

Economic Issues in Shared Infrastructures

Costas Courcoubetis

Department of Computer Science, Athens University of
Economics and Business

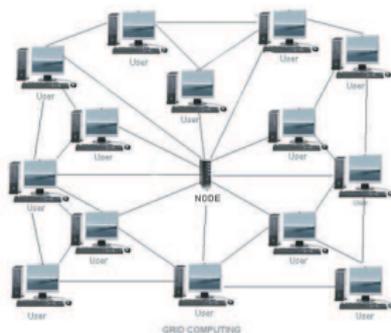
(work with Richard Weber, Statistical Laboratory, University of
Cambridge)

CALTECH, August 2009

Resource management in virtual facilities

The topic of this talk: efficient **resource management** within **shared infrastructures**.

E.g., Grid computing



This is complex because of the details of technology specificities. Mathematics/economics can help to highlight some key issues.

Motivation: The GridEcon Project

Recent advances in high speed Internet access, coupled with the ever increasing performance and reducing cost of computers, have made this vision of Grid computing a reality.

Some well known companies such as Amazon, IBM, and Sun are already providing basic Grid services.

Motivation: The GridEcon Project

Recent advances in high speed Internet access, coupled with the ever increasing performance and reducing cost of computers, have made this vision of Grid computing a reality.

Some well known companies such as Amazon, IBM, and Sun are already providing basic Grid services.

GridEcon is a European community funded project, exploring the perceived economic barriers to the adoption of Grid Computing.

Motivation: The GridEcon Project

Recent advances in high speed Internet access, coupled with the ever increasing performance and reducing cost of computers, have made this vision of Grid computing a reality.

Some well known companies such as Amazon, IBM, and Sun are already providing basic Grid services.

GridEcon is a European community funded project, exploring the perceived economic barriers to the adoption of Grid Computing.

It sets out to develop economic models that would allow a market to be developed that would exploit the benefits of Grid computing, and to develop the components that would make such a market a reality.

The key issue in this talk

Agents (users) have **private information** (about the value of the tasks they wish to carry out).

This creates a problem for efficiently sharing resources.

- ▶ Agents will attempt to **free-ride**.
- ▶ Obvious policies (like 'internal market', or 'equal sharing') may not be suitable.

The key issue in this talk

Agents (users) have **private information** (about the value of the tasks they wish to carry out).

This creates a problem for efficiently sharing resources.

- ▶ Agents will attempt to **free-ride**.
- ▶ Obvious policies (like 'internal market', or 'equal sharing') may not be suitable.

How one chooses to share a facility's resources will influence what agents reveal of their private information.

We would like agents to truthfully reveal their privately held information since then we can operate the facility more efficiently.

A model of a managed shared infrastructure

- ▶ An infrastructure as composed of **resources**.
(links, servers, buffers, etc.)

A model of a managed shared infrastructure

- ▶ An infrastructure as composed of **resources**.
(links, servers, buffers, etc.)
- ▶ It can be **operated in various ways**, $\omega_1, \omega_2, \dots$
(by scheduling, routing, bandwidth allocation, etc.)

A model of a managed shared infrastructure

- ▶ An infrastructure as composed of **resources**.
(links, servers, buffers, etc.)
- ▶ It can be **operated in various ways**, $\omega_1, \omega_2, \dots$
(by scheduling, routing, bandwidth allocation, etc.)
- ▶ The subset of agents who wish to use resources of the infrastructure, say S , differs from day to day.
This occurs with frequency $\alpha(S)$.

A model of a managed shared infrastructure

- ▶ An infrastructure as composed of **resources**.
(links, servers, buffers, etc.)
- ▶ It can be **operated in various ways**, $\omega_1, \omega_2, \dots$
(by scheduling, routing, bandwidth allocation, etc.)
- ▶ The subset of agents who wish to use resources of the infrastructure, say S , differs from day to day.
This occurs with frequency $\alpha(S)$.
- ▶ If operated in manner ω on day t then agent i has **utility**

$$\theta_{i,t} u_i(\omega)$$

$u_i(\cdot)$ is public knowledge, but **only agent i knows $\theta_{i,t}$** .

A model of a managed shared infrastructure

- ▶ An infrastructure as composed of **resources**.
(links, servers, buffers, etc.)
- ▶ It can be **operated in various ways**, $\omega_1, \omega_2, \dots$
(by scheduling, routing, bandwidth allocation, etc.)
- ▶ The subset of agents who wish to use resources of the infrastructure, say S , differs from day to day.
This occurs with frequency $\alpha(S)$.
- ▶ If operated in manner ω on day t then agent i has **utility**

$$\theta_{i,t} u_i(\omega)$$

$u_i(\cdot)$ is public knowledge, but **only agent i knows $\theta_{i,t}$** .

