Promoting Collaboration in Peer-to-Peer Computational Grids*

Francisco Brasileiro
fubica@dsc.ufcg.edu.br

Universidade Federal de Campina Grande, Brasil
Centro de Engenharia Elétrica e Informática
Departamento de Sistemas e Computação
Laboratório de Sistemas Distribuídos
http://www.lsd.ufcg.edu.br/

*Joint work with Nazareno Andrade, Walfredo Cirne and Miranda Mowbray developed in collaboration with HP Brazil R&D the context of the OurGrid project

Agenda

- Why peer-to-peer computational grids?
- How can incentives for collaboration be provided in such grids?
- How does the Network of Favors work?
- How good is it?
- Practical use of the Network of Favors

Why peer-to-peer computational grids?

e-Science

- Computers are changing scientific research
  - Enabling collaboration
  - As investigation tools
    - Data analysis (eg. data mining)
    - Data generation (eg. simulations)
  - As a result, many research labs around the world are now computation hungry
- Buying more computers is just part of the answer
- Sharing resources though a grid is another

The Virtual Organization

- Most widespread grid architecture
- Entrace is negotiated (by humans)
  - Security issues are handled by conventional A/A/A mechanisms
  - Limits scale
- Flexible and powerful policy enforcement policies
  - Complex middleware (Globus, gLite, etc)
  - Requires skilled support team

Voluntary Computing

- Use large quantities of idle resources in the edges of the Internet
  - Berkley's SETI@home, Stanford's Folding@home
- Able to haverst significant amounts of computational power
  - Open grid for resource contributors
  - Simple instalation of the client software
- Entrance barrier is even higher
  - High visibility project
  - Non-trivial marketing effort
  - Prestigious application provider
  - Skilled upport team to manage the "server"

Peer-to-Peer Grid

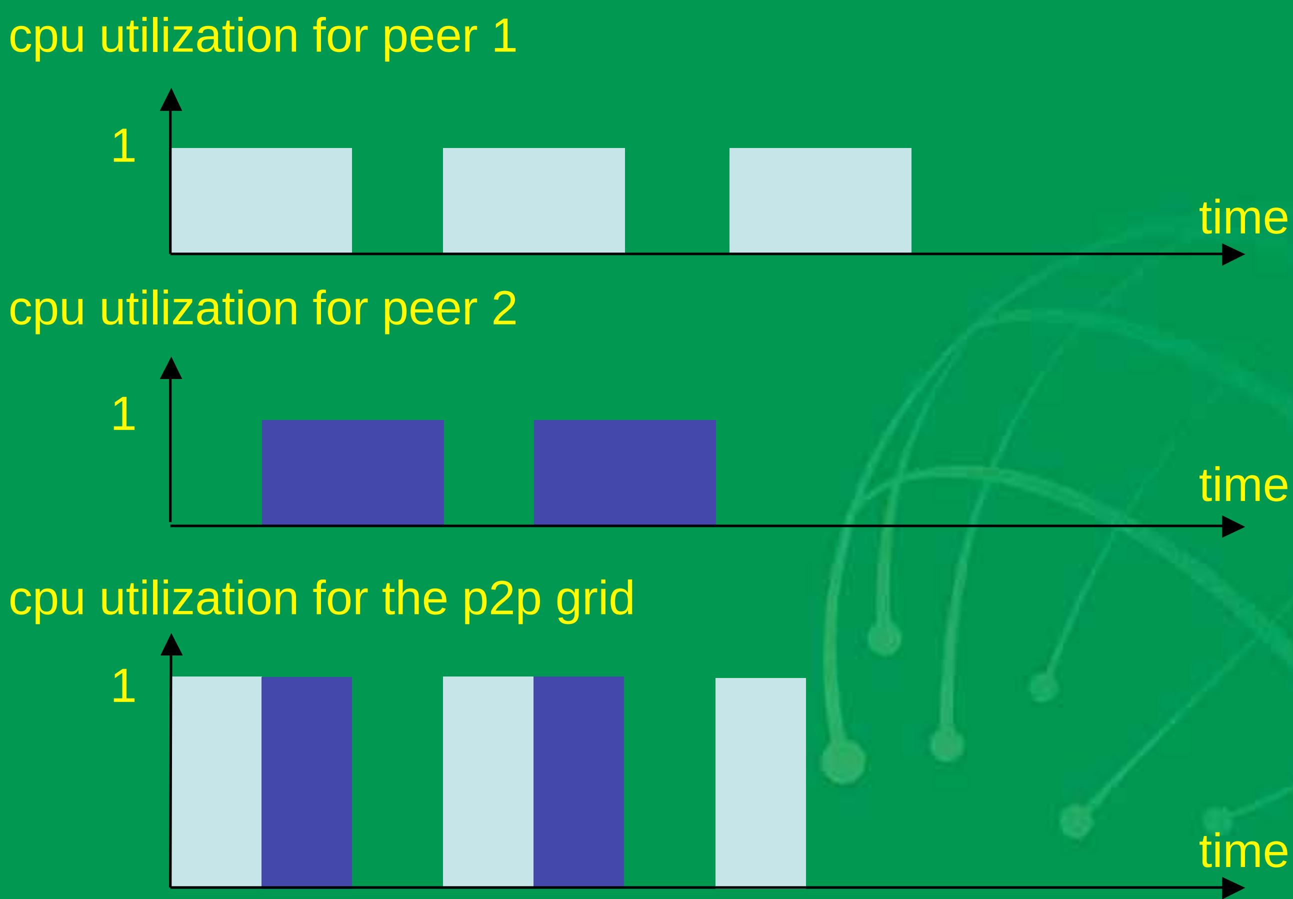

Peer-to-Peer Grids

- Peers join the grid at their will
  - No paperwork
  - No central authority
    - Open grid for resource providers and resource consumers
- Shared deployment and maintenance cost
- Potentially simpler middleware
- This makes it a solution affordable to most users

Peer-to-Peer Grids

- But ...
  - No trust among peers raises new and important security issues
    - Protecting against malicious applications
    - Protection against malicious resources
  - No support for complex sharing policies
  - Free riding severely reduces system efficiency, and may even lead the system to collapse
    - Must provide some incentive for collaboration

How can incentives for collaboration be provided in P2P grids?

