Chapter 1
Multiagent Pursuit Evasion, or Playing Kabaddi

Kyle Klein and Subhash Suri

Abstract We study a version of pursuit evasion where *two or more pursuers* are required to capture the evader because the evader is able to overpower a single defender. The pursuers must coordinate their moves to fortify their approach against the evader while the evader maneuvers to disable pursuers from their unprotected sides. We model this situation as a game of *Kabaddi*, a popular South Asian sport where two teams occupy opposite halves of a field and take turns sending an *attacker* into the other half, in order to win points by tagging or wrestling members of the opposing team, while holding his breath during the attack. The game involves team coordination and movement strategies, making it non-trivial to formally model and analyze, yet provides an elegant framework for the study of multiagent pursuit-evasion, for instance, a team of robots attempting to capture a rogue agent. Our paper introduces a simple discrete (time and space) model for the game, offers analysis of winning strategies, and explores tradeoffs between maximum movement speed, number of pursuers, and locational constraints.[1]

1.1 Introduction

Pursuit-evasion games provide an elegant and tractable framework for the study of various algorithmic and strategic questions with relevance to exploration or monitoring by autonomous agents. Indeed, there is a rich literature on these games under various names [13], such as *man-and-the-lion* [12, 16, 9], *cops-and-robber* [6, 1, 8, 3, 4], *robot-and-rabbit* [6], and *pursuit-evasion* [14, 5, 7], just to name a few.

Kyle Klein
Computer Science, UC Santa Barbara, CA 93106, USA, e-mail: kyleklein@cs.ucsb.edu

Subhash Suri
Computer Science, UC Santa Barbara, CA 93106, USA, e-mail: suri@cs.ucsb.edu

[1] A 4-page abstract announcing some of these results, without proofs and without differential speed, was presented at the 15th Canadian Conference on Computational Geometry, August 2010.

In this paper we study a (discrete time, discrete space) version of pursuit evasion where *two or more pursuers* are required to capture the evader because the evader is able to overpower a single (and isolated) defender. These situations arise in pursuit of a *rogue non-cooperative agent*, which could be a malfunctioning robot, a delirious evacuee, or a noncooperative patient. Thus, the pursuers are forced to coordinate their moves to fortify their approach against the evader while the evader maneuvers to disable pursuers from their unprotected sides. In the basic formulation of the game, all agents have the same capabilities including the maximum movement speed, but we also derive some interesting results when one side can move faster than the other.

In modeling our pursuit-evasion scenario, we draw inspiration from the game of *Kabaddi*, which is a popular South Asian sport. The game involves two teams occupying opposite halves of a field, each team taking turns to send an "attacker" into the other half, in order to win points by tagging or wrestling members of the opposing team [17]. The attacker must hold his breath during the entire attack and successfully return to his own half—the attacker continuously chants "kabaddi, kabaddi, $\cdots$" to demonstrate holding of the breath. There are several elements of this game that distinguish it from the many other pursuit games mentioned above, but perhaps the most significant difference is that it typically requires two or more defenders to capture the opponent, while the attacker is able to capture a single isolated defender by itself. This *asymmetry* in the game adds interesting facets to the game and leads to interesting strategies and tradeoffs.

While the use of multiple pursuers is common in many existing pursuit evasion games, the main concern in those settings is to simply distribute pursuers in the environment to keep the evader from visiting or reentering a region. This is indeed the case in all graph searching [2, 11, 14] or visibility based pursuit evasion [5, 7, 15]. In the lion-and-the-man game also there are known results that show that multiple lions can capture the man when the man lies inside the convex hull of the lions [10]. By contrast, the main question in Kabaddi is whether the defenders can ever *force* the attacker inside their convex hull, perhaps even by sacrificing some of their agents. The other games such as the *cops-and-robber* differ from kabaddi in the way capture occurs as well as the *information* about the evader's position. For instance, the current position of all the players is public information in kabaddi while the position of the robber or evader is often assumed to be unknown to cops or pursuers. Furthermore, it is also typically assumed that each cop (robot) follows a fixed trajectory that is *known* to the robber (rabbit). This makes sense in situations where the defenders (cops) have fixed patrol routes, but not in interactive games like kabaddi. The problems and results in the graph searching literature are also of a different nature than ours [2, 11], although variations using differential speed [4] and capture from a distance [3] have been considered in graphs as well.

Finally, in the *visibility-based* pursuit-evasion games, the evader is often assumed to have infinite speed, and the capture is defined as being "seen" by some defender—both infinite visibility or limited-range visibility models have been considered [5, 8]. By contrast, kabaddi involves equal speed agents and requires a physical capture that leads to a very different set of strategies and game outcomes. With this background, let us now formalize our model of kabaddi.

1.1.1 The Standard Model

We consider a *discrete* version of the game, in which both time and space are discrete: the players take alternating turns, and move in discrete steps. In particular, the game is played on a $n \times n$ grid S, whose cells are identified as tuples (i, j), with $i, j \in \{1, 2, \ldots, n\}$. We will mainly use the Kabaddi terminology, namely, attacker and defenders, with the former playing the role of the evader and the latter the pursuers. Our main analysis will focus on the case of one attacker and two defenders, although in the latter part of the paper, we do derive some results for the case of d defenders, for any $d > 1$.

