# Friendly frogs, stable marriage, and the magic of invariance

Maria Deijfen, Alexander E. Holroyd, James B. Martin

# Friendly frogs, stable marriage, and the magic of invariance

Maria Deijfen[1]    Alexander E. Holroyd[2]    James B. Martin[3]

11 April 2016

### Abstract

We introduce a two-player game involving two tokens located at points of a fixed set. The players take turns to move a token to an unoccupied point in such a way that the distance between the two tokens is decreased. Optimal strategies for this game and its variants are intimately tied to Gale-Shapley stable marriage. We focus particularly on the case of random infinite sets, where we use invariance, ergodicity, mass transport, and deletion-tolerance to determine game outcomes.

*Keywords:* Combinatorial game; random game; stable marriage; Poisson process.

AMS 2010 Subject Classification: 91A46; 60D05; 60G55

## 1   Friendly Frogs

Here is a simple two-player game, which we call **friendly frogs**. A pond contains several lily pads. (Their locations form a finite set $L$ of points in Euclidean space $\mathbb{R}^d$). There are two frogs. The first player, **Alice**, chooses a lily pad and places a frog on it. The second player, **Bob**, then places a second frog on a distinct lily pad. The players then take turns to move, starting with Alice. A move consists of jumping either frog to another lily pad, in such a way that the distance between the two frogs is strictly decreased, but they are not allowed to occupy the same lily pad. (The frogs are friends, so do not like to be moved further apart, but a lily pad is not large enough to support them both.) A player who cannot move **loses** the game (and the other player **wins**). See Figure 1 for an example game.

We are interested in optimal play. A **strategy** for a player is a map that assigns a legal move (if one exists) to each position, and a **winning** strategy is one that results in a win for that player whatever strategy the other player uses. (In friendly frogs, a position consists of the locations of 0, 1 or 2 frogs.) If there exists a winning strategy for a player, we say that the game is a **win** for that player (and a **loss** for the other player).

Since there are only finitely many possible positions, and the distance between the frogs decreases on each move, the game must end after a finite number of moves. Consequently, for any set $L$, the game is a win for exactly one player. Is it Alice or Bob? Surprisingly, the answer depends only on the size of $L$.

[1]Department of Mathematics, Stockholm University; mia@math.su.se
[2]Microsoft Research, Redmond, USA; holroyd@microsoft.com
[3]Department of Statistics, University of Oxford, UK; martin@stats.ox.ac.uk
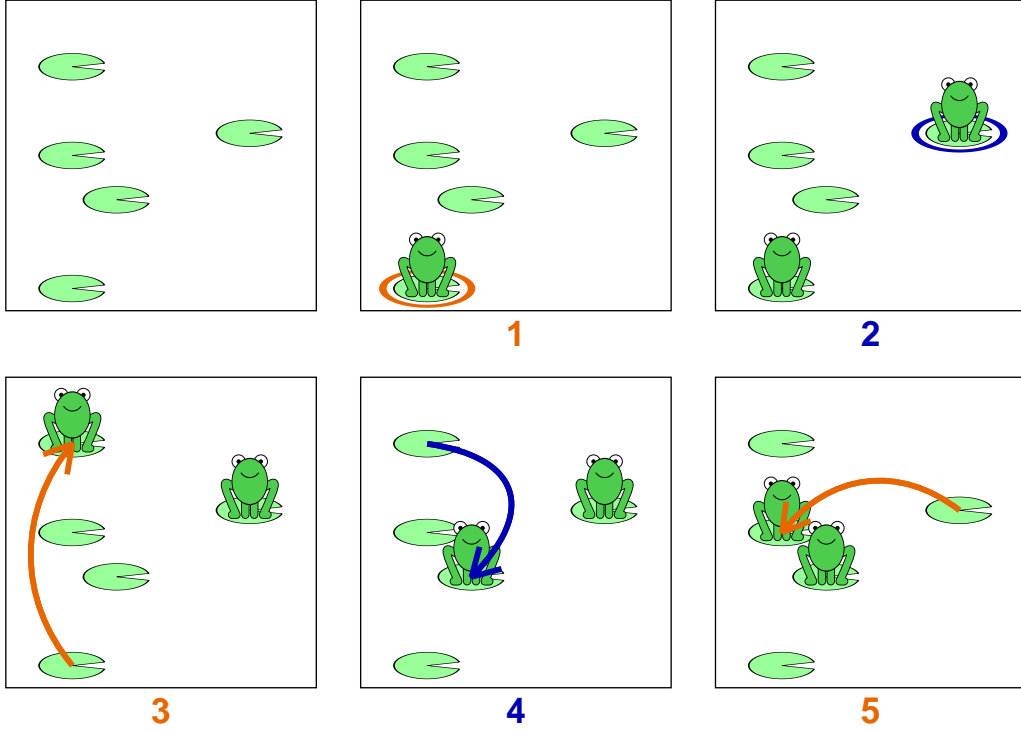
Figure 1: A game of friendly frogs on a set $L$ of size 5. Alice starts. Alice's moves are shown in amber, Bob's in blue. After move 5, Bob has no legal move, so Alice wins.

**Theorem 1.** *Consider friendly frogs played on a finite set $L \subset \mathbb{R}^d$ of size $n$ in which all pairs of points have distinct distances. The game is a win for Alice if $n$ is odd, and a win for Bob is $n$ is even.*

*Proof.* Let $M$ be the set of all unordered pairs $\{x, y\}$ in $L$ such that the game started with two frogs at $x$ and $y$ is a loss for the next player. The key ingredient is a simple algorithm that identifies $M$. (We postpone consideration of the two opening moves, in which the frogs are placed). In fact $M$ will form a partial matching on $L$. We construct this matching iteratively as follows. The idea is to work backwards from positions where the outcome is known. Order the set of all $\binom{n}{2}$ pairs in $L$ in increasing order of distance between the pair. Then for each pair in turn, match the two points to each other if and only if neither is already matched. The algorithm ends with at most one point not matched. See Figure 2 for an example.

To show that this $M$ has the claimed property for the game, we need to check that from any position in $M$, it is impossible to move to another position in $M$, while from a position not in $M$, it is possible to move to a position in $M$. The former is immediate because $M$ is a partial matching (and a move consists of moving only one frog). For the latter, suppose the frogs are located at $x$ and $y$, and that $x$ and $y$ are not matched to each other. Since $x$ and $y$ were not matched by the algorithm, at least one of them was matched to a closer point; without loss of generality, $x$ is matched to $w$, where $|x - w| < |x - y|$. (Here and subsequently, $|\cdot|$ is the Euclidean norm on $\mathbb{R}^d$.) Therefore we can move a frog from $y$ to $w$.

If $n$ is odd then there is exactly one point that is not matched, so Alice wins by placing the first frog there; wherever Bob places the second frog, the two frog locations are not matched to each other. If $n$ is even then the matching $M$ is perfect (i.e. every point is matched). Therefore, wherever Alice places the first frog, Bob wins by placing the second on its partner in $M$. $\qquad\square$
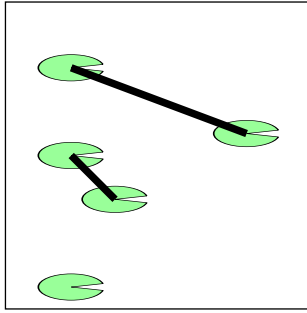
Figure 2: The matching $M$ of the set $L$ in Figure 1.