Market-based Mechanisms

- Well known mechanisms for regulating access to resources
- Require services provided by trusted institutions
  - Currency distribution
  - Banking
  - Auditing
- Complex to use
  - Price resources provided
  - Plan budget for consuming resources

Reciprocation-based mechanisms

- Reward participantes based on previous behavior
- A reputation system is a way to store information about peers' behavior
- Aggregated opinion
  - Prone to collusions, which is easy if identities are cheap to obtain
  - Rely on specialized secure score management mechanisms

Pair-wise Reciprocation

- Uses only first hand information locally computed from the pair-wise interactions among peers
- It does not work in all settings
  - See “Robust Incentive Techniques for Peer-to-Peer Networks” by Michal Feldman et al.
- It is not successful when the interactions between the same pairs of peers is not frequent enough
  - As it is the case for many P2P file sharing systems
- But it has been quite efficient in a few settings
  - BitTorrent being the most popular system to use it
  - OurGrid, as I will show shortly, is another success case

How does the Network of Favors work?

Basic Functioning

- Assume that any peer can autonomously and accurately value:
  - the amount of work it has received from other peer
  - the amount of work it provides to another peer
- For two peers P and Q, let $S_p(Q)$ be the score of Q in the eyes of P
- Initially $S_p(Q) = 0$ for any P and Q
- If $S_p(Q) = x$ and P provides Q with "favors" of value *v*, then P update $S_p(Q)$ to:
  - $S_p(Q) = \max(x - v, 0)$
- If $S_p(Q) = y$ and Q provides P with "favors" of value *v*, then P update $S_p(Q)$ to:
  - $S_p(Q) = y + v$

Basic Functioning

- Resource allocation is performed as follows
  - Whenever P's idle resources are contended by more than one peer, P allocates them proportionally to the local scores of the requesting peers
    - It works equally well if resources are allocated only to the peer with highest score
  - If only peers with scores equal to zero are contending for P's idle resources, then P shares them among requesters randomly chosen

Important features

- $S_p(Q)$ – the upper bound on the favors that P owns to Q - is an indication of the priority Q has on P's eyes
- The only way Q may increase its priority is by providing favors to P
- Whitewashers gain nothing from creating new identities to interact with the system
- No special bootstrap mechanism is needed
  - Newcomers, free-riders and indebted collaborators are all treated the same

How good is it?

Methodology

- We started with an idealized P2P grid
  - We analyzed in which conditions a perfect reciprocation mechanism could provide incentives for collaboration
- Then, we identified representative scenarios and used simulations to compare the Network of Favors (NoF) against this perfectly informed reciprocation mechanism
- Finally, we run experiments in a controlled grid using the NoF

System model

- We assume a grid comprised of collaborators and free-riders
- At any time t, a peer is either consuming or donating resources from/to the grid
- When donating, collaborators donate all their resources, while free-riders go idle
- Resources are consumed up to the limit of system consumption
  - Excess resources are not used

System model

- Design parameters:
  - C is the maximum amount of favors that a peer can consume from the system
  - Each peer has an independent probability ρ of being in consuming state
  - D is the maximum amount of favors that a peer can donate to the system
  - The utility lost by donating a favor is a constant factor v, $0<v<1$, of the utility gained by the peer that receives the favor
  - N is the total number of peers in the system and f_t is the proportion of free riders at time t

Analysis

- The system may be in three possible states regarding the amount of resources available

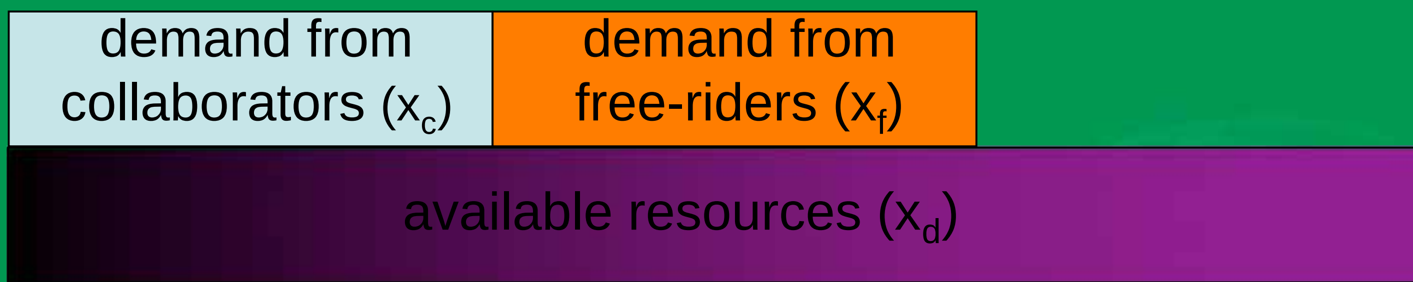

Strong contention ($x_d \leq x_c$)

Weak contention ($x_c < x_d < x_c+x_f$)

No contention ($x_d \geq x_c+x_f$)

- We measure the average advantage to collaborators (AC):
  - AC = Mean utility of collaborators – Mean utility of free-riders
- We say the system works at time t, if there is a disincentive for collaborators to become free riders, ie. AC>0

Analysis

- The system works under strong contention, since free riders utility is zero
- The system does not work under no contention, since collaborators utility increases if they turn into free riding

Analysis

- Under weak contention the advantage to collaborators is:

$$(x_c - v \cdot x_d)/(1 - f_t \cdot N) - (x_d - x_c)/(f_t \cdot N)$$

- The system works if this expression is positive
- We can estimate if the system will work at a time t by determining whether the system will work for the mean values of x_d, x_c and x_f, which can be expressed as:

$$\overline{x_d} = (1-\rho) \cdot \overline{D} \cdot (1 - f_t) \cdot N$$

$$\overline{x_c} = \rho \cdot \overline{C} \cdot (1 - f_t) \cdot N$$

$$\overline{x_f} = \rho \cdot \overline{C} \cdot f_t \cdot N$$

The dynamics of the system

- What if peers change their strategy?

