

Chapter 26

Dominating Set

In this chapter we present another randomized algorithm that demonstrates the power of randomization to break symmetries. We study the problem of finding a small dominating set of the network graph. As it is the case for the MIS, an efficient dominating set algorithm can be used as a basic building block to solve a number of problems in distributed computing. For example, whenever we need to partition the network into a small number of local clusters, the computation of a small dominating set usually occurs in some way. A particularly important application of dominating sets is the construction of an efficient backbone for routing.

Definition 26.1 (Dominating Set). *Given an undirected graph $G = (V, E)$, a dominating set is a subset $S \subseteq V$ of its nodes such that for all nodes $v \in V$, either $v \in S$ or a neighbor u of v is in S .*

Remarks:

- It is well-known that computing a dominating set of minimal size is NP-hard. We therefore look for approximation algorithms, that is, algorithms which produce solutions which are optimal up to a certain factor.
- Note that every MIS (cf. Chapter 7) is a dominating set. In general, the size of every MIS can however be larger than the size of an optimal minimum dominating set by a factor of $\Omega(n)$. As an example, connect the centers of two stars by an edge. Every MIS contains all the leaves of at least one of the two stars whereas there is a dominating set of size 2.

All the dominating set algorithms that we study throughout this chapter operate in the following way. We start with $S = \emptyset$ and add nodes to S until S is a dominating set. To simplify presentation, we color nodes according to their state during the execution of an algorithm. We call nodes in S *black*, nodes which are covered (neighbors of nodes in S) *gray*, and all uncovered nodes *white*. By $W(v)$, we denote the set of white nodes among the direct neighbors of v , including v itself. We call $w(v) = |W(v)|$ the *span* of v .

26.1 Sequential Greedy Algorithm

Intuitively, to end up with a small dominating set S , nodes in S need to cover as many neighbors as possible. It is therefore natural to add nodes v with a large span $w(v)$ to S . This idea leads to a simple greedy algorithm:

Algorithm 26.2 Greedy Algorithm

- 1: $S := \emptyset$;
 - 2: **while** there are white nodes **do**
 - 3: $v := \{v \mid w(v) = \max_u \{w(u)\}\}$;
 - 4: $S := S \cup v$;
 - 5: **end while**
-

Theorem 26.3. *The Greedy Algorithm computes a $\ln \Delta$ -approximation, that is, for the computed dominating set S and an optimal dominating set S^* , we have*

$$\frac{|S|}{|S^*|} \leq \ln \Delta.$$

Proof. Each time, we choose a new node of the dominating set (each greedy step), we have cost 1. Instead of letting this node pay the whole cost, we distribute the cost equally among all newly covered nodes. Assume that node v , chosen in line 3 of the algorithm, is white itself and that its white neighbors are v_1, v_2, v_3 , and v_4 . In this case each of the 5 nodes v and v_1, \dots, v_4 get charged $1/5$. If v is chosen as a gray node, only the nodes v_1, \dots, v_4 get charged (they all get $1/4$).

Now, assume that we know an optimal dominating set S^* . By the definition of dominating sets, to each node which is not in S^* , we can assign a neighbor from S^* . By assigning each node to exactly one neighboring node of S^* , the graph is decomposed into stars, each having a dominator (node in S^*) as center and non-dominators as leaves. Clearly, the cost of an optimal dominating set is 1 for each such star. In the following, we show that the amortized cost (distributed costs) of the greedy algorithm is at most $\ln \Delta + 2$ for each star. This suffices to prove the theorem.

Consider a single star with center $v^* \in S^*$ before choosing a new node u in the greedy algorithm. The number of nodes that become dominated when adding u to the dominating set is $w(u)$. Thus, if some white node v in the star of v^* becomes gray or black, it gets charged $1/w(u)$. By the greedy condition, u is a node with maximal span and therefore $w(u) \geq w(v^*)$. Thus, v is charged at most $1/w(v^*)$. After becoming gray, nodes do not get charged any more. Therefore first node that is covered in the star of v^* gets charged at most $1/(d(v^*) + 1)$. Because $w(v^*) \geq d(v^*)$ when the second node is covered, the second node gets charged at most $1/d(v^*)$. In general, the i^{th} node that is covered in the star of v^* gets charged at most $1/(d(v^*) + i - 2)$. Thus, the total amortized cost in the star of v^* is at most

$$\frac{1}{d(v^*) + 1} + \frac{1}{d(v^*)} + \dots + \frac{1}{2} + \frac{1}{1} = H(d(v^*) + 1) \leq H(\Delta + 1) < \ln(\Delta) + 2$$

where Δ is the maximal degree of G and where $H(n) = \sum_{i=1}^n 1/i$ is the n^{th} number. \square

Remarks:

- One can show that unless $\text{NP} \subseteq \text{DTIME}(n^{O(\log \log n)})$, no polynomial-time algorithm can approximate the minimum dominating set problem better than $(1 - o(1)) \cdot \ln \Delta$. Thus, unless $\text{P} \approx \text{NP}$, the approximation ratio of the simple greedy algorithm is optimal (up to lower order terms).

