

A distributed blossom algorithm for minimum-weight perfect matching

Eric C. Peterson
Eigenware
Berkeley, CA, USA
peterson.eric.c@gmail.com

Peter J. Karalekas
Eigenware
Berkeley, CA, USA
peter@karalekas.com

ABSTRACT

We describe a distributed, asynchronous variant of Edmonds’s exact algorithm for producing perfect matchings of minimum weight [Edm65a]. The development of this algorithm is driven by an application to online error correction in quantum computing, first envisioned by Fowler [FWH12]; we analyze the performance of our algorithm as applied to this domain in a sequel [PK].

1 INTRODUCTION

Edmonds’s now-classic results on maximum matchings lie at the intersection of computer science, combinatorics, and integer linear programming: starting from a known polynomial-time algorithm for producing maximum matchings in bipartite graphs [CCPS09, Proposition 5.7], he showed first that a polynomial-time modification could be used to handle a non-bipartite graph [Edm65b], then that in the presence of edge weights another polynomial-time modification could be used to produce a minimum-weight representative among maximum matchings [Edm65a]. These are respectively called the “blossom algorithm” and the “weighted blossom algorithm”. These landmark results set in motion broad research programs in several domains: there are theoretical consequences in both computer science and mathematics; the algorithmic technique itself admits both generalizations and efficiency improvements; and it opened the door to a host of applications.

As an example of such an application, minimum-weight perfect matchings (MWPMs) attracted the attention of quantum computer scientists, who showed that an MWPM solver can be used as an approximation algorithm for decoding syndromes appearing in quantum error correction, with approximation ratio dependent on the physical properties of the underlying quantum device [DKLP02]. This idea has taken such hold with designers of quantum computers that it has appeared in a variety of surveys on the subject (see, e.g., [FMMC12, CNAA’ 20]) as a solved problem. However, in order to deploy this on a live quantum device, Fowler showed that one must make use of a “parallelized” MWPM solver [FWH12], and work has stopped short of producing (or referencing) such an algorithm.

Careful consideration of the intended application indicates that the “parallelized” implementation must actually be *distributed* with only local information available to each worker, *online* so as to cope with a dynamic problem graph, and ideally *asynchronous* to best match lab hardware. Meanwhile, though state of the art in MWPM solvers has advanced substantially since the ’60s, they have had other concerns top of mind: it

is easy to show that the worst-case complexity of an exact solution to the matching problem on a cycle graph has runtime polynomial in the diameter [Lin92], which has encouraged the development of approximation algorithms instead ([WW04], [LPR09], [LPP15], and many others).

In this paper, driven by the extra structure available on the problem graphs in our intended application, we return to the exact setting: we describe an exact, asynchronous distributed blossom algorithm suitable for fulfilling Fowler’s claim and prove its correctness. As part of extending Edmonds’s algorithm to operate on alternating forests rather than trees, we draw the reader’s attention to a new, naturally-occurring forest operation which we call *multireweight*, which does not arise during serial execution and which is crucial to the correctness of the distributed algorithm. We also provide an implementation of the algorithm, *anatevka* [ana, Ale49], as part of a simulation testbed for distributed systems described in a previous paper [PK20, aet]. We make no reference to quantum computing outside of this introduction, since the existence and behavior of this algorithm is entirely a matter of distributed computing. Instead, we direct interested readers to the sequel paper [PK] for the further modifications necessary to the application and the performance analysis in that context.

2 THE SERIAL ALGORITHM

Our distributed algorithm is most easily cast as a piecewise modification of Edmonds’s serial blossom algorithm, with one extra operation. To facilitate such a description, and to put the unfamiliar reader at ease, we first review the details of the serial algorithm. The inputs and output of the problem are:

Definition 1. A *matching* M on a graph G is a set of edges with no repeated vertices. A matching is *maximum* when it is of maximum cardinality. A maximum matching on G is *perfect* when G has an even number of vertices. For edge-weighted G , a matching is said to be *minimum weight* if there is no equal-sized matching with smaller edge weight sum.

The goal is to produce minimum-weight perfect matchings.

Remark 2 (Standing assumptions on G). Initially, we will assume G to be unweighted and bipartite, though we will drop these assumptions as our discussion progresses. Between any pair of vertices in G , we permit there to be no edges, one edge, or several edges—but since a loop can never be a match edge, it is harmless to assume that G is loopless. For the purposes of our description, it is convenient to permit the case of multiple edges, but it is not necessary: the algorithm will behave as if there is at most one edge between any pair of vertices (viz., the one of least weight). In the weighted setting, one can also

All work was done prior to joining Amazon Web Services.

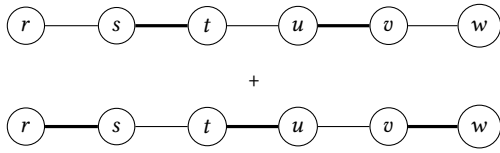


Figure 1: The graph on top illustrates an augmenting path joining r to w : neither r nor w is matched, and the edges between them alternate between not belonging and belonging to the matching. The graph below shows the effect of augmenting along this path: whether an edge is or is not a member of the matching reverses, and the size of the matching increases by one edge.

model missing edges by edges of infinite weight, so that the entire algorithm need only be described for a complete graph.

2.1 Augment and graft

Suppose that we are given a matching M , perhaps not yet maximum. The core mechanism of the algorithm is to identify an augmenting path:

Definition 3. We say that an edge is *matched* if it belongs to M or otherwise that it is *unmatched*, and we say that a vertex is *matched* if it is the endpoint of any matched edge or otherwise that it is unmatched. Thus, an *alternating chain* or *alternating path* is a sequence of adjoining edges which alternate between being matched and unmatched. An *augmenting path* is an alternating path whose first and last vertices are both unmatched.

With such a path in hand, one can produce a new matching by inverting which edges in the path belong to the matching. The number of edges in the new matching is one larger than that of the old.

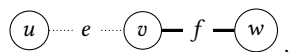
Definition 4. This inversion procedure is called *augmenting* M along the augmenting path.

Example 5. See Figure 1 for a depiction of augmentation.

The meat of the algorithm is then a search for augmenting paths, which begin and end at unmatched vertices. The data structure which powers this is an *alternating tree*:

Definition 6. An *alternating tree* is a tree whose root is unmatched, and whose edges alternate between unmatched and matched as they descend from the root. We refer to the even- and odd-depth tree vertices respectively as *positive* and *negative*, and we write T_+ and T_- for these subsets of vertices.¹

Definition 7. The inductive operation used to assemble such an alternating tree is called *grafting*.² Let M be an intermediate matching, and let T be an alternating subtree of the ambient graph G . Select a pair of edges e and f , as in



¹Some authors refer to positive, negative, and unmatched vertices respectively as *outer*, *inner*, and *exposed*.

²Some authors call this operation *grow* [Kol09].

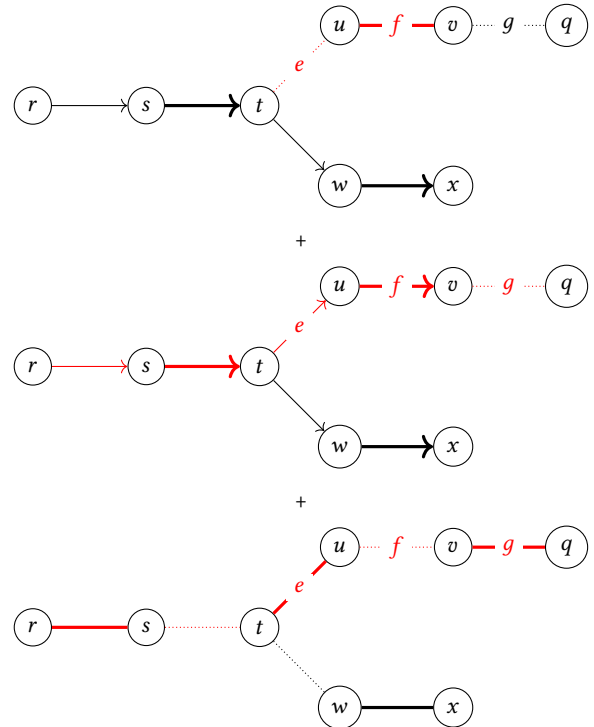


Figure 2: Begin by grafting a matched edge f along edge e onto an alternating tree T rooted at r . This creates an augmenting path (middle, red) formed from a branch of T and an edge g not in T . Then augment through this path to produce a maximum matching. Note that edges participating in the tree carry arrow heads (pointing toward the leaves), edges participating in the partial matching are bold, and the remaining edges are dotted.

with the additional properties that

- (1) u belongs to T , but v and w do not.
- (2) If u has a parent in T , the edge to that parent is in M .
- (3) f belongs to M , but e does not.

We then define the *graft* of this edge pair onto T to be the union $T^0 = T \cup \{e, f, g\}$. The first property of the edge pair ensures that T^0 is a tree, and the others ensure that T^0 is alternating.

Remark 8 ([Kuh10], [CCPS09, Proposition 5.7]). Used together, these operations make up the Hungarian algorithm, Algorithm 1, for producing maximum matchings on bipartite graphs.

Example 9. We illustrate using an alternating tree to find a maximum matching on a bipartite graph in Figure 2.