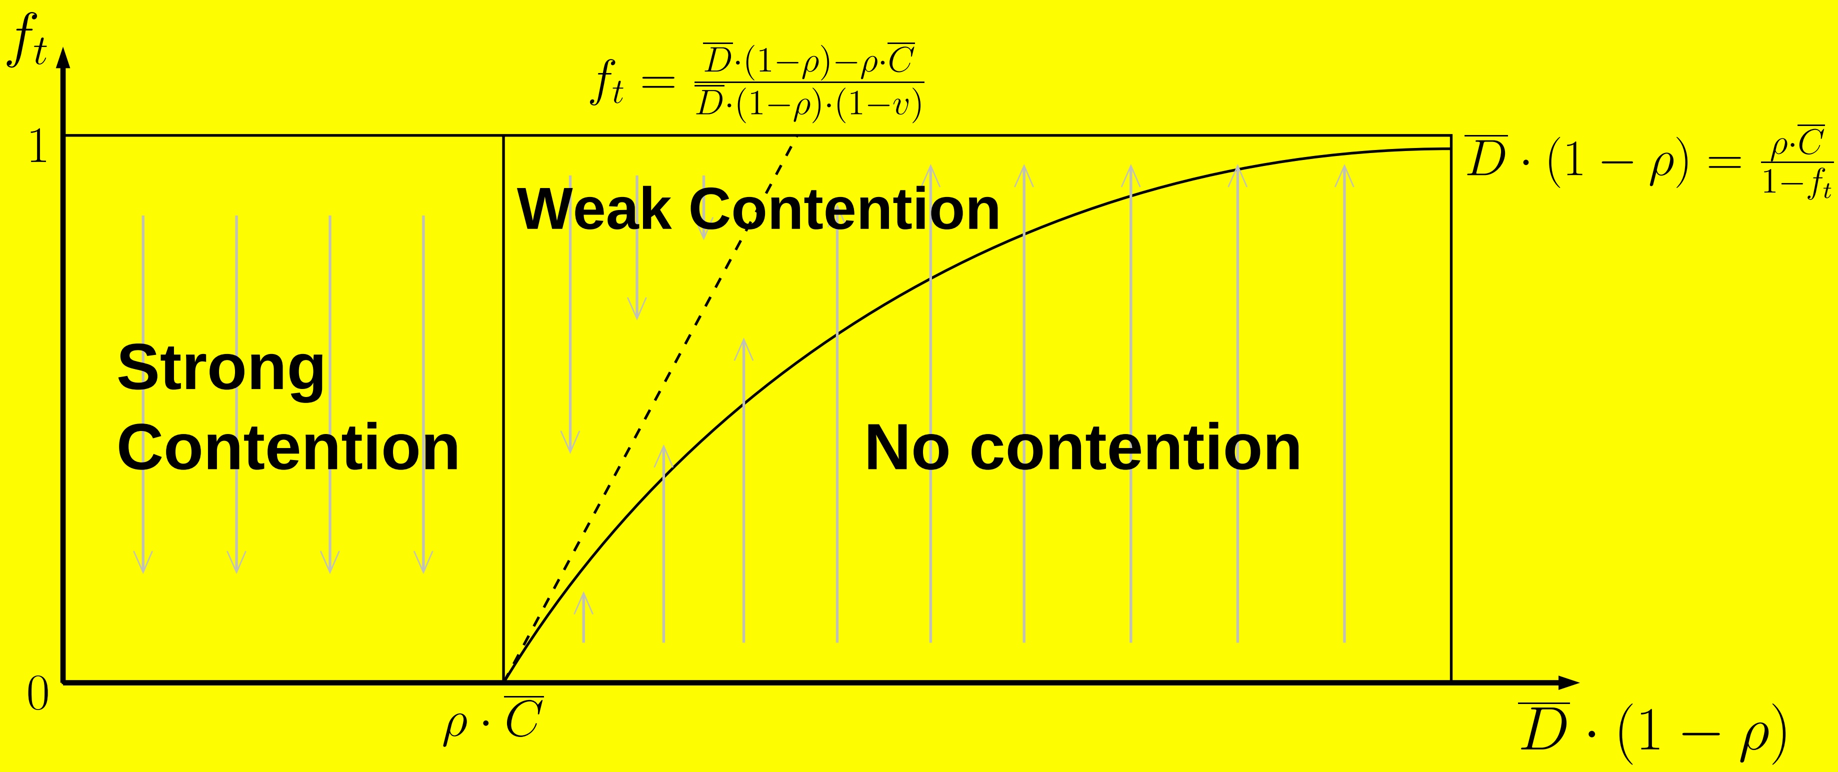

Simulation Scenarios

- 54 scenarios in which:
  - $N = 10{,}000$
  - $D = 10$
  - C is either D/10, D, or 9D
  - ρ is either 0.1, 0.5, or 0.9
  - f_t is either 0.25, 0.5, or 0.75
  - v is either 0.1 or 0.4
- These cover low and high values and include scenarios in the borderline of the different contention scenarios
- The timeline is in turns
  - Each simulation executes 2,000 turns

Simulations Results

- For both incentive mechanisms, the advantage to collaborators was positive for the 36 scenarios in which our analysis had predicted that it would be
  - Interestingly, the performance is worse for a system with less peers (will come to this later)
- For most scenarios there was little difference between the two mechanisms
- The difference was noticeable only for the scenarios in the border from strong to weak contention
  - In these scenarios the NoF was in average 22% worse than a perfectly informed mechanism

