

How to measure efficiency?

Arnaud Legrand and Corinne Touati
CNRS-INRIA, LIG laboratory, MESCAL project
Grenoble, France
arnaud.legrand@imag.fr, corinne.touati@imag.fr

ABSTRACT

In the context of applied game theory in networking environments, a number of concepts have been proposed to measure both efficiency and optimality of resource allocations, the most famous certainly being the price of anarchy and the Jain index. Yet, very few have tried to question these measures and compare one to another, in a general framework, which is the aim of the present article.

Keywords

Game theory, Nash equilibrium, Pareto optimality, Braess-like paradox, Price of anarchy.

1. INTRODUCTION

The networking community has witnessed an impressive amount of work based on applications of game theory concepts. This paper focuses here on the ones dealing with characterizations of performance of general policies.

We do not deal here with the choice of users utility functions. We consider some general utilities u , may they represent throughput, experienced delays... or any utility function and study in this paper different allocation policies. We distinguish in particular two kinds of policies:

- those who are index-based, that is to say that result on the optimization of a given function, as for example the Nash Bargaining Solution (also called proportional fairness), that maximizes the product of the users' utilities, or the social utility (maximizing their sum).
- general policy optimization. Those do not optimize a specific function. The most common example being the Nash equilibrium.

While many definitions of efficiency measure can be found in the literature, at the present day, it seems that no fully satisfactory concept is available. The goal of this article is to present and study various commonly used characterizations of the performance of policies.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

GameComm '07, October 22, 2007, Nantes, France
Copyright 2007 ICST 978-963-9799-00-4.

After introducing some general notations (Section 2), we present some qualitative characterization of the allocations (Section 3): in particular the notion of Pareto efficiency (a general notion of efficiency), of index-optimization (that would reflect some particular property of efficient points, as for example fairness), and Braess-like paradoxes (a particularly non-desirable property of allocation policies). We then analyze properties of allocations, in particular regarding continuity (which ensures some stability of the allocation for slight changes of the resources) and monotonicity (which ensures that an adding of resources will always be beneficial to the users).

Then, in Section 4, we consider quantitative measures of efficiency. In particular, we discuss the concepts of Jain index, Price of anarchy (and more generally of index-optimizing based metrics) and the recently introduced SDF (Selfishness Degradation Factor).

2. NOTATIONS

We consider an n -player game, each of them having a utility function whose values belong to \mathbb{R}_+ . A utility set U is thus a subset of \mathbb{R}_+^n . Let $\mathcal{H}(\mathbb{R}_+^n)$ denote the set of non-empty compact sets of \mathbb{R}_+^n and $\mathcal{C}(\mathbb{R}_+^n)$ denote the set of non-empty compact and convex sets of \mathbb{R}_+^n . In the rest of this article, we assume that \mathcal{U} the set of all utility sets is either equal to $\mathcal{H}(\mathbb{R}_+^n)$ or $\mathcal{C}(\mathbb{R}_+^n)$. Any negative result regarding $\mathcal{C}(\mathbb{R}_+^n)$ also applies to $\mathcal{H}(\mathbb{R}_+^n)$.

We define in this section the two kinds of allocation studied (index-based or not) and two concepts that will turn useful for the analytical study, namely the Hausdorff metric and some canonical partial orders.

Definition 1 (Policy function). *A policy function $\alpha : \mathcal{U} \rightarrow \mathbb{R}_+^n$ is a function such that for all $U \in \mathcal{U}$, $\alpha(U) \in U$.*

Policy functions defined on $\mathcal{H}(\mathbb{R}_+^n)$ are said to be general policy function and policy functions defined on $\mathcal{C}(\mathbb{R}_+^n)$ are said to be convex policy function.

Note that in this framework, we do not consider policy optimization that depend on previous states of the system. Such systems can occur for instance when considering dynamic systems where Nash equilibria adjusts to the system evolution. In the event of multiple equilibria, the initial conditions have an impact on the convergence point.

Definition 2 (Index-optimizing). *An index function f is a function from \mathbb{R}_+^n to \mathbb{R}_+ . A policy function α is said to be f -optimizing if for all $U \in \mathcal{U}$, $f(\alpha(U)) = \sup_{u \in U} f(u)$.*

Index may also be called aggregation operators [2].

