Seeing the (game) trees for the forest

Brant Jones (James Madison University)

Spring MAA MD-DC-VA Section Meeting
April 13, 2019

(Transcript available at `bit.ly/gametrees`)

What is a game?

- moves (made by players)
- outcomes
    - intermediate board positions;
    - eventually win, lose, or tie

Examples of “games?”

chess, checkers, go, mancala, tic-tac-toe
rock-paper-scissors
poker
card games such as hearts, spades or war
jigsaw puzzles
pencil games such as sudoku
lottery
pub trivia, trivial pursuit
video games with a range of motion, race car driving
picking a winning stock
soccer, football
battleship, stratego

Taxonomy?

- There may be 1, 2, or many players.
- Moves may be alternating or simultaneous.
- Moves may be finite or continuous.
- Outcomes may be finite or continuous.
- Games can be deterministic or non-deterministic.
- Games can have perfect information or partial information.

Taxonomy?

- There may be **1, 2**, or many players.
- Moves may be **alternating** or simultaneous.
- Moves may be **finite** or continuous.
- Outcomes may be **finite** or continuous.
- Games can be **deterministic or non-deterministic**.
- Games can have **perfect information or partial information**.

Math problem?

- A **strategy** is a *complete* guide for how a player should move in *every possible* situation.

  (In common language, we tend to blur this with **heuristics**: incomplete advice about how to move in certain situations.)

- A **winning strategy** guarantees a win for the player using it, *regardless of how the other player moves*.

21-flags

- Begin with 21 flags.
- Two players take turns removing 1, 2, or 3 flags.
- The player that takes the last flag (whether alone or part of a group) wins.

Does it matter how many are taken initially? Is there a winning strategy? For which player?

21-flags

- *Moves*: remove **1**, **2**, or **3** flags. *Goal*: remove last flag.

If current player faces	then they should remove	to achieve
1 flag	1 flag	win
2 flags	2 flags	win
3 flags	3 flags	win
4 flags	doesn't matter (1→3, 2→2, 3→1)	all lose

21-flags

- *Moves*: remove **1**, **2**, or **3** flags. *Goal*: remove last flag.

If current player faces	then they should remove	to achieve
5 flags	1 flag	opp. faces 4 flags
6 flags	2 flags	opp. faces 4 flags
7 flags	3 flags	opp. faces 4 flags
8 flags	doesn't matter ($1 \to 7$, $2 \to 6$, $3 \to 5$)	all (eventually) lose
$4k + i$	$i \neq 0$	opp. faces $4k$ flags
$4k$	doesn't matter	all lose
21 flags	1 flag	opp. faces 20 flags

Generalizations

21-flags has a winning *corner* strategy,

$\longrightarrow$ Same for **Nim**, but corners are defined with binary/xor (Bouton, 1901),

$\longrightarrow$ Same for *any* (alternating, finite, perfect information) *impartial move* game! (Sprague, 1935; Grundy, 1939).

Example

Benesh–Ernst–Sieben, 2018, consider the game in which two players alternately select elements from a finite group until their union generates the group. Found Nim correspondence (hence, winning strategy) for cyclic, abelian, dihedral, symmetric, alternating, nilpotent, ...

Combinatorial game theory (Berlekamp–Guy–Conway, 1970's), tries to generalize these ideas to *partisan move* games.

Game trees

To analyze a game we:

- *play forward*, recording all possible outcomes/moves

$$21 \to 20 \to 19 \to 18 \to \cdots \to 5 \to 4 \to 3 \to 2 \to 1 \to 0$$

- *solve backwards*, finding best outcome for current player

$$21 \to 20 \to 19 \to 18 \to \cdots \to 5 \to 4 \to 3 \to 2 \to 1 \to 0$$

Theorem (Zermelo, 1913; von Neumann–Morgenstern, 1944)

This procedure finds the optimal strategy for any *deterministic game with alternating finite moves and perfect information.*

https://xkcd.com/832/

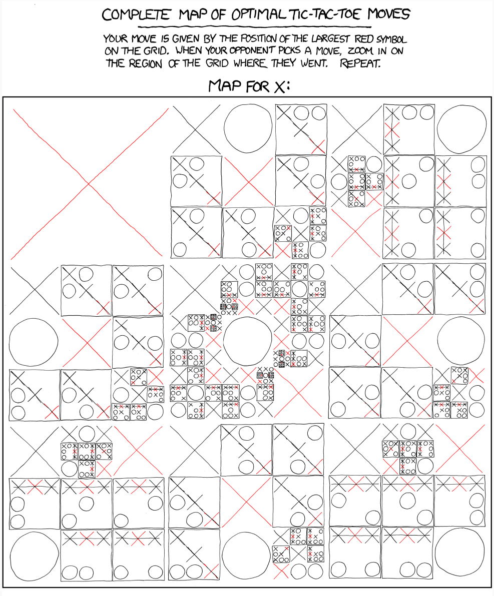

Computational results

Game	Optimal strategy	Team who computed the game tree
Tic-tac-toe	tie	
21-flags, Nim	win for player not in a corner state	Bouton (1901)
Connect Four	win for first player	Allen and Allis (1988)
Checkers	tie	Schaeffer (2007)
Mancala	win for the first player	Irving, Donkers and Uiterwijk (2000), Carstensen (2011)
Chess	?	
Go	?	