Simulation Results

- C=D and v=0.4

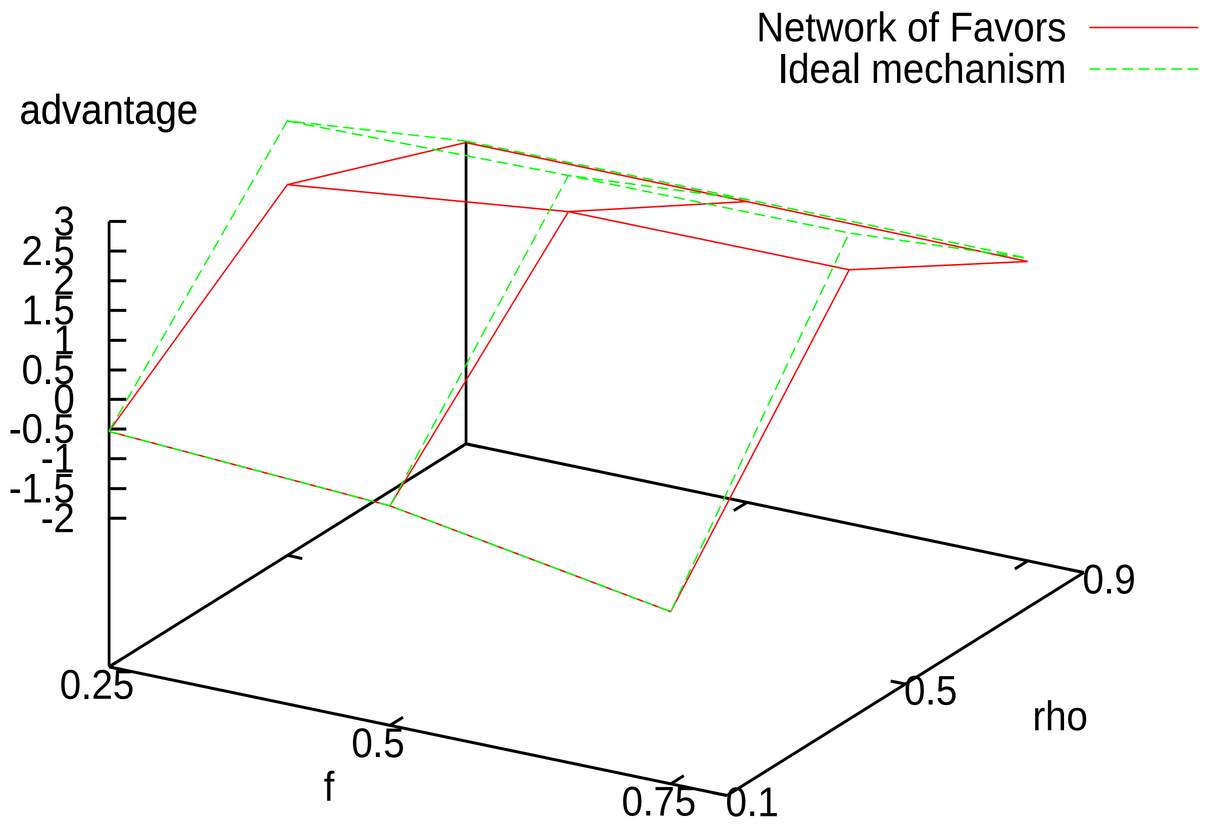

Simulation Results

- C=9D and v=0.1

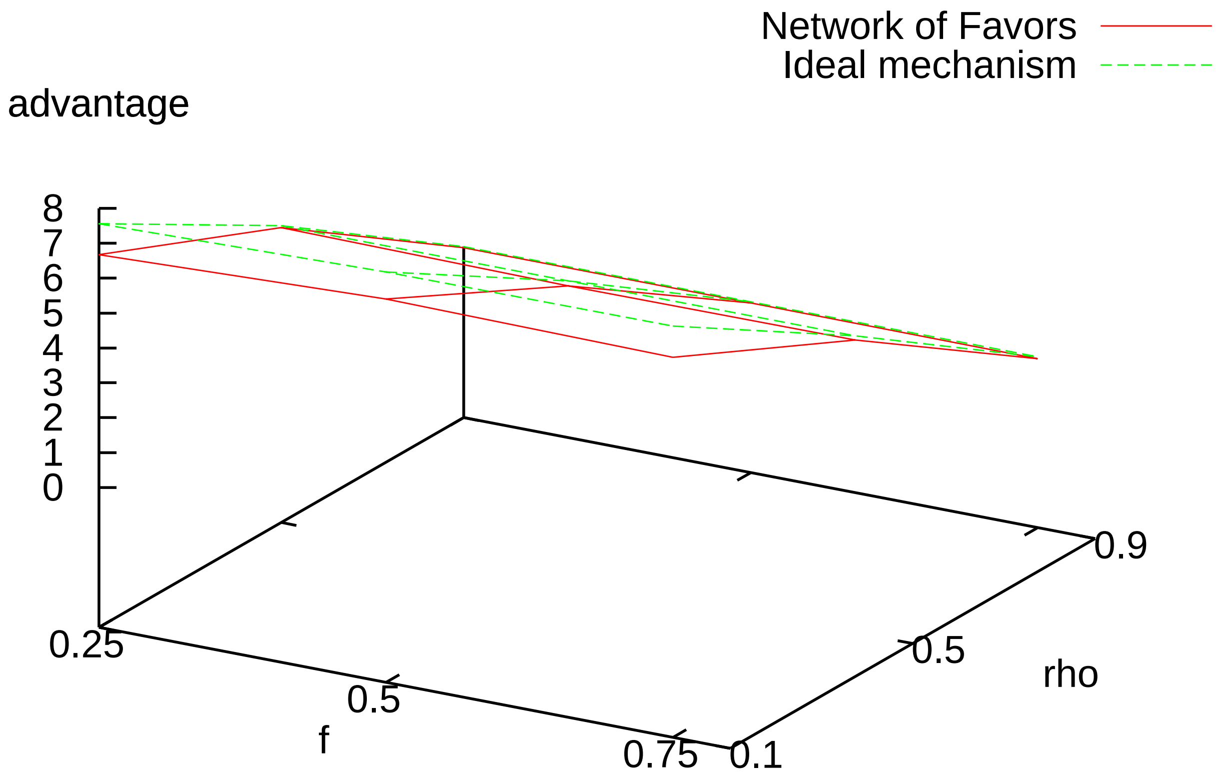

How quick free riders are marginalized?