We use the letters A and D to denote the *attacker* and a *defender*, respectively. When there are multiple defenders, we use subscripts such as D_1, D_2, etc. We need the concepts of *neighborhood*, *moves*, and *capture* to complete the description of the game. Throughout, we assume that precisely two defenders are required to capture the attacker.

Neighborhood. The *neighborhood* $N(p)$ of a cell $p = (i, j)$ is the set of (at most) 9 cells, including p itself, adjacent to p, or equivalently the set of all cells with L_∞ distance at most 1 from p. In Figure 1.1, the neighborhood of A is shown with a box around it. Slightly abusing the notation, we will sometimes write $N(A)$ or $N(D)$ to denote the neighborhood of the current position of A or D.

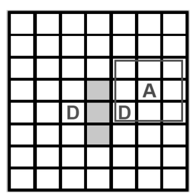

Fig. 1.1 The standard model of kabaddi. A can capture the defender closer to it, which is inside $N(A)$. The defenders can capture A at any position in the shaded region, which is the common intersection of their neighborhoods.

Moves. The attacker and the defenders take turns making their moves, with the attacker moving first. In one step, the attacker and the defenders can move to any cell in their neighborhood. *All the defenders can move simultaneously in one step.*

Capture. *A captures* a defender D if it is the unique defender lying inside the neighborhood of A. That is, with two defenders, D_1 is captured when $D_1 \in N(A)$ and

$D_2 \notin N(A)$. (Notice that A only needs to enclose a defender within its neighborhood to capture it.)

Conversely, the defenders capture the attacker, when A lies in the common intersection of the two defenders' neighborhoods. That is, $A \in (N(D_1) \cap N(D_2))$.

Game Outcome. The attacker *wins* the game if he can capture one or more defenders, and the defenders win the game if they can capture the attacker. If neither side wins, then the game is a tie.

This particular form of capture has a tendency to make defenders always stick together, and fails to model the real world phenomenon where defenders try to "surround" the attacker—see figure above. We therefore introduce a *minimum separation* condition on the defenders in the following way:

no defender can be inside the neighborhood of another defender.

These rules together define our *standard model* of kabaddi. Other models can be obtained by varying the definition of the *neighborhood* and relaxing the separation condition for defenders, and we obtain some results to highlight the impact of these modeling variables.

Safe Return and Holding of the Breadth. In Kabaddi, the attacker must hold his breath during the attack, and after the attack successfully return to his side. These are non-trivial issues to model tractably, and we exclude them from our current model, instead relying on the following interpretation: *the worst-case number of moves before the game's outcome serves as a proxy for the breath, and the attacker can conservatively decide at some point to return to his side.* However, if this duration is known to the defenders, then they can attempt to interfere with his return. We leave these interesting, but complicated, issues for future work. One could argue that these issues are not important in the multiagent pursuit-evasion problem.

1.1.2 Our Results

In the case of a single attacker A against a single defender D, the game resembles the discrete version of the man-and-the-lion. We include a simple analysis of this case for two reasons: first, it serves as a building block for the multi-defender game; and second it allows us to highlight the impact of player's *speed* on the game outcome, which we believe is a new direction in pursuit evasion problems. Unsurprisingly, in the single defender case, the attacker can always capture the defender D in $O(n)$ number of steps, which is clearly optimal, upto a constant factor, in the worst-case.

We show that a speed of $1+\Theta(1/n)$ is both necessary and sufficient for the defender to indefinitely evade the attacker. In particular, a defender with the maximum speed $1+5/(\frac{n}{4}-3)$ can evade the attacker indefinitely, but a defender with the maximum speed of $1+1/n$ can be captured.

The game becomes more challenging to analyze with two defenders, where the attacker continuously runs the risk of being captured himself, or have the defenders evade him forever. Our main result is to show that the attacker has a winning strategy in worst-case $O(n)$ moves. One important aspect of the standard model is the *separation* requirement for the defenders—each must remain outside the neighborhood of the other. Without this restriction, we show that the two defenders, whom we call *strong defenders* to distinguish from the standard ones, can force a draw: neither the attacker nor a defender can be captured. A further modification of the model, which disallows the *diagonal* moves, tips the scale further in the favor of strong defenders, allowing them to capture the attacker in $O(n^2)$ steps.

Extending the analysis to more than two players is a topic for ongoing and future work, and seems non-trivial. Surprisingly, for the standard model, it is not obvious that even $\Theta(n)$ defenders can capture the attacker, nor it is obvious that the attacker can win against k defenders, for $k > 2$. (The definition of capture remains the same: two defenders are enough to capture the attacker.)

However, if we endow the agents with different speeds, then we can obtain some interesting results, as in the case of the single defender mentioned earlier. In particular, if the attacker can make $\min\{10, d-1\}$ single steps in one move, then it can avoid capture indefinitely against d defenders, and if it can make $\min\{11, d\}$ steps per move, then it can capture all d defenders in time $O(dn)$. Thus, the attacker has a winning or non-losing strategy with $O(1)$ speed against an unbounded number of players, assuming a safe initial position.

