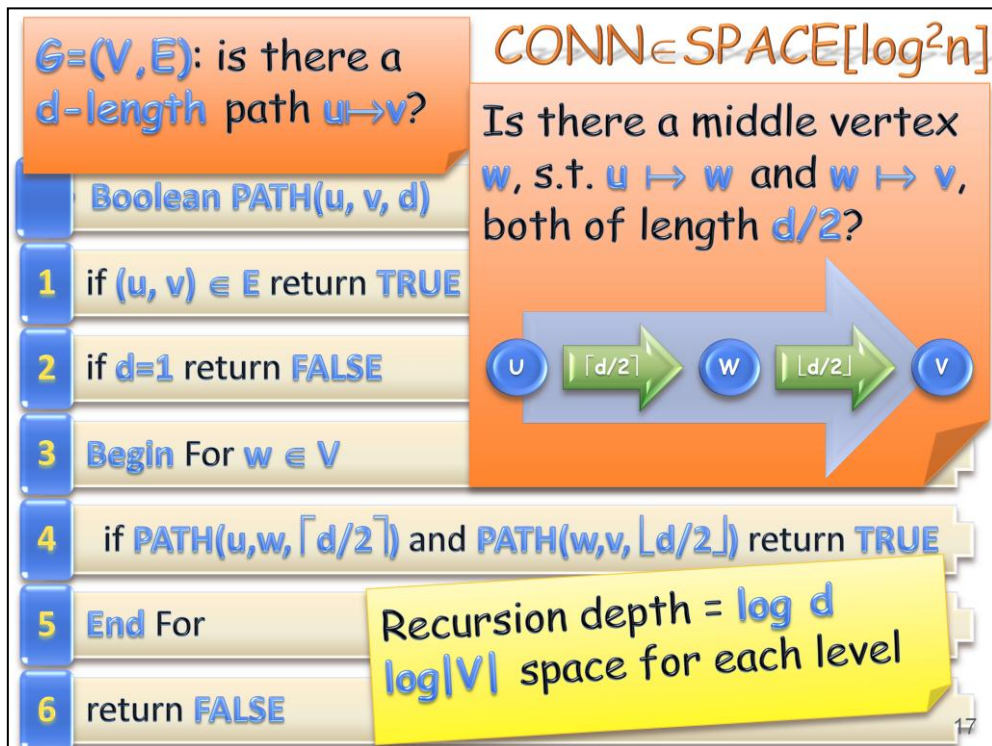


### Proof:

- Suffice to show  $\text{CONN} \in \text{DSPACE}[\log^2 n]$

16


The first is a theorem by Savitch concerning the overhead involved in converting a non-deterministic computation to a deterministic one. It turns out that the overhead in terms of space is not that large, it is in fact quadratic. To prove that theorem, we will start with the special case of NL, and proceed to show a general technique of how to extend such statements for smaller classes to larger classes.



Savitch's deterministic simulation algorithm for connectivity is recursive: to decide if there is a path of length  $d$ , it goes over all possible vertexes for the middle of the path, and call itself to decide whether the appropriate paths of half the lengths exist. Namely, one from the start vertex to the middle vertex, and another from the middle of vertex to the target vertex.

The recursion depth is logarithmic in the length of the path, and at each level the algorithm maintains a pointer to one vertex.

## Example of Savitch's algorithm



```
boolean PATH(a,b,d) {  
  if there is an edge from a to b then  
    return TRUE  
  else {  
    if (d=1) return FALSE  
    for every vertex v (not a,b) {  
      if PATH(a,v, ⌈d/2⌉) and  
        PATH(v,b, ⌊d/2⌋) then  
        return TRUE  
    }  
    return FALSE  
  }  
}
```

(a,b,c)=Is there a path from a to b, that takes no more than c steps.

(1,4,3) TRUE

$3\log_2(d)$

Complexity

18

Here is a simulation of the algorithm on a simple example.

## $O(\log^2 n)$ -Space DTM for NL

### Proof (Lemma, end):

- To solve **CONN**: call **PATH**( $s, t, |V|$ )

### Have:

- $NL \subseteq SPACE(\log^2 n)$

### Want:

- $\forall S(n) \geq \log n$   
 $NSPACE(S(n)) \subseteq SPACE(S^2(n))$

19

To solve connectivity, one can simply apply the algorithm with a number of vertexes and the length of the path.

