

Mechanism Design for Scheduling with Uncertain Execution Time.

Angelina Vidali

Joint work with

Vincent Conitzer (Duke University)

and

ongoing work with Carmine Ventre

Crowdsourcing




- Accomplish a complicated very big job using many agents!
- Give incentives to the agents to complete the tasks!


Mechanism Design

A single choice for the whole society

Voting
Public Projects
Auctions
Scheduling



We need to construct a function that takes as input the **preferences of many different individuals** and „**amalgamates**” (/aggregates) them into a **single** preference or choice.




Mechanism Design

Design a game whose **outcome** is an **equilibrium** for the players.

truthful

We want mechanisms,
that don't give players an incentive to lie.

When a player has an incentive to lie...



$$v(\text{EELV}) > v(\text{PS}) > v(\text{UMP})$$

Can the captain achieve a more preferable outcome by lying
about his political preferences?

Hazard Rate:

Probability a painter finishes the painting at time t given that he hasn't finished it until time $t-1$

$$\varphi_i(t) = \frac{P(T_i = t)}{(1 - P[T_i < t])}$$

Monotone (non-increasing) hazard rate assumption:

the more time a painter takes

the less likely is he to finish at the next time step

The Queen wants a painting for her palace.
Every day she decides which painters will draw.

Goal: minimize $E[\text{painting time}]$




Creativity is unpredictable!

Painters don't know how much time it's going to take them and need incentives to draw!



Time painter i needs to finish the job $\sim$ distribution f_i
painter i knows f_i but not t_i and f_{-i}



Players are selfish
want to maximize:
utility = $E[-\text{time spent painting} + \text{payment}]$

Crowdsourcing Application

We want to solve a problem (e.g. a SAT instance), by running a crowdsourcing contest.

Payment 0 if you fail: Only the winner (agent who solves the problem first) will get paid.

Uncertainty: The agents don't know how much time they will need.

For SAT there are some good heuristics so the probability of finding a solution in the beginning is rather high, but if these don't work then it might take forever... (MHR assumption makes sense)

The greedy algorithm is optimal!



-“Off with their heads!”

Objective: Minimize the
 $E[\text{sum of processing times}]$

Greedy=OPT:
assign at each time step
the job to the machine
with maximum hazard rate
i.e. the machine more likely to finish!

To prove this we need:
Monotone hazard rates assumption

“Sort the hazard rates!”

Objective: Minimize the $E[\text{sum of processing times}]$

OPT: assign at each time step the job to the machine with maximum hazard rate

Input

	$\varphi(1)$	$\varphi(2)$	$\varphi(3)$	$\varphi(4)$	$\varphi(5)$	...
	0.9	0.8	0.5	0.1	0.1	
	0.7	0.4	0.3	0.3	0.3	
	0.6	0.2	0.1	0.1	0.1	

OPT

							
0.9	0.8	0.7	0.6	0.5	0.4	0.3	0.3

Consistency Property

If we remove one player e.g.  to get OPT for the rest of the players we just need to remove the player from the schedule

Input

	$\varphi(1)$	$\varphi(2)$	$\varphi(3)$	$\varphi(4)$	$\varphi(5)$
	0.9	0.8	0.5	0.1	0.1
	0.7	0.4	0.3	0.3	0.3
	0.6	0.2	0.1	0.1	0.1

OPT

							
0.9	0.8	0.7	0.6	0.5	0.4	0.3	0.3

Input

	$\varphi(1)$	$\varphi(2)$	$\varphi(3)$	$\varphi(4)$	$\varphi(5)$
	0.7	0.4	0.3	0.3	0.3
	0.6	0.2	0.1	0.1	0.1

OPT

				
0.7	0.6	0.4	0.3	0.3

We focus on **direct revelation mechanisms**

Input: true types of the players
here: the distributions f_i


Output: allocation(which machine processes
at each time step) and
payments

We have a **revelation principle**

The Vickrey-Clark-Groves(VCG) mechanism
maximizes the welfare (sum of the valuations of the players)
and is truthful

The highest bid v_i wins and

pays 2nd highest bid



Auctions

“The payments align the incentives of the players with the objective of the mechanism”

Thus no incentive to lie=truthfull

**OBJECTIVE of
the mechanism:
welfare**

$$\max \left(\sum \text{valuations of the players} \right)$$

$$\sum \text{valuations of the players} = \text{valuation of } i + \sum_{j \neq i} \text{valuations of the other players}$$

$$\text{Utility of selfish player } i = \max \left(\text{valuation of } i - \boxed{\begin{array}{c} ? \\ \text{payment} \end{array}} \right)$$

$$\text{Groves payment of player } i = - \left(\sum_{j \neq i} \text{valuations of the other players} \right)$$

But the Groves payment is negative.