1.2 One on One Kabaddi

We begin with the simple case of the attacker playing against a single defender. Besides being of interest in its own right, it also serves as building block for the more complex game against two defenders. We show that in this case the attacker always has a winning strategy in $O(n)$ moves.

Throughout the paper, we assume that the grid is aligned with the axes, and use $\Delta x = |D_x - A_x|$ and $\Delta y = |D_y - A_y|$, resp., for the x (horizontal) and the y (vertical) distance between A and D.

Theorem 1. *The attacker can always capture a single defender in a $n \times n$ game of kabaddi in $O(n)$ moves.*

Proof. The attacker's basic strategy is to chase the defender towards a wall, keeping him trapped inside a continuously shrinking rectangular region. Specifically, as long as $\min\{\Delta x, \Delta y\} > 0$ on its move, the attacker makes the (unique) diagonal move towards the defender, reducing both Δx and Δy by one. Because the grid is $n \times n$, the attacker can make at most n such moves before either Δx or Δy becomes zero. Without loss of generality, suppose $\Delta x = 0$. From now on, the attacker always moves to maintain $\Delta x = 0$ while reducing Δy by one in each move. Because Δy can be initially at most n, the attacker can reduce to it one in at most $n-1$ moves, at which point it has successfully captured the defender because both Δx and Δy are at most 1. This completes the proof.

1.3 Attacker Against Two Defenders

The game is more complex to analyze against two defenders. We begin by isolating some necessary conditions for the game to terminate, or for the next move to be safe. We then discuss the high level strategy for the attacker, and show that it can pursue the defenders using that strategy *without being captured* itself. Together with a bound for the duration of the pursuit, this yields our main result of $O(n)$ steps win for A in the standard model. We denote the two defenders by D_1 and D_2, and use D to refer to a non-specific defender when needed. Throughout the game, *we ensure that whenever A makes a move, it is safe in the sense that it cannot be captured by the defenders in their next move.*

Lemma 1. *On A's turn, if* $\max\{\Delta x, \Delta y\} \leq 2$ *for at least one of the defenders, then A can capture a defender in one step. Conversely, on the defenders' turn, if* $\max\{\Delta x, \Delta y\} > 2$ *for one of the defenders, then they cannot capture A on their move.*

Proof. We first observe that neither defender can be inside the neighborhood of A, namely, $N(A)$. This holds because a single defender inside $N(A)$ must have been captured in A's last move and if both the defenders are inside $N(A)$, then they would have captured A in their last move. Thus, we must have $\max\{\Delta x, \Delta y\} \geq 2$ for both the defenders.

Let D_1 be the defender that satisfies the conditions of the lemma, meaning that $\max\{\Delta x, \Delta y\} = 2$. If both the defenders satisfy the condition, then let us choose the one for which $\Delta x + \Delta y$ is smaller; in case of a tie, choose arbitrarily. Without loss of generality, assume that D_1 lies in the upper-left quadrant from A's position (i.e. north-west of A). We now argue that A can always capture D_1 as follows. See Figure 1.2.

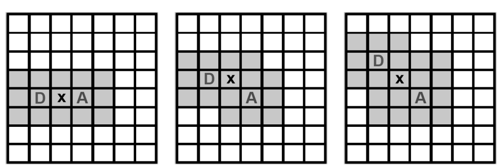

(a) (b) (c)

Fig. 1.2 Illustrating the three cases in Lemma 1: $\Delta x + \Delta y = 2$ (a), 3 (b) and 4 (c). The shaded area is the region that cannot contain the second defender.

If $\Delta x + \Delta y = 2$, then we must have either $\Delta x = 2, \Delta y = 0$ or $\Delta x = 0, \Delta y = 2$. In the former case, A can capture D_1 by moving to its x-neighbor (shown in Figure 1.2 (a)), and in the latter by moving to its y-neighbor. Since the second defender must lie outside $N(A) \cup N(D_1)$, this move cannot cause A to be captured. Similarly, if $\Delta x + \Delta y = 3$ (shown Figure 1.2 (b)), then we have either $\Delta x = 2, \Delta y = 1$, or $\Delta x = 1, \Delta y =$

2. In both cases, A captures D_1 by moving to its north-west neighbor $(A_x - 1, A_y + 1)$. Observe that, by the minimum separation rule, if there is a defender at $(A_x - 2, A_y + 1)$, then there can't be one at $(A_x - 1, A_y + 2)$, and vice versa ensuring the safety of this move—there also cannot be a defender at $(A_x, A_y + 2)$ because that would contradict the choice of the closest defender by distance.

