

Lecture 19

November 15, 2023

Instructor: Sepehr Assadi

Disclaimer: *These notes have not been subjected to the usual scrutiny reserved for formal publications. They may be distributed outside this class only with the permission of the Instructor.*

Topics of this Lecture

1	Multiplicative Weight Update for the Matching Linear Program	1
1.1	The Setup and the Oracle LP	2
1.2	The Algorithm	3
1.3	The Analysis	4
1.4	Runtime Analysis	6
2	A Faster MWU Algorithm for Bipartite Matching	6
2.1	A Faster Algorithm via a Better Oracle	7
2.2	An Even (Slightly) Better Algorithm for Additive Approximation	8

1 Multiplicative Weight Update for the Matching Linear Program

In the last lecture, we studied several MWU-type algorithms and saw how to analyze them. We will now use the same set of ideas to design algorithms for solving (some) LPs *approximately*. The approach we use in this section (at least in its first part) is quite general and can be applied to many other LP (and “LP-type”) problems. However, to keep things concrete and for providing more intuition, we will focus on the *maximum matching* LP that we have worked on a lot with already in this course.

The LP Relaxation of Matching. For any graph $G = (V, E)$, we can write the following LP as a relaxation for the maximum matching problem in G :

$$\begin{aligned}
 & \max_{x \in \mathbb{R}^E} && \sum_{e \in E} x_e \\
 \text{subject to} &&& \sum_{e \ni v} x_e \leq 1 \quad \forall v \in V \\
 &&& x_e \geq 0 \quad \forall e \in E.
 \end{aligned} \tag{1}$$

Recall that in Lecture 8, we showed that this LP relaxation for *bipartite* graphs have integrality gap, i.e., we can round any fractional solution to this LP to an integral one without decreasing its size. The LP is well-defined for all graphs (not only bipartite ones), although for general graphs, its integrality gap is $2/3$ (think of a triangle). For now, we only focus on solving this LP near-optimally and ignore any rounding aspects; hence, we will just work with general graphs.

1.1 The Setup and the Oracle LP

The idea behind the MWU algorithm for solving any LP is to repeatedly reduce the problem to solving a much simpler LP, often called the *oracle LP*. The oracle LP is obtained from the original LP by partitioning the constraints of the LP into “easy” and “hard” ones, and then focusing only on solving a *convex combination* of the “hard” constraints instead of all of them. This may sound too abstract and vague, so let us try to clarify this further.

Suppose we have some weights $w_v > 0$ on each vertex $v \in V$. Define $W := \sum_{v \in V} w_v$. Consider the following simpler LP called the **oracle LP** with weights $\{w_v\}_{v \in V}$:

$$\begin{aligned} & \max_{x \in \mathbb{R}^E} \quad \sum_{e \in E} x_e \\ & \text{subject to} \quad \sum_{v \in V} w_v \sum_{e \ni v} x_e \leq W \\ & \quad \quad \quad x_e \geq 0 \quad \forall e \in E. \end{aligned} \tag{2}$$

Notice that this LP is simpler than the original LP in that instead of the n “hard” constraints of the original LP—one for each vertex of the graph—we now only have a single constraint for their convex combination.

Observation 1. *For any choice of weights $w_v \geq 0$ for $v \in V$, the optimal objective value of LP (2) is at least as large as that of LP (1).*

Proof. Any feasible solution x to LP (1) satisfies the following equation for every $v \in V$:

$$\sum_{e \ni v} x_e \leq 1$$

Thus if we multiply each side by $w_v \geq 0$, we will still get

$$w_v \cdot \sum_{e \ni v} x_e \leq w_v;$$

and, if we further sum up both sides on all vertices $v \in V$, we get

$$\sum_v w_v \cdot \sum_{e \ni v} x_e \leq \sum_v w_v = W.$$

Thus, any feasible solution to LP (1) is also a feasible solution to LP (2). As both LPs are maximization LPs with the same objective, we obtain the result. \square