26.2 Distributed Greedy Algorithm

For a distributed algorithm, we use the following observation. The span of a node can only be reduced if any of the nodes at distance at most 2 is included in the dominating set. Therefore, if the span of node v is greater than the span of any other node at distance at most 2 from v , the greedy algorithm chooses v before any of the nodes at distance at most 2. This leads to a very simple distributed version of the greedy algorithm. Every node v executes the following algorithm.

Algorithm 26.4 Distributed Greedy Algorithm (at node v):

```

1: while  $v$  has white neighbors do
2:   compute span  $w(v)$ ;
3:   send  $w(v)$  to nodes at distance at most 2;
4:   if  $w(v)$  largest within distance 2 (ties are broken by IDs) then
5:     join dominating set
6:   end if
7: end while

```

Theorem 26.5. *Algorithm 26.4 computes a dominating set of size at most $\ln \Delta + 2$ times the size of an optimal dominating set in $\mathcal{O}(n)$ rounds.*

Proof. The approximation quality follows directly from the above observation and the analysis of the *greedy algorithm*. The time complexity is at most linear because in every iteration of the while loop, at least one node is added to the dominating set and because one iteration of the while loop can be implemented in a constant number of rounds. \square

The approximation ratio of the above distributed algorithm is best possible (unless $\text{P} \approx \text{NP}$ or unless we allow local computations to be exponential). However, the time complexity is very bad. In fact, there really are graphs on which in each iteration of the while loop, only one node is added to the dominating set. As an example, consider a graph as in Figure 26.6. An optimal dominating set consists of all nodes on the center axis. The *distributed greedy algorithm* computes an optimal dominating set, however, the nodes are chosen sequentially from left to right. Hence, the running time of the algorithm on the graph of Figure 26.6 is $\Omega(\sqrt{n})$. Below, we will see that there are graphs on which Algorithm 26.4 even needs $\Omega(n)$ rounds.

The problem of the graph of Figure 26.6 is that there is a long path of descending degrees (spans). Every node has to wait for the neighbor to the left. Therefore, we want to change the algorithm in such a way that there




Figure 26.6: Distributed greedy algorithm: Bad example




Figure 26.7: Distributed greedy algorithm with rounded spans: Bad example

are no long paths of descending spans. Allowing for an additional factor 2 in the approximation ratio, we can round all spans to the next power of 2 and let the greedy algorithm take a node with a maximal rounded span. In this case, a path of strictly descending rounded spans has at most length $\log n$. For the distributed version, this means that nodes whose rounded span is maximal within distance 2 are added to the dominating set. Ties are again broken by unique node IDs. If node IDs are chosen at random, the time complexity for the graph of Figure 26.6 is reduced from $\Omega(\sqrt{n})$ to $\mathcal{O}(\log n)$.

Unfortunately, there still is a problem remaining. To see this, we consider Figure 26.7. The graph of Figure 26.7 consists of a clique with $n/3$ nodes and two leaves per node of the clique. An optimal dominating set consists of all the $n/3$ nodes of the clique. Because they all have distance 1 from each other, the described distributed algorithm only selects one in each while iteration (the one with the largest ID). Note that as soon as one of the nodes is in the dominating set, the span of all remaining nodes of the clique is 2. They do not have common neighbors and therefore there is no reason not to choose all of them in parallel. However, the time complexity of the simple algorithm is $\Omega(n)$. In order to improve this example, we need an algorithm that can choose many nodes simultaneously as long as these nodes do not interfere too much, even if they are neighbors. In Algorithm 26.8, $N(v)$ denotes the set of neighbors of v (including v itself) and $N_2(v) = \bigcup_{u \in N(v)} N(u)$ are the nodes at distance at most 2 of v . As before, $W(v) = \{u \in N(v) : u \text{ is white}\}$ and $w(v) = |W(v)|$. It is clear that if Algorithm 26.8 terminates, it computes a valid dominating set. We will now show that the computed dominating set is small and that the algorithm terminates quickly.