- ▶ ω is to be chosen as a function of S and of the declared $\theta_t = (\theta_{1,t}, \dots, \theta_{n,t})$.

Agents pay for operating cost

Agent i is charged a fee $p_i(S, \theta)$.

Fees are to cover a daily operating cost, c , so we require

$$E_{S,\theta} [p_1(S, \theta) + \cdots + p_n(S, \theta)] \geq c$$

Agents pay for operating cost

Agent i is charged a fee $p_i(S, \theta)$.

Fees are to cover a daily operating cost, c , so we require

$$E_{S, \theta} [p_1(S, \theta) + \cdots + p_n(S, \theta)] \geq c$$

Agent i wishes to maximize his expected net benefit

$$nb_i(\theta_i) = E_{S, \theta_{-i}} [\theta_i u_i(\omega(S, \theta)) - p_i(S, \theta)]$$

Agents pay for operating cost

Agent i is charged a fee $p_i(S, \theta)$.

Fees are to cover a daily operating cost, c , so we require

$$E_{S, \theta} [p_1(S, \theta) + \dots + p_n(S, \theta)] \geq c$$

Agent i wishes to maximize his expected net benefit

$$nb_i(\theta_i) = E_{S, \theta_{-i}} [\theta_i u_i(\omega(S, \theta)) - p_i(S, \theta)]$$

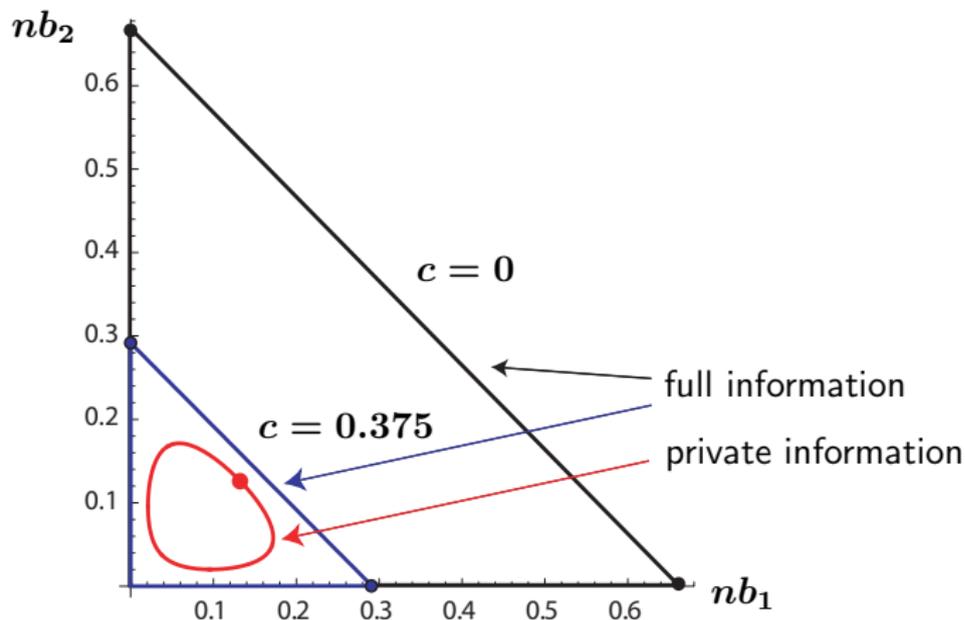
In some situations we may take the fee as money.

In others we may wish to take the fee as a contribution to the pool of resources that is available in the infrastructure.

The efficient frontier

We wish to find Pareto optimal points of the vector

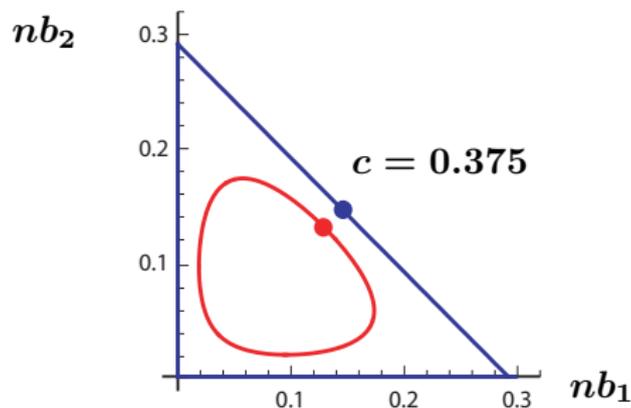
$$(nb_1, \dots, nb_n) = E_{\theta} [nb_1(\theta_1), \dots, nb_n(\theta_n)]$$



Maximum social welfare

Suppose we wish to find the particular point that maximizes

$$nb_1 + \dots + nb_n = E_{S,\theta} [\theta_1 u_1(\omega(S)) + \dots + \theta_n u_n(\omega(S))] - c$$



We call this the 'social welfare'.