- New simulations with:
  - 1,000 peers
  - 75% of free riders
- We measured:
  - The proportion of the available resources donated to free riders in the last 50 runs (ε)
  - The relation between the amount of resources consumed and donated by each peer (FR)

How quick free riders are marginalized?

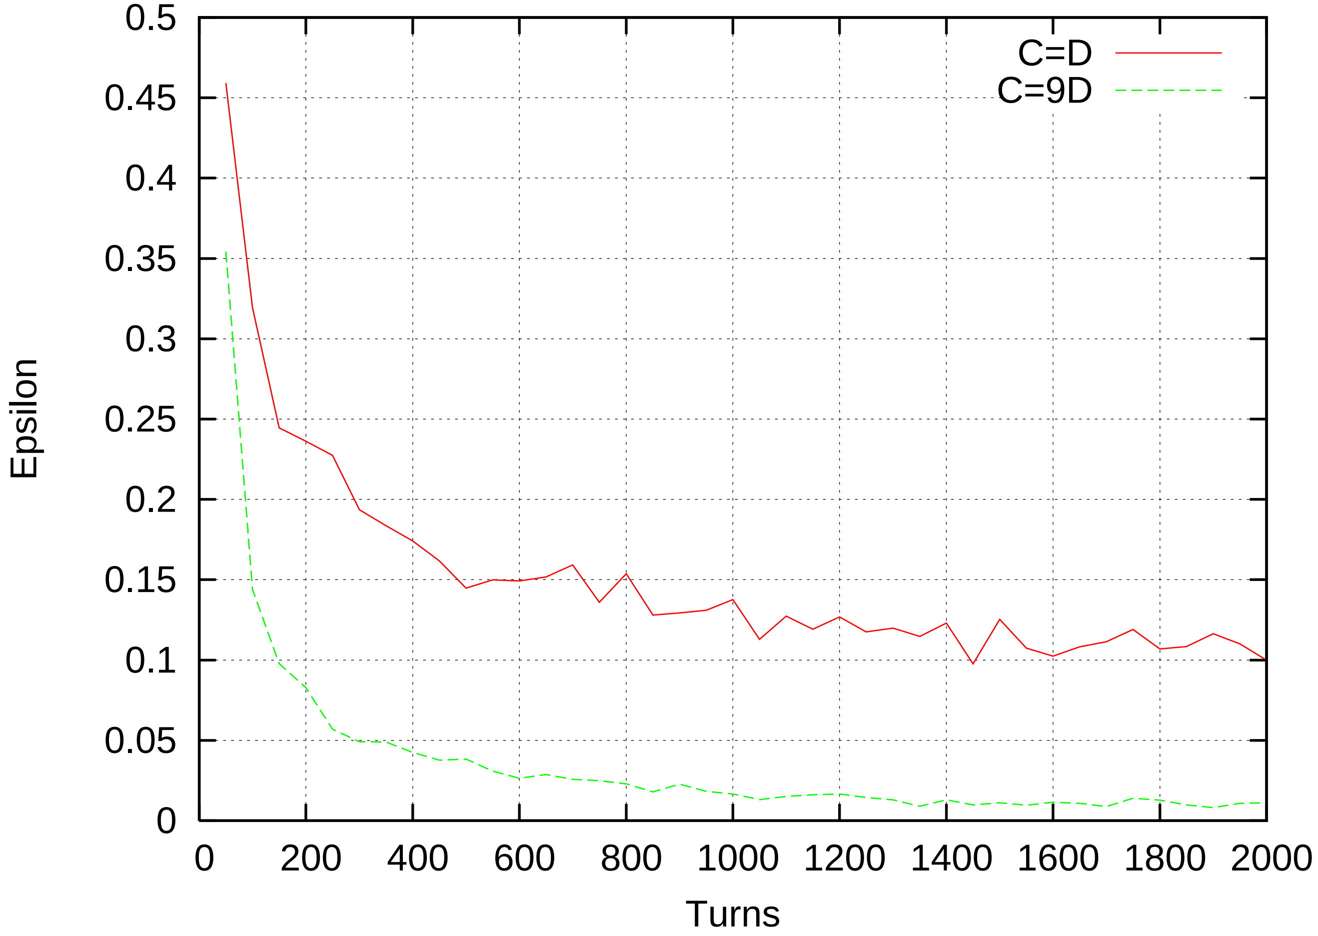

Equity Among Collaborators

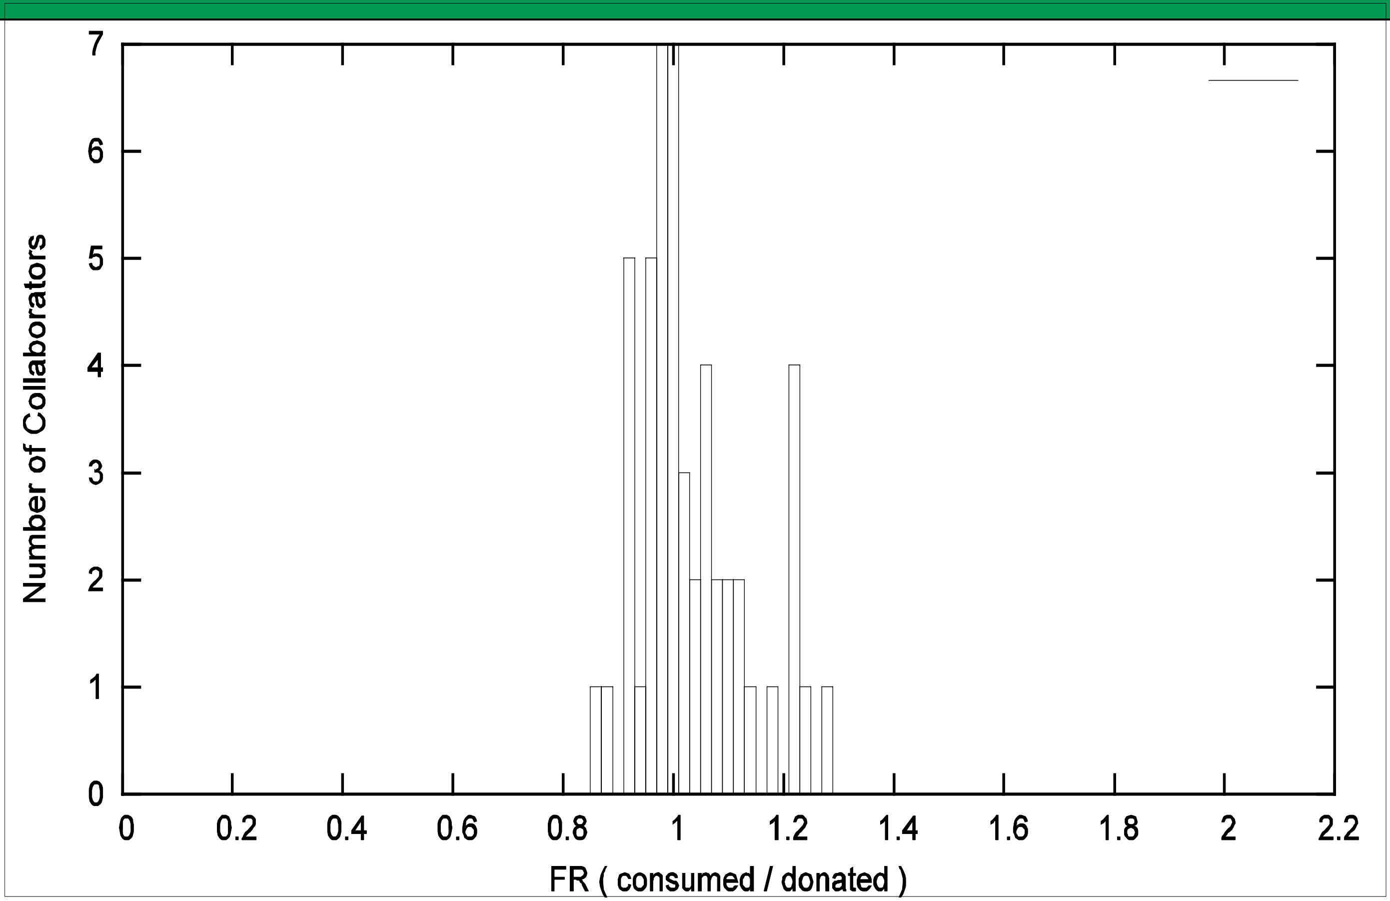

Experiments with a Controlled Grid

- We used a 4-peer grid, each peer with 4 machines
- At each peer jobs arrive with a uniform distribution U(1,20) minutes
- Only one job is scheduled at a time in each peer
  - jobs wait in a queue if other job is already running
- No checkpointing
- Each peer receives 60 jobs of 40 1-minute tasks
- We measure the job makespan
  - If a peer uses only local resources it would complete a job in 10 minutes (disregarding queuing and other overheads)

Experiment Results

- With peers acting in isolation
  - Average makespan was 26.18 minutes
- Peers in a P2P grid
  - Average makespan dropped to 7.41 minutes
- 4-peer grid plus a free rider
  - Average makespan of collaborators was 7.21 minutes (with larger variance when compared to the previous scenario)
  - Average makespan of free rider was 12.15 minutes

How fast free riders are marginalized?