Moreover, solving this LP is also much simpler than the original one. For every edge $e = (u, v) \in E$, define the *artificial weight* of the edge e to be

$$aw(e) := w_u + w_v.$$

We have,

$$\sum_{v \in V} w_v \cdot \sum_{e \ni v} x_e := \sum_{e=(u,v) \in E} x_e \cdot (w_u + w_v) = \sum_{e \in E} aw(e) \cdot x_e.$$

Hence, the single main constraint of LP (2) is equivalent to:

$$\sum_{e \in E} aw(e) \cdot x_e \leq W. \tag{3}$$

But then it means that to obtain the optimal solution to LP (2), we simply need to find

$$e^* := \arg \min_{e \in E} aw(e)$$

and then set

$$x_{e^*} := \frac{W}{aw(e^*)} \quad \text{and} \quad x_e = 0 \quad \text{for all other } e \neq e^*. \quad (4)$$

This clearly maximizes $\sum_{e \in E} x_e$ (because “moving” any value from x_{e^*} on any other x_e can only make the constraint violated without helping with the objective value). We summarize this as follows.

Observation 2. *The optimal solution to LP (2) is given by Eq (4).*

1.2 The Algorithm

So now how can we use this oracle LP for solving the original LP? Notice that if we put a *large* weight on a vertex, we will “move” the optimum solution of the oracle LP away from that vertex; thus, by adjusting the weights of the constraints, we can make them more important or reduce their importance, hence getting the oracle LP to “focus” on different parts of the graph. Then, by taking the *average* of all the solutions we found, we will hopefully get a solution which is handling every constraint of LP (1) (approximately). We show how we can update the weights using MWU to achieve this goal.

The algorithm is formally as follows (the algorithm has two parameters $\eta > 0$ and $T \geq 1$ that will be determined later).

Algorithm 1. A MWU algorithm for the Matching LP.

1. For every vertex $v \in V$, let $w_v^{(1)} = 1$.
2. For $t = 1$ to T iterations:
 - (a) Let $x^{(t)}$ be the optimal solution to the oracle LP with weights $w^{(t)}$ according to Eq (4).
 - (b) For every vertex $v \in V$, update:

$$w_v^{(t+1)} = \left(1 + \eta \cdot \sum_{e \ni v} x_e^{(t)} \right) \cdot w_v^{(t)}.$$

3. Return the final solution

$$\bar{x} := \frac{1}{T} \cdot \sum_{t=1}^T x^{(t)}.$$

A note that the update rule basically means in each iteration the two vertices u, v incident on the optimal edge of Eq (4) are having their weight increased (somewhat proportional to the “violation” of their matching constraint in the original LP) and all other vertices retain their original weight. However, it will be easier to work with the given formulation both for the current analysis, as well as future improvements.

Let $\varepsilon > 0$ be a parameter and recall that our goal is to obtain a $(1 + \varepsilon)$ -approximation to LP (1) via Algorithm 1. The choice of ε will play a role in determining the values of η and T . The following is the main lemma for this algorithm.

Lemma 3. *For a proper choice of η and T , the output \bar{x} of Algorithm 1 satisfies the following:*

$$\sum_{e \in E} \bar{x}_e \geq \text{OPT} \quad (5)$$

$$\sum_{e \in v} \bar{x}_e \leq 1 + \varepsilon \quad \text{for all } v \in V \quad (6)$$

where OPT is the optimum value of LP (1).

Notice that the statement of this lemma does *not* imply \bar{x} is feasible for LP (1), even though its value is as good as the optimal one. However, we can simply rescale \bar{x} with $(1 + \varepsilon)$, i.e., consider $\bar{x}/(1 + \varepsilon)$ as our solution which is now feasible and a $(1 + \varepsilon)$ -approximation to LP (1)

1.3 The Analysis

It thus remains to prove [Lemma 3](#) to finalize the whole proof.

Proof of Eq (5) for all choices of η, T . We start with the easy part first which is to prove [Eq \(5\)](#). By [Observation 1](#), we have that for every $t \geq 1$, the optimum solution $x^{(t)}$ satisfies

$$\sum_{e \in E} x_e^{(t)} \geq \text{OPT}.$$

Thus, the average solution \bar{x} also satisfies this equation, hence proving [Eq \(5\)](#) no matter what is the choice of η and T .

