Hypergraph Games

Lectured by I. B. Leader, Easter Term 2008

Course description

Many natural games may be modelled as follows. We have a board (an arbitrary finite set), and some of its subsets are designated as "winning". Two players take it in turn to play on an (unoccupied) place on the board, and the first player to complete a winning subset is declared the winner. If when the board is full no player has won then the game is a draw.

The above is called the "strong" version of the game. In the "weak" version, also called "maker-breaker", the second player's aim is not to occupy a winning set but just to prevent the first player from doing so. The interest is both for general theorems about games and also in particular games of interest, like the Hales-Jewett game (multi-dimensional noughts & crosses). Roughly speaking, a fair amount is known for maker-breaker while nothing at all is known for strong games.

The aim of the course is to give an introduction to this area. There are some very elegant and appealing results, and there are also many open problems. We hope to cover the following material.

General games

Basic examples. Weight functions; the Erdős-Selfridge theorem. The local lemma for games. Biased games.

The Hales-Jewett game

Pairing strategies; the Pairing Conjecture and the Ratio Conjecture. Beck's big game / little game approach.

Prerequisites

It would be helpful to have a basic knowledge of the Lovász Local Lemma, but this is by no means essential. It would be of similar use to have met the Hales-Jewett Theorem.

Last updated: Sat 18^{th} Jun, 2011

Please let me know of any corrections: glt1000@cam.ac.uk

Hypergraph Games

Let X be a set (finite unless otherwise stated), and let $H \subset \mathbb{P}(X)$. (H is the "winning lines".)

H is a *hypergraph* on X. Often all $A \in H$ have the same size n, and then H is an *n-graph*.

In a play of the *game* on H, two players P1 and P2 take turns to mark a previously unmarked $x \in X$, with P1 playing first. The first player to occupy all of some $A \in H$ is the winner. If X is filled with no $A \in H$ occupied by one player, then the game is a draw.

Example. "Noughts & Crosses".

o	o	o
×	o	
	×	×

Or: three-dimensional; on a $3 \times 3 \times 3$ board; or on a $4 \times 4 \times 4$ board (known as "Qubic").

In general, for the $[n]^d$-game, the board is $[n]^d = \{1, 2, \ldots, n\}^d$, and the winning lines are the "lines", where:

- a *combinatorial line* is a set of the form

$$\left\{ x = (x_1, \ldots, x_d) \in [n]^d : \begin{array}{ll} x_i = x_j & \forall\, i, j \in I, \\ x_i = a_i & \forall\, i \notin I \end{array} \right\},$$

  where $I \subset [d]$, $I \neq \emptyset$, and $a_i \in [n]$ for each $i \notin I$.

  Say I is the set of *active coordinates.*

  E.g., in $[4]^3$, $\{(1, x, 3) : 1 \leqslant x \leqslant 4\}$ or $\{(1, y, y) : 1 \leqslant y \leqslant 4\}$.

- a *line* is a set of the form

$$\left\{ x = (x_1, \ldots, x_d) \in [n]^d : \begin{array}{ll} x_i = x_j & \forall\, i, j \in I, \\ x_i = x_j & \forall\, i, j \in J, \\ x_i = n + 1 - x_j & \forall\, i \in I, j \in J, \\ x_i = a_i & \forall\, i \notin I \cup J \end{array} \right\},$$

  where $I, J \subset [d]$, $I \cup J \neq \emptyset$, $I \cap J = \emptyset$ and $a_i \in [n]$ for each $i \notin I \cup J$.

  E.g., in $[4]^4$, $\{(x, 3, x, 5 - x) : 1 \leqslant x \leqslant 4\}$.

So, the number of lines is $\dfrac{(n+2)^d - n^d}{2}$.

(Explanation: $(n+2)^d$ – each coordinate is "up", "down" or "constant at $1, 2, \ldots$ or n"; n^d – but not *all* constant; $1/2$ – we have named everything twice.)

Example. "Ramsey Game". Played on $X = \{\text{edges of a } K_N\}$, winning lines are each K_s.

E.g., $s = 3$, trying to make a triangle.

Example. "Binary tree game". $X =$ binary tree of depth n, $H =$ all paths from root to leaves.

In any game, every position has a *value*: one of “P1 wins”, “P2 wins”, and “draw”, obtained by backtracking.

E.g., in Noughts & Crosses, with Noughts first:

o	o	o
×		o
	×	×

P1 win

backtrack:

	o	o
×		o
	×	×

P1 win, as ∃ O-move giving a P1 win

backtrack:

	o	o
×		o
	×	

P1 win, as ∀ X-moves, ∃ a P1 win

A *winning strategy (for P1)* is a function from positions to moves such that in any play of the game in which P1’s moves are determined by this function, P1 wins. (Similarly for *drawing strategy* and *P2 winning strategy*.)

For any game, exactly one of the following ($\star$) holds:

(i) P1 has a winning strategy
(ii) P2 has a winning strategy
(iii) each player has a drawing strategy

(Proof: look at the value of the start position.)

Example. “Unrestricted n-in-a-row”.

Here, $X = \mathbb{Z}^2$, $H = n$ points in a row, horizontally, vertically, or diagonally. If the entire sequence of moves results in no winning line for either player, the game is a draw.

For any such infinite game, with each winning set **finite**, condition ($\star$) above still holds. Indeed, it is enough to show that “P1 has no winning strategy” implies “P2 has a drawing strategy”.

Suppose then that P1 does not have a winning strategy. Here is a way for P2 to play that guarantees at least a draw.

– P1 makes some move x_1
– then P1 has no forced win
– so there exists a move x_2 for P2 giving a position where P1 has no forced win
– continue

We obtain a play of the game in which at no time t did P1 have a forced win. So, in particular, at no time t did P1 occupy a winning line. But winning lines are finite, so P1 did not win in that play of the game.