We will consider various extensions of the friendly frogs game, including versions where frogs and/or points are player-specific (available only to one player), where certain moves are forbidden, and where different winning criteria apply. Notwithstanding the humble beginning of Theorem 1, we will be led into some very intriguing waters. For concreteness we will focus throughout on points in $\mathbb{R}^d$, although many arguments carry over to more general metric spaces (and, for instance, the above proof extends even to any injective symmetric distance function on $L$). We will continue to assume that all inter-point distances are distinct. (Relaxing this assumption is also quite natural, but we choose instead to pursue other directions.)

Matters become particularly interesting when we allow the set of points (lily pads) $L$ to be *infinite*, and especially a *random* countable set. The "losing" two-frog positions will still form a matching, and this matching is most naturally interpreted as a version of the celebrated *stable marriage* of Gale and Shapley, the topic of the 2012 Nobel prize in economics (awarded to Shapley and Roth). We will make crucial use of *invariance* of the probability distribution of $L$ under symmetries of $\mathbb{R}^d$. This powerful tool permits remarkably simple and elegant proofs of facts apparently not amenable to other arguments. In games involving points of several types, we will see an example of a *phase transition*, as well as a situation in which existence of a phase transition is an open question. We will also analyze play of simultaneous games by making a connection to the remarkable theory of Sprague-Grundy values (or "nimbers").

The article contains a mixture of original research and expository material. We use the friendly frogs game partly as a vehicle to showcase some beautiful known ideas, and we assume a minimum of technical background. The game and its analysis are novel, so far as we know. Stable marriage [10] and its variants have been extensively studied, but the connection to games appears to be new. Many of the results that we use on matchings of random point sets are taken from [17]. We will review the necessary background and give proofs where appropriate. The general theory of combinatorial games is highly developed (see e.g. [3]). We will explain the relevant parts of the theory as they apply in our context. Other recent work on games in random settings appears for example in [1, 14, 15] and the review [19]. In a different direction, certain games in infinite spaces have intimate connections with general topology [21].

## 2 Infinite point sets

Theorem 1 shows that the outcome of friendly frogs on a finite set $L$ is determined solely by the parity of the number of points (lily pads). What happens when $L$ is infinite? Is $\infty$ odd or even? The answer now depends on the choice of set; we will focus especially on the behaviour of *typical* (i.e. random) infinite sets.

Let $L$ be an infinite subset of $\mathbb{R}^d$. As before, we assume that all distances between pairs of points in $L$ are distinct. We call a sequence of points $x_1, x_2, \ldots$ a **descending chain** if
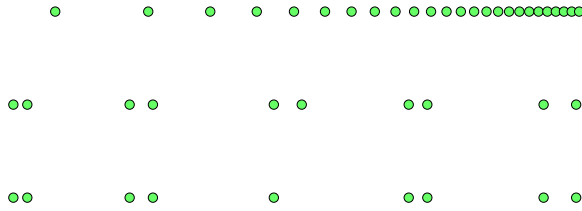
Figure 3: Examples of infinite sets $L \subset \mathbb{R}$. Top: play continues forever, owing to an infinite descending chain. Middle: Bob wins. Bottom: Alice wins.

the distances $(|x_i - x_{i+1}|)_{i \geq 1}$ form a strictly decreasing sequence. If there exists an infinite descending chain $x_1, x_2, \cdots \in L$, then it is possible for the game to last forever. See Figure 3. Therefore we make the additional assumption that $L$ has no infinite descending chains. This implies in particular that $L$ is discrete, i.e. any bounded set contains only finitely many points.

It is easy to construct examples of infinite $L$ satisfying the above conditions for which either player wins friendly frogs; see Figure 3. Firstly, in dimension 1, place exactly *two* points in each of the intervals $[3i, 3i + 1]$ for $i \in \mathbb{Z}$. (A simple way to make all inter-point distances distinct is to choose each point uniformly at random in the appropriate interval, independently of all others.) Then Bob wins by placing a frog at the unique point in the same interval as Alice's initial frog. Secondly, suppose the points are as above except that the interval $[0, 1]$ now contains only one point. Then Alice wins by placing the first frog on this point; whichever point Bob chooses for the second frog, Alice can then move the first frog to the "partner" of that point in the appropriate unit interval.

As in the previous section, the key to analyzing the game for general $L$ is to identify those positions from which the game is a loss for the player whose turn it is to move. Following standard conventions of combinatorial game theory (see e.g. [3]), such positions are called **P-positions** to indicate that the [P]revious player wins, while all other positions are called **N-positions**, since the [N]ext player wins. Since terminal positions are P-positions, the P- and N-positions satisfy the following.

(N) From every N-position, there is at least one possible move to a P-position.

(P) From every P-position, every possible move is to an N-position.

Since the game terminates in a finite number of moves, it follows by induction that these properties are sufficient to characterize the P- and N-positions. That is, to check that a claimed partition of the positions into P- and N-positions is correct, it suffices to check that it satisfies (N) and (P).

In many games, characterizing the set of P-positions is a difficult problem requiring experimentation and insight. In contrast, checking via (N) and (P) that such a characterization is correct may be essentially mechanical.

In friendly frogs, the two-frog P-positions are given by a matching. Here is some notation. Let $L \subseteq \mathbb{R}^d$. A **matching** of $L$ is a set $M$ of unordered pairs of distinct points in $L$ such that each point of $L$ is included in at most one pair. The matching is **perfect** if each point is included in exactly one pair. For $x \in L$, we write $M(x)$ for the **partner** of $x$, i.e. the unique point $y$ such that $\{x, y\} \in M$, or, if there is no such $y$, we set $M(x) := \infty$ and say that $x$ is **unmatched**.

As in the case of finite $L$ in the last section, we will construct the relevant matching iteratively. Now, however, there may be no closest pair of points, so we need a local version of the algorithm.

4

The following abstraction will prove very useful. Imagine that each point of $L$ "prefers" to be matched to a partner that is as close as possible. Given a matching $M$ of $L$, a pair of points $x, y \in L$ is called **unstable** if they both strictly prefer each other over their own partners, i.e. if $|x - M(x)|$ and $|y - M(y)|$ both strictly greater than $|x - y|$ (where $|x - M(x)| := \infty$ if $M(x) = \infty$, so that any partner is preferable to being unmatched). A matching $M$ is called **stable** if there are no unstable pairs. Note that any stable matching of $L$ has at most one unmatched point.

Stable matching can be applied to a wide variety of settings involving agents each of which has preferences over the others. The concept was introduced in a celebrated paper of Gale and Shapley [10], who considered the setting of $n$ heterosexual marriages between $n$ girls and $n$ boys, each of whom has an arbitrary preference order over those of the opposite sex. Gale and Shapley gave a beautiful algorithm proving the existence of a stable matching in this case. (They showed however that stable matchings are not necessarily unique, and may not exist in the same-sex "room-mates" variant). As mentioned earlier, the 2012 Nobel prize in economics was awarded on the basis of this and ensuing work, to Lloyd S. Shapley for theoretical advances, and to Alvin E. Roth for practical applications. Our setting differs from the standard Gale-Shapley same-sex matching problem in that the set $L$ is infinite; on the other hand, our preferences are very special, since they are based on distance. This case was studied in [17].