Finally, if $\Delta x + \Delta y = 4$ (shown Figure 1.2 (c)), then A captures D_1 by moving to $(A_x - 1, A_y + 1)$. This is a safe move because the only position for D_2 that can capture A is at $(A_x, A_y + 2)$, but in that case D_2 is the defender with the minimum value of $\Delta x + \Delta y$, contradicting our choice of the defender to capture. This completes the first claim of the lemma. For the converse, suppose that $\Delta x > 2$ for defender D_1. Then, after the defenders' move, A is still outside the neighborhood of D_1, and so A is safe. This completes the proof.

The attacker initiates its attack by first aligning itself with one of the defenders in either x or y coordinate, without being captured in the process. The following two technical lemmas establish this.

Lemma 2. *A can move to the boundary in $O(n)$ moves without being captured.*

Proof. By assumption, A is currently safe. We first check whether A can capture a defender in the next move: if so, he wins. Otherwise, by Lemma 1, we must have that $\max\{\Delta x, \Delta y\} > 2$ for both D_1 and D_2. The attacker A now (arbitrarily) chooses a defender, say, D_1 and moves so as to increase both its x and y distances to that defender by one—this is always possible unless A is already on the boundary. Because this always maintains $\max\{\Delta x, \Delta y\} > 2$ with respect to D_1, by Lemma 1, the defenders cannot capture A, and is A guaranteed to reach the boundary in $O(n)$ steps.

Lemma 3. *By moving along the boundary, A can always force either $\Delta x = 0$ or $\Delta y = 0$ for one of the defenders in $O(n)$ moves, without being captured.*

Proof. Without loss of generality, assume that A is on the bottom boundary, and that at least one of the defenders, say, D_1 lies in its upper-right quadrant (i.e. has larger x coordinate). Then, A's strategy is to always moves right on its turn, and is guaranteed to achieve $\Delta x = 0$ with D_1 at some point. We only need to show that A cannot be captured during this phase. But if A were captured at position $(i, 0)$, then the defenders must be at positions $(i-1, j_1)$ and $(i+1, j_2)$, for $j_1, j_2 \in \{0, 1\}$—these are the only positions whose neighborhoods contain the cell $(i, 0)$ in common. However, the position of A one move earlier was $(i-1, 0)$, so the first defender would necessarily satisfy the conditions of Lemma 1 and would have been captured by A already.

1.3.1 The Second Phase of the Attack

Having reached the starting position for this second phase of the game, we assume without loss of generality that A is at the bottom boundary, and that after A's last move, $\Delta x = 0$ for one of the defenders. From now on, A will always ensure that $\Delta x \leq 1$ for one of the defenders *after each of A's moves*. The x-distance can become $\Delta x = 2$ *after the defenders' move* but A will always reduce it to 1 in its next move.

By Lemma 1, if both Δx and Δy are at most 2, then A can win the game. On the other hand, if the players are too far apart, then both sides are safe for the next move. Thus, all the complexity of the game arises when the distance between A and the defenders is 3, requiring careful and strategic moves by both sides. We show that A can always follow an attack strategy that ensures a win in $O(n)$ steps, while avoiding capture along the way.

In order to measure the *progress* towards A's win, we use the distance from A's current position to the top boundary of the grid *while ensuring that* $\Delta x \leq 1$ *continues to hold*. In particular, define $\Phi(A)$ as the gap between the current y position of A and the top boundary. That is, $\Phi(A) = (n - A_y)$, where this gap is exactly $n-1$ when the second phase begins with A on the bottom boundary. We say that A *makes progress* if $\Phi(A)$ shrinks by at least 1, while Δx remains at most 1 for some defender. Clearly, when the $\Phi(A)$ reaches one, A has a guaranteed win (by Lemma 1). If the attacker succeeds in capturing a defender, then we consider that also progress for the attacker.

The overall plan for our analysis is the following:

1. If $\max\{\Delta x, \Delta y\} \leq 2$ for at least one defender, then the attacker wins in one move (Lemma 1). If $\Delta y > 3$ for some defender, then A can move to reduce Δy by one, while keeping $\Delta x \leq 1$, and this move is safe by Lemma 1. Thus, the only interesting cases arise when $\Delta y = 3$; these are handled as follows.
2. If $\Delta y = 3$ and $\Delta x = 0$ for some defender, then Lemma 4 below shows that A makes progress in $O(1)$ number of moves.
3. If $\Delta y = 3$ and $\Delta x = 1$ or 2 for some defender, then Lemmas 5 and 6 show that A can make progress in $O(n)$ number of moves.

In the following, we use the notation $N^2(p)$ to denote the 2-*neighborhood* of a cell p, meaning all the positions that can be reached from p in two moves.

Lemma 4. *On A's move, if* $\Delta y = 3$ *and* $\Delta x = 0$ *holds for some defender, then A makes progress in one move.*

Proof. Figure 1.3 (a) illustrates the game configuration for this case, where the defender satisfying the distance condition $\Delta x = 0, \Delta y = 3$ is shown as D. There are three positions for A to advance and make progress, and they are marked as x in the figure—in each case, the y distance reduces by 1, while Δx remains at most 1. We only need to show that A can move to one of these positions without being captured itself.

