PARALLEL NASH EQUILIBRIA IN BIMATRIX GAMES

ISAAC ELBAZ
CSE633 FALL 2012
INSTRUCTOR: DR. RUSS MILLER

WHAT *IS* GAME THEORY?

- **Branch of mathematics that deals with the *analysis of situations involving parties with conflicting interests*.**
- **There are mainly two branches of Game Theory: *Cooperative* and *Non Cooperative*.**
- **Non cooperative game theory deals with how rational individuals interact with e/o in an effort to achieve *their own goals* (in other words, with no regard for social welfare).**
- ***The single most important* idea of non cooperative games is the *solution concept* (i.e., a prediction of how the game will be played).**

PRISONER’S DILEMMA

An NYPD officer arrested two suspects, *A* and *B*. The problem is, the officer does not have enough evidence to convict either suspect for the crimes committed.

Instead, the officer locks both suspects in separate rooms, and offer the following identical deal to each:

“If you confess your crime, and your partner doesn’t, you go free, and your partner stays in jail for ten years. If you don’t confess, and your partner does, you go to jail for ten years, and your partner walks free. If you both confess, you both receive a reduced sentence”.

PRISONER'S DILEMMA

We can model the previous game by using the following payoff matrix (also called a *normal form* representation).

	Confess A	Stay quiet A
Confess B	6 6	10 0
Stay quiet B	0 10	2 2

http://www.answers.com/topic/prisoner-s-dilemma

THE (KNOWN) OUTCOME.

- **Players are *always better off* choosing to confess to improve their own payoff.**
- **The *only* stable solution to this game is when both players choose to confess.**
- ***This is not by chance!* Many games (including this one) are designed so that the outcome could be predicted.**
    - See Mechanism Design.

SAME GAME IN DISGUISE …

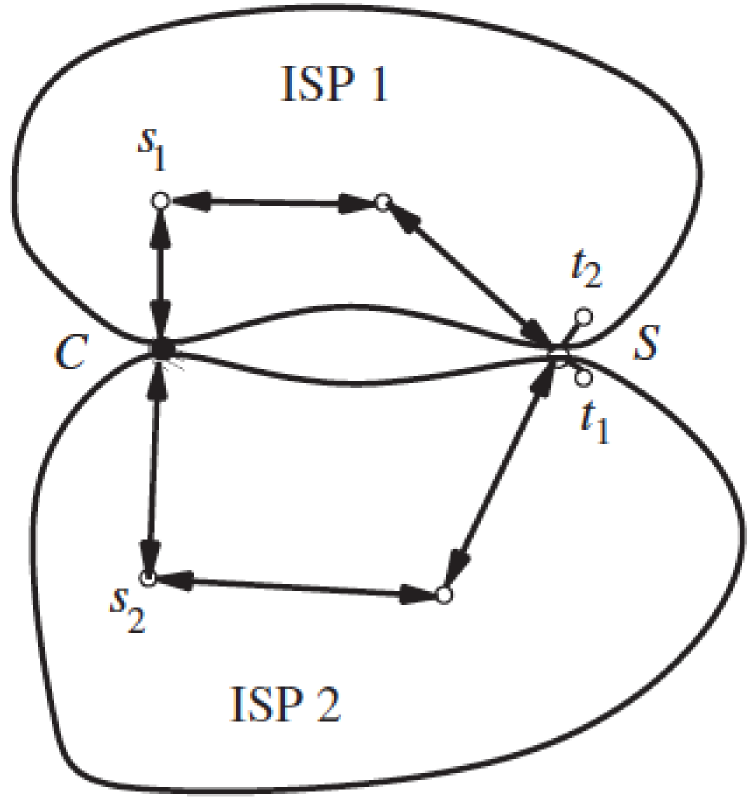

Figure 1.1. The ISP routing problem.

Algorithmic Game Theory, Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay V. Vazirani, editors, Cambridge University Press, Cambridge, 2007.