**Proposition 2** ([17]). *Suppose $L \subset \mathbb{R}^d$ has all pairwise distances distinct and has no infinite descending chains. Then there exists a unique stable matching of $L$.*

*Proof.* We will show that the following algorithm leads to a stable matching. First match all mutually closest pairs of points. Then remove them and match all mutually closest pairs in the remaining point set. Repeat indefinitely (i.e. for a countably infinite sequence of stages), and take as the final matching the set of all pairs that are ever matched.

By induction over the stages in the algorithm, every pair that is matched by the algorithm must be matched in any stable matching.

Furthermore, at most one point can be left unmatched by the algorithm. To see this, assume that there are at least two unmatched points. Since there are no descending chains, the set of unmatched points then contains at least one pair of points that are mutually closest in this set and, since $L$ is discrete, this pair must have been mutually closest at some finite stage of the algorithm. However, then they should have been matched to each other, which is a contradiction.

Finally, we need to confirm that the resulting matching is in fact stable. To this end, assume that there exist $x, y \in L$ with $|x - M(x)|$ and $|y - M(y)|$ both strictly greater than $|x - y|$. By the previous argument, at least one of $x$ and $y$ is matched, so consider the earliest stage at which one of them was matched by the algorithm. Since both $x$ and $y$ were unmatched prior to this stage, we obtain a contradiction. $\qquad\square$

**Proposition 3.** *Suppose $L \subset \mathbb{R}^d$ has all pairwise distances distinct and has no infinite descending chains. Let $M$ be the stable matching of $L$ and consider friendly frogs on $L$. The position with the two frogs at $x$ and $y$ is a P-position if and only if $x$ is matched to $y$ in $M$.*

*Proof.* Since $L$ has no infinite descending chains, the game terminates. Therefore, it suffices to check the conditions (N) and (P) above. For (N), if $\{x, y\} \notin M$, then $x$ (or $y$) must have a partner that is closer than $y$ (or respectively $x$), since otherwise $x$ and $y$ would constitute an unstable pair. Without loss of generality, $M(x) = w$ where $|x - w| < |x - y|$, and we can then move a frog from $y$ to $w$. The claim (P) is immediate, since $M$ is a matching. $\qquad\square$

As before, if $M$ has one unmatched point then Alice wins by placing the first frog at that point. If the matching is perfect then Bob wins by placing the second frog at the partner of Alice's initial move. As we have seen, both situations are possible for suitable infinite sets $L$.
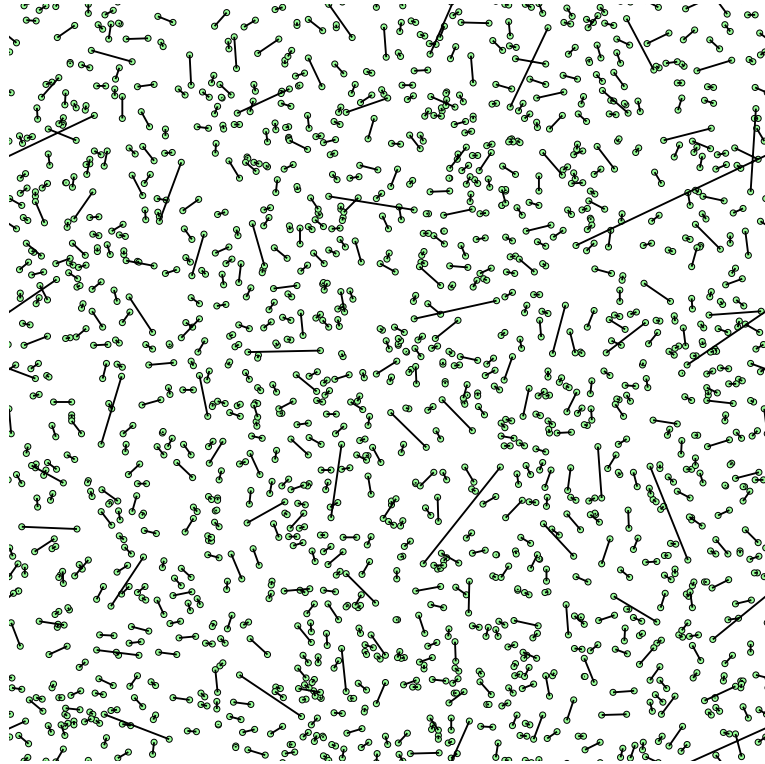
Figure 4: The stable matching of random points on a two-dimensional torus.

## 2.1   Random infinite sets

It is natural to ask what happens for a *typical* infinite set of points. A natural and canonical way to formalize this notion is the Poisson point process, which is defined as follows. Fix $\lambda > 0$. Let any Borel set of finite volume contain a random number of points with a Poisson distribution of mean equal to $\lambda$ times its volume, and let disjoint sets contain independent numbers of points. These conditions characterize the distribution of the set of points, and the resulting random set is called a (homogeneous) **Poisson (point) process** with **intensity** $\lambda$ on $\mathbb{R}^d$. It is a countable infinite set with probability 1. (The Poisson process has other equivalent definitions – for instance it may be constructed as a limit as $n \to \infty$ of $n$ uniformly random points in a ball of volume $n/\lambda$ around the origin, or as a limit as $\epsilon \to 0$ of a grid of cubes of volume $\epsilon$ each of which contains a point with probability $\epsilon\lambda$ independently.) If $L$ is a Poisson process of intensity 1 then $\{\lambda^{1/d}x : x \in L\}$ is a Poisson process of intensity $\lambda$ – the intensity parameter will be unimportant for us until we consider several Poisson processes together. See e.g. [6] for background. It is straightforward to check that with probability 1, all pairs of points have distinct distances, and that there are no descending chains. See e.g. [13] or [5] for proofs. The process is **translation-invariant**, which is to say, its distribution is invariant under the action of any translation of $\mathbb{R}^d$.

**Theorem 4.** *Let $L$ be a Poisson point process on $\mathbb{R}^d$. With probability 1, friendly frogs on $L$ is a win for Bob.*

*Proof.* By Proposition 2, there is a unique stable matching $M$ of $L$. It suffices to check that this matching is perfect with probability 1. The matching has at most one unmatched point.
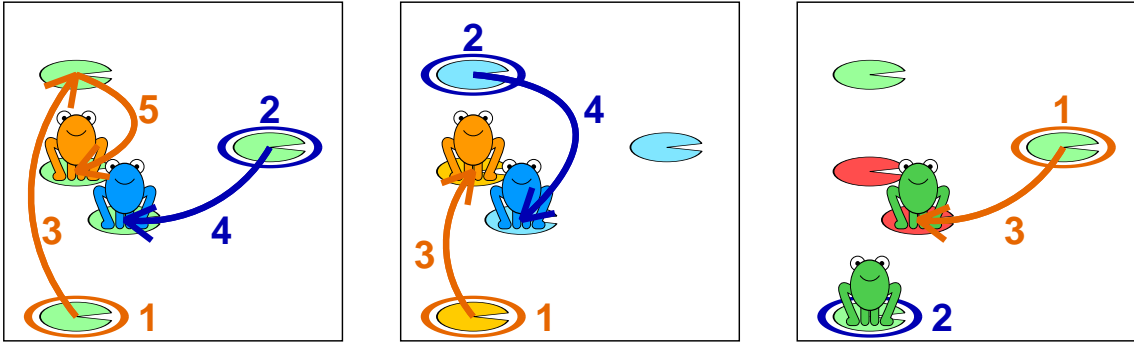
Figure 5: Three variant games: (a) colored friendly frogs, in which each player may only move their own frog; (b) colored friendly frogs on colored points, where in addition a frog may only occupy a point of its own color; (c) fussy frogs, in which the two frogs may not both occupy red points.