2.2 Contract and expand blossom

We now trade the bipartite assumption for two new tree operations. Consider the situation of Figure 3. The alternating tree T is “maximally grafted”, but no edge emanating from its positive vertices reaches an unmatched vertex outside of T , so the algorithm of Algorithm 1 cannot make progress. Nonetheless, an augmenting path exists: starting at r , one can proceed

Algorithm 1: Hungarian algorithm

Data: Bipartite graph G , intermediate matching M

Result: Maximum matching on G

while true do

$G_L [G_R$ a vertex 2-coloring of G ;

T an unmatched vertex in G_L ;

while true do

if there is an $e = (v, w)$ with $v \in T \setminus G_L$,

$w \in G_R, e \notin M$ **then**

 augment T along e ;

break;

else if there is an $e = (v, w)$ with $v \in T \setminus G_L$,

$w \in G_R \setminus M$ **then**

m match edge for w ;

 graft m onto T using e ;

else

return;

end

end

end

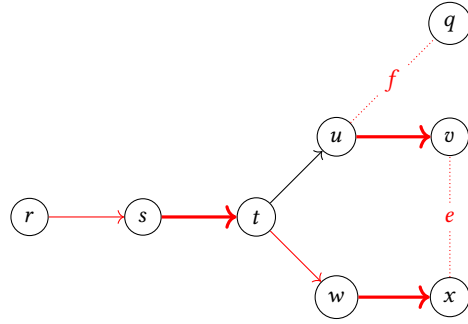


Figure 3: A non-bipartite scenario, where edge e joins two positive vertices. From the view of Algorithm 1, it is illegal to augment along the edge f because u is not positive, hence the path through T to its root r is not alternating. However, the edge e could be used to produce an augmenting path, in red.

down the lower branch, cross vertically along e to the upper branch, walk backwards through the tree to u , and finally cross f to q . These cycles, where an edge not in T joins two of its positive vertices, are the essential new complication of the non-bipartite case. Edmonds’s first fundamental observation was that *all* of the vertices within such a cycle are well-suited to constructing an augmenting path, and the second was that incorporating this into the search algorithm permits one to use it on an arbitrary graph.

Definition 10. Let M be an intermediate matching on G , and let $v \in G$ be a vertex. By an *alternating cycle rooted at v* , we mean an odd-length alternating path $C \subseteq G$ of distinct edges leading from v and returning to v .³ We say that we *contract a blossom* from C when we contract (the full subgraph spanned by) C to a point to produce a new graph $G^0 = G/C$. We refer to the vertex B in G^0 which is the image of C as the *blossom* or the *macrovertex*.⁴ The graph G^0 inherits a matching M^0 , defined through three cases:

- (1) If v is matched in M , B inherits that match in M^0 .
- (2) The matched edges in M internal to C are discarded, as they’ve been contracted out of G^0 .
- (3) All other matched edges in M do not interact with C , so M^0 inherits them verbatim.

Example 11. In Figure 4 we illustrate macrovertex contraction.

Remark 12. Note that, even if G is singly edged (i.e., is not a multigraph), this need not be the case for the derived graph G^0 after contracting a cycle $C \subseteq G$ to a macrovertex $B \in G^0$. If there are two distinct vertices $c, c' \in C$ both with edges to a third vertex $d \notin C$, then B inherits two distinct edges with target d .

³In particular, C begins and ends with unmatched edges, lest v participate in two matched edges.

⁴Some authors refer instead to the alternating cycle as the blossom.

Definition 13. Conversely, suppose that we are given an intermediate matching N^0 on a graph G^0 , where G^0 was originally contracted from a graph G along an alternating cycle C . The matching N^0 can then always be lifted to a matching N on G , called the *expansion* of N^0 along C . Namely, note that the macrovertex participates in at most one matched edge in N^0 , which can be identified uniquely with an edge in G incident on some vertex $w \in C$. By rotating the alternating pattern of matches within C so that the successive pair of unmatched edges is joined at w , and otherwise inheriting the matched edges from N^0 , we obtain a matching on N .

Example 14. We illustrate match expansion in Figure 5.

Remark 15 ([Edm65b], [CCPS09, Theorem 5.10]). As promised, we use these operations to extend Algorithm 1 to cover non-bipartite graphs. We also modify the termination condition: given a maximally-grafted alternating tree $T \subseteq G$, if it does not admit any exiting edges along which we may augment, we additionally search for unmatched edges between positive vertices in T . If such an edge is present, then we use it to form a minimum alternating cycle C within T , and contract that cycle to form a new graph G^0 with matching M^0 . By contracting T along C , we also produce a new alternating tree T^0 in G^0 , and we proceed to run the matching algorithm from this new state.⁵ When this recursion returns, we either lift the modified matching on G^0 to a modified matching on G (via macrovertex expansion) and restart the outer loop, or we proceed to try the next unmatched vertex as a root.

Remark 16. Representing a contracted graph in memory is somewhat onerous. In particular, the algorithm described in Remark 15 may involved nested contractions, where a vertex becomes contracted into a macrovertex, which in turn becomes contracted into another macrovertex, and so on. We defer discussion of this point to Section 3.4, where we will describe it in full in the setting most relevant to this paper.

⁵Some authors call this the *derived graph* and *derived state*.

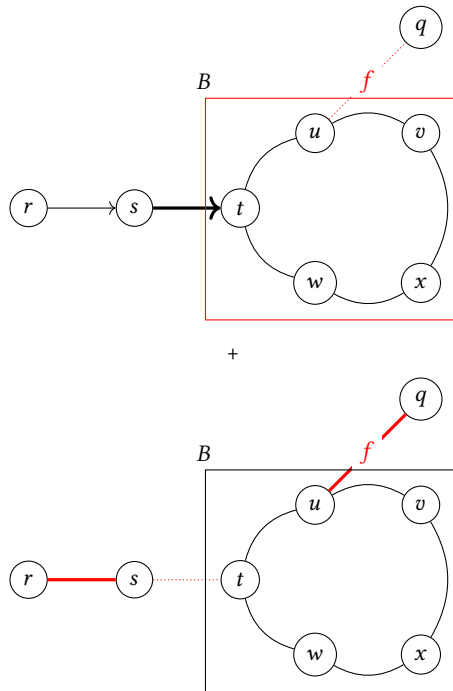


Figure 4: Continuing from Figure 3, we contract the cycle formed by the edge e into a macrovertex B . This causes f to be attached to a positive vertex in G^0 , and hence it participates in an augmenting path. Note that edges in the contracted cycle are curved solid lines.

2.3 Reweight

Finally, Remark 15 can be extended to produce maximum matchings of minimum weight by housing it in a “primal-dual update” scheme [DFF56]. Speaking very loosely, the dual step sorts the edges not yet considered by the amount of weight their participation would incur, and the primal step consists of Remark 15 as applied to these minimally-sifted graphs.

Definition 17. The *internal weight* of a (macro)vertex is a numeric value managed by the dual step. The *adjusted weight* of an edge is its weight after subtracting the internal weights of its two endpoints and those of any macrovertices to which they belong. The *weightless subgraph* G is then the maximal subgraph with the same vertices but retaining only edges of adjusted weight zero.

Remark 18 ([CCPS09, Theorem 5.20]). The internal weight of a vertex may be negative. However, if the edge weights of a graph are all nonnegative, then so are the algorithm’s internal weights and adjusted edge weights.

Definition 19. Let G be an edge-weighted graph with internal vertex weights, M an intermediate matching on G , and T an alternating tree in G . The *reweighting* of T is an update to the internal weights of G given by increasing the internal weights of the positive vertices T_+ and decreasing the internal weights of the negative vertices T_- by the amount given by the minimum of the following three sets:

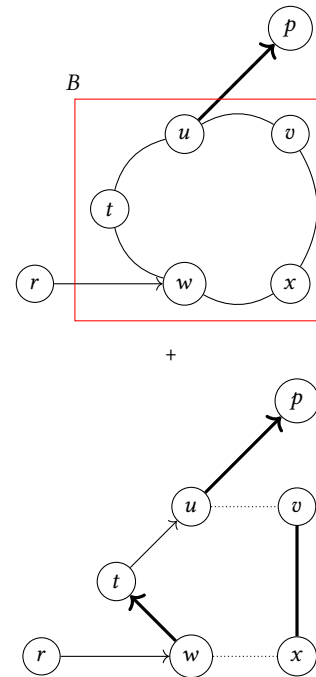


Figure 5: Beginning with a scenario similar to that of Figure 4, but with B in a negative position, we demonstrate the effect of macrovertex expansion. The edges within the cycle are tagged as matched edges so as to provide an alternating path from the root to the leaf, and edges not along that path are ejected from the tree.

Graft/augment candidates The adjusted edge weights of edges in G joining vertices in T_+ to vertices not in T .

Contract candidates The adjusted edge weights of edges in G , scaled down by half, joining vertices in T_+ to each other.

Expand candidates The internal weights of vertices in T_- which are macrovertices.

This enlarges G while maintaining the weightlessness of T .

The edge-weighted blossom algorithm is then given by applying the non-bipartite blossom algorithm to G , with the additional step that the alternating subtree should attempt a reweight operation before the loop gives up and moves on to the next candidate root vertex.

Example 20. We illustrate a sample run of this algorithm in Figure 6.

2.4 The main loop

Altogether, these operations make up the serial blossom algorithm, Algorithm 2, for producing minimum-weight maximum matchings on edge-weighted graphs.

Remark 21. To simplify the outer loop slightly, we have assumed while writing Algorithm 2 that G is fully connected with an even number of vertices. This means that *all* of the vertices will participate in a maximum (perfect) matching, and

A distributed blossom algorithm for minimum-weight perfect matching

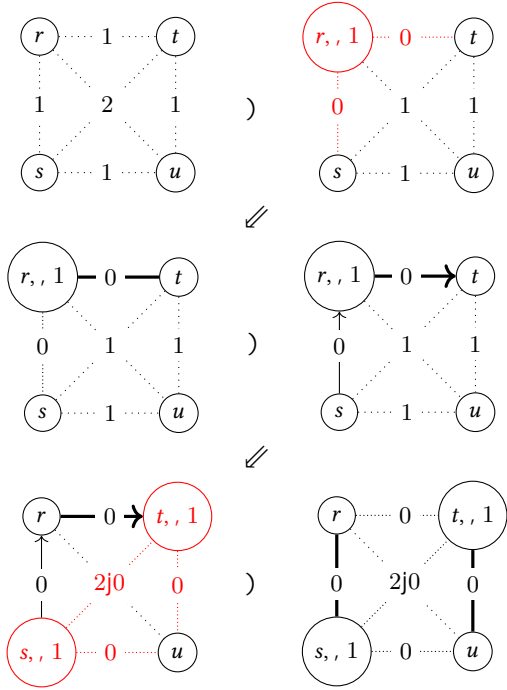


Figure 6: A demonstration of the weighted blossom algorithm. Beginning with the root vertex r , we first minimally increase its weight so some of its edges become weightless (the densely dotted edges). We then select one such edge along which to augment the matching. We choose s to be our next root, to which we graft the previous matched edge. We are out of weightless edges on which to act, so we minimally reweight the tree to produce more. We choose the augmenting path $utrs$, along which we augment to produce the final matching.

the objective is only to minimize the total weight. Matching all of the vertices can be used as a stopping condition.