Proposition 1 ("Strategy Stealing"). A hypergraph game cannot be a P2 win.

Proof. Suppose P2 has a winning strategy, S. Then P1 also has a winning strategy (which is therefore a contradiction) as follows.

P1 plays any move x_1, and then pretends x_1 is empty and that he is the second player, playing with strategy S. If S calls on P1 to move to x_1, he picks some other move instead, etc. Then P1 wins this play of the game – as an extra square filled in cannot hurt you. □

Remarks.
1. Proof is completely non-constructive – it does not give a clue as to how P1 should play to guarantee at least a draw.
2. Proof also applies if X infinite.
3. Alternative version: if no end position is a draw, then the game is a P1 win.

Note that strategy stealing applies to any kind of game in which the starting position is symmetric (between P1 and P2) and an extra move doesn't hurt.

Example. "Hex"

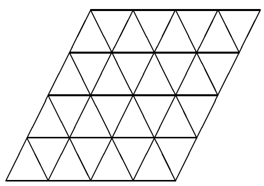

P1 plays vertices, wanting a path from left to right.
P2 plays vertices, wanting a path from top to bottom.
A draw is impossble – nice (but not trivial!) exercise.

Example. "Bridg-it"

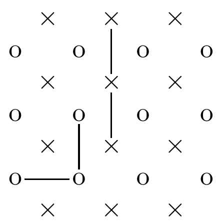

P1 connects adjacent "o"s, wanting a path from left to right.
P2 connects adjacent "×"s, wanting a path from top to bottom.
Edges are not allowed to cross.
Again, a draw is impossible.

Which of the earlier examples are P1 wins?

Examples.

1. Ramsey Game: (K_N, K_s).

Ramsey's Theorem says that any 2-coloured K_N contains a monochromatic K_s if N is large enough. The least such N is called $R(s)$, known to satisfy $\sqrt{2}^s \leqslant s \leqslant 4^s$.

Thus the (K_N, K_s) game is a P1 win if N is large enough (e.g., $N \geqslant 4^s$).

Is there an explicit winning strategy?

$\underline{s = 3}$

– P1 picks some vertex and plays an edge from it.
– P2 plays some edge, which may or may not touch P1's edge. Either way, there exists an endpoint of P1's first move untouched by P2's first move.
– P1 now plays another edge from the vertex he chose.
– P2 must join the endpoints of P1's first two edges together to avoid immediate defeat.
– Then P1 plays yet another edge from the vertex he chose. There are now two edges P2 would need to fill to avoid defeat. Thus P1 wins.

To illustrate:

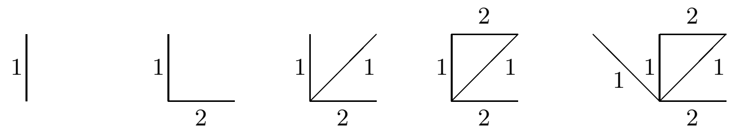

$\underline{s = 4}$

No explicit strategy is known!

The trouble is: (a) P1 must create a K_4 **and** block P2 from an earlier K_4; (b) a winning strategy "corresponds" to a statement of the form

"∃ P1 move : ∀ P2 moves, ∃ P1 move : … " – quantifiers to a large depth.

2. Binary tree game.

An easy P1 win, even doable in time n, by P1 taking the root, and then taking the child of the root in the half of the tree where P2 didn't play, etc.

3. $[n]^d$-game.

Hales-Jewett theorem says that any 2-coloured $[n]^d$ has a monochromatic line (even a combinatorial line – more abstract, and do not depend on an ordering of $[n]^d$) if d is large enough.

So $[n]^d$-game is a P1 win, for d large.

Write $HJ(n)$ for the least d that works. Unfortunately, the best bound known on $HJ(n)$, due to Shelah, is $HJ(n) \leqslant f_4(n+4)$, where

$$
\begin{aligned}
f_1(n) &= 2n \\
f_2(n) &= \text{iterate } f_1 \ n \text{ times} = f_1(f_1(\ldots f_1(1)\ldots)) = 2^n \\
f_3(n) &= \text{iterate } f_2 \ n \text{ times} = \left.2^{2^{\cdot^{\cdot^{\cdot^{2}}}}}\right\} n \text{ times} \\
f_4(n) &= \text{iterate } f_3 \ n \text{ times}
\end{aligned}
$$

E.g., $f_4(1) = 2$, $f_4(2) = 2^2 = 4$, $f_3(2) = 2^{2^{2^2}} = 65536$, $f_4(4) = \left.2^{2^{\cdot^{\cdot^{\cdot^{2}}}}}\right\} 65536 \text{ times}$

The f_k are called the "Grzegorcyk" or "Ackermann" hierarchy.

Warning. Adding a winning line to a P1-win game may not result in another P1-win game: "non-monotonicity".

E.g.,

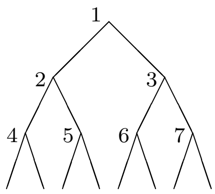

Binary tree of depth 4, plus the winning line $\{3, 6, 5\}$.

– P1 plays 1 (for if not, P2 plays 1 and P1 cannot win).
– Then P2 plays 3. P1 must play 2 (or else P2 will).
– Then P2 plays 5. P1 must play 6 (or else P2 wins).
– Then P2 plays 4, and the game is a draw.

What about the *disjoint union* of games? That is, G and H – with disjoint ground sets.

If G, H are both draws, then $G \sqcup H$ is a draw. (P2 just follows P1 into G or H respectively, playing his drawing move.)

If G is a draw and H a P1 win, then $G \sqcup H$ is a P1-win. (P1 plays out his win in H, always following P2 into G if P2 plays there.)

