A Local Approximation Algorithm for Maximum Weight Matching

Tim Nieberg
Research Institute for Discrete Mathematics
University of Bonn

Overview

- introduction
- $\mathcal{LOCAL}$ model for distributed communication networks
  - locality of graph structures
- weighted matchings
  - matchings, connected ℓ-augmentations, and their gain
  - ℓ-augmentation graph
  - Algorithm: ImproveMatching M
- $(1-\varepsilon)$-approximation
- wireless communication networks
  - preprocessing: colored cluster-graphs
  - decreasing the runtime
- conlusions

Matchings

Definition

A *matching* in a graph $G = (V, E)$ is a subset $M \subseteq E$ such that no two edges in M share a common node.

We look at the weighted version, where each edge is given a (non-negative) weight.

- ... and seek a matching of largest weight.

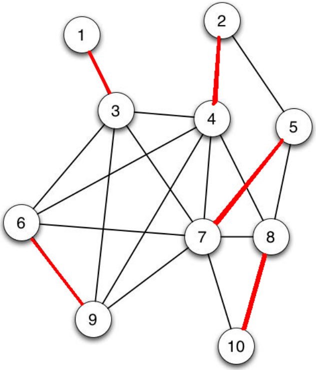

Local Communication Model

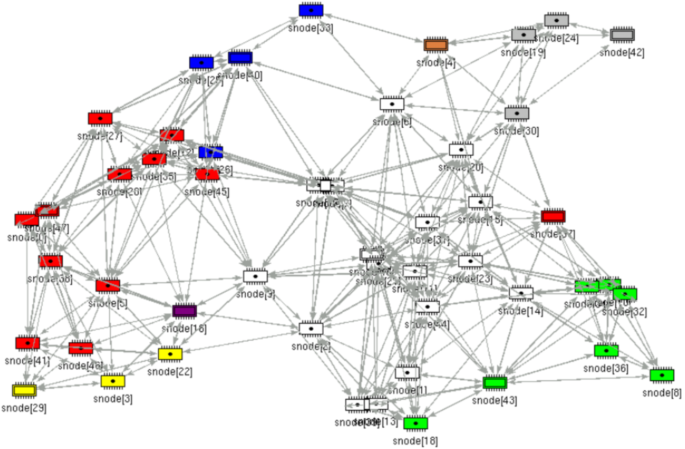

Distributed Network Model

Consider a network $G = (V, E)$, where each node $v \in V$ is eqipped with

- CPU, memory, and
- communication capabilities (e.g. wireless transceiver).

Let E denote the possible communication links.

- Each node is independent, and can locally participate in a distributed algorithm.

We now want to characterize distributed algorithms such that we can make statements about protocols running in the network.

- Note: we communicate *in* and we optimize *for* G !

The $\mathcal{LOCAL}$ Model

Networking operates in *global communication rounds*.
In each round, a node can

- communicate with its direct neighbors (Phase 1), and
- perform some local computations (Phase 2).

The order, in which the message packets are sent is not specified (assume simultaneously).
Simple consequences of the $\mathcal{LOCAL}$ model:

- consider two nodes $u, v \in V$ with $d(u, v) = k$: it takes at least k rounds for a message from v to arrive at node u!
- it takes $O(r)$ rounds for a node to learn about its r-neighborhood.

Complexity Measures

There are three complexity measures for local, distributed algorithms in the $\mathcal{LOCAL}$ model:

- *time complexity*
  - number of rounds until all nodes have terminated the algorithm
- *message complexity*
  - number of messages sent during execution of the algorithm
  - usually given with respect to a single node in the network
- *maximum message size*
  - largest message packet sent in a round
  - gives the amount of information exchanged
  - $\Omega(\log n)$

Locality of Graph Structures

The $\mathcal{LOCAL}$ model is also interesting in terms of theory:

- exploit *locality* of the graph structures
- focus on a fraction of the instance
- typical question:
  - What type of *local* information is necessary and/or sufficient to create/decide on a *global* solution?
- note: many greedy-approaches are based on local decisions
- trivial: allow $O(n)$-neighborhoods

Maximum Weight Matching:

- global perspective
  - see [Edmonds 1965]
- local perspective
  - local information only not sufficient!
  - (closer look at matching-polytope)

Augmentations of Matchings

Definition

Given a matching $M \subset E$, we call another matching $S \subseteq E \setminus M$ an *augmentation* for M.