THEOREM 22 ([EDM65A], [CCPS09, THEOREM 5.16]). *Algorithm 2 is correct: it always terminates, and on termination it emits a perfect matching of minimum weight.* \square

Remark 23. There is a variant of the algorithm in which the main loop resets the tree T after each primal and dual state update, i.e., each update other than grafting. This variant is less strict than Algorithm 2, in the sense that it admits more execution paths: it can always recreate the state preserved here which it has erased, but Algorithm 2 is not always able to erase grafting decisions it has made. In trade, Algorithm 2 performs less recomputation as the state evolves, giving it better runtime properties.

3 THE DISTRIBUTED ALGORITHM

In this section, we rebuild the serial weighted blossom algorithm from Section 2 in a concurrent framework. Our intention is to alleviate a fundamental bottleneck in the serial algorithm: the outer loop fixes an unmatched vertex v to use as the root

Algorithm 2: Serial blossom algorithm

Data: Edge-weighted fully-connected graph

$G = \langle V, E \rangle$ with $|V|$ even

Result: Minimum-weight perfect matching on G
weight all the vertices in G to 0;

while G has unmatched vertices **do**

T an unmatched vertex in G ;

while true do

T_e vertices of T of even depth;

T_o vertices of T of odd depth;

G subgraph of weightless edges;

if there is an $e = \langle v, w \rangle \in G$ with $v \in T_e$,

$w \notin T$ **then**

augment T along e ;

break;

else if there is an $e = \langle v, w \rangle \in G$ with $v \in T_e$,

$w \in T$ **matched then**

m match edge for w ;

graft m onto T using e ;

else if there is an $e = \langle v, w \rangle \in G$ with

$v, w \in T_o$ **then**

a nearest common ancestor of v and w ;

C alternating cycle rooted at a ;

contract C into macrovertex B using e ;

else if there is a weightless macrovertex $B \in T_o$

then

expand B ;

else if $r \neq 0$ for r the candidate reweight

amount for T **then**

reweight T by r ;

else

break;

end

end

end

while G contains macrovertices **do**

B a macrovertex in G ;

expand B ;

end

of an alternating tree T , and all operations proceed in view of T . One can imagine a variant of the algorithm which instead constructs a forest of alternating trees, and one can further imagine a decentralized variant where each tree in the forest⁶ is responsible for “managing itself”. This is the manner of algorithm which we now describe.

We adopt the language of the DECOUPLED model of computing [Lin92, CDF’19, DFFR19], in which we consider an actor-based programming model [HBS73] which rides atop a message-passing system with guaranteed, ordered delivery of messages. In a prequel to this paper, we described a specific such system [PK20] as well as an emulator for it [aet], and as a companion to this paper we provide an implementation of our algorithm within that framework [ana].

⁶Indeed, each vertex in the tree.

3.1 The distributed environment

Our preferred model of distributed computing consists of a family of actors, called *processes*, each with a public address, an interrupt table of handlers which service messages arriving at the public address, and a continuation representing a current computational state. On each computational step, the process walks every message received at its public address using the message handlers, then evaluates its current continuation, which either produces a new continuation or terminates the process. These components are subject to some basic guarantees:

- A non-terminated process will eventually act again.
- A message sent between processes will eventually arrive.
- Messages sent from the same originating process to the same destination address will arrive in order.

In practice, the precise timing of these operations is influenced by many factors (e.g., network pressure), and we assume no guarantees about synchronicity or bounded delay. As is common in distributed graph algorithms, we will spawn a process for each vertex in the problem graph, and their communication will run along edges in the problem graph. The job of these processes, then, is to coordinate with one another, decide which vertices are matched to which others, and signal when they have finished working.

Rather than embed the problem instance directly into the processes, we instead provide an oracle service, called the *dryad*,⁷ which responds to API requests with details about the problem graph. These messages are:

- message-discover** The dryad replies with a list of vertices which are connected by edges to the querying vertex.
- message-sprout** The sender vertex announces that it has begun to participate in the intermediate matching.

3.2 The blossom main loop

As in the serial case, the operations available to an alternating tree are: graft an external matched edge, augment through another alternating tree,⁸ form a macrovertex from a cycle, expand a macrovertex, reweight its vertices, and do nothing (which we call “pass” or “hold”). We will first discuss how an alternating tree selects among these operations, called a *scan*, turning to the distributed implementation of these operations only in subsequent sections.

Since this procedure makes use of the alternating tree structure, we will need to understand how this state is maintained by the algorithm. We store the relevant information in the following slots on each blossom process, where a “blossom” may represent a vertex or a family of (macro)vertices contracted into a macrovertex:⁹

addr The public address of the blossom. Messages sent to this address will be handled by the blossom.¹⁰

⁷As in: a being which attends to the health and welfare of flora.

⁸In the serial case, the other tree is a lone vertex, joined to a positive vertex in the tree by an unmatched edge.

⁹The convention we use going forward of having “blossom” mean either vertex or macrovertex is our own, and differs from the one used in Section 2.

¹⁰Unless explicitly stated, when we refer to blossoms we are referring to their addresses, not the processes themselves. We use Greek letter and calligraphic variables to distinguish addresses from other data.

- id** The internal name of the blossom. Blossoms also have at their disposal a function called `edge-weight` which takes in two `ids` and returns the weight of the edge between them.
- match-edge** The edge connecting this blossom to the blossom to which it is matched, if any.
- parent** The optional edge connecting this blossom to its parent in an alternating tree, if it is not the root.
- children** A list of edges connecting this blossom to its children in an alternating tree, if any.
- positive?** Parity of the distance to the root of the alternating tree: true if even, false if odd.
- pistil** If this blossom has been contracted as part of a cycle, this holds the address of its immediate “parent” macrovertex.
- petals** If this blossom is a macrovertex, this stores the cyclic list of edges which were contracted into it.
- internal-weight** The internal weight of the blossom.
- pingable** A flag indicating whether the blossom responds to inbound ping requests (see immediately below) or it allows them to queue for later perusal.
- paused?** If toggled, continue to respond to messages but take no other actions.
- dryad** The address of this blossom’s dryad.

Most operations that an alternating tree might perform correspond to the edges emanating from the vertices in the tree. However, the action to which an edge corresponds depends on the states of *both* of the edge’s vertices: the “target” vertex may belong to the same tree, a different tree (and at even or odd height), or to no tree at all. To discern the appropriate action, the (positive) source vertex sends a `message-ping` to a target vertex with its half of the data, then listens for a `message-pong` in reply with the calculated operation. Attached to most calculated operations is the sponsoring edge, which in this context is a data structure with 4 address slots: `source-blossom`, `source-vertex`, `target-vertex`, and `target-blossom`. The vertex-suffixed slots are self explanatory; the blossom-suffixed slots return the address of the topmost macrovertex associated with this vertex (if one exists), and otherwise return the vertex address. In this way edges keep track of the vertices and topmost blossoms associated with each endpoint of the directed edge, which will later make it easier to manage macrovertex contraction and expansion. The procedure for calculating sponsored actions and edges is described in Algorithm 3, and the information provided by the `message-ping` includes:

- root** The root blossom of the source alternating tree.
- blossom** The topmost macrovertex to which the source belongs. If the source is not contained within a macrovertex, this is (the address of) the source vertex itself.
- weight** The sum of the source vertex’s internal weight and the internal weights of all macrovertices to which it belongs.
- hold-cluster** A set of roots that will have future implications on how operations are processed (see Section 3.6).¹¹
- addr** The address of the source vertex, used to build the sponsoring edge and send a reply.
- id** The name of the source vertex, used to calculate the underlying edge weight joining it to the recipient.

¹¹We call this `internal-roots` or `internal-root-set` in our implementation, but have renamed it here to more clearly link it to its intended purpose.

Algorithm 3: message-ping handler

Data: a (target) vertex process P , a message-ping m
Result: a message-pong reply containing an action a
 ρ the root of the tree to which P belongs;
 β the topmost macrovertex to which P belongs;
 w the macrovertex-adjusted internal weight of P ;
 H hold-cluster(m);
 e edge(blossom(m), addr(m), addr(P), β);
 w_e edge-weight(id(P), id(m));
 w_e^0 $w_e - w - \text{weight}_{m^0}$;

if P is matched, and match-edge(P) = e **then**
 | a pass;
else if β = blossom(m) **then**
 | a pass;
else if β is negative **then**
 | a hold(ρ);
else if w_e^0 is nonzero **then**
 | **if** ρ = root(m), or $\rho \notin H$ and root(m) $\notin H$ **then**
 | w_e^0 $w_e^0/2$;
 end
 | a reweight(w_e^0);
else if $\rho \neq$ root(m), and ρ is matched **then**
 | a graft(e);
else if $\rho \neq$ root(m), and ρ is unmatched **then**
 | a augment(e);
else if ρ = root(m) **then**
 | a contract(e);
end
send message-pong(a) to addr(m);

Remark 24. The slots on a message-ping and on a blossom process are locally accessible during the execution of Algorithm 3. But, of course, not all data that we use during the course of an algorithm are locally computable. For example, P might send messages to other processes to compute the addresses β and ρ (and to compute properties of their underlying processes, e.g. whether they are matched), in Algorithm 3.