To study the continuity of policy functions, we need a topology on \mathcal{U} . That is why in the following, we use the classical metric on compact sets.

Definition 3 (Hausdorff metric). *Considering a metric function d on \mathbb{R}_+^n , one can define the distance from x to the compact B as:*

$$d(x, B) = \min\{d(x, y) | y \in B\}$$

The distance from the compact A to the compact B as:

$$d(A, B) = \max\{d(x, B) | x \in A\}$$

The Hausdorff distance between two compacts A and B can thus be defined as:

$$h(A, B) = \max(d(A, B), d(B, A))$$

$(\mathcal{H}(\mathbb{R}_+^n), h)$ and $(\mathcal{C}(\mathbb{R}_+^n), h)$ are complete metric spaces [1] and we can thus study the continuity of policy functions under pretty clean conditions.

Definition 4 (Canonical partial orders). *We consider the following orders as being canonical.*

- The canonical partial \preceq order on \mathbb{R}_+^n is defined by:

$$u \preceq v \Leftrightarrow \forall k : u_k \leq v_k$$

- The canonical partial order on $\mathcal{H}(\mathbb{R}_+^n)$ is the classical inclusion order: \subseteq .

The two classical strict partial order \prec and \subset are defined accordingly. ($u \prec v \Leftrightarrow u \preceq v$ and $u \neq v$.)

- We also define an additional strict partial order \ll on \mathbb{R}_+^n , namely the strict Pareto-superiority, by:

$$u \ll v \Leftrightarrow \forall k : u_k < v_k$$

3. QUALITATIVE CHARACTERIZATIONS

In this section, we focus on qualitative characterizations of performance of allocations. Of particular interest are:

- The notion of Pareto optimality: a concepts that define the set of points of U that are globally optimal,
- Index or aggregation operators: they reflect the optimality of a point with respect to a particular criterion,
- Braess-like paradoxes: reflects whether an increase of the system resource can be detrimental to all users concurrently.

The rest of this section is organized as follows: after defining these three fundamental concepts, we study the link between Pareto-optimality and index optimization, the continuity of allocations and their monotonicity.

3.1 Common Definitions

We recall here the definitions of Pareto optimality, index-optimizing function and Braess-like paradoxes.

Definition 5 (Pareto optimality). *A choice $u \in U$ is said to be Pareto optimal if*

$$\forall v \in U, \exists i, v_i > u_i \Rightarrow \exists j, v_j < u_j.$$

In other words, u is Pareto optimal if it is maximal in U for the canonical partial order on \mathbb{R}_+^n .

A policy function is said to be Pareto-optimal if for all $U \in \mathcal{U}$, $\alpha(U)$ is Pareto-optimal.

The key idea here is that Pareto optimality is a global notion. Even in systems that consists of independent elements, the Pareto optimality cannot be determined on each independent subsystem. Such phenomena has been exhibited in [6]. The considered system is a master-slave platform in which the master can communicate with as many slaves as it needs at any time. The master holds a infinite number of tasks corresponding to N applications, and each of them can be executed on any slave. The authors study the system at the Nash equilibrium (each application competing with each other for both resource and CPU). Although the problems associated with each machine is independent, the authors show that for any system with one slave the equilibrium is Pareto optimal, while Pareto inefficiency can occur in multiple slave systems.

Definition 6 (f -increasing). *A policy α is said to be f -increasing if $f \circ \alpha$ is monotone. Any f -optimizing policy is thus f -increasing.*

Definition 7 (Common Indexes). *Many different indexes have been proposed in the literature. We present a few ones:*

- **Arithmetic mean:** $\sum_i u_i$.

- **Minimum:** $\min_i u_i$.

- **Maximum:** $\max_i u_i$.

- **Geometric Mean:** also called Nash Bargaining Solution or proportional fairness $\prod_i u_i$.

- **Harmonic Mean:** $\frac{1}{\sum_i 1/u_i}$.