For such an augmentation S,

- denote by $M(S) \subset E$ all edges in M that have a node in common with an edge from S
- $\Rightarrow$ $(M \setminus M(S)) \cup S$ again is a matching
  - M augmented by S
- the *size* is given by the number of edges in S
- S *connected* $\iff$ $M(S) \cup S$ is single component in G
  - connected augmentation is either single path or cycle in G
- $\text{gain}_M(S)$ is the difference in weight between M and $(M \setminus M(S)) \cup S$
  - $\text{gain}_M(S) = w(S) - w(M(S))$

Augmentations

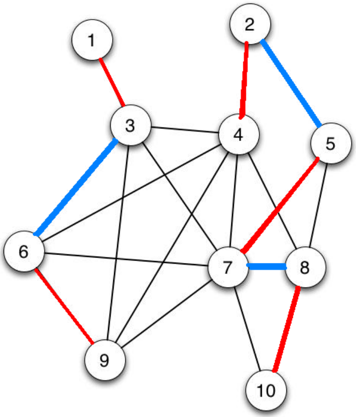

Denote by

- $w_{\max} := \max\{w_e \mid e \in E\}$
- $\mathrm{gain}^{\ell}_{\max} := \max\{\mathrm{gain}_M(S) \mid S \text{ augmentation of size at most } \ell\}$

l-Augmentation Graph

Let $\ell \in \mathbb{N}$ be some constant.

Definition

The *ℓ-augmentation graph* $G' = (V', E')$ (of a graph G w.r.t. a matching M) is defined as the intersection graph of connected augmentation of size at most ℓ in G:

- the nodes V' are all connected augmentations of size $\leq \ell$,
- two nodes are connected if the respective augmentations share a common node.

For each augmentation in G', we call the node with the lowest identifier in the augmentation its *representative*:

- this maps G' to G
- communication along an edge in G' takes $O(\ell) = O(1)$ rounds in G

We can easily and locally construct G' in $O(1)$!
It is: $|V'| = O(n^{2\ell})$.

Algorithm to Improve a Matching

We now restrict our attention to the ℓ-augmentation graph G'. Given a matching M (possibly empty), we improve M by

- looking at all favorable augmentations
  - that is, with *high gain*
- selecting those that can be used in a parallel approach
  - that is, they do not overlap
- augment M in parallel

We then repeat this improvement algorithm to receive the final algorithm.

Algorithm to Improve a Matching

Algorithm: Improve Matching M

Construct ℓ-augmentation graph $G' = (V', E')$
$\mathcal{A} := \varnothing$
$V^{(1)} := V'$
for $t := 1$ to $\lceil \log_2 \ell^2 n \rceil$ **do**
  $W := \{v \in V^{(t)} \mid$
    $\Gamma(v) \cap \{u \in V^{(t)} \mid \text{gain}(u) > 2\text{gain}(v)\} = \varnothing\}$
  Calculate MIS I in $G'(W)$
  $\mathcal{A} := \mathcal{A} \cup I$
  $V^{(t+1)} := V^{(t)} \setminus \Gamma(I)$
end for
$M' := M$ augmented by $\mathcal{A}$

Algorithm to Improve a Matching

Lemma

The set $\mathcal{A}$ is an independent set in G'.

- follows from construction

Theorem

The set M' computed in the algorithm is a matching in G.

- no two augmentations in $\mathcal{A}$ overlap
- note: M' constructed in parallel

Theorem

Let T_{MIS} denote the distributed time to construct a MIS. Then, the algorithm has runtime $O(\ell + \log(\ell^2 n) \cdot T_{MIS}(n^{O(\ell)}))$.

Gain of M' over M

Lemma

After $c = \lceil \log_2 \ell^2 n \rceil$ *iterations of the for-loop,*

$$\max\{gain(v) \mid v \in V^{(c+1)}\} < \frac{w_{\max}}{\ell n}$$

holds.

- $\max\{\text{gain}_M(v) \mid v \in V'\} \leq \text{gain}^{\ell}_{\max} \leq \ell \cdot w_{\max}$
- claim follows by induction:

$$(1/2)^{\log_2(\ell^2 n)} = 1/(\ell^2 n) \text{ and } w_{\max} \leq \text{gain}_{\max}.$$

Overall Gain

Theorem

$$w(M') \geq w(M) + \frac{1}{8\ell}\left(\frac{\ell-1}{\ell}w(M^*) - w(M)\right),$$

where M^* is an optimal solution.

- $w(M') - w(M) = \text{gain}(\mathcal{A})$
- split $M(M^*)$ into multiple, connected ℓ-augmentations
- use charging argumentation on G' to compare ℓ-augmentations $M(M^*)$ with $M(M')$

Corollary

A single invocation of the algorithm Improve Matching with $M = \varnothing$ yields a constant-factor approximation for the Maximum Weight Matching problem.