We coordinate the different possible pings within a tree using the procedure described in Algorithm 4. This algorithm walks over the edges attached to the vertices participating in the tree, asking each to sponsor an operation. Because these operations carry a preference order, in the sense that the availability of one operation can preclude the consideration of another, we can then unify the responses into a single course of action according to the following ordered set of rules:

- (1) It is possible for an edge to recuse itself from sponsoring an operation, typically from a misguided ping (e.g., if a vertex should ping itself). We call this situation a pass. Given a choice between a pass and any other option, we will always prefer to act by the other option.
- (2) A hold operation arises when the ping targets a negative vertex. In the serial algorithm there are no interactions between positive and negative vertices, but there are implications in the distributed setting which will require us

Algorithm 4: message-scan handler

Data: a blossom process B , a message-scan m
Result: (optional) a supervisor sponsoring an action a
 ρ root(m);
 β blossom(m) **if** pistil(B) **else** addr(B);
 w weight(m), internal-weight(B);
 H hold-cluster(m);
 α addr(B);
 δ dryad(B);
 E children(B);

if positive?(B) **then**
 | a pass;
 | E $E \sqcup$ petals(B);
else if B is a macrovertex **then**
 | a expand(α);
end

if positive?(B) and B is not a macrovertex **then**
 send message-discover to δ , store response in N ;
 foreach $v \in N$ **do**
 | p message-ping($\rho, \beta, w, H, \alpha, \text{id}(B)$);
 | send p to v , store response in a^0 ;
 | a unify-pongs(a, a^0, H);
 end

end

foreach $e \in E$ **do**
 | τ target-blossom(e);
 | s message-scan($\rho, \beta, w, H, \alpha$);
 | send s to τ , store response in a^0 ;
 | a unify-pongs(a, a^0, H);
end

if addr(m) **then**
 | send message-pong(a) to addr(m);
else if B is still an eligible root, and $a \neq$ pass **then**
 | pause B and spawn a supervisor sponsoring a ;
end

to process holds distinctly from passes (see Section 3.6). For now, we will treat both operations the same.

- (3) As seen in Algorithm 3, hold operations are initialized with the root ρ of the target vertex. The operation stores this root in a slot called root-bucket, which is a set (of length one, to start). Given a choice between two holds, we instead hold with the union of the root-buckets of each individual hold.
- (4) Given a choice between a reweight and any operation other than a reweight, we prefer the other operation.
- (5) Given a choice between two reweight operations, we prefer to reweight by the lesser amount.
- (6) Finally, given a choice between two nontrivial operations, we arbitrarily prefer one over the other.

Altogether, these rules define the method unify-pongs referenced in Algorithm 4.

Finally, there are the matters of initiating a scan over a tree and acting on the resulting sponsored operation. We delegate

the responsibility of initiation to the root: each unmatched blossom whose `paused?` flag is not set to `true` checks whether it is a root in the forest and, if so, sends itself a `message-scan`, then awaits a reply. The slot specification for a `message-scan` is identical to that of a `message-ping`, but without an `id` slot. When sending the initial `message-scan` to itself, the root fills out the message slots as follows: `root` is the root’s address, `blossom` is null, `weight` is zero, `addr` is null, and `hold-cluster` is a set containing just the root’s address. If the result of this scan is not a pass, and if the root has not changed state since the start of the scan, it spawns a *supervisor* process seeded with the result and sets its own `paused?` flag to `true`. It is then the supervisor’s responsibility to carry out the operation and to unpause the root when complete.

Remark 25. In our implementation, we describe the operation sponsored by an edge as a pair of an atom (`pass`, `graft`, `augment`, `contract`, `expand`, `hold`) and a numeric value which records the adjusted weight of the edge. In particular, we do not use a separate atom for reweighting: a `reweight` directive is instead recorded by an edge of nonzero weight. The atom is then determined by the rules in Algorithm 3 and Algorithm 4.

3.3 Augment and graft

We turn now to modifying the grafting and augmenting operations for distributed use. These are the simplest operations, which lets us ease into the problem while devoting extra time to the general scaffolding we re-use in future operations.

Recall that each operation is enacted by a supervisor, which is spawned by a root blossom at the conclusion of a scan. Each such supervisor operation follows the same outline:

- (1) Acquire (recursive) locks on the trees involved [PK20]. This has the effect of preventing all locked blossoms from starting scans. As part of locking, we additionally set the `pingable` flag to `false` to prevent locked blossoms from responding to pings, in order to avoid exposing incomplete state. If it is not possible to establish any of the locks, abort.
- (2) Check the root blossoms advertised to the supervisor. If they are no longer unmatched roots, abort.
- (3) Check that the sponsored action is still valid. For most actions (`graft`, `augment`, `contract`), this means sending another `message-ping` to check that the edge responsible for sponsoring the action still sponsors the same action. For `expand`, this means checking that the sponsoring blossom is still a negative macrovertex. We will defer discussion on checking the validity of `reweight` and `hold` to later sections. If the sponsored action is no longer valid, abort.
- (4) Enter the critical section. Send messages instructing the locked vertices to make the appropriate state changes.
- (5) Exit the critical section, release the locks, restore ping responsiveness by setting the `pingable` flag to `true`, and terminate the supervisor process.

For the `graft` operation, the supervisor is required to lock the tree from which it came, as well as both blossoms participating in the isolated matched edge to be grafted.¹² The critical section is then constituted of three pairs of set instructions:

¹²In our implementation, we refer to an isolated matched edge as a *barbell*. Each blossom in a barbell considers itself to be the root.

Algorithm 5: message-augment handler

Data: a blossom process B , a message-augment m
Result: nothing
 e preceding-edge(m);
 δ dryad(B);
if $e = \text{match-edge}(B)$ **then**
 | match-edge(B) parent(B);
else
 | match-edge(B) e ;
end
if B is the tree root **then**
 | send a message-sprout to δ ;
 | send a “done” message to $\text{addr}(m)$;
else
 | e reverse(parent(B));
 | π target-blossom(parent(B));
 | send message-augment($\text{addr}(m)$, e) to π ;
end

- Set `positive?` on both blossoms in the matched edge.
- Set `parent` on both blossoms in the matched edge.
- Set `children` on both blossoms in the (unmatched) edge which sponsored the action.

The `augment` operation is more complex. After locking both the source tree and target tree, the supervisor sends `message-augment` to the two blossoms in the sponsoring edge. This message has a slot `addr` containing the supervisor’s address and a slot `preceding-edge`, initially containing the sponsoring edge. This causes the blossoms to modify their matches and forward the `message-augment` toward the roots of their respective trees; details are given in Algorithm 5.¹³ An additional twist is that, when releasing the recursive lock, we instruct all of the locked blossoms—not just the ones participating in the `message-augment` chain—to reset their `parent` and `children` fields. This disassembles the tree structures, which are no longer alternating.

3.4 Contract and expand blossom

We turn to even more complicated tree operations: first, the construction of macrovertices; and later, their dissolution.

As indicated in the previous section, macrovertex construction follows the same general template of locking and re-checking before entering the critical section, whose behavior we now specify. To begin, the supervisor spawns a new blossom process which will serve as the macrovertex and immediately acquires a lock on it to prevent it from acting on its own (viz., initiating a scan). It then takes the following steps:

Petal calculation The supervisor computes the paths in the alternating tree from the sponsoring blossoms to the root, trims edges which are common to both paths, and joins them through the sponsoring edge. This is the minimal alternating cycle containing the sponsoring edge, and is

¹³We introduce here a function `reverse` to reverse the direction of an edge. Note that the supervisor needs to reverse the sponsoring edge before sending it via `message-augment` to the target tree.

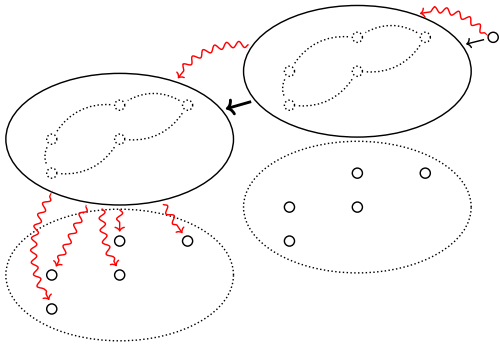


Figure 7: Illustration of the propagation of scan messages, denoted by wavy red arrows, along an alternating tree with macrovertices. Scan messages are forwarded first among blossoms along the edges that constitute the alternating tree, and then each positive macrovertex which receives a message propagates it to the blossoms which make up the cycle which it encloses.

used to set the petals slot on the new macrovertex.¹⁴ The blossoms in the cycle set their pistil to the macrovertex.

Parent calculation The final such edge trimmed above, if any, connects to the parent of the macrovertex. The supervisor substitutes the new macrovertex into that parent’s children, and also sets the parent of the macrovertex.

Children calculation Where appropriate, the blossoms in the cycle transfer the parent-child relationship of their children to the macrovertex.

Match calculation The macrovertex inherits the match edge of its contracted blossom which is nearest to the root, if it is matched. The blossoms in the cycle erase their match edges.

Unpause Finally, the macrovertex is unpause.

Remark 26. In this way, alternating trees with macrovertices become “two-dimensional trees”: there is a parent-child relation which captures potential alternating paths in the contracted graph, as well as a parent-child relation which captures the subsumption of blossoms by macrovertices. We illustrate this structure in Figure 7.

Expansion is performed “in reverse”: all of the relations that went into macrovertex formation are recovered whenever they can be, and they are reverted to a neutral state otherwise. The amounts to the following operations:

External parent calculation The parent of the macrovertex holds an edge to the macrovertex in its list of children. The supervisor “unwraps” the target of that edge by one layer: it replaces the edge by an edge with *target-blossom* set to the penultimate blossom ancestor of its *target-vertex*. It does this by asking *target-vertex* to recursively message its *pistil* to find the process whose *pistil* is equal to

target-blossom (which, if *target-vertex* has only been part of one macrovertex contraction, would be itself).¹⁵

External match calculation The match of the macrovertex holds a pointer to the macrovertex in its match edge. In the same fashion as external parent calculation, the supervisor descends that edge by one layer of macrovertex contraction.

Internal match calculation The match edges within the cycle are set between neighbors so that the unique cycle blossom receiving an external match edge does not participate in a match within the cycle.

