

LOCAL LIMIT OF PRIM'S ALGORITHM

BENOÎT CORSINI, ROWEL GÜNDLACH, AND REMCO VAN DER HOFSTAD

ABSTRACT. We study the local evolution of Prim's algorithm on large finite weighted graphs. When performed for n steps, where n is the size of the graph, Prim's algorithm will construct the minimal spanning tree (MST). We assume that our graphs converge locally in probability to some limiting rooted graph. In that case, Aldous and Steele already proved that the local limit of the MST converges to a limiting object, which can be thought of as the MST on the limiting infinite rooted graph.

Our aim is to investigate *how* the local limit of the MST is reached *dynamically*. For this, we take $tn + o(n)$ steps of Prim, for $t \in [0, 1]$, and, under some reasonable assumptions, show how the local structure interpolates between performing Prim's algorithm on the local limit when $t = 0$, to the full local limit of the MST for $t = 1$. Our proof relies on the use of the recently developed theory of *dynamic local convergence*. We further present several examples for which our assumptions, and thus our results, apply.

CONTENTS

1. Introduction	1
2. Background work and preliminary results	9
3. Convergence of Prim's algorithm	18
4. Examples and counter-examples	23
Appendix A. List of symbols and notation	26
References	27

1. INTRODUCTION

1.1. Background and novelty. In this paper we analyse the evolution of the local structure of the subtree discovered by Prim's algorithm [54] as the number of steps in the algorithm varies. This process is one of the most commonly used greedy algorithms to find the minimal spanning tree (MST) of a rooted weighted graphs. However, Prim's algorithm is different from other greedy methods to construct the MST, such as Borůvka's algorithm [52] and Kruskal's algorithm [42], as it discovers the tree *locally* starting from a given root. More precisely, it sequentially adds the neighbouring edge with lowest weight that does not form a cycle with the current component. When running Prim's algorithm to completion, the outcome is known to be the MST of the whole finite graph. In that case, the local limit has already been analysed in [6, Theorem 5.6], where the authors show local convergence towards the minimal spanning forest of the infinite graph.

The *novelty* of our paper is the study of the outcome of Prim's algorithm when only considering a partial number of steps. More precisely, we provide a characterisation of the limit of Prim's algorithm as a function of t when it runs for $tn + o(n)$ steps on a graph with n vertices, for $t \in [0, 1]$. This provides a process on rooted weighted graphs and fits within the recent research field of dynamic local convergence [26, 47]. This process convergence provides

key insights on the evolution of the structure of Prim's algorithm when the number of steps grows.

One of the noteworthy interpretations of our main result (Theorem 1.1), is *how* Prim's algorithm discovers the local neighbourhood of the MST around the root. Indeed, our result states that it first spends a sub-linear time in the local neighbourhood, during which it creates the *invasion percolation cluster*, after which it explores far away parts of the graph, until it reaches a linear number of steps, where it regularly returns to the local neighbourhood and adds edges of the minimum spanning forest of larger and larger weights. Our result further shows that the local structure of the graph does not change as soon as we run Prim's algorithm for $n - o(n)$ steps, allowing for a sub-linear correction. In particular, this means that the last $o(n)$ steps of the algorithm always explore faraway neighbourhoods of the graph.

1.2. Main result. For the rest of this work, we consider a sequence $(G_n)_{n \geq 1}$ of graphs on n vertices, which converges locally in probability towards a rooted graph (G, o) . Local convergence was defined independently by Benjamini and Schramm [12] and Aldous and Steele [6]; see [36, Chapter 2] for an extensive introduction to local convergence. We further endow $(G_n)_{n \geq 1}$ and G with independent $\text{UNIFORM}([0, 1])$ edge weights and sample a random vertex o_n from G_n ; we write $\mathbb{G} = (G, o, w)$ and $\mathbb{G}_n = (G_n, o_n, w_n)$ for the corresponding rooted weighted graphs and call these the *standard extensions* of $(G_n)_{n \geq 1}$ and (G, o) . We are now interested in the local limit of $\mathcal{P}_{k_n}(\mathbb{G}_n)$, defined as the subtree obtained after k_n steps of Prim's algorithm, when $n \rightarrow \infty$ and for arbitrary increasing sequences $(k_n)_{n \geq 1}$.

As discussed above, when $k_n = n$, then $\mathcal{P}_{k_n}(\mathbb{G}_n)$ is exactly the MST of (G_n, w_n) , which converges locally towards the minimum spanning forest $\mathbb{F}^{\mathbb{G}}$ of \mathbb{G} [6, Theorem 5.4] (see Section 2.2 for exact definitions). Using this result, we know that there exist growing radii $(r_n)_{n \geq 1}$ such that $(\mathbb{G}_n, \mathcal{P}_n(\mathbb{G}_n))$ can asymptotically be coupled with $(\mathbb{G}, \mathbb{F}^{\mathbb{G}})$ on balls of radius r_n .

We say that the sequence of functions $(k_n(t))_{n \geq 1, t \in [0, 1]}$ is a *linearly growing sequence* with respect to $(G_n)_{n \geq 1}$ if $k_n(0) = r_n$ and $k_n(t)/n \rightarrow t$ as $n \rightarrow \infty$ for all $t \in [0, 1]$ (see Section 2.5 for a more precise definition of linearly growing sequences). Our main result characterises the behaviour of Prim's algorithm for any arbitrary linearly growing sequence. In particular, we investigate the convergence of the rooted graph obtained by applying Prim $k_n(t)$ steps as a stochastic process on rooted graphs, using the recently developed theory of dynamic local convergence [47].

On the infinite graph \mathbb{G} , we note that the minimum spanning forest $\mathbb{F}^{\mathbb{G}}$ contains the subtree $\mathcal{P}_k(\mathbb{G})$ obtained by running Prim's algorithm on $\mathbb{G} = (G, o, w)$ for any arbitrarily chosen number of steps k . Thus, $\mathbb{F}^{\mathbb{G}}$ also contains the limit $\mathcal{P}_\infty(\mathbb{G})$ of Prim's algorithm, corresponding to the invasion percolation cluster. We now generalise this definition to the *expanded invasion percolation cluster* at level p , which we denote by $\mathbb{F}_+^{\mathbb{G}}(p)$. We define this as the union of the invasion percolation cluster and the edges in the minimal spanning forest $\mathbb{F}^{\mathbb{G}}$ with weight less than p (see Section 2.2 for a more precise definition). This tree is fundamental in characterising the limit of Prim's algorithm and we now provide the two assumptions required for our main theorem to hold.

Given a rooted weighted graph $\mathbb{G} = (G, o, w)$, write $\mathcal{C}_o^{\mathbb{G}}(p)$ for the component containing o when only keeping edges with weight less than p and let $\theta(p) = \theta^{\mathbb{G}}(p)$ be the (annealed) survival probability of o at level p in \mathbb{G} , i.e.,

$$(1) \quad \theta(p) = \mathbb{P}\left(|\mathcal{C}_o^{\mathbb{G}}(p)| = \infty\right).$$

We further define its critical value by $p_c = \inf\{p : \theta(p) > 0\}$ and its (pseudo-)inverse, for $t \in [0, 1)$, by

$$(2) \quad \theta^{-1}(t) = \inf \left\{ p : \theta(p) > t \right\}.$$

For convenience, we define $\theta^{-1}(1) = 1$. When $\theta(p) > 0$, it is possible for the root of \mathbb{G} to be in an infinite component at level p . It is thus natural to wonder what the rest of the graph looks like in the case where $\theta(p) > 0$ but the root does not belong to an infinite component. Our first assumption states that, whenever $\theta(p) > 0$, there exists a component in $\mathbb{F}^{\mathbb{G}}$ which is infinite:

Assumption 1 (Percolation function and infinite percolation components). *The (infinite) rooted and weighted graph \mathbb{G} is said to satisfy the local percolation of a giant assumption if, for any p , the following equivalence holds:*

$$\theta(p) > 0 \iff \mathbb{P}(\exists v \in V(\mathbb{F}^{\mathbb{G}}) : |\mathcal{C}_v^{\mathbb{G}}(p)| = \infty) = 1.$$

Assumption 1 entails that even when the root does not percolate at level p , if p is such that $\theta(p) > 0$, then there are vertices close to the root that percolate. It plays a key role in ensuring that the giant component in the finite graph is locally visible, and can thus be reached by Prim's algorithm after a bounded number of steps.

In Section 4, we discuss why this assumption is required and provide examples of (random) graphs that do not satisfy it. We further conjecture that this assumption can be omitted and the main result would remain true, since we did not find counterexamples to prove otherwise.

Besides the relation between the probability to have an infinite component and the existence of an infinite component in the minimum spanning forest, we need a further assumption on θ . Indeed, since θ plays a key role in the evolution of the limit (see Theorem 1.1), it is also important for this function to be smooth enough, as stated in the following assumption:

Assumption 2 (Smoothness of the percolation function). *The (infinite) rooted weighted graph \mathbb{G} is said to satisfy the smooth percolation assumption if $p \mapsto \theta(p)$, considered as a function on $(p_c, 1]$, is continuous and strictly monotone.*

We emphasize that \mathbb{G} is possibly random, and that $\theta(p)$ is the *annealed* percolation probability, where we take the average w.r.t. the random graph \mathbb{G} and the edge statuses on it. Note that Assumption 2 implies that $t \mapsto \theta^{-1}(t)$ is continuous and strictly increasing as a function from $[0, 1]$ to $[p_c, 1]$, with $\theta^{-1}(0) = p_c$, which, in fact, is the main property that θ needs to satisfy for our proofs.

While we believe Assumption 2 holds in most cases, it is still possible to create counterexamples where θ jumps multiple times (see Section 4). However, these examples tend to also not satisfy Assumption 3 below, and so it is unclear to us whether Assumption 2 is redundant or not. Without proof of the contrary, we include it and keep its removal as an open problem.

The previous two assumptions only pose conditions on the local limit of the graph sequence. We next state an assumption to make sure that the sequence itself behaves as desired. Given the relation between the MST and connected components at level p in the percolated graph, we next state an assumption linking the survival probability and size of the giant, as introduced in [36, Theorem 2.28] for general (unpercolated) graph sequences:

Assumption 3 (Percolation giant is almost local). *The sequence of rooted weighted graphs $(\mathbb{G}_n)_{n \geq 1}$ is said to satisfy the percolation giant is almost local assumption if, for any p ,*

$$\lim_{k \rightarrow \infty} \limsup_{n \rightarrow \infty} \frac{1}{n^2} \mathbb{E} \left[\left| \left\{ u, v \in V(\mathbb{G}_n) : |\mathcal{C}_u^{\mathbb{G}_n}(p)| \geq k, |\mathcal{C}_v^{\mathbb{G}_n}(p)| \geq k, \mathcal{C}_u^{\mathbb{G}_n}(p) \cap \mathcal{C}_v^{\mathbb{G}_n}(p) = \emptyset \right\} \right| \right] = 0.$$

Assumption 3 allows us to relate the size of the largest component in \mathbb{G}_n at level p to $\theta^{\mathbb{G}}(p)$ and we invite the reader to take a look at Propositions 2.5 and 2.6 to gain a first idea of the importance of this assumption. This assumption plays an essential role in our main theorem, as it is a key property required to connect the behaviour of percolated components between finite and infinite graphs

With these assumptions in hand, we now have all the ingredients to state our main theorem. The following result relies on the *local process convergence*, which corresponds to the natural extension of the local topology on graphs to the process topology (see Section 2.6 for more details):

Theorem 1.1 (Local limit of Prim's algorithm). *Let $(G_n)_{n \geq 1}$ be converging locally in probability to (G, o) . Let $(\mathbb{G}_n)_{n \geq 1}$ and \mathbb{G} be the respective standard expansions of $(G_n)_{n \geq 1}$ and (G, o) . Define θ as in (1) and let $k = (k_n(\cdot))_{n \geq 1}$ be a linearly growing sequence with respect to $(G_n)_{n \geq 1}$ as defined in Section 2.5. If \mathbb{G} satisfies Assumption 1, $\theta(p)$ satisfies Assumption 2, and $(\mathbb{G}_n)_{n \geq 1}$ satisfies Assumption 3, then*

$$\left(\mathcal{P}_{k_n(t)}(\mathbb{G}_n) \right)_{t \in [0,1]} \xrightarrow{\text{lpc}} \left(\mathbb{F}_+^{\mathbb{G}}(\theta^{-1}(t)) \right)_{t \in [0,1]},$$

where the convergence occurs according to the local process convergence defined in Section 2.6.

The proof of Theorem 1.1 can be found in Section 3.3. It can be reduced to three main steps. First, we prove the one-dimensional convergence of Prim's algorithm towards the correct limit. This proof boils down to the fact that the largest component in $\mathbb{G}_n(p)$ has size approximately $n\theta(p)$, and so Prim's algorithm applied $n\theta(p)$ steps will first quickly reach the largest component and then remain in that component, explaining the structure of $\mathbb{F}_+^{\mathbb{G}}(p)$. Second, we prove that the previous one-dimensional convergence can be extended to the finite-dimensional distributions, and show that the limit satisfies the assumptions of the local process convergence as defined in Section 2.6. Finally, we prove the tightness of the process (as defined in Section 2.6), and complete the proof of Theorem 1.1, in Section 3.3.

It is worth mentioning that the previous convergence not only occurs as a process convergence but also jointly with the convergence of \mathbb{G}_n towards \mathbb{G} . This is a direct consequence of the proof method from Section 3, where we start by coupling \mathbb{G}_n and the minimal spanning tree on it, to \mathbb{G} and the minimal spanning forest on it, and show the convergence of Theorem 1.1 under this coupling. We decided not to state it as a joint convergence here for clarity.

Organisation. This paper is organised as follows. In Section 2 we give some background on the different definitions from the introduction and state some useful preliminary results. In Section 3 we then apply the previous results in order to prove Theorem 1.1. Finally, in Section 4, we consider a few interesting specific cases and examples. We also invite the reader to take a look at Appendix A for a list of often-used symbols and notations. We now provide a more detailed review of the literature relevant to this work.

1.3. Related literature. The introduction of local convergence [6, 12] at the start of this millennium has formalised the connection between analysing graphs and graph processes on different scales. The power of this technique is demonstrated in our results for Prim's algorithm, but has also been demonstrated for many different processes in the literature. We present an overview of how the introduction of local convergence gave rise to a formal local description of graphs and how this was used to locally analyse global graph properties and graph processes. Next, we emphasise the importance of the MST and present the state-of-the-art results for the MST on random graphs, both from a local and global perspective.