How to make it positive:

$$p_i = \text{Groves} + h_i(\text{types of the other players})$$

e.g. h_i = welfare of best allocation if player i was absent


Player i cannot affect the **$h()$ part**
since it does not depend on the value he reports

Thus still no incentive to lie!

also true but hard proof:

any other truthful payment is of this form for
some function $h_i(\text{types of the other players})$

The highest bid v_i wins and



Auctions


gets payed the 2nd lowest bid

The fastest machine (lowest bid)
 t_i wins and

The Vickrey mechanism

min sum of processing times

truthful



Scheduling

Solution Concept: Ex-post equilibrium

If the other players are telling the truth,
then the best thing for me to do is to tell the truth,
for **any** private information the players might.

Dominant $\subseteq$ Ex-post $\subseteq$ Bayes Nash

$$\text{Groves payment of player } i = -\left(E\left[\sum_{j \neq i} \text{valuations of the other players}\right]\right)$$

In our setting
the Expected Groves mechanism
is not truthful.

Groves Realized

OPT

REALIZED TIMES
OF THE OTHER
PLAYERS

FINISH
TIME



After completing the task we have the
realized running times of the players

$$\text{Groves payment} = - \left(\text{Player 1} + \text{Player 2} \right) \quad \text{--“sum of the realized times of the other players”}$$

“The payments align the incentives of the players with the objective of the mechanism”

OPT

FINISH
TIME



Thus no incentive
to lie or miscompute!

Groves payment = $-(\text{[Portrait of man with mustache]} + \text{[Portrait of man with large white beard]})$

OBJECTIVE of
the mechanism

min ($\text{[Portrait of man with large white beard]} + \text{[Portrait of man with large white beard]} + \text{[Portrait of man with mustache]} + \text{[Portrait of man with large white beard]} + \text{[Portrait of man with large white beard]}$)

Utility of
selfish player:

max $-(\text{[Portrait of man with large white beard]} + \text{[Portrait of man with large white beard]} + \text{[Portrait of man with large white beard]} + \text{[Portrait of man with mustache]} + \text{[Portrait of man with large white beard]})$

valuation

Groves
payment

Valuations are **interdependent**

The valuation of a player depends on
whether another player
has already finished before him.

If we want our mechanism to have an
 $h_i(\text{types of the other players})$ part
(useful for getting properties like IR, etc)
 $p_i = \text{Groves} + h_i(\text{types of the other players})$

We have to consider the situation
when **the player who finishes** isn't there.

“How much does the player who finish contribute to the social welfare?”

OPT

FINISH
TIME



OPT_



What if the player who finishes wasn't there?

CONSISTENCY
PROPERTY

REALIZED

IN EXPECTATION



h part of the mechanism

Vickrey Variations

T_N := how long it takes a group N to finish the task (random variable)

r_N := realized value of T_N h_i (types of the other players) part

$p_i = E[T_N - T_i] + 0$		Expected Pure Groves (EPG)
$p_i = E[T_N - T_i] + E[T_{N \setminus \{i\}}]$		Expected Clarke (EC)
$p_i = -(r_N - r_i) + 0$		Realized Pure Groves (RPG)
$p_i = -(r_N - r_i) + E[T_{N \setminus \{i\}}]$		Clarke h (ChE) in Expectation
$p_i = -(r_N - r_i) + (r_N - r_i) + E[T_{N \setminus \{i\}} - r_{N \setminus \{i\}} T_{N \setminus \{i\}} \geq r_{N \setminus \{i\}}]$		(ChpE) Clarke h partially in Expectation

Groves
part

$h_i()$

This rewriting uses the
consistency property!

Properties of different Mechanisms

	efficient	truthful in dominant strategies	ex-post truthful	IR	no incentive to miscompute	payment 0 if fail
Clarke in Expectation (CE)	✓	✗	✗	✓	✗	✗
Pure Realized Groves (PRG)	✓	✗	✓	✗	✓	✗
Clarke h (ChE) in Expectation	✓	✗	✓	✓	✓	✗
Clarke h partially in Expectation (ChpE)	✓	✗	✓	✓	✓	✓

Main Theorem:

There exists an (ex-post) truthful mechanism that:

- (a) doesn't pay players who did not complete the task
- (b) satisfies IR in expectation
- (c) has positive payments
- (d) it is to the best interest of the players to exert full effort.