Now that we have proven the Theorem for NL, we need to extend it to general classes. Namely, show that for every space bound, the cost of translating a non-deterministic algorithm to a deterministic one is quadratic.

## Scale up

### Claim:

- For any two space constructible functions  $s_1(n)$ ,  $s_2(n) \geq \log n$ ,  $e(n) \geq n$ :

### If

- $\text{NSPACE}[s_1(n)] \subseteq \text{SPACE}[s_2(n)]$

### Then

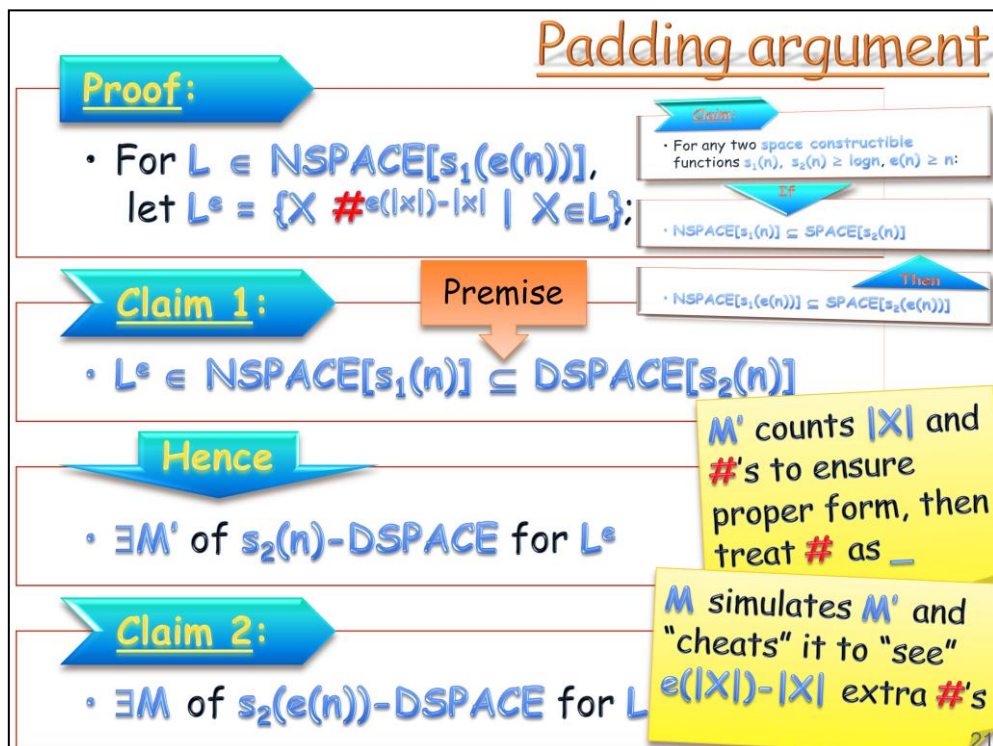
- $\text{NSPACE}[s_1(e(n))] \subseteq \text{SPACE}[s_2(e(n))]$