Proof of Eq (6) for proper choices of η, T . We now prove the main part. We follow a similar *potential function* argument as the one for the expert problem in the previous lecture. Our potential function is again $W^{(t)} := \sum_{v \in V} w_v^{(t)}$, namely, the total weights of the vertices. The first part is to show that the potential does not increase too much.

Claim 4. *For every choice of $T \geq 1$, we have that at the end of the last iteration*

$$W^{(T+1)} \leq \exp(\eta \cdot T + \ln n).$$

Proof. For every $t \geq 1$,

$$\begin{aligned} W^{(t+1)} &= \sum_{v \in V} w_v^{(t+1)} && \text{(by the definition of } W^{(t+1)}) \\ &= \sum_{v \in V} 1 + \eta \cdot \sum_{e \ni v} x_e^{(t)} \cdot w_v^{(t)} && \text{(by the update rule of the algorithm)} \\ &= \sum_{v \in V} w_v^{(t)} + \eta \cdot \left(\sum_{v \in V} w_v^{(t)} \cdot \sum_{e \ni v} x_e^{(t)} \right) && \text{(by rearranging the terms)} \\ &= W^{(t)} + \eta \cdot \left(\sum_{v \in V} w_v^{(t)} \cdot \sum_{e \ni v} x_e^{(t)} \right) && \text{(by the definition of } W^{(t)}) \\ &\leq W^{(t)} + \eta \cdot W^{(t)} && \text{(by the feasibility of } x^{(t)} \text{ in the oracle LP (2))} \\ &= (1 + \eta) \cdot W^{(t)}. \end{aligned}$$

As a result, and since $W^{(1)} = n$, we obtain that after the T -th iteration:

$$W^{(T+1)} \leq (1 + \eta)^T \cdot n \leq \exp(\eta \cdot T) \cdot n = \exp(\eta \cdot T + \ln n),$$

where we used the inequality $1 + x \leq e^x$ for all $x > 0$. □

The second part is to show that if there is a vertex v that even at the end of the T iterations does not satisfy [Eq \(6\)](#), its weight should have increased by a lot. Specifically, we have a vertex $v \in V$ such that after T iterations, it still does not satisfy [Eq \(6\)](#), i.e.,

$$\sum_{e \ni v} \bar{x}_e > (1 + \varepsilon).$$

This equivalently means that

$$\sum_{t=1}^T \sum_{e \ni v} x_e^{(t)} > (1 + \varepsilon) \cdot T. \tag{7}$$

Claim 5. For every choice of $T \geq 1$ if a vertex v satisfies Eq (7) at the end of the last iteration, then,

$$w_v^{(T+1)} \geq \exp\left(\eta \cdot \left(1 + \frac{\varepsilon}{4}\right) \cdot T\right),$$

as long as $\varepsilon < 1/2$ and $\eta \leq \varepsilon/2\rho_v$ for ρ_v defined as:

$$\rho_v := \max_{t \geq 1} \sum_{e \ni v} x_e^{(t)}.$$

Proof. We have,

$$\begin{aligned} w_v^{(T+1)} &= \prod_{t=1}^T \left(1 + \eta \cdot \sum_{e \ni v} x_e^{(t)}\right) && \text{(by the update rule and since } w_v^{(1)} = 1) \\ &\geq \prod_{t=1}^T \exp\left(\eta \cdot \sum_{e \ni v} x_e^{(t)} - \eta^2 \cdot \left(\sum_{e \ni v} x_e^{(t)}\right)^2\right), \end{aligned}$$

as long as $\eta < 1/\rho_v$ because for such choice of η , we can apply the inequality $1 + y \geq e^{y-y^2}$ which holds for $y \in (0, 1)$. Continuing the above equations, we get,