Theorem 26.9. *Algorithm 26.8 computes a dominating set of size at most*

Algorithm 26.8 Fast Distributed Dominating Set Algorithm (at node v):

```

1:  $W(v) := N(v)$ ;  $w(v) := |W(v)|$ ;
2: while  $W(v) \neq \emptyset$  do
3:    $\tilde{w}(v) := 2^{\lfloor \log_2 w(v) \rfloor}$ ; // round down to next power of 2
4:    $\hat{w}(v) := \max_{u \in N_2(v)} \tilde{w}(u)$ ;
5:   if  $\tilde{w}(v) = \hat{w}(v)$  then  $v.active := \mathbf{true}$  else  $v.active := \mathbf{false}$  end if;
6:   compute support  $s(v) := |\{u \in N(v) : u.active = \mathbf{true}\}|$ ;
7:    $\hat{s}(v) := \max_{u \in W(v)} s(u)$ ;
8:    $v.candidate := \mathbf{false}$ ;
9:   if  $v.active$  then
10:     $v.candidate := \mathbf{true}$  with probability  $1/\hat{s}(v)$ 
11:   end if;
12:   compute  $c(v) := |\{u \in W(v) : u.candidate = \mathbf{true}\}|$ ;
13:   if  $v.candidate$  and  $\sum_{u \in W(v)} c(u) \leq 3w(v)$  then
14:     node  $v$  joins dominating set
15:   end if
16:    $W(v) := \{u \in N(v) : u \text{ is white}\}$ ;  $w(v) := |W(v)|$ ;
17: end while

```

$(6 \cdot \ln \Delta + 12) \cdot |S^*|$, where S^* is an optimal dominating set.

Proof. The proof is a bit more involved but analogous to the analysis of the approximation ratio of the greedy algorithm. Every time, we add a new node v to the dominating set, we distribute the cost among v (if it is still white) and its white neighbors. Consider an optimal dominating set S^* . As in the analysis of the greedy algorithm, we partition the graph into stars by assigning every node u not in S^* to a neighbor v^* in S^* . We want to show that the total distributed cost in the star of every $v^* \in S^*$ is at most $6H(\Delta + 1)$.

Consider a node v that is added to the dominating set by Algorithm 26.8. Let $W(v)$ be the set of white nodes in $N(v)$ when v becomes a dominator. For a node $u \in W(v)$ let $c(u)$ be the number of candidate nodes in $N(u)$. We define $C(v) = \sum_{u \in W(v)} c(u)$. Observe that $C(v) \leq 3w(v)$ because otherwise v would not join the dominating set in line 15. When adding v to the dominating set, every newly covered node $u \in W(v)$ is charged $3/(c(u)w(v))$. This compensates the cost 1 for adding v to the dominating set because

$$\sum_{u \in W(v)} \frac{3}{c(u)w(v)} \geq w(v) \cdot \frac{3}{w(v) \cdot \sum_{u \in W(v)} c(u)/w(v)} = \frac{3}{C(v)/w(v)} \geq 1.$$

The first inequality follows because it can be shown that for $\alpha_i > 0$, $\sum_{i=1}^k 1/\alpha_i \geq k/\bar{\alpha}$ where $\bar{\alpha} = \sum_{i=1}^k \alpha_i/k$.

Now consider a node $v^* \in S^*$ and assume that a white node $u \in W(v^*)$ turns gray or black in iteration t of the while loop. We have seen that u is charged $3/(c(u)w(v))$ for every node $v \in N(u)$ that joins the dominating set in iteration t . Since a node can only join the dominating set if its span is largest up to a factor of two within two hops, we have $w(v) \geq w(v^*)/2$ for every node $v \in N(u)$ that joins the dominating set in iteration t . Because there are at most $c(u)$ such nodes, the charge of u is at most $6/w(v^*)$. Analogously to the sequential greedy

algorithm, we now get that the total cost in the star of a node $v^* \in S^*$ is at most

$$\sum_{i=1}^{|N(v^*)|} \frac{6}{i} \leq 6 \cdot H(|N(v^*)|) \leq 6 \cdot H(\Delta + 1) = 6 \cdot \ln \Delta + 12.$$

□

To bound the time complexity of the algorithm, we first need to prove the following lemma.