But if there is an unmatched point then its location is a translation-invariant random variable on $\mathbb{R}^d$, which is impossible. More precisely, by translation-invariance of the Poisson process and uniqueness of the stable matching, every unit cube in $\mathbb{R}^d$ has equal probability $p$ of containing an unmatched point. We can partition $\mathbb{R}^d$ into unit cubes indexed by $\mathbb{Z}^d$, so the probability that there exists an unmatched point is $\sum_{z \in \mathbb{Z}^d} p$. Since this sum must be finite, $p = 0$, whence the sum is 0. $\qquad\square$

Despite the simplicity of the above proof, there is something subtle and mysterious about the argument. What probability-one property of the Poisson process does it use? In other words, is there some easily described set $\mathcal{A}$ of subsets of $\mathbb{R}^d$ such that (a) the Poisson process lies in $\mathcal{A}$ with probability 1, and (b) Bob wins on any $L \in \mathcal{A}$? We do not know of such a set, except for unsatisfying choices such as $\mathcal{A} = \{L : L \text{ has a perfect stable matching}\}$ or $\mathcal{A} = \{L : \text{ Bob wins}\}$. As we have seen, the set of $L$ with distinct inter-point distances and no descending chains satisfies (a) but not (b). The proof of Theorem 4 uses translation-invariance of the Poisson process in a fundamental way that apparently cannot be easily reduced to such a probability-one property. Many elegant arguments in probability theory involve an appeal to some symmetry or invariance property of this kind. In the next section we will use stronger probabilistic properties of Poisson processes – deletion-tolerance and ergodicity.

In fact, the algorithm in the proof of Proposition 2 leads to a perfect stable matching for a large class of translation-invariant point processes on $\mathbb{R}^d$ – see [17, Proposition 9]. The conclusion of Theorem 4 hence remains valid for this class of processes.

The article [17] is also concerned with the distribution of the distance from a point to its partner in the stable matching. These distances are potentially relevant to issues of computational complexity and length of the game. For instance, if Alice is required to place her first frog within distance $r$ of the origin, how difficult can she make it for Bob to win? We leave these interesting questions for future investigation.

## 3   Colored frogs, colored points

In this section we consider variants of friendly frogs in which frogs and/or lily pads have multiple colors, and the allowed moves are correspondingly restricted. Throughout we take $L$ to be an

infinite set satisfying the assumptions of Proposition 2.

## 3.1 Colored frogs

First we introduce the **colored friendly frogs** game. Here, Alice starts by placing an **amber** frog on some point of $L$, then Bob places a **blue** frog on a different point. Subsequently, the game proceeds exactly as before, except that Alice may only move the amber frog, and Bob may only move the blue frog. As before, a player who cannot move loses.

A two-frog position can now be specified by an *ordered* pair $(x, y)$, where $x$ is the location of the frog of the previous player to move, and $y$ the location of the frog of the next player.

Rather than requiring an entirely new analysis, it turns out that the P-positions can again be described in terms of the stable matching $M$ of $L$. If $|x - y| \leq |x - M(x)|$ then we say that $x$ **desires** $y$. (This terminology is natural given the interpretation of preferences described earlier.) Note the use of the weak inequality $\leq$, so that a point desires its own partner. Here is the analogue of Proposition 3 for colored friendly frogs.

**Proposition 5.** *Suppose $L \subset \mathbb{R}^d$ has all pairwise distances distinct and has no infinite descending chains. Let $M$ be the stable matching of $L$, and consider the colored friendly frogs game on $L$. The position $(x, y)$ is a P-position if and only if $x$ desires $y$.*

*Proof.* Again it suffices to check the conditions (N) and (P). For (N), if $|x - y| > |x - M(x)|$, then the frog at $y$ can be moved to $M(x)$. On the other hand, for (P), if $|x - y| \leq |x - M(x)|$, then there cannot exist $z \in L$ with $|x - z| < |x - y|$ and $|x - z| \leq |z - M(z)|$, since in that case $x$ and $z$ would constitute an unstable pair. Hence moving the frog at $y$ must result in a position $(z, x)$ with $|x - z| > |z - M(z)|$. $\square$

Note that an unmatched point in the stable matching is not desired by any other point, since that pair would be unstable. Hence, if the stable matching of $L$ has one unmatched point, then Alice wins colored friendly frogs by placing her amber frog at the unmatched point. If the matching is perfect, then Bob wins, e.g. by placing his blue frog at the partner of Alice's initial point. The outcome in hence the same as in the original friendly frogs game. In particular, we have the following.

**Corollary 6.** *Let $L$ be a Poisson process on $\mathbb{R}^d$. With probability $1$, colored friendly frogs on $L$ is a win for Bob.*

Indeed, Bob may use the same strategy in the colored and uncolored games, always moving to a matched pair. Does this mean that the games are essentially identical? No. To highlight an interesting difference, let us modify the rules in a way that favors Alice. In **shy** friendly frogs, we fix a constant $c > 0$, and stipulate that Bob, on his opening move, cannot place the second frog within distance $c$ of the first frog. (But we place no such restriction on subsequent moves.) Shy colored friendly frogs is defined analogously. Surprisingly, the outcome now differs between the two variants; the proof will employ an interesting probabilistic argument.

**Theorem 7.** *Let $L$ be a Poisson process on $\mathbb{R}^d$, and fix $c > 0$. With probability $1$:*

(i) *shy friendly frogs is a win for Alice;*

(ii) *shy colored friendly frogs is a win for Bob.*

*Proof.* For (i), Alice places the first frog on any point $x$ whose partner $M(x)$ is at least distance $c$ away. Such a point exists, since the stable matching $M$ is perfect, but the Poisson process has points whose nearest neighbor is at least distance $c$ away.

Turning to (ii), we claim that with probability 1, every point is desired by infinitely many others. This implies in particular that whatever Alice's opening move $x$, there exists a point $y$ with $|x - y| > c$ that desires $x$, so Bob wins by placing his frog there, by Proposition 5. The claim follows from [7, Theorem 1.3 (i)]. Since the proof in our case is short, we include it.