- **Quasi-arithmetic Mean:** $f^{-1}(\frac{1}{n} \sum_{i=1}^n f(u_i))$ where f is a strictly monotone continuous function on $[0, +\infty]$. The particular case where f is defined by $f : x \rightarrow x^\delta$ has been widely studied [7]. The five previous index are particular case of this index for particular values of δ (respectively, $1, -\infty, +\infty, 0$ and -1).

- **Jain:** $\frac{(\sum u_i)^2}{n \sum u_i^2}$ (see [3]).

- **Ordered Weighted Averaging:** $OWA(u_1, \dots, u_n) = \sum_i w_i \cdot u_{\sigma(i)}$ where σ is a permutation such that $u_{\sigma(1)} \leq u_{\sigma(2)} \leq \dots \leq u_{\sigma(n)}$.

All these indexes are continuous, however, some of them are not strictly monotone.

Definition 8 (Braess-paradox). *A policy function α is said to have Braess-paradoxes if there exists U_1 and U_2 such that*

$$U_1 \subset U_2 \text{ and } \alpha(U_1) \gg \alpha(U_2)$$

with \gg defined as in definition 4. A policy function such that there is no Braess-paradox is called Braess-paradox-free.

3.2 Pareto-optimality and Index Optimization

Pareto optimality and monotonicity of the index optimization are closely related, as illustrated in the following results.

Theorem 1. *Let α be an f -optimizing policy. If f is strictly monotone then α is Pareto-optimal.*

Proof. Suppose that α is not Pareto optimal. Then, there exists U such that $\alpha(U)$ is not Pareto optimal. Hence, there exists $v \in U$ such that $\alpha(U) \prec v$, and hence $f(\alpha(U)) < f(v)$, which contradicts the definition of $\alpha(U)$. \square

Theorem 2. Let α be an f -optimizing policy. If α is Pareto-optimal then f is monotone.

Proof. Suppose that f is not monotone. Then there exists $u \prec v$ such that $f(v) \prec f(u)$. Consider $U = \{u, v\}$. As $u \prec v$ and α is Pareto-optimal, then $\alpha(U) = v$ which is in contradiction with $f(v) \prec f(u)$. \square

Arithmetic and geometric mean indexes are examples or strictly monotone index. On the other hand, the Jain index is an example of non-monotone index. The min and the max index are also not strictly monotone, which is why, max-min fairness or min-max fairness are recursively defined in the literature.

3.3 Continuity

Let us assume that a configuration r made of p resources is modeled as an element r of \mathbb{R}_+^p . A set of configurations R is thus a compact of $\mathcal{H}(\mathbb{R}_+^p)$. Let us assume that utility of users g are continuous functions from \mathbb{R}_+^p to \mathbb{R}_+ . Then utility sets U are built with the help of R and g .

$$U : \begin{cases} (\mathcal{H}(\mathbb{R}_+^p), C_0(\mathbb{R}_+^p, \mathbb{R}_+^n)) & \rightarrow \mathcal{H}(\mathbb{R}_+^n) \\ (R, g) & \mapsto \{g(r) | r \in R\} \end{cases}$$

The mapping U being continuous, $\alpha \circ U$ represents the sensibility of the allocation with respect to resources and utility functions. Continuity of the allocation α is thus an essential feature. Indeed, it ensures that a slight change in the system resources would not significantly affect the allocation. In dynamically changing systems, this ensures a certain stability. It also ensures that a slight error in utility functions does not affect too much the allocation.

Theorem 3. The Pareto set of a convex utility set is not necessarily compact.

The function $\bar{\mathcal{P}}$ from $\mathcal{C}(\mathbb{R}_+^n)$ to $\mathcal{H}(\mathbb{R}_+^n)$ that associates to U the closure of its Pareto set is not continuous.

Proof. Let us first exhibit a convex utility set whose Pareto set is not closed. Let $C = \{(x, y, z) \in \mathbb{R}_+^3 | x^2 + y^2 \leq 1, 0 \leq z \leq 1 - \frac{1}{2} \frac{x^2 - 2y^2 + 1}{1-y}\}$. The set C is depicted on Figure 1. We have $\bar{\mathcal{P}}(C) = \{(x, y, 1 - \frac{1}{2} \frac{x^2 - 2y^2 + 1}{1-y}) | x^2 + y^2 \leq 1, x + y \geq 1, x > 0\} \cup \{(0, 1, 1)\}$, which is not closed.