20

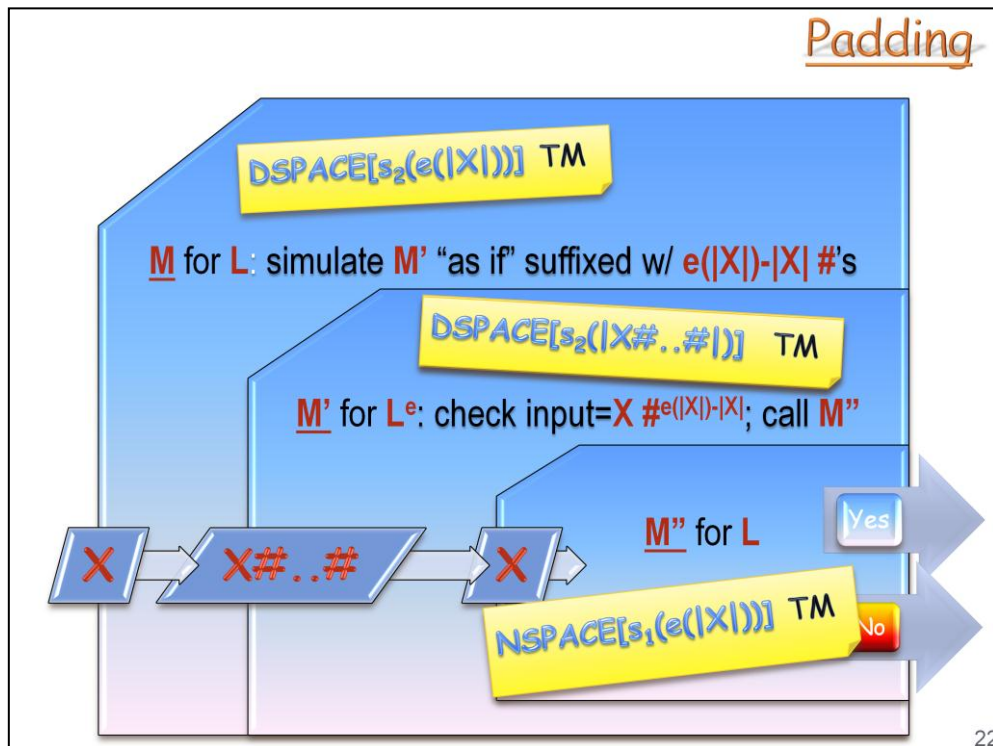
We show a more general principle, that any such relation between models and bounds can be scaled up with a super linear extension function. The extension function scales up both bounds.

This technique is simple yet tricky and is referred to as the padding argument.