If G, H are both P1 wins, must $G \sqcup H$ be a P1 win? (See Proposition 3, later.)

$[n]^d$-game : small values

$[3]^2$-game and $[4]^2$-game are draws (by case analysis).

$[5]^2$-game is a draw, as follows.

Consider the labelling on the grid to the right.

We have two points on each of the 12 winning lines, disjointly, so it is a draw (for P2), by playing in the paired square to where P1 plays.

					↙ 11
12	1	8	1	11	← 1
6	2	2	9	10	← 2
3	7		9	3	← 3
6	7	4	4	10	← 4
11	5	8	5	12	← 5
↑ 6	↑ 7	↑ 8	↑ 9	↑ 10	↖ 12

Such a strategy is called a *pairing strategy.*

$[6]^2$-game has a similar pairing strategy.

Hence there is a pairing strategy for the $[n]^2$-game for all $n \geqslant 5$, because we can pass from n to $n+2$, as shown to the right.

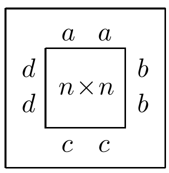

$[3]^3$-game is a P1 win (by easy case analysis).
$[4]^3$-game ("Qubic") is a P1 win (by a huge case analysis - Patashnik).
$[5]^3$-game is believed to be a draw.
$[5]^4$-game is believed to be a P1 win.

What about unrestricted n-in-a-row?

3- and 4-in-a-row are easy P1 wins.

30-in-a-row? Tile $\mathbb{Z}^2$ with copies of $[5]^2$ as shown to the right.

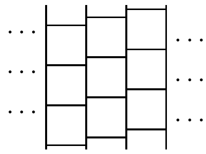

Any line of 30 meets some $[5]^2$ in a winning line for that $[5]^2$.
So a draw, by P2 using pairing in each $[5]^2$.

9-in-a-row? This time, tile $\mathbb{Z}^2$ with copies of 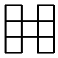 as shown below.

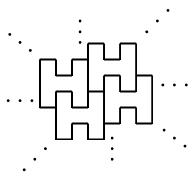

Easy to check that each 9-in-a-row meets some H in one of

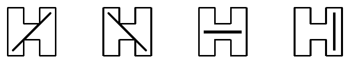

and P2 can stop these being filled by P1, in any H.

8-in-a-row is known to be a draw (similar).

5-in-a-row is unknown.

Warning. On an infinite board, we can have a game that is a P1 win, but not in bounded time. (I.e., there is no k for which P1 can guarantee to win in $\leqslant k$ moves.)

We can even have this when all winning lines have size $\leqslant n$, for some n, and bounded maximum degree (i.e., each point is in at most K winning lines, for some K).

Let B = binary tree game of depth 4.

Let T_n be 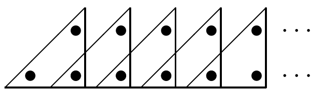 (n times). (The triangles mark the winning lines.)

We can allow our opponent to play in T_n twice before we need to respond. We might then have to keep responding n times, but it is then a draw.

The game is the disjoint union of $B, T_1, T_2, T_3, \ldots$

(Easy to check that it is a P1 win, but not in bounded time.)

Open question. Could it be that unrestricted 5-in-a-row is a P1 win, but not in bounded time?

Hex. No winning strategy for P1 is known.

Bridg-it. A winning strategy is known, from ...

Shannon Switching Game. Have a graph G. P1 and P2 take turns marking (distinct) edges. P2 wins if he marks all edges of a spanning tree (i.e., if his edges connect all of the vertices). Otherwise it is a P1 win.

E.g., G has a bridge $\Rightarrow$ P1 win: 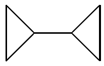 – P1 plays the bridge.

E.g., 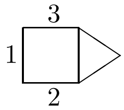

– P1 plays 1; P2 plays one of 2 or 3, but then P1 plays the other.

When does P2 win?

Proposition 2 (Lehmann). P2 wins the Shannon Switching Game on G iff G has two edge-disjoint spanning trees.

Note. Slightly easier to prive this for all multigraphs (i.e., allowed more than one edge between two vertices).

Proof. ($\Rightarrow$) P1 plays any first move, then applies strategy stealing (i.e., ignores his first move, then pretends he is following P2's winning strategy). Then P1 makes a spanning tree, and it is disjoint from P2's spanning tree.

($\Leftarrow$) Whenever P1 takes an edge from one of the two spanning trees, thus splitting it into two components A and B, P2 chooses an edge e in the other spanning tree that joins A to B. Contracting e, we have a graph G' that still has two disjoint spanning trees.

So P2, by induction, can make a spanning tree in G'. Hence, in G, the presence of e tells us that P2 has a spanning tree. $\square$

E.g.

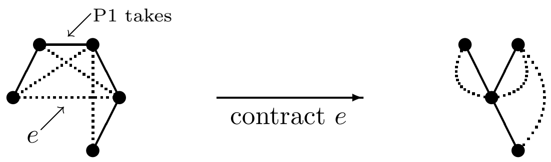

Applying this to Bridg-it.

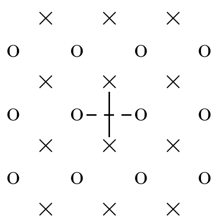

P1 wants to connect left to right. Each P2 move just deletes one allowed edge from the P1 graph. So we may simply play on the P1 vertices, with P2 claiming edges.

Consider the P1 graph:

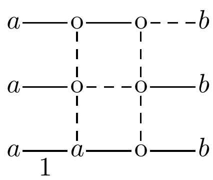

Suppose P1 starts with edge 1. Since P1 wants to connect left to right, we may regard the vertices labelled a as a single vertex, and those labelled b as a single vertex. So P1 wants to connect a to b. In fact, P1 can connect the whole graph (by Proposition 2), because there exists two edge-disjoint spanning trees (shown above as solid and dashed lines). So P1 wins.