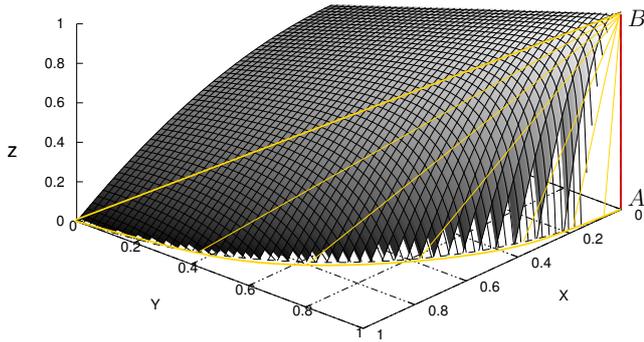


Figure 1: Convex set whose Pareto set is not closed. The segment $[A, B]$ does not belong to the Pareto set.

Let us consider $\bar{\mathcal{P}}$ from $\mathcal{C}(\mathbb{R}_+^n)$ to $\mathcal{H}(\mathbb{R}_+^n)$ that associate to U the closure of its Pareto set. Figure 2 depicts a converging

sequence of convex C_n such that $\bar{\mathcal{P}}(C_n)$ does not converge to $\bar{\mathcal{P}}(C_\infty)$. \square

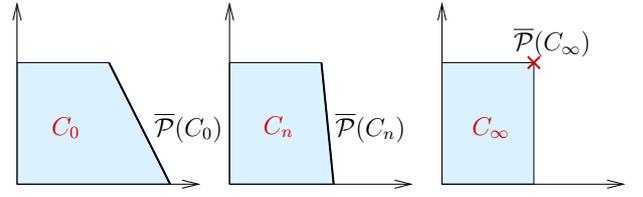


Figure 2: $\bar{\mathcal{P}}$ is not continuous.

Theorem 4. Let α be a general Pareto-optimal policy function. α is not continuous.

Proof. We prove that α cannot be continuous with the simple instances depicted on Figure 3. The only Pareto-optimal points are A and B . Therefore α has to choose in the first set between A and B . If A_1 is chosen, then by moving A_1 to A_0 , the choice has to “jump” to B , hence α is not continuous. \square

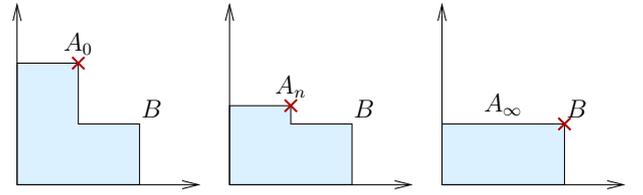


Figure 3: General Pareto-optimal policies are discontinuous: a path leading to discontinuity.

Remark 1. There exists continuous and non-continuous convex Pareto-optimal policy functions.

Proof. Let us consider a policy function α optimizing the sum of utilities. The two convex sets on Figure 4 show that α is not continuous around the set $K = \{(x, y) | x + y \leq 1\}$. This discontinuity is due to the fact that many different points of K simultaneously optimize the sum.

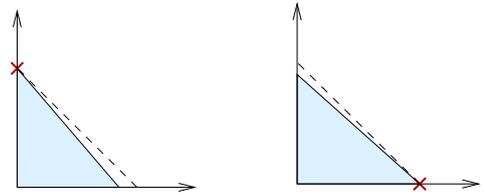


Figure 4: Optimizing \sum : a discontinuous convex policy.

The policy function α optimizing the product \prod of utilities is continuous though. As \prod is strictly monotone, α is Pareto-optimal. Moreover, as for any c , $I_c = \{x \in \mathbb{R}_+^n | \prod x_i \geq c\}$ is strictly convex, for any convex, there is a single point optimizing the \prod . Let us assume by contradiction that α is not continuous at the point C . Then there exists C_n converging to C and such that $x_n = \alpha(C_n)$ converges to

$x_\infty \neq \alpha(C)$. As our sets are compact, there exists a sequence $y_n \in C_n$ such that y_n converges to $\alpha(C)$. By definition, we have $\forall n, \prod(y_n) \leq \prod(x_n)$. Therefore $\prod(\alpha(C)) \leq \prod(x_\infty)$, which is absurd as $\alpha(C)$ is optimal in C for \prod and $\alpha(C) \neq x_\infty$. \square

3.4 Monotonicity

We state in this sub-section two results on monotonicity of index and policy functions. The first one emphasizes that index-functions only measures a specific characteristic of performance measure, and are hence not compatible. This explains why allocations that are efficient (optimizing the arithmetic mean) cannot (in general) also be fair (optimizing the geometric mean).