Local convergence of graphs. The early applications of local convergence show locally tree-like behaviour for popular sparse random graphs. Key examples include the Erdős-Rényi graph [22], the uniform random graph and the configuration model [22], inhomogeneous random graph [17] and [36, Chapter 3], and the preferential attachment model [13, 30]. Recent years have seen an increasing interest in this line of work and have extended the results to more complex graphs such as uniform planar graphs [58], spatial inhomogeneous random graphs [39], Gibbs random graphs [28], and random intersection graphs [43]. For more extensive overviews, we refer to [10, 36]. While the field of local convergence for random graphs developed, others investigated how to optimally make use of these novel results. Let us mention two main applications of local convergence results: First, local convergence proved itself to be useful in analysing global graph properties, such as the existence, size and uniqueness of the giant. Second, processes on random graphs can often be analysed from a local perspective. Next, we give a short overview of both applications.

Local analysis of global objects. Many sparse random graphs are characterised by their local limit to such a strong degree, that their structure can be used to analyse global graph properties. [39, Chapter 2.5] illustrates this in detail for different graphs measures, such as the global cluster coefficient and assortativity coefficient for graphs that have a uniformly integrable second and third moment for the degree of a uniformly chosen vertex respectively. Related is the PageRank distribution [31] that is also well studied from a local perspective.

Such results also extend to an analysis of the number of spanning trees for large sparse graphs [45, 55], to the spectral measure for locally tree-like graphs [14, 18] and to graph colourings [20]. A similar analysis for the critical percolation probability for transitive graphs has seen much attention over the years. First conjectured that this was indeed a local property (see [11] for details), many works have since been extending the class of graphs for which the critical percolation probability was indeed local, see [21] and the references therein. The conjecture was finally proven in full generality in [27]. Finally, recent work shows that, under some conditions on the graph, the giant component can be completely characterised from a local perspective [35, 36, 37]. In this case, the giant was also described as 'almost local'.

We finally remark that these results find direct practical applications in the mathematical physics literature via, for example, the factor model [23] and the Ising model. The latter model has received a significant amount of attention. Here one assigns labels (or spins) $\{-1, 1\}$ to a vertices of a random graph according to the Boltzmann distribution, which forms a central representation for ferromagnetic fields. The works of [22, 23, 49] show that for a multitude of locally tree-like graphs the spin distributions can be analysed from the local limit. These works have later been extended to different graphs and to different properties of the Ising model from both a theoretical and practical point of view, see [4, 8, 23, 25, 59].

Processes on a random graph. A second key development initiated by local convergence is the translation of processes on graphs to their local limit. Ideally one hopes that analysing

a process on a random graph is similar to analysing the process on the local limit. We first observe such a phenomenon on dynamic random graphs, for example the Erdős-Rényi and inhomogeneous random graph [26] and the random intersection graph [47]. Interacting particle systems, where states of vertices develop via a Markov or diffusion process based on their neighbours, have been extensively analysed, see [44] and the references therein, where the authors show that the vertex interaction translates to the local limit. Related examples have also been studied, such as the contact process [50], SIR epidemic models [48], and voter models [40].

Algorithms for the minimum spanning tree. There are many classical construction algorithms for the MST dating back to the early 20th century, such as Borůvka’s algorithm [52], Kruskal’s algorithm [42] and Prim’s algorithm [54], which generally run in $O(E \log(V))$ time complexity [53]. From a discrete optimisation point of view, the main interest in the MST is to improve the efficiency of search algorithms, see for example [41] and the reference therein.

For this work, the geometry of the MST in disordered networks is of significant importance. To the best of our knowledge, interest in the topic arose in the mid-eighties for properties of MST on the complete graph with random edge weights [29] and for the Euclidean minimum spanning tree [5, 56] and the references therein. This was later further generalised in [7]. Shortly after the introduction of local weak convergence in the beginning of the new millennium, Aldous and Steele [57] examined the MST on the weighted complete graph, and related that to a minimum spanning forest on the Poisson weighted infinite tree (PWIT). This work was later extended to general graphs in [6].

The geometry and other properties of the MST find many applications in applied physics, see for example [24] and the references therein. A classic example of such geometrical property is the diameter of the MST, often used as an upper bound for the typical distance between uniform vertices. These results hold for many classes of disordered networks [19], due to the universality of the MST [24].

Local limits for the minimum spanning tree. So far, local convergence of the MST (i.e., $k_n = n$) has been shown for general graphs, see [6, Theorem 5.4] and more in depth for the complete graph with random edge weights [1, 51].

Global limits for the minimum spanning tree. The global perspective of the MST is usually captured by its *scaling limit*. This was first derived for the complete graph in [2], and is closely related to the Erdős-Rényi random graph. For other random graph models, the global geometry of the MST has also been analysed before, such as the 3-regular graph [3] and the inhomogeneous random graph with power-law degrees [15]. Interestingly, while the scaling limit of these MSTs are intimately related to large critical percolation cluster on them, the Hausdorff dimension of the MST is generally larger than that of critical percolation clusters. Indeed, for the complete graph, the Hausdorff dimension of critical percolation clusters equals 2 (as for the closely related continuum random tree), while for the scaling limit of the MST it is 3 [2].

Quenched versus annealed. As one can consider this model as random disorder (the uniform edge weights) in a random environment (the random graph), we emphasise that our results are for the *annealed setting*, that is, where one averages out over the random environment. This is in contrast with the *quenched setting*, where one shows results for an almost sure realisation of the environment. This distinction is important as it is known that results can differ between regimes, see for example [9] and the references therein.

1.4. Discussion of the results. In the following, we discuss the implications of Theorem 1.1. We first illustrate how this result completely describes how Prim's algorithm discovers a random graph from a local perspective. We provide simulations of Prim's algorithm on large graphs and show how it follows patterns found in our theoretical results. Finally, we show that our results also provide a good description of the addition times of edges.

Discovering the MST. Analysing the local limit of the MST on (random) graphs via Prim's algorithm involves two limits: one relating to the graph size (i.e., n) and the other relating to number of steps of Prim's algorithm (i.e., k). Except for the clear condition $n \geq k$, there is some freedom in how these limits are taken. Theorem 1.1 illustrates the behaviour for different choices of the relation between n and k . First, by taking $k = o(n)$, we locally obtain the invasion percolation cluster of the local limit $\mathbb{F}_+^{\mathbb{G}}(p_c) = \mathcal{P}_\infty(\mathbb{G})$. This agrees with simply taking these limits *sequentially*, i.e., taking $n \rightarrow \infty$ followed by $k \rightarrow \infty$.

While clearly the local neighbourhood of the MST is not fully explored in this setting (consider the example of a single vertex attached to the root with a very high weight), this directly implies that k growing (arbitrary slowly) in n is not sufficient to discover the full local neighbourhood of the MST (at least when $p_c < 1$). Indeed, for some small $\varepsilon > 0$, at time εt , by Proposition 3.1, Prim's algorithm is discovering the largest percolation component of the graph percolated at $\theta^{-1}(\varepsilon) > p_c$, which is located outside of the local neighbourhood. By Proposition 2.6, the largest component is of linear size when percolated at $p > p_c$, we find that only after a linear number of steps, Prim's algorithm is able to *return* to the local neighbourhood. Indeed, our results state that, in the time interval $[cn, (c + \delta)n]$, Prim's extends its search in the local neighbourhood by connecting edges with weights in between $\theta^{-1}(c)$ and $\theta^{-1}(c + \delta)$.

The local limit of the MST is therefore constructed in two *phases*. In the *first phase*, consisting of $o(n)$ steps, Prim's algorithm basically explores the invasion percolation cluster on the local limit. In the *second phase*, after spending εn steps away from the root, it returns to the neighbourhood of the root, and adds edges and groups of vertices that belong to the MST. This was also noted in [1] for the complete graph.

Simulations. Our results imply that there are vertices in the local neighbourhood that are explored very early by Prim's algorithm (first phase) and some that are explored very late (second phase). Such behaviour implies that it takes around $n(1 - o(1))$ steps for Prim's algorithm to fully discover the local limit of the MST, as we confirm in a simulation study.

In Figures 1, 2, and 3 we ran Prim's algorithm on large grids, triangular lattices, and 3-regular graphs respectively. While all figures show about 1000 vertices of the graph, the simulations were run on a larger graph extending past the boundary of the images, in order to better reproduce the effect of a *local limit*. In all cases, we can see that some edges close to the root are added late in the algorithm (i.e., they are red). Finally, in Figure 4, we provide a representation on a much larger scale with a 2000×2000 grid; in this case, we once again see that late components, in red, are added in every neighbourhood around the root.

Addition and completion times. We next use our results to describe the *addition times* for edges in the neighbourhood of the root, and the *completion time* of the neighbourhood of the root in the MST. For $r \geq 0$, we define the *addition time*

$$\tau_n(r, m) = \min \left\{ k : |\mathcal{P}_k(\mathbb{G}_n) \cap B_r(\mathbb{G}_n)| = m \right\}$$




FIGURE 1. Simulation of Prim's algorithm on a large grid, starting from the root (in turquoise, in the middle). Vertices and edges are coloured based on the step at which Prim's algorithm adds them to the current tree (from green to yellow, and finally red) and the three images correspond to running Prim's algorithm for $n/3$, $2n/3$, and n steps.




FIGURE 2. Simulation of Prim's algorithm on a large triangular lattice, starting from the root (in turquoise, in the middle). Vertices and edges are coloured based on the step at which Prim's algorithm adds them to the current tree (from green to yellow, and finally red) and the three images correspond to running Prim's algorithm for $n/3$, $2n/3$, and n steps.



FIGURE 3. Simulation of Prim's algorithm on a large 3-regular graph, starting from the root (in turquoise, in the middle). Vertices and edges are coloured based on the step at which Prim's algorithm adds them to the current tree (from green to yellow, and finally red) and the three images correspond to running Prim's algorithm for $n/3$, $2n/3$, and n steps.

to be the time that the m^{th} vertex in $B_r(\mathbb{G}_n)$ is added. Additionally, for fixed $r \geq 0$, we define the *completion time* $C_n(r)$ by

$$C_n(r) = \min \left\{ k : B_r(\mathcal{P}_k(\mathbb{G}_n)) = B_r(\mathcal{P}_n(\mathbb{G}_n)) \right\}$$




FIGURE 4. Simulation of Prim's algorithm on a 2000×2000 grid, starting from the root (in turquoise, in the middle). Each pixel corresponds to a single vertex and vertices and edges are coloured based on the step at which Prim's algorithm adds them to the current tree (from green to yellow, and finally red, with extra shaded fluctuations to highlight the local behaviour).

i.e., $C_n(r)$ denotes the first time that Prim's algorithm has discovered the complete r -neighbourhood of the MST. Our next theorem concerns the convergence in distribution of the addition and completion times:

Theorem 1.2 (Limits of addition and completion times). *Under the assumptions of Theorem 1.1, as $n \rightarrow \infty$, the process*

$$\left(\frac{\tau_n(r, m)}{n}, \frac{C_n(r)}{n} \right)_{m, r \geq 0}$$

converges in distribution to the process of addition and completion times on the local limit \mathbb{G} .

Note that $\tau_n(r, m)/n$ converges to zero *precisely* for the edges that are reached by the invasion percolation cluster on the local limit. Both limits can be made more explicit using only the structure of the limit. In particular, one can check that

$$\frac{C_n(r)}{n} \xrightarrow{n \rightarrow \infty} \max \left\{ \theta(w(e)) : e \in E(B_r(\mathbb{F}^{\mathbb{G}})) \setminus E(\mathcal{P}_{\infty}(\mathbb{G})) \right\},$$

since any edge in $\mathcal{P}_{\infty}(\mathbb{G})$ is explored in $o(n)$ steps and all edges connected to $\mathcal{P}_{\infty}(\mathbb{G})$ via a path of weight less than p are explored in the first $\theta(p)n + o(n)$ steps of Prim's algorithm.

2. BACKGROUND WORK AND PRELIMINARY RESULTS

In this section we provide further details on the definitions from Section 1.2. We start with background notations and definitions regarding graphs in Section 2.1. We then define the minimum spanning forest, the minimum spanning tree, the invasion percolation cluster, and the expanded invasion percolation cluster in Section 2.2. In Section 2.3, we study properties related to the levels in an infinite rooted weighted graph, more precisely, the time it takes for Prim's algorithm to reach an infinite component at level p , and the regularity of the survival probability. In Section 2.4, we introduce properties of the local convergence and discuss its

consequences in Section 2.5. Finally, in Section 2.6 we extend the local convergence to the local process convergence from Theorem 1.1.

2.1. Graph notations. A *locally finite graph* (or simply *graph*) $G = (V(G), E(G))$ is a pair of sets, where $V(G)$ is the arbitrary set of *vertices*, $E(G) \subseteq \binom{V(G)}{2} = \{\{u, v\} : u, v \in V(G)\}$ is the set of *edges*, composed of (unordered) pairs of vertices, and for any $v \in V(G)$, the set of neighbors $\{u \in V(G) : \{u, v\} \in E(G)\}$ is finite.

For any graph $G = (V(G), E(G))$, a pair (G, o) with $o \in V(G)$ is called a *rooted graph*, a pair (G, w) with $w : E(G) \rightarrow \mathbb{R}$ injective (or $w : \mathcal{E} \rightarrow \mathbb{R}$ with $E(G) \subseteq \mathcal{E}$) is called a *weighted graph*, and the subsequent triplet (G, o, w) is called a *rooted weighted graph*. We further consider (G, o, w) , $((G, o), w)$, and $((G, w), o)$ to be equivalent, allowing us to directly go from rooted or weighted graphs to rooted weighted graphs. In these settings, o is referred to as the *root* and w as the *weights*. We often denote weighted graphs with a curly capital letter \mathcal{G} and rooted weighted graphs with a double bar capital letter \mathbb{G} .

We let the size of G be its number of vertices $|G| = |V(G)|$, and say that G is finite when $|G| < \infty$. For any two graphs G and G' , we define their union and intersection by $G \otimes G' = (V(G) \otimes V(G'), E(G) \otimes E(G'))$ where $\otimes \in \{\cup, \cap\}$. We then say that G and G' intersect when $V(G \cap G') \neq \emptyset$ (note that it might have zero edges). Finally, we extend these notations by allowing one of the two graphs to be rooted, weighted, or rooted weighted, and by letting the union or the intersection be itself rooted, weighted, or rooted weighted with the same parameters. Observe that the union requires the weights to be defined on both edge sets and that the intersection requires the root to be part of both vertex sets, which is always the case here.