Checkers game tree has

$$500,000,000,000,000,000,000 \text{ nodes}$$

which required running a program on more than 50 computers for over 18 years.

Hex (Nash–Hein, 1950’s)

- Moves = label hexagon with your initial
- First player win = path connecting top and bottom
- Second player win = path connecting left and right

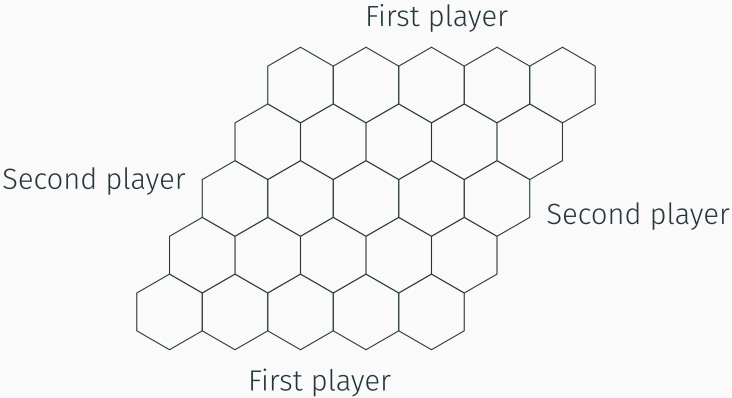

Hex (Nash–Hein, 1950's)

- Moves = label hexagon with your initial
- First player win = path connecting top and bottom
- Second player win = path connecting left and right

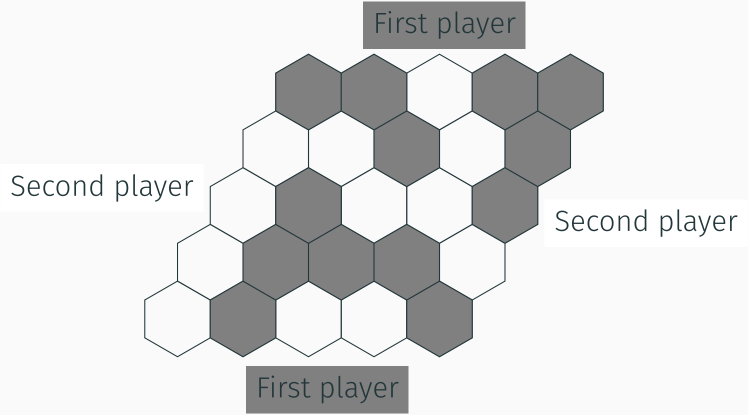

Hex (Nash–Hein, 1950's)

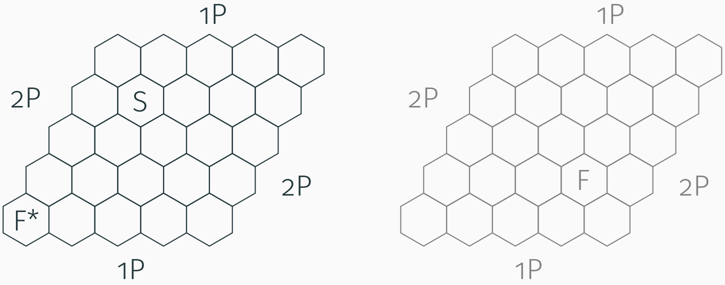

There are no ties, so winning strategy exists for one player.

Suppose (for contradiction) it is Second player.

Then, First player can make dual board (ignore first F; transpose all positions) and "steal" Second's strategy.

Mancala (North American Kalah)

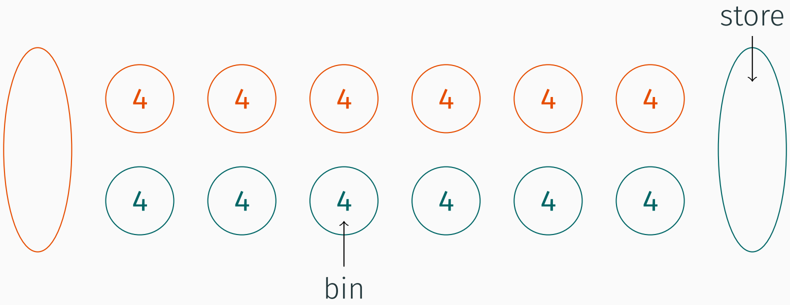

In Mancala play, pick up all the seeds in a bin and **sow** them, placing one seed in each bin to the right.

If the last seed lands in the store then play again.

Win if you have more seeds in your store than your opponent.