Let $X$ be the (random) point of $L$ closest to the origin. It suffices to show that infinitely many points desire $X$. Let $D$ be the set of points that desire $X$. Modify the set $L$ as follows. Whenever $D$ is finite, delete all points of $D$ and their partners, except for $X$ itself (which is the partner of a point in $D$). It is easy to check that the stable matching of the modified set is simply the restriction of $M$ to the points that remain. In particular, if $D$ was finite then $X$ is now unmatched. However, the Poisson process is **deletion-tolerant**, which is to say: deleting any finite set of points, even in a way that depends on the process, results in a point process whose distribution is absolutely continuous with respect to the original distribution. (See e.g. [17, Lemma 18] or [18].) That is, the deletion cannot cause any event of zero probability to have positive probability. (Intuitively, the picture after deletion is still plausible.) Since the stable matching of the Poisson process is perfect with probability 1, we deduce that $D$ was infinite with probability 1. □

## 3.2 Colored points

There is a further natural variant of colored friendly frogs in which the two frogs are restricted to different point sets. Let $L_A$ and $L_B$ be two disjoint subsets of $\mathbb{R}^d$ whose union satisfies the assumptions of Proposition 2. We refer to points of $L_A$ and $L_B$ as amber and blue, respectively. We stipulate that Alice's amber frog can only occupy an amber point, and Bob's blue frog can only occupy a blue point. Otherwise the rules are as for colored friendly frogs. We call this game **colored friendly frogs on colored points**. The P-positions in this case are given by a two-color variant of stable matching.

A two-color matching of $(L_A, L_B)$ is a set $M$ of pairs of points $(x, y) \in L_A \times L_B$ such that each point is contained in at most one pair. As in the one-color case, the matching is perfect if each point of $L_A \cup L_B$ is included in a pair. A two-color matching $M$ of $(L_A, L_B)$ is **stable** if and only if there do not exist $x \in L_A$ and $y \in L_B$ with $|x - M(x)|$ and $|y - M(y)|$ both strictly greater than $|x - y|$.

Proposition 2 and Proposition 5 remain true for this game, with $L$ replaced by $(L_A, L_B)$, "stable matching" replaced by "stable two-color matching", and a revised definition of desire under which a point can only desire a point of the opposite color (see [17] for more detail). The same proofs apply with only minor adjustments. Specifically, in the algorithm described in the proof of Proposition 2, points of the same color cannot be matched to each other. Therefore, instead of leaving at most one point unmatched, it follows from the same arguments that all unmatched points must be of the same color.

Note that an unmatched point desires all points of the other color, and an unmatched point cannot be desired by any point of the other color, since they would be an unstable pair. If the two-color stable matching has unmatched amber points, then Alice wins by placing her frog at one of these points. If not, Bob wins by placing his frog on an unmatched blue point (if one exists), or on the partner of Alice's opening move.

**Theorem 8.** *Let $L_A$ and $L_B$ be two independent Poisson processes on $\mathbb{R}^d$, with respective intensities $\alpha$ and $\beta$. Consider colored friendly frogs with colored points on $(L_A, L_B)$. The game*

*is a win for Bob if $\alpha \leq \beta$, and a win for Alice if $\alpha > \beta$.*

The probabilistic setup of Theorem 8 is equivalent to that of a single Poisson process of intensity $\alpha + \beta$ in which each point is independently declared amber or blue with respective probabilities $\alpha/(\alpha + \beta)$ and $\beta/(\alpha + \beta)$. (See e.g. [6].) The conclusion of Theorem 8 is an example of a *phase transition*: an abrupt qualitative change of behavior as a parameter crosses a critical value.

To prove Theorem 8, we need a property that is stronger than translation invariance. A point process is said to be **ergodic** if every event that is invariant under translations has probability 0 or 1. For example, the event that there is no point within distance 1 of the origin is not translation-invariant, but the event that there are infinitely many disjoint balls of radius 1 that contain no points is translation-invariant. A Poisson process is ergodic (and so is the two-color process made up of two independent Poisson processes) – this can be deduced using the independence of the process on disjoint subsets of the space. (See e.g. [6].)

*Proof of Theorem 8.* First let us consider the case $\alpha = \beta$. The set of unmatched points in the stable matching is either empty, or consists only of amber points or only of blue points. Applying ergodicity, one of these three events must have probability 1, and the others probability 0. But by symmetry the probabilities of unmatched amber points and of unmatched blue points must be equal. Hence they are both 0, and with probability 1 the matching is perfect, giving a win for Bob.

When the two intensities are different, it is natural to expect that we cannot match amber points to blue points in a translation-invariant way without leaving some of the higher-intensity set unmatched. Making this intuition rigorous may at first appear tricky. We might compare the numbers of points in a large ball, but perhaps many points have their partners outside the ball. And where should we use translation-invariance? Since $L_A$ and $L_B$ are countable infinite sets, there certainly exists *some* perfect matching between them.

In fact, there is a clean solution, using a simple but powerful tool, the *mass transport principle*. (See [2, 12] for background.) Consider any function $f : \mathbb{Z}^d \times \mathbb{Z}^d \to [0, \infty]$ that is translation-invariant in the sense that $f(s,t) = f(s+u, t+u)$ for all $s, t, u \in \mathbb{Z}^d$. Then note that $\sum_{t \in \mathbb{Z}^d} f(0,t) = \sum_{t \in \mathbb{Z}^d} f(-t, 0) = \sum_{s \in \mathbb{Z}^d} f(s, 0)$. It is sometimes helpful to think of $f(s,t)$ as the mass sent from $s$ to $t$.

Now suppose $\alpha < \beta$. For $s \in \mathbb{Z}^d$, let $Q_s$ be the unit cube $s + [0,1)^d$ in $\mathbb{R}^d$. Define $f(s,t)$ to be the expected number of amber points in $Q_s$ that are matched to blue points in $Q_t$. This $f$ is translation-invariant in the sense of the previous paragraph, because of translation-invariance of the Poisson processes. Thus, $\sum_s f(s, 0)$, which is the expected number of matched blue points in $Q_0$, is equal to $\sum_t f(0, t)$, which is the expected number of matched amber points in $Q_0$. The latter is at most $\alpha$, the expected total number of amber points in $Q_0$. But the expected number of blue points in $Q_0$ is $\beta$, so the expected number of unmatched blue points in $Q_0$ is at least $\beta - \alpha$. In particular, the probability that there exists an unmatched blue point is positive. Applying ergodicity again shows that this probability is therefore 1. Thus Bob wins.

Similarly, if $\alpha > \beta$ then with probability 1 there are unmatched amber points, leading to a win for Alice. □

Once again, the above proof uses invariance and ergodicity in a subtle and fundamental way that cannot easily be reduced to probability 1 properties of the point process. What property of $(L_A, L_B)$ guarantees Bob wins when $\alpha = \beta$? It is not that $L_A$ and $L_B$ have equal asymptotic density. Modifying the example in Figure 3, that holds if $L_A$ consists of one point in every

interval $[3i, 3i + 1]$ for $i \in \mathbb{Z}$ while $L_B$ has one point in each such interval except $[0, 1]$. But here Alice wins.

Again, the conclusion of Theorem 8 remains valid for a large class of translation-invariant point processes; see [17] for details of the corresponding results for stable matchings.

## 3.3 Fussy Frogs

Despite the relatively complete analysis in the last two cases, we need not go far to reach an unsolved problem. In **fussy friendly frogs**, the points again have two colors, now **green** and **red**, denoted by sets $L$ and $L_R$ respectively. The rules are as in the original friendly frogs game (in particular, the two frogs are once again identical and can be moved by either player), except that it is not permitted that *both* frogs simultaneously occupy red points.