$$\begin{aligned} w_v^{(T+1)} &\geq \exp\left(\eta \cdot \sum_{t=1}^T \sum_{e \ni v} x_e^{(t)} - \eta^2 \cdot \sum_{t=1}^T \left(\sum_{e \ni v} x_e^{(t)}\right)^2\right) && \text{(by using } e^\alpha \cdot e^\beta = e^{\alpha+\beta}) \\ &\geq \exp\left(\eta \cdot \sum_{t=1}^T \sum_{e \ni v} x_e^{(t)} - \eta^2 \cdot \rho_v \cdot \sum_{t=1}^T \sum_{e \ni v} x_e^{(t)}\right) && \text{(by replacing } \sum_{e \ni v} x_e^{(t)} \leq \rho_v) \\ &\geq \exp\left(\eta \cdot \sum_{t=1}^T \sum_{e \ni v} x_e^{(t)} - \frac{\varepsilon}{2} \cdot \eta \cdot \sum_{t=1}^T \sum_{e \ni v} x_e^{(t)}\right) && \text{(by the assumption on the value of } \eta) \\ &= \exp\left(\eta \left(1 - \frac{\varepsilon}{2}\right) \cdot \sum_{t=1}^T \sum_{e \ni v} x_e^{(t)}\right) \\ &\geq \exp\left(\eta \cdot \left(1 - \frac{\varepsilon}{2}\right) \cdot (1 + \varepsilon) \cdot T\right) && \text{(by Eq (7))} \\ &\geq \exp\left(\eta \cdot \left(1 + \frac{\varepsilon}{4}\right) \cdot T\right), && \text{(as } \varepsilon < 1/2) \end{aligned}$$

concluding the proof. \square

Since all the weights are positive, for any vertex $v \in V$, as long as we pick $\eta < \varepsilon/2\rho_v$ as prescribed by Claim 5, we obtain that for v to satisfy Eq (7) we need to have,

$$\exp\left(\eta \cdot \left(1 + \frac{\varepsilon}{4}\right) \cdot T\right) \leq w_v^{(T+1)} \leq W^{(T+1)} \leq \exp(\eta \cdot T + \ln n),$$

where the LHS inequality is by Claim 5 and the RHS inequality is by Claim 4. This implies that

$$\eta \cdot \left(1 + \frac{\varepsilon}{4}\right) \cdot T \leq \eta \cdot T + \ln n,$$

which in turn means

$$T \leq \frac{4 \ln n}{\eta \cdot \varepsilon}.$$

Thus, after these many iterations, all vertices violated Eq (7) and thus in turn satisfy Eq (6).

We state all these—for future reference—in the following lemma (whose proof already appeared above).

Lemma 6. For any $\varepsilon \in (0, 1/2)$, if we run *Algorithm 1* for $T > \frac{8\rho \cdot \ln n}{\varepsilon^2}$ iterations for

$$\rho \geq \max_{v \in V} \max_{t \geq 1} \sum_{e \ni v} x_e^{(t)},$$

with parameter $\eta = \frac{\varepsilon}{2\rho}$, then the output solution \bar{x} satisfies *Eq (6)*.

Finally, we emphasize that in the proof of *Lemma 8*, we never once used the fact that each $x^{(t)}$ returned by the algorithm is *optimal* for the oracle LP (2), only that it is *feasible* (this is the exact opposite of the case when proving *Eq (5)*).

1.4 Runtime Analysis

Unfortunately, we are still not done entirely, because it is not clear how to pick the parameter ρ needed in *Lemma 8*. Notice that ρ is basically the *largest* value we will ever assign to any edge e in any $x_e^{(t)}$ in any of the iterations (technically, it is the largest value of $\sum_{e \ni v} x_e^{(t)}$ but since the algorithm only assigns value to a single edge in each iteration, these two are the same). But how can we bound this quantity?

The key observation is that we do *not* need to solve each oracle LP (2) optimally! We only need to solve it as good the value of the *original* LP (1); this is because, even in that case, we can still prove *Eq (5)* exactly as it is (and lower bound the value of each $x^{(t)}$ by OPT still, which is what we did anyway), and *Eq (6)* does not have anything to do with the optimality of $x^{(t)}$, only its feasibility.