Proposition 3. If G and H are P1 wins, then so is $G \sqcup H$.

Notation. For a hypergraph G, the *game-with-pass* is the usual game on G except that P2 can pass whenever he likes. For a game G a P1 win, the *delay* of G is the least k for which P1 has a winning strategy for the game-with-pass on G such that whenever he follows that strategy P2 passes at most k times.

Example.

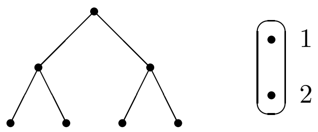

Binary tree of depth 3, disjoint union a pair of points (with that pair being a winning line).

P1 plays 1, then P2 must play 2. But then P2 can pass twice before P1 wins.

Thus this game has delay 2.

Proof of Proposition 3. Let G have delay a, H have delay b, with $a \leqslant b$. Let S be a strategy for the first player in G-with-pass that guarantees a win with the other player passing at most a times. And let T be a strategy for the second player in H-with-pass that guarantees that he can pass at least b times before losing.

Let P1 act as follows. He plays in G, according to S, unless P2 plays in H, in which case P1 replies in H, using T – but if the T strategy calls on him to pass, he plays in G instead (according to S, pretending P2 had just passed there). Continue.

If P2 wins this play of the game on $G \sqcup H$, then P1 passed at least b times in H (by definition of T), so P2 passed at least b times in G, and P1 has not yet won in G. Thus P2 could pass now in G, giving at least $b + 1$ passes in G – contradicting the definition of S. $\square$

Remark. This interpretation of "delay" is due to N. Bowler.

Some open questions.

(1) If $[n]^d$ is a draw, must $[n+1]^d$ be a draw?

(2) More interestingly: if $[n]^d$ is a P1 win, must $[n]^{d+1}$ be a P1 win?

(3) If (K_N, K_s) is a P1 win, must (K_{N+1}, K_s) be a P1 win?

(4) We know that (K_N, K_s) is a P1 win for all N sufficiently large, e.g. $N \geqslant R(s)$. But is (K_ω, K_s) a P1 win? (The board is the complete graph on, say, $\mathbb{N}$.)

(5) What about an infinite disjoint union of (K_N, K_s) games (say, for $N \geqslant R(s)$) – must it be a P1 win?

Warning. Look at 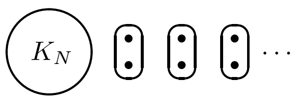

– Each finite part is a P1 win, but the whole game is a draw.

(6) For fixed s, does P1 win (K_N, K_s) in bounded time? I.e., is there a t such that for all $N \geqslant R(s)$, P1 wins (K_N, K_s) in time $\leqslant t$?

In fact:

Proposition 4. Answer to **(4)** yes ⇔ answer to **(6)** yes.

Note. In any play of the (K_N, K_s) game or (K_ω, K_s) game, wlog the k^{th} move is on $K_{[2k]}$. (Proof: induction on k.)

Proof. **(6)** ⇒ **(4)**. We have a bound t, say. In particular, P1 wins the game (K_{2t}, K_s) in at most t moves (and choose $t > \frac{1}{2}R(s)$ so that $2t$ is indeed at least $R(s)$). So let P1 follow a winning strategy for (K_{2t}, K_s) on the K_ω board. This is well-defined, as all the first t moves in K_ω are in $K_{[2t]}$. Thus P1 has won on the K_ω board, in time t.

(4) ⇒ **(6)**. Suppose **(6)** is false, so for every t, there is a game $(K_{f(t)}, K_s)$ in which P2 can stay alive for time at least t, say, following strategy S_t. In K_ω, let P2 play as follows.

P1 makes a move, which has to be $\overset{1}{\bullet}\!\!-\!\!\!-\!\!\overset{2}{\bullet}$. P2 has two choices: join two new points, or join a new point to an old point. Infinitely many of the S_t agree on what P2 should do, so P2 does that. Then P1 moves. P2 chooses a reply that infinitely many of "those infinitely many S_t" agree on. (Again, there are only finitely many choices.)

At time T, P2 has followed, exactly, infinitely many of the S_t, so in particular, has followed some S_t $(t > T)$. So we have lost by time t. □

Remark. This is called a "compactness" argument. (Like: in a metric space, every sequence has a convergent subsequence.)

Question. Is it the case that **(5)** is true iff there is a winning strategy for P1 in $K_N \sqcup K_N$ in which P1 starts in one copy of K_N, keeps playing there, and can even play there for the move immediately after P2's first move to the other copy?

Maker-Breaker Games

In the *maker-breaker* game on hypergraph H, we have the same rules as before, except: P1 (*maker*) wins if he occupies a winning line, and P2 (*breaker*) wins otherwise.

From now on, the previous "win-win" game will be called the *strong* game on H.

Example. $[3]^2$ is a maker win:

3	2	2′
5	1	*
*		4

P1 plays 1. P2 plays 2 or 2′. P3 plays 3. P2 must play 4 or P1 wins. P1 plays 5. Then wherever P2 goes, P1 can take one of * – and we don't care if P2 gets a line.

Why study maker-breaker?

1. Cleaner version of the strong game – P1 now only tries to make a line, not "make and block" at the same time.

2. Gives information about the strong game – e.g., breaker wins ⇒ strong game is a draw.

3. We now have monotonicity – adding a line cannot hurt maker.

So, e.g., for all n, there is a "breakpoint" d_0 with $[n]^d$ a breaker win if $d \leqslant d_0$, but a maker win if $d > d_0$. And for all s, there is a breakpoint N_0 with (K_N, K_s) a breaker win if $N \leqslant N_0$, but a maker win if $N > N_0$.