**Open Problem.** *Let $L$ and $L_R$ be independent Poisson processes on $\mathbb{R}^d$ with respective intensities $1$ and $\rho$. Do there exist $d \geq 1$ and $\rho > 0$ for which Bob wins fussy friendly frogs with positive probability?*

Fussy friendly frogs again has an associated matching, the analogue of stable matching under the restriction that red points cannot be matched to each other. This matching can be constructed iteratively as in the proof of Proposition 2, and Bob wins if and only if it is perfect. Ergodicity shows that this has probability 0 or 1 for each $\rho$ and $d$. When $\rho > 1$ (and even when $\rho > 1 - \epsilon$ for some $\epsilon = \epsilon(d) > 0$), it is not difficult to show that there are unmatched red points (so Alice wins); the question is whether this holds for every positive $\rho$. This is not known for any dimension $d$, although in [16] it is proved that for any fixed $\rho > 0$, there exists $d_0 = d_0(\rho)$ such that there are unmatched red points for all $d \geq d_0$.

# 4 Variations on a theme

In this section we consider some further variant games, in which the rules are modified in more fundamental ways.

## 4.1 Playing to Lose

We consider a **misère** version of friendly frogs. In general, a game is said to be played under misère rules if the legal moves are the same, but a player who cannot move now *wins* the game instead of losing it. This means that a player tries to avoid moving to positions where the next player cannot move. Specifically, in misère friendly frogs, a player wants to avoid having to move to a mutually closest pair.

Let $L$ satisfy the assumptions of Proposition 2. The P-positions in the misère game are given by a variant of the stable matching of $L$ with the added restriction that mutually closest points cannot be matched. A matching $\widetilde{M}$ of $L$ is said to be stable subject to this restriction if there do not exist $x, y \in L$ that are not mutually closest and with $|x - \widetilde{M}(x)|$ and $|y - \widetilde{M}(y)|$ both strictly greater than $|x - y|$. The unique matching with this property is obtained by the following modification of the iterative procedure used to construct the unrestricted stable matching. Call $x$ and $y$ potential partners of each other if they are both unmatched and they are not mutually closest points of $L$; then match all pairs $x$ and $y$ that are each others' mutually closest potential partner. Repeat indefinitely. The resulting matching has at most two unmatched points (and if there are two such points, they must be mutually closest points of $L$).

**Proposition 9.** *Let $L \subset \mathbb{R}^d$ have distinct distances and no infinite descending chains. Let $\widetilde{M}$ the stable matching of $L$ subject to the restriction that mutually nearest neighbors cannot be matched. In misère friendly frogs, the position with two frogs at $x$ and $y$ is a P-position if and only if $x$ is matched to $y$ in $\widetilde{M}$.*

*Proof.* With misère rules, all terminal positions are N-positions, and the characterization of N-positions and P-positions is modified by replacing condition (N) with:

(N') From every N-position that is not terminal, there is at least one move to a P-position.

Assume that $x$ and $y$ are not mutually closest and are not matched in $\widetilde{M}$ (so that they hence define an N-position that is not terminal). If both $|x - \widetilde{M}(x)| > |x - y|$ and $|y - \widetilde{M}(y)| > |x - y|$, then $x$ and $y$ would constitute an unstable pair in $M$. Hence either the frog at $x$ could be moved to $\widetilde{M}(y)$, or the frog at $y$ could be moved to $\widetilde{M}(x)$. The property (P) follows since $\widetilde{M}$ is a matching. $\square$

This argument shows that the misére friendly frogs is a win for Alice if and only if the restricted stable matching has exactly one unmatched point.

**Corollary 10.** *Let $L$ be a Poisson process on $\mathbb{R}^d$. Misère friendly frogs is a win for Bob with probability $1$.*

*Proof.* The argument in the proof of Theorem 4 shows that the matching $\widetilde{M}$ is perfect with probability $1$ – it is impossible for the unmatched points to form a non-empty finite translation-invariant random set. $\square$

## 4.2 Blocking and multi-matching

The games can be modified by allowing moves to be blocked. Consider colored friendly frogs, but suppose that in addition to the two frogs, there are $k$ **stones**. After moving or placing their frog, a player then places the stones on any $k$ points (lily pads). The other player is then forbidden from moving their frog to any of those $k$ points on the next move. (Equivalently, we can imagine that the next player tries to make a move, but the previous player can reject it and request that they try a different move, up to $k$ times. The chess variants *compromise chess* and *refusal chess* are similar; see e.g. [22].) The rules are otherwise as in colored friendly frogs. A player loses if they cannot move, perhaps because all possible moves are blocked by stones. We call this game $k$-**stone colored friendly frogs**.

The P-positions are related to stable multi-matchings, which were introduced and studied in [7, 8]. Let $L$ be an infinite set satisfying the assumptions of Proposition 2. Let $m \geq 1$. An $m$-**multi-matching** or $m$-**matching** is defined analogously to a matching, except that each point may be matched to up to $m$ other points. The $m$-matching is **perfect** if each point is matched to exactly $m$ points. For an $m$-matching of $L$, let $D(x)$ denote the distance to the most distant partner of $x$, with $D(x) = \infty$ if $x$ has strictly fewer than $m$ partners. The matching is **stable** if and only if there do not exist $x, y \in L$ that are not matched to each other with $D(x)$ and $D(y)$ both strictly greater than $|x - y|$. A pair of points violating this is called unstable. A point $x$ **desires** $y$ if $|x - y| \leq D(x)$.

Proposition 2 extends to stable $m$-matchings. The following modification of the iterative procedure in its proof leads to the unique stable $m$-matching of $L$. Call two points potential partners if they are not already matched to each other and if neither is already matched to $m$ other points. Match all mutually closest potential partners. Repeat indefinitely. See Figure 6 (left) for an example.
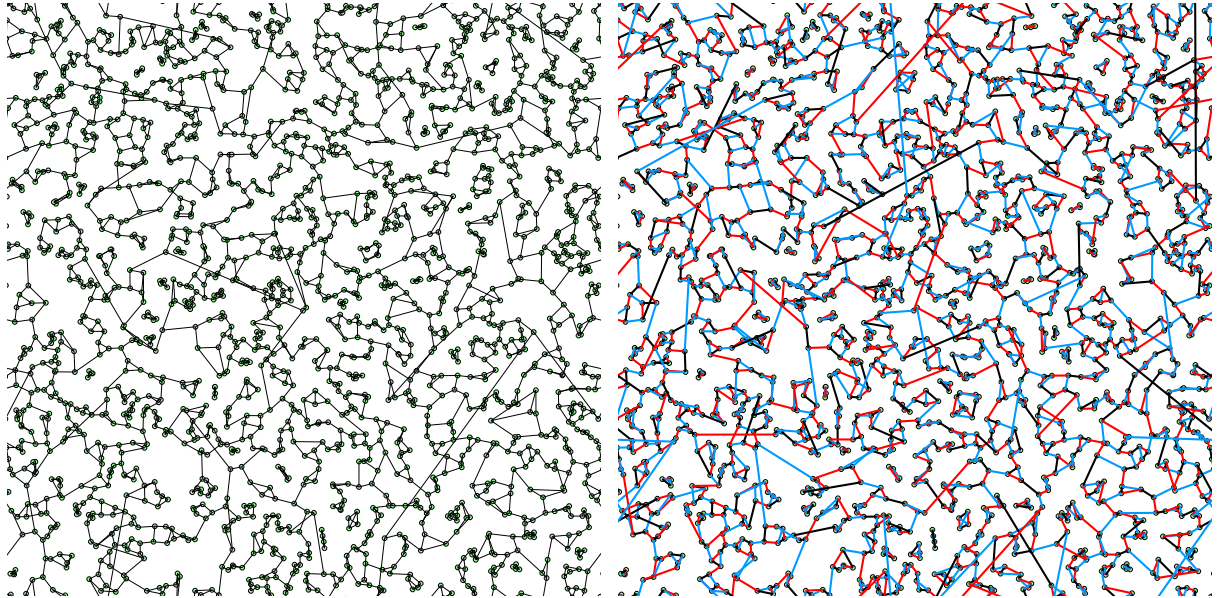
Figure 6: Left: the stable 3-multi-matching of random points in a torus. Right: pairs having friendly frogs Sprague-Grundy values 0 (black), 1 (red), and 2 (blue), for the same points.