(Also generalizes to many tasks.)

No incentive to miscompute

Maybe the painters **reported their true distributions** but in the end decided it is **to their best interest to take a break** instead of painting!

If **no realized values are used** then players can just sit and compute nothing!



“The payments align the incentives of the players with the objective of the mechanism”

OPT

FINISH
TIME



Thus no incentive to lie or miscompute!

Groves payment = $-(\text{[Portrait of man with mustache]} + \text{[Portrait of man with large white beard]})$

OBJECTIVE of the mechanism

min ($\text{[Portrait of man with large white beard]} + \text{[Portrait of man with large white beard]} + \text{[Portrait of man with mustache]} + \text{[Portrait of man with large white beard]} + \text{[Portrait of man with large white beard]}$)

Utility of selfish player:

max $-(\text{[Portrait of man with large white beard]} + \text{[Portrait of man with large white beard]} + \text{[Portrait of man with large white beard]} + \text{[Portrait of man with mustache]} + \text{[Portrait of man with large white beard]})$

valuation

Groves payment

No incentive to lie or miscompute in ChpE

(Clarke h partially in Expectation)

Proof is more involved. Idea:

$$\text{payment}_i = -(r_N - r_i) + \underbrace{(r_N - r_i) + E[T_{N \setminus \{i\}} - r_{N \setminus \{i\}} | T_{N \setminus \{i\}} \geq r_{N \setminus \{i\}}]}_{h_i()}$$



When player i finishes the task he determines which part is taken in expectation and which not. Suppose that he had even more power and could “cut” h at any point then he could still not affect the Expectation of h by miscomputing!



**“Would you tell me, please,
which way I ought to go from
here?”**

**“That depends a good deal on
where you want to get to.”**

“I don't much care where –”

**“Then it doesn't matter which
way you go.”** Lewis Carol, *Alice in Wonderland*

Angelina Vidali

Duke University
Department of Computer Science



Paris Dauphine, 2014

Notice biographique

Postes

Postdoc: **Université Duke**, États-Unis, groupe du Prof. **Vincent Conitzer**, Septembre 2010-2011

Postdoc: **Université de Vienne**, Institut: Théorie et Applications d'Algorithmes, de Prof. **Monika Henzinger** 2010-12

Séjour de recherche, Institut for Advanced Studies, 23/5- 19/6/2011, **Université hébraïque de Jérusalem**, invitée par **Noam Nisan**.

Institut Max-Planck pour l'informatique, Chercheur Postdoctoral, Saarbrücken, Allemagne, 2009-10, Groupe: Algorithms de Prof. **Kurt Mehlhorn**

Formation

Doctorat en Informatique de l'Université d'Athènes (2009), (Très bien), **"Analyse des réseaux par la théorie des jeux."**

Directeur de thèse: **Elias Koutsoupas**

Master's (2005) (Très bien), Directeur de thèse: **Yiannis Moschovakis**, „Inter-university Programm: Logic, Algorithms and Computation” (MPLA)

Enseignements

Université de Saarland:

"Théorie de jeux algorithmique (Algorithmic Game Theory)"

Université d'Athènes:

"Algorithmique et Complexité (Algorithms and Complexity)",

"Sujets de théorie de jeux (Topics in Game Theory and Computation)",

"Calculabilité (Computability, Recursive Functions)",

"Mathématiques pour l'informatique et les Télécommunications
(Mathematics for Informatics and Telecommunications)"

"Algorithmes Probabilistes (Probabilistic Algorithms)"

Supervision des Thèses

Seyed Majid Zahedi, Doctorat, Université de Duke

Xiaoming Nate Xu, Doctorat, Arbitrage d'articles:

Hendrik Molter, Thèse de Bachelor, Université de Saarland 2010

Emmanouil Pountourakis, Thèse de Bachelor, Université d'Athènes 2008

Aikaterini Iatropoulou, Thèse de Master, Université d'Athènes 2007

paper in **ESA'10**,

Algorithmica special issue for ESA

Membre du Comité de lecture:

ACM conference Economics and Computation EC '13

Arbitrage d'articles pour les Conférences:

AAAI (Conference on **Artificial Intelligence**)

FOCS (IEEE Symposium on Foundations of

Computer Science), STOC (ACM Symposium on **Theory** of Computing),

SODA (ACM-SIAM Symposium on Discrete Algorithms),

ICALP (International Colloquium on Automata, Languages and Programming),

ESA (European Symposium on **Algorithms**),

EC (ACM Conference on **Electronic Commerce**),

ISSAC (International Symposium on **Symbolic and Algebraic Computation**),

SPAA (ACM Symposium on **Parallelism** in Algorithms and Architectures),

SWAT (Scandinavian Symposium and Workshops on Algorithm Theory),

WINE (**Internet & Network Economics**), CiE (Computability in Europe)

SAGT (International Symposium on **Algorithmic Game Theory**),

Arbitrage d'articles pour les revues: GEB (Games and **Economic** Behavior),

TOCS (Theory of Computing Systems), IPL (Information Processing Letters),

Journal of **Symbolic Computation**,

Maths of **Operations Research**, Journal of **Scheduling**.

Administration

Organisation de séminaires

Séminaire interdisciplinaire (Economique, Informatique, Business School)

sponsorisé par **Yahoo!** Université de Duke 2013-2014

Séminaire de théorie algorithmique des jeux, Université de Vienne, 2011-

Séminaire étudiant de théorie, Université d'Athènes, 2008-2009.

Cours intensif sur Bayésiennes et approximatifs mécanismes d'incitation,

Décembre 2010-Janvier 2011 Université d'Athènes, département d'informatique

Préparation des rapports de recherche pour les projets AEOLUS,

FRONTS, PENED **présentation des objectifs de recherche** de

l'Université d'Athènes entre le projet FRONTS (Lubeck, Allemande, 2008).

Représentante des étudiantes au conseil d'administration du

programme de master MPLA 2003-2005

Étudiante volontaire, organisation de: "3rd Athens Colloquium on

Algorithms and Complexity 2008" Université d'Athènes

Some of the main topics addressed by my publications

Revenue maximizing auctions



Scheduling selfish unrelated machines

Mechanism design with uncertain execution time




Strongly truthful mechanisms, externality resistant mechanisms

Cost-sharing group strategyproofness




Mechanisms with verification (restricting which lies agents can report)




Mechanism Design

A single choice for the whole society

Voting
Public Projects
Auctions
Scheduling



We need to construct a function that takes as input the **preferences of many different individuals** and „**amalgamates**” (/aggregates) them into a **single** preference or choice.




Mechanism Design

Design a game whose **outcome** is an **equilibrium** for the players.

truthful

We want mechanisms,
that don't give players an incentive to lie.

When a player has an incentive to lie...



Mille millions de mille sabords!
 $v(\text{EELV}) > v(\text{PS}) > v(\text{UMP})$


Can the captain **achieve a more preferable outcome** by lying
about his political preferences?

- avec Monika Henzinger under revision TOCS
- avec Elias Koutsoupias Algorithmica 2012
- avec Emmanouil Pountourakis Algorithmica 2012 special issue for ESA'10
- avec George Christodoulou, Elias Koutsoupias Algorithmica 2009

Conferences internationales (avec comité de lecture et actes)

- avec Vincent Conitzer AAAI'14 to appear
- avec Amos Fiat, Anna Karlin, Elias Koutsoupias ITCS'13 (citations: 6)
- WINE'11 (citations: 1)
- avec Monika Henzinger ESA'11 (citations: 2)
- avec Emmanouil Pountourakis ESA'10 (citations: 5)
- WINE '09 (citations: 2)
- avec George Christodoulou, Elias Koutsoupias ESA '08 (citations: 18)
- avec Elias Koutsoupias MFCS '07 (citations: 41)
- avec George Christodoulou, Elias Koutsoupias SODA'07 (citations: 61)

Scheduling Selfish unrelated Machines an auction for selling jobs to agents



📖 „A lower bound for scheduling mechanisms.” E. Koutsoupias, G. Christodoulou and A. V. Goldberg, conference version: 18th ACM-SIAM Symposium on Discrete Algorithms SODA'07.

journal: Algorithmica '09 (citations: 63)

📖 „A lower bound of $1+\varphi$ for scheduling mechanisms.” E. Koutsoupias, and A. Vidali, conference version: 32nd International Symposium on Mathematical Foundations of Computer Science MFCS '07. journal: Algorithmica '12 (citations: 28)