(So far we know: $d_0 \leqslant$ some iterated tower, and $N_0 \leqslant 4^s$. How about lower bounds?)

4. "A. Thomason Philosophy" – strong game is too delicate to be interesting, as P1 always has an advantage of only one move.

Pairing Strategies

Recall: a *pairing strategy* consists of a choice of two points in each line, distinctly. Then the game is of course a breaker win. (E.g., $[3]^2$.)

In general, when is there a pairing strategy? Clearly need: any k lines contain at least $2k$ points. This is sufficient, by polygamous Hall.

Hall's Theorem says: to marry off some boys, each to a girl he knows, it is necessary and sufficient that for all k, any k boys know at least k girls.

To deduce polygamous Hall ("can marry off each boy at **two** girls he knows iff any k boys know at least $2k$ girls"), just replace each boy with two clones.

Remark. If all lines have size n, it is enough to have maximum degree $\leqslant n/2$. I.e., each point is on $\leqslant n/2$ lines. Indeed, given k lines seeing ℓ points:

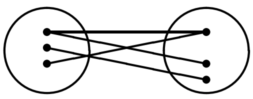

The number of edges:
$= kn$ (counting from the left)
$\leqslant \ell n/2$ (counting from the right)
and so $\ell \geqslant 2k$.

Thus:

Proposition 5. In $[n]^d$ game, breaker has a pairing strategy for $d \leqslant \dfrac{\log n - \log 2}{\log 3}$.

Proof. Maximum degree is $\leqslant 3^d$ (as each coordinate is increaing, decreasing, or constant). So done if $3^d \leqslant n/2$. □

I.e., $d \leqslant c \log n$.

How far could pairing ever work? We certainly need: $2\times$ #lines $\leqslant$ #points.

In other words: $(n+2)^d - n^d \leqslant n^d$ (see page 1). I.e., $(1+2/n)^d \leqslant 2$, so $d \log(1+2/n) \leqslant \log 2$.

For n large, $\log(1+2/n) \approx 2/n$. Thus there is no pairing strategy beyond $d = \dfrac{\log 2}{2} n$.

We only got $d = c \log n$ because in $[n]^d$ the maximum degree is much more than the average degree. So we'll try to reduce the maximum degree using *truncation*: throw out a few points from each line (e.g., those with many coordinates equal).

Proposition 6. In the $[n]^d$ game ($n \geqslant 12$), we can choose for each line L a subset L', with $|L'| \geqslant \frac{n}{2} - 1$, such that no point belongs to more than $2d^{4d/n}$ of the various L'.

(I.e., the hypergraph formed by the L' has maximum degree $\leqslant 2d^{4d/n}$.)

We will then have:

Corollary 7. In the $[n]^d$ game, breaker has a pairing strategy if $d \leqslant n/8$ (for $n \geqslant 12$).

Proof. We'll find a pairing for the L' (instead of the L). Put $d = n/8$.

Maximum degree is $\leqslant 2d^{4d/n}$, and each line has $\geqslant \frac{n}{2} - 1$ points, so done if $2d^{4d/n} \leqslant \frac{1}{2}(\frac{n}{2} - 1)$. I.e., if $d^{4d/n} \leqslant \frac{n}{8} - \frac{1}{4}$. With $d = n/8$, we get $\sqrt{n/8} \leqslant \frac{n}{8} - \frac{1}{4}$, which is correct. □

So we can get $d = cn$.

Hales-Jewett conjecture. There exists a pairing strategy iff #lines $\leqslant \frac{1}{2}$#points. (I.e., roughly, $d \leqslant \frac{\log 2}{2} n$.) Unknown, but see later.

Gammill's conjecture. Maker wins if #lines $\leqslant$ #points. (I.e., roughly, $d \leqslant \frac{\log 3}{2} n$.)

Citrenbaum's conjecture. Maker wins if $d > n$. See later.

Proof of Proposition 6. Given a line L, say $k \in [n]$ is *overused* if the number of constant coordinates equal to k or $n + 1 - k$ is greater than $4d/n$. (We expect each choice to appear about $1/n$ of the time, and so this pair about $2/n$ of the time. Twice that counts as overuse.)

Note that the number of overused k is $\leqslant n/2$. (For if $> n/2$, then we have $> (n/4)(4d/n) = d$ constant coordinates – contradiction.)

Let $L' = \{x \in L : x \text{ is the } k^{\text{th}} \text{ point in } L, \text{ where } k \text{ not overused}, k \neq (n+1)/2\}$. This is well-defined. Certainly $|L'| \geqslant n/2 - 1$.

Given $x \in [n]^d$, in how many lines do we have: x is k^{th} point of L and $x \in L'$?

Let $N_k = \{i : x_i = k \text{ or } n + 1 - k\}$. Then active coordinates are in N_k, and are all but at most $4d/n$ of N_k, else k is overused. Thus #lines $\leqslant \left|N_k^{(\leqslant 4d/n)}\right|$.

$$
\begin{aligned}
\text{So degree of} \quad x &\leqslant \sum_k \left|N_k^{(\leqslant 4d/n)}\right| \quad \leqslant \quad d^{(\leqslant 4d/n)} \\
&= \binom{d}{4d/n} + \binom{d}{4d/n - 1} + \ldots + \binom{d}{1} + \binom{d}{0} \\
&\leqslant 2\binom{d}{4d/n} \text{ (as } 4d/n \leqslant d/3) \\
&\leqslant 2d^{4d/n}
\end{aligned}
$$

□

Hales-Jewett conjecture said:

$$\text{if \#lines} \leqslant \frac{1}{2}\text{\#points (in } [n]^d\text{) then for all } S \subset [n]^d \text{ we have \#lines in } S \leqslant \frac{1}{2}|S|.$$

Equivalently, average degree $\leqslant n/2 \Rightarrow$ average degree on $S \leqslant 1/2$.