Internal parent and children calculation The blossom algorithm makes two guarantees when expanding macrovertices: a macrovertex is always matched; and it either does not participate in an alternating tree (viz., when expanding a blossom at the termination of the algorithm) or, if it does, it lies in a negative position with zero internal weight (viz., when expanding a blossom during the bulk of the algorithm). In the second case, it is thus guaranteed to have at most one child (viz., its match) in addition to having a single parent. Accordingly, by choosing between clockwise and counterclockwise traversal, the supervisor can find a unique path through the petals which continues the alternating pattern from the ambient tree. These edges are assigned parent-child relationships, and all edges not along the path are ejected from the tree.

3.5 Reweight

Finally, we consider the most complex of the serial algorithm’s tree operations. Although reweighting only involves setting values internal to a tree, the *validity* is influenced by an indeterminate number of trees. It is not feasible to ensure the validity of a reweighting operation by acquiring locks, since pessimistically we would have to lock the entire platform. Instead, we perform a tentative change, then undo (or “rewind”) if that change is interrupted or seen to be invalid.¹⁶

The possibility of an invalid reweighting is a wrinkle unique to the distributed setting: a change can only become invalid if two trees, which are mutually influencing one another’s scans, both elect to reweight at the same time. They can also only detect this invalidity by re-scanning—which, without intervention, would result in deadlock, as both trees refuse to reply to inbound pings in their respective critical sections. In light of this, we introduce a new kind of ping and new state of ping responsiveness, called *soft pings*, which are used only to establish weight validity and not to sponsor new operations.

pingable This existing flag is extended from a simple on/off switch to three values: *all*, meaning the blossom responds to all pings (replaces *true*); *none*, meaning the blossom defers all ping responses (replaces *false*); and *soft*, meaning the blossom only handles soft pings and defers others.

With this new pingability mode in place, we can define how a supervisor checks the validity of a reweight (which we deferred during step 3 in the supervisor outline in Section 3.3).

¹⁴In Figure 3, the edge $e = „v, x”$ would sponsor a macrovertex formation, and the two paths from the root would read as $rstuv$ and $rstwx$. After trimming the common prefix, one resulting alternating cycle is $tuwx$.

¹⁵By recursively message we mean send a message to its *pistil*, and then have that process send a message to its *pistil*, and so on.

¹⁶The simplest form of validation is to apply Remark 18 and check for negatively-weighted edges.

First, the supervisor tells the source and target roots to set themselves and their trees to only respond to soft pings.¹⁷ Then, the supervisor instructs the source root to initiate a *soft scan* (a scan that generates soft pings). If the resulting recommendation differs from what was originally sponsored, the supervisor aborts.

This change is sufficient to maintain consistent state, but it can still result in livelock [Ash75]: two trees competing for the same reweighting operation in perfect synchrony can forever tentatively reweight, regret, rewind, and repeat. To avoid this scenario, we use a priority scheme to break the symmetry: if a tree in the process of reweighting receives a (soft) ping from a higher-priority source which evidences that it has changed its weight by too much, it aborts its own reweight in order to make room for its superior to perform the operation instead.¹⁸

With all this in mind, we describe steps involved in the critical section of the reweighting operation:

- (1) Set the source tree's pingability to none.
- (2) Modify the internal weights of the top-level blossoms in the source tree: the positive blossoms are increased by the desired amount, and negative blossoms are decreased.
- (3) Set the source tree's pingability to soft t.
- (4) Tell the source root to perform a soft scan to determine the minimum edge weight emanating from this tree to another. This message contains the same arguments as described in Section 3.2, except *addr* is the supervisor's address.
- (5) If this minimum edge weight is negative, rewind.¹⁹

Example 27. See Figure 8 for an illustration of this resolution.

Remark 28. Rather than preventing livelock through priorities, an alternative strategy when two trees conflict through simultaneous reweighting is to have each tree rewind by half of the overshoot,²⁰ check again, and rewind the rest of the way if there is still a problem. This is an imperfect form of resolution, but it nonetheless seems to be useful in practice. See Figure 9 for an example of this conflict resolution.

3.6 Multireweight

Finally, we come to the crucial new feature that arises in the distributed algorithm. The algorithm described so far suffers from an additional deadlock which arises not because of communication conflicts, but because of decentralized action. Where the serial algorithm could exhaust the possibilities for a tree to act, *release the active tree*, and move on to the next possible root, the decentralized algorithm builds multiple trees at once and *has no release mechanism*. Since the hold operation previously indicated that the serial algorithm should move on, this results in a possible failure mode when handling it in the distributed setting.

¹⁷We additionally modify the validity checks for *graft*, *augment*, *contract* to do the same before sending their verification ping.

¹⁸It is possible, but not necessary, to employ this symmetry-breaking mechanism with the locks acquired by other operations.

¹⁹This is the point at which one could only rewind halfway and check again. In the priority-preferenced scheme, only a higher-priority vertex or a vertex which has finalized its critical section can emit a negative-weight reply, and in both these cases we must rewind.

²⁰The factor of 1/2 comes from each 1 edge being attached to 2 vertices.

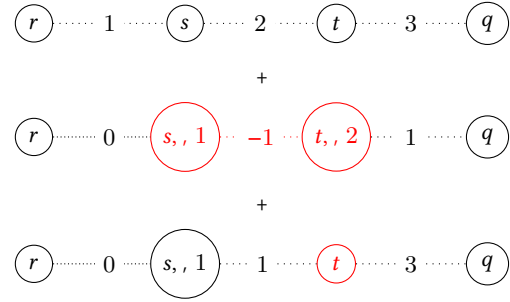


Figure 8: A demonstration of a reweight conflict. Two vertices s and t both attempt to reweight according to their local understanding of the largest value by which they can reweight. Before exiting the critical section, they both notice that the edge connecting them has negative adjusted weight, which is illegal. Taking t to have lower priority than s , t reverses its reweight whereas s retains it. This repairs the violation, and the algorithm has made progress by making the edge „ r, s “ weightless.

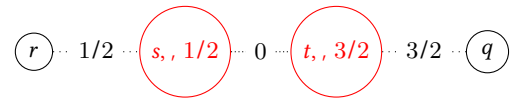


Figure 9: Resolving the conflict in the second step in Figure 8 by partial rewinding instead. This time the resulting weightless edge is „ s, t “.

Example 29. For an example configuration, see the second step in Figure 10. These two trees are arranged so that they both have some positive weight on their negative vertices, and they both have positive vertices with weightless edges to each others' negative vertices. This means that the trees abut along edges that sponsor hold operations, and neither one can perform any other operation, including a reweight.

However, the two trees *could* make progress if they coordinated their action to simultaneously transfer weight from their negative vertices to their positive vertices. This kind of coordinated reweighting is what is pictured in the transition to the third step of Figure 10.

To codify this kind of coordination, we introduce a variant of the reweighting operation, called a *multireweight*, which is triggered by the presence of a hold cluster:

Definition 30. A *hold cluster* is a set of trees whose scans all elect to hold and all of whose edges sponsoring those holds have both endpoints in the set.²¹

A supervisor responding to a sponsored hold operation deviates from the typical recipe described in Section 3.3. Namely, it performs the preflight validation checks for the operation (step 3) *before* locking and checking the roots (steps 1 and 2). For a hold, this validation check consists of verifying that the tree is actually part of a nontrivial hold cluster. It does this by messaging all the roots in the sponsoring operation's

²¹In particular, there are no valid hold clusters of size 1.

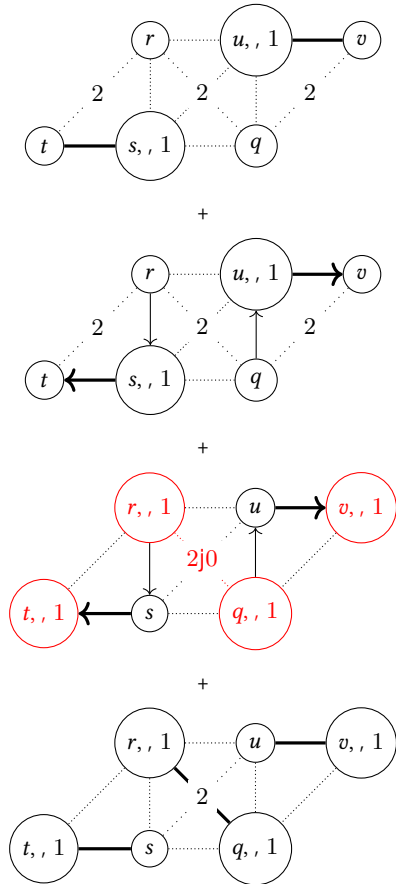


Figure 10: An example of a multireweight operation. In this graph, the edge weights are determined by Manhattan distance. First, r and q act as roots, and they each graft the indicated matched edges onto their respective trees. This causes a serial impasse: the only remaining weightless edge connects two negative vertices, and neither tree can reweight alone because of the weightless edge connecting its root to the other tree’s negative vertex. However, the two trees can together perform a multireweight by 1, which redistributes the internal weight off of the negative vertices and onto the positive ones, resulting in a weightless edge connecting the two roots which can be augmented to produce a perfect matching.

root-bucket to ask if they are held (and by whom), and aggregates the responses. If new mutually-held trees (beyond the root-bucket) are encountered, the message is forwarded along to them. This has the effect of incrementally growing a set of trees along edges which sponsor hold operations. Once no new trees are proposed, then this set is a valid hold cluster; however, if a new tree is proposed that sponsors a non-hold operation, then the cluster is invalid and the supervisor aborts.

Once the supervisor has established a valid hold cluster, and if it was spawned by the root with the highest priority among the cluster, it proceeds through steps 1 and 2 of the standard recipe. This locks all the trees in the hold cluster and

then checks that their roots have not changed since the spawn of the supervisor, otherwise aborting. Then, skipping step 3 (as we performed validation already), the supervisor reaches the `multireweight` critical section. Before describing the way in which a supervisor enacts a `multireweight` operation, we must first (as promised in Section 3.2) make a modification to how holds are processed by `unify-pongs`. Rather than always treating it like a pass and preferring the other operation, `unify-pongs` only discards a hold when its root is part of the hold cluster, and otherwise considers it a nontrivial operation.

