Stressed web environments as strategic games: Risk profiles and Weltanschauung

J. Gabarro [(1)] P. Kilpatrick [(2)] M. Serna [(1)] A. Stewart [(2)]

(1) ALBCOM Research Group
Universitat Politècnica de Catalunya (Barcelona Tech)

(2) School of Computer Science. The Queen's University of Belfast.

{gabarro}@lsi.upc.edu

J. Gabarró and M. Serna are partially supported by
FET pro-active Integrated Project 15964 (AEOLUS), TIN-2007-66523 (FORMALISM), TIN-2005-25859-E and SGR 2009-2015 (ALBCOM).

Summary

- Introduction
- Resource allocation games, KP games
- Risk Profiles and Weltanschauungs
- Angel-daemon games
- Snapshots and anonymous pure Nash equilibria
- On pure Nash equilibria
- Discussion

Introduction

A web-based computation involves the discovery and utilisation of services. It is often the case that a service is made available by a number of providers.

- The performance of a provider can vary over time (SLAs may provide information about "normal" expected performance). Brokers are often used to provide an interface to the "best" current provider.
- It is usually the case that the performance of a provider deteriorates as demand increases (although "elastic" providers may call on extra servers in times of peak demand – thus, in such (stressed) situations performance can conceivably improve).
- The goal of this paper is to study the behaviour of a set of service providers in a stressed environment with the hope that a clearer understanding of stressed behaviour may aid the design of intelligent brokers.
- Here we adopt the point of view that providers should be treated in toto since web users alternate between providers in times of high usage. We also assume that users behave in a non-cooperative way.

Resource Allocation Games

A resource allocation game, also called load balancing game, is a tuple

$$C = \langle N, R, (w_i)_{i \in N}, (d_r)_{r \in R}, (A_i)_{i \in N}, \rangle$$

- The set of players is $N = \{1, \ldots, n\}$.
- Player $i \in N$ has to execute a job (or work) w_i.
- The set of resources is $R = \{1, \ldots, k\}$.
- Each resource $r \in R$ has a delay function d_r.
- For $1 \leq i \leq n$, A_i is the set of possible allocations for work w_i, with $A_i \subseteq R$.

- A strategy s_i for a player i is an element of A_i (player i chooses a resource).
- A strategy profile is $s = (s_1, \ldots, s_n)$.
- Given $s = (s_1, \ldots, s_n)$, the set of players using resource $r \in R$ is $L_r(s) = \{i \in N \mid r = s_i\}$ and the load of resource r is $\ell_r(s) = \sum_{i \in L_r(s)} w_i$.
- The cost for player i is $c_i(s) = d_{s_i}(\ell_{s_i}(s))$.
- The social cost of a strategy profile s is defined additively as $c_{\mathfrak{s}}(s) = \sum_{i \in N} c_i(s)$.

A pure Nash equilibrium, PNE, is a strategy profile s^* such that for any player $i \in N$ and for any strategy $s_i \in A_i$, $c_i(s^*) \leq c_i(s^*_{-i}, s_i)$, where (s^*_{-i}, s_i) denotes the strategy profile in which s^*_i is replaced by s_i.

We are interested in delay functions like $d_r(x) = d_r\, x$ with $0 < d_r < \infty$.

Example: A $Fortran\&MPI\_Servers$

- Servers 1 and 3 offer `Fortran` and `MPI` services.
- Server 2 offers `Fortran` services.
- Server 4 offers `MPI` services.

To get integer values, performance delays are defined $d_r(x) = \lceil d_r x \rceil$ where

$$d_1 = 1/2,\ d_2 = 1/4,\ d_3 = 1/4,\ d_4 = 1/8$$

There are 5 jobs to be executed (superindices denote the type of a job).

$$w_1^{\texttt{MPI}} = 10,\ w_2^{\texttt{F}} = 5,\ w_3^{\texttt{F}} = 6,\ w_4^{\texttt{MPI}} = 15,\ w_5^{\texttt{F}} = 3$$

An initial allocation $s = (1, 1, 1, 4, 3)$ is displayed.