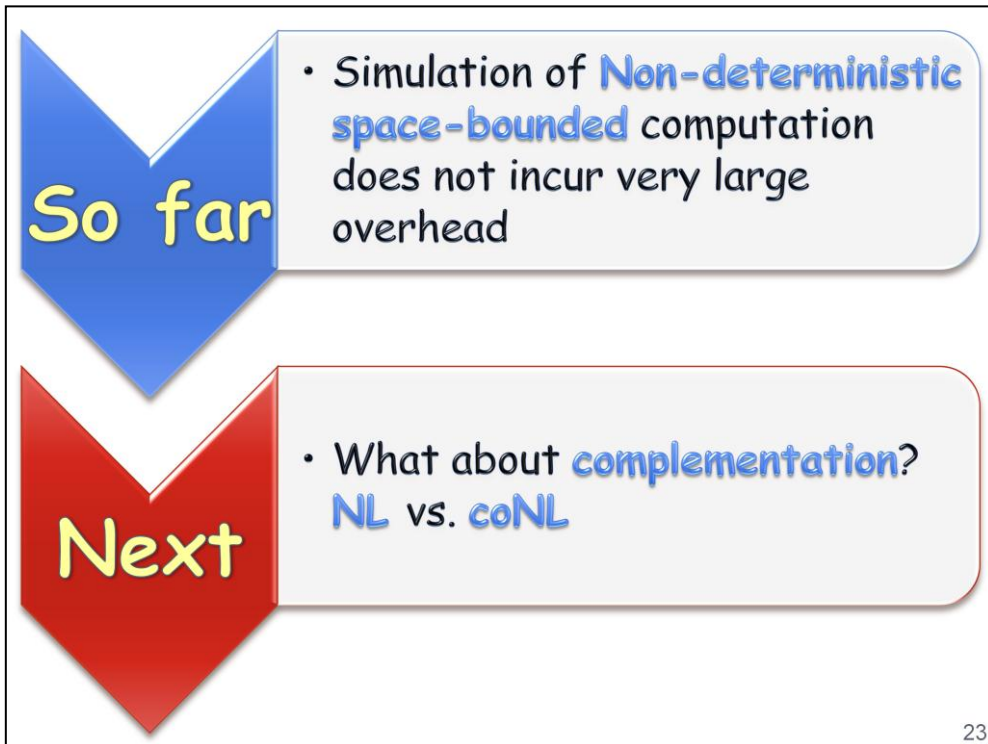


The padding argument goes as follows: given a language  $L$ , accepted by a non-deterministic TM, define the language  $L^e$  that comprises all strings in  $L$  padded with the appropriate number of  $\#$ . That padding makes the language  $L^e$  in the appropriate non-deterministic class. Now, one can apply the containment of the premise and obtain a determined TM for  $L^e$ . This deterministic TM verifies that the number of  $\#$ 's is appropriate with respect to the size of the "real" input. One can in turn, given only the real input, simulate this machine maintaining a counter of the number of  $\#$ 's, and letting the TM work as if the appropriate number of  $\#$ 's is appended to the real input.



Here's an illustration of the construction: we start with a TM  $M''$  for  $L$ , which can be converted into a TM for  $L_e$  (checking that the number of  $\#$ 's is appropriate can be carried out in LOGSPACE), which by the assumption of the premise can be made deterministic --- that's the TM  $M'$ .  $M$  is a TM for  $L$  of appropriate space that simulates  $M'$ , and if  $M'$  wonders off to the  $\#$  section, it maintains a pointer (it has enough space to do that) to where it is and simulates it as if the  $\#$ 's are there.

This completes the proof of Savitch's theorem.

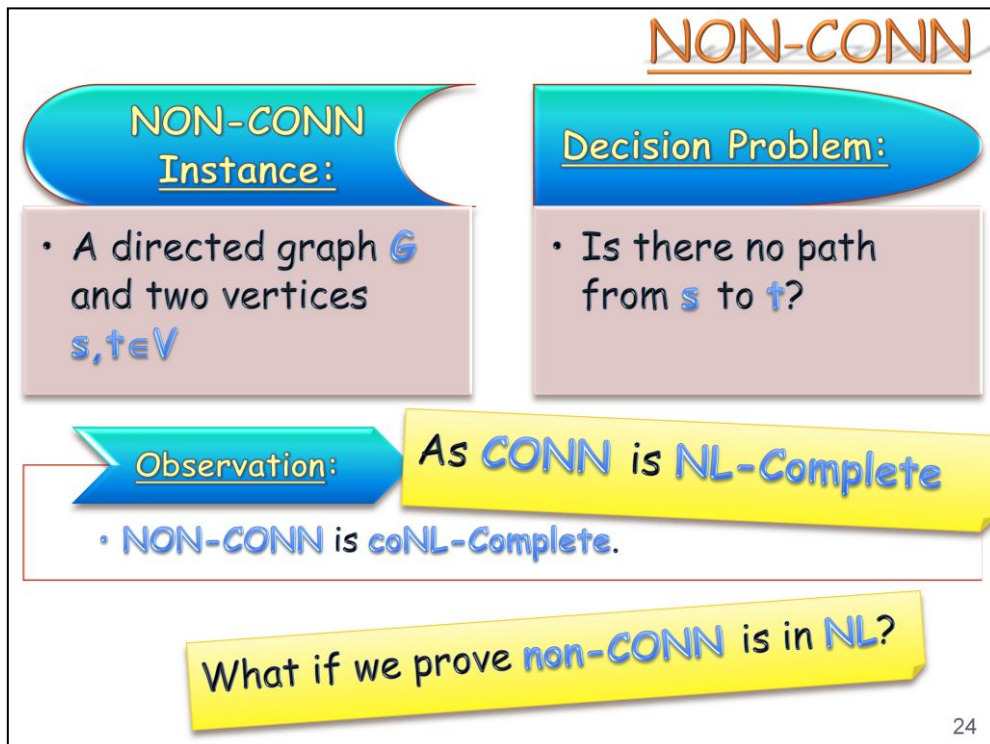


- Simulation of **Non-deterministic space-bounded** computation does not incur very large overhead
- What about **complementation?**  
**NL** vs. **coNL**

23

We have just seen that enhancing space-bounded computation with non determinism does not make it so much stronger.

Next, we look at another aspect by which non determinism for space bounded computations has a limited effect.



Let us first define the non-connectivity problem, which is simply the complement of the connectivity problem.

