given any pos, p want

$E[p]$ = expected payout given that game has reached p

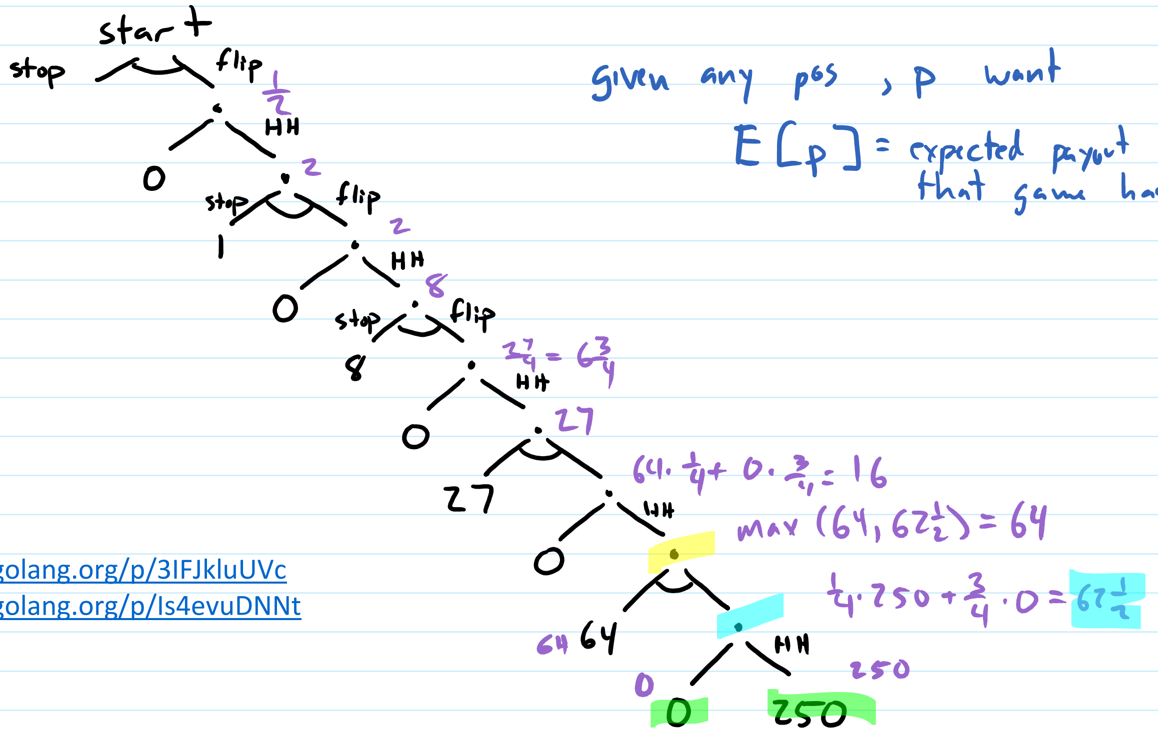

https://play.golang.org/p/3IFJkluUVc
https://play.golang.org/p/Is4evuDNNt

Analysis of 1-player Finite Probabilistic Games

$E[pos]$ = expected winnings having reached position pos

For final positions pos, $E[pos]$ determined by rules

For non-final choice positions

$$E[pos] = \max_{\text{choice } c} E[\underbrace{next(pos, c)}_{\text{position reached by making choice } c \text{ from position } pos}]$$

For non-final random event positions

$$E[pos] = \sum_{\text{outcome } \sigma} P(\sigma) \cdot E[next(pos, \sigma)]$$

```
for every terminal position pos
    E[pos] ← payoff(pos)

for every nonterminal position pos   in reverse order of topological sort
    if pos is a choice position
        max ← -∞
        argmax ← NIL
        for every choice c
            e ← E[next(pos,c)]
            if e > max
                max ← e
                argmax ← c
        E[pos] ← max
        OPT[pos] ← argmax
    else
        e ← 0.0
        for every outcome σ
            e ← e + P(σ) · E[next(pos,σ)]
        E[pos] ← e
```

$E[start]$ = value of game

Coino: Start with n coins.
On each turn, flip as many of your remaining coins as you wish.
If #T ≥ #H, lose all the T
Else earn #H points
Win at X points

If # T ≥ # H, lose all the T
Else earn # H points
Win at X points
Lose if no coins left and < X points

anchor: position at start of turn (before 1st roll)

component: positions reachable from one anchor w/o going through another

number of anchors: $\underbrace{7^6}_{\text{7 ways for each of 1s,...,6s}} \cdot \underbrace{28^2}_{\text{unused or 0,5...30 for 3k,4k}} \cdot \underbrace{27}_{\text{C: same as but can't be 0}} \cdot \underbrace{3^3}_{\text{unused, 0, nonzero: FH, SS, LS}} \cdot \underbrace{15}_{\text{Y = unused, 0, 50, ..., 1250}} \gtrsim 1 \text{ trillion anchors}$

$\cdot$ 1600 pos/anchor

$\parallel$

$\sim$ 1.6 quadrillion positions

modification: $E[pos]$ = expected future score from pos

For nonterminal choice positions

$$E[pos] = \max_{\text{choice } c} E[next(pos, c)] + score(pos, c)$$

For nonterminal random event positions

$$E[pos] = \sum_{\text{outcome } \sigma} P(\sigma) \cdot \left(E[next(pos, \sigma)] + score(pos, \sigma) \right)$$

anchors = $\underbrace{2^{12}}_{\text{used/unused}} \cdot \underbrace{3}_{\text{Y: unused, 0, nonzero}} \cdot \underbrace{64}_{\text{upper total}} \gtrsim 750$ million anchors

$\cdot$ 1600

$\gtrsim$ 2 billion

"position" = minimal info that affects strategy

so < 1 hour if 1 million pos/sec

Yahtzee Graph