Fortran&MPI_Servers	Resources			
	1	2	3	4
	Services			
	`F`, `MPI`	`F`	`F`, `MPI`	`MPI`
	Delay			
	1/2	1/4	1/4	1/8
Initial strategy				
s	$w_1^{\texttt{MPI}} = 10$ $w_2^{\texttt{F}} = 5$ $w_3^{\texttt{F}} = 6$		$w_5^{\texttt{F}} = 3$	$w_4^{\texttt{MPI}} = 15$
Cost	11		1	2

Moves	Resources			
	1 F, MPI	2 F	3 F, MPI	4 MPI
$1 \xrightarrow{w_1^{\texttt{MPI}}} 3$	$w_2^{\texttt{F}} = 5$ $w_3^{\texttt{F}} = 6$		$w_1^{\texttt{MPI}} = 10$ $w_5^{\texttt{F}} = 3$	$w_4^{\texttt{MPI}} = 15$
Cost	6		4	2
$1 \xrightarrow{w_3^{\texttt{F}}} 3$	$w_2^{\texttt{F}} = 5$		$w_1^{\texttt{MPI}} = 10$ $w_3^{\texttt{F}} = 6$ $w_5^{\texttt{F}} = 3$	$w_4^{\texttt{MPI}} = 15$
Cost	3		5	2

Moves	Resources			
	1 F, MPI	2 F	3 F, MPI	4 MPI
$3 \xrightarrow{w_5^{\mathtt{F}}} 1$	$w_2^{\mathtt{F}} = 5$ $w_5^{\mathtt{F}} = 3$		$w_1^{\mathtt{MPI}} = 10$ $w_3^{\mathtt{F}} = 6$	$w_4^{\mathtt{MPI}} = 15$
Cost	4		4	2
$1 \xrightarrow{w_2^{\mathtt{F}}} 2$	$w_5^{\mathtt{F}} = 3$	$w_2^{\mathtt{F}} = 5$	$w_1^{\mathtt{MPI}} = 10$ $w_3^{\mathtt{F}} = 6$	$w_4^{\mathtt{MPI}} = 15$
Cost	2	2	4	2
$3 \xrightarrow{w_3^{\mathtt{F}}} 2$	$w_5^{\mathtt{F}} = 3$	$w_2^{\mathtt{F}} = 5$ $w_3^{\mathtt{F}} = 6$	$w_1^{\mathtt{MPI}} = 10$	$w_4^{\mathtt{MPI}} = 15$
Cost	2	3	3	2

KP games

In some cases (*uniform*) all the resources have the same capability,

- this forces $A_i = R$ for $1 \leq i \leq n$ and
- all the resources have the same affine delay function $d(x) = dx$.

We distinguish the case where we have $w_i = 1$ for all $i \in N$.

Let $KP_{n,k}$ denote the unit resource allocation game with n players and k resources with *uniform affine delays*.

Risk Profiles

Risk profiles are used to describe stressed environments in which two competing forces act on resources.

Assume that R is partitioned into $\mathcal{A}$ and $\mathcal{D}$ such that $R = \mathcal{A} \cup \mathcal{D}$ and $\mathcal{A} \cap \mathcal{D} = \emptyset$.

- Subset $\mathcal{A}$ is controlled by an agent $\mathfrak{a}$ called the angel. When r is selected by $\mathfrak{a}$ it runs under the angelic delay function $d_r^{\mathcal{A}}$.

  The angel forces angelic behaviour on $f_{\mathcal{A}}$ resources.

  The angel's objective is to improve system as much as possible.

- Subset $\mathcal{D}$ is controlled by another agent $\mathfrak{d}$ called the daemon. When r is selected by $\mathfrak{d}$ under the daemonic delay function $d_r^{\mathcal{D}}$.

  Daemon can affect only $f_{\mathcal{D}}$ resources.

  The daemon's objective is to worse the system as much as possible.

Definition. A risk profile for C is a tuple $\mathcal{R} = \langle C, \mathcal{A}, \mathcal{D}, f_{\mathcal{A}}, f_{\mathcal{D}}, (d_r^{\mathcal{A}})_{r \in \mathcal{A}}, (d_r^{\mathcal{D}})_{r \in \mathcal{D}} \rangle$.

Weltanschauungs

Some instances of angel and daemon delay functions are considered and the resulting situations (world views) are analysed using the notion of Weltanschauung.

- The angel and daemon have two sensitivities with respect to the environment: an *extreme*, $\mathbb{E}$ or *moderate*, $\mathbb{M}$.

  Combined sensitivities are $\mathbb{S} = \{\mathbb{E}\text{-}\mathbb{E}, \mathbb{E}\text{-}\mathbb{M}, \mathbb{M}\text{-}\mathbb{E}, \mathbb{M}\text{-}\mathbb{M}\}$.
- Both the angel and the daemon have a joint psychological view of the environment called here the moral.

  The set of morals is: $\mathsf{M} = \{\mathsf{Crash}, \mathsf{Benevolent}, \mathsf{Polarized}, \mathsf{Schizophrenic}\}$

The set of Weltanschauungs is $\mathfrak{W} = \mathbb{S} \times \mathsf{M}$.