But now, we do know that $\text{OPT} \leq n$ always because it is optimal solution to a matching LP. This implies that in Line (2a) of *Algorithm 1*, we can in fact replace $x^{(t)}$ computed in *Eq (4)* by the following ($e^* := \arg \min_{e \in E} aw(e)$ is as before):

$$x_{e^*} := \min \left(n, \frac{W}{aw(e^*)} \right) \quad \text{and} \quad x_e = 0 \quad \text{for all other } e \neq e^*,$$

i.e., we are now taking the minimum of the previously chosen value and $n \geq \text{OPT}$. This ensures that every $x^{(t)}$ still have value at least OPT while also continue to be feasible. Thus, we obtain both *Eq (5)* and *Eq (6)* as before with no other changes to the proof.

The benefit of this is that we now can say $\rho \leq n$ because we will never assign a larger value in any iteration of the algorithm. This allows us to pick a choice of $\eta \leq \varepsilon/2n$ as well which gives us an actual algorithm (with known parameters throughout).

Finally, the number of iterations of this algorithm is

$$O\left(\frac{n \ln n}{\varepsilon^2}\right)$$

by *Lemma 8*. Each iteration also requires us finding the edge with *minimum* artificial weight $aw(e)$ (based on the weight function $w^{(t)}$). It is easy to see that each update only changes the weight of $O(n)$ other edges and thus we can maintain the edges in a min-heap and run each iteration of the algorithm in $O(n \log n)$ time or a Fibonacci heap and $O(n)$ time (in fact, here simpler solutions with $O(n)$ time also exist). Thus, the total runtime of the algorithm is

$$O\left(\frac{n^2}{\varepsilon^2} \cdot \log n\right)$$

time.

2 A Faster MWU Algorithm for Bipartite Matching

Let us see how we can use the MWU method differently to reduce the number of iterations dramatically. As stated in *Lemma 8* and earlier in the proof, “*all*” we need to be able to use MWU *efficiently* is the following:

- We should be able to solve the oracle LP (2) as well as the *original* LP (1) and not at all necessarily optimally – this ensures our output will also be (nearly) optimal;
- We would like to reduce the width parameter ρ (defined in Lemma 8) to be as small as possible – this ensures the number of iterations will be as small as possible.

While the strategy we outlined previous to this more or less works for any LP, we can often use the underlying properties of our specific LP to obtain even faster algorithms. To show case this, we will consider the **bipartite matching** problem now and show how we can make the above algorithm even faster.

2.1 A Faster Algorithm via a Better Oracle

As we proved earlier in the course (in Lecture 8), the integrality gap of LP (1) on bipartite graphs is one. Thus, we can assume that there is always a matching M_{OPT} of size OPT in the in the input graph G . On the other hand, recall the Oracle LP (2) problem, when stated with the artificial weights of edges:

$$\begin{aligned} & \max_{x \in \mathbb{R}^E} \sum_{e \in E} x_e \\ & \text{subject to} \quad \sum_{e \in E} x_e \cdot aw(e) \leq W \\ & \quad \quad \quad x_e \geq 0 \quad \forall e \in E. \end{aligned} \tag{8}$$

Thus, we are saying that there is a matching M_{OPT} in the graph G with the total artificial weight at most W (no matter what weights we are choosing for the vertices and thus, in turn, the artificial weight of edges).

Now, consider our previous approach of picking the edge e^* with minimum weight and setting $x_{e^*} = \text{OPT}$ and $x_e = 0$ for all $e \neq e^*$ (we actually put a different value because we do not know OPT , but you can see that *had* we known OPT , using this choice would have worked as well). We can think of this strategy as follows: while M_{OPT} has OPT edges with weight at most W , we can find a single edge with weight at most W/OPT – thus, on this single edge, we are *competing* quite favorably with the optimum in terms of “bang for the buck”: how much we contribute to objective value versus the main constraint of the Oracle LP. The problem with this approach was that we needed to put a very large value on x_{e^*} and thus get a large width ρ , which in turn led to a large number of iterations in the algorithm.

But, now suppose could pick k edges with weight W/k still for some $k > 1$. Can we again compete favorably with the optimum solution? *Yes*; do we get a lower width? *Not necessarily* because the width is determined by the largest x -value we put on a single *vertex* and not an edge. Thus, just picking k edges is not enough. But what if these k edges form a matching? Then, indeed we can reduce the width to OPT/k also by putting an x -value of OPT/k over each of these edges. This is precisely what our better oracle does.