SOME DEFINITIONS …

A game consists of a set $P = \{1, 2, \ldots, n\}$ of n players. $\forall i \in P, S_i$ is a set of possible strategies for player i.

We define $s = (s_1, s_2, \ldots, s_n)$ to denote the vector of strategies, and S to be set of all possible vector strategies.

Each player i must be able to rank its strategies. In other words, $\forall i \in P$, given possible strategies $s_j, s_k \in S_i$, i must be able to estabilish a preference over s_j or s_k. Mathematically, for each player i, we need to give a complete, transitive, reflexive binary relation on the set S_i.

In this presentation, we will assign for each player, a value to each outcome. $u_i: S \to \mathbb{R}$. It is important to notice that u_i maps from $S \to \mathbb{R}$ and not from $S_i \to \mathbb{R}$. The latter case would simply be an n independent optimiation problems. Instead, in a game, the payoff depends not only on a player's strategy but also on the strategies chosen by all other players.

For a strategy vector $s \in S$, we use s_i to denote the ith player strategy, and s_{-i} to denote the $(n-1)$ dimensional vector of the strategies played by the others.

NASH EQUILIBRIUM DEFINED.

$s \in S = (s_1, s_2, \ldots, s_n), \forall i \in P, \forall s_i' \in S_i, u_i(s_i, s_{-i}) \geq u_i(s_i', s_{-i}) \implies s$ is a nash equilibrium.

- **A solution vector s is a Nash Equilibrium if no player can *unilaterally* change its strategy in order to improve his payoff.**
- **Nash Equilibrium is not always (socially) optimal.**
  - **Price of Anarchy.**
- ***If both stayed quiet, the payoff would be substantially better***
- ***no matter how much effort is put in coordinating such play, both side would be tempted to deviate*, and would end up confessing.**

POLLUTION GAME, AKA, N-PLAYER PRISONER'S DILEMMA.

Assume there are *N* countries in the world. Each country is faced with the decision of passing legislation for pollution control.

If a country decides to pass such legislation, ***there is a cost of 5 associated with it***; but ***each country that pollutes adds 1 to the cost of all countries***. Notice that ***polluting is much cheaper than controlling pollution*** (when thinking selfishly).

If k countries choose to ignore pollution control, the cost of each of these countries is k on the other hand, each of the other n-k countries have a cost of k+5.

MULTIPLE NASH EQUILIBRIA?

Battle of Sexes. This is a classic "coordination game" (that is, players choose between two options, wanting to choose the same).

Can you identify the Nash Equilibrium? More than one?

		Woman	
		Baseball	Ballet
Man	Baseball	(3, 2)	(1, 1)
	Ballet	(0, 0)	(2, 3)

http://catalog.flatworldknowledge.com/bookhub/13?e=mcafee-ch16_s02

DO WE ALWAYS HAVE A NASH EQUILIBRIUM?

The answer to this question is yes and no. So far, all the Nash Equilibria we've seen were *pure strategy*; that is, each player *deterministically* plays his chosen strategy.

***If we limit ourselves to pure strategy Nash Equilibria, then it is not the case that every game has an equilibrium*.**

Column

Row	Heads	Tails
Heads	(1, –1)	(–1, 1)
Tails	(–1, 1)	(1, –1)

http://www.web-books.com/eLibrary/NC/B0/B59/096MB59.html

MIXED STRATEGIES TO THE RESCUE.