Publications [2/2]

Conferences internationales (sans comité de lecture et actes)

- GREAT: Economic and Algorithmic Theory Week, 2014 (**invitée**)
- INFORMS 2013 in Minneapolis, MN, held on Oct. 2013 (**invitée**)
- Stony Brook Workshop on Computational Game Theory, July 2013 (**invitée**)
- ISMP'12, International Symposium on Mathematical Programming, Berlin, (**invitée**)
- Summer school on Algorithmic Game Theory, Samos, 14-21/7 2012.(**invitée**)
- GREAT: Economic and Algorithmic Theory Week, 2011 (**invitée**)
- Innovations in Algorithmic Game Theory workshop, May 2011, Hebrew University of Jerusalem (**invitée**)
- Bertinoro Workshop on Frontiers of Algorithmic Mechanism Design, 2010 (**invitée**)

Conferences nationales (sans comité de lecture et actes)

- ACAC'13, ACAC'11, ACAC'10, ACAC'09, ACAC'08

Posters

ACM EC(Conference on Economics and Computation)' 14
New York Computation and Economics Day '13


Theses

- Game-theoretic Analysis of Networks
- Continued Fractions and the Euclidean Algorithm

Scheduling under Uncertain Execution Time

joint work with Vince Conitzer to appear in
Twenty-Eighth AAAI Conference on Artificial Intelligence AAAI'14

The Queen wants a painting for her palace.
Every day she decides which painters will draw.
Goal: minimize $E[\text{painting time}]$



Creativity is unpredictable!

Painters don't know how much time it's going to take them and need incentives to draw!

Time painter i needs to finish the job $\sim$ distribution f_i

painter i knows f_i but not t_i and f_{-i}

Agents are selfish
want to maximize:

$t_1 \sim f_1$ utility = $E[\text{time spent painting} - \text{payment}]$



$t_1 \sim f_1$



$t_2 \sim f_2$



$t_3 \sim f_3$



job to process

Crowdsourcing



Challenges:

- Split a complicated very big job into many simpler tasks!
- Give incentives to the agents to complete the tasks!
- Ensure solution quality

Crowdsourcing Application

We want to solve a problem (e.g. a SAT instance), by running a crowdsourcing contest.

Payment 0 if you fail: Only the winner (agent who solves the problem first) will get paid.

Uncertainty: The agents don't know how much time they will need.

Main Theorem:

There exists an (ex-post) truthful mechanism that:

- (a) doesn't pay players who did not complete the task
- (b) the agents expected utility is positive (IR in expectation)
- (c) has positive payments
- (d) it is to the best interest of the players to exert full effort.

(Also generalizes to many tasks.)

Ongoing and future projects

Mechanism Design without money

Fair Division, Cake cutting
Kidney Exchange,
Facility Location,
Scheduling Problems

Truthful in Expectation algorithms,
Envy free solutions

Mechanisms with verification

Virtual money

Characterization of Truthful Mechanisms




Mechanism design without money: **Fair and efficient division of resources of Computer Architecture**

How to:

Fairly assign to tasks

- multiple cores
- memory



Efficiently divide a computer
into many virtual machines



Scheduling: Splitting a job into tasks and parallelizing

(ongoing work with Vincent Conitzer, Benjamin C Lee (Duke ECE) and Seyed Majid Zahedi)

Combinatorial Voting

Voting on what the meal is going to include:



Voting on the order in which tasks are going to be scheduled

Preferences without intensity



>



Preferences with intensity

$$v(\text{glass of red wine}) = 18$$

$$v(\text{glass of white wine}) = 9$$

Immediate collaborations in LAMSADE:

Pole: Aide à la décision

Projet: Agents intelligents pour la décision et le raisonnement

Jérôme Lang (Voting, Fair Division, Judgement aggregation)

Gabriella Pigozzi (Social Choice, Judgment Aggregation)

Stefano Moretti (Cost sharing, auctions)

Stéphane Airiau (Resource allocation)

aussi:

Projet: Optimisation combinatoire multicritère

Mohamed Ali Aloulou (Scheduling)

Cristina Bozgan (approximation algorithms, complexity)

Laurent Gourvès, Jérôme Monnot, (scheduling, Fair allocation, Price of Anarchy)

aussi?

Projet: Modélisation des préférences et aide multicritère à la décision

Denis Bouyssou, Jérôme Lang

Merci beaucoup!