Given a (possibly random) finite graph G , the *standard extension* of G is the random rooted weighted graph $\mathbb{G} = (G, o, w)$ where o is chosen uniformly at random among $V(G)$ and $(w(e))_{e \in E(G)}$ is a family of independent $\text{UNIFORM}([0, 1])$. Similarly, the standard extension of a rooted (not necessarily finite) graph (G, o) is the random rooted weighted graph $\mathbb{G} = (G, o, w)$ where $(w(e))_{e \in E(G)}$ is a family of independent $\text{UNIFORM}([0, 1])$ random variables. In these definitions, the weights in w are always independent of (G, o) .

For any weighted graph $\mathcal{G} = (G, w)$, we denote the graph obtained by only keeping edges of weight less than p by $\mathcal{G}(p)$, that is, $\mathcal{G}(p) = (G_p, w)$ with $V(G_p) = V(G)$ and $E(G_p) = \{e \in E(G) : w(e) \leq p\}$; we similarly define $\mathbb{G}(p) = ((G, w)(p), o)$ when $\mathbb{G} = (G, o, w)$ is a rooted weighted graph. Observe that, if G is a graph and \mathbb{G} is its standard extension, then G_p as previously defined has the distribution of G percolated at level p , however we also implicitly couple all such $(G_p)_{p \in [0, 1]}$ here so that $p \mapsto G_p$ is càdlàg.

For any weighted graph $\mathcal{G} = (G, w)$ and any $v \in V(\mathcal{G})$, we denote the connected component of v in $\mathcal{G}(p)$ by $\mathcal{C}_v^{\mathcal{G}}(p)$. Moreover, if \mathcal{G} is finite, we denote the connected components of $\mathcal{G}(p)$ in decreasing order of size by $\mathcal{C}_{(1)}^{\mathcal{G}}(p), \mathcal{C}_{(2)}^{\mathcal{G}}(p), \dots$ (breaking ties arbitrarily). We also extend these notations to rooted weighted graphs. In both cases (weighted or rooted weighted), the components are considered as standard graphs and we drop the weights (and the root) when considering these components. Observe that, if G is a graph and \mathbb{G} is its standard extension, then $\mathcal{C}_v^{\mathbb{G}}(0) = (\{v\}, \emptyset)$ and $\mathcal{C}_v^{\mathbb{G}}(1) = \mathcal{C}_{(1)}^{\mathbb{G}}(1) = G$ for any $v \in V(G)$.

2.2. Minimum spanning forest and expanded invasion percolation cluster. Let $\mathcal{G} = (G, w)$ be a weighted graph. We call the graph with vertex set $V(\mathbb{F}^{\mathcal{G}}) = V(G)$ and edge set defined by $e = \{u, v\} \in E(\mathbb{F}^{\mathcal{G}})$ if and only if the two components $\mathcal{C}_u^{\mathcal{G} \setminus \{e\}}(w(e))$ and $\mathcal{C}_v^{\mathcal{G} \setminus \{e\}}(w(e))$ are disjoint and not both infinite (where $\mathcal{G} \setminus \{e\}$ denotes the graph \mathcal{G} where the

edge e is removed) the *minimum spanning forest* of \mathcal{G} , and denote it $F^{\mathcal{G}}$. For a rooted weighted graph $\mathbb{G} = (G, o, w)$, we let $F^{\mathbb{G}} = F^{(G, w)}$ be the minimum spanning forest of \mathbb{G} and further define $\mathbb{F}^{\mathbb{G}} = (F^{\mathbb{G}}, o, w)$ to be the rooted weighted minimum spanning forest of \mathbb{G} . Since this definition may yield a disconnected forest when G is infinite or disconnected, and we are only interested in its local limit, we actually reduce $\mathbb{F}^{\mathbb{G}}$ to the unique tree containing the root. We still refer to it as a *forest*, although it is actually a tree.

In the case where \mathbb{G} (or \mathcal{G}) is finite and connected, the forest $F^{\mathbb{G}}$ is always connected and usually referred to as the *minimum spanning tree* of \mathbb{G} , which we denote $T^{\mathbb{G}}$ here (or $\mathbb{T}^{\mathbb{G}}$ for its rooted weighted counterpart). There exist several efficient algorithms to compute the minimum spanning tree and we now present one of them, Prim's algorithm [54], while not providing the proof that it indeed outputs the minimum spanning tree.

Given a connected rooted weighted graph $\mathbb{G} = (G, o, w)$, we recursively define a sequence of trees as follows. First, we start from $T_1 = (\{o\}, \emptyset)$. Then, given $T_{k-1} = (V_{k-1}, E_{k-1})$ with $k \leq |\mathbb{G}|$, let $\{u_k, v_k\}$ be the edge of minimal weight among all edges with exactly one end in T_{k-1} . Assume without loss of generality that $u_k \in V_{k-1}$ and $v_k \notin V_{k-1}$. Then, we let $T_k = (V_{k-1} \cup \{v_k\}, E_{k-1} \cup \{u_k, v_k\})$ be the tree obtained by adding this edge to T_{k-1} . This defines a sequence of growing trees such that $T_{|\mathbb{G}|}$ is the minimum spanning tree of \mathbb{G} when the graph is finite. If that is the case, we also let $T_k = T_{|\mathbb{G}|}$ when $k \geq |\mathbb{G}|$. For any $k \geq 1$, we now let $\mathcal{P}_k(\mathbb{G})$ be the rooted weighted outcome of this process after k steps, that is

$$\mathcal{P}_k(\mathbb{G}) = (T_k, o, w).$$

We further extend this definition to $k = \infty$ by letting

$$T_{\infty} = \left(\bigcup_{k \geq 0} V_k, \bigcup_{k \geq 0} E_k \right)$$

and setting $\mathcal{P}_{\infty}(\mathbb{G}) = (T_{\infty}, o, w)$ as before. We observe that $\mathcal{P}_{\infty}(\mathbb{G})$ is always a subtree of $\mathbb{F}^{\mathbb{G}}$.

When \mathbb{G} is infinite, $\mathcal{P}_{\infty}(\mathbb{G})$ is actually a commonly studied object, usually referred to as the *invasion percolation cluster*. For this reason, we now define the *expanded invasion percolation cluster*, written as $\mathbb{F}_{+}^{\mathbb{G}}(p)$, to be the union of the invasion percolation cluster and the edges in the minimal spanning forest of weight at most p , denoted by $E(\mathbb{F}^{\mathbb{G}}(p))$, obtained by only keeping edges in $\mathbb{F}^{\mathbb{G}}$ of weight less than p . More precisely, $\mathbb{F}_{+}^{\mathbb{G}}(p)$ is the rooted weighted graph which has the same vertex set, root, and weights as \mathbb{G} , and with edge set defined by

$$E(\mathbb{F}_{+}^{\mathbb{G}}(p)) = E(\mathcal{P}_{\infty}(\mathbb{G})) \cup E(\mathbb{F}^{\mathbb{G}}(p)).$$

We again observe that the expanded invasion percolation cluster is likely not connected, and thus only consider the component containing the root, making $\mathbb{F}_{+}^{\mathbb{G}}(p)$ a tree again.

2.3. Components in an infinite graph. In this section, we consider an infinite rooted weighted graph $\mathbb{G} = (G, o, w)$. Following the definition of $\theta = \theta^{\mathbb{G}}$ in (1), we are interested in the time it takes for Prim's algorithm to reach an infinite component from the root. More precisely, for any $p \in [0, 1]$, we let $\mathcal{K}(p) = \mathcal{K}^{\mathbb{G}}(p)$ be the random variable

$$(3) \quad \mathcal{K}(p) = \min \left\{ k : \exists v \in V(\mathcal{P}_k(\mathbb{G})), |\mathcal{C}_v(p)| = \infty \right\} - 1.$$

From this definition, we observe that $\mathcal{K}(p) = \infty$ when Prim's algorithm *never* explores an infinite component at percolation level p , and $\mathcal{K}(p) = 0$ is equivalent to the root o already belonging to an infinite component at percolation level p . In particular, it follows that $p \mapsto \mathcal{K}(p)$ is decreasing (under the coupling of $\mathbb{G}(p)$ using its edge weights) with $\mathcal{K}(0) = \infty$ and

$\mathcal{K}(1)$ equals 0 if \mathbb{G} is infinite, ∞ otherwise. Finally, when $\mathcal{K}(p) < \infty$, we let $v(p) = v^{\mathbb{G}}(p)$ be the $(\mathcal{K}(p) + 1)$ -th vertex explored during Prim's algorithm and we note that this implies that $|\mathcal{C}_{v(p)}^{\mathbb{G}}(p)| = \infty$; in the case $\mathcal{K}(p) = \infty$, for convenience, we let $v(p) = \infty$ be a “boundary” vertex and define $\mathcal{C}_{v(p)}^{\mathbb{G}}(p) = (\emptyset, \emptyset)$ to be the empty graph.

In order to see the importance of $\mathcal{K}(p)$, we start with a result showing the relation between its value and Assumption 1:

Proposition 2.1 (Infinite components reached by Prim). *Let \mathbb{G} be a rooted weighted graph. Then, for any $p \in [0, 1]$,*

$$\mathcal{K}(p) = \infty \iff \forall v \in \mathbb{F}^{\mathbb{G}} : |\mathcal{C}_v^{\mathbb{G}}(p)| < \infty.$$

In words, Prim's algorithm never reaches an infinite component at level p if and only if all components at level p reachable from o in the minimum spanning forest have finite sizes.

Assumption 1 states that the event on the right-hand side has probability zero when $\theta(p) > 0$. Thus, in this case, we obtain that $\mathcal{K}(p) < \infty$ a.s.

Proof. Note first that $\mathcal{C}_v^{\mathbb{G}}(p) \cap \mathbb{F}^{\mathbb{G}} = \mathbb{F}^{\mathbb{G}}(p)$ so that $|\mathcal{C}_v^{\mathbb{G}}(p)| = |\mathbb{F}^{\mathbb{G}}(p)|$ (if they are both infinite, they are not necessarily equal as $\mathbb{F}^{\mathbb{G}}$ is only the connected component of the root). We now observe that the vertices and edges found by Prim's algorithm starting from o in \mathbb{G} are all in $\mathbb{F}^{\mathbb{G}}$.

From this observation, the left implication directly follows, since there is no infinite component for Prim's algorithm to encounter. For the converse implication, assume that $\mathbb{F}^{\mathbb{G}}$ has infinite components at level p and let \mathbb{F}^{∞} be their union. We now show that Prim's algorithm eventually reaches \mathbb{F}^{∞} .

If $o \in \mathbb{F}^{\infty}$, then it is trivially true, so we now assume that $o \notin \mathbb{F}^{\infty}$. Let $o = v_0, \dots, v_k$ be a shortest path from o to \mathbb{F}^{∞} and let $e = \{v_i, v_{i+1}\}$ be the edge of maximal weight in this path. Since we know that $o \notin \mathbb{F}^{\infty}$ and e is the edge of maximal weight, it follows that $w(e) > p$. But then, in that case $\mathcal{C}_{v_{i+1}}^{\mathbb{G}}(w(e))$ contains the infinite component at level p included in \mathbb{F}^{∞} , so it is infinite. By the definition of the minimum spanning forest from Section 2.1, this implies that $|\mathcal{C}_{v_i}^{\mathbb{G}}(w(e))| = |\mathcal{C}_o^{\mathbb{G}}(w(e))| < \infty$, and so Prim's algorithm eventually visits the edge e . Using this argument inductively, we see that Prim's algorithm eventually reaches \mathbb{F}^{∞} , as desired. \square

We now use $\mathcal{K}(p)$ to characterise the expanded invasion percolation cluster, and highlight its relation to the invasion percolation cluster (IPC) and Prim's algorithm:

Lemma 2.2 (Decomposition of expanded IPC). *Let \mathbb{G} be a rooted weighted graph, $\mathbb{F}^{\mathbb{G}}$ its rooted weighted minimum spanning forest, and $\mathbb{F}_+^{\mathbb{G}}(p)$ its expanded invasion percolation cluster as defined in Section 2.2. Further define $\mathcal{K}(p) = \mathcal{K}^{\mathbb{G}}(p)$ as in (3) and $v(p) = v^{\mathbb{G}}(p)$ as the $(\mathcal{K}(p) + 1)$ -th vertex explored in Prim's algorithm. Then, for any $p \in [0, 1]$,*

$$\mathbb{F}_+^{\mathbb{G}}(p) = \mathcal{P}_{\mathcal{K}(p)+1}(\mathbb{G}) \cup \left(\mathbb{F}^{\mathbb{G}} \cap \mathcal{C}_{v(p)}^{\mathbb{G}}(p) \right),$$

where, when $\mathcal{K}(p) = \infty$, this formula equals

$$\mathbb{F}_+^{\mathbb{G}}(p) = \mathcal{P}_{\infty}(\mathbb{G}).$$

Proof. Consider first the case $\mathcal{K}(p) = \infty$. We already know from its definition in Section 2.2 that $\mathcal{P}_{\infty}(\mathbb{G})$ is a subtree of $\mathbb{F}_+^{\mathbb{G}}(p)$. For the other inclusion, let $e = \{u, v\}$ be an edge on the boundary of $\mathcal{P}_{\infty}(\mathbb{G})$ in $\mathbb{F}^{\mathbb{G}}$, so that u is explored by Prim's algorithm, while v is not. Using the

definition of Prim's algorithm, we observe that $\mathcal{C}_u^{\mathbb{G} \setminus \{e\}}(w(e))$ contains an infinite component of $\mathcal{P}_\infty(\mathbb{G})$ and so is infinite. Using the result from Proposition 2.1, we see that this implies that $w(e) > p$ and so $e \notin E(\mathbb{F}_+(p))$ as desired.