The second result states that, even when restricted to convex utility sets, policy functions cannot be monotone. This infers that even in Braess-free systems, an increase in the resource can be detrimental to some users.

Theorem 5. *Let f and g be two monotone index functions. A g -optimizing policy α_g is f -increasing if and only if α_g is f -optimizing.*

Proof. If α_g is f -optimizing, then α_g is clearly f -increasing.

Let us assume that α_g is not f -optimizing. We define the partial order \prec_f (resp. \prec_g) on \mathbb{R}_+^n by $x \prec_f y$ iff $f(x) \leq f(y)$. We have $\prec_f \neq \prec_g$, otherwise α_g would be f -optimizing. Thus there exists \bar{x}_1 and \bar{x}_2 such that: $\bar{x}_1 \prec_f \bar{x}_2$ and $\bar{x}_2 \prec_g \bar{x}_1$. Considering $U = \{x_1\}$ and $U' = \{x_1, x_2\}$, shows that α_g is not f -increasing. \square

In other words, a policy optimizing an index f is always non-monotone for a *distinct* index g .

Theorem 6. *Even if convex, policy functions cannot be monotone.*

Proof. Let us consider α a monotone convex policy function and let us consider the three following convex sets $U_1 = \{(0, 1)\}$, $U_2 = \{(1, 0)\}$, and $U_3 = \{(x, 1-x) | 0 \leq x \leq 1\}$ (see Figure 5).

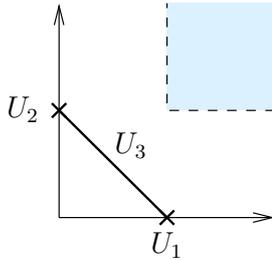


Figure 5: Even convex policy functions cannot be monotone.

We necessarily have $\alpha(U_1) = (0, 1)$ and $\alpha(U_2) = (1, 0)$. As U_1 and U_2 are subsets of U_3 , we have $\alpha(U_3) \succ (1, 1)$, which is absurd because no such point belongs to U_3 . \square

3.5 Conclusion

In this section, we have established the following results:

- Indexes should be strictly monotone to ensure Pareto-Optimality.

- Continuity (of allocations) is only possible when considering convex utility sets.
- It is impossible to ensure that the growth of the utility set does not incur the decrease of the utility of some player (i.e. policy functions cannot be monotone, even when restricting to convex utility sets).
- A policy optimizing a given index f leads to erratic values of an other index g when growing utility sets (unless f and g induce the same optimization).

Note that even though being Braess-paradox-free does not lead to bad properties, it does not give any information on the *efficiency* of such policies. For example, an allocation α that would be defined as returning 1/1000 of the NBS to all users would obviously be Braess-paradox-free but is very inefficient. This calls for more quantitative characterization of efficiency.

4. QUANTITATIVE CHARACTERIZATIONS

How to measure the efficiency of a given policy is still an open question. Many approaches have been proposed in the literature but we will see in Section 4.1 that none of the previously proposed approach is fully satisfying. We discuss in particular the most two popular ones: the *Jain index* [3] and the *Price of Anarchy* [5]. Then in Section 4.2, we propose a new metric based on a more topological point of view and explain how it relates to the notion of ε -approximation [8].

4.1 Discussion

4.1.1 Jain index

The Jain efficiency measure (or Jain index) [3] of a choice u is defined as $\frac{(\sum u_i)^2}{n \sum u_i^2}$. The Jain index is thus the ratio of the first to the second moment of the choice u . Hence, it is a good measure of a choice fairness (as defined by max-min fairness). The Jain index has many interesting properties:

- It is independent of the number of users.
- It remains unchanged if the utility set is linearly scaled.
- It is bounded (by $1/n$ and 1).
- It is continuous.

