
***On Compact Representations of
All-Pairs-Shortest-Path
Distance Matrices***

Igor Nitto and Rossano Venturini

Dipartimento di Informatica

University of Pisa, Italy

Outline

- Problem statement
- A simple information-theoretic lower bound
- A first compact representation
- A more space-efficient representation

Problem Statement

- **Given:** an unweighted undirected graph
- **Build:** a static data structure encoding distance matrix **D**
 - **D**[u,v] := shortest path distance from u to v
- **Using:** $O(\text{minimum possible number of bits})\dots$
- **Supporting:** access to any **D**[u,v] in $O(1)$ time
- **Model:** Word RAM

Problem Statement

- **Given:** an unweighted undirected graph with n nodes and m edges. The graph is sparse, i.e., $m = O(n)$. The graph is connected. The graph is undirected.
- **Build:** a static data structure that stores the graph and supports the following operations:
 - $D[u,v]$:= shortest path distance from u to v
- **Using:** $O(\text{minimum possible number of bits})$ bits of space. The space complexity is measured in terms of the information-theoretic minimum $\frac{n(n-1)}{2} \sim O(1)$ bits per element.
- **Supporting:** access to any $D[u,v]$ in $O(1)$ time
- **Model:** Word RAM

Comparison with related works

- Approximate Distance Oracles
 - Distributed data structures
 - Fast **estimate** of distances in small space
 - $O(n^{1+1/t} \log n)$ bits and $O(t)$ multiplicative error
[Thorup-Zwick, 2001]
- Our result: **exact** distances in $O(1)$ time and $O(n^2)$ bits
 - First result offering such space/time tradeoff

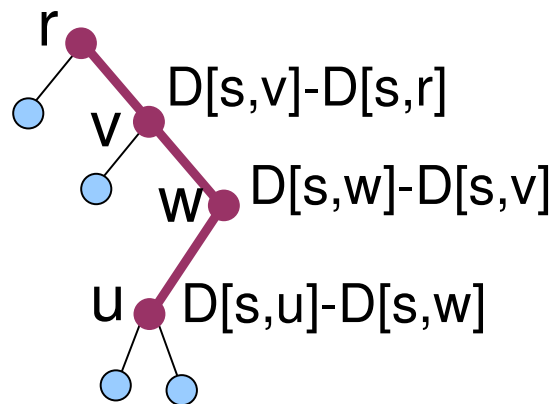
Single Source Distances Encoding

Let T be a spanning tree, rooted at an arbitrary node r

- (assume the graph connected)

Store separately each row $\mathbf{D}[s, \dots]$ in $O(n)$ bits:

- Define the following labelling \mathbf{T}_s of T :
 - $\mathbf{T}_s(r) = \mathbf{D}[s, r]$
 - $\mathbf{T}_s(u) = \mathbf{D}[s, u] - \mathbf{D}[s, f(u)]$ ($\in \{-1, 0, 1\}$) for $u \neq r$
- $\mathbf{D}[s, u]$ equals to *path-sum* at u in \mathbf{T}_s




Single Source Distances Encoding

Let T be a spanning tree, rooted at an arbitrary node r

- (assume the graph connected)

Store separately each row $\mathbf{D}[s, \dots]$ in $O(n)$ bits:




- Define the following labelling \mathbf{T}_s of T :
 - $\mathbf{T}_s(r) = \mathbf{D}[s, r]$
 - $\mathbf{T}_s(u) = \mathbf{D}[s, u] - \mathbf{D}[s, f(u)]$ ($\in \{-1, 0, 1\}$) for $u \neq r$
- $\mathbf{D}[s, u]$ equals to *path-sum* at u in \mathbf{T}_s



Encoding $\mathbf{D}[s, \dots]$
 \Leftrightarrow
Store \mathbf{T}_s with fast path-sum