Next, consider the case where $\mathcal{K}(p) < \infty$. The left inclusion once again directly follows from the definition of $\mathbb{F}_+^{\mathbb{G}}(p)$, so we focus on the right inclusion. First of all, since Prim's algorithm reaches the component $\mathcal{C}_{v(p)}^{\mathbb{G}}(p)$ after $\mathcal{K}(p) + 1$ steps and this component is infinite, it remains within it for the rest of the time. Using that Prim's algorithm is a subtree of the minimum spanning forest, we see that

$$\mathcal{P}_\infty(\mathbb{G}) \subseteq \mathcal{P}_{\mathcal{K}(p)+1}(\mathbb{G}) \cup \left(\mathbb{F}^{\mathbb{G}} \cup \mathcal{C}_{v(p)}^{\mathbb{G}}(p) \right).$$

Write T for the tree on the right-hand side of the previous equation (i.e., the target object from the lemma) and consider an edge $e = \{u, v\}$ from $\mathbb{F}^{\mathbb{G}}$ such that u is in T and v is not. Then either $u \in \mathcal{C}_{v(p)}^{\mathbb{G}}(p)$ and since v does not belong to T , it follows that $w(e) > p$ so e does not belong to $\mathbb{F}_+^{\mathbb{G}}(p)$. Otherwise, $u \notin \mathcal{C}_{v(p)}^{\mathbb{G}}(p)$ and so $\mathcal{K}(p) \geq 1$. In that case, there is an edge between o and $\mathcal{C}_{v(p)}^{\mathbb{G}}(p)$ in $\mathbb{F}^{\mathbb{G}}$ whose weight is larger than p . But then, the edge e has to be heavier than p , otherwise it would have been incorporated in Prim's algorithm. In both cases, we observe that $w(e) > p$ and so e does not belong to $\mathbb{F}_+^{\mathbb{G}}(p)$, proving the second inclusion. \square

2.4. Local limits. For any $r \geq 0$ and any rooted graph (G, o) , the ball of radius r of (G, o) denoted $B_r(G, o)$ is the rooted subgraph of G composed of all vertices that can be reached from o using only r edges; for example, $B_0(G, o) = (\{o\}, o)$ and $B_1(G, o)$ is defined by $V(B_1(G, o)) = \{o\} \cup N$, where N is the set of neighbours of o in G , and

$$E(B_1(G, o)) = \left\{ \{o, v\} : v \in N \right\} \cup \left(E(G) \cap \binom{N}{2} \right).$$

For a rooted weighted graph $\mathbb{G} = (G, o, w)$, we extend the notations and denote $B_r(\mathbb{G}) = (B_r(G, o), w)$ for the rooted weighted subgraph of \mathbb{G} composed of vertices at distance r from o .

Two graphs G and G' are said to be *isomorphic* if there exists a bijection $\varphi: V(G) \rightarrow V(G')$ such that $e = \{u, v\} \in E(G)$ if and only if $\varphi(e) = \{\varphi(u), \varphi(v)\} \in E(G')$. Two rooted graphs (G, o) and (G', o') are said to be isomorphic if G and G' are isomorphic and the previous bijection φ can be chosen to satisfy $\varphi(o) = o'$. Finally, while this definition could directly be extended to rooted weighted graphs, since the weights considered here are continuous, we prefer using an approximation, further explained below.

Consider two rooted weighted graphs $\mathbb{G} = (G, o, w)$ and $\mathbb{G}' = (G', o', w')$. Then, φ is said to be a ε -isomorphism between \mathbb{G} and \mathbb{G}' if φ satisfies the previous isomorphic condition between (G, o) and (G', o') and

$$|w(e) - w'(\varphi(e))| \leq \varepsilon$$

for any $e \in E(G)$. We then write $\mathbb{G} \equiv_r \mathbb{G}'$ if there exists a $1/r$ -isomorphism between $B_r(\mathbb{G})$ and $B_r(\mathbb{G}')$. If that is the case, this means that the unweighted versions of \mathbb{G} and \mathbb{G}' are isomorphic on a ball of radius r and that such isomorphism can be chosen so that edges are mapped to edges with weights that are at most $1/r$ apart. We also naturally extend this notation to unweighted graphs by saying that $(G, o) \equiv_r (G', o')$ if and only if $(G, o, \mathbf{1}) \equiv_r (G', o', \mathbf{1})$, where $\mathbf{1}$ is the constant function equal to 1.

For two rooted weighted graphs \mathbb{G} and \mathbb{G}' , define their distance to be $1/(r+1)$, where r is the maximal radius on which they are equivalent, that is

$$(4) \quad d(\mathbb{G}, \mathbb{G}') = \left(1 + \sup \left\{ r \geq 0 : \mathbb{G} \equiv_r \mathbb{G}' \right\} \right)^{-1}.$$

Then, the metric space of rooted weighted graphs (up to isomorphisms), endowed with the distance d , is a Polish space. Indeed, both the space of rooted graphs endowed with d (by giving constant weights equal to 1) and $[0, 1]$ endowed with the distance in absolute value are Polish, and so it is not difficult to verify that the space of rooted weighted graphs is Polish. Thus, we can consider convergence, random variables, and random processes on it. In particular, this allows us to define the standard local weak topology, as introduced in [12], as follows.

A sequence of (random) weighted graphs $(\mathcal{G}_n)_{n \geq 1}$ converges *locally weakly* to the rooted weighted graph \mathbb{G} if, for any bounded continuous function h with respect to d ,

$$\mathbb{E}[h(\mathcal{G}_n, o_n)] \longrightarrow \mathbb{E}[h(\mathbb{G})],$$

where o_n is a vertex from $V(\mathcal{G}_n)$ chosen uniformly at random. When that is the case, we denote it by $\mathcal{G}_n \xrightarrow{\text{lwc}} \mathbb{G}$. We further extend this notation to allow the convergence to occur not only in expectation, but also in probability, i.e., for random graphs \mathcal{G}_n ,

$$\mathbb{E}[h(\mathcal{G}_n, o_n) \mid \mathcal{G}_n] \xrightarrow{\mathbb{P}} \mathbb{E}[h(\mathbb{G})],$$

where now the expectation is with respect to the random root only (i.e., now with respect to the randomness of the graph). When that is the case, we say that $(\mathcal{G}_n)_{n \geq 1}$ converges *locally in probability* towards \mathbb{G} and denote it by $\mathcal{G}_n \xrightarrow{\text{lcP}} \mathbb{G}$. Similarly, we extend these two notations to rooted weighted graphs (that is, $\mathbb{G}_n \xrightarrow{\text{lwc}} \mathbb{G}$ and $\mathbb{G}_n \xrightarrow{\text{lcP}} \mathbb{G}$), where it is always the case that the root is chosen uniformly at random among the vertex set of G_n . We also later apply a function on \mathbb{G}_n (namely, Prim's algorithm) but the root is always chosen in advance, and uniformly in $V(G_n)$. Finally, when considering unweighted graphs, we simply add constant edge weights equal to 1, and use the same notations and terminology. Interestingly, the two notions of weighted and unweighted local weak convergence are strongly related, as specified by the following lemma:

Lemma 2.3 (Local convergence of random graphs with i.i.d. weights). *Let $(G_n)_{n \geq 1}$ be a sequence of graphs and (G, o) a rooted graph. Further endow these graphs with independent and identically distributed edge weights to obtain $(\mathcal{G}_n)_{n \geq 1} = (G_n, w_n)_{n \geq 1}$ and $\mathbb{G} = (G, o, w)$, where the weights have the same distributions. Then weighted and unweighted local weak convergence are equivalent, that is*

$$G_n \xrightarrow{\text{lwc}} (G, o) \iff \mathcal{G}_n \xrightarrow{\text{lwc}} \mathbb{G}.$$

Proof. We give a direct proof, but this can also be proved using a coupling argument instead. Fix $r \geq 0$ and a rooted weighted graph $\mathbb{H} = (H, \rho, \xi)$. Further write W for a random variable distributed according to the common weight distribution. Then, given the random graph $\mathbb{G} = (G, o, w)$ and using the definition of \equiv_r , the event $\{\mathbb{G} \equiv_r \mathbb{H}\}$ can be split according to the existence of the isomorphism and the weights on this isomorphism being close to each

other:

$$\begin{aligned} \mathbb{P}(\mathbb{G} \equiv_r \mathbb{H}) &= \mathbb{P}(\mathbb{G} \equiv_r \mathbb{H} \mid (G, o) \equiv_r (H, \rho)) \mathbb{P}((G, o) \equiv_r (H, \rho)) \\ &= \left(\prod_{e \in E(B_r(\mathbb{H}))} \mathbb{P}\left(|W - \xi(e)| \leq \frac{1}{r}\right) \right) \mathbb{P}((G, o) \equiv_r (H, \rho)). \end{aligned}$$

Denote the product of probabilities in the last equation by $P(\mathbb{H})$. The same argument as before can then be applied to (\mathcal{G}_n, v) for any $n \geq 1$ and $v \in V(\mathcal{G}_n)$, leading to

$$\begin{aligned} &\frac{1}{|\mathcal{G}_n|} \sum_{v \in V(\mathcal{G}_n)} \mathbb{P}((\mathcal{G}_n, v) \equiv_r \mathbb{H}) - \mathbb{P}(\mathbb{G} \equiv_r \mathbb{H}) \\ &= P(\mathbb{H}) \cdot \left(\frac{1}{|G_n|} \sum_{v \in V(G_n)} \mathbb{P}((G_n, v) \equiv_r (H, \rho)) - \mathbb{P}((G, o) \equiv_r (H, \rho)) \right). \end{aligned}$$

To conclude this proof, use [36, Theorem 2.16] to restrict our study to \mathbb{H} such that $P(\mathbb{H}) > 0$, thus proving that the left-hand side converges to 0 if and only if the right-hand side does as well. It now suffices to notice that the left-, respectively, right-hand sides converge to 0 for all such \mathbb{H} if and only if the weighted, respectively unweighted, graph sequence converges, which is exactly the desired statement. \square

2.5. Properties of the local limit. In this section we explore some of the known properties regarding the local limits of graph sequences. As a first step, and because it is an essential part of our work, we provide the following theorem stating that the local limit of the minimum spanning tree is the minimum spanning forest:

Theorem 2.4 ([6, Theorem 5.4]). *Let $(\mathcal{G}_n)_{n \geq 1}$ be a sequence of weighted graphs which converges locally weakly towards \mathbb{G} . Then, under the standard product topology,*

$$(\mathcal{G}_n, \mathbb{T}^{\mathcal{G}_n}) \xrightarrow{\text{lwc}} (\mathbb{G}, \mathbb{F}^{\mathbb{G}}).$$

It is easy to see that, when \mathbb{G}_n is the standard extension of G_n , the previous result implies that the paired rooted weighted convergence also occurs, i.e.,

$$(\mathbb{G}_n, \mathbb{T}^{\mathbb{G}_n}) \xrightarrow{\text{lwc}} (\mathbb{G}, \mathbb{F}^{\mathbb{G}}).$$

Furthermore, using a similar argument as that of the original proof, the previous convergence of measures can be extended, thanks to Skorohod's representation theorem, to an almost-sure convergence on an appropriate probability space. Under this coupling, for any $r \geq 1$, there exists n_0 such that, for $n \geq n_0$, we have $\mathbb{G}_n \equiv_r \mathbb{G}$ and $\mathbb{T}^{\mathbb{G}_n} \equiv_r \mathbb{F}^{\mathbb{G}}$. Naturally, this can be extended to $r = r_n$ growing slowly enough. We say that a sequence of functions $k = (k_n(\cdot))_{n \geq 1}$ is a *linearly growing sequence* with respect to $(\mathbb{G}_n)_{n \geq 1}$ (or equivalently with respect to $(G_n)_{n \geq 1}$) if, for any $t \in [0, 1]$, we have $\lim_{n \rightarrow \infty} k_n(t)/n = t$, and further $(k_n(0))_{n \geq 1}$ is a slowly growing radius diverging to infinity such that there exists a coupling satisfying

$$(5) \quad \mathbb{P}(\mathbb{G}_m \equiv_{k_n(0)} \mathbb{G} \text{ and } \mathbb{F}^{\mathbb{G}_m} \equiv_{k_n(0)} \mathbb{F}^{\mathbb{G}}) = 1$$

for all $m \geq n$. In other words, this definition means that $k_n(t) = nt + o(n)$ where the $o(n)$ term when $t = 0$ is diverging at a rate allowing us to couple $(\mathbb{G}_n, \mathbb{T}^{\mathbb{G}_n})$ with $(\mathbb{G}, \mathbb{F}^{\mathbb{G}})$.

After having stated the fundamental result regarding the convergence of the minimum spanning tree, and having defined linearly growing sequences, we now complete this section with applications of Assumption 3. The first result states that the largest component of

the finite sequence of graphs asymptotically covers the same proportion of vertices as the probability for the graph to be infinite in the limit:

Proposition 2.5 ([36, Theorem 2.28]). *Let $(G_n)_{n \geq 1}$ be locally converging in probability to (G, o) . Let $(\mathbb{G}_n)_{n \geq 1}$ and \mathbb{G} be the respective standard expansions of $(G_n)_{n \geq 1}$ and (G, o) , and define θ as in (1). Then, if $(G_n)_{n \geq 1}$ satisfies Assumption 3, for any $p \in [0, 1]$,*

$$\frac{|\mathcal{C}_{(1)}^{\mathbb{G}_n}(p)|}{n} \xrightarrow{\mathbb{P}} \theta(p) \quad \text{and} \quad \frac{|\mathcal{C}_{(2)}^{\mathbb{G}_n}(p)|}{n} \xrightarrow{\mathbb{P}} 0.$$

In fact, if Assumption 3 holds for some $p \in [0, 1]$, then also Proposition 2.5 holds for that p .