Sowing move:

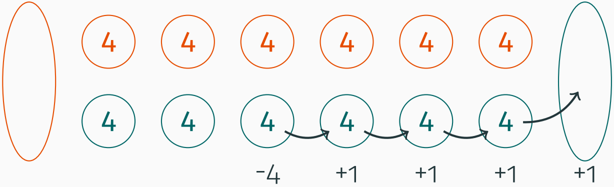

Mancala

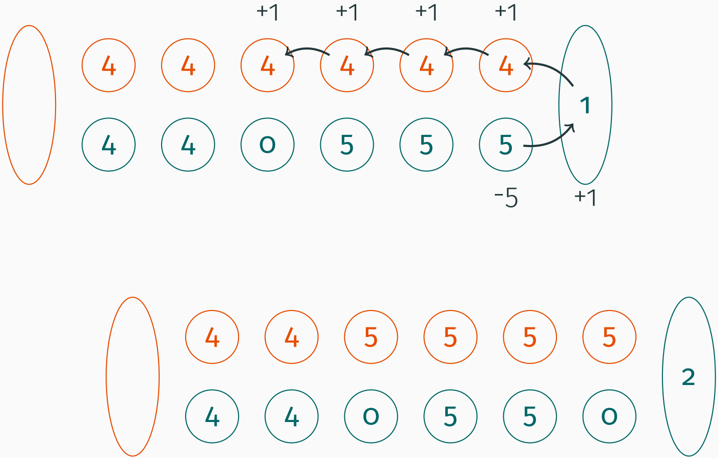

Mancala

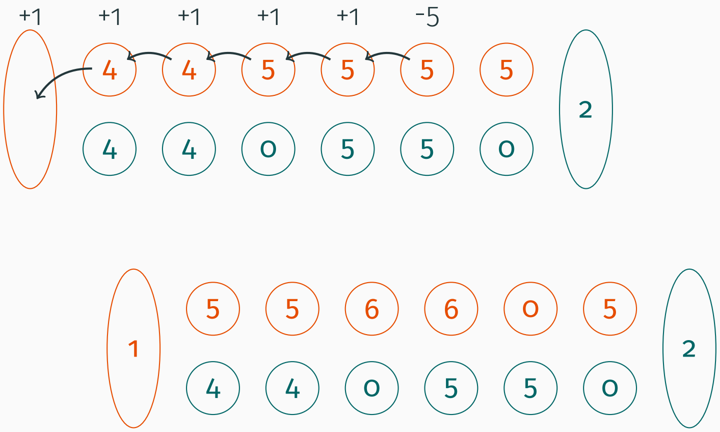

Tchoukaillon (Gautheron, 1970's)

Sweep moves:

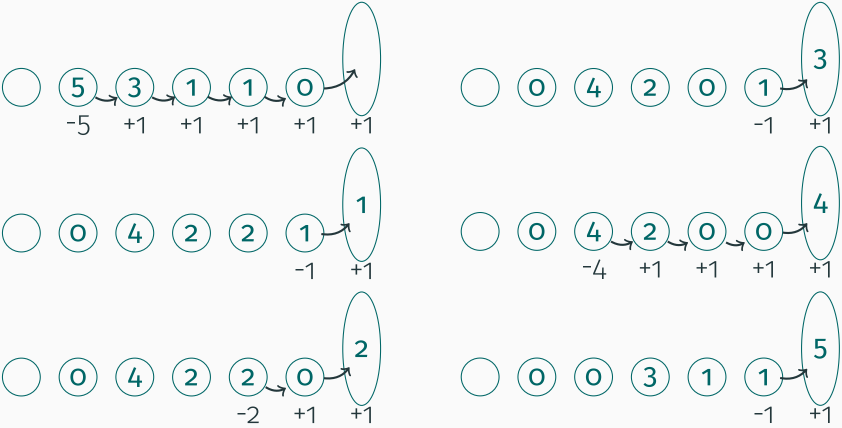

Tchoukaillon (Gautheron, 1970's)

Sweep moves:

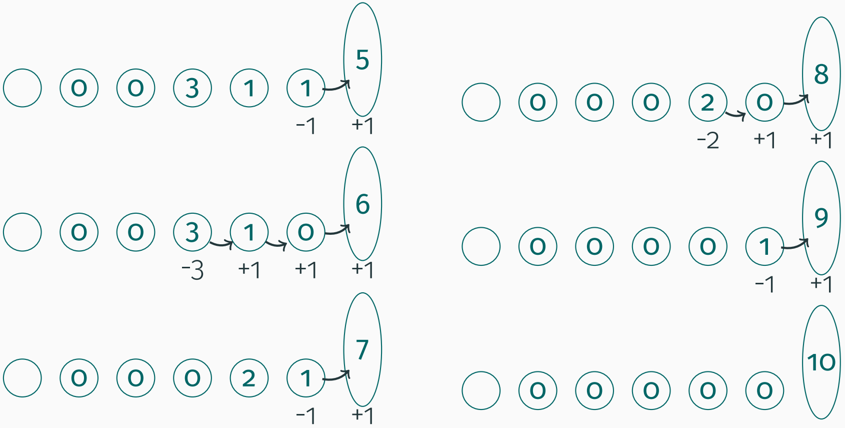

Tchoukaillon

> *"In any mancala game that includes the rule that a player can move again if a sowing ends in [their] own store, these [Tchoukaillon] positions are important. These games include Kalah, Dakon, Ruma Tchuka and many others. … Also mancala games that use the 2-3 capture rule and have no stores (like Wari and Awale) benefit from Tchoukaillon positions."*

-Jeroen Donkers, Jos Uiterwijk, Alex de Voogt, "Mancala games - Topics in Mathemathics and Artificial Intelligence", *The Journal of Machine Learning Research*, 2001