Our infrastructure optimization problem

The problem of economics for infrastructure optimization is as follows.

- ▶ Say how the infrastructure will be operated for possible subset of users S .

Our infrastructure optimization problem

The problem of economics for infrastructure optimization is as follows.

- ▶ Say how the infrastructure will be operated for possible subset of users S .
- ▶ Say what fees will be collected from users.

Our infrastructure optimization problem

The problem of economics for infrastructure optimization is as follows.

- ▶ Say how the infrastructure will be operated for possible subset of users S .
- ▶ Say what fees will be collected from users.

Do the two things above, as function of declared θ_i , so that:

1. Users find it in their best interest to truthfully reveal their θ_i .
2. Users will see positive expected net benefit from participation.
3. Expected total fees cover the daily running cost, say c .
4. Expected social welfare (total net benefit) is maximized

Example: a Bridge

A bridge may or may not be built. There are 2 potential users.



Mathematical Bridge, Queens' College, Cambridge

Example: a Bridge

A bridge may or may not be built. There are 2 potential users.



Mathematical Bridge, Queens' College, Cambridge

If it is built (at cost \$1) then user i benefits by θ_i . Knowing θ_1 and θ_2 , we should build the bridge if $\theta_1 + \theta_2 > 1$.

Example: a Bridge

A bridge may or may not be built. There are 2 potential users.



Mathematical Bridge, Queens' College, Cambridge

If it is built (at cost \$1) then user i benefits by θ_i . Knowing θ_1 and θ_2 , we should build the bridge if $\theta_1 + \theta_2 > 1$.

If we build the bridge we must charge for the cost. Suppose we decide to charge user i a fee of $\theta_i / (\theta_1 + \theta_2)$. The problem is that user i will have an incentive to under-report his true value of θ_i .

Example: a Bridge

A bridge may or may not be built. There are 2 potential users.



Mathematical Bridge, Queens' College, Cambridge

If it is built (at cost \$1) then user i benefits by θ_i . Knowing θ_1 and θ_2 , we should build the bridge if $\theta_1 + \theta_2 > 1$.

If we build the bridge we must charge for the cost. Suppose we decide to charge user i a fee of $\theta_i / (\theta_1 + \theta_2)$. The problem is that user i will have an incentive to under-report his true value of θ_i .

What is the best fee mechanism, $p_1(\theta_1, \theta_2)$ and $p_2(\theta_1, \theta_2)$?

Fees should incentivize users to truthfully reveal θ_1, θ_2 , with

$p_1(\theta_1, \theta_2) + p_2(\theta_1, \theta_2) = 1$ or 0 , as bridge is built or not built.

Example: A Bridge

In the example above the bridge is a **nonrivalrous good**.

Example: A Bridge

In the example above the bridge is a **nonrivalrous good**.

But what if its use is **rivalrous**?

E.g., suppose that on each given day only one person may use the bridge. On each day, users 1 and 2 wish to use the bridge, independently, with probabilities α_1 and α_2 .

Example: A Bridge

In the example above the bridge is a **nonrivalrous good**.

But what if its use is **rivalrous**?

E.g., suppose that on each given day only one person may use the bridge. On each day, users 1 and 2 wish to use the bridge, independently, with probabilities α_1 and α_2 .

Now we must decide

- ▶ whether or not the bridge is built;
- ▶ what contributions the users should make towards its cost;
- ▶ who gets to use the bridge on those days that both users say that they wish to do so.

All this is to be decided as some function of the initially declared θ_1, θ_2 .

Motivation

Similarly, in grid computing: how do we incentivize agents to participate and contribute computational resource? what size of computational resource will be installed and shared? what contributions should agents make towards its cost — or what amounts of resource should they be willing to contribute? how should the resource be shared?

How do we address these questions?

Motivation

Similarly, in grid computing: how do we incentivize agents to participate and contribute computational resource? what size of computational resource will be installed and shared? what contributions should agents make towards its cost — or what amounts of resource should they be willing to contribute? how should the resource be shared?

How do we address these questions?

Are auction theory and Mechanism Design appropriate? And under what assumptions on our model are these applicable?

Motivation

Similarly, in grid computing: how do we incentivize agents to participate and contribute computational resource? what size of computational resource will be installed and shared? what contributions should agents make towards its cost — or what amounts of resource should they be willing to contribute? how should the resource be shared?

How do we address these questions?

Are auction theory and Mechanism Design appropriate? And under what assumptions on our model are these applicable?

What is fundamentally new in this problem?

Motivation

Similarly, in grid computing: how do we incentivize agents to participate and contribute computational resource? what size of computational resource will be installed and shared? what contributions should agents make towards its cost — or what amounts of resource should they be willing to contribute? how should the resource be shared?

How do we address these questions?