This provides a first characterisation of the sizes of components in the finite graphs. It is interesting to note that the fact that the second largest component has a size of a smaller order than n also means that this is true for every subsequent component. We now provide a more local characteristic of the largest and other components:

Proposition 2.6 ([36, Theorem 2.32]). *Let $(G_n)_{n \geq 1}$ be locally converging in probability to (G, o) . Let $(\mathbb{G}_n)_{n \geq 1}$ and \mathbb{G} be the respective standard expansions of $(G_n)_{n \geq 1}$ and (G, o) , and define θ as in (1). Then, if $(G_n)_{n \geq 1}$ satisfies Assumption 3, for any $p \in [0, 1]$,*

$$\mathcal{C}_{(1)}^{\mathbb{G}_n}(p) \xrightarrow{\text{lcP}} \mathcal{C}_o^{\mathbb{G}}(p) \mid |\mathcal{C}_o^{\mathbb{G}}(p)| = \infty \quad \text{and} \quad \bigcup_{k \geq 2} \mathcal{C}_{(k)}^{\mathbb{G}_n}(p) \xrightarrow{\text{lcP}} \mathcal{C}_o^{\mathbb{G}}(p) \mid |\mathcal{C}_o^{\mathbb{G}}(p)| < \infty.$$

This second proposition is very useful as it states that we can approximately know whether the root belongs to the largest component or not simply from observing the neighbourhood of a node. This is a key point in the proof of Proposition 2.7 below, stating that the only large component in a neighbourhood is the largest component:

Proposition 2.7. *Let $(G_n)_{n \geq 1}$ be locally converging in probability to (G, o) . Let $(\mathbb{G}_n)_{n \geq 1}$ and \mathbb{G} be the respective standard expansions of $(G_n)_{n \geq 1}$ and (G, o) . Then, if $(G_n)_{n \geq 1}$ satisfies Assumption 3, for any $p \in [0, 1]$ and $r \geq 0$,*

$$\lim_{k \rightarrow \infty} \limsup_{n \rightarrow \infty} \mathbb{P}(\exists v \in B_r(\mathbb{G}_n) \setminus \mathcal{C}_{(1)}^{\mathbb{G}_n}(p) : |\mathcal{C}_v^{\mathbb{G}_n}(p)| \geq k) = 0.$$

In other words, for any $\varepsilon > 0$, there exists k_0 and $n_0 = n_0(k_0)$ such that, for all $k \geq k_0$ and $n \geq n_0$,

$$\mathbb{P}(\exists v \in B_r(\mathbb{G}_n) \setminus \mathcal{C}_{(1)}^{\mathbb{G}_n}(p) : |\mathcal{C}_v^{\mathbb{G}_n}(p)| \geq k) \leq \varepsilon.$$

Proof. Fix $\varepsilon > 0$. First observe that, for $M > 0$ large enough and for all $n \geq 1$,

$$\mathbb{P}(|B_{2r}(\mathbb{G}_n)| > M) \leq \varepsilon.$$

It follows that

$$\begin{aligned} & \mathbb{P}(\exists v \in B_r(\mathbb{G}_n) \setminus \mathcal{C}_{(1)}^{\mathbb{G}_n}(p) : |\mathcal{C}_v^{\mathbb{G}_n}(p)| \geq k) \\ & \leq \varepsilon + \mathbb{P}(\exists v \in B_r(\mathbb{G}_n) \setminus \mathcal{C}_{(1)}^{\mathbb{G}_n}(p) : |\mathcal{C}_v^{\mathbb{G}_n}(p)| \geq k, |B_{2r}(\mathbb{G}_n)| \leq M) \\ & \leq \varepsilon + \mathbb{P}(\exists v \in B_r(\mathbb{G}_n) : v \in \bigcup_{i \geq 2} \mathcal{C}_{(i)}^{\mathbb{G}_n}(p), |\mathcal{C}_v^{\mathbb{G}_n}(p)| \geq k, |B_r(\mathbb{G}_n, v)| \leq M), \end{aligned}$$

where we have used that if the ball of radius $2r$ has size at most M , then so does any ball of radius r of a vertex at distance at most r from the original root. Furthermore, using that

$v \in B_r(\mathbb{G}_n) = B_r(G_n, o_n, w_n)$ if and only if $o_n \in B_r(G_n, v)$, and by summing over all possible $v \in V(\mathbb{G}_n)$, the previous bound becomes

$$\begin{aligned} & \mathbb{P}\left(\exists v \in B_r(\mathbb{G}_n) \setminus \mathcal{C}_{(1)}^{\mathbb{G}_n}(p) : |\mathcal{C}_v^{\mathbb{G}_n}(p)| \geq k\right) \\ & \leq \varepsilon + \sum_{v \in V(G_n)} \mathbb{P}\left(o_n \in B_r(G_n, v), v \in \bigcup_{i \geq 2} \mathcal{C}_{(i)}^{\mathbb{G}_n}(p), |\mathcal{C}_v^{\mathbb{G}_n}(p)| \geq k, |B_r(G_n, v)| \leq M\right). \end{aligned}$$

Observe that only the first event of the probability depends explicitly on o_n . Thus, using that o_n is uniform over the set of vertices, it follows that

$$\begin{aligned} & \mathbb{P}\left(\exists v \in B_r(\mathbb{G}_n) \setminus \mathcal{C}_{(1)}^{\mathbb{G}_n}(p) : |\mathcal{C}_v^{\mathbb{G}_n}(p)| \geq k\right) \\ & \leq \varepsilon + \sum_{v \in V(G_n)} \mathbb{E}\left[\frac{|B_r(G_n, v)|}{|V(G_n)|} \cdot \mathbb{1}_{\left\{v \in \bigcup_{i \geq 2} \mathcal{C}_{(i)}^{\mathbb{G}_n}(p), |\mathcal{C}_v^{\mathbb{G}_n}(p)| \geq k, |B_r(G_n, v)| \leq M\right\}}\right] \\ & \leq \varepsilon + \frac{M}{|V(G_n)|} \sum_{v \in V(G_n)} \mathbb{P}\left(v \in \bigcup_{i \geq 2} \mathcal{C}_{(i)}^{\mathbb{G}_n}(p), |\mathcal{C}_v^{\mathbb{G}_n}(p)| \geq k\right). \end{aligned}$$

To conclude the proof, we simply note that Proposition 2.6 tells us that

$$\lim_{n \rightarrow \infty} \frac{1}{|V(G_n)|} \sum_{v \in V(G_n)} \mathbb{P}\left(v \in \bigcup_{i \geq 2} \mathcal{C}_{(i)}^{\mathbb{G}_n}(p), |\mathcal{C}_v^{\mathbb{G}_n}(p)| \geq k\right) = \mathbb{P}\left(|\mathcal{C}_o^{\mathbb{G}}(p)| \geq k \mid |\mathcal{C}_o^{\mathbb{G}}(p)| < \infty\right),$$

and the right-hand side converges to 0 as $k \rightarrow \infty$. This proves the desired convergence. \square

2.6. Process convergence. We conclude Section 2 with a short background on process convergence for sequences of stochastic processes as used in Theorem 1.1. The main strategy is to use the standard *Skorohod topology* (also called *Skorohod J_1 topology*) to allow for processes on Polish spaces (like the space of rooted weighted graphs) to be defined. We hereafter simplify the process convergence to three conditions and refer to [16, Chapter 3] for more details on this topology.

Let $(\mathbb{G}_n(t))_{n \geq 1, t \in [0,1]}$ be a sequence of cadlag processes on the space of rooted weighted graphs and recall the definition of d from (4). Then we say that it converges towards the cadlag process $\mathbb{G} = (\mathbb{G}(t))_{t \in [0,1]}$, and write it $(\mathbb{G}_n(t))_{t \in [0,1]} \xrightarrow{\text{lpc}} (\mathbb{G}(t))_{t \in [0,1]}$ if and only if it satisfies the following three conditions.

▷ The multi-dimensional process converges: for all integer $r \geq 1$, sets of time t_1, \dots, t_k , and rooted weighted graphs H_1, \dots, H_k , we have

$$(6) \quad \lim_{n \rightarrow \infty} \mathbb{P}\left(\mathbb{G}_n(t_1) \equiv_r H_1, \dots, \mathbb{G}_n(t_k) \equiv_r H_k\right) = \mathbb{P}\left(\mathbb{G}(t_1) \equiv_r H_1, \dots, \mathbb{G}(t_k) \equiv_r H_k\right).$$

▷ The limiting process is continuous at 1: for all $\varepsilon > 0$, we have

$$(7) \quad \lim_{\delta \rightarrow 0} \mathbb{P}\left(d(\mathbb{G}(1), \mathbb{G}(1 - \delta)) \geq \varepsilon\right) = 0$$

▷ The sequence does not have high increments: for all $\varepsilon > 0$, there exists $\delta_n \in (0, 1)$ such that

$$(8) \quad \lim_{n \rightarrow \infty} \mathbb{P}\left(\exists t < t' < t'' < t + \delta_n : d(\mathbb{G}_n(t), \mathbb{G}_n(t')) \geq \varepsilon \text{ and } d(\mathbb{G}_n(t'), \mathbb{G}_n(t'')) \geq \varepsilon\right) = 0.$$

This definition for the local process convergence arises from [16, Theorem 13.3] and the latter two conditions (7) and (8) can be understood as the graph process being tight. We refer to the extensive discussion of dynamic local convergence in [47].

It is worth noting that while this notion of convergence can be seen as very generic, the structure of $\mathbb{G}_n(t)$ is specific in our case. Indeed, it is always of the form $f_t(G_n, o_n, w_n)$ where $(G_n)_{n \geq 1}$ is a generic sequence of graphs, o_n is a uniformly chosen vertex, w_n are independent $\text{UNIFORM}([0, 1])$ edge weights, and f_t is a function mapping rooted weighted graphs unto rooted trees (more precisely, Prim's algorithm applied a partial number of steps).

3. CONVERGENCE OF PRIM'S ALGORITHM

The goal of this section is to prove Theorem 1.1. Let us start by recalling some notations that we use throughout this section.

For the rest of the section, we assume that $(G_n)_{n \geq 1}$ converges locally in probability towards (G, o) and we let $(\mathbb{G}_n)_{n \geq 1}$ and \mathbb{G} be the respective standard expansions of $(G_n)_{n \geq 1}$ and (G, o) . Using Skorohod's representation theorem, we assume without loss of generality that $(\mathbb{G}_n, \mathbb{T}^{\mathbb{G}_n})$ converges almost-surely towards $(\mathbb{G}, \mathbb{F}^{\mathbb{G}})$ and we let $k = (k_n(\cdot))_{n \geq 1}$ be a linearly growing sequence satisfying (5). We let $\theta = \theta^{\mathbb{G}}$, θ^{-1} , and $\mathcal{K} = \mathcal{K}^{\mathbb{G}}$ be defined as in (1), (2), and (3) respectively. Finally, assume that \mathbb{G} satisfies Assumption 1 and $(\mathbb{G}_n)_{n \geq 1}$ satisfies Assumption 3 (Assumption 2 is not necessary for the first results).

The rest of the section is organised as follows. In Section 3.1, we cover a special convergence result for Prim's algorithm which we then use in Section 3.2 to first prove the necessary one-dimensional convergence for Theorem 1.1. In turn, we then extend one-dimensional to multi-dimensional convergence as in (6). Finally, we conclude with the proof that Prim's algorithm and its limit satisfy the two conditions (7) and (8), thus proving convergence as a process.

3.1. Prim's algorithm with an appropriate number of steps. In this section, we provide a special case for the convergence of Prim's algorithm towards the expanded invasion percolation cluster, which is the first key step in the proof of Theorem 1.1. We state this result as follows:

Proposition 3.1. *Let Assumptions 1 and 3 hold. Using the notations from the beginning of Section 3, for any $r \geq 1$ and p such that $\theta(p) > 0$,*

$$\lim_{n \rightarrow \infty} \mathbb{P} \left(\mathcal{P}_{\mathcal{K}(p) + |\mathcal{C}_{(1)}^{\mathbb{G}_n}(p)|}(\mathbb{G}_n) \equiv_r \mathbb{F}_+^{\mathbb{G}}(p) \right) = 1.$$

Proof. Fix $r \geq 1$, $\varepsilon > 0$ and p satisfying $\theta(p) > 0$. Recall that $k_n = k_n(0)$ diverges and that $\mathbb{G}_n \equiv_{k_n} \mathbb{G}$ and $\mathbb{T}^{\mathbb{G}_n} \equiv_{k_n} \mathbb{F}^{\mathbb{G}}$ almost-surely. Using that $\theta(p) > 0$ along with Assumption 1 and Proposition 2.1, $\mathcal{K}(p)$ is finite so there exists n_0 such that, for all $n \geq n_0$,

$$\mathbb{P}(\mathcal{K}(p) > k_n) \leq \frac{\varepsilon}{3}.$$

Moreover, by applying Proposition 2.7 with $r = k_{n_0}$, we know that there exists n_1 and k_0 such that, for all $n \geq n_1$ and $k \geq k_0$,

$$\mathbb{P} \left(\exists v \in B_{k_{n_0}}(\mathbb{G}_n) \setminus \mathcal{C}_{(1)}^{\mathbb{G}_n}(p) : |\mathcal{C}_v^{\mathbb{G}_n}(p)| \geq k \right) \leq \frac{\varepsilon}{3}.$$

Further, set n_2 such that, for all $n \geq n_2$, we have $k_n \geq k_{n_0} + k_0$. Finally, using that \mathbb{G} is locally finite, define n_3 such that, for all $n \geq n_3$,

$$\mathbb{P} \left(\inf \left\{ |w_1 - w_2| : w_1 \neq w_2 \in \{p\} \cup \{w(e) : e \in E(B_{k_{n_2}}(\mathbb{G}))\} \right\} \leq \frac{1}{2k_n} \right) \leq \frac{\varepsilon}{3}.$$

In words, the distance between all edge weights in $B_{k_{n_2}}(\mathbb{G})$ as well as the distance between these edge weights and p is at least $1/2k_{n_3}$.

For all n , we write E_n for the union of all the bad events considered previously, that is $\mathcal{K}(p) > k_{n_0}$, $B_{k_{n_0}}(\mathbb{G}_n)$ has a component of size at least k_0 which is not its largest component, and the weights of the edges of $B_{k_{n_2}}(\mathbb{G})$ are at a distance less than $1/(2k_{n_3})$ from each other and from p . The previous definitions of n_0, n_1, n_2, n_3 imply that, for any $n \geq \max\{n_0, n_1, n_2, n_3\}$, $\mathbb{P}(E_n) \leq \varepsilon$. We now show that, on the event E_n^c , the event stated in the proposition holds.

We recall that $v(p) = v^{\mathbb{G}}(p)$ was defined in Section 2.3 as the first vertex in an infinite component explored by Prim's algorithm on \mathbb{G} . In particular, on the event E_n^c , we have that $\mathcal{K}(p) \leq k_{n_0} < \infty$ and so $|\mathcal{C}_{v(p)}^{\mathbb{G}}(p)| = \infty$. Furthermore, if $\mathcal{K}(p) \leq k_{n_0}$ and since $k_{n_2} \geq k_{n_0} + k_0 \geq k_0 + \mathcal{K}(p)$, we see that $\mathcal{C}_{v(p)}^{\mathbb{G}}(p)$ intersects $B_{k_{n_2}}(\mathbb{G}, o)$ in at least k_0 vertices.