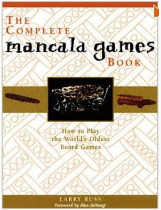

Tchoukaillon

Some data:

b_7	b_6	b_5	b_4	b_3	b_2	n	ℓ
0	0	0	0	0	1	1	1
0	0	0	0	2	0	2	2
0	0	0	0	2	1	3	2
0	0	0	3	1	0	4	3
0	0	0	3	1	1	5	3
0	0	4	2	0	0	6	4
0	0	4	2	0	1	7	4
0	0	4	2	2	0	8	4
0	0	4	2	2	1	9	4
0	5	3	1	1	0	10	5
0	5	3	1	1	1	11	5
6	4	2	0	0	0	12	6

Easy facts from “playing” backwards:

- There **exists** a **unique** vector for every nonnegative n.

$\text{Tchoukaillon}(n) = \frac{d}{dx}\left[n \bmod x\right]$

Unexpected pattern with attractive slogan!

Example

Say $n = 17$. Find representatives for $n \bmod x$ that are *increasing*:

x	2	3	4	5	6	7	8	9	$\cdots$
$n \bmod x$	1	2	1	2	5	3	1	8	
$n \ {}_{\mathrm{mo}}\mathrm{d}\ x$	1	2	5	7	11	17	17	17	
$\Delta\left[n \ {}_{\mathrm{mo}}\mathrm{d}\ x\right]$	1	1	3	2	4	6	0	0	

Data check:

b_7	b_6	b_5	b_4	b_3	b_2	n
0	0	0	3	1	1	5
6	4	2	0	0	0	12
6	4	2	3	1	1	17

Tchoukaillon

b_7	b_6	b_5	b_4	b_3	b_2	n	ℓ	c_7	c_6	c_5	c_4	c_3	c_2
0	0	0	0	0	1	1	1	1	1	1	1	1	1
0	0	0	0	2	0	2	2	2	2	2	2	2	0
0	0	0	0	2	1	3	2	3	3	3	3	3	1
0	0	0	3	1	0	4	3	4	4	4	4	1	0
0	0	0	3	1	1	5	3	5	5	5	5	2	1
0	0	4	2	0	0	6	4	6	6	6	2	0	0
0	0	4	2	0	1	7	4	7	7	7	3	1	1
0	0	4	2	2	0	8	4	8	8	8	4	2	0
0	0	4	2	2	1	9	4	9	9	9	5	3	1
0	5	3	1	1	0	10	5	10	10	5	2	1	0
0	5	3	1	1	1	11	5	11	11	6	3	2	1
6	4	2	0	0	0	12	6	12	6	2	0	0	0

Theorem (J.–Taalman–Tongen)

For all n, $b_i(n) = c_i(n) - c_{i-1}(n)$, *and* $c_i(n) = \sum_{j=2}^{i} b_j(n)$.

Corollary

For all i, $(b_2(n), b_3(n), \ldots, b_i(n))_{n=0}^{\infty}$ *is periodic with period* $\operatorname{lcm}(2, 3, \ldots, i)$.

Theorem (Broline–Loeb)

As $n \to \infty$, $\ell(n) \sim \sqrt{\pi n}$.

Open: Circular/Affine Tchoukaillon?

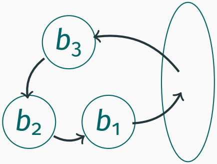

b_3	b_2	b_1	n
0	0	1	1
0	2	0	2
0	2	1	3
3	1	0	4
3	1	1	5
			6
2	0	5	7
2	2	4	8
			9
			10
0	10	1	11

Example

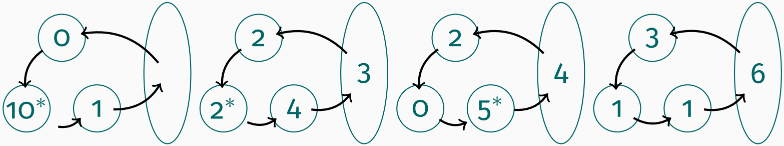

Best choice problem

Given a **uniformly random** permutation of N, with entries revealed **sequentially**, choose the **best (maximum) value**.

Example

$$\pi = \underbrace{\underbrace{\underbrace{\underbrace{2}_{1}\ 5}_{12}\ 1\ 6}_{2314}\ 3\ 7}_{241536}\ 4$$

What is the **optimal strategy** and **probability of success**?

Theorem (Lindley, 1961; Flood–Robbins, 1950's)

Reject the first N/e entries, and select the next left-to-right maximum. This succeeds $1/e \approx 37\%$ of the time.

Best choice problem

Given a **uniformly random** permutation of N, with entries revealed **sequentially**, choose the **best (maximum) value**.

Example

$$\pi = \underbrace{\underbrace{\underbrace{\underbrace{2}_{1}\ 5}_{12}\ 1\ 6}_{2314}\ 3\ 7}_{241536}\ 4$$

What is the **optimal strategy** and **probability of success**?