We remark that in the stable $m$-matching there may be more than one point that has strictly fewer than $m$ partners, but there cannot be more than $m$ of them (otherwise there would be two that are not matched to each other).

**Proposition 11.** *Let $L \subset \mathbb{R}^d$ have distinct distances and no infinite descending chains, and assume that the stable m-matching of $L$ is perfect. Consider k-stone colored friendly frogs. Suppose that the frogs are at $x$ and $y$, with $y$ being the frog of the next player. This position is a P-position if and only if, in the stable $(k + 1)$-matching, $x$ desires $y$, and all partners of $x$ that are closer than $y$ are blocked by stones.*

*Proof.* We check (N) and (P). For (P), suppose the given conditions hold. Since $x$ desires $y$ we have $D(x) \geq |x - y|$. Thus, if the next player moves their frog from $y$ to $z$, then $x$ also desires $z$. But $z$ is not a partner of $x$, because we assumed that all possible such $z$ are blocked. Therefore $z$ does not desire $x$ (otherwise they would be unstable), so the new position is an N-position (regardless of where the player moves the stones). We now check (N). If $x$ desires $y$ but some closer partner $z$ of $x$ is not blocked, then the next player can move to $z$. On the other hand, if $x$ does not desire $y$ then all the partners of $x$ are closer than $y$, and at least one of them, $z$ say, is not blocked, so the next player moves there. In either case, this player then blocks all $k - 1$ partners of $z$ other than $y$. $\square$

**Theorem 12.** *Let $L$ be a Poisson process on $\mathbb{R}^d$ and let $k \geq 1$. With probability 1, k-stone colored friendly frogs on $L$ is a win for Bob.*

*Proof.* We claim that the stable $(k+1)$-matching is perfect with probability 1. Indeed, there are at most $k + 1$ incompletely matched points. But the invariance argument of Theorem 4 shows that a translation invariant random set of points cannot have a positive finite number of points with positive probability.

By Proposition 11, Bob wins by placing his frog on an unblocked partner of Alice's opening frog. $\square$

Previous works on stable multi-matching [7–9] have considered questions about connectivity of the graph (many of which remain open). We do no know whether such questions have natural game interpretations.

## 4.3   Multiple Ponds and Bitwise XOR

Finally we address how to play several games of friendly frogs simultaneously. Consider $k$ sets $L_1, \ldots, L_k \subset \mathbb{R}^d$, each assumed to have no infinite descending chains and all distances distinct. (We imagine $k$ disjoint ponds, each with its own set of lily pads). In a position of $k$-**pond friendly frogs**, each set $L_i$ has two frogs on two distinct points. (We discuss the opening moves, in which the frogs are placed, below). Alice and Bob take turns, and a move consists of jumping one frog in one set $L_i$ to a different point in the same set $L_i$ according to the usual rules: the two frogs in $L_i$ must get strictly closer, but may not occupy the same point. A player loses if they have no legal move in any of the sets $L_i$.

The above game is an example of a general construction; it is known as the *disjunctive sum* of $k$ copies of friendly frogs. A remarkable theory of such sums of games was developed independently by Sprague [20] and Grundy [11], building on Bouton's analysis of the game of Nim [4] (also see [3] for an exposition as well as many far-reaching extensions). It turns out that this theory fits perfectly with friendly frogs, enabling us to show that Bob can win even with a substantial handicap in the opening moves.

**Theorem 13.** *Fix $k \geq 1$ and let $L_1, \ldots, L_k$ be independent Poisson processes on $\mathbb{R}^d$. Consider a game of $k$-pond friendly frogs, in which Alice first places two frogs in each of $L_1, \ldots, L_{k-1}$ and one frog in $L_k$, then Bob places the final frog in $L_k$, and Alice moves next. With probability $1$, Bob wins.*

In fact Bob has a unique good opening move that depends in an intricate way on Alice's $2k - 1$ initial frogs. The key to the proof is the following result extending stable matching to an integer-valued labeling of all pairs of points. Write $\mathbb{N} := \{0, 1, 2 \ldots\}$. For $S \subsetneq \mathbb{N}$, let $\operatorname{mex} S := \min(\mathbb{N} \setminus S)$ be the **minimum excluded value**. For a set $L \subset \mathbb{R}^d$ and an unordered pair of distinct points $x, y$ of $L$, let $F(x, y)$ be the set of positions to which one can legally move in friendly frogs, i.e. pairs that are strictly closer to each other than $x, y$ and share exactly one point with $x, y$.

**Proposition 14.** *Let $L$ be a Poisson process on $\mathbb{R}^d$. With probability $1$, there exists a map $G$ assigning an element of $\mathbb{N}$ to each unordered pair of $L$, with the following properties.*

*(i) For every $x \in L$ and $k \in \mathbb{N}$ there is a unique $y \neq x$ such that $G(x, y) = k$.*

*(ii) For each pair $x, y$ we have $G(x, y) = \operatorname{mex}\{G(u, v) : \{u, v\} \in F(x, y)\}$.*

*Proof.* As before, we construct the map via an iterative algorithm. Start with $G(x, y)$ undefined for all $x, y$. We say that each point of $x \in L$ **looks at** the closest other point $y$ for which $G(x, y)$ is currently undefined. For every pair $x, y$ that are looking at each other, set $G(x, y)$ to equal the smallest non-negative integer that is not currently assigned to any pair containing $x$ or $y$. Now repeat indefinitely.

We first check that the resulting $G$ assigns an integer to every pair of points. Indeed, if $G(x, y)$ is undefined then $x, y$ never looked at each other, and so one of them, say $y$, must have a closer point $z$ for which $G(y, z)$ is undefined. Passing to the closest such $z$ and iterating gives an infinite descending chain, a contradiction.