Lemma 26.10. *Consider an iteration of the while loop. Assume that a node u is white and that $2s(u) \geq \max_{v \in \mathcal{C}(u)} \hat{s}(v)$ where $\mathcal{C}(u) = \{v \in N(u) : v.\text{candidate} = \mathbf{true}\}$. Then, the probability that u becomes dominated (turns gray or black) in the considered while loop iteration is larger than $1/9$.*

Proof. Let $D(u)$ be the event that u becomes dominated in the considered while loop iteration, i.e., $D(u)$ is the event that u changes its color from white to gray or black. Thus, we need to prove that $\Pr[D(u)] > 1/9$. We can write this probability as

$$\Pr[D(u)] = \Pr[c(u) > 0] \cdot \Pr[D(u)|c(u) > 0] + \underbrace{\Pr[c(u) = 0] \cdot \Pr[D(u)|c(u) = 0]}_{=0}.$$

It is therefore sufficient to lower bound the probabilities $\Pr[c(u) > 0]$ and $\Pr[D(u)|c(u) > 0]$. We have $2s(u) \geq \max_{v \in \mathcal{C}(u)} \hat{s}(v)$. Therefore, in line 10, each of the $s(u)$ active nodes $v \in N(u)$ becomes a candidate node with probability $1/\hat{s}(v) \geq 1/(2s(u))$. The probability that at least one of the $s(u)$ active nodes in $N(u)$ becomes a candidate therefore is

$$\Pr[c(u) > 0] > 1 - \left(1 - \frac{1}{2s(u)}\right)^{s(u)} > 1 - \frac{1}{\sqrt{e}} > \frac{1}{3}.$$

We used that for $x \geq 1$, $(1-1/x)^x < 1/e$. We next want to bound the probability $\Pr[D(u)|c(u) > 0]$ that at least one of the $c(u)$ candidates in $N(u)$ joins the dominating set. We have

$$\Pr[D(u)|c(u) > 0] \geq \min_{v \in N(u)} \Pr[v \text{ joins dominating set} | v.\text{candidate} = \mathbf{true}].$$

Consider some node v and let $C(v) = \sum_{v' \in W(v)} c(v')$. If v is a candidate, it joins the dominating set if $C(v) \leq 3w(v)$. We are thus interested in the probability $\Pr[C(v) \leq 3w(v) | v.\text{candidate} = \mathbf{true}]$. Assume that v is a candidate. Let $c'(v') = c(v') - 1$ be the number of candidates in $N(v') \setminus \{v\}$. For a node $v' \in W(v)$, $c'(v')$ is upper bounded by a binomial random variable $\text{Bin}(s(v') - 1, 1/s(v'))$ with expectation $(s(v') - 1)/s(v')$. We therefore have

$$\mathbb{E}[c(v') | v.\text{candidate} = \mathbf{true}] = 1 + \mathbb{E}[c'(v')] = 1 + \frac{s(v') - 1}{s(v')} < 2.$$

By linearity of expectation, we hence obtain

$$\begin{aligned} \mathbb{E}[C(v) | v.\text{candidate} = \mathbf{true}] &= \sum_{v' \in W(v)} \mathbb{E}[c(v') | v.\text{candidate} = \mathbf{true}] \\ &< 2w(v). \end{aligned}$$

We can now use Markov's inequality to bound the probability that $C(v)$ becomes too large:

$$\Pr[C(v) > 3w(v) | v.\text{candidate} = \mathbf{true}] < \frac{2}{3}.$$

Combining everything, we get

$$\begin{aligned} & \Pr[v \text{ joins dom. set} | v.\text{candidate} = \mathbf{true}] \\ &= \Pr[C(v) \leq 3w(v) | v.\text{candidate} = \mathbf{true}] > \frac{1}{3} \end{aligned}$$

and hence

$$\Pr[D(u)] = \Pr[c(u) > 0] \cdot \Pr[D(u) | c(u) > 0] > \frac{1}{3} \cdot \frac{1}{3} = \frac{1}{9}.$$

□

Theorem 26.11. *In expectation, Algorithm 26.8 terminates in $\mathcal{O}(\log^2 \Delta \cdot \log n)$ rounds.*

Proof. First observe that every iteration of the while loop can be executed in a constant number of rounds. Consider the state after t iterations of the while loop. Let $\tilde{w}_{\max}(t) = \max_{v \in V} \tilde{w}(v)$ be the maximal span rounded down to the next power of 2 after t iterations. Further, let $s_{\max}(t)$ be the maximal support $s(v)$ of any node v for which there is a node $u \in N(v)$ with $w(u) \geq \tilde{w}_{\max}(t)$ after t while loop iterations. Observe that all nodes v with $w(v) \geq \tilde{w}_{\max}(t)$ are active in iteration $t + 1$ and that as long as the maximal rounded span $\tilde{w}_{\max}(t)$ does not change, $s_{\max}(t)$ can only get smaller with increasing t . Consider the pair $(\tilde{w}_{\max}, s_{\max})$ and define a relation \prec such that $(w', s') \prec (w, s)$ iff $w' < w$ or $w = w'$ and $s' \leq s/2$. From the above observations, it follows that

$$(\tilde{w}_{\max}(t), s_{\max}(t)) \prec (\tilde{w}_{\max}(t'), s_{\max}(t')) \implies t > t'. \quad (26.11.1)$$