Equivalently, $\dfrac{\text{\#lines}}{\text{\#points}} \leqslant \dfrac{1}{2} \;\Rightarrow\; \dfrac{\text{\#lines in } S}{|S|} \leqslant \dfrac{1}{2}$.

This is a special case of the **ratio conjecture**:

$$\forall\, n, d,\ \forall\, S \subset [n]^d : \frac{\text{\#lines in } S}{|S|} \leqslant \frac{\text{\#lines in } [n]^d}{n^d}.$$

More generally, is it true that $|S| = n^e \Rightarrow \dfrac{\text{\#lines in } S}{|S|}$ is maximised (over $S \subset [n]^d$) when S is an e-dimensional hypercube?

This is false! D. Christofides found a set of size 3^5 in $[3]^6$ with more lines than $[3]^5$ has. Even worse, he found that the ratio conjecture is false for every n.

Proof for $n = 3$. Here, allow lines counted each way, and allow constant lines.

So #lines in $[3]^d$ is 5^d. So $\dfrac{\text{\#lines}}{\text{\#points}} = (5/3)^d$.

We will remove from $[3]^d$ the set $T = \{x \in [3]^d : \text{\#coordinates} \in \{1, 3\} \text{ is } k\}$. So $T = 2^k \dbinom{d}{k}$.

How many lines meet T?

How many with midpoint $x \in T$? E.g., $x : \underbrace{1111333}_{k}\underbrace{222222}_{\text{active coordinates live in here}}$

Thus, we have $\dbinom{d}{k} 2^k 3^{d-k}$.

(I.e., choose "1 or 3" k times, and "constant", "up", or "down" $d-k$ times.)

How many with midpoint $x \notin T$? (So each endpoint $\in T$.) E.g., $x : \underbrace{1111333\,\overbrace{222}^{\text{active}}}_{k}222$

Thus, we have $\dbinom{d}{k} 4^k - \dbinom{d}{k} 2^k$.

(I.e., choose 1, 3, "up", "down" k times, and avoid counting constant lines in both counts.)

So we're done if $\dfrac{\binom{d}{k}2^k 3^{d-k} + \binom{d}{k}4^k - \binom{d}{k}2^k}{\binom{d}{k}2^k} < \left(\dfrac{5}{3}\right)^d$. I.e., if $3^{d-k} + 2^k - 1 < (5/3)^d$.

Put $k = \alpha d$ (α fixed, d large). We want $3^{(1-\alpha)d} + 2^{\alpha d} - 1 < (5/3)^d$.

So it is enough to have $3^{1-\alpha}\, 2^\alpha < 5/3$. I.e., we need $\alpha < \dfrac{\log 5/3}{\log 2}$ and $\alpha > \dfrac{\log 9/5}{\log 3}$.

So we're done if $\dfrac{\log 9/5}{\log 3} < \dfrac{\log 5/3}{\log 2}$, i.e. if $\dfrac{\log 9/5}{\log 3/2} < \dfrac{\log 3}{\log 2}$, which is true. □

So far, the breakpoint for the $[n]^d$-game: $c\,n \leqslant d_0 \leqslant$ iterated tower.

And for the (K_N, K_s)-game: $s+1 \leqslant N_0 \leqslant 4^s$.

(Why $s+1$? #lines $= \binom{N}{s}$, #points $= \binom{N}{2}$. So need $\binom{N}{s} \leqslant \frac{1}{2}\binom{N}{2}$ – false for $N = s+2$.)

Question. In a hypergraph game H with all the lines of size n (i.e., "H is an n-graph"), how **few** lines can we have in a maker win?

E.g., can achieve 2^{n-1} in the binary tree game.

E.g., also in 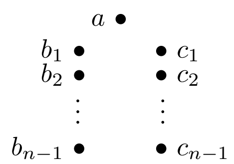 with winning lines $\{a,\ (b_1 \text{ or } c_1),\ (b_2 \text{ or } c_2), \ldots\}$.

Can we get smaller than 2^{n-1}?

Theorem 8 (Erdős-Selfridge Theorem). H an n-graph with $|H| < 2^{n-1}$ implies the game on H is a breaker win.

Remarks.
1. Best possible as there exists H, $|H| = 2^{n-1}$ that is a maker win.
2. Given $|H| < 2^{n-1}$ there **exists** a draw position – i.e., each point occupied by P1 or P2 with neither having a winning line. This is because of the following.

Consider assigning each point to P1 or P2 at random (i.e., independently, with probability 1/2). For a line $L \in H$, $P(L \text{ monochromatic}) = 1/2^{n-1}$. (As $= 2 \times (1/2)^n$.)

So $P(\exists\,\text{monochromatic line}) \leqslant |H|/2^{n-1} < 1$.

So there exists an assignment of the points in such a way that no line is monochromatic. □

Proof of Erdős-Selfridge. For a position P (say, maker has played $x_1, \ldots, x_k$; breaker has played $y_1, \ldots, y_l$) define the *danger* of P as $D(P) = \sum_{L \in H} D(l)$, where

$$D(L) = \begin{cases} 0 & \text{if breaker has played in } L \\ 2^{-\#\,\text{empty in } L} & \text{if not} \end{cases}$$

So $D(P)$ is the expected number of winning lines for maker, if the rest of the board is filled in, each point at random. (If the board is large, this is the same as saying "if players subsequently play at random".)

Let breaker always move so as to minimise the danger. Suppose it is breaker's turn to move, in position P, with danger $D(P) = D$.

Claim. After one move each, danger $\leqslant D$.