In order to prove this, we observe that (1) neither defender is currently inside $N^2(A)$ because that is a winning configuration for A by Lemma 1; (2) the second defender is not in $N(D)$, as required by the separation rule for defenders. Thus, the second defender must be outside $N^2(A) \cup N(D)$. But in order to foil A's move to all three x positions, the second defender must also be within the 2-neighborhood of *all* the x positions. That, however, is impossible, as is readily confirmed by inspection

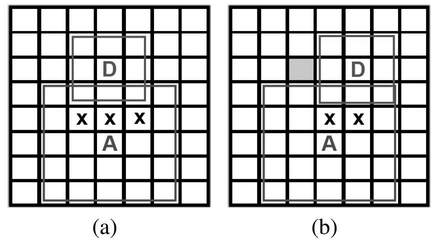

Fig. 1.3 Proofs of Lemmas 4 (a) and 5 (b).

of Figure 1.3 (a). Thus, A can safely move to one of the positions marked as x, and guarantee progress. We note that when A and D are on the boundary, there are two x positions instead of three, and in that case A can always move to the x directly north and make progress.

Lemma 5. *On A's move, if $\Delta y = 3$ and $\Delta x = 1$ holds for some defender, then A makes progress in $O(n)$ number of moves.*

Proof. Figure 1.3 (b) illustrates the game configuration for this case, where the defender satisfying the distance condition $\Delta x = 1, \Delta y = 3$ is shown as D. (We assume without loss of generality that $D_x = A_x + 1$ because the case $D_x = A_x - 1$ is entirely symmetric.) In this case, there are two positions marked x that allow A to make progress by reducing Δy. In order to foil A's move, the second defender must be positioned so as to cause A's capture at both these positions. Reasoning as in the previous lemma, however, D_2 has to lie outside both $N(D)$ as well as $N^2(A)$. It is easy to see that there is precisely one position for D_2, shown as the shaded cell, that threatens A's capture at both the x positions.

This is a case where A cannot ensure progress in a single step, and instead a multi-step argument is needed. In particular, A moves to its right neighboring cell, at location $(A_x + 1, A_y)$, which does not improve $\Phi(A)$, but we show that $\Phi(A)$ *will* improve in $O(n)$ steps. Consider the next move of the defenders. The defender labeled D must move to a cell within $N(D)$, and we analyze the progress by A as follows: (i) if D moves up, making its distance from A equal to $\Delta y = 4$, then the next move of A makes a guaranteed progress by moving to make $\Delta y = 3$ and $\Delta x \leq 1$. This move is safe for A by Lemma 1. (ii) if D moves down, making its distance from A equal to $\Delta y = 2$, then, A has a guaranteed win according to Lemma 1. (iii) if D stays in its current cell, then we have $\Delta y = 3$ and $\Delta x = 0$ on A's move, for which Lemma 4 guarantees progress in one move.

Thus, the only interesting cases are if D moves to its left or right neighbor. If D moves left, causing $\Delta y = 3$ and $\Delta x = 1$, then A can immediately make progress because both the defenders are on the left side of A's position (recall that A was forced to make a move without progress because the second defender was in the shaded cell), and so A can safely move diagonally to reduce both Δx and Δy distances to D. In this case we have progress in a total of 2 moves.

On the other hand, if D moves to its right neighbor, then the situation of impasse can persist, because both positions marked x where A can make progress can cause A to be captured. This forces A to continue to mimic D's rightward move by moving to its right neighbor. However, this impasse can continue only for $O(n)$ moves because as soon as D reached the right boundary of the field, he is forced to move up, down, or left, giving A a chance to make progress. This completes the proof of the lemma.

Lemma 6. *If $\Delta y = 3$ and $\Delta x = 2$ for some defender say D_1 then A may make progress in $O(n)$ moves.*

Proof. The proof is similar to the proof of Lemma 5, and omitted due to lack of space.

1.3.2 Completing the Analysis

We can now state our main theorem.

Theorem 2. *In the standard model of kabaddi on a $n \times n$ grid, the attacker can capture both the defenders in $O(n)$ worst-case moves.*

Proof. We show that, starting from an initial safe position, the attacker always has a move that keep him safe for the next move of the defenders, and that after $O(n)$ moves the attacker can place itself on a boundary with either $\Delta x = 0$ or $\Delta y = 0$ for some defender. Without loss of generality, suppose the attacker reaches the bottom boundary, with $\Delta x = 0$ (Lemmas 2, 3). In the rest of the game, the attacker always maintains $\Delta x \leq 1$ after each of its moves. The attacker's next move is described as follows:

1. If $\max\{\Delta x, \Delta y\} \leq 2$ for some defender, then the attacker can capture a defender in 1 move (Lemma 1, and the remaining defender in $O(n)$ moves.
2. If $\Delta y \geq 4$, then the attacker always moves to reduce Δy and Δx by one, unless Δx is already zero.
3. If $\Delta y = 3$, then depending on whether $\Delta x = 0, 1$ or 2, the attacker's strategy is given by Lemma 4, 5 or 6, respectively.