Theorem (Lindley, 1961; Flood–Robbins, 1950's)

Reject the first N/e *entries, and select the next left-to-right maximum. This succeeds* $1/e \approx 37\%$ *of the time.*

$N = 4$ prefix tree

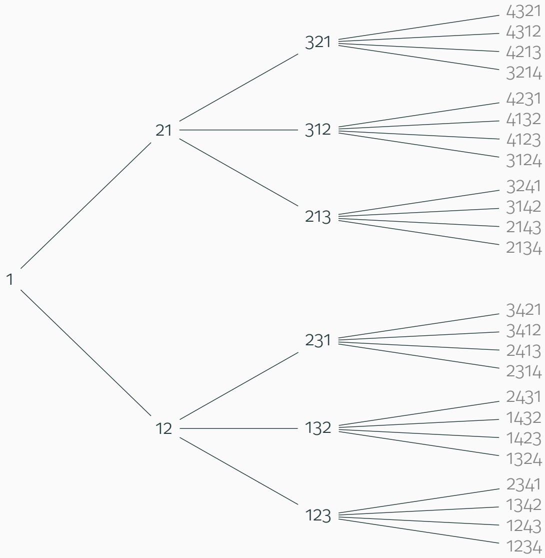

$N = 4$ prefix tree

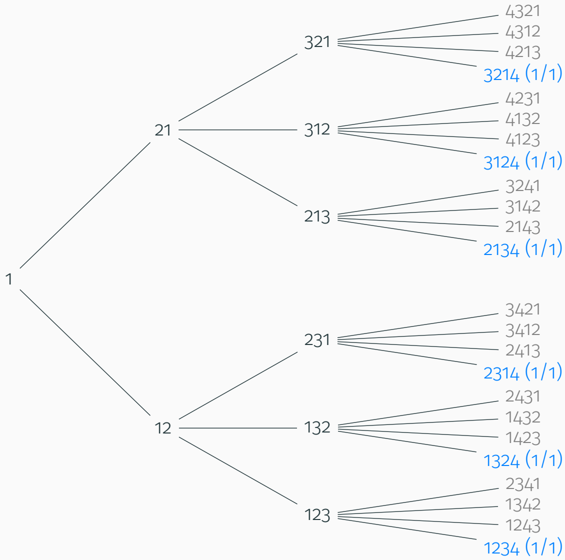

$N = 4$ prefix tree

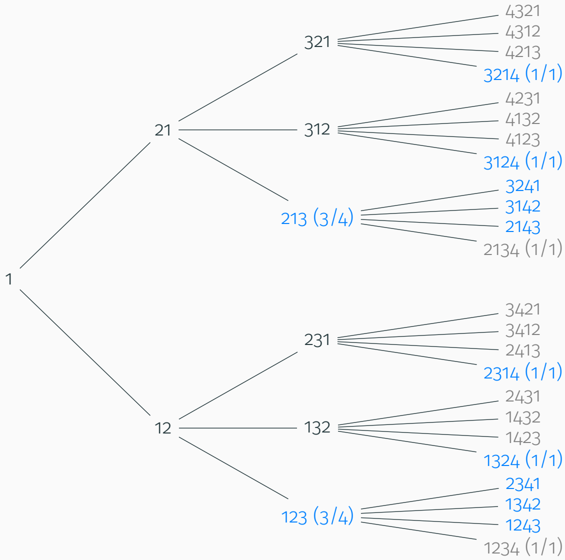

$N = 4$ prefix tree

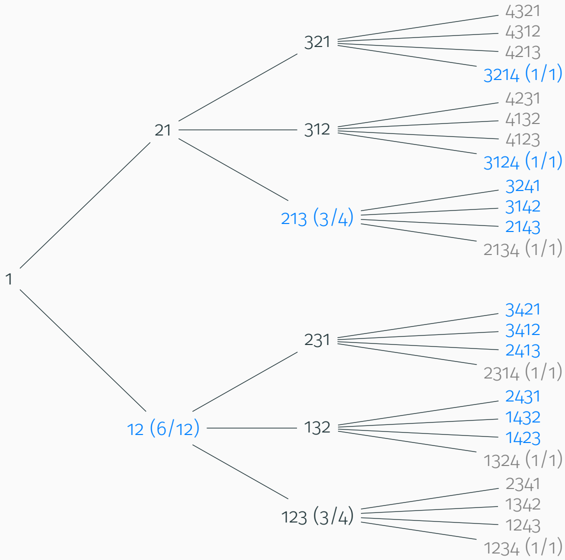

$N = 4$ prefix tree

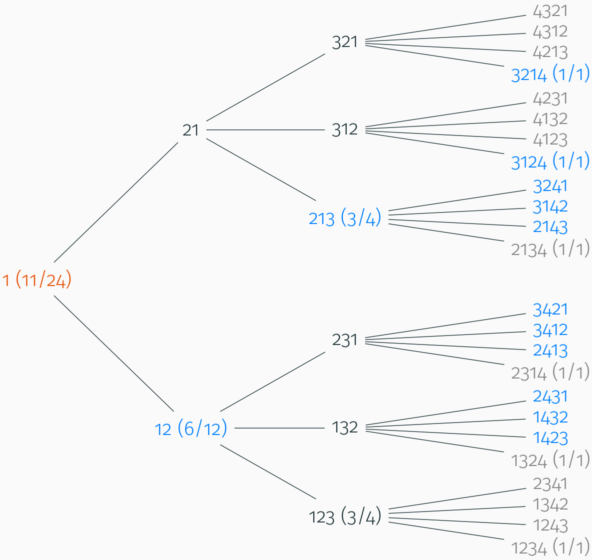

New directions

Consider **non-uniform** distributions on permutations:

- Avoid a permutation pattern (of size 3).