On the event E_n^c , and since the edge weights of \mathbb{G}_n and \mathbb{G} are at a distance at most $1/(2k_n)$, we observe that the ordering of the edges on the ball of radius k_{n_2} is the same between \mathbb{G}_n and \mathbb{G} . In particular, as long as Prim's algorithm does not leave the ball of radius k_{n_2} , it explores the same vertices and in the same order in \mathbb{G}_n and in \mathbb{G} . Using that, on the event E_n^c , $\mathcal{K}(p) \leq k_{n_0}$, it follows that Prim's algorithm behaves in the same way on \mathbb{G}_n and on \mathbb{G} , at least until it reaches $v(p)$. However, then, since $\mathcal{C}_{v(p)}^{\mathbb{G}}(p)$ intersects $B_{k_{n_2}}(\mathbb{G})$ on at least k_0 vertices and since no weight in \mathbb{G} larger than p would be smaller than p in \mathbb{G}_n , it follows that $|\mathcal{C}_{v(p)}^{\mathbb{G}_n}(p)| \geq k_0$ and, by using that we are on E_n^c again, it follows that $\mathcal{C}_{v(p)}^{\mathbb{G}_n}(p) = \mathcal{C}_{(1)}^{\mathbb{G}_n}(p)$.

Consider now the tree $\mathcal{T} = \mathcal{P}_{\mathcal{K}(p)+|\mathcal{C}_{(1)}^{\mathbb{G}_n}(p)|}(\mathbb{G}_n)$. We aim to prove that, on E_n^c ,

$$\mathcal{T} \equiv_r \mathbb{F}_+^{\mathbb{G}}(p).$$

For all n such that $r \leq k_n$, since the weights on both balls are at least $1/2k_n$ away from each other thanks to the coupling between \mathbb{G} and \mathbb{G}_n , it only suffices to check that the two graphs are isomorphic to each other until depth r . To do so, we first observe that, since Prim's algorithm reaches $\mathcal{C}_{(1)}^{\mathbb{G}_n}(p)$ after $\mathcal{K}(p) + 1$ steps, it follows that

$$\mathcal{T} = \mathcal{P}_{\mathcal{K}(p)+1}(\mathbb{G}_n) \cup \mathbb{T}^{(\mathcal{C}_{(1)}^{\mathbb{G}_n}(p), w_n)},$$

where we recall that $w_n = (w_n(e))_{e \in E(\mathbb{G}_n)}$ are the edge weights on \mathbb{G}_n , and we note that \mathcal{T} is the first part of Prim's algorithm combined with the minimum spanning tree on its largest component at level p . Using that

$$\mathbb{T}^{(\mathcal{C}_{(1)}^{\mathbb{G}_n}(p), w_n)} = \mathbb{T}^{\mathbb{G}_n} \cap \mathcal{C}_{(1)}^{\mathbb{G}_n}(p),$$

along with the properties of the coupling, it follows that

$$\mathcal{T} \equiv_r \mathcal{P}_{\mathcal{K}(p)+1}(\mathbb{G}) \cup \left(\mathbb{F}^{\mathbb{G}} \cap \mathcal{C}_{v(p)}^{\mathbb{G}}(p) \right),$$

as long as $k_n \geq r$ and we are on E_n^c . However, this last term is exactly the target structure thanks to Lemma 2.2. Thus, we have proved that the desired equivalence is true, which occurs with probability at least $1 - \varepsilon$ for n large enough. This is exactly the desired convergence statement. \square

3.2. Marginal convergence. In this section, we use Proposition 3.1 to show that the multi-dimensional distributions converge as stated in Theorem 1.1. We start with the one-dimensional distribution:

Proposition 3.2 (One-dimensional convergence). *Let Assumptions 1 and 3 hold. Further, assume that θ satisfies Assumption 2. Using the notations from the beginning of Section 3, for any $r \geq 1$ and $t \in [0, 1]$,*

$$\lim_{n \rightarrow \infty} \mathbb{P}(\mathcal{P}_{k_n(t)}(\mathbb{G}_n) \equiv_r \mathbb{F}_+^{\mathbb{G}}(\theta^{-1}(t))) = 1.$$

Proof. Fix $\varepsilon > 0$, $r \geq 1$ and $t \in [0, 1]$. Let $\delta > 0$ be such that

$$(9) \quad \mathbb{P}\left(\inf \left\{ |\theta^{-1}(t) - w(e)| : e \in E(B_r(\mathbb{G})) \right\} \leq \delta \right) \leq \frac{\varepsilon}{5}.$$

We start with the case $t = 0$.

When $t = 0$, using the fact that $\mathbb{G}_n \equiv_{k_n(0)} \mathbb{G}$, it follows that

$$\mathcal{P}_{k_n(0)}(\mathbb{G}_n) \equiv_{k_n(0)} \mathcal{P}_{k_n(0)}(\mathbb{G}).$$

Then, since $k_n(0) \rightarrow \infty$, it also follows that

$$\lim_{n \rightarrow \infty} \mathbb{P}(\mathcal{P}_{k_n(0)}(\mathbb{G}) \equiv_r \mathcal{P}_{\infty}(\mathbb{G})) = 1.$$

It only remains to use that $\mathcal{P}_{\infty}(\mathbb{G}) = \mathbb{F}_+^{\mathbb{G}}(0)$, as stated in Lemma 2.2, to see that the desired result holds when $t = 0$.

Assume now that $t \in (0, 1)$. Thanks to Assumption 2 and the definition from (2), θ^{-1} is continuous and strictly increasing from $[0, 1]$ to $[p_-, 1]$, with $\theta^{-1}(0) = p_-$, and we recall that $\theta^{-1}(1) = 1$. Thus, we can find $p_- \leq \theta^{-1}(t) \leq p_+$ such that $\theta(p_-) < t < \theta(p_+)$ and $|p_+ - \theta^{-1}(t)|, |p_- - \theta^{-1}(t)| \leq \delta$. As a first step, we observe that $\theta(p_+) > t > 0$, and therefore we can apply Proposition 3.1 to find n_+ such that, for all $n \geq n_+$,

$$\mathbb{P}\left(\mathcal{P}_{\mathcal{K}(p_+)+|\mathcal{C}_{(1)}^{\mathbb{G}_n}(p_+)|}(\mathbb{G}_n) \equiv_r \mathbb{F}_+^{\mathbb{G}}(p_+)\right) \geq 1 - \frac{\varepsilon}{5}.$$

Moreover, using that $|p_+ - \theta^{-1}(p)| \leq \delta$ along with (9), the error in probability arising from replacing p_+ with $\theta^{-1}(p)$ in $\mathbb{F}_+^{\mathbb{G}}$ is $\varepsilon/5$, leading to

$$\mathbb{P}\left(\mathcal{P}_{\mathcal{K}(p_+)+|\mathcal{C}_{(1)}^{\mathbb{G}_n}(p_+)|}(\mathbb{G}_n) \equiv_r \mathbb{F}_+^{\mathbb{G}}(\theta^{-1}(p))\right) \geq 1 - \frac{2\varepsilon}{5}.$$

We next note that there are two possibilities: We either have that $\theta(p_-) > 0$, in which case we can apply the same method and obtain that, for all $n \geq n_-$,

$$\mathbb{P}\left(\mathcal{P}_{\mathcal{K}(p_-)+|\mathcal{C}_{(1)}^{\mathbb{G}_n}(p_-)|}(\mathbb{G}_n) \equiv_r \mathbb{F}_+^{\mathbb{G}}(\theta^{-1}(p))\right) \geq 1 - \frac{2\varepsilon}{5}.$$

Alternatively, $\theta(p_-) = 0$, and, in that case, we use the definition of k_n to see that

$$\mathbb{P}\left(\mathcal{P}_{k_n(0)}(\mathbb{G}_n) \equiv_r \mathcal{P}_{\infty}(\mathbb{G})\right) = 1 \geq 1 - \frac{\varepsilon}{5}.$$

Moreover, using Assumption 1 and Proposition 2.1, we see that $\theta(p_-) = 0$ implies that $\mathbb{P}(\mathcal{K}(p_-) = \infty)$, which, combined with Lemma 2.2, means that $\mathcal{P}_{\infty}(\mathbb{G}) = \mathbb{F}_+^{\mathbb{G}}(p_-)$ almost-surely. Again using the definition of δ and p_- , it follows that

$$\mathbb{P}\left(\mathcal{P}_{k_n(0)}(\mathbb{G}_n) \equiv_r \mathbb{F}_+^{\mathbb{G}}(\theta^{-1}(p))\right) \geq 1 - \frac{2\varepsilon}{5}.$$

Combing the previous results, we have proved that

$$\mathbb{P}\left(\mathcal{P}_{L_n}(\mathbb{G}_n) \equiv_r \mathbb{F}_+^{\mathbb{G}}(\theta^{-1}(p)) \equiv_r \mathcal{P}_{U_n}(\mathbb{G}_n)\right) \geq 1 - \frac{4\varepsilon}{5},$$

where

$$U_n = \mathcal{K}(p_+) + \left| \mathcal{C}_{(1)}^{\mathbb{G}_n}(p_+) \right|,$$

and

$$L_n = \begin{cases} \mathcal{K}(p_-) + \left| \mathcal{C}_{(1)}^{\mathbb{G}_n}(p_-) \right| & \text{if } \theta(p_-) > 0, \\ k_n(0) & \text{otherwise.} \end{cases}$$

It remains to prove that, with high probability,

$$(10) \quad L_n \leq k_n(t) \leq U_n.$$

To show that the bounds in (10) hold, use Proposition 2.5, as well as the fact that \mathcal{K} is almost-surely finite when θ is positive by Assumption 1 and Proposition 2.1, to obtain that

$$\frac{U_n}{n} \xrightarrow{\mathbb{P}} \theta(p_+) > t,$$

and

$$\frac{L_n}{n} \xrightarrow{\mathbb{P}} \begin{cases} \theta(p_-) < t & \text{if } \theta(p_-) > 0, \\ 0 & \text{otherwise.} \end{cases}$$

Recalling that $k_n(t)/n \rightarrow t > 0$, there exists $n_0 \geq \max\{n_+, n_-\}$ such that, for all $n \geq n_0$,

$$\mathbb{P}\left(L_n \leq k_n(t) \leq U_n\right) \geq 1 - \frac{\varepsilon}{5}.$$

Thus, for any $n \geq n_0$, and $k_n \geq r$,

$$\mathbb{P}\left(\mathcal{P}_{k_n(t)}(\mathbb{G}_n) \equiv_r \mathbb{F}_+^{\mathbb{G}}(\theta^{-1}(t))\right) \geq 1 - \varepsilon,$$

where we have used that Prim's algorithm is a monotone process. This is exactly the desired result when $t \in (0, 1)$. Thus, it only remains to cover the case $t = 1$. However, this case uses the exact same proof as that of the case $0 < t < 1$, except that the coupling directly gives us that

$$\mathbb{T}^{\mathbb{G}_n} = \mathcal{P}_n(\mathbb{G}_n) \equiv_r \mathbb{F}_+^{\mathbb{G}}(\theta^{-1}(1)) = \mathbb{F}^{\mathbb{G}},$$

which we use as the upper bound instead of Prim's algorithm performed U_n steps. Since the lower bound remains true in that case as well, the desired result follows. \square

We have proved that the one-dimensional distribution converges with high probability to the desired limit, and extending this proof to the multi-dimensional distribution directly follows from the fact that the intersection of finitely many high-probability events is also a high-probability event:

Corollary 3.3 (Multi-dimensional convergence). *Let Assumptions 1 and 3 hold. Further, assume that θ satisfies Assumption 2. Using the notations from the beginning of Section 3, for any $r \geq 1$ and $t, \dots, t_k \in [0, 1]$,*

$$\lim_{n \rightarrow \infty} \mathbb{P}\left(\forall \ell \in [k] : \mathcal{P}_{k_n(t_\ell)}(\mathbb{G}_n) \equiv_r \mathbb{F}_+^{\mathbb{G}}(\theta^{-1}(t_\ell))\right) = 1.$$

This concludes the section on marginal convergence, and in the next section, we use these results to prove that the process convergence occurs as stated in Theorem 1.1.

3.3. Process convergence. In this section we conclude the first part of this work by providing the proof of Theorem 1.1. To do so, we simply verify that the three assumptions (6), (7), and (8) from Section 2.6 are satisfied.

Proof of Theorem 1.1. Observe first that the convergence corresponding to (6) directly follows from the definition of the coupling at the beginning of Section 3 along with Corollary 3.3. For the second condition, we observe that, if all edges in $B_r(\mathbb{G})$ have weight less than $1 - \delta$, then $\mathbb{F}_+^{\mathbb{G}}(1 - \delta) \equiv_r \mathbb{F}_+^{\mathbb{G}}(1)$. Using that \mathbb{G} is locally finite, this implies that, for all $r \geq 1$,

$$\lim_{\delta \rightarrow 0} \mathbb{P} \left(d \left(\mathbb{F}_+^{\mathbb{G}}(1 - \delta), \mathbb{F}_+^{\mathbb{G}}(1) \right) > \frac{1}{r+1} \right) = 0,$$

and this exactly corresponds to (7). We now prove that the final condition in (8) holds.