It can be straight-forwardly adapted to any measure of fairness when considering the ratio of the first and second moment of z where for all i , $z_i = u_i/v_i$ where v_i is the fair considered point. Another interpretation of the Jain index is to write it as: $1/n \sum_i (u_i/u_f)$ where $u_f = (\sum u_i^2)/(\sum u_i)$. Then each u_i/u_f represents the ratio of the choice with the fair allocation. The Jain is then the mean of these values. The index is therefore considered a useful “distance” measure to a given fair point.

The interest of the Jain factor is to determine which users are discriminated, and which are favored in a given allocation. Transfer of share from favored to discriminate users always increase the index, while the opposite reduces it.

However, as we have seen in Section 3.2 the Jain index is non-monotone (see Figure 6(a)), hence optimal solutions for the Jain index may not be Pareto-optimal. Even worse, some max-min fair allocations (that are as “fair” as possible) may have sub-optimal Jain index. Such an example is given on Figure 6(b).

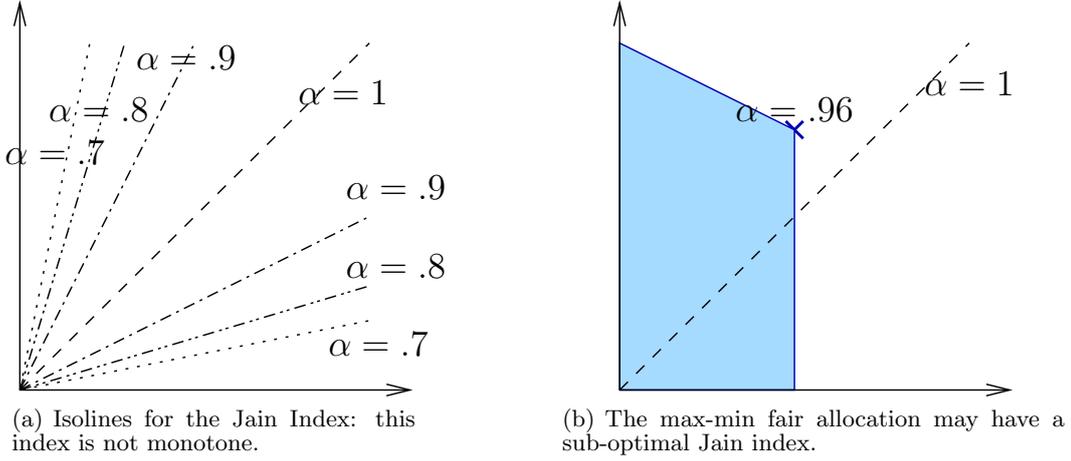


Figure 6: Highlighting Jain's index flaws.

4.1.2 Price of Anarchy and Index-Optimizing Based Metrics

Index-optimizing based metrics are easy to compute, continuous and generally conserve Pareto-superiority (under some mild conditions). It is thus natural to select an index f and to try to compare an allocation to the optimal one for f . Papadimitriou [5] introduced the now popular measure “price of anarchy” that we will study in this section.

For a given index f , let us consider $\alpha^{(f)}$ a f -optimizing policy function. We define the inefficiency $I_f(\beta, U)$ of the allocation $\beta(U)$ for f as

$$I_f(\beta, U) = \frac{f(\alpha^{(f)}(U))}{f(\beta(U))} \geq 1 = \max_{u \in U} \frac{f(u)}{f(\beta(U))}. \quad (1)$$

Papadimitriou focuses on the arithmetic mean Σ defined by $\Sigma(u_1, \dots, u_k) = \sum_{k=1}^K u_k$. The price of anarchy ϕ_Σ is thus defined as the largest inefficiency:

$$\phi_\Sigma(\beta) = \sup_{U \in \mathcal{U}} I_f(\beta, U) = \sup_{U \in \mathcal{U}} \frac{\sum_k \alpha^{(\Sigma)}(U)_k}{\sum_k \beta(U)_k}$$

In other words, $\phi_\Sigma(\beta)$ is the *approximation ratio* of β for the objective function Σ . This measure is very popular and rather easy to understand. However, we will see that it may not reflect what people have in mind when speaking about “price of anarchy”.