Are auction theory and Mechanism Design appropriate? And under what assumptions on our model are these applicable?

What is fundamentally new in this problem?

Can we describe optimal policies?

Assumptions

Two possibilities:

1. The facility already exists, has fixed size Q and known operating cost, or

Assumptions

Two possibilities:

1. The facility already exists, has fixed size Q and known operating cost, or
2. The facility does not exist, its size will be determined by contributions of the participants.

Assumptions

Two possibilities:

1. The facility already exists, has fixed size Q and known operating cost, or
2. The facility does not exist, its size will be determined by contributions of the participants.

How to share resources and recover costs?

Assumptions

Two possibilities:

1. The facility already exists, has fixed size Q and known operating cost, or
2. The facility does not exist, its size will be determined by contributions of the participants.

How to share resources and recover costs?

- ▶ Easy when we know utilities of participants.

Assumptions

Two possibilities:

1. The facility already exists, has fixed size Q and known operating cost, or
2. The facility does not exist, its size will be determined by contributions of the participants.

How to share resources and recover costs?

- ▶ Easy when we know utilities of participants.
- ▶ But in practice agents' utilities are private information. We must design the system to operate well, under the constraint that each agent will reveal information in a manner that is to his best advantage.

Example: scheduling a server

- ▶ Suppose n agents share a single server. Agent i generates a jobs as a Poisson process of rate λ_i , whose service times are exponentially distributed with parameter 1 .

Example: scheduling a server

- ▶ Suppose n agents share a single server. Agent i generates a jobs as a Poisson process of rate λ_i , whose service times are exponentially distributed with parameter 1.
- ▶ Initially, agents contribute resource amounts y_1, \dots, y_n . This results in a server of rate $\sum_k y_k$. Under FCFS scheduling all jobs have mean waiting time $1/(\sum_k y_k - \sum_k \lambda_k)$.

Example: scheduling a server

- ▶ Suppose n agents share a single server. Agent i generates a jobs as a Poisson process of rate λ_i , whose service times are exponentially distributed with parameter 1.
- ▶ Initially, agents contribute resource amounts y_1, \dots, y_n . This results in a server of rate $\sum_k y_k$. Under FCFS scheduling all jobs have mean waiting time $1/(\sum_k y_k - \sum_k \lambda_k)$.
- ▶ Agent i suffers delay cost, so his net benefit is, say,

$$nb_i = \lambda_i r - \theta_i \lambda_i \frac{1}{\sum_k y_k - \sum_k \lambda_k} - y_i.$$

θ_i is private information of agent i , but it has an *a priori* distribution that is public information.

Optimal queue scheduling

Instead of declaring contributions they are willing to make, we can imagine that agents (equivalently) declare their θ_i .

Suppose $\theta_1 < \theta_2 < \dots < \theta_n$.

As a function of these declarations we take contributions of the form $y(\theta_i)$ from some subset of agents $i = 1, \dots, j$ (a set with smallest θ_i).

Optimal queue scheduling

Instead of declaring contributions they are willing to make, we can imagine that agents (equivalently) declare their θ_i .

Suppose $\theta_1 < \theta_2 < \dots < \theta_n$.

As a function of these declarations we take contributions of the form $y(\theta_i)$ from some subset of agents $i = 1, \dots, j$ (a set with smallest θ_i).

We employ a priority scheduling policy in which priority is always given to the current job belonging to the agent with greatest θ_i .

Optimal queue scheduling

Instead of declaring contributions they are willing to make, we can imagine that agents (equivalently) declare their θ_i .

Suppose $\theta_1 < \theta_2 < \dots < \theta_n$.

As a function of these declarations we take contributions of the form $y(\theta_i)$ from some subset of agents $i = 1, \dots, j$ (a set with smallest θ_i).

We employ a priority scheduling policy in which priority is always given to the current job belonging to the agent with greatest θ_i .

Under this scheme, an agent with too great a θ_i will find unprofitable to consider participating.

$y_i(\theta_i)$ is increasing in θ_i , and is determined by an incentive compatibility condition.

Example: a scalar resource

- ▶ 2 participants, both present on all days.

Example: a scalar resource

- ▶ 2 participants, both present on all days.
- ▶ On day t , agent i has utility for resource of $\theta_{i,t}u(x)$, assumed known to be distributed $U[0, 1]$.