Algorithm 2. A Better Oracle for the Bipartite Matching LP.

1. Sort the edges in the *increasing* order of their artificial weights (determined by the input weights on the vertices).
2. Let $M = \emptyset$. While the artificial weight of M is less than or equal to $W/2$:
 - Pick the lowest weight available edge e in M if both its endpoints are unmatched, and otherwise skip this edge.
3. Return $x = 2 \cdot \mathbb{1}_M$ where $\mathbb{1}_M$ denotes the characteristic vector of the matching $M \subseteq E$.

Lemma 7. *Size of M in Algorithm 2 is at least $\text{OPT}/2$ for all choices of the artificial weights on the edges.*

Proof. Let $e_1, e_2, \dots, e_{\text{OPT}/2}$ be the first $\text{OPT}/2$ edges of M in the order added to M and $o_1, o_2, \dots, o_{\text{OPT}}$ be the edges of M_{OPT} sorted in the increasing order of their artificial weight. We claim inductively that for

every $i \in [\text{OPT}/2]$,

$$\underbrace{\sum_{j=1}^i aw(e_j)}_{aw(e_{\leq i})} \leq \frac{1}{2} \cdot \underbrace{\sum_{j=1}^{2i} aw(o_j)}_{aw(o_{\leq 2i})};$$

in words, the total weight of the first i edges in M , denoted by $aw(e_{\leq i})$ is at most half the total weight of the first $2i$ edges in M_{OPT} , denoted by $aw(o_{\leq 2i})$.

The base case for $i = 1$ holds because $aw(e_1)$ is smaller than all other edges in the graph and thus in particular is half of the $aw(o_1) + aw(o_2)$.

For induction case, suppose we are adding the edge e_i in this iteration. Since e_1, \dots, e_{i-1} is a matching, each of these edges can be incident on at most two edges in M_{OPT} . Thus, by the pigeonhole principle, among the edges $o_1, o_2, \dots, o_{2i-1}, o_{2i}$, there are at least two edges that are not incident on any of these edges. Hence, when picking e_i , we could have alternatively picked any of these two edges, which means that weight of e_i is at most equal to the weight of each of these two edges. In particular, we definitely have,

$$aw(e_i) \leq \frac{1}{2} \cdot (aw(o_{2i-1}) + aw(o_{2i})).$$

Moreover, by induction, we have that

$$aw(e_{\leq i-1}) \leq \frac{1}{2} \cdot aw(o_{\leq 2i-2}).$$

Combining these two implies the induction hypothesis.

The lemma now follows from this because we have the weight of the first $\text{OPT}/2$ edges of M is at most $W/2$ and thus the algorithm picks at least $\text{OPT}/2$ edges before terminating. \square

As a result, if we use [Algorithm 2](#) in Line (2a) of [Algorithm 1](#), we will still obtain a feasible solution to the Oracle LP (2) with value as large as OPT . Plugging this in [Lemma 8](#), we also obtain that

$$\rho = \max_{v \in V} \max_{t \geq 1} \sum_{e \ni v} x_e^{(t)} = 2,$$

because in each iteration, each vertex is incident on at most one edge of M (because it is a matching) and the value we put on edges of M is only 2. Thus, the number of iterations of our algorithm is now only

$$O\left(\frac{\ln n}{\varepsilon^2}\right),$$

namely, a factor of n smaller than what it was before. Each iteration also takes $O(m \log m)$ time if we sort the edges directly using any sorting algorithm. We can in fact implement each iteration in $O(m)$ time also because the artificial weight of an edge is simply determined by the number of times each of its endpoints is matched so far, which is an integer; so, we can simply use Radix sort (or any other fast integer sorting algorithm) to sort their weight in $O(m)$ time. All in all, this gives us an algorithm with runtime

$$O\left(m \cdot \frac{\ln n}{\varepsilon^2}\right).$$

2.2 An Even (Slightly) Better Algorithm for Additive Approximation

Finally, to show case an important feature (or rather “weakness” in the way we analyzed [Algorithm 1](#)), let us show that if our goal is to settle for a slightly weaker approximation guarantee of outputting a matching of size $\text{OPT} - \varepsilon \cdot n$, instead of $(1 - \varepsilon) \cdot \text{OPT}$, we can reduce the number of iterations to be independent of n , i.e., remove the $\ln n$ -term.