Depending on the sensitivity and moral, the angel and the daemon stress the delay functions of a resource in the following way

	Crash	Benevolent	Polarized	Schizophrenic
$\mathbb{E}$-$\mathbb{E}$	$d_r^{\mathcal{A}} = \infty$ $d_r^{\mathcal{D}} = \infty$	$d_r^{\mathcal{A}} = 0$ $d_r^{\mathcal{D}} = 0$	$d_r^{\mathcal{A}} = 0$ $d_r^{\mathcal{D}} = \infty$	$d_r^{\mathcal{A}} = \infty$ $d_r^{\mathcal{D}} = 0$
$\mathbb{E}$-$\mathbb{M}$	$d_r^{\mathcal{A}} = \infty$ $d_r^{\mathcal{D}} = \beta d_r$	$d_r^{\mathcal{A}} = 0$ $d_r^{\mathcal{D}} = \beta d_r$	$d_r^{\mathcal{A}} = 0$ $d_r^{\mathcal{D}} = \beta d_r$	$d_r^{\mathcal{A}} = \infty$ $d_r^{\mathcal{D}} = \beta d_r$
$\mathbb{M}$-$\mathbb{E}$	$d_r^{\mathcal{A}} = \alpha d_r$ $d_r^{\mathcal{D}} = \infty$	$d_r^{\mathcal{A}} = \alpha d_r$ $d_r^{\mathcal{D}} = 0$	$d_r^{\mathcal{A}} = \alpha d_r$ $d_r^{\mathcal{D}} = \infty$	$d_r^{\mathcal{A}} = \alpha d_r$ $d_r^{\mathcal{D}} = 0$
$\mathbb{M}$-$\mathbb{M}$	$d_r^{\mathcal{A}} = \alpha d_r$ $d_r^{\mathcal{D}} = \beta d_r$	$d_r^{\mathcal{A}} = \alpha d_r$ $d_r^{\mathcal{D}} = \beta d_r$	$d_r^{\mathcal{A}} = \alpha d_r$ $d_r^{\mathcal{D}} = \beta d_r$	$d_r^{\mathcal{A}} = \alpha d_r$ $d_r^{\mathcal{D}} = \beta d_r$
	$1 < \alpha < \infty$ $1 < \beta < \infty$	$0 < \alpha < 1$ $0 < \beta < 1$	$0 < \alpha < 1$ $1 < \beta < \infty$	$1 < \alpha < \infty$ $0 < \beta < 1$

Definition. Let A risk profile is a tuple $\mathcal{R} = \langle C, \mathcal{A}, \mathcal{D}, f_{\mathcal{A}}, f_{\mathcal{D}}, \mathfrak{w} \rangle$.

Example: *Stop&Slow*

Consider *Fortran&MPI_Servers*.

Suppose $\mathfrak{a}$ controls 1 and 2 (i.e. $\mathcal{A} = \{1, 2\}$) and $\mathfrak{d}$ controls 3 and 4 ($\mathcal{D} = \{3, 4\}$).

Both, $\mathfrak{a}$ and $\mathfrak{d}$ have limited capacity, $f_{\mathcal{A}} = f_{\mathcal{D}} = 1$.

Consider the following scenario:

- The angel $\mathfrak{a}$ controls a number of servers: assume that one of these servers fails when put under stress i.e. $d_r^{\mathcal{A}} = \infty$.
- The daemon $\mathfrak{d}$ controls a number of robust servers. However, one server's performance is degraded under stress and so $d_r^{\mathcal{D}} = \beta d_r$ where $\beta = 3/2$.

This is a Crash Weltanschauung $\mathfrak{w} = (\infty, 3/2)$ of type $\mathbb{E}$-$\mathbb{M}$.

$\mathit{Stop\&Slow} = \langle \mathit{Fortran\&MPI\_Servers}, \{1, 2\}, \{3, 4\}, 1, 1, (\infty, 3/2)\rangle$.

Angel-daemon games

Definition. Given a risk profile $\mathcal{R} = \langle C, \mathcal{A}, \mathcal{D}, f_{\mathcal{A}}, f_{\mathcal{D}}, \mathfrak{w} \rangle$, the angel-daemon game associated to $\mathcal{R}$ is

$$\Gamma(\mathcal{R}) = \langle N \cup \{\mathfrak{a}, \mathfrak{d}\}, (A_i)_{i \in N}, A_{\mathfrak{a}}, A_{\mathfrak{d}}, (c_i)_{i \in N}, c_{\mathfrak{a}}, c_{\mathfrak{d}} \rangle$$

where $A_{\mathfrak{a}} = \{a \subseteq \mathcal{A} \mid |a| = f_{\mathcal{A}}\}$ and $A_{\mathfrak{d}} = \{b \subseteq \mathcal{D} \mid |b| = f_{\mathcal{D}}\}$.