Example: a scalar resource

- ▶ 2 participants, both present on all days.
- ▶ On day t , agent i has utility for resource of $\theta_{i,t}u(x)$, assumed known to be distributed $U[0, 1]$.
- ▶ The infrastructure provides a single resource, parameterized by a number (such as computing cycles), so operating methods correspond to allocations:

$$\{\omega\} \equiv \{x_1, x_2 : x_1 + x_2 \leq 1\}$$

Example: a scalar resource

- ▶ 2 participants, both present on all days.
- ▶ On day t , agent i has utility for resource of $\theta_{i,t}u(x)$, assumed known to be distributed $U[0, 1]$.
- ▶ The infrastructure provides a single resource, parameterized by a number (such as computing cycles), so operating methods correspond to allocations:

$$\{\omega\} \equiv \{x_1, x_2 : x_1 + x_2 \leq 1\}$$

Suppose $u_i(x) = x$. Focus on one day t ; with $\theta_i = \theta_{i,t}$.

$$E_{\theta_1, \theta_2} \left[\max_{\substack{x_1, x_2 \\ x_1 + x_2 \leq 1}} \{\theta_1 u_1(x_1) + \theta_2 u_2(x_2)\} \right] = E[\max\{\theta_1, \theta_2\}] = \frac{2}{3}$$

We call this the '**first best**'.

The second-best solution

A 'second-best' is with fee structure:

$$p_i(\theta_i) = \begin{cases} 0, & \theta_i < \theta_0 \\ \frac{1}{2}(\theta_i^2 + \theta_0^2), & \theta_i \geq \theta_0 \end{cases}$$

Agent i will not wish to participate if $\theta_i < \theta_0$, since his net benefit cannot be positive.

The entire resource is allocated to the agent declaring the greatest θ_i , provided this is $> \theta_0$.

Thus, the resource is given wholly to one agent, but perhaps to neither.

But both agents may pay.

The 'solution'

This solution has the advantages that

- ▶ Agents are incentivized to truthful.

The 'solution'

This solution has the advantages that

- ▶ Agents are incentivized to truthful.
- ▶ The sum of the expected payments is

$$E\left[p_1(\theta_1) + p_2(\theta_2)\right] = 1/3 + \theta_0^2 - (2/3)\theta_0^3.$$

The 'solution'

This solution has the advantages that

- ▶ Agents are incentivized to truthful.
- ▶ The sum of the expected payments is

$$E\left[p_1(\theta_1) + p_2(\theta_2)\right] = 1/3 + \theta_0^2 - (2/3)\theta_0^3.$$

- ▶ The expected social welfare is decreasing in θ_0 .

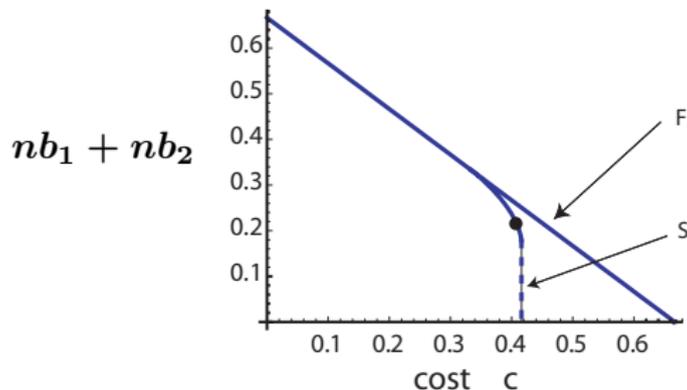
But by taking $1/3 + \theta_0^2 - (2/3)\theta_0^3 = c$ we maximize the social welfare of

$$nb_1 + nb_2 = E\left[\sum_{i=1}^2 \theta_i u_i(x_i) - p_i(\theta_i)\right]$$

subject to covering cost c .

Second-best versus first-best

Expected social welfare as a function of c , compared to the first-best value.



For $c \in [0.333, 0.416]$ the second-best falls short of the first-best. There is no way to cover a cost greater than $\frac{5}{12} = 0.416$.

Other mechanisms can also work.

Other mechanisms can also work.

- ▶ We can ensure $p_1(\theta_1, \theta_2) + p_2(\theta_1, \theta_2) = c$.

$$p_1(\theta_1, \theta_2) = \frac{1}{2}c + \frac{1}{2}(\theta_1^2 + \theta_0^2)1_{\{\theta_1 > \theta_0\}} - \frac{1}{2}(\theta_2^2 + \theta_0^2)1_{\{\theta_2 > \theta_0\}}$$

We call this 'ex-post' cost-covering.

Other mechanisms can also work.

- ▶ We can ensure $p_1(\theta_1, \theta_2) + p_2(\theta_1, \theta_2) = c$.

$$p_1(\theta_1, \theta_2) = \frac{1}{2}c + \frac{1}{2}(\theta_1^2 + \theta_0^2)1_{\{\theta_1 > \theta_0\}} - \frac{1}{2}(\theta_2^2 + \theta_0^2)1_{\{\theta_2 > \theta_0\}}$$

We call this 'ex-post' cost-covering.