With this modification in place, we can finally describe the `multireweight` critical section:

- (1) Set the pingability of the hold cluster to `soft`.
- (2) Perform a soft scan across the entire hold cluster. For each root in the cluster, this soft scan resembles that of the `reweight` critical section in Section 3.5, with one notable difference: the `hold-cluster` slot is populated with the hold cluster motivating our `multireweight`. This has the effect of treating the trees as if they were temporarily grafted onto a common root—the resulting sponsored action is guaranteed to be a `reweight`, with weight equal to half the distance between the closest trees in the cluster.²²
- (3) Set the pingability of the hold cluster to `none`.
- (4) Reweight the hold cluster by the weight from step 2. The supervisor interleaves the `reweight` operations for each tree in the cluster, which has the effect of reweighting the entire cluster simultaneously.
- (5) Set the pingability of the hold cluster to `soft`.
- (6) Perform another soft scan across the entire hold cluster, as in step 2, but no longer filling out the `hold-cluster` slot. This determines the minimum weight edge between the trees in the hold cluster and their surroundings. If this minimum edge weight is negative, rewind the hold cluster.

Example 31. We demonstrate the `multireweight` operation in Figure 10. This operation is a consequence of constructing an alternating *forest* rather than a lone alternating tree. Starting from the first step of Figure 10, we illustrate how the serial algorithm would have progressed differently in Figure 11.

3.7 The dryad main loop

Finally, we consider the apparatus for starting and stopping the blossom algorithm. The dryad, first mentioned in Section 3.1 as a directory for the solver’s various components, is a natural deposit for this responsibility. The interface that it exposes between the solver and the outside world consists of the following messages:

- message-sow** Instructs the dryad to inject a new vertex into the solver. Carries an `id` which is used to refer uniquely to the vertex and to calculate edge weights in the graph.
- message-reap** Sent by the dryad to `match-address` (a slot on the dryad) to announce a single matched edge in the solution, in the form of a pair of `ids`.

Example 32. We describe an implementation of a dryad with some major simplifying assumptions:

²²Compare to the result in Algorithm 3 when two positive blossoms sharing the same root elect to `reweight`.

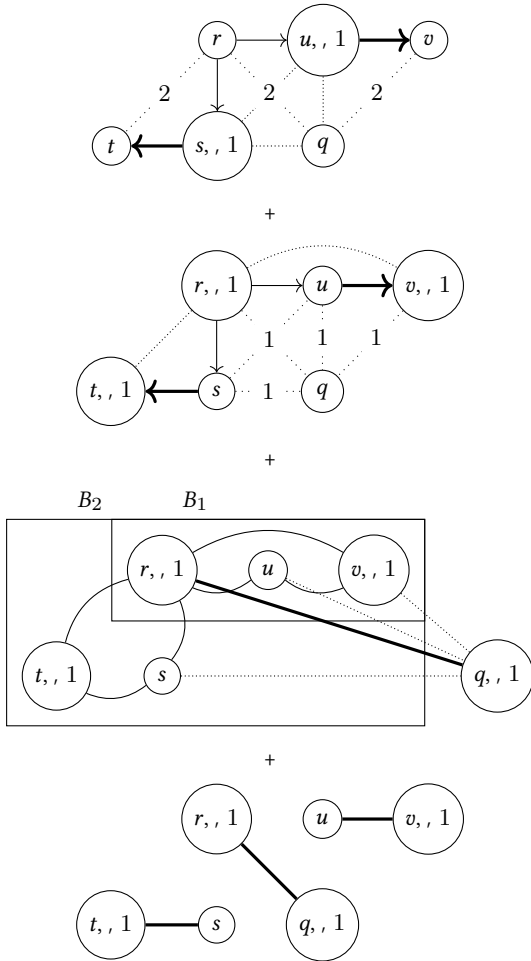


Figure 11: First, taking r as the active root, r grafts the two available match edges and reweights. While q is still disconnected in G , there are two cycles available from which it forms macrovertices B_1 and B_2 . In G^0 , B_2 (or, equivalently, q) can be reweighted, and q can then form a match edge to any of r, s, u , and v . Finally, expanding B_2 and B_1 lifts the perfect matching on G^0 to one on G .

- The dryad is “monolithic”, meaning there is a single process responsible for servicing the API.
- The problem graph is complete, has an even number of vertices, and is fixed for the duration of the solver. Such a graph is guaranteed to have a perfect matching in which all vertices participate, simplifying the termination condition.
- Rather than provide a separate message which installs a plurality of vertices at once, we expect that the plurality of message-sows arrive together.²³

On initialization, the dryad reads the pending message-sows, spawns a process for each, and waits. Each vertex process sends a message-sprout to the dryad when it joins the partial matching; once when every vertex makes that announcement,

²³This assumption has no effect on the functionality of the algorithm if the entire problem is known at initialization time.

the solver has finished. The dryad then queries the vertices for their match edges to announce the solution. If any vertex replies that it does not have a match edge, it is because it is currently wrapped in a macrovertex. Since the solver has finished, the dryad can safely send a directive for that macrovertex to expand, producing matches for all of its petals. This eventually terminates, at which time the dryad announces the set of matched edges by emitting message-reaps.

4 ANALYSIS

We now turn to the efficacy and efficiency of our variant of the blossom algorithm.

4.1 Correctness

As our algorithm is a modification of Edmonds’s algorithm, our proof of correctness also hews closely to his, so we briefly recall the steps he takes. He first analyzes the matching polytope so as to give a combinatorial recognition principle for minimum-weight maximum matchings [Edm65a, Section 4, Theorem (M)]. Namely, a matching $M = M_0$ on a graph $G = G_0$ qualifies when it can be extended to a sequence of graphs and matchings $fM_j \quad G_j g_{j=0}^n$ with the following properties:

- Each G_j is related to the next by a blossom contraction through an alternating cycle.
- Each G_j is assigned “admissible” internal vertex weights.
- The vertex weights differ between G_j and $G_{j,1}$ only at the new macrovertex.
- The contraction edges in each G_j become weightless for these internal weights.
- All matched edges in G_n are weightless, and all unmatched vertices in G_n have zero internal weight.

To show these conditions suffice, Edmonds uses a polytope which encodes the maximum matching problem, with the following executive summary:

[Edm65a, Section 2] A matching gives rise to a vertex in this polytope, and a minimum-weight maximum matching is a polytope vertex which extremizes a linear functional encoding the edge weights.

[Edm65a, Section 5] In general, polytope vertices can be modeled by sequences $fM_j \quad G_j g_j$ of the above type for some choice of edge weights, hence are solutions to some minimum-weight maximum matching problem.

[Edm65a, Section 6] Given a polytope vertex extremizing a fixed edge weight functional, a sequence $fM_j \quad G_j g_j$ for that particular functional can be constructed.

[Edm65a, Section 7] Finally, he provides a description of a primal-dual solver with verifiable progression toward such a combinatorial sequence—essentially, Algorithm 2.

Since we are solving the same problem, we can reuse his recognition principle as-is. In fact, we can reuse most of his proof that his algorithm produces the desired witnessing sequence of combinatorial steps: the sequence $fM_j \quad G_j g_j$ is read off from the steps in the pre-termination while loop from Algorithm 2, and with the same procedure we can read off such a sequence from the final state of the distributed algorithm. We need only show that our algorithm makes progress toward this same goal.

A distributed blossom algorithm for minimum-weight perfect matching

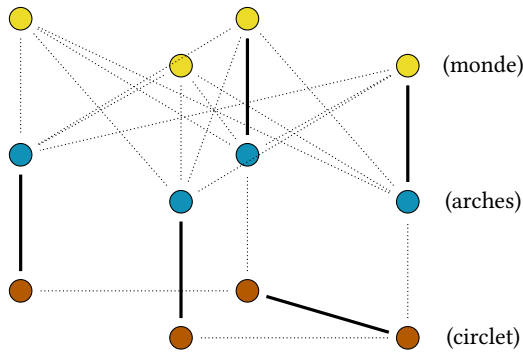


Figure 12: The crown on a four-vertex linear graph. The bottom vertices form the (linear, four-vertex) circllet, the middle vertices and downward edges the arches, and the top vertices and upward edges the monde. The matched edges depict a regal matching extending the matching consisting of just the rightmost circllet edge.

4.1.1 Primal correctness. For primal updates (i.e., for all steps save reweighting), we deduce correctness of the distributed algorithm by realizing its behavior as a special case of that of the serial algorithm. Specifically, the primal updates in the distributed algorithm operating on a graph G behave like the serial algorithm applied to a slightly larger graph $\mathbb{C}G$, which we call the *crown*. We proceed as follows:

- (1) We introduce normalizing conditions on the steps taken by the distributed and classical algorithms. Essentially, a run of the distributed algorithm is said to be “rooted” when its tree operations do not interleave with one another, and a run of the classical algorithm is said to be “good” when it does not involve the new crown graph vertices “too much”.
- (2) We show that every run of either algorithm can be modified so that the respective normalization conditions hold: operations in the distributed algorithm can be judiciously delayed so as to become non-interleaved, and the classical algorithm can be discouraged from manipulating the new vertices in the crown graph without preventing progress.
- (3) In the presence of the normalization conditions, we show that a rooted run of the distributed algorithm corresponds exactly to a good run of the classical algorithm. We do this by directly comparing individual steps in one algorithm to individual steps in the other.
- (4) We use this correspondence to transport Edmonds’s witnessing sequence for a good run of the classical algorithm on $\mathbb{C}G$ to a witnessing sequence for a rooted run of the distributed algorithm on G . Altogether, this establishes primal correctness.

The idea behind the construction $\mathbb{C}G$ is to incorporate the distributed algorithm’s supervisor processes as actual vertices in the problem graph. This then enables us to identify the supervisor actions in the distributed algorithm with the behavior of trees rooted at these vertices in the serial algorithm. The following definition makes this precise:



Figure 13: St. Edward’s crown [SRFR19] with spherical monde at top connected by arches to circllet at bottom.