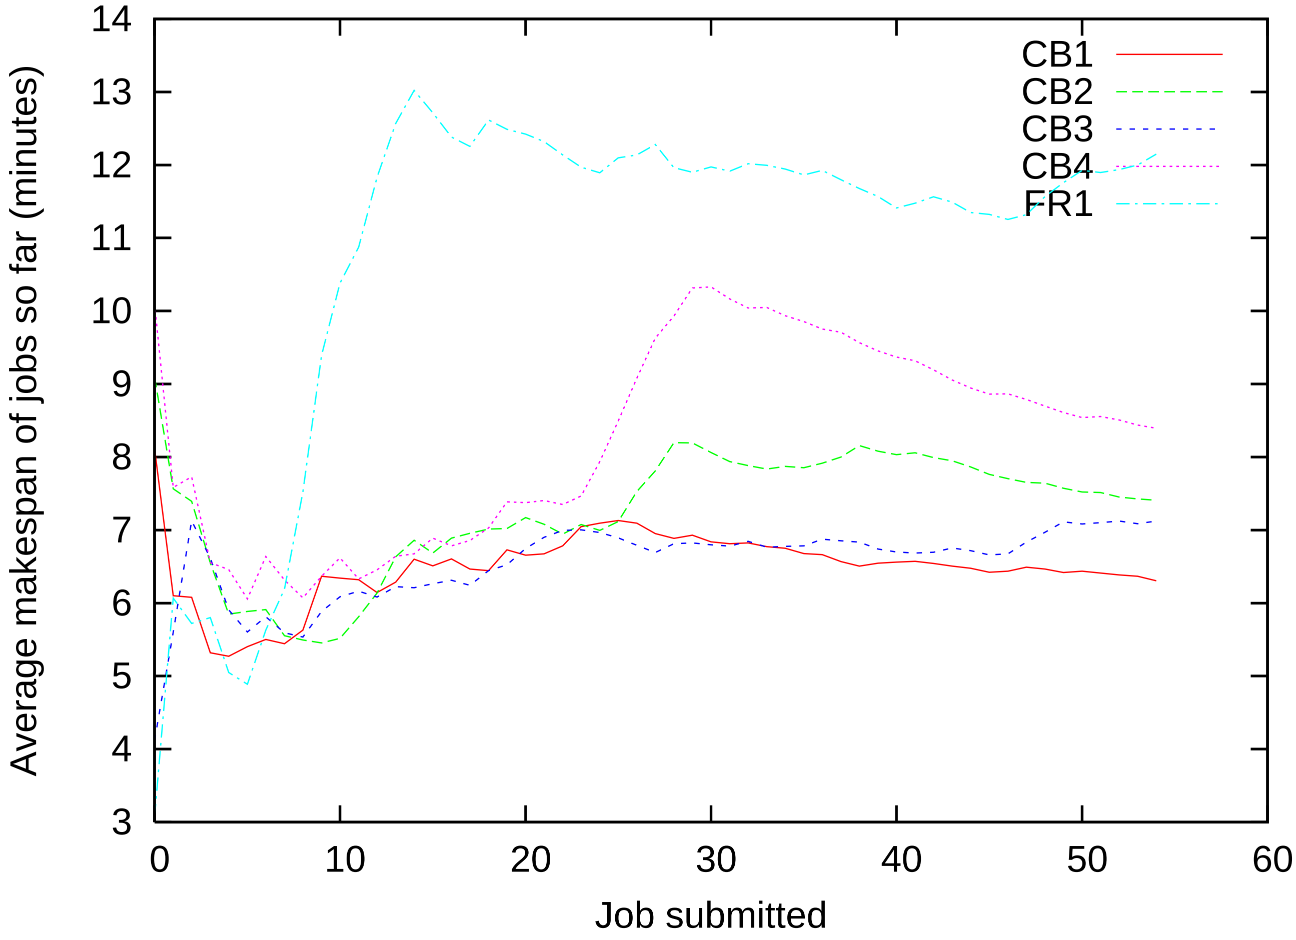

Can this be applied elsewhere?

- Can free riding in file-sharing be prevented with the Network of Favors?
- Feldman and others have shown that for any reciprocation mechanism to work, peers that have interacted once must have a high probability of interacting again
  - High churn and asymmetry of interests rule out the possibility of using the Network of Favors in this setting

How come it works for CPU-sharing?

- Each peer represents a site and has an incentive to be in the system for a long time
- More symmetry of interests
- Many-to-may interactions
- Score function leads to increased interaction
  - After the first interaction between any two collaborators, there will always be one that feels indebted to the other, ie. $S_P(Q)+S_Q(P)>0$