Given $(a, d) \in A_{\mathfrak{a}} \times A_{\mathfrak{d}}$ the cost functions of a resource r is defined as follows.

$$d_r[a, d] = \begin{cases} d_r^{\mathcal{A}} & \text{if } r \in a. \\ d_r^{\mathcal{D}} & \text{if } r \in d. \\ d_r & \text{if } r \notin (a \cup d). \end{cases}$$

Given a strategy profile $\sigma = (s, a, d)$,

- player i incurs a cost $c_i(\sigma) = d_{s_i}[a, d](\ell_{s_i}(s))$,
- the angel cost is $c_{\mathfrak{a}}(\sigma) = \sum_{i \in N} c_i(\sigma)$ and
- the daemon cost is $c_{\mathfrak{d}}(\sigma) = -c_{\mathfrak{a}}(\sigma)$.

Example: $\Gamma(\mathit{Stop\&Slow})$

Consider game $\Gamma(\mathit{Stop\&Slow})$

The set of players is $N \cup \{\mathfrak{a}, \mathfrak{d}\}$, $f_{\mathcal{A}} = f_{\mathcal{D}} = 1$, $A_{\mathfrak{a}} = \{\{1\}, \{2\}\}$, $A_{\mathfrak{d}} = \{\{3\}, \{4\}\}$.

Suppose that jobs are allocated using the schedule $s = (3, 2, 3, 4, 1)$.

The angel $\mathfrak{a}$ closes site 1 ($a = \{1\}$) while the daemon $\mathfrak{d}$ chooses to deteriorate the performance of site 3 ($d = \{3\}$).

The stressed delay functions $d_r[a, d]$ are:

$$d_1[\{1\}, \{3\}] = \infty, \;\; d_2[\{1\}, \{3\}] = d_2 = 1/4,$$
$$d_3[\{1\}, \{3\}] = \beta d_3 = 3/4 \times 1/4, \;\; d_4[\{1\}, \{3\}] = d_4 = 1/8$$

Profile $((3, 2, 3, 4, 1), \{1\}, \{3\})$ is not a PNE because $w_5^{\mathtt{F}}$ can improve his situation.

Profile $((3, 2, 2, 4, 2), \{1\}, \{3\})$ is a PNE.

Snapshots and anonymous pure Nash equilibria

When we need to make explicit the $\mathfrak{w}$ we replace $d_r[a, d]$ with $d_r[a, d, \mathfrak{w}]$.

Definition. Let $\Gamma(\mathcal{R})$ be an angel-daemon resource allocation game, $\mathcal{R} = \langle C, \mathcal{A}, \mathcal{D}, f_{\mathcal{A}}, f_{\mathcal{D}}, \mathfrak{w}\rangle$, with profile $\sigma = (s, a, d)$. For a resource $r \in R$ define two properties:

$$\mathsf{affiliation} = \begin{cases} \mathfrak{a} \text{ if } r \in \mathcal{A} \\ \mathfrak{d} \text{ otherwise} \end{cases} \qquad \mathsf{selected} = \begin{cases} \mathsf{y} \text{ if } r \in a \cup d \\ \mathsf{n} \text{ otherwise} \end{cases}$$

Property affiliation indicates whether r is controlled by the angel or the daemon. Property selected denotes whether r's behaviour is abnormal or not (i.e. whether r has been chosen by either angel or daemon). A snapshot is a tuple $\delta(\sigma, \mathfrak{w}) = (\delta_1(\sigma, \mathfrak{w}) \mid \cdots \mid \delta_k(\sigma, \mathfrak{w}))$ which provides information about the current state of each of the resources where $\delta_r(\sigma, \mathfrak{w}) = (\mathsf{affiliation}_{\mathsf{selected}}, \ell_r(s), d_r[a, d, \mathfrak{w}])$

Snapshots provide a clear picture of

- the occupancy of resources,
- the strategies of $\mathfrak{a}$ and $\mathfrak{d}$, and
- the delay functions applicable to each resource.