Fix $r \geq 1$ and $\varepsilon > 0$. Our goal is to show that there exists n_0 and $\delta_0 > 0$ such that, for all $n \geq n_0$ and $\delta \leq \delta_0$, we have

$$\mathbb{P} \left(\exists t < t' < t'' < t + \delta : \begin{cases} d \left(\mathcal{P}_{k_n(t)}(\mathbb{G}_n), \mathcal{P}_{k_n(t')}(\mathbb{G}_n) \right) > \frac{1}{r+1} \\ d \left(\mathcal{P}_{k_n(t')}(\mathbb{G}_n), \mathcal{P}_{k_n(t'')}(\mathbb{G}_n) \right) > \frac{1}{r+1} \end{cases} \right) \leq \varepsilon.$$

Using first that θ^{-1} is continuous and strictly monotone from $[0, 1]$ to $[p_c, 1]$, as implied by Assumption 2, we observe that

$$\Delta(\theta, \delta) = \sup \left\{ \left| \theta^{-1}(t + \delta) - \theta^{-1}(t) \right| : t \in [0, 1 - \delta] \right\}$$

is decreasing in δ , strictly positive whenever $\delta > 0$, and converges to 0 as $\delta \rightarrow 0$. Combining this with the fact that \mathbb{G} is locally finite, we can set δ_0 so that

$$\mathbb{P} \left(\min \left\{ |w(e) - w(e')| : e \neq e' \in E(B_r(\mathbb{G})) \right\} \leq \Delta(\theta, 3\delta_0) \right) \leq \frac{\varepsilon}{2}.$$

We observe that the complement of the previous event implies that no segment of length less than $\Delta(\theta, 2\delta_0)$ contains two edge weights. From this observation, the previous definition of δ_0 implies that

$$\mathbb{P} \left(\exists t, \exists e \neq e' \in E(B_r(\mathbb{G})) : \theta^{-1}(t) \leq w(e), w(e') \leq \theta^{-1}(t + 2\delta_0) \right) \leq \frac{\varepsilon}{2}.$$

Consider now a sequence of closed intervals $([t_\ell^-, t_\ell^+])_{1 \leq \ell \leq k}$ of length $2\delta_0$ such that any interval of length δ_0 is fully included into one of these intervals. One can for example choose $t_\ell^- = (\ell - 1)\delta_0$ with $t_\ell^+ = t_\ell^- + 2\delta_0$ for $\ell \leq 1/\delta - 1$ and add the extra interval $[1 - 2\delta_0, 1]$ at the end. Using Corollary 3.3, we can find n_0 such that, for all $n \geq n_0$,

$$\mathbb{P} \left(\forall \ell \in [k] : \mathcal{P}_{k_n(t_\ell^\pm)}(\mathbb{G}_n) \equiv_r \mathbb{F}_+^{\mathbb{G}}(\theta^{-1}(t_\ell^\pm)) \right) \geq 1 - \frac{\varepsilon}{2}.$$

We conclude this proof by showing that δ_0 and n_0 satisfy the desired property. We assume without loss of generality that n_0 is large enough so that $\mathbb{G}_n \equiv_r \mathbb{G}$ for all $n \geq n_0$.

Consider $n \geq n_0$ and $\delta \leq \delta_0$. Consider three times $t < t' < t'' < t + \delta$. Then t, t', t'' are all contained in an interval of length δ_0 and so they all belong to one of the interval $[t_\ell^-, t_\ell^+]$ for some ℓ . However, now, if we assume that $\mathcal{P}_{k_n(t_\ell^\pm)}(\mathbb{G}_n) = \mathbb{F}_+^{\mathbb{G}}(\theta^{-1}(t_\ell^\pm))$ and that there exists no pair of edges whose difference of weights is less than $\Delta(\theta, 2\delta_0)$, the existing coupling between \mathbb{G}_n and \mathbb{G} as well as the monotony of Prim's algorithm implies that, for all $\tilde{t} \in \{t, t', t''\}$,

$$\mathcal{P}_{k_n(\tilde{t})}(\mathbb{G}_n) \equiv_r \mathbb{F}_+^{\mathbb{G}}(t_\ell^-),$$

or

$$\mathcal{P}_{k_n(\bar{i})}(\mathbb{G}_n) \equiv_r \mathbb{F}_+^{\mathbb{G}}(t_\ell^+).$$

It follows that

$$\min \left\{ d\left(\mathcal{P}_{k_n(t)}(\mathbb{G}_n), \mathcal{P}_{k_n(t')}(\mathbb{G}_n)\right), d\left(\mathcal{P}_{k_n(t')}(\mathbb{G}_n), \mathcal{P}_{k_n(t'')}(\mathbb{G}_n)\right) \right\} \leq \frac{1}{r+1},$$

and thus

$$\begin{aligned} & \mathbb{P} \left(\exists t < t' < t'' < t + \delta : \begin{cases} d\left(\mathcal{P}_{k_n(t)}(\mathbb{G}_n), \mathcal{P}_{k_n(t')}(\mathbb{G}_n)\right) > \frac{1}{r+1} \\ d\left(\mathcal{P}_{k_n(t')}(\mathbb{G}_n), \mathcal{P}_{k_n(t'')}(\mathbb{G}_n)\right) > \frac{1}{r+1} \end{cases} \right) \\ & \leq \mathbb{P} \left(\exists \ell \in [k] : \mathcal{P}_{k_n(t_\ell^\pm)}(\mathbb{G}_n) \not\equiv_r \mathbb{F}_+^{\mathbb{G}}(\theta^{-1}(t_\ell^\pm)) \right) \\ & \quad + \mathbb{P} \left(\exists t, \exists e \neq e' \in E(B_r(\mathbb{G})) : \theta^{-1}(t) \leq w(e), w(e') \leq \theta^{-1}(t + 2\delta_0) \right) \leq \varepsilon, \end{aligned}$$

as desired. This proves that the third condition in (8), as stated in Section 2.6, is satisfied and thus concludes the proof of Theorem 1.1. \square

4. EXAMPLES AND COUNTER-EXAMPLES

In Section 4.1, we first show that Theorem 1.1 applies to four common graph sequences and limits: the multi-dimensional grid, the Erdős-Rényi graph, the configuration model, and the preferential attachment graph. In Section 4.2, we then provide two interesting examples of graphs satisfying exactly one of Assumption 1 and Assumption 3 and state the conjecture that Assumption 1 is not strictly necessary.

4.1. Example graph sequences. In this section, we discuss important examples of graph sequences to which our results apply.

Grids. The minimal spanning tree on grids is closely related to percolation on \mathbb{Z}^d , which is a topic of intense research. See [32] for overviews of the field, and [34] for an overview of percolation in high dimensions. Assumptions 1, 2 and 3 are all well known. In particular, Assumption 1 holds since the infinite component is *unique* (see [32, Theorem 8.1]), so that the minimal spanning forest (and even Prim's algorithm) needs to enter it at some point (and then in fact will never leave it). Assumption 2 is [32, Theorem 8.8]. Finally, Assumption 3 follows from the fact that $|\mathcal{C}_{(1)}^{\mathbb{G}_n}(p)|/|V(G_n)| \xrightarrow{\mathbb{P}} \theta(p)$, while $|\mathcal{C}_{(2)}^{\mathbb{G}_n}(p)|/|V(G_n)| \xrightarrow{\mathbb{P}} 0$ (see e.g., [38]), and the fact that this is equivalent to Assumption 3 by [35]. Thus, our results apply to this important case.

Complete graph. Percolation on the complete graph is the Erdős-Rényi random graph, and the connections between the minimal spanning tree and the critical Erdős-Rényi random graph are tight [2]. Our results apply, see also below for the discussion on the configuration model. The setting is slightly different though, as we now need to take $p = \lambda/n$, and consider λ as the parameter instead of $p \in [0, 1]$.

Configuration model. Assumption 1 is true as the local limit is a unimodular branching-process tree. Assumption 3 is true by [36, Section 4.3.1], while continuity of θ , which implies Assumption 2, follows from [46]. The above references can easily be extended to the Erdős-Rényi random graph.

Preferential attachment models. Again our results apply. Assumption 1 is true by [33], which also proves the continuity of $p \mapsto \theta(p)$ in Assumption 2. Assumption 3 again holds since $|\mathcal{C}_{(1)}^{\mathbb{G}_n}(p)|/|V(G_n)| \xrightarrow{\mathbb{P}} \theta(p)$, while $|\mathcal{C}_{(2)}^{\mathbb{G}_n}(p)|/|V(G_n)| \xrightarrow{\mathbb{P}} 0$ by [33], and the fact that this is equivalent to Assumption 3 by [35].

We conclude that there is a wide range of models to which our results apply, and where the dynamic local limit of Prim's algorithm gives insight into how the minimal spanning tree is formed.

4.2. A discussion on the assumptions. In this section, we provide some examples of graph sequences which do not satisfy at least one of the assumptions and study the behaviour of Prim's algorithm in these cases.

Union graphs. For the rest of the section, we let $(\mathbb{G}_n^1)_{n \geq 1}$ and $(\mathbb{G}_n^2)_{n \geq 1}$ be two sequences of graphs locally converging to \mathbb{G}^1 and \mathbb{G}^2 , with percolation functions θ^1 and θ^2 and both satisfying all the assumptions from Theorem 1.1. One can for example consider $(\mathbb{G}_n^1)_{n \geq 1}$ and $(\mathbb{G}_n^2)_{n \geq 1}$ to be sequences of random d_1 - and d_2 -regular graphs. We further assume that both \mathbb{G}_n^1 and \mathbb{G}_n^2 have size n and let $(\mathbb{G}_n)_{n \geq 1}$ be the graph obtained by taking the union of $\mathbb{G}_{\lfloor n/2 \rfloor}^1$ and $\mathbb{G}_{\lfloor n/2 \rfloor}^2$, so that \mathbb{G}_n also has size n . We call such graphs *union graphs*. They are of interest since they are the simplest examples where we do have local convergence, but our assumptions may be false. It may be convenient to assume that the two graphs are connected by adding an arbitrary edge between the two components (or even an arbitrary, but finite, number of edges), and we note that this does not affect the local limit.

Before observing the effect of this construction on the different assumptions, we first observe that the local limit of $(\mathbb{G}_n)_{n \geq 1}$ is the graph \mathbb{G} which is distributed as \mathbb{G}^1 with probability $1/2$, and as \mathbb{G}^2 otherwise. This implies that $\theta = \theta^{\mathbb{G}} = (\theta^1 + \theta^2)/2$.

Satisfying the assumptions for union graphs. To start this section, we observe in which cases $(\mathbb{G}_n)_{n \geq 1}$ and its limit satisfy the different assumptions.

If we denote by p_c^1 and p_c^2 the critical percolation for \mathbb{G}^1 and \mathbb{G}^2 , we observe that Assumption 1 is satisfied if and only if $p_c^1 = p_c^2$. Indeed, by conditioning the probability from Assumption 1 according to whether \mathbb{G} is distributed as \mathbb{G}^1 or \mathbb{G}^2 , we observe that

$$\mathbb{P}\left(\exists v \in V(\mathbb{F}^{\mathbb{G}}) : |\mathcal{C}_v^{\mathbb{G}}(p)| = \infty\right) = \frac{1}{2} \left(\mathbb{1}_{\{\theta^1(p) > 0\}} + \mathbb{1}_{\{\theta^2(p) > 0\}} \right).$$

From this formula, and using that $\theta = (\theta^1 + \theta^2)/2$, we see that Assumption 1 holds if and only if θ^1 and θ^2 become positive at the same time, i.e., if and only if $p_c^1 = p_c^2$.

To test whether Assumption 2 holds here, we first observe that, if θ^1 and θ^2 both satisfy this assumption, then θ also satisfies it when $p_c^1 = p_c^2$. In the case where $p_c^1 \neq p_c^2$, since $p_c = \min\{p_c^1, p_c^2\}$, we see that Assumption 2 is satisfied if and only if the graph with larger critical value has a *continuous* percolation function. For example, if we assume that $p_c^1 < p_c^2$, then we want θ^2 to be continuous on the whole interval $[0, 1]$ (or equivalently on $(p_c^1, 1]$ since $\theta^2(p) = 0$ whenever $p < p_c^2$).

Finally, we can see that Assumption 3 does not hold if both $\theta^1(p)$ and $\theta^2(p)$ are positive for some $p < 1$. Indeed, in that case both $\mathbb{G}_{\lfloor n/2 \rfloor}^1$ and $\mathbb{G}_{\lfloor n/2 \rfloor}^2$ have components of linear sizes,

and using that only one edge connects the two graphs, we see that, for any $k \geq 0$,

$$\begin{aligned} \liminf_{n \rightarrow \infty} \frac{1}{n^2} \mathbb{E} \left[\left| \left\{ u, v \in V(\mathbb{G}_n) : |\mathcal{C}_u^{\mathbb{G}_n}(p)| \geq k, |\mathcal{C}_v^{\mathbb{G}_n}(p)| \geq k, \mathcal{C}_u^{\mathbb{G}_n}(p) \cap \mathcal{C}_v^{\mathbb{G}_n}(p) = \emptyset \right\} \right| \right] \\ \geq (1-p) \frac{\theta^1(p) \theta^2(p)}{4}. \end{aligned}$$

We further observe that Assumption 3 is still not satisfied when we remove the edge connecting the two graphs, or add a finite number more of such edges.

Satisfying our theorem for union graphs. In this part, we now evaluate in which case the sequence of union graphs $(\mathbb{G}_n)_{n \geq 1}$ satisfies Theorem 1.1. Interestingly, while it does not always satisfy it in its current form, it actually does by adapting the formula, leading us to believe that a more general result can be extracted from these observations.

First of all, since the result applies to $(\mathbb{G}_n)_{n \geq 1}$ when all assumptions are met, we now focus on the cases where one of the assumptions is *not* satisfied. We thus assume without loss of generality that $0 < p_c^1 < p_c^2 < 1$. In that case, if the root of \mathbb{G}_n lies within the half of the vertices corresponding to \mathbb{G}^i , then, on any ball of finite radius, it seems that $\mathcal{P}_k(\mathbb{G}_n)$ behaves similarly to $\mathcal{P}_k(\mathbb{G}_{\lfloor n/2 \rfloor}^i)$, whose local limit is $\mathbb{F}_+^{\mathbb{G}^i}(p)$ with p satisfying an equation of the form $\theta^i(p) \simeq 2k/n$.

The exact behaviour is actually a lot more complicated if there is an edge connecting the two graphs. Indeed, depending on when the endpoint of this edge is explored, and what the edge weight is, the exploration of the current graph \mathbb{G}_m^i might be suddenly interrupted in order to explore the other graph. Naturally, this behaviour becomes more and more complicated as we add more and more edges between the two graphs, while keeping the local limit intact.

For simplicity, we now assume that the two graphs are actually not connected to each other. Forgetting about the process convergence for now, the previous observations lead, in that case, to the following convergence:

$$\mathcal{P}_{k_n(t)}(\mathbb{G}_n) \rightarrow \begin{cases} \mathbb{F}_+^{\mathbb{G}^1}((\theta^1)^{-1}(2t)) & \text{w.p. } 1/2 \\ \mathbb{F}_+^{\mathbb{G}^2}((\theta^2)^{-1}(2t)) & \text{w.p. } 1/2, \end{cases}$$

meaning that Theorem 1.1 can still hold without all the assumptions being satisfied, by adapting the statement. Note that a similar formula can be given for more complex union graphs, where we combine more than two graphs and do not necessarily have equal proportions of vertices in each graph.