Consider the utility set $S_{M,N} = \{u \in \mathbb{R}_+^N \mid u_1/M + \sum_{k=2}^N u_k \leq 1\}$ depicted in Fig 7. As the roles of the u_k , $k \geq 2$ are symmetric, we can freely assume that $u_2 = \dots = u_N$ for metrical index-optimizing policies.

Remark 2. This example was taken from the master-slave scheduling problem of [6].

It is then easy to compute the following index optimizing allocation:

- $\alpha^{(\Sigma)}(S_{M,N}) = (M, 0, \dots, 0)$ corresponds to the allocation optimizing the average utility;
- $\alpha^{(\min)}(S_{M,N}) = \left(\frac{1}{N-1+1/M}, \dots, \frac{1}{N-1+1/M}\right)$ corresponds to the max-min fair allocation [9];

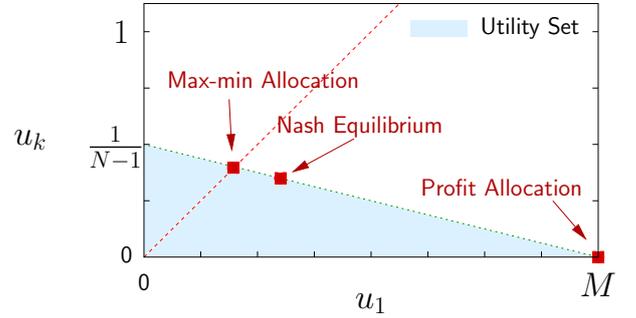


Figure 7: Utility set and allocations for $S_{M,N}$ ($N = 3, M = 2$), with $u_2 = \dots = u_N$.

- $\alpha^{(\Pi)}(S_{M,N}) = \left(\frac{M}{N}, \frac{1}{N}, \dots, \frac{1}{N}\right)$ corresponds to the proportionally fair allocation which is a particular Nash Bargaining Solution [9].

Note that, $\alpha^{(\Sigma)}$, $\alpha^{(\min)}$, and $\alpha^{(\Pi)}$ are Pareto optimal by definition. One can easily compute the price of anarchy of the Nash Bargaining solution:

$$I_\Sigma(\alpha^{(\Pi)}, S_{M,N}) = \frac{M}{\frac{M}{N} + \frac{N-1}{N}} \xrightarrow{M \rightarrow \infty} N.$$

The price of anarchy is therefore unbounded. However, the fact that this allocation is Pareto-optimal and has interesting properties of fairness (it corresponds to a Nash Bargaining Solution [9]) questions the relevance of the *price of anarchy* notion as a Pareto efficiency measure.

Likewise, the inefficiency of the max-min fair allocation is equivalent to M for large values of M (as opposed to K for the non-cooperative equilibrium). It can hence be unbounded even for bounded number of applications and machines. This seems even more surprising as such points generally result from complex cooperations and are hence Pareto optimal. These remarks raise once more the question of the measure of Pareto inefficiency.

These are due to the fact that a policy optimizing an index f is always non-monotone for a distinct index g (from Theorem 5). Hence any policy (including Pareto optimal ones)

optimizing a distinct index from the arithmetic mean will experience a bad price of anarchy. Note that the previous problems are not specific to the efficiency measure arithmetic mean. The same kind of behavior can be exhibited when using the min or the product of the throughputs for instance.

That is why we think that Pareto inefficiency should be measured as the *distance* to the Pareto border and not to a specific point.

4.1.3 Selfishness Degradation Factor

To quantify the degradation of Braess-like Paradoxes (the degree of Paradox), Kameda [4] introduced the Pareto-comparison of α and β as $\varrho(\alpha, \beta) = \min_k \frac{\alpha_k}{\beta_k}$. Therefore, α is strictly superior to β iff $\varrho(\alpha, \beta) > 1$. Intuitively ϱ represents the performance degradation between α and β . Using this definition, the following definition of Pareto inefficiency, named Selfishness Degradation Factor (SDF), was proposed [6]:

$$I_{SDF}(\