- ▶ We can ensure ex-post incentive compatibility and rationality. I.e., that an agent only pays if he gets resource, and is happy after-the-fact with the θ_i he declared.

$$p_1(\theta_1, \theta_2) = \max(\theta_0, \theta_2)1_{\{\theta_1 > \max(\theta_0, \theta_2)\}}$$

A concave utility

Suppose $u_i(x) = \sqrt{x}$

Now the resource is shared differently.

The optimal policy is found by solving a Lagrangian dual problem

$$\min_{\lambda \geq 0} \left\{ E_{\theta_1, \theta_2} \left[\max_{\substack{x_1, x_2 \geq 0 \\ x_1 + x_2 \leq 1}} \sum_{i=1}^2 h_\lambda(\theta_i) u_i(x_i) \right] - (1 + \lambda)c \right\}.$$

where $h_\lambda(\theta_i) = (\theta_i + \lambda(2\theta_i - 1))$ and

$$x_i(\theta_1, \theta_2) = \frac{h_\lambda(\theta_i)^2}{\sum_{j=1}^2 h_\lambda(\theta_j)^2}$$

As λ increases the fee structure changes, so that greater cost can be covered. The social welfare decreases, but is maximized subject to the constraint of covering the cost.

The role of operating policy

The resource is not allocated in the 'most efficient' way.

That would be

$$x_i(\theta_1, \theta_2) = \frac{\theta_i^2}{\theta_1^2 + \theta_2^2} \quad \left(\text{not } \frac{h_\lambda(\theta_i)^2}{h_\lambda(\theta_1)^2 + h_\lambda(\theta_2)^2} \right)$$

The role of operating policy

The resource is not allocated in the 'most efficient' way.

That would be

$$x_i(\theta_1, \theta_2) = \frac{\theta_i^2}{\theta_1^2 + \theta_2^2} \quad \left(\text{not } \frac{h_\lambda(\theta_i)^2}{h_\lambda(\theta_1)^2 + h_\lambda(\theta_2)^2} \right)$$

This is a key lesson:

If one wishes to optimally incentivize participation in shared infrastructures, and to make the most of the resources available, then both the (i) fees, and (ii) policies for ' resource sharing, must play a part in providing the correct incentives to users.

Building a facility from scratch

A different model: facility of size Q , costing $c(Q) = Q$ (per slot), is formed by **initial contributions** of agents. **These are incentivized to contribute because their contribution will affect the amount of resources they will get at run time.** Probably a good model for virtual Grid infrastructures.

Building a facility from scratch

A different model: facility of size Q , costing $c(Q) = Q$ (per slot), is formed by **initial contributions** of agents. **These are incentivized to contribute because their contribution will affect the amount of resources they will get at run time.** Probably a good model for virtual Grid infrastructures.

- ▶ $\theta_{i,t} = \theta_i$ for all t (**private information**).

Building a facility from scratch

A different model: facility of size Q , costing $c(Q) = Q$ (per slot), is formed by **initial contributions** of agents. **These are incentivized to contribute because their contribution will affect the amount of resources they will get at run time.** Probably a good model for virtual Grid infrastructures.

- ▶ $\theta_{i,t} = \theta_i$ for all t (**private information**).
- ▶ Agent i is 'on-off' w.p. $\alpha_i, 1 - \alpha_i$. (**public information**).

Building a facility from scratch

A different model: facility of size Q , costing $c(Q) = Q$ (per slot), is formed by **initial contributions** of agents. **These are incentivized to contribute because their contribution will affect the amount of resources they will get at run time.** Probably a good model for virtual Grid infrastructures.

- ▶ $\theta_{i,t} = \theta_i$ for all t (**private information**).
- ▶ Agent i is 'on-off' w.p. $\alpha_i, 1 - \alpha_i$. (**public information**).
- ▶ Sharing policy is $x_i(\theta, S)$, S = set of agents 'on'.

Building a facility from scratch

A different model: facility of size Q , costing $c(Q) = Q$ (per slot), is formed by **initial contributions** of agents. **These are incentivized to contribute because their contribution will affect the amount of resources they will get at run time.** Probably a good model for virtual Grid infrastructures.

- ▶ $\theta_{i,t} = \theta_i$ for all t (**private information**).
- ▶ Agent i is 'on-off' w.p. $\alpha_i, 1 - \alpha_i$. (**public information**).
- ▶ Sharing policy is $x_i(\theta, S)$, S = set of agents 'on'.
- ▶ System planner posts how he will compute agents' contributions and the $x_i(\theta, S)$ as functions of the θ_i s that they declare.

Building a facility from scratch

A different model: facility of size Q , costing $c(Q) = Q$ (per slot), is formed by **initial contributions** of agents. **These are incentivized to contribute because their contribution will affect the amount of resources they will get at run time.** Probably a good model for virtual Grid infrastructures.