These cases exhaust all the possibilities, and as argued earlier, the attacker can reduce $\Phi(A)$ by one in $O(n)$ moves. Since the maximum possible value of $\Phi(A)$ is initially $n-1$, and it monotonically decreases, we must reach $\Phi(A) = 1$ in worst-case $O(n^2)$ moves, terminating in a win by A.

We now argue that the $O(n^2)$ bound is pessimistic and that $O(n)$ moves suffice. The key idea is that once the attacker forces $\Delta x = 0$, it only moves to the three cells above it and the one to its right. The three upward moves clearly cause progress, so we only need to argue that the rightward moves happen $O(n)$ times. This follows because the grid has width n, and therefore after at most $n-1$ rightward moves, every additional rightward move must be preceded by some leftward move. Since the attacker always moves upward in its left-directed moves, it makes progress in each of those moves. Then due to the fact A only needs $n-2$ upward moves, there can be at most $2n-3$ right moves (the initial $n-1$ moves plus the $n-2$ moves corresponding to upward moves), and thus at worst $3n-5$ total moves. Thus the attacker captures both defenders in $O(n)$ worst case moves. This completes the proof of the theorem.

1.4 Strong Defenders

In the standard model, each defender must remain outside the neighborhood of other defenders; that is, $D_i \notin N(D_j)$, for all i, j. The defenders become more powerful when this requirement is taken away. Let us call these *stronger* defenders. In this case we explore what happens when we remove the stipulation that the defenders cannot be within each other's neighborhoods. This creates two stronger defenders and as a result creates a game where ideal play means not only can the attacker not win, but the defenders cannot either. We assume that play starts with defenders already in a side-by-side position, that is, $\Delta x + \Delta y = 1$ with respect to D_1 and D_2's coordinates.

Theorem 3. *Under the* strong *model of defenders, there is a strategy for the defenders to avoid capture forever. At the same time, the attacker also has a strategy to avoid capture.*

Proof. We first argue that the attacker can evade capture. Suppose that the defenders were to capture A in their *next* move. If neither defender is inside $N^2(A)$, then A is clearly safe in its current position for the defenders' next move, so at least one of the defenders, say D_1 is inside $N^2(A)$. Unless $D_2 \in N(D_1)$, by Lemma 1, then A can capture D_1 in its next move. Thus, D_1, D_2 must be adjacent, namely, in each other's 1-neighborhoods.

We now argue that all defender positions from which they can capture A in the next move are *unsafe*, meaning the attacker can capture one of the defenders in its current move. There are only two canonical positions for the defenders with one or both of them in the outer cells of $N^2(A)$: either side-by-side, or diagonal from one another. In the first case, the defenders only threaten the cells in front of them but not those that are diagonal, so A can move to one of those diagonal spaces. In the second case, A can capture by moving to any space diagonal from a defender.

Similarly, we can show that defenders can also avoid capture. Figure 1.4 shows a representative situation just before the defenders' move. Suppose that the attacker were to capture one of the defenders in its *next move*. We claim that the cells marked as A in the figure are the only places (upto symmetry) for the attacker's *current* position—i.e. these are the positions where A is not captured currently but can capture a defender by moving to the cells shown shaded. This is found by taking the union of the 1-neighborhoods of the three shaded spots (the only places A captures a defender) to find all possible places A may move to capture from, then removing all those that the defenders could capture. This result in a list of spots D cannot capture but must avoid capture from. However, the defenders can avoid this capture by simply "flipping" their orientation, as shown by arrows in the figure. Notice that after the flip the attacker now cannot capture with its move. Also the flip does not rely on the position of the boundary, as the defenders move up only if the attacker is above them, and move down only if the attacker is below them. Thus this can be performed regardless of location.

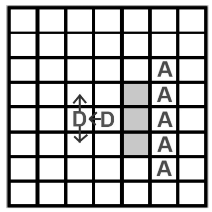

Fig. 1.4 Illustrates Theorem 3.

1.5 Strong Defenders with Manhattan Moves

Thus, in the standard model but with strong defenders, we have a tie, and neither side can guarantee a win. In the following, we show that if we disallow the *diagonal moves*, permitting a player to move only to its left, right, up, and down neighbors, then the defenders have a winning strategy. That is, the movement metric is Manhattan metric—a player can only move to a cell within the L_1 distance of 1 from its current cell. The definition of the capture, however, remains the same as in the standard model. Due to lack of space, the proof of the following theorem is omitted from this extended abstract.

Theorem 4. *Two strong defenders playing under the Manhattan moves model can always capture the attacker in $O(n^2)$ moves.*

1.6 Differential Speed Pursuit Evasion

So far, we have assumed that all players have the same (unit) speed. While we are unable to resolve the outcome of these games when the attacker plays against more than two defenders, we show below that *differential speed* leads to some interesting results. We model the speed as the *number* of unit-step moves a player can make on its turn—each step is the same elementary move used in the standard model. In particular, *on its turn, a player with speed s can repeat the following s times, starting at a cell* $p = p_0$*:*

$$\text{move to any cell } p' \in N(p), \text{ and set } p = p'.$$

We allow the speed to be any rational number. Thus, a player with movement speed $s + \frac{p}{q}$ can make s unit step moves on each turn *plus* it can make $s+1$ steps on every $\lfloor q/p \rfloor$th turn. Please note that this definition is not the same as being able to move to a cell at distance at most s—specifically, our attacker has a chance to visit, and possibly capture, s defenders in a single move. *However, during his turn, if the attacker is ever in the common intersection of two defenders' neighborhoods, then it is captured (as in the standard model).*

We first consider the minimum speed advantage needed by a single defender to escape the attacker forever.