Non-connectivity is clearly coNL-complete, therefore, it represents the entire coNL class.

It follows, that if we show non-connectivity is in NL, we've proven  $NL = coNL$ .



How do we  
know  
everyone's  
in  
attendance

When your name's called

All those present...

Raise your hands!

Count how many



## Immerman/Szelepcsényi: $\text{coNL} = \text{NL}$

**Theorem:**

- $\text{Non-CONN} \in \text{NL}$

**Proof:**

**Def**  $\text{reachable}(G) = \{ v \mid s \rightarrow^* v \}$

**Def** let  $G_{-t} = (V, E - V \times \{t\})$

**Witness:**  $\# \text{reachable}(G) = \# \text{reachable}(G_{-t})$

**Suffice:**  $\# \text{reachable}(G) = r$

**Def**  $\text{reachable}_l = \{ v \mid s \rightarrow^* v \text{ of length } \leq l \}$

**Induction:**  $r_l = \# \text{reachable}_l$  Base:  $r_0 = 1$

Describe an NL-verifiable witness **W** that there is no  $s \rightarrow^* t$

**W#r#**

26

To show that non-connectivity is in NL, we can use the witness formulation of NL, where the TM can read a witness of membership, from left to write, and verify it indeed prove the input is in a given language.

We define the set of reachable vertexes, namely those that can be reached by a directed path from the start vertex.

To show there is no path from start to target, we can show that the size of the reachable set is the same for the graph and for the same graph only where all edges going into the target are removed.

Hence, it is enough to verify a proof showing what is the number of reachable vertexes of a given graph (first have a proof for the graph, store that number, then verify a proof for the altered graph, and compare the two numbers).

To verify that indeed the number of reachable vertexes is as claimed, the witness can be constructed inductively, over the length of the path.

There is obviously exactly one vertex reachable within 0 steps.

We'll next see how to extend a witness, proving the number of reachable vertexes after  $l$  steps is  $r_l$ , into a witness for  $l+1$ , and so that if the prefix can be verified by a LOGSPACE TM then so is the entire witness.



- W#r<sub>1</sub>#**



The answer is the crux of the entire argument and is as follows: the NL TM verifies that the number of vertexes listed as reachable within  $l$  steps is exactly

$R_l$ , the number proven in  $W$  to be the number of reachable vertexes within  $l$  steps!

## N.D. Algorithm for $\text{reach}_s(v, l)$

$\text{reach}_s(v, l)$

1. **length** =  $l$ ;  $u = s$
2. **while** (**length** > 0) {
  3. **if**  $u = v$  **return** 'YES'
  4. **else, for all** ( $u' \in V$ ) {
    5. **if** ( $u, u' \in E$ ) **nondeterministic switch**:
      - 5.1  $u = u'$ ; **--length**; **break**
      - 5.2 **continue**
6. **return** 'NO'


Takes up logarithmic space

This N.D. algorithm might never stop

## N.D. Algorithm for $CR_s$

$CR_s(d)$

```
1. count = 0
2. for all  $u \in V$  {
  3.  $count_{d-1} = 0$ 
  4. for all  $v \in V$  {
    5. nondeterministic switch:
      5.1 if reach( $v, d - 1$ ) then ++countd-1 else fail
          if  $(v, u) \in E$  then ++count; break
      5.2 continue
  }
  6. if  $count_{d-1} < CR_s(d-1)$  fail
}
7. return count
```



Assume  $(v, v) \in E$

Recursive call!

## N.D. Algorithm for CR

$CR_s(d, C)$

1.  $count = 0$

2. **for all**  $u \in V$  {

3.  $count_{d-1} = 0$

4. **for all**  $v \in V$  {

5. **nondeterministic switch:**

5.1 **if**  $reach(v, d - 1)$  then  $++count_{d-1}$  **else** fail

**if**  $(v, u) \in E$  then  $++count$ ; **break**

5.2 **continue**

}

6. **if**  $count_{d-1} < C$  **fail**

}

7. **return**  $count$

Main Algorithm:

$CR_s$

$C = 1$

**for**  $d = 1..|V|$

$C = CR(d, C)$

**return**  $C$

**parameter**

Complexity

0

## Corollary

- Space-bounded computation classes **closed** under **complementation**:  
 $\forall s(n) \geq \log(n)$ :  
 $\text{NSPACE}(s(n)) = \text{coNSPACE}(s(n))$

padding argument

## Next

- A basic problem **complete** for **PSPACE**



## TQBF

### Instance:

- a fully quantified Boolean formula  $\phi$

### Decision Problem:

- Is  $\phi$  true?