While the previous convergence already has a major caveat in the fact that we might completely change the previous behaviour by simply adding one edge between the two graphs, it still gives the impression that Theorem 1.1 can easily be extended to more complex graph structures. However, a very similar case as the previous example, but with a distinctly different behaviour, is the sequence of graphs $(\mathbb{G}_n)_{n \geq 1}$ where instead of taking the union between two graphs, we flip a (fair) coin and choose the sequence accordingly. More precisely, $(\mathbb{G}_n)_{n \geq 1}$ is exactly $(\mathbb{G}_n^1)_{n \geq 1}$ with probability 1/2, and $(\mathbb{G}_n^2)_{n \geq 1}$ otherwise. Observe that this sequence of random graphs has the same limit as the union graph: it is \mathbb{G}^1 with probability 1/2 and \mathbb{G}^2 otherwise. We emphasize, though, that the behaviour of Prim's algorithm is quite different, as the exploration will have to go through the whole graph and thus look more like

$$\mathcal{P}_{k_n(t)}(\mathbb{G}_n) \rightarrow \begin{cases} \mathbb{F}_+^{\mathbb{G}^1}((\theta^1)^{-1}(t)) & \text{w.p. } 1/2, \\ \mathbb{F}_+^{\mathbb{G}^2}((\theta^2)^{-1}(t)) & \text{w.p. } 1/2. \end{cases}$$

Once again, in this case it is possible to extend Theorem 1.1 without satisfying all the assumptions.

Because of the previous examples, we believe that Theorem 1.1 can be extended past the assumptions we have made here. However, the complexity of possible behaviours, as well as the strong dependency on the *exact details* of the construction sequence of graphs considered, and not only their local limit, made us prefer its current format. It is still worth noting that, while it is possible to create much more intricate local behaviour than taking unions of converging graphs, or choosing sequence of converging graphs at random, these two interesting cases can actually still be well studied with our main result by simply adapting the limits accordingly.

Acknowledgement. This work is supported by the Netherlands Organisation for Scientific Research (NWO) through Gravitation-grant NETWORKS-024.002.003. The work of RvdH is further supported by the National Science Foundation under Grant No. DMS-1928930 while he was in residence at the Simons Laufer Mathematical Sciences Institute in Berkeley, California, during the spring semester 2025. The work of BC is further supported by the European Union's Horizon 2020 Research and Innovation Programme under the Marie Skłodowska-Curie grant agreement No. 101034253.



APPENDIX A. LIST OF SYMBOLS AND NOTATION

MST	Minimum spanning tree
IPC	Invasion percolation cluster
w.h.p.	With high probability
i.i.d.	Independent and identically distributed
$\text{UNIFORM}([0, 1])$	Standard uniform distribution on $[0, 1]$
$G = (V(G), E(G))$	Graph
$\mathcal{G} = (G, w)$	Weighted graph
(G, o)	Rooted graph with root o
$\mathbb{G} = (G, o, w)$	Rooted weighted graph with i.i.d. $\text{UNIFORM}([0, 1])$ edge weights
$\mathcal{G}(p)$	Graph \mathcal{G} obtained by only keeping edges of weight less than p
$\mathbb{G}(p)$	Rooted weighted graph obtained by only keeping edges of weight less than p
$\mathbf{F}^{\mathbb{G}} = \mathbf{F}^{(G, w)}$	Minimum spanning forest of \mathcal{G}
$\mathbb{F}^{\mathbb{G}} = (\mathbf{F}^{\mathbb{G}}, o)$	Rooted minimum spanning forest of \mathbb{G}
$\mathcal{P}_k(\mathbb{G})$	Rooted subtree of \mathbb{G} obtained by performing k steps of Prim's algorithm
$\mathbb{F}_+^{\mathbb{G}}(p)$	Expanded invasion percolation cluster at level p (see Section 2.2)
$\theta(p)$	Percolation survival probability
p_c	Critical percolation coefficient
$\mathcal{C}_o^{\mathbb{G}}(p)$	Component of o when only keeping edges with weight less than p
$\mathcal{C}_{(1)}^{\mathbb{G}_n}(p)$	Largest component when only keeping edges with weight less than p
$\mathcal{C}_{(2)}^{\mathbb{G}_n}(p)$	Second largest component when only keeping edges with weight less than p

$\xrightarrow{\mathbb{P}}$	Convergence in probability
$\xrightarrow{\text{lwc}}$	Local weak convergence (see Section 2.4)
$\xrightarrow{\text{lcP}}$	Local convergence in probability (see Section 2.4)
$\xrightarrow{\text{lpc}}$	Local process convergence (see Section 2.6)

REFERENCES

- [1] L. Addario-Berry. The local weak limit of the minimum spanning tree of the complete graph. Preprint available at arXiv:1301.1667v2 [math.PR], 2013.
- [2] L. Addario-Berry, N. Broutin, C. Goldschmidt, and G. Miermont. The scaling limit of the minimum spanning tree of the complete graph. *The Annals of Probability*, 45(5):3075–3144, 2017.
- [3] L. Addario-Berry and S. Sen. Geometry of the minimal spanning tree of a random 3-regular graph. *Probability Theory and Related Fields*, 180(3-4):553–620, 2021.
- [4] M. Albenque, L. Ménard, and G. Schaeffer. Local convergence of large random triangulations coupled with an Ising model. *Transactions of the American Mathematical Society*, 374:175–217, 2021.
- [5] D. Aldous and J. M. Steele. Asymptotics for Euclidean minimal spanning trees on random points. *Probability Theory and Related Fields*, 92(2):247–258, jun 1992.
- [6] D. Aldous and J. M. Steele. The objective method: Probabilistic combinatorial optimization and local weak convergence. *Probability on Discrete Structures*, 110:1–72, 2004.
- [7] K. S. Alexander. Percolation and minimal spanning forests in infinite graphs. *The Annals of Probability*, 23(1):347–370, jan 1995.
- [8] A. Anandkumar and R. Valluvan. Latent graphical model selection: Efficient methods for locally tree-like graphs. In *Advances in Neural Information Processing Systems*, volume 25, pages 1052–1060. Curran Associates, Inc., 2012.
- [9] G. B. Arous, S. Molchanov, and A. F. Ramírez. Transition from the annealed to the quenched asymptotics for a random walk on random obstacles. *The Annals of Probability*, 33(6):2149–2187, 2005.
- [10] S. Banerjee, S. Bhamidi, J. Shen, and S. P. Young. Local weak convergence and its applications. Preprint available at arXiv:2403.01544 [math.PR], 2024.
- [11] I. Benjamini, A. Nachmias, and Y. Peres. Is the critical percolation probability local? *Probability Theory and Related Fields*, 149(1):261–269, 2011.
- [12] I. Benjamini and O. Schramm. Recurrence of distributional limits of finite planar graphs. *Electronic Journal of Probability*, 6:1–13, 2001.
- [13] N. Berger, C. Borgs, J. T. Chayes, and A. Saberi. Asymptotic behavior and distributional limits of preferential attachment graphs. *The Annals of Probability*, 42(1):1–40, 2014.
- [14] S. Bhamidi, S. N. Evans, and A. Sen. Spectra of large random trees. *Journal of Theoretical Probability*, 25(3):613–654, 2012.
- [15] S. Bhamidi and S. Sen. Geometry of the minimal spanning tree in the heavy-tailed regime: New universality classes. *Probability Theory and Related Fields*, 188(3–4):729–804, 2024.
- [16] P. Billingsley. *Convergence of Probability Measures*. Wiley Series in Probability and Statistics. John Wiley & Sons, Inc., second edition, 1999.
- [17] B. Bollobás, S. Janson, and O. Riordan. The phase transition in inhomogeneous random graphs. *Random Structures and Algorithms*, 31(1):3–122, 2007.
- [18] C. Bordenave, M. Lelarge, and J. Salez. The rank of diluted random graphs. *The Annals of Probability*, 39(3):1097–1121, 2011.
- [19] L. A. Braunstein, S. V. Buldyrev, R. Cohen, S. Havlin, and H. E. Stanley. Optimal paths in disordered complex networks. *Physical Review Letters*, 91(16):1–4, 2003.
- [20] A. Coja-Oghlan, C. Efthymiou, and N. Jaafari. Local convergence of random graph colorings. *Combinatorica*, 38(2):341–380, 2018.
- [21] D. Contreras, S. Martineau, and V. Tassion. Locality of percolation for graphs with polynomial growth. *Electronic Communications in Probability*, 28(1):1–9, 2023.
- [22] A. Dembo and A. Montanari. Ising models on locally tree-like graphs. *The Annals of Applied Probability*, 20(2):565–592, 2010.
- [23] A. Dembo, A. Montanari, and N. Sun. Factor models on locally tree-like graphs. *The Annals of Probability*, 41(6):4162–4213, 2013.

- [24] R. Dobrin and P. M. Duxbury. Minimum spanning trees on random networks. *Physical Review Letters*, 86(22):5076–5079, 2001.
- [25] S. Dommers, C. Giardinà, and R. van der Hofstad. Ising critical exponents on random trees and graphs. *Communications in Mathematical Physics*, 328(1):355–395, 2014.
- [26] L. Dort and E. Jacob. Local weak limit of dynamical inhomogeneous random graphs. Preprint available at arXiv:2303.17437v1 [math.PR], 2023.
- [27] P. Easo and T. Hutchcroft. The critical percolation probability is local. Preprint available at arXiv:2310.10983 [math.PR], 2023.
- [28] E. O. Endo and D. Valesin. Local limits of spatial Gibbs random graphs. *Alea (Rio de Janeiro)*, 17(1):51–63, 2020.
- [29] A. M. Frieze. On the value of a random minimum spanning tree problem. *Discrete Applied Mathematics*, 10(1):47–56, jan 1985.
- [30] A. Garavaglia, R. S. Hazra, R. van der Hofstad, and R. Ray. Universality of the local limit of preferential attachment models. Preprint available at arXiv:2212.05551v2 [math.PR], 2023.
- [31] A. Garavaglia, R. van der Hofstad, and N. Litvak. Local weak convergence for PageRank. *The Annals of Applied Probability*, 30(1):40–79, 2020.
- [32] G. Grimmett. *Percolation*. Springer, New-York, second edition, 1999.
- [33] R. S. Hazra, R. van der Hofstad, and R. Ray. Percolation on preferential attachment models. Preprint available at arXiv:2312.14085, 2023.
- [34] M. Heydenreich and R. van der Hofstad. *Progress in high-dimensional percolation and random graphs. CRM Short Courses Series*. Springer, Cham, (2017).
- [35] R. van der Hofstad. The giant in random graphs is almost local. Preprint available at arXiv:2103.11733v2 [math.PR], 2021.
- [36] R. van der Hofstad. *Random Graphs and Complex Networks, Volume 2*. Cambridge University Press, 2024.
- [37] R. van der Hofstad and M. Pandey. Are giants in random digraphs ‘almost’ local? Preprint available at arXiv:2403.02137 [math.PR], 2024.
- [38] R. van der Hofstad and F. Redig. Maximal clusters in non-critical percolation and related models. *J. Statist. Phys.*, **122**(4):671–703, 2006.
- [39] R. van der Hofstad, P. van der Hoorn, and N. Maitra. Scaling of the clustering function in spatial inhomogeneous random graphs. *Journal of Statistical Physics*, 190(110):1–43, 2023.
- [40] R. Huo and R. Durrett. Latent voter model on locally tree-like random graphs. *Stochastic Processes and their Applications*, 128(5):1590–1614, 2018.
- [41] A. Khan, A. A. Aesha, and J. Sarker. A new algorithmic approach to finding minimum spanning tree. *4th International Conference on Electrical Engineering and Information and Communication Technology, iCEEiCT 2018*, pages 590–594, 2019.
- [42] J. B. Kruskal. On the shortest spanning subtree of a graph and the traveling salesman problem. *Proceedings of the American Mathematical Society*, 7:48–50, 1956.
- [43] V. Kurauskas. On local weak limit and subgraph counts for sparse random graphs. *Journal of Applied Probability*, 59(3):755–776, 2022.
- [44] D. Lacker, K. Ramanan, and R. Wu. Local weak convergence for sparse networks of interacting processes. *The Annals of Applied Probability*, 33(2):843 – 888, 2023.
- [45] R. Lyons. Asymptotic enumeration of spanning trees. *Combinatorics Probability and Computing*, 14(4):491–522, 2005.
- [46] M. Michelen, R. Pemantle, and J. Rosenberg. Quenched survival of Bernoulli percolation on Galton–Watson trees. *Journal of Statistical Physics*, 181(4):1323–1364, 2020.
- [47] M. Milewska, R. van der Hofstad, and B. Zwart. Dynamic random intersection graph: Dynamic local convergence and giant structure. *Random Structures & Algorithms*, 66(1):e21264, 2025.
- [48] M. Milewska, R. van der Hofstad, and B. Zwart. Sir on locally converging dynamic random graphs. Preprint available at arXiv:2501.09623 [math.PR], 2025.
- [49] A. Montanari, E. Mossel, and A. Sly. The weak limit of Ising models on locally tree-like graphs. *Probability Theory and Related Fields*, 152:31–51, 2012.
- [50] T. Mountford, D. Valesin, and Q. Yao. Metastable densities for the contact process on power law random graphs. *Electronic Journal of Probability*, 18(11201150), 2013.
- [51] A. Nachmias and P. Tang. The wired minimal spanning forest on the poisson-weighted infinite tree. *The Annals of Applied Probability*, 34(2):2415–2446, 2024.

- [52] J. Nešetřil, E. Milková, and H. Nešetřilová. Otakar Borůvka on minimum spanning tree problem: Translation of both the 1926 papers, comments, history. *Discrete Mathematics*, 233(1-3):3–36, 2001.
- [53] S. Pettie and V. Ramachandran. An optimal minimum spanning tree algorithm. *Journal of the ACM*, 49(1):16–34, 2002.
- [54] R. C. Prim. Shortest connection networks and some generalizations. *Bell System Technical Journal*, 36(6):1389–1401, 1957.
- [55] J. Salez. Weighted enumeration of spanning subgraphs in locally tree-like graphs. *Random Structures and Algorithms*, 43(3):377–397, 2013.
- [56] J. M. Steele. Growth rates of Euclidean minimal spanning trees with power weighted edges. *The Annals of Probability*, 16(4):347–370, oct 1988.
- [57] J. M. Steele. Minimal spanning trees for graphs with random edge lengths. In *Mathematics and Computer Science II*, pages 223–245. Birkhäuser Basel, Basel, 2002.
- [58] B. Stuffer. Local convergence of random planar graphs. *Journal of the European Mathematical Society*, 25(1):1–73, 2021.
- [59] Q. Yu and Y. Polyanskiy. Ising model on locally tree-like graphs: Uniqueness of solutions to cavity equations. *IEEE Transactions on Information Theory*, 70(3):1913–1938, 2023.