Practical use of the Network of Favors

OurGrid Architecture

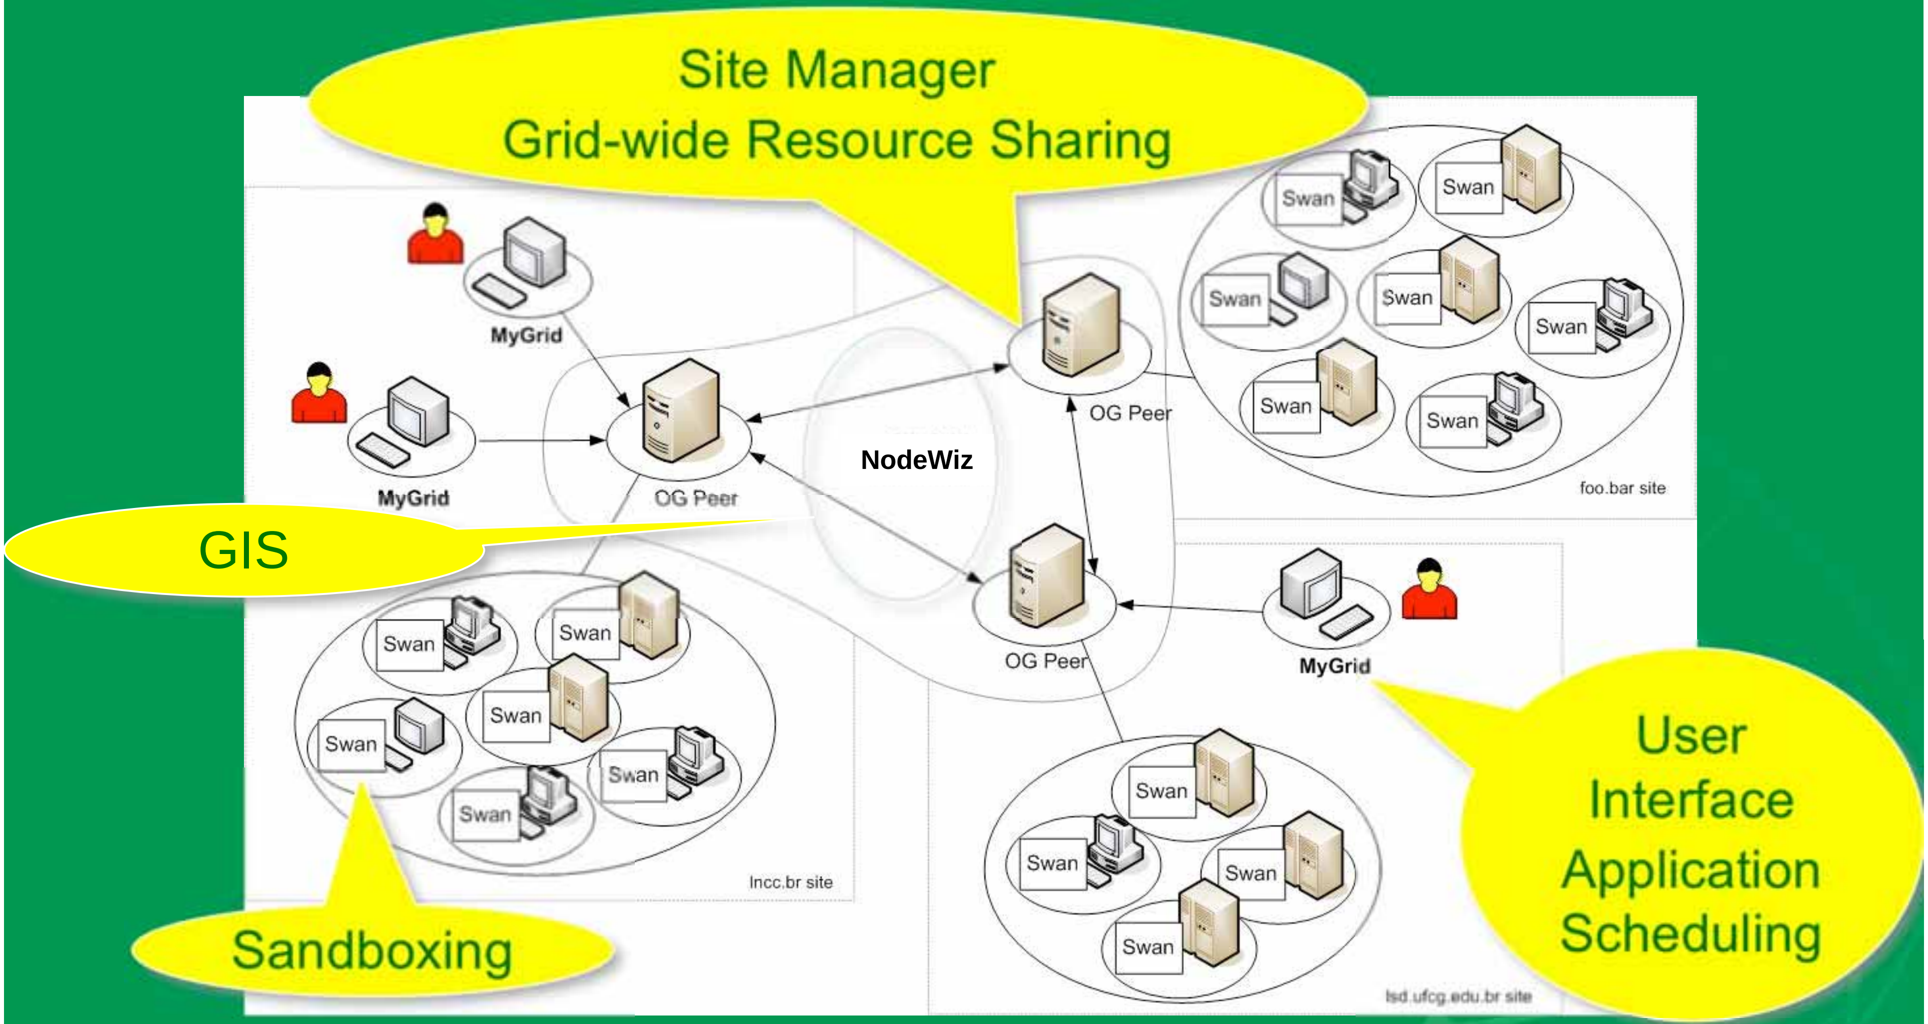

http://status.ourgrid.org/

OurGrid Web Status 3.3

OurGrid STATUS

Statistics	
Server time:	Tue May 22 11:42:44 BRT 2007
Last snapshot time:	Tue May 22 11:42:16 BRT 2007
Peers connected:	18
Consumers connected:	3
Machines on grid:	265
Idle Machines:	20
Machines in use:	17

Online Peers

Peer Name	Version	Local Consumers	Local Machines					Received Machines
			Total	Idle	In use	Donated	Unavailable	
aptech.sstu.ru	3.3.2	0	18	0	0	1	17	0
ciram.epagri.rct-sc.br	3.3.2	0	10	0	0	0	10	0
copad.lsd.ufcg.edu.br	3.3.2	0	10	0	0	0	10	0
cpad.pucrs.br	3.3.2	0	34	0	0	17	17	0
dca.ufcg.edu.br	3.3.2	0	18	7	0	0	11	0
glidlab0.di.unipmn.it	3.3.2	0	6	0	0	6	0	0
hidraulica.hidro.ufcg.edu.br	3.3.2	0	13	5	0	0	8	0
labarc-peer.sytes.net	3.3.2	0	2	1	0	0	1	0
lcc.ufcg.edu.br	3.3.2	0	42	0	0	0	42	0
lmrs-semarh.ufcg.edu.br	3.3.2	0	18	0	0	0	18	0
localhost	3.3.2	0	7	0	0	0	7	0
maspohn.dsc.ufcg.edu.br	3.3.2	0	1	0	0	1	0	0
peer.gmf.ufcg.edu.br	3.3.2	0	13	5	0	0	8	0