For a given time t , let $T(t)$ be the first time for which

$$(\tilde{w}_{\max}(T(t)), s_{\max}(T(t))) \prec (\tilde{w}_{\max}(t), s_{\max}(t)).$$

We first want to show that for all t ,

$$\mathbb{E}[T(t) - t] = O(\log n). \quad (26.11.2)$$

Let us look at the state after t while loop iterations. By Lemma 26.10, every white node u with support $s(u) \geq s_{\max}(t)/2$ will be dominated after the following while loop iteration with probability larger than $1/9$. Consider a node u that satisfies the following three conditions:

- (1) u is white
- (2) there is a node $v \in N(u) : w(v) \geq \tilde{w}_{\max}(t)$
- (3) $s(u) \geq s_{\max}(t)/2$.

As long as u satisfies all three conditions, the probability that u becomes dominated is larger than $1/9$ in every while loop iteration. Hence, after $t+\tau$ iterations (from the beginning), u is dominated or does not satisfy (2) or (3) with probability larger than $(8/9)^\tau$. Choosing $\tau = \log_{9/8}(2n)$, this probability becomes $1/(2n)$. There are at most n nodes u satisfying Conditions (1) – (3). Therefore, applying union bound, we obtain that with probability more than $1/2$, there is no white node u satisfying Conditions (1) – (3) at time $t + \log_{9/8}(2n)$. Equivalently, with probability more than $1/2$, $T(t) \leq t + \log_{9/8}(2n)$. Analogously, we obtain that with probability more than $1/2^k$, $T(t) \leq t + k \log_{9/8}(2n)$. We then have

$$\begin{aligned} \mathbb{E}[T(t) - t] &= \sum_{\tau=1}^{\infty} \Pr[T(t) - t = \tau] \cdot \tau \\ &\leq \sum_{k=1}^{\infty} \left(\frac{1}{2^k} - \frac{1}{2^{k+1}} \right) \cdot k \log_{9/8}(2n) = \log_{9/8}(2n) \end{aligned}$$

and thus Equation (26.11.2) holds.

Let $t_0 = 0$ and $t_i = T(t_{i-1})$ for $i = 1, \dots, k$, where $t_k = \min_t \tilde{w}_{\max}(t) = 0$. Because $\tilde{w}_{\max}(t) = 0$ implies that $w(v) = 0$ for all $v \in V$ and that we therefore have computed a dominating set, by Equations (26.11.1) and (26.11.2) (and linearity of expectation), the expected number of rounds until Algorithm 26.8 terminates is $\mathcal{O}(k \cdot \log n)$. Since $\tilde{w}_{\max}(t)$ can only have $\lfloor \log \Delta \rfloor$ different values and because for a fixed value of $\tilde{w}_{\max}(t)$, the number of times $s_{\max}(t)$ can be decreased by a factor of 2 is at most $\log \Delta$ times, we have $k \leq \log^2 \Delta$. \square

Remarks:

- It is not hard to show that Algorithm 26.8 even terminates in $\mathcal{O}(\log^2 \Delta \cdot \log n)$ rounds with probability $1 - 1/n^c$ for an arbitrary constant c .
- Using the median of the supports of the neighbors instead of the maximum in line 8 results in an algorithm with time complexity $\mathcal{O}(\log \Delta \cdot \log n)$. With another algorithm, this can even be slightly improved to $\mathcal{O}(\log^2 \Delta)$.
- One can show that $\Omega(\log \Delta / \log \log \Delta)$ rounds are necessary to obtain an $\mathcal{O}(\log \Delta)$ -approximation.
- It is not known whether there is a fast deterministic approximation algorithm. This is an interesting and important open problem. The best deterministic algorithm known to achieve an $\mathcal{O}(\log \Delta)$ -approximation has time complexity $2^{\mathcal{O}(\sqrt{\log n})}$.

Chapter Notes

See [JRS02, KW05].

Bibliography

- [JRS02] Lujun Jia, Rajmohan Rajaraman, and Torsten Suel. An efficient distributed algorithm for constructing small dominating sets. *Distributed Computing*, 15(4):193–205, 2002.
- [KW05] Fabian Kuhn and Roger Wattenhofer. Constant-Time Distributed Dominating Set Approximation. In *Springer Journal for Distributed Computing, Volume 17, Number 4*, May 2005.