Definition 33. For an unweighted graph G , we define its *crown* $\mathbb{C}G$ as the following iterated pushout:

$$\begin{array}{ccc}
 & G^{\text{disc}} & \longrightarrow K_{G^{\text{disc}}, G^{\text{disc}}} \\
 & \downarrow & \downarrow \\
 G^{\text{disc}} & \longrightarrow \text{Cyl}_n G^{\text{disc}} & \\
 \downarrow & \searrow & \downarrow \\
 G & \longrightarrow & \mathbb{C}G,
 \end{array}$$

where G^{disc} is the discretization, and

$$\text{Cyl}_n = \text{“} \square \text{”} \square \text{“} \text{! } 1g$$

is the *cylinder graph*, constructed using the graph “Cartesian product”. Concretely, the vertices of $\mathbb{C}G$ partition into three sets, each separately isomorphic to the vertices of G : the *circllet*, the *arches*, and the *monde*. For a vertex $v \in G$, we refer to the corresponding vertices in the circllet, arch, and monde as its *avatars*. The edges joining these regions are as follows:

- Arches–arches:** None.
- Arches–circllet:** Each pair of avatars of v is joined by an edge.
- Arches–monde:** Every arch vertex is joined to every monde vertex by an edge.
- Circllet–circllet:** These are the same as the edges of G .
- Circllet–monde:** None.
- Monde–monde:** None.

Example 34. Figure 12 depicts the crown $\mathbb{C}G$ of a four-vertex linear graph G . The circllet, arches, and monde are vertically arranged. Additionally, we illustrate a regal matching on $\mathbb{C}G$ (see Definition 37) which extends a partial matching on G consisting of a single edge. In Figure 13 we provide the inspiration behind the naming of $\mathbb{C}G$, and in Figure 14 we separately indicate how the crown is constructed as an iterated pushout.

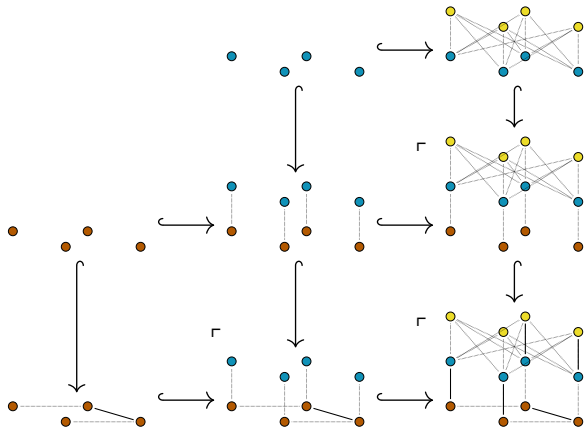


Figure 14: The presentation of the crown on a four-vertex linear graph as an iterated pushout, including the two intermediate pushouts. Each square describes its lower-right corner as two graphs, those in the upper-right and lower-left, glued together along the common subgraph in the upper-left.

We show that these conditions permit the following correspondence, whose proof we briefly defer:

THEOREM 35. *Rooted Lamport orderings of the distributed blossom algorithm acting on G biject with good runs of the serial blossom algorithm acting on $\mathbb{C}G$, modulo the choice of monde vertex at the start of each serial segment. This bijection preserves the sequence of state updates.*

To support this claim, we show in some Lemmas that these conditions are not so limiting. Our first Lemma shows that by deferring irrelevant grafting operations, we can benignly re-ordering the events in a Lamport ordering so that the ordering becomes rooted.

To make precise the relationship between matchings on $\mathbb{C}G$ and matchings on G , we need three auxiliary definitions which help normalize the indeterminacy on both sides of our purported correspondence. We begin with a condition on distributed runs:

Definition 36. In a Lamport ordering [Lam78] of the events of a run of the distributed algorithm, we refer to the sequence of grafting events between two adjacent state updates, as well as the later state update, as a *segment*. A segment is said to be *rooted* if all of its roots belong to the set of roots which act in its final bookending state update. The entire Lamport ordering is said to be *rooted* if all of its segments are rooted.

Our intent with Definition 36 is that rooted orderings display a kind of focused attention. Rather than many tree operations happening in parallel in a forest, all the operations within a maximally reasonable block of time pertain to one tree only.

Next, we have conditions for the serial runs on the crown:

Definition 37. A matching on $\mathbb{C}G$ is said to be *regal* if all of the vertices in the circllet and arches are matched (but perhaps not to each other, i.e., some vertices in the arches may be

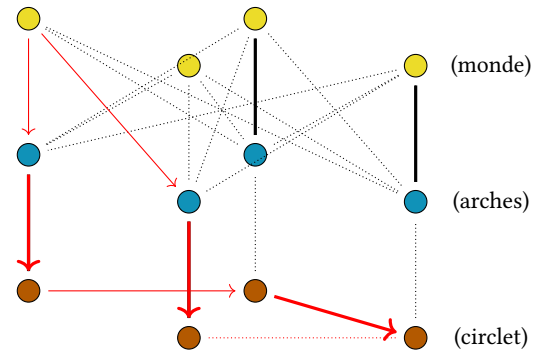


Figure 15: An alternating tree directed along the crown.

matched to vertices in the monde). The *initial regal matching* is the matching where each vertex in the circllet is matched to its avatar in the arches.

Definition 38. A run of the serial algorithm on $\mathbb{C}G$ is said to be *good* if it satisfies the following properties:

- The initial matching is regal.
- Every macrovertex formation in which a monde vertex participates is directly followed by an augment to another monde vertex and an expand operation on the macrovertex.
- Trees are *directed downwards along the crown*: No circllet vertex is permitted to graft an arch vertex as a child.

Our intent behind Definition 37 and Definition 38 is to limit the roles which the new monde and arch vertices can play in the serial algorithm, facilitating a comparison between the distributed algorithm acting directly on G and the serial algorithm's behavior as understood through the circllet.

Example 39. In Figure 15, we then examine an example on the same $\mathbb{C}G$ of an alternating tree directed along the crown. The two circllet vertices without circllet matches engender two unmatched crown vertices, at which the tree roots. Restricted to the circllet (i.e., from the viewpoint of the distributed blossom algorithm), this tree has found an augmenting path between two unmatched circllet nodes, terminating with the foremost edge. In the larger crown, this alternating tree is ready to contract into a macrovertex, at which point it can augment with the other unmatched monde vertex, then re-expand.

LEMMA 40. *Every run of the distributed algorithm admits a rooted Lamport ordering.*

PROOF. Fixing any Lamport ordering of a run, we inductively “smooth” it to produce a rooted ordering. Consider the earliest segment within the run which is not rooted, and consider the last grafting operation within that segment which violates the rooted property. Because the root which is enacting this graft is disjoint from all of the operations’ roots from this point to the state update which terminates the segment, this graft operation can be commuted to occur just after the state update. Continuing in this way produces the desired rooted ordering. \square

Next, we show that the *good run* condition is not limiting: any partial good run which still admits progress as a run also admits progress as a good run.

LEMMA 41. *Every partial good run of the serial algorithm operating on $\mathbb{C}G$ can be extended to a longer good run.*

PROOF. We have two objectives:

- (1) We must show that a contraction which includes a monde vertex can be followed by a monde-monde augmentation and an expansion.
- (2) We must show that an augmenting path which does not respect the directedness property can be abbreviated to one that does.

For the first objective, consider such a macrovertex and the cycle used to produce it. The contracted cycle includes an arch vertex, hence the macrovertex inherits all of its edges to the vertices in the monde. Any edge to an unmatched monde vertex can be used to perform the augmentation. The macrovertex can then be directly expanded.

For the second objective, consider an augmenting path which violates directedness. First, we may assume that the augmenting path visits the monde only at its endpoints as follows. By regalness, the augmenting path must begin and end in the monde, so that any intermediate visit to the monde is through a matched vertex. As the augmenting path begins and ends at positive vertices, there must be two adjacent monde vertices whose intervening path begins and ends with an unmatched edge. Select that subpath, then replace its terminating edges with visits to the original terminating monde vertices. Second, we may even assume that the terminating monde vertices are the same, since any arch vertex is reachable by any monde vertex. Together, these observations produce an alternating cycle accessible by a crown-directed tree. \square

With these Lemmas in hand, we turn to the proof of the main correspondence between serial and distributed runs, stated as Theorem 35.

PROOF OF THEOREM 35. Given a rooted Lamport ordering of a run of the distributed blossom algorithm, select a segment for which we will construct a corresponding segment of steps in the serial algorithm. Select an unmatched vertex from the monde to serve as the root of the serial segment. For each root participating in the distributed segment, graft the corresponding circlet-arches matched edge to the chosen root in the monde. Then, each non-augment maneuver in the distributed segment corresponds to an identical maneuver in the serial segment. Finally, an augment in the distributed segment corresponds on the serial side to contraction, augmentation, and expansion. Since we have assumed an interest only in good runs, this correspondence can also be applied in reverse. \square

COROLLARY 42. *Ignoring internal weights, the distributed algorithm can be used to produce the same $\{M_j \setminus G_j\}$ data as in Edmonds's original proof [Edm65a, Section 7].* \square

4.1.2 *Dual correctness.* Unfortunately, the dual step must be treated more manually: the multitude of monde vertices prevent us from fruitfully applying the weighted algorithm to

$\mathbb{C}G$. Fortunately, there is altogether less to show. Instead, let us examine the conditions for performing a reweight in the distributed setting.

Positive-positive: A tree for which no more primal updates are possible has no positive vertices connected by weightless edges to other positive vertices, whether in this tree or in another. Hence, this case is empty. This is as in the dual update step of the serial weighted algorithm.

Negative-any: No edge emanating from a negative vertex in this tree affects the reweight calculation. This is also as in the dual update step of the serial weighted algorithm.

Positive-negative internal: An edge emanating from a positive vertex to a negative vertex in the same tree is not included in the reweight amount calculation. This is also as in the dual update step of the serial weighted algorithm.