peer.lsd.ufcg.edu.br	3.3.2	2	54	2	14	0	38	1
peer.unisantos.br	3.3.2	0	11	0	0	11	0	0
piraiba.gsm.unir.br	3.3.2	1	3	0	3	0	0	0
public.lsd.ufcg.edu.br	3.3.2	0	2	0	0	0	2	0
sergiodbe.sytes.net	3.3.2	0	3	0	0	0	3	0
Totals			**265**	**20**	**17**	**36**	**192**	

Legend:

- Idle machines
- Machines in use by local requests
- Machines in use received from the community
- Machines donated to the community
- Machines that are either off-line or in use by their owners (not idle)

\- Information not available (Old version)

OurGrid Web Status 3.3.1
OurGrid Project

Contact information

- Francisco Brasileiro (fubica@dsc.ufcg.edu.br)
- LSD/UFCG (http://www.lsd.ufcg.edu.br)
- OurGrid project (http://www.ourgrid.org)
- Related projects
  - ShareGrid (http://dcs.di.unipmn.it/)
  - SegHidro (http://seghidro.lsd.ufcg.edu.br/)
  - Bio Pauá (https://www.biopaua.lncc.br/ENGL/index.php)
  - SmartPumping (http://www.sp.lsd.ufcg.edu.br/)
  - GridUnit (http://gridunit.sourceforge.net/)
  - Portal GIGA (http://portalgiga.unisantos.edu.br/)

Does contention arises in practice?