From path- to prefix-sum




- Solution: “linearize” T_S in a sequence L_S of length $2n$



$r[] =$ starting times in DFS

From path- to prefix-sum

- Solution: “linearize” T_S in a sequence L_S of length $2n$




$r[]$ = starting times in DFS

Path-sum at node u in T_S boils down to a prefix-sum at $r[u]$ in L_S

From path- to prefix-sum

- **Solution:** “linearize” T_S in a sequence L_S of length $2n$



$r[]$ = starting times in DFS

- Build *wavelet tree* [Grossi et al., 2003] on L_S (root label apart)
 - **Space:** $2n\gamma + o(n)$ bits, $\gamma = \log_2 3$ (magic constant...)
 - Prefix-sums in $O(1)$ time via *rank* computations


A First Compact Encoding

$$\gamma = \log_2 3$$

- Encoding of distance matrix $\mathbf{D}[1..n,1..n]$:
 - Array $\mathbf{r}[]$: $O(n \log n)$ bits
 - Prefix-sums on each \mathbf{L}_s : $(2\gamma)n + o(n)$ bits
 - **Space**: $\gamma n^2 + o(n^2)$ bits, **Access time**: $O(1)$
- This does not match information-theoretic optimum:
 - γn^2 vs $n^2/2$ bits
 - Linearizations waste space!
 - Need to store two elements per label
- Improved approach: reduce γ to $\gamma/2$
 - Combine above encoding with *macro-micro tree partition*


Macro-Micro Tree Partition *[Bender et al.,2000]*

- Depending on a parameter α ($= c \log n$, with $c < 1$), partition the nodes of T into:



Macro-Micro Tree Partition [Bender et al., 2000]


- Depending on a parameter α ($= c \log n$, with $c < 1$), partition the nodes of T into:



- Jump-Nodes: maximally deep nodes having more than α descendants

Macro-Micro Tree Partition [Bender et al., 2000]


- Depending on a parameter α ($= c \log n$, with $c < 1$), partition the nodes of T into:



- Jump-Nodes: maximally deep nodes having more than α descendants
- Macro-Nodes: any ancestor of a jump-nodes

Macro-Micro Tree Partition [Bender et al., 2000]

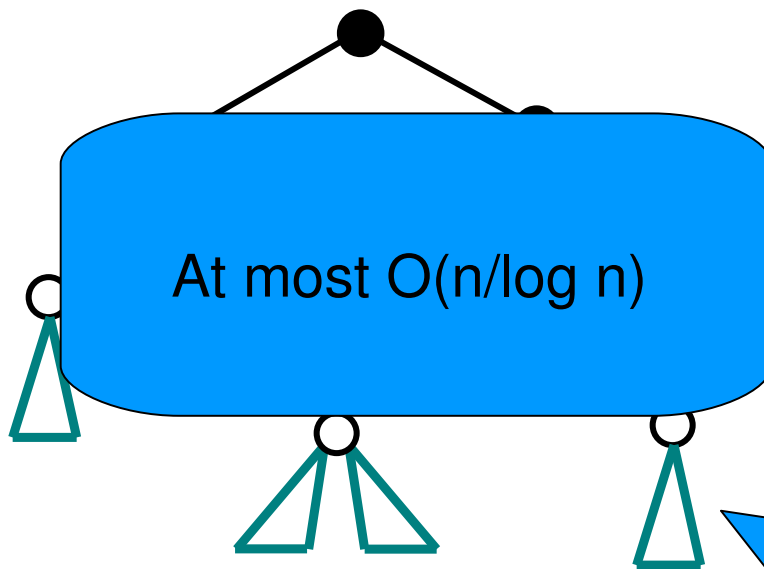
- Depending on a parameter α ($= c \log n$, with $c < 1$), partition the nodes of T into:



- Jump-Nodes: maximally deep nodes having more than α descendants
- Macro-Nodes: any ancestor of a jump-nodes
- Micro-Nodes: any other node

Macro-Micro Tree Partition [Bender et al., 2000]