Note. Then done. The initial danger $< 1/2$, so after maker's first move, we have danger < 1 (as the absolute worst case is that his move is in every line, doubling the danger of each). So after **each** move of maker, we have danger < 1, so maker did not win (or else danger $\geqslant 1$).

Proof of claim. Say breaker moves to y, and then maker moves to x.

Breaker's move: danger decreases by: $\displaystyle\sum_{L:\ y\in L} D(l) = \sum_{\substack{L:\ y\in L,\\ y \text{ is breaker's}\\ \text{first move in } L}} 2^{-\#\text{ empty in } L \text{ in position } P}.$

Maker's move: danger increases by: $\displaystyle\sum_{L:\ x\in L} D(l) = \sum_{\substack{L:\ x\in L,\\ \text{breaker had not played}\\ \text{in } L \text{ at position } P,\\ \textbf{and } y\notin L}} 2^{-\#\text{ empty in } L \text{ in position } P}.$

But breaker's decrease $\geqslant$ maker's increase, by the choice of y. (And the fact that the "**and**" bit may be a further restriction on the sum.) □

Remarks.
1. Proof involves the idea of randomness ("randomisation") in our choice of the "danger" or "potential" of a position.
2. It also involves "derandomisation" – the final proof does not involve a strategy of random moves.
3. Call the n-graph H *extremal* if $|H| = 2^{n-1}$ and the game is a maker win.

Beck's Conjecture. In every extremal game, maker can win in time n.

This is **false**. A. Sanders found extremal games that can takes exponential time for maker to win.

Corollary 9. Breaker wins the (K_N, K_s)-game if $N \leqslant 2^{s/2}$.

Proof. We want $\displaystyle\binom{N}{s} < 2^{\binom{s}{2}-1}$.

We have $\displaystyle\binom{N}{s} \leqslant \frac{N^s}{s!} \sim \frac{N^s}{s^s e^{-s}}$, so want $\displaystyle\frac{N}{s/e} \leqslant 2^{s/2}$, i.e., $\displaystyle N \leqslant 2^{s/2}\,\frac{s}{e}$. □

Thus breakpoint N_0 has: $2^{s/2} \leqslant N_0 \leqslant 4^s$.

What about the $[n]^d$-game?

We want $\displaystyle\frac{(n+2)^d - n^d}{2} < 2^{n-1}$. Note that $n^{d-1} \leqslant (n+2)^d - n^d \leqslant (n+2)^d$.

So want $d \log n < n \log 2$, i.e., $\displaystyle d < \log 2\,\frac{n}{\log n}$ – i.e., **worse** than pairing.

Say H is *simple* if any two lines meet in at most one point.

E.g., lines of $[n]^d$, but not lines of (K_N, K_s).

Theorem 10 (maker form of Erdős-Selfridge). Let H be a simple n-graph on X. Then $|H| > 2^{n-3}|X| \Rightarrow$ maker wins the game on H.

Proof. Using the same "danger" as before, let maker always move to maximise the danger. Initially, danger $> |X|/8$.

Claim. A maker move, followed by a breaker move, decreases the danger by $\leqslant 1/4$.

Note. Then done, as game consists of $\leqslant |X|/2$ such pairs, so final danger > 0, so final position must be a maker win.

Proof of claim. Exactly as before, except we lose $\leqslant 1/4$, from the $\leqslant 1$ line in which both maker and breaker play. (If maker plays x in L, then breaker play y in L, then then danger changes from $1/2^t$ to $1/2^{t-1}$ to 0, for some $t \geqslant 2$.) □

Corollary 11. Maker wins the $[n]^d$-game if $d \geqslant \frac{\log 2}{2} n^2$.

Proof. We want $\frac{(n+2)^d - n^d}{2} > 2^{n-3} n^d$, i.e., $\left(1 + \frac{2}{n}\right)^d - 1 > 2^{n-2}$.

Taking logs, $\frac{2d}{n} > (n-2)\log 2$, i.e., $d > n(n-2)\frac{\log 2}{2}$. □

Thus breakpoint d_0 has: $\frac{n}{8} \leqslant d_0 \leqslant \frac{\log 2}{2} n^2$.

We had: #lines $< 2^{n-1} \Rightarrow$ there exists a draw position.

Lovász local lemma tells us: if each line meets $< 2^{n-3}$ other lines, then there exists a draw position.

Neighbourhood Conjecture. Is there a $c > 0$ such that if each line meet $< c\,2^n$ other lines, the breaker wins?

Equivalently, does maximum degree $< c\,\frac{2^n}{n} \Rightarrow$ breaker win?

Weaker question: is there a $c > 1$ such that max degree $< c^n \Rightarrow$ breaker win? (This is the "game-theoretic local lemma question".)

All these are unknown.

Turning it around: how small can the max degree be in a maker win?

What do we know?

$$\left.\begin{array}{rcll} n/2 & : & \text{no.} & \text{(pairing draw)} \\ 2^{n-1} & : & \text{yes.} & \text{(binary tree)} \end{array}\right\} \text{ Nothing better is known!}$$

E.g., $n = 2$. Max degree 1 – no.
Max degree 2 – yes: 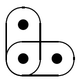

E.g., $n = 3$. Max degree 1 – no.
Max degree 2 – yes:

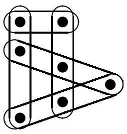 is a P1 win: 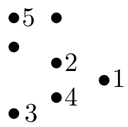

so take: 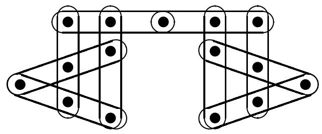

E.g., $n = 4$. Max degree 2 – no.
Max degree 4 – yes ("double example for $n = 3$").
Max degree 3 – unknown!

Recall the (weaker) **Neighbourhood Conjecture**: is there a $c > 1$ such that max degree $< c^n \Rightarrow$ breaker win?