14

We now check the claimed properties. For (i), it is immediate that no two pairs containing $x$ are assigned the same integer. It remains to check that some pair containing $x$ has the label $k$. Let $U_k$ be the set of points $x$ that are not contained in any pair with label $G(x, y) = k$. By invariance, if $U_k$ is non-empty then it is infinite. Let $W \subseteq U_k$ be any set of size $k + 2$. By the pigeon-hole principle there exist $u, v \in W$ with $G(u, v) > k$. But this is a contradiction: the algorithm should instead have assigned $u, v$ a value $\leq k$.

To check (ii), note that, during the stages of the algorithm, a given point looks at other points of $L$ in order of increasing distance (perhaps looking at the same point for multiple consecutive stages). Therefore, when the algorithm assigns a value to the pair $x, y$, all pairs in $F(x, y)$ have been assigned values, while all other pairs that share a point with $x, y$ have not. Therefore $G(x, y)$ is assigned the mex as claimed. $\qquad\square$

It is easy to see that the set of pairs $\{x, y\}$ with $G(x, y) = 0$ is precisely the stable matching. However, the set of pairs with $G(x, y) \leq m$ does not in general coincide with the $m$-matching considered earlier. See Figure 6. It should also be noted that the analogue of property (i) in Proposition 14 does not hold in general for finite sets $L$ – it is possible that for some $x$ the set $\{G(x, y) : y \in L \setminus \{x\}\}$ is not the interval $\{0, \dots, L - 2\}$.

*Proof of Theorem 13.* Let $\oplus$ denote bitwise XOR of binary expansions, so if $a = \sum_{j \in \mathbb{N}} \alpha_j 2^j$ and $b = \sum_{j \in \mathbb{N}} \beta_j 2^j$ with $\alpha_j, \beta_j \in \{0, 1\}$ then $a \oplus b := \sum_{j \in \mathbb{N}} \sigma_j 2^i$ where $\sigma_j \in \{0, 1\}$ satisfies $\sigma_j \equiv \alpha_j + \beta_j \pmod{2}$. Consider a position of $k$-pond friendly frogs with two frogs in each pond, at locations $x_i, y_i \in L_i$. We claim that it is a P-position if and only if $\bigoplus_{i=1}^k G_i(x_i, y_i) = 0$, where $G_i$ is the map given by Proposition 14 for $L_i$. This remarkable fact follows immediately from the general theory (see [3, 11, 20]), given condition Proposition 14 (ii) on $G$ and the fact that friendly frogs is an *impartial* game (i.e. the same moves are available to each player) and has no infinite lines of play. Since the proof is quite simple (given the highly non-trivial insight of what to prove), we will summarize it below.

Given this characterization of P-positions, Bob's winning move is easy to describe. He computes $h := \bigoplus_{i=1}^{k-1} G_i(x_i, y_i)$, and places the final frog on the unique point $y_k \in L_k$ for which $G_k(x_k, y_k) = h$, which exists by Proposition 14 (i). Since $h \oplus h = 0$, this gives a P-position.

Finally, we explain how to prove the claim. As usual, this amounts to checking conditions (N) and (P). Let $g_i = G_i(x_i, y_i)$ and $g = \bigoplus_{i=1}^k g_i$. For (N), suppose that $g \neq 0$. Write $g = \sum_{j \in \mathbb{N}} \gamma_j 2^j$, and let $k$ be maximal such that $\gamma_k = 1$ (the most significant bit of $g$). Choose $i$ such that $g_i$ also has $k$th bit equal to 1, and note that $g_i \oplus g < g_i$. By Proposition 14 (ii), we can move a frog in $L_i$ to reduce $G_i(x_i, y_i)$ to $g_i \oplus g$, resulting in a P-position. On the other hand, for (P), if $g = 0$ then by Proposition 14 (ii), any move changes one of the $G_i(x_i, y_i)$, giving an N-position. $\qquad\square$

# References

[1] R. Basu, A. E. Holroyd, J. B. Martin, and J. Wästlund. Trapping games on random boards. *Ann. App. Prob.* To appear.

[2] I. Benjamini, R. Lyons, Y. Peres, and O. Schramm. Group-invariant percolation on graphs. *Geom. Funct. Anal.*, 9(1):29–66, 1999.

[3] E. R. Berlekamp, J. H. Conway, and R. K. Guy. *Winning ways for your mathematical plays. Vol. 1.* A K Peters, Ltd., Natick, MA, second edition, 2001.

[4] C. L. Bouton. Nim, a game with a complete mathematical theory. *Ann. of Math. (2)*, 3(1-4):35–39, 1901/02.

[5] D. J. Daley and G. Last. Descending chains, the lilypond model, and mutual-nearest-neighbour matching. *Adv. in Appl. Probab.*, 37(3):604–628, 2005.

[6] D. J. Daley and D. Vere-Jones. *An introduction to the theory of point processes. Vol. II.* Probability and its Applications. Springer, New York, second edition, 2008. General theory and structure.

[7] M. Deijfen, O. Häggström, and A. E. Holroyd. Percolation in invariant Poisson graphs with i.i.d. degrees. *Ark. Mat.*, 50(1):41–58, 2012.

[8] M. Deijfen, A. E. Holroyd, and Y. Peres. Stable Poisson graphs in one dimension. *Electron. J. Probab.*, 16(44):1238–1253, 2011.

[9] M. Deijfen and F. M. Lopes. Bipartite stable Poisson graphs on $\mathbb{R}$. *Markov Process. Related Fields*, 18(4):583–594, 2012.

[10] D. Gale and L. S. Shapley. College admissions and the stability of marriage. *Amer. Math. Monthly*, 69(1):9–15, 1962.

[11] P. M. Grundy. Mathematics and games. *Eureka*, 2:6–8, 1939.

[12] O. Häggström. Infinite clusters in dependent automorphism invariant percolation on trees. *Ann. Probab.*, 25(3):1423–1436, 1997.

[13] O. Häggström and R. Meester. Nearest neighbor and hard sphere models in continuum percolation. *Random Structures Algorithms*, 9(3):295–315, 1996.

[14] A. E. Holroyd, I. Marcovici, and J. B. Martin. Percolation games, probabilistic cellular automata, and the hard-core model. 2015. arXiv:1503.05614.

[15] A. E. Holroyd and J. B. Martin. Galton-Watson games. In preparation.

[16] A. E. Holroyd, J. B. Martin, and Y. Peres. Asymmetric stable matchings in high dimensions. In preparation.

[17] A. E. Holroyd, R. Pemantle, Y. Peres, and O. Schramm. Poisson matching. *Ann. Inst. Henri Poincaré Probab. Stat.*, 45(1):266–287, 2009.

[18] A. E. Holroyd and T. Soo. Insertion and deletion tolerance of point processes. *Electron. J. Probab.*, 18(74):1–24, 2013.

[19] M. Krivelevich. Positional games. In *Proc. International Congress Math.*, volume 4, pages 355–379, 2014.

[20] R. P. Sprague. Über mathematische Kampfspiele. *Tôhoku Math. J.*, 41:438–444, 1935/6.

[21] R. Telgársky. Topological games: on the 50th anniversary of the Banach-Mazur game. *Rocky Mountain J. Math.*, 17(2):227–276, 1987.

[22] J. Wästlund. Replica symmetry of the minimum matching. *Ann. of Math. (2)*, 175(3):1061–1091, 2012.