Positive-negative foreign: An edge between a positive vertex in this tree and a negative vertex in another constrains the reweight amount. **Since the serial algorithm uses a tree rather than a forest, it has no concept of a "foreign" target and hence has no analogue of this case.**

Our observation is that this last constraint is always a phantom: even when it is weightless, (non-local) progress can always be made. There are two subcases to this final point: according to whether the tree in question belongs to a hold cluster.

Tree belongs to a hold cluster: The multireweighting procedure circumnavigates the weightless hold constraints. In fact, it is guaranteed to reweight by a (nonzero) amount giving a minimal enlargement of the weightless subgraph.

Tree does not belong to a hold cluster: The tree responsible for breaking the hold cluster is necessarily free to act in some other way, whether by a primal update or by a (multi)reweight of its own.

In Edmonds's language, both subcases (hence all cases) afford algorithmic progress.

Since the correctness of the distributed blossom algorithm comes down only to the combinatorial properties of the primal step, together with the ability of the second step to make progress, we have therefore proven the following theorem:

THEOREM 43. *The distributed blossom algorithm is correct: it always terminates, and upon termination it emits a minimum-weight perfect matching.* \square

4.2 Timing

THEOREM 44. *The runtime of the distributed blossom algorithm is $O_p(n^4)$, with n the number of vertices in G .*

PROOF. The serial algorithm does not backtrack, and using a priority mechanism the distributed algorithm *also* does not backtrack. Hence, the runtime of the distributed algorithm is bounded by the duration of the serial run to which it corresponds under Theorem 35. Since the original algorithm runs in time $O_p(n^4)$ and $\mathbb{C}G$ is not asymptotically larger than G , we learn the same for the distributed algorithm. \square

Remark 45. In particular, there is never an asymptotic penalty to using this algorithm over the (original) serial version. Of course, modern implementations of the serial algorithm reduce

the asymptotic runtime substantially compared to Edmonds's original.

5 CLOSING COMMENTS

We point out some of the stones we have left unturned.

Remark 46 (Structured data improvements to runtime). The worst case time complexity bound given in Theorem 44 is likely to be achievable for the algorithm presented here: when working with a fully-connected graph, reweighting operations rooted at high-priority nodes have a particularly bothersome ability to prevent lower-priority reweighting operations from succeeding. However, improvements to the serial blossom algorithm have appeared since its invention in the 1960s [CR99, Kol09, MV80] which more intelligently deploy data structures to track which proposals are worth querying, lowering the time complexity bound. It is open whether these serial blossom variants can be imported to the distributed setting, where the maintenance of delicate large-scale structure has the potential to destroy locality.

Remark 47 (Geometric improvements to runtime). Our intended application rests heavily on the observation that geometric structure in the problem graphs can be leveraged in the dryad to give dramatic improvements in runtime [PK]. A study of this phenomenon is well worth pursuing, not least because of the further application to quantum error correction.

Remark 48 (Lower bounds on runtime). It is also of interest to find classes of problem graphs which cannot be quickly matched, thereby putting lower bounds on the runtime of the distributed algorithm. For instance, it is a well-known result that 2-coloring a line takes o, n^n time in the size of the line [Lin92], which is a special case of a perfect matching: given an enumeration of the vertices in the line (with no regard for their ordering) and a perfect matching, one can produce a 2-coloring by partitioning vertices into those whose matches are higher- or lower-valued than they are. Producing a family of examples of this form, as well as understanding their interactions with Remark 47, would be very valuable.

Remark 49 (Online variants). In our intended application, we implement a dryad that supports, to a limited but extremely useful extent, injection and ejection of nodes from an actively running algorithm. It is of interest to understand how far this can be pushed and what applications this unlocks.

ACKNOWLEDGEMENTS

We would like to thank both Charles Hadfield and Austin Fowler for indirectly suggesting this problem: we would not have thought to look in this direction without Charles, nor would we have known where to begin without Fowler's groundwork. We would also like to thank Zac Cranko for pointing out that it was possible to make a time complexity argument. Lastly, the first author would like to thank Samrita Dhindsa for all manner of support as this work was carried out.

REFERENCES

[aet] aether - Distributed system emulation in Common Lisp. <https://github.com/dtqec/aether>. Accessed: 2021-01-11.

- [Ale49] Sholem Aleichem. *Tevye and his Daughters*. Crown Publishers, 1949.
- [ana] anatevka - A distributed blossom algorithm. <https://github.com/dtqec/anatevka>. Accessed: 2021-01-11.
- [Ash75] E.A. Ashcroft. Proving assertions about parallel programs. *Journal of Computer and System Sciences*, 10(1):110–135, 1975.
- [CCPS09] William J Cook, WH Cunningham, WR Pulleyblank, and A Schrijver. *Combinatorial optimization*, volume 5. Springer, 2009.
- [CDF⁺19] Armando Castañeda, Carole Delporte-Gallet, Hugues Fauconnier, Sergio Rajsbaum, and Michel Raynal. Making local algorithms wait-free: the case of ring coloring. *Theory Comput. Syst.*, 63(2):344–365, 2019.
- [CNAA⁺20] Christopher Chamberland, Kyungjoo Noh, Patricio Arrangoiz-Arriola, Earl T. Campbell, Connor T. Hann, Joseph Iverson, Harold Putterman, Thomas C. Bohdanowicz, Steven T. Flammia, Andrew Keller, Gil Refael, John Preskill, Liang Jiang, Amir H. Safavi-Naeini, Oskar Painter, and Fernando G. S. L. Brandão. Building a fault-tolerant quantum computer using concatenated cat codes, 2020.
- [CR99] William J. Cook and André Rohe. Computing minimum-weight perfect matchings. *INFORMS J. Comput.*, 11(2):138–148, 1999.
- [DFF56] George Bernard Dantzig, Lester R Ford, and Delbert Ray Fulkerson. *A Primal-Dual Algorithm for Linear Programs*. Princeton University Press, 1956.
- [DFFR19] Carole Delporte-Gallet, Hugues Fauconnier, Pierre Fraigniaud, and Mikael Rabie. Distributed computing in the asynchronous LOCAL model. *CoRR*, abs/1904.07664, 2019.
- [DKLP02] Eric Dennis, Alexei Kitaev, Andrew Landahl, and John Preskill. Topological quantum memory. *Journal of Mathematical Physics*, 43(9):4452–4505, 2002.
- [Edm65a] Jack Edmonds. Maximum matching and a polyhedron with 0, 1-vertices. *Journal of research of the National Bureau of Standards B*, 69(125-130):55–56, 1965.
- [Edm65b] Jack Edmonds. Paths, trees, and flowers. *Canadian Journal of mathematics*, 17:449–467, 1965.
- [FMMC12] Austin G. Fowler, Matteo Mariantoni, John M. Martinis, and Andrew N. Cleland. Surface codes: Towards practical large-scale quantum computation. *Phys. Rev. A*, 86:032324, Sep 2012.
- [FWH12] Austin G. Fowler, Adam C. Whiteside, and Lloyd C. L. Hollenberg. Towards practical classical processing for the surface code: Timing analysis. *Phys. Rev. A*, 86:042313, Oct 2012.
- [HBS73] Carl Hewitt, Peter Boehler Bishop, and Richard Steiger. A universal modular ACTOR formalism for artificial intelligence. In Nils J. Nilsson, editor, *Proceedings of the 3rd International Joint Conference on Artificial Intelligence. Stanford, CA, USA, August 20-23, 1973*, pages 235–245. William Kaufmann, 1973.
- [Kol09] Vladimir Kolmogorov. Blossom V: a new implementation of a minimum cost perfect matching algorithm. *Math. Program. Comput.*, 1(1):43–67, 2009.
- [Kuh10] Harold W. Kuhn. The hungarian method for the assignment problem. In Michael Jünger, Thomas M. Liebling, Denis Naddef, George L. Nemhauser, William R. Pulleyblank, Gerhard Reinelt, Giovanni Rinaldi, and Laurence A. Wolsey, editors, *50 Years of Integer Programming 1958-2008 - From the Early Years to the State-of-the-Art*, pages 29–47. Springer, 2010.
- [Lam78] Leslie Lamport. Time, clocks, and the ordering of events in a distributed system. *Commun. ACM*, 21(7):558–565, 1978.
- [Lin92] Nathan Linial. Locality in distributed graph algorithms. *SIAM J. Comput.*, 21(1):193–201, 1992.
- [LPP15] Zvi Lotker, Boaz Patt-Shamir, and Seth Pettie. Improved distributed approximate matching. *J. ACM*, 62(5):38:1–38:17, 2015.
- [LPR09] Zvi Lotker, Boaz Patt-Shamir, and Adi Rosén. Distributed approximate matching. *SIAM J. Comput.*, 39(2):445–460, 2009.
- [MV80] Silvio Micali and Vijay V. Vazirani. An $o(\sqrt{V})$ algorithm for finding maximum matching in general graphs. In *21st Annual Symposium on Foundations of Computer Science, Syracuse, New York, USA, 13-15 October 1980*, pages 17–27. IEEE Computer Society, 1980.
- [PK] Eric C. Peterson and Peter J. Karalekas. Distributed online decoding for large-scale quantum computers. To appear.
- [PK20] Eric C. Peterson and Peter J. Karalekas. aether: Distributed system emulation in common lisp. *CoRR*, abs/2011.06180, 2020.
- [SRFR19] Sgt. Rupert Frere RLC/MOD. Photo of St. Edward's Crown, September 2019. [https://commons.m.wikimedia.org/wiki/File:St_Edward%27s_Crown_-_Royal_Collection_Trust_\(1\).jpg](https://commons.m.wikimedia.org/wiki/File:St_Edward%27s_Crown_-_Royal_Collection_Trust_(1).jpg).
- [WW04] Mirjam Wattenhofer and Roger Wattenhofer. Distributed weighted matching. In Rachid Guerraoui, editor, *Distributed Computing, 18th International Conference, DISC 2004, Amsterdam, The Netherlands, October 4-7, 2004, Proceedings*, volume 3274 of *Lecture Notes in Computer Science*, pages 335–348. Springer, 2004.