- **Analyzing the previous game quickly reveals that playing a deterministic strategy is *not a good idea* for any player.**
  - **Randomly *pick a strategy. That way we can perhaps 'fool' the other player*.**
  - **This leads to the notion of Mixed Strategy.**
    - **Allow each player to pick a probability distribution over his set of possible strategies.**

MIXED STRATEGIES DEFINED

Definition 3. Mixed strategies of players 1 and 2 are the vectors of probabilities $\boldsymbol{p}, \boldsymbol{q}$ for which the following conditions hold:

$$\boldsymbol{p} = (p_1, p_2, \dots p_m); \quad p_i \geq 0, \quad p_1 + p_2 + \cdots + p_m = 1,$$

$$\boldsymbol{q} = (q_1, q_2, \dots q_n); \quad q_i \geq 0, \quad q_1 + q_2 + \cdots + q_n = 1.$$

Here p_j (q_j) expresses the probability of choosing the j-th strategy from the strategy space S (T).

Definition 4. Expected payoffs are defined by the relations:

$$\begin{aligned} \text{Player 1:} \qquad & \pi_1(\boldsymbol{p}, \boldsymbol{q}) = \sum_{i=1}^{m} \sum_{j=1}^{n} p_i q_j a_{ij} \\ \text{Player 2:} \qquad & \pi_2(\boldsymbol{p}, \boldsymbol{q}) = \sum_{i=1}^{m} \sum_{j=1}^{n} p_i q_j b_{ij} \end{aligned} \tag{4.2}$$

http://euler.fd.cvut.cz/predmety/game_theory/games_bim.pdf

A NOBEL THEOREM IN GAME THEORY.

"*Any game with a finite set of players and finite set of strategies has a Nash Equilibrium of mixed strategies*"

This theorem was proved by John F. Nash in 1949.

Since the graph is closed and since the image of each point under the mapping is convex, we infer from Kakutani's theorem[1] that the mapping has a fixed point (i.e., point contained in its image). Hence there is an equilibrium point.

http://web.mit.edu/linguistics/events/iap07/Nash-Eqm.pdf

MIXED STRATEGY EQUILIBRIUM.

Definition 6 *A mixed strategy equilibrium is a mixed strategy profile* $(\alpha_1^*, \ldots, \alpha_n^*)$ *such that, for all* $i = 1, \ldots, n$

$$\alpha_i^* \in \arg\max_{\alpha_i \in \Delta(A_i)} u_i\left(\alpha_i, \alpha_{-i}^*\right)$$

or

$$\alpha_i^* \in B_i\left(\alpha_{-i}^*\right).$$

http://portal.ku.edu.tr/~lkockesen/teaching/uggame/lectnotes/uglect4.pdf

- **This reveals a very interesting property which can guide the process of finding a mixed strategy Nash Equilibrium.**
- **A mixed strategy profile is an equilibrium iff for each player i, *each action on the support of its mixed strategy is a best response* to every other mixed strategy in the strategy profile.**

COMPUTING THE NASH EQUILIBRIA.

- **There are a total of Binomial[m+n,n] – 1 possible pairs of supports (where n <= m)**
- **Each will produce (n + m) + 2 equations.**
- **The systems of equations can be solved in O((n+m)^3) using Gaussian Elimination.**
    - INTEL MKL / CLAPACK / BLAS
- **By using Sterling's Approximation the total runni**

  can be simplified to ***O(4^n * n^3)***
- ***Not pretty.***

SEQUENTIAL ALGORITHM.

```
for i in (1 ... n)
        supports = GenerateAllSupportsOfSize(n);
        for each (x, y) in supports
                x' = MixedStrategy(x,v);
                y' = MixedStrategy(y,u);

                if IsNashEquilibrium(x',y')
                        output (x',y');
                end if
        end for each
end for
```

PARALLEL ALGORITHM.

```
comm_size = GetMPICommSize(COMM_WORLD);
rank = GetMPIRank(COMM_WORLD);
max = min{ actions(playerA), actions(playerB) };
for i in [1 ... max]
        supports = GenerateSupportsOfSize(n,rank)
        for each (x, y) in supports
                        x' = MixedStrategy(x,v);
                        y' = MixedStrategy(y,u);

                        if IsNashEquilibrium(x',y')
                                output (x',y');
                        end if
        end for each
end for
```

PARALLEL ALGORITHM (II)

The function MixedStrategy(a) generates a mixed strategy for a player given a support a. This can be achieved by solving the following systems of equations

$$\sum_{i \in M_x} x_i b_{ij} = v, \quad \text{for } j \in N_y$$

$$\sum_{i \in M_x} x_i = 1$$

and

$$\sum_{j \in N_y} y_j a_{ij} = u \quad \text{for } i \in M_x$$

$$\sum_{j \in N_y} y_j = 1$$

http://www.cs.wayne.edu/~dgrosu/pub/ispdc08.pdf

WORKED OUT EXAMPLE.