- ▶ $\theta_{i,t} = \theta_i$ for all t (**private information**).
- ▶ Agent i is 'on-off' w.p. $\alpha_i, 1 - \alpha_i$. (**public information**).
- ▶ Sharing policy is $x_i(\theta, S)$, S = set of agents 'on'.
- ▶ System planner posts how he will compute agents' contributions and the $x_i(\theta, S)$ as functions of the θ_i s that they declare.
- ▶ Agents declare θ_i s and system runs according to posted policy.

Analysis of proportional sharing

How do sharing policies affect incentives for agents to contribute?

Consider simple case of 2 identical agents $\theta_1 = \theta_2 = 1$,

$$c(Q) = Q.$$

Analysis of proportional sharing

How do sharing policies affect incentives for agents to contribute?

Consider simple case of 2 identical agents $\theta_1 = \theta_2 = 1$,

$c(Q) = Q$.

Agent i contributes q_i . Agent 1 has net benefit

$$\alpha_1(1 - \alpha_2)u(x_1^{\{1\}}) + \alpha_1\alpha_2u(x_1^{\{1,2\}}) - q_1 .$$

Analysis of proportional sharing

How do sharing policies affect incentives for agents to contribute?

Consider simple case of 2 identical agents $\theta_1 = \theta_2 = 1$,

$c(Q) = Q$.

Agent i contributes q_i . Agent 1 has net benefit

$$\alpha_1(1 - \alpha_2)u(x_1^{\{1\}}) + \alpha_1\alpha_2u(x_1^{\{1,2\}}) - q_1 .$$

Consider 4 possible sharing disciplines:

Analysis of proportional sharing

How do sharing policies affect incentives for agents to contribute?

Consider simple case of 2 identical agents $\theta_1 = \theta_2 = 1$,

$c(Q) = Q$.

Agent i contributes q_i . Agent 1 has net benefit

$$\alpha_1(1 - \alpha_2)u(x_1^{\{1\}}) + \alpha_1\alpha_2u(x_1^{\{1,2\}}) - q_1.$$

Consider 4 possible sharing disciplines:

- ▶ **Acting alone:** $x_i^{\{i\}} = x_i^{\{1,2\}} = q_i$.

Analysis of proportional sharing

How do sharing policies affect incentives for agents to contribute?

Consider simple case of 2 identical agents $\theta_1 = \theta_2 = 1$,

$c(Q) = Q$.

Agent i contributes q_i . Agent 1 has net benefit

$$\alpha_1(1 - \alpha_2)u(x_1^{\{1\}}) + \alpha_1\alpha_2u(x_1^{\{1,2\}}) - q_1.$$

Consider 4 possible sharing disciplines:

- ▶ **Acting alone:** $x_i^{\{i\}} = x_i^{\{1,2\}} = q_i$.
- ▶ **Equal sharing:** $x_i^{\{i\}} = q_1 + q_2$ and $x_i^{\{1,2\}} = \frac{1}{2}(q_1 + q_2)$.

Analysis of proportional sharing

How do sharing policies affect incentives for agents to contribute?

Consider simple case of 2 identical agents $\theta_1 = \theta_2 = 1$,

$$c(Q) = Q.$$

Agent i contributes q_i . Agent 1 has net benefit

$$\alpha_1(1 - \alpha_2)u(x_1^{\{1\}}) + \alpha_1\alpha_2u(x_1^{\{1,2\}}) - q_1.$$

Consider 4 possible sharing disciplines:

- ▶ **Acting alone:** $x_i^{\{i\}} = x_i^{\{1,2\}} = q_i$.
- ▶ **Equal sharing:** $x_i^{\{i\}} = q_1 + q_2$ and $x_i^{\{1,2\}} = \frac{1}{2}(q_1 + q_2)$.
- ▶ **Proportional sharing:**

$$x_i^{\{i\}} = q_1 + q_2, \quad x_i^{\{1,2\}} = \frac{q_i}{q_1 + q_2}(q_1 + q_2).$$

Analysis of proportional sharing

How do sharing policies affect incentives for agents to contribute?

Consider simple case of 2 identical agents $\theta_1 = \theta_2 = 1$,

$$c(Q) = Q.$$

Agent i contributes q_i . Agent 1 has net benefit

$$\alpha_1(1 - \alpha_2)u(x_1^{\{1\}}) + \alpha_1\alpha_2u(x_1^{\{1,2\}}) - q_1.$$

Consider 4 possible sharing disciplines:

- ▶ **Acting alone:** $x_i^{\{i\}} = x_i^{\{1,2\}} = q_i$.
- ▶ **Equal sharing:** $x_i^{\{i\}} = q_1 + q_2$ and $x_i^{\{1,2\}} = \frac{1}{2}(q_1 + q_2)$.
- ▶ **Proportional sharing:**

$$x_i^{\{i\}} = q_1 + q_2, \quad x_i^{\{1,2\}} = \frac{q_i}{q_1 + q_2}(q_1 + q_2).$$