EG

$\forall x \exists y \forall z [(x \vee \neg y \vee z) \wedge (\neg x \vee y)]$

### Theorem:

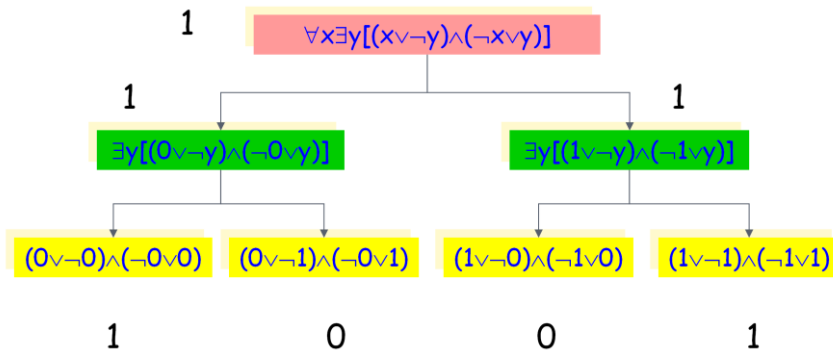
- $\text{TQBF} \in \text{PSPACE}$

### Proof:

- A poly-space algorithm  $A$  that evaluates  $\phi$ :
  - if  $\phi$  is quantifier-free return its **value**
  - if  $\phi = \forall x. \psi(x, \dots)$  return  $A(\psi(0, \dots)) \wedge A(\psi(1, \dots))$
  - if  $\phi = \exists x. \psi(x, \dots)$  return  $A(\psi(0, \dots)) \vee A(\psi(1, \dots))$

32

# Algorithm for TQBF



# TQBF is PSPACE-Complete

## Theorem:


- TQBF is PSPACE-hard

## Proof:

- For a TM  $M$ , start with a BF  $\text{transition}_M(\underline{u}, \underline{v}) \Leftrightarrow$  on  $\underline{u}$   $M$  moves to  $\underline{v}$   
Construct, inductively, the BF  $\phi_M(\underline{u}, \underline{v}, d) \Leftrightarrow$  on  $\underline{u}$ ,  $M$  arrives at  $\underline{v}$  in  $\leq 2^d$  steps:  
 $\phi_M(\underline{u}, \underline{v}, 0) = \text{transition}_M(\underline{u}, \underline{v}) \vee \underline{u} = \underline{v}$   
 $\phi_M(\underline{u}, \underline{v}, d) = \exists \underline{w} \forall \underline{u}' \forall \underline{v}'$   
 $[ ((\underline{u}' = \underline{u} \wedge \underline{v}' = \underline{w}) \vee (\underline{u}' = \underline{w} \wedge \underline{v}' = \underline{v})) \Rightarrow \phi_M(\underline{u}', \underline{v}', d-1) ]$
- $f(M, x) = \phi_M(\text{start}[x], \text{accept}, \lceil \lg \# \text{ of config.} \rceil)$

$\underline{u}, \underline{v}$  vectors describing configurations  
 $\text{transition}_M(\underline{u}, \underline{v})$   
 almost equality

## Synopsis



Defined space-complexity, in particular, the complexity classes: **L**, **NL**, **coNL**, **PSPACE**.  
Proved:


## Completeness:

CONN for NL; TQBF for PSPACE

## Savitch's theorem ( $NL \subseteq SPACE(\log^2)$ )

## The padding argument (scaling up)

## Immerman's theorem ( $NL=coNL$ )



Space Complexity

Savitch's Theorem

WWinindex

Log Space Reductions

Immerman's Theorem



Savitch, Walter

TQBF



Immerman, Neil

Complexity Classes

PSPACE



**Róbert Szelepcsényi**

L

NL