- **Back to the game of matching pennies.**
- **Already established no point in considering pure strategy.**
- **Consider the support (x1,x2) and (y1,y2)**
- **We generate the following equations:**
- **-x1 + x2 = v; x1 – x2 = v; x1 + x2 = 1;**
  - Leads to (x1,x2) = (1/2,1/2);
- **y1 – y2 = u; y1 – y2 = u; y1 + y2 = 1;**
  - Leads to (y1,y2) = (1/2,1/2)
- **Now, we must decide whether (((1/2),(1/2)),((1/2),(1/2))) is a mixed strategy equilibrium.**
- **We do this by calculating the expected payoff for playing each pure strategy in the support. Indeed, the expected payoff for playing this mixed strategy, is 0. As expected, we don't win, or lose.**

http://gamut.stanford.edu/

TEST PLAN.

1. **Test games were generated by GAMUT, using the following commands:**
   1. java –jar gamut.jar -int_payoffs -output TwoPlayerOutput -players 2 –actions N -g MinimumEffortGame -random_params
   2. Chose the following number of actions: 10, 12, 14, 16
2. **Progressively tested the games on the following number of cores:**
   1. 1, 2, 4, 8, 16, 32, 64

PUTTING IT TO THE TEST I

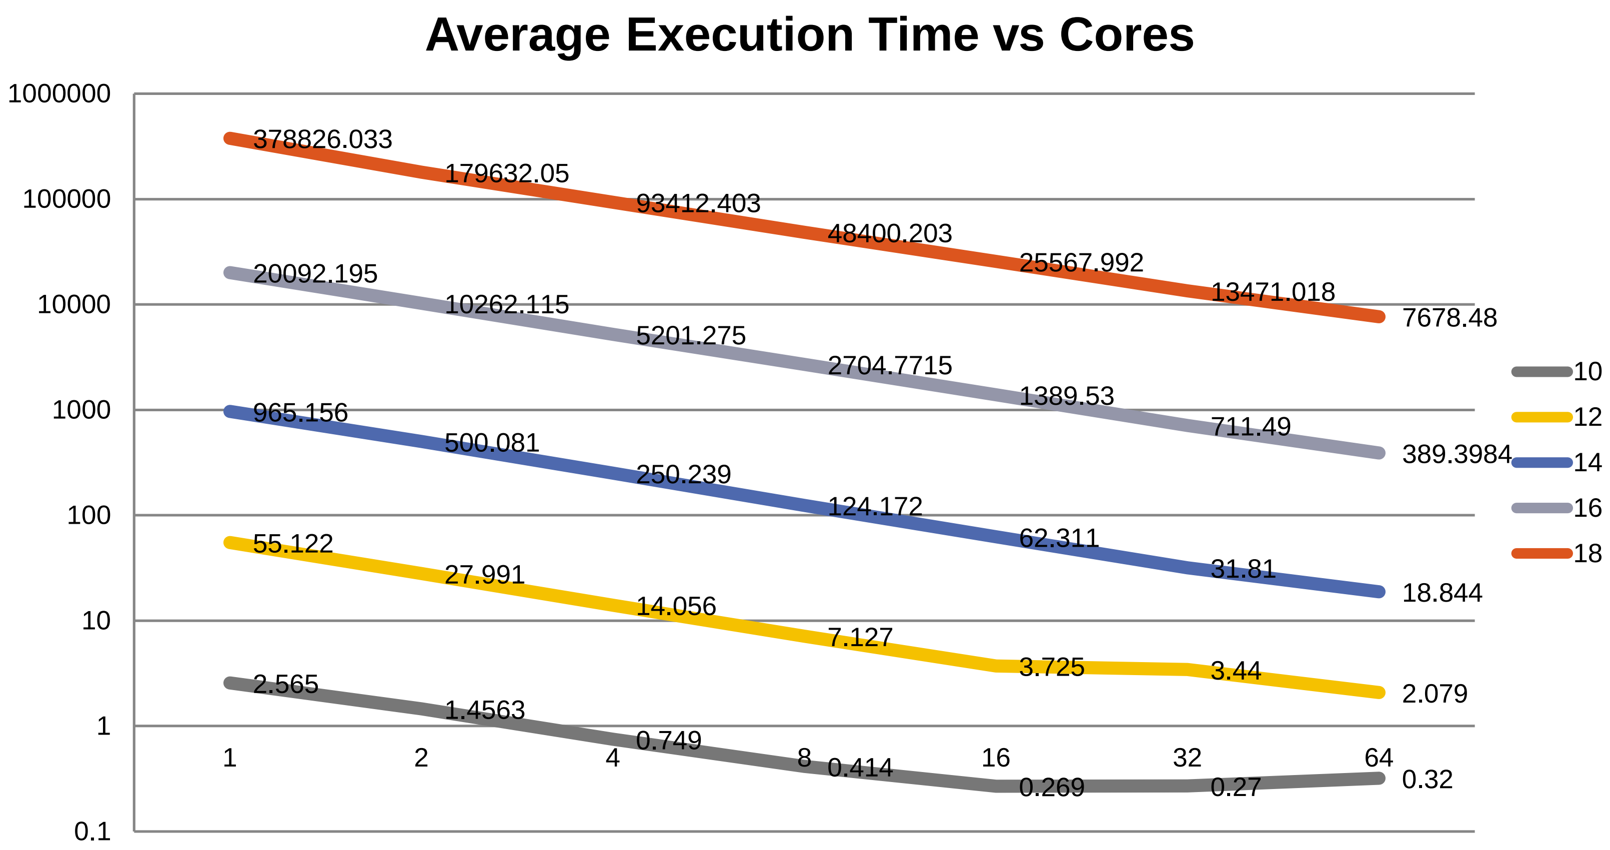

PUTTING IT TO THE TEST II

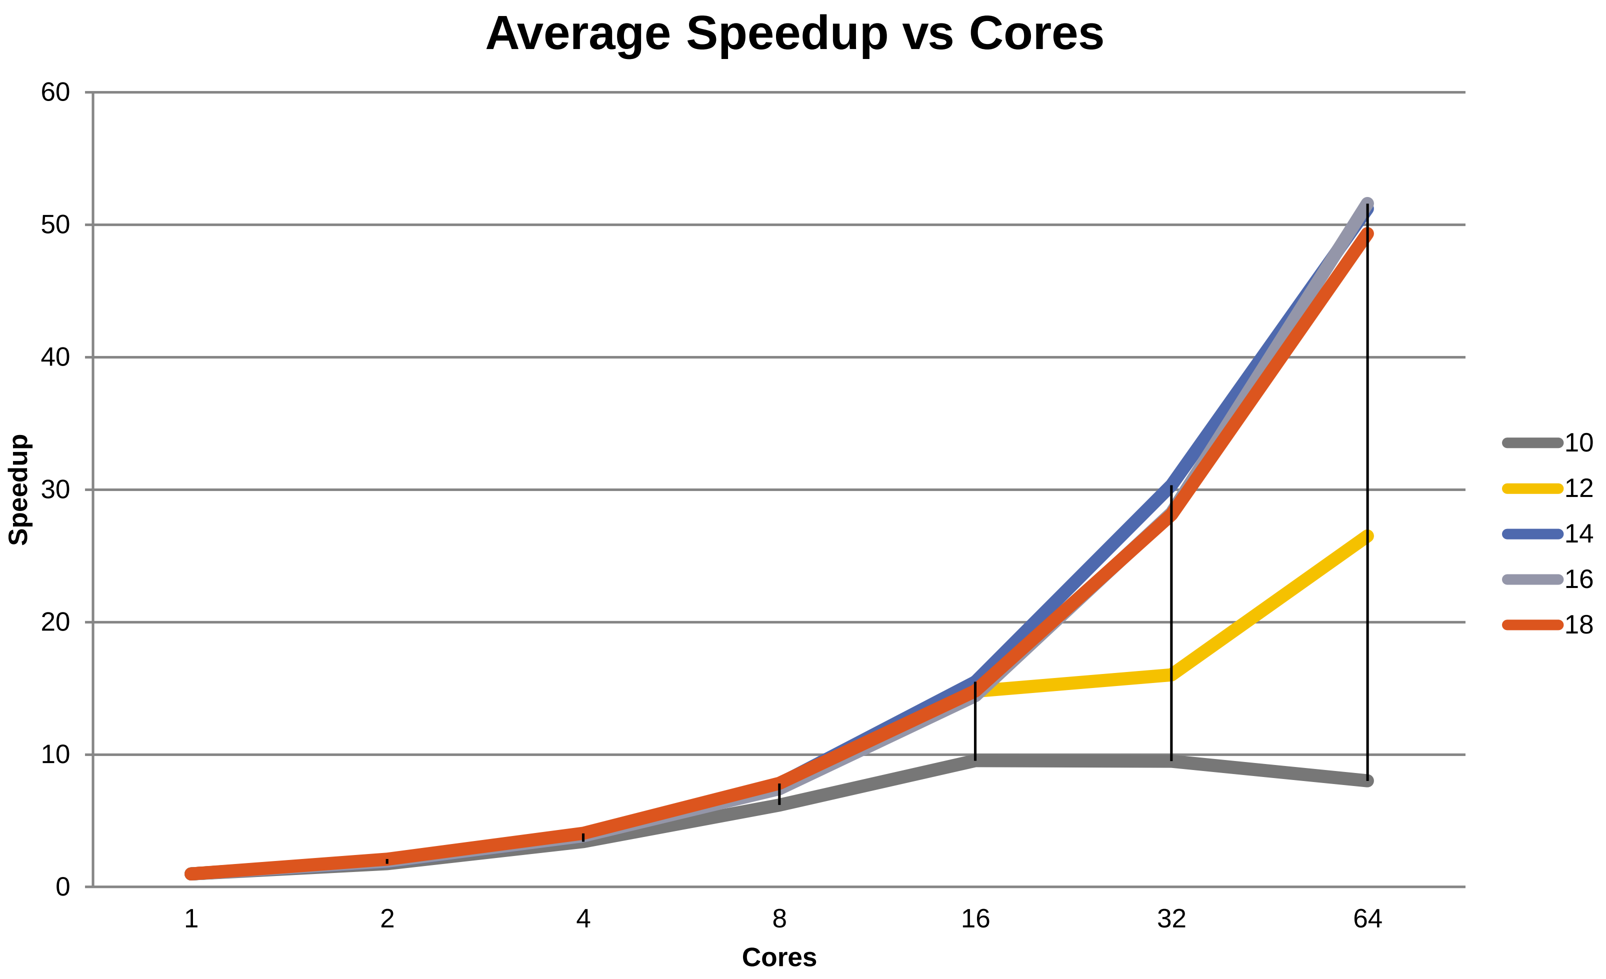

PUTTING IT TO THE TEST III

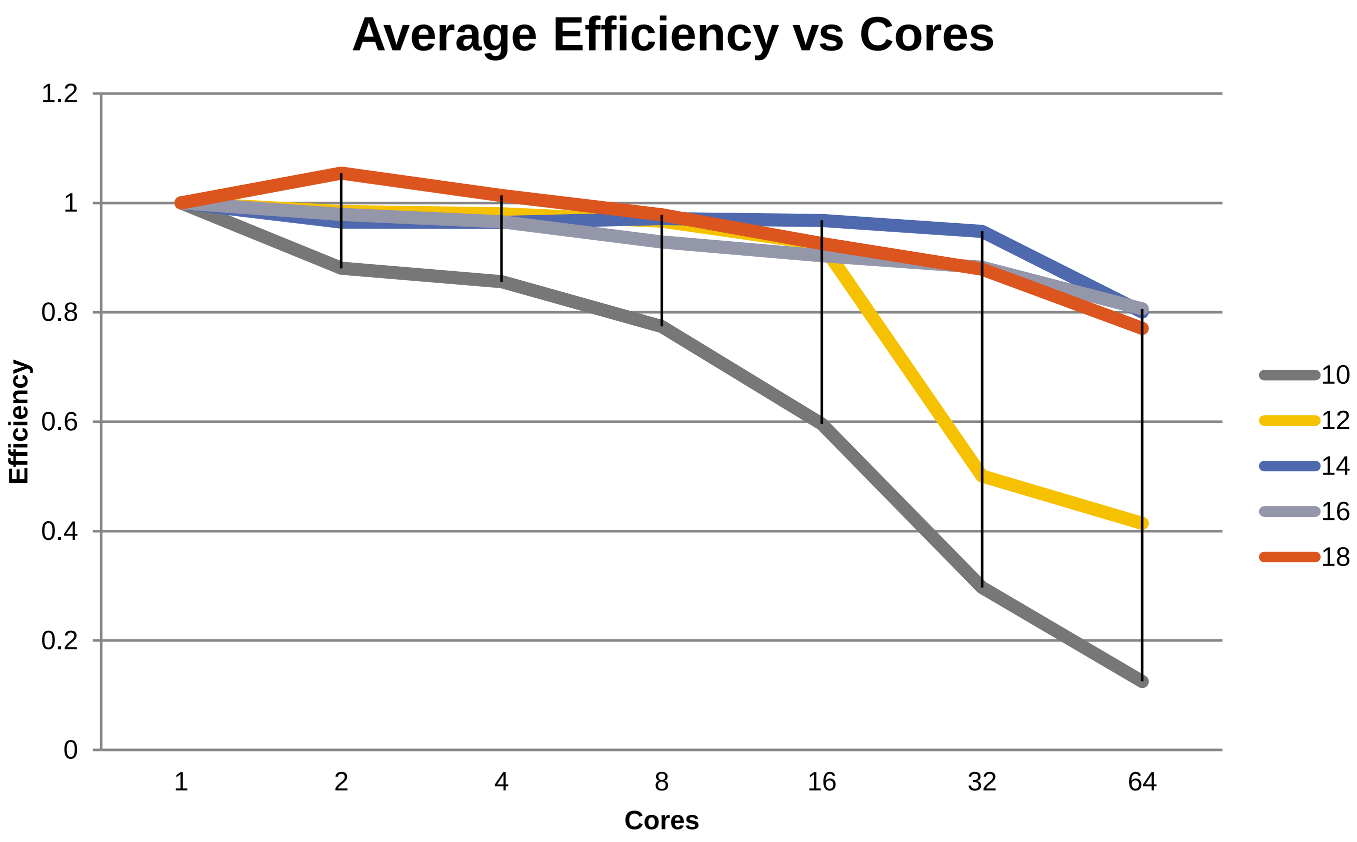

TRY IT YOURSELF.

```
git clone git@github.com:script3r/nash.git &&
  make && mpirun -np N ./nash [sample-game]
```

THE THEORY THREAD.