$(1-\varepsilon)$-Approximation

Theorem

Let $\ell \in \mathbb{N}$. Calling algorithm Improve Matching ℓ times returns a machting M of weight at least $(1 - O(1/\ell)) \cdot w(M^*)$.

- $M_0 = \varnothing$ and M_i matching of i-th call
- $\Rightarrow$ recursive improvement of $w(M_i) \geq w_i \cdot w(M^*)$ with $w_0 = 0$ and

$$w_{i+1} = w_i + \frac{1}{8\ell}\left(\frac{\ell - 1}{\ell} - w_i\right) w(M^*)$$

- solving the recurrence relation yields

$$w(M_i) \geq \frac{\ell - 1}{\ell}\left(1 - \left(1 - \frac{1}{8\ell}\right)^i\right) w(M^*)$$

- $\Rightarrow$ $i = O(\ell)$ results in claim.

Wireless Communication Topologies

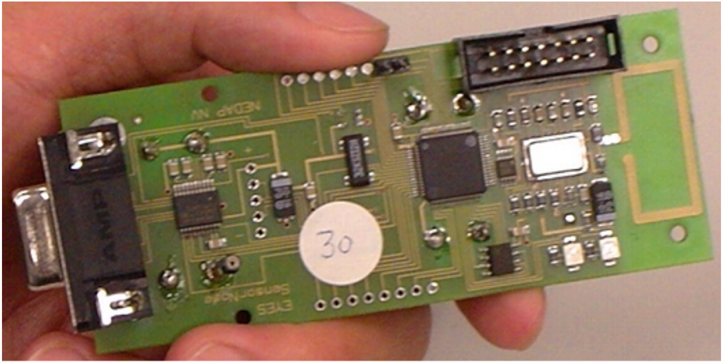

Geometric Intersection Graphs

A geometric intersection graph is given by a collection V of nodes, and for each $v \in V$,

- $f(v)$ center position of node v
- A_v area covered by v's transmitter

Containment Model

$$(u, v) \in E \iff f(u) \in A_v$$

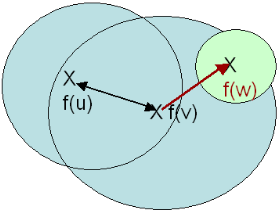

Intersection Model

$$(u, v) \in E \iff A_u \cap A_v \neq \varnothing$$

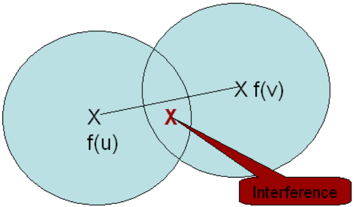

Bounded Growth Graphs

Definition

Let $G = (V, E)$ be a graph. If there exists a function$f(.)$ such that every r-neighborhood in G contains at most $f(r)$ independent vertices, then G is f-*growth-bounded*. In this case, we call f the growth function.

- if the growth function is a polynomial of bounded degree, we say that G has polynomially bounded growth
- note that the growth function only depends on the radius of the neighborhood, and *not* on the number of vertices in the graph
- definition does *not* depend on any geomtric data (e.g. representation)
- bounded growth is closed under taking vertex-induced subgraphs

Matching: The Wireless Case

The above algorithm depends on time to construct MIS on G'.

- on a wireless graph of bounded growth, we can do some preprocessing:
  - construct MIS $\mathcal{I}$ and create clusters ($O(\log \Delta \log^* n)$)
  - color these clusters according to $\Gamma_{4\ell+8}(v), v \in \mathcal{I}$ ($O(\log n)$)
  - $\Rightarrow$ $O(f(4\ell+8)) = O(1)$ colors, and two clusters of same color are non-overlapping w.r.t. ℓ-augmentation they contain
- during algorithm *Improve Matching*: use coloring to construct MIS $\mathcal{A}$ in parallel $\Rightarrow$ $O(1)$ rounds

Overall runtime: $O(\log n \log^* n)$ rounds.

Conclusions

$(1-\varepsilon)$-approximation of Maximum Weight Matching by local, distributed approach:

- $O(\frac{1}{\varepsilon} \log n \cdot T_{\text{MIS}}(n^{O(1/\varepsilon)}))$ communication rounds
  - randomized MIS-construction in $O(\log n)$ [Luby86]
  - $\Rightarrow$ $O(\log^2 n)$ randomized algorithm
- wireless communication networks (bounded growth)
  - preprocessing $\Rightarrow$ $O(\log n \log^* n)$ deterministic algorithm

Construction based on local structure

- connected ℓ-augmentation
- ℓ gives trade-off between locality and quality of solution

EOF

- Thanks for your attention!

nieberg@or.uni-bonn.de