1.6.1 One on One Game with Speedier Defender

The following theorem shows that a speed of $1 + 1/n$ is not enough for the single defender to evade capture by the attacker.

Theorem 5. *A defender with maximum speed* $1 + \frac{1}{n}$ *can be captured in* $O(n)$ *moves by an unit-speed attacker on the* $n \times n$ *grid.*

Proof. The attacker's strategy is the same as in Theorem 1. We simply observe that despite the speed disadvantage the attacker still reduces either Δx or Δy to zero within n moves. Without loss of generality, assume that Δx becomes zero. After that, the attacker can also enforce $\Delta y = 0$ within n moves. In these n moves, the defender gains only one extra move, which only increases Δx to 1, but is still sufficient for the capture. Thus the defender is captured in $O(n)$ moves by the attacker.

Surprisingly, it turns out that a speed of $1 + \Theta(1/n)$ suffices for the defender to escape, as shown in the following theorem.

Theorem 6. *A defender with maximum speed* $1 + 5/(\frac{n}{4} - 3)$ *can indefinitely evade the attacker on an* $n \times n$ *grid.*

Proof. Assume an initial placement of the two agents in which (1) the defender D is at least distance $n/4$ from its closest boundary, which we assume to be the bottom boundary, (2) A is distance $\frac{n}{4} + 3$ from the same boundary, and (3) $\Delta x + \Delta y = 3$. (The defender can easily enforce the condition $\Delta x = 0$, and the remaining conditions are to achieve a safe initial separation between the attacker and the defender.) We argue that the defender can successfully maintain these conditions, and when needed use its extra moves to reestablish them with respect to a different boundary.

The defender's strategy now is to simply mimic the moves of the attacker as long as it can do this without running into a boundary. During these moves, the defender is safe because of the condition $\Delta y = 3$ or $\Delta x = 3$ (cf. Lemma 1).

Since the defender D is at least $n/4$ away from the boundary that is opposite the attacker, its speed advantage guarantees it 5 extra steps before it can no longer mimic a move of the attacker—which can only happen due to running into a boundary. We now assert that the 5 extra moves are sufficient for D to reestablish the starting conditions without being captured. This is illustrated in Figure 1.5 (a), where only a small portion of the grid surrounding the players is shown for clarity. With its 5 moves (shown labeled $1,2,\ldots,5$), the defender is able to restore the initial condition with respect to the right boundary. During this maneuver, the defender maintains a safe distance from A, and therefore is not captured.

Of course, the defender earns its five extra moves *gradually*, and not at once, but it is easy to see that the defender can plan and execute these extra moves (amortize, so to speak) during the at least $n/4$ moves it makes mimicking A, as it earns them. In particular, D always "rotates" around the attacker in the direction of the farther of the two boundaries, which must be at distance at least $n/2$. D cannot run into a boundary because the closest one is at least $n/4$ away and it completes its rotation in $\frac{n}{4}-3$ turns, during which the 5 extra moves will never decrease the defender's distance to that boundary. The new target boundary is at least $n/2$ away and once the attacker finishes its rotation must still be $n/4$ away. This is because there are at most $\frac{n}{4}-3$ moves in this direction resulting from moves mimicking A and the three additional moves from the rotation. Thus, after the rotation, the defender is $n/4$ away from a boundary and the attacker is $\frac{n}{4}+3$ from the same boundary, with A and D both in the same row or column. Thus the defender can continue this strategy forever and avoid capture.

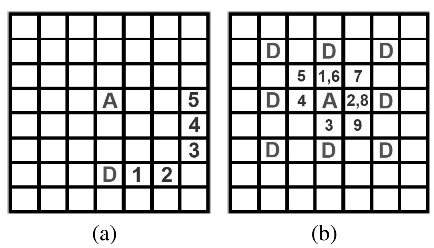

Fig. 1.5 Figure (a) illustrates the proof of Theorem 6: the defender uses 5 extra moves to reestablish the initial conditions. Figure (b) illustrates Theorem 7: the attacker can capture seven of the maximum possible eight defenders using 9 steps, and return to its original position in the 10th step.

1.6.2 Speedier Attacker Against Multiple Defenders

We now consider the speed advantage of attacker against multiple defenders. We showed earlier that in the standard model, the unit-speed attacker wins against two unit-speed defenders. However, the game against more than two defenders remains unsolved. In the following we show that with a constant factor speed advantage, a single attacker can win against any number of defenders.