Von Neumann

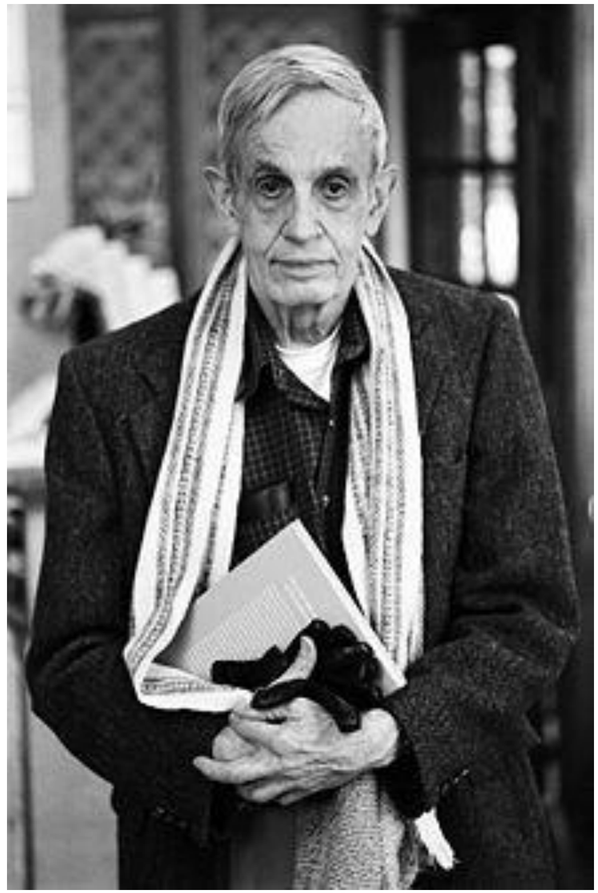

Nash

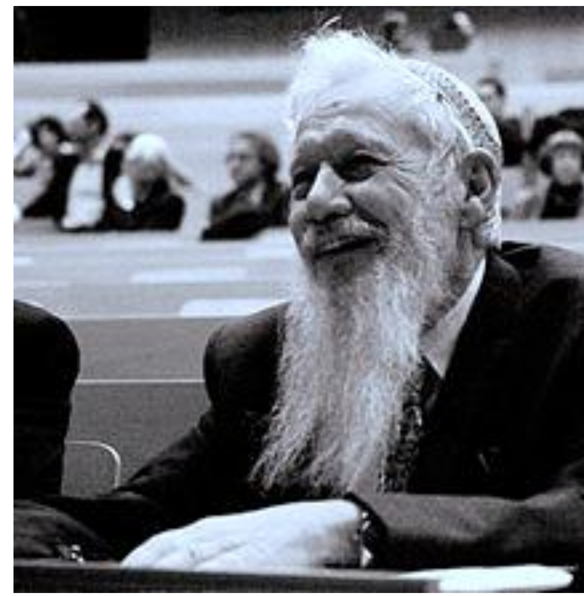

Robert Aumann

THE COMPUTATIONAL THREAD.

Tim Roughgarden

Eva Tardos

Noam Nisan

Vazirani

INTERESTED IN THE TOPIC?

http://www.cambridge.org/journals/nisan/downloads/Nisan_Non-printable.pdf

Free book on the subject.

REFERENCES

Nash, John (1950) "Equilibrium points in n-person games" Proceedings of the National Academy of Sciences 36(1):48-49.

"16.3. Mixed Strategies." 16.3. Mixed Strategies. N.p., n.d. Web. 11 Oct. 2012. <http://www.web-books.com/eLibrary/NC/B0/B59/096MB59.html>.

Nisan, Noam. Algorithmic game theory. Cambridge: Cambridge University Press, 2007. Print.

Jonathan Widger and Daniel Grosu. (2008). Computing Equilibria in Bimatrix Games by Parallel Support Enumeration.