- ▶ **s-Proportional sharing:**

$$x_i^{\{i\}} = q_1 + q_2, \quad x_i^{\{1,2\}} = \frac{q_i^s}{q_1^s + q_2^s}(q_1 + q_2).$$

Results for $\alpha_i = \alpha = 0.8$, $u(x) = 10 - 1/x$

scheme	social welfare	values of q_1, q_2
Acting alone	$r\alpha - 2\sqrt{\alpha}$ 6.21115	$\sqrt{\alpha}$ 0.894427
Equal sharing $s = 0$	$r\alpha - \frac{3}{2}\sqrt{\alpha(1+\alpha)}$ 6.2	$\frac{1}{2}\sqrt{\alpha(1+\alpha)}$ 0.6
Proportional sharing $s = 1$	$r\alpha - \frac{\sqrt{\alpha}(3+5\alpha)}{2\sqrt{1+3\alpha}}$ 6.30225	$\frac{1}{2}\sqrt{\alpha(1+3\alpha)}$ 0.824621
Central planner $s = \frac{1}{2}(1 + 1/\alpha)$	$r\alpha - \sqrt{2\alpha(1+\alpha)}$ 6.30294	$\sqrt{\alpha(1+\alpha)}/2$ 0.848528

Results for $\alpha_i = \alpha = 0.8$, $u(x) = 10 - 1/x$

scheme	social welfare	values of q_1, q_2
Acting alone	$r\alpha - 2\sqrt{\alpha}$ 6.21115	$\sqrt{\alpha}$ 0.894427
Equal sharing $s = 0$	$r\alpha - \frac{3}{2}\sqrt{\alpha(1+\alpha)}$ 6.2	$\frac{1}{2}\sqrt{\alpha(1+\alpha)}$ 0.6
Proportional sharing $s = 1$	$r\alpha - \frac{\sqrt{\alpha}(3+5\alpha)}{2\sqrt{1+3\alpha}}$ 6.30225	$\frac{1}{2}\sqrt{\alpha(1+3\alpha)}$ 0.824621
Central planner $s = \frac{1}{2}(1 + 1/\alpha)$	$r\alpha - \sqrt{2\alpha(1+\alpha)}$ 6.30294	$\sqrt{\alpha(1+\alpha)}/2$ 0.848528

Open problem: how do these results generalize? How much information is needed to determine the optimal s ?

Conclusions

- ▶ In most realistic resource allocation problems there is private information to participants.

Conclusions

- ▶ In most realistic resource allocation problems there is private information to participants.
- ▶ Resource allocation policies need to take account of need to give right incentives. To encourage agents who value the resource more to say so, and so be willing to contribute more towards the cost, we need to reward them better than an internal market would do. But figuring out exactly how to do this is not a simple task!

Conclusions

- ▶ In most realistic resource allocation problems there is private information to participants.
- ▶ Resource allocation policies need to take account of need to give right incentives. To encourage agents who value the resource more to say so, and so be willing to contribute more towards the cost, we need to reward them better than an internal market would do. But figuring out exactly how to do this is not a simple task!
- ▶ Simple-minded sharing policies (like proportional sharing) may not to produce sufficient incentives for participants to contribute resources.

Conclusions

- ▶ In most realistic resource allocation problems there is private information to participants.
- ▶ Resource allocation policies need to take account of need to give right incentives. To encourage agents who value the resource more to say so, and so be willing to contribute more towards the cost, we need to reward them better than an internal market would do. But figuring out exactly how to do this is not a simple task!
- ▶ Simple-minded sharing policies (like proportional sharing) may not to produce sufficient incentives for participants to contribute resources.
- ▶ Many new interesting problems!!!

A market solution

One possible approach is to form a market for computation. In this market providers (sellers) and consumers (buyers) of computing resources go to trade.

For instance, an organization might go to the market and say that it needs 10 virtual machines of a certain type for 8 hours and state that the maximum price it is willing to pay is 100 euros. This corresponds to a 'bid' in this market. Similarly, an organization can post in the market its excess computing resources with an 'ask' of the minimum price at which it is willing to sell. The market matches the asks and bids, just as in the stock market, and allocates resources accordingly.

A market solution

One possible approach is to form a market for computation. In this market providers (sellers) and consumers (buyers) of computing resources go to trade.

For instance, an organization might go to the market and say that it needs 10 virtual machines of a certain type for 8 hours and state that the maximum price it is willing to pay is 100 euros. This corresponds to a 'bid' in this market. Similarly, an organization can post in the market its excess computing resources with an 'ask' of the minimum price at which it is willing to sell. The market matches the asks and bids, just as in the stock market, and allocates resources accordingly.

Our approach differs. We provide rules for sharing the resource pool.