Towards this:

Theorem 12 (Beck). Let H be a simple n-graph on X. Then $|X| \leqslant 2^{n^2/100}$, and max degree $\leqslant 1.5^n \Rightarrow$ breaker wins.

(Really: $\forall\, c < 2 \exists\, \delta : |X| \leqslant 2^{n^2 \delta}$, max degree $< c^n \Rightarrow$ breaker wins.)

We'll need:

The *shutout game* (with parameter b) on a hypergraph H: maker wins if he can occupy b points of some line $L \in H$ when breaker has occupied no points of L.

Lemma 13 (Shutout form of Erdős-Selfridge). If H is a hypergraph with #lines $\leqslant 2^{b-1}$, then breaker wins the b-shutout game on H.

Proof. Let $D(P) = \sum_{L \in H} D(L)$, where $D(L) = \begin{cases} 0 & \text{if breaker has played in } L \\ 2^{-\#\,\text{empty}+|L|-b} & \text{if not} \end{cases}$

Have initial danger $< 1/2$, and maker wins $\Rightarrow$ danger at that time $\geqslant 1$. So done exactly as in Erdős-Selfridge. □

The *shrinking shutout game* on H: at any time after a maker move, maker can declare some empty squares as illegal for all later turns (for both players).

Note. This could hurt breaker: 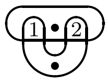

with $b = 2$. Maker plays 1 than declares 2 illegal.

Lemma 14 (Shrinking shutout form of Erdős-Selfridge). If H is a hypergraph, with #lines $< 2^{b-1}$, then breaker wins the shrinking b-shutout game on H.

Proof. Same as for Lemma 13. □

Proof of Theorem 12. Fix an integer k (will be $n/5$).

Idea: "regard $L \in H$ as dangerous if maker has played in it $n-k$ times, without breaker playing in it."

At any time, we will have a big board B and small board S, so $X = B \sqcup S$. After each maker move, some (unoccupied) points may move from B to S. (I.e., S grows, B shrinks.)

Breaker will always play in B if maker does, and in S if maker does. After a maker move, if there exists $L \in H$ such that $L \cap B$ has $n-k$ maker points and no breaker points, move the empty k points of L into S. (Some of these k points might already be in S.) Call that k-set "dangerous".

On S, breaker will play (if he can) a pairing strategy, by (hopefully) each dangerous set having two "private" points that are in no other dangerous set.

So, on B, breaker must play in such a way that neither of the following can ever happen.

(1) Pairing fails.

Dangerous $A, A_1, \ldots, A_{k-1}$ such that $A = \{x_1, \ldots, x_{k-1}\}$, $x_i \in A_i$.

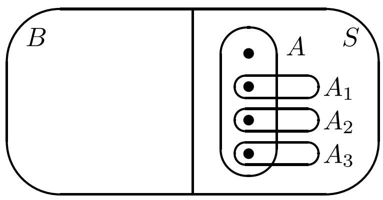

Say, $A_i \subset L_i$, $A \subset L$, $(L, L_1, \ldots, L_{k-1} \in H)$. Then $\{L_1, \ldots, L_{k-1}\}$ are *linked* (there is a line meeting them all), and in big board, maker occupied all but $k(k-1)$ points in $L_1 \cup \ldots \cup L_{k-1}$. But $|L_1 \cup \ldots \cup L_{k-1}| \geqslant (k-1)n - k(k-1)$ (as H simple).

So in big board, maker occupied $\geqslant k(n-2k)$ points of $L_1 \cup \ldots \cup L_{k-1}$.

(2) Maker fills a line $L \in H$.

A has $\geqslant k$ points in S, or else we would have made a dangerous set out of L at some earlier time.

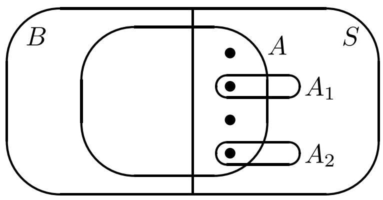

So we get **exactly** the configuration from (a) (and a bit more).

So done by the shrinking shutout Erdős-Selfridge, applied to the hypergraph of all linked $\{L_1, \ldots, L_{k-1}\}$ with $b = k(n-2k)$, provided that #linked sets $\geqslant 2^{k(n-2k)-1}$.

$$\text{But \#linked sets} \leqslant |H|\binom{n}{k-1}D^{k-1} \leqslant |X|Dn^{n-1}D^{k-1} = |X|n^{k-1}D^k$$

(where D = max degree – the D^{k-1} factor is the key improvement on Erdős-Selfridge).

$$\text{So done if } |X| \leqslant \frac{2^{k(n-2k)-1}}{n^{k-1}D^k} \leqslant \left(\frac{2^{n-2k}}{nD}\right)^k \leqslant \left(\frac{2^{3n/5}}{n(1.5)^n}\right)^{n/5} \leqslant \left(2^{n/20}\right)^{n/5} \qquad \square$$

Corollary 15. The $[n]^d$-game is a breaker win if $d \leqslant \dfrac{1}{5000}\dfrac{n^2}{\log n}$ (so Citrenbaum false).

Proof. Apply the above to the set of truncated lines: we had size $\geqslant n/2 - 1$, and the max degree was $2d^{4d/n}$.

Done if $|X| = n^d \leqslant 2^{\frac{n^2/4}{100}}$, and $2d^{4d/n} \leqslant (1.5)^n$. Take logs:

$$\left.\begin{array}{ll} \text{want} & d\log n \leqslant \dfrac{n^2}{400}\log 2 \\ \text{and} & \dfrac{4d}{n}\log d \leqslant n\log 1.5 \end{array}\right\} \quad \text{True if } d = \frac{n^2}{\log n}\,\frac{1}{5000}$$

□