- Depending on a parameter α ($= c \log n$, with $c < 1$), partition the nodes of T into:



- Jump-Nodes: maximally deep nodes having more than α descendants


- Macro-Nodes: any ancestor of a jump-nodes

- Micro-Nodes: the nodes of a micro-tree

Micro-tree: maximal connected component of micro-nodes. It contains at most α nodes

Path-Sum on Macro-Micro Partition


- Some additional definitions:



- Branching-Nodes: more than one macro or jump nodes among children.
- Unary-Nodes: any other macro node

Path-Sum on Macro-Micro Partition

- Some additional definitions:




- Branching-Nodes: more than one macro or jump nodes among children.
- Unary-Nodes: any other macro node

In every maximal path of unary nodes, *sample* one out of $\log n$ nodes.

Path-Sum on Macro-Micro Partition

- Some additional definitions:



At most $O(n/\log n)$

- Branching-Nodes: more than one macro or jump nodes among children.
- Unary-Nodes: any other macro node

At most $O(n/\log n)$

In every maximal path of unary nodes, *sample* one out of $\log n$ nodes.

Path-Sum on Macro-Micro Partition

- Some additional definitions:

Breaking nodes: the collection of branching, jump and sampled unary nodes.


At most $O(n/\log n)$

- Branching-Nodes: more than one macro or jump nodes among children.
- Unary-Nodes: any other macro node


At most $O(n/\log n)$

In every maximal path of unary nodes, *sample* one out of $\log n$ nodes.

Path-Sums on Breaking Nodes




Path-Sums on Breaking Nodes




- Build T_F from T by collapsing paths of unary non-breaking nodes
- “Contract” labeling of T onto T_F
- Use a variant of the previous solution to provide path-sums on T_F .
- Space required : $O((n/\log n) \log \log n) = o(n)$ bits

Path-Sums on Unary Nodes



- \mathbf{M}_T : sequence of macro-nodes labels sorted according to DF- visit
- To compute path-sum(u):
 - u' : closest breaking ancestor of u
 - Sum in $O(1)$ time labels from u to u'
 - By prefix-sums on \mathbf{M}_T


Path-Sums on Unary Nodes



- \mathbf{M}_T : sequence of macro-nodes labels sorted according to DF- visit
- To compute path-sum(u):
 - u' : closest breaking ancestor of u
 - Sum in $O(1)$ time labels from u to u'
 - By prefix-sums on \mathbf{M}_T
- Add to path-sum(u')

Path-Sums on Unary Nodes

Single label per element,
gain factor 2 in space




■ \mathbf{M}_T : sequence of macro-nodes labels sorted according to DF- visit

To compute path-sum(u):

- u' : closest breaking ancestor of u
- Sum in $O(1)$ time labels from u to u'
 - By prefix-sums on \mathbf{M}_T


■ Add to path-sum(u')

Path-Sums on Micro Nodes




- Tabulate in C all path-sums:
 - for any possible microtree
 - for any possible labeling in $\{-1,0,1\}$
- Table C requires $o(n)$ bits
- Store the sequence L of micro-tree labelings
 - $O(1)$ time access to them

Path-Sums on Micro Nodes



- To answer $\text{path-sum}(u)$:
 - q : root of micro-tree containing u
 - Sum of labels from u to q in $O(1)$ time using C and L

Path-Sums on Micro Nodes



- To answer $\text{path-sum}(u)$:
 - q : root of micro-tree containing u
 - Sum of labels from u to q in $O(1)$ time using C and L
 - Add $\text{path-sum}(f(q))$

Conclusions

- A storage scheme for all-pairs-shortest-path-distance matrices of unweighted graphs:
 - space: $(\log_2 3)/2 n^2 + o(n^2)$ bits, access time: $O(1)$
- Open:
 - match the information-theoretic lower bound
 - space dependent on the number of edges
 - Impact on Web and Social Networks applications

Thanks.