Theorem 7. *An attacker with speed s can indefinitely avoid capture against $s+1$ defenders, for $s < 10$. An attacker with speed $s = 10$ can avoid capture against any number d of defenders.*

Proof. Let us first consider $s < 10$. The attacker follows a lazy strategy, which is to sit idly unless it is in danger of being captured in the defenders' next turn. Specifically, if no defenders are in $N^2(A)$, the attacker is safe (by Lemma 1). If some defenders enter $N^2(A)$, then the attacker can capture the defender closest to it using Lemma 1, in a single elementary step, with $s-1$ steps (and at most s defenders) remaining before his turn is up. We repeat the argument from the new location of A, until either A is safe for the next turn of the defenders, or it has captured all but one defenders. Thus, either A can remain safe indefinitely, or if only one defender remains it can win.

When $s \geq 10$, we note that due to the minimum separation constraint among the defenders, at most 8 defenders can simultaneously exist inside $N^2(A)$—clearly, no defender lies in $N(A)$ because that is already a captured position, and there are 16 cells in $N^2(A) - N(A)$, and no two consecutive ones can have defenders in them. Figure 1.5 (b) shows A's strategy to capture seven of the maximum possible eights defenders in nine steps, and then return to its original position in the 10th step. It is easy to check that the attacker can achieve a similar result for any configuration of fewer than eight defenders.

The following theorem, whose proof is omitted due to lack of space, shows that an additional increase of speed allows the attacker to capture, and not just evade, any number of defenders.

Theorem 8. *An attacker with speed $s \leq 10$ can capture s or fewer defenders in $O(sn)$ turns. An attacker with speed $s = 11$ can capture any number d of defenders in $O(dn)$ turns.*

1.7 Discussion

We considered a pursuit-evasion game in which two pursuers are required to capture an evader. We modeled this game after Kabaddi, which introduces a new and challenging game of physical capture for mathematical analysis. We believe that Kabaddi offers an elegant and useful framework for studying attack and defensive moves against a team of opponents who can strategically coordinate their counterattacks. Our analysis shows that even with two defenders the game reveals significant complexity and richness.

Our work poses as many open questions as it answers. Clearly, in order to obtain our initial results, we have made several simplification in the game of Kabaddi.

While these simplifications do not affect the relevance of our results to multiagent pursuit-capture, they are crucial for a proper study of kabaddi. The most significant among them is the proper modeling of "holding the breadth" and "safe return." Among the more technical questions, analyzing the game for more then two defenders remains open in the standard model. The minimum separation rule leads to some pesky modeling problems because the attacker could sit in a corner cell and not be captured. So some modification is needed in the rules to avoid such deadlocks. Finally, we have not addressed the game when more than two defenders are required for the capture.

References

1. M. Aigner and M. Fromme. A game of cops and robbers. *Discrete Applied Mathematics*, 8(1):1 – 12, 1984.
2. D. Bienstock and P. Seymour. Monotonicity in graph searching. *J. Algorithms*, 12(2):239–245, 1991.
3. A. Bonato, E. Chiniforooshan, and P. Pralat. Cops and robbers from a distance. *Theoretical Computer Science*, In Press., 2010.
4. A. Frieze, M. Krivelevich, and P.-S. Loh. Variations on cops and robbers. Unpublished manuscript.
5. L. J. Guibas, J.-C. Latombe, S. M. LaValle, D. Lin, and R. Motwani. Visibility-based pursuit-evasion in a polygonal environment. *IJCGA*, 9(5):471–494, 1999.
6. B. Halpern. The robot and the rabbit–a pursuit problem. *The American Mathematical Monthly*, 76(2):140–145, 1969.
7. V. Isler, S. Kannan, and S. Khanna. Randomized pursuit-evasion with local visibility. *SIAM J. Discrete Math.*, 20(1):26–41, 2006.
8. V. Isler and N. Karnad. The role of information in the cop-robber game. *TCS*, 399(3):179 – 190, 2008.
9. N. Karnad and Volkan Isler. Lion and man game in the presence of a circular obstacle. In *IROS*, pages 5045–5050, 2009.
10. S. Kopparty and C. V. Ravishankar. A framework for pursuit evasion games in r^n. *Information Processing Letters*, 96(3):114 – 122, 2005.
11. A. LaPaugh. Recontamination does not help to search a graph. *J. ACM*, 40(2):224–245, 1993.
12. J. E. Littlewood. *Littlewood's Miscellany*. Cambridge University Press, 1986.
13. P. J. Nahin. *Chases and Escapes*. Princeton University, 2007.
14. T. D. Parsons. Pursuit-evasion in a graph. In Y. Alavi and D. R. Lick, editors, *Theory and Application of Graphs*, pages 426–441. Springer-Verlag, 1976.
15. S. Sachs, S. Rajko, and S. M. LaValle. Visibility-based pursuit-evasion in an unknown planar environment. *International Journal of Robotics Research*, 23(1):3–26, 2004.
16. J. Sgall. Solution of david gale's lion and man problem. *Theor. Comput. Sci.*, 259(1-2):663–670, 2001.
17. Wikipedia. Kabaddi. http://en.wikipedia.org/wiki/Kabaddi, 2009.