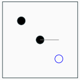

321, "avoiding disappointment"

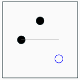

231, "raising the bar"

- Weight permutations by a statistic: $\frac{\theta^{c(\pi)}}{\sum_{\pi \in \mathfrak{S}_N} \theta^{c(\pi)}}$, $\theta \in (0, \infty)$
  - Mallows: $c(\pi) = \#$ of inversions in π
  - Ewens: $c(\pi) = \#$ of left-to-right maxima in π
  - Opportunity cost: $c(\pi) =$ position of N in π
    $= \#$ "wasted" interviews

Weighted games of best choice

Choose $\pi \in \mathfrak{S}_N$ with probability $\frac{\theta^{c(\pi)}}{\sum_{\pi \in \mathfrak{S}_N} \theta^{c(\pi)}}$ $\quad (\theta \in \mathbb{R}_+)$.

If change in $c(\pi)$ from *permuting a prefix* remains the same when we *restrict to the prefix*, say c is **sufficiently local**.

Example

$c(\pi) = \#$ inversions (pairs $\pi_i > \pi_j$ with $i < j$) are suff. local:
Consider $\pi = \pi_1\pi_2 \cdots \pi_k | \pi_{k+1}\pi_{k+2} \cdots \pi_m$.

But $c(\pi) = \#321$-instances is not: $c(2468|1357) = 0$ yet $c(4268|1357) = 1$ even though $c(2468) = 0 = c(4268)$.

Theorem (J.)

For a weighted game of best choice defined using a sufficently local statistic, the optimal strategy is always positional (reject r and accept next best).

In the stacks...

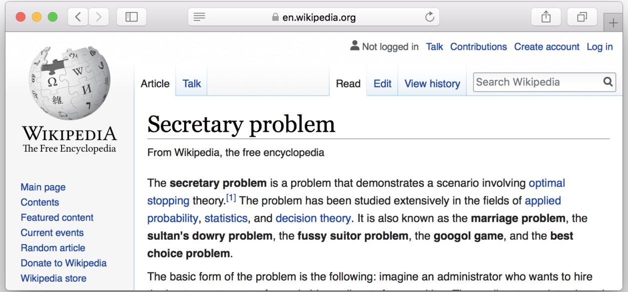

→ Merrill R. Flood, letter written in 1958, a copy of which can be found in the Martin Gardner papers at Stanford University Archives, series 1, box 5, folder 19.

(Many thanks to JMU librarian Alyssa Valcourt and special collections at Stanford!)

In the stacks...

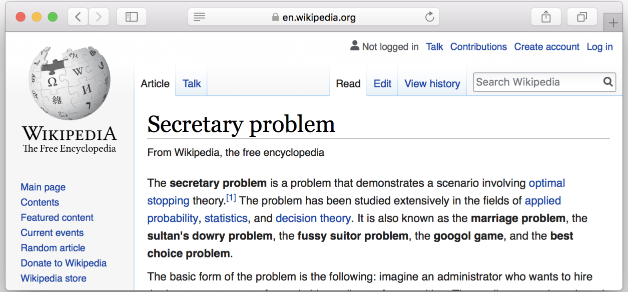

→ Merrill R. Flood, letter written in 1958, a copy of which can be found in the Martin Gardner papers at Stanford University Archives, series 1, box 5, folder 19.

(Many thanks to JMU librarian Alyssa Valcourt and special collections at Stanford!)

In the stacks...

THE UNIVERSITY OF MICHIGAN
ENGINEERING RESEARCH INSTITUTE
ANN ARBOR, MICHIGAN

TELEPHONE:
ANN ARBOR, NORMANDY 3-1511

5 May 1958

ADDRESS REPLY TO
YPSILANTI, MICHIGAN

Professor Leonard Gilman
Department of Mathematics
Purdue University
Lafayette, Indiana

Dear Professor Gilman:

Harry Goode brought back to Michigan the decision problem that you had posed to him. Harry suggested that I write to you regarding my solution of the problem. I became interested in it also because of possible applications.

Problem. $I \equiv (i_1 i_2 \text{ --- } i_n)$ is a random permutation on the first n positive integers. A game is played in which the player attempts to identify the position of the integer n.

On the first move the referee asks the player if $i_1 = n$. The payoff is 1 if the player says yes and $i_1 = n$. The payoff is 0 if the player says yes and $i_1 \neq n$, or if the player says no and $i_1 = n$. If the player says no and $i_1 \neq n$, then the referee asks next if $i_2 = n$ but also tells the player whether

In the stacks...

All that remains to establish the solution is to show that the optimal strategy is of the stated form, namely consisting of a sequence of no y times followed by a yes at the next large integer. No proof of this is given here, for I have none, but Max Woodbury assures me that a general theorem of Kuhn on behavior strategies settles this point neatly; I also suspect that it really is obvious, and I am chagrined to admit that it is not yet obvious to me. An outline proof of this by R. Palermo is attached.