So, why did we need the $\ln n$ -term in the analysis of [Algorithm 1](#) and in particular [Lemma 8](#)? This was essentially because we were comparing the weight of a single constraint v which were violated, against the entire potential function (a combination of n different weights). In other words, we should have run the algorithm long enough such that even one violated constraint could outweigh all n constraints. But, what if we set our goal to make sure there are $< \varepsilon \cdot n$ violated constraints, namely,

$$\sum_{e \ni v} \bar{x}_e > (1 + \varepsilon),$$

for $< \varepsilon \cdot n$ vertices $v \in V$. Why is this good enough for $\Theta(\varepsilon n)$ additive approximation? Notice that in [Algorithm 2](#), we never violated a constraint with more than 2 so we always have

$$\sum_{e \ni v} \bar{x}_e \leq 2,$$

(because $\rho = 2$). Thus, after we have $< \varepsilon \cdot n$ violated constraints, we can simply remove *all* values of x incident on their edges so they become feasible. But this means that we reduce the value of \bar{x} with at most $2\varepsilon n$ in total. Thus, after this modification, we obtain a solution which satisfies all constraints up to a $(1 + \varepsilon)$ factor and its value is at least $\text{OPT} - 2\varepsilon n$. This implies that we obtain a solution which is a $\text{OPT} - 3\varepsilon \cdot n$ after re-scaling \bar{x} to become $\bar{x}/(1 + \varepsilon)$. Finally, to obtain an additive εn approximation (instead of $3\varepsilon n$, we can simply start this algorithm with ε replaced by $\varepsilon/3$).

The final question is how many iterations we need to reduce the number of violated constraints to $< \varepsilon \cdot n$? This is handled by the following lemma.

Lemma 8. *For any $\varepsilon \in (0, 1/2)$, if we run [Algorithm 1](#) for $T > \frac{8\rho \cdot \ln(1/\varepsilon)}{\varepsilon^2}$ iterations for*

$$\rho \geq \max_{v \in V} \max_{t \geq 1} \sum_{e \ni v} x_e^{(t)},$$

with parameter $\eta = \frac{\varepsilon}{2\rho}$, then the output solution \bar{x} only violates [Eq \(6\)](#) for at most $\varepsilon \cdot n$ constraints.

Proof. Recall that by [Claim 5](#), if a vertex v violates [Eq \(6\)](#)—or equivalently satisfies [Eq \(7\)](#)—at the end of the last iteration, then,

$$w_v^{(T+1)} \geq \exp\left(\eta \cdot \left(1 + \frac{\varepsilon}{4}\right) \cdot T\right).$$

Let S be the set of all violating vertices and suppose towards a contradiction that $|S| \geq \varepsilon \cdot n$. We thus have,

$$\sum_{v \in S} w_v^{(T+1)} \geq \varepsilon \cdot n \cdot \exp\left(\eta \cdot \left(1 + \frac{\varepsilon}{4}\right) \cdot T\right) = \exp\left(\eta \cdot \left(1 + \frac{\varepsilon}{4}\right) \cdot T + \ln \varepsilon + \ln n\right).$$

On the other hand, by [Claim 4](#),

$$W^{(T+1)} \leq \exp(\eta \cdot T + \ln n).$$

Given that the LHS of the first equation is still upper bounded by the LHS of the next equation, we have

$$\eta \cdot \left(1 + \frac{\varepsilon}{4}\right) \cdot T + \ln \varepsilon + \ln n \leq \eta \cdot T + \ln n,$$

which implies that

$$T \leq \frac{4 \cdot \ln(1/\varepsilon)}{\eta \cdot \varepsilon}.$$

Replacing η with its value finalizes the proof. □

This implies that after running the algorithm for only $O(\ln(1/\varepsilon)/\varepsilon^2)$ iterations, we obtain the desired solution. In general, the ideas in this part can be very helpful for reducing the number of iterations to something independent of n .