Example. Consider a $KP_{2,4}$ with delay functions $d(x) = dx$ for $r \in R$. Suppose $\mathcal{A} = \{1, 2\}$, $\mathcal{D} = \{3, 4\}$ and $f_{\mathcal{A}} = f_{\mathcal{D}} = 1$. Given $\sigma = s, a, d) = ((2, 4), \{1\}, \{4\})$ and $\mathfrak{w} = (\infty, \infty)$, the snapshot is

$$\delta(\sigma, (\infty, \infty)) = (\mathfrak{a}_{\mathsf{y}}, 0, \infty \mid \mathfrak{a}_{\mathsf{n}}, 1, d(x) \mid \mathfrak{d}_{\mathsf{n}}, 0, d(x) \mid \mathfrak{d}_{\mathsf{y}}, 1, \infty)$$

Lemma 1 *Let* $\mathcal{R} = \langle C, \mathcal{A}, \mathcal{D}, f_{\mathcal{A}}, f_{\mathcal{D}}, \mathfrak{w}\rangle$, *be a risk profile associated with a unit resource allocation game* C *and Weltanschauung* $\mathfrak{w}$ *and let* $\sigma = (s, a, d)$, $\sigma' = (s', a', d')$ *be two strategy profiles of the angel-daemon game* $\Gamma(\mathcal{R})$. *Then* $\delta(\sigma, \mathfrak{w}) = \delta(\sigma', \mathfrak{w})$ *iff* $a = a'$ *and* $d = d'$ *and there is a permutation* $\pi : N \to N$ *such that* $\pi(L_r(s)) = L_r(s')$ *for any* $r \in R$.

Lemma 2 *Given* $\Gamma(\mathcal{R})$ *for* $\mathcal{R} = \langle C, \mathcal{A}, \mathcal{D}, f_{\mathcal{A}}, f_{\mathcal{D}}, \mathfrak{w}\rangle$, *and* $\sigma = (s, a, d)$, $\sigma' = (s', a, d)$ *such that* $\delta(\sigma, \mathfrak{w}) = \delta(\sigma', \mathfrak{w})$ *then* σ *is a PNE iff* σ' *is a PNE.*

Definition. A snapshot δ is called an anonymous Nash equilibrium iff there exists a PNE σ such that $\delta(\sigma) = \delta$.

Lemma 3 *For any* $\mathfrak{w} \in \mathfrak{W}$, *the game* $\Gamma(\mathcal{R}_{\mathfrak{w}})$ *corresponding to the profile* $\mathcal{R}_{\mathfrak{w}} = \langle KP_{1,2}, \emptyset, \{1, 2\}, 0, 1, \mathfrak{w}\rangle$ *has no pure Nash equilibria.*

On pure Nash equilibria

Theorem 4 *There are tuples* $\langle KP_{n,k}, \mathcal{A}, \mathcal{D}, f_{\mathcal{A}}, f_{\mathcal{D}} \rangle$ *such that for any* $\mathfrak{w} \in \mathcal{W}$*, the profile* $\mathcal{R}_{\mathfrak{w}} = \langle KP_{n,k}, \mathcal{A}, \mathcal{D}, f_{\mathcal{A}}, f_{\mathcal{D}}, \mathfrak{w} \rangle$ *describes a game* $\Gamma(\mathcal{R}_{\mathfrak{w}})$ *with no pure Nash equilibria.*

Theorem 5 $\mathcal{R}_{\mathfrak{w}} = \langle KP_{n,k}, \mathcal{A}, \mathcal{D}, f_{\mathcal{A}}, f_{\mathcal{D}}, \mathfrak{w} \rangle$ *such that* $f_{\mathcal{A}} + k \leq \#\mathcal{A}$*, then* $\Gamma(\mathcal{R}_{\mathfrak{w}})$ *has always a pure Nash equilibrium.*

Discussion

Conventionally the behaviour of a web service is captured by treating the service in isolation. A service level agreement (SLA) might provide information about the expected behaviour of a service.

In this paper an alternative view of web services is presented; here the behaviour of a set of services within a stressed web environment is modelled by a strategic angel daemon game.

Two different abstractions are presented.

- Risk profiles partition web services into angel and daemon sets.
- Weltanschauungs consider the various scenarios that can arise when angels and daemons have abnormal delay functions.

The use risk profiles and Weltanschauungs raises a number of questions:

1. How can a set of services be partitioned into $\mathfrak{a}$ and $\mathfrak{d}$ controlled sets?
2. How can an service's performance improve under stress?
3. How can abnormal bounds for the angel and the daemon be set?

We may interpret the angel/daemon partition is by means of a cost model: sets of low cost services are liable to be severely affected by stressed environments – such sets may be controled by $\mathfrak{d}$. Expensive services may be responsive even when the surrounding environment becomes stressed.

Some services may be implemented on elastic clouds – as demand increases a service may call upon more servers to facilitate ongoing requests. In stressed conditions the behavior may even be better.

It seems likely that the choice of the abnormal bound parameters would have to be made on the basis of experimental evidence.

Thanks!