I asked Max Woodbury about this problem, and its likely origins, when I saw him in Cleveland recently. He tells me that Herbert Robbins, of Columbia University, has solved it and that it has sometimes been known as the "secretary problem". I suspect that Herb told me about the problem a few years ago when I posed my "husband hunting problem" to him, namely how should a young girl decide whether to marry her fiance or to find and try a new boy to see if he is better. It may even be that Herb solved this mathematical problem as one representing my husband hunting problem, since I first posed it in a talk in 1949. At any rate, I am interested now in other possible

In the stacks...

Addendum by R. Palermo

This is an outline of a proof that the optimal strategy is among the class of strategies considered in the letter.

Let y ≡ "y randomly chosen integers" (selected from {1, 2, ..., n})

Let x ≡ "the integer under consideration"

If x > max y, then there exists $P_{x,y}$ where

$$P_{x,y} = \text{probability that } x \text{ is } n$$

$$= P_y\ (x = n)$$

Note: (1) $P_{x,y}$ is the probability of winning if the integer x is played

(2) $P_{x,y}$ is a monotone increasing function of y

(3) If we let $\overline{P}_{x,y}$ = the probability of winning if x is not played

then $\overline{P}_{x,y}$ is a monotone decreasing function of y.

Opportunity cost model

Set $\theta < 1$ and weight each π by $\theta^{\pi^{-1}(N)-1}$ to obtain uniform game with varying payoffs:

hire best immediately, $\theta^0 = 1$
hire best after one interview, $\theta^1 = 0.95$
hire best after two interviews, $\theta^2 = 0.9025$, etc.

Define

$$W_{N,r}(\theta) = \sum_{r\text{-winnable } \pi \in \mathfrak{S}_N} \theta^{\pi^{-1}(N)-1}$$

Find a **recurrence**, **solve** it, divide by the full **generating function**, and then **take the limit**, to get asymptotic **success probability** for the strategy that initially rejects r candidates:

$$P_r(\theta) := \lim_{N\to\infty} W_{N,r}(\theta) = r(1-\theta)\sum_{i=r}^{\infty} \frac{\theta^i}{i}.$$

Opportunity cost model

Set $\theta < 1$ and weight each π by $\theta^{\pi^{-1}(N)-1}$ to obtain uniform game with varying payoffs:

hire best immediately, $\theta^0 = 1$
hire best after one interview, $\theta^1 = 0.95$
hire best after two interviews, $\theta^2 = 0.9025$, etc.

Define

$$W_{N,r}(\theta) = \sum_{r\text{-winnable } \pi \in \mathfrak{S}_N} \theta^{\pi^{-1}(N)-1}$$

Find a **recurrence**, **solve** it, divide by the full **generating function**, and then **take the limit**, to get asymptotic **success probability** for the strategy that initially rejects r candidates:

$$P_r(\theta) := \lim_{N\to\infty} W_{N,r}(\theta) = r(1-\theta)\sum_{i=r}^{\infty} \frac{\theta^i}{i}.$$

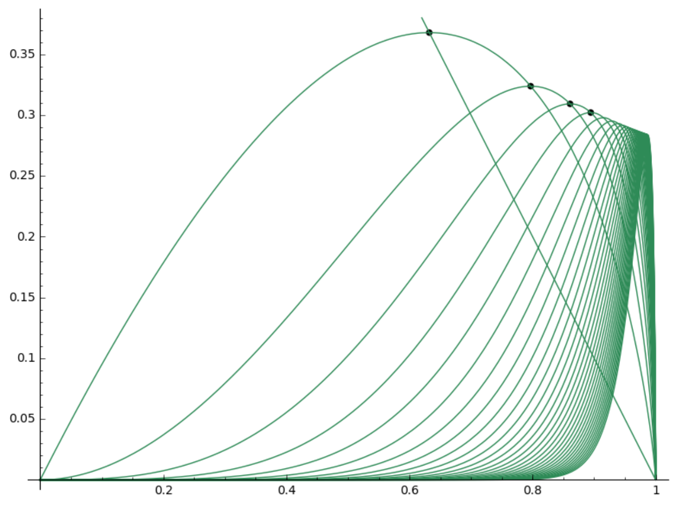

$r =$	is optimal for $\theta \in$	success probability
0	(0, 0.6321]	$1/e \approx 0.36788$
1	[0.6321, 0.7968]	0.323805
2	[0.7968, 0.8609]	0.309256
3	[0.8609, 0.8945]	0.302113

Lemma. For each r, the intersection of P_{r-1} and P_r coincides with the maximum value of P_r.

Opportunity cost model

Let

$$E_1(x) = \int_x^\infty \frac{e^{-t}}{t}\, dt$$

Consider $F(x) = xE_1(x)$ on $(0, \infty)$.

Define α and β be defined by $F'(\alpha) = 0$ and $F(\alpha) = \beta$. Then, $\alpha \approx 0.4348182150039929 3$ and $\beta \approx 0.2814936299569167 4$.

Theorem (Crews–J.–Myers–Taalman–Urbanski–Wilson)
As $\theta \to 1^-$, the optimal strategy is $\left(\frac{\alpha}{1-\theta}\right)$-positional. This strategy has a success probability of β.

To optimize $P_r(\theta) = r(1-\theta)\sum_{i=r}^{\infty}\frac{\theta^i}{i}$ we estimate the series,

$$\int_{t=r}^{\infty}\frac{\theta^t}{t}\,dt < \sum_{i=r}^{\infty}\frac{\theta^i}{i} < \int_{t=r}^{\infty}\frac{\theta^{t-1}}{t-1}\,dt = \int_{t=r-1}^{\infty}\frac{\theta^t}{t}\,dt.$$

So $\widetilde{P}_r(\theta) = r(1-\theta)\int_{t=r}^{\infty}\frac{\theta^t}{t}\,dt$ has error less than $r(1-\theta)\frac{\theta^{r-1}}{r-1} < 4(1-\theta)\theta^r$.

Next, we change variables from r to $c = (1-\theta)r$, and from t to $u = (1-\theta)t$ in the integral. We obtain $du = (1-\theta)\,dt$ so

$$\widetilde{P}_c(\theta) = c\int_{u=c}^{\infty}\frac{\left(\theta^{1/(1-\theta)}\right)^u}{u}\,du.$$

and our error estimate for $|P-\widetilde{P}|$ becomes $4\theta^{c/(1-\theta)}(1-\theta)$.

Take $\theta \to 1$, using $\lim_{\theta\to 1}\theta^{1/(1-\theta)} = 1/e$.

Origins

of G/Q is generated in degree 2, i.e., $I(G/Q)$ is generated as an ideal by

There are similar results for multi-cones over Schubert varieties [81],

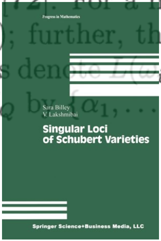

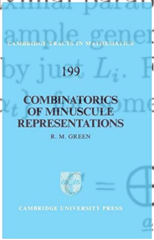

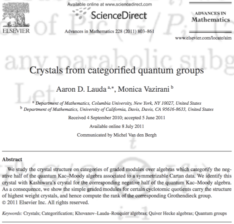

Available online at www.sciencedirect.com
ScienceDirect
Advances in Mathematics 228 (2011) 803–861
ADVANCES IN Mathematics
www.elsevier.com/locate/aim

Crystals from categorified quantum groups

Aaron D. Lauda [a,*], Monica Vazirani [b]

[a] Department of Mathematics, Columbia University, New York, NY 10027, United States
[b] Department of Mathematics, University of California, Davis, Davis, CA 95616-8633, United States

Received 4 September 2010; accepted 5 June 2011
Available online 8 July 2011
Communicated by Michel Van den Bergh

Abstract

We study the crystal structure on categories of graded modules over algebras which categorify the negative half of the quantum Kac–Moody algebra associated to a symmetrizable Cartan data. We identify this crystal with Kashiwara's crystal for the corresponding negative half of the quantum Kac–Moody algebra. As a consequence, we show the simple graded modules for certain cyclotomic quotients carry the structure of highest weight crystals, and hence compute the rank of the corresponding Grothendieck group.

Keywords: Crystals; Categorification; Khovanov–Lauda–Rouquier algebras; Quiver Hecke algebras; Quantum groups

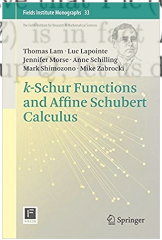

$$C_Q(w) = \bigoplus H^0(X_Q(w), \bigotimes^t L_i^{a_i}),$$

The problem

Successful **Ph.D. thesis research** is conducted:	Successful **undergraduate research** is conducted:
over years	over months
with complete dedication, fast-paced	in the context of other classes/undergrad life
with support of grad program	limited institutional support
for original results, juried by top experts	for helping student evolve learning/thinking, to keep UG faculty engaged

Possible responses:

- give up research entirely, or change fields completely
- hyperfocus on very specific aspect of thesis research
- hypergeneralize to very abstract aspect of thesis research

An alternative?

Successful **Ph.D. thesis research** is conducted:	Successful **undergraduate research** is conducted:
over years	over months
with complete dedication, fast-paced	in the context of other classes/undergrad life
with support of grad program	limited institutional support
for original results, juried by top experts	for helping student evolve learning/thinking, to keep UG faculty engaged

Pivot to research that is more *accessible* but *leverages* existing skills (e.g. programming and algorithms, solving recurrences, enumeration, generating functions) or skills you want to learn anyway (e.g. probability applied to combinatorial structures).

An alternative?

Successful **Ph.D. thesis research** is conducted:	Successful **undergraduate research** is conducted:
over years	over months
with complete dedication, fast-paced	in the context of other classes/undergrad life
with support of grad program	limited institutional support
for original results, juried by top experts	for helping student evolve learning/thinking, to keep UG faculty engaged

Pivot to research that is more *accessible* but *leverages* existing skills (e.g. programming and algorithms, solving recurrences, enumeration, generating functions) or skills you want to learn anyway (e.g. probability applied to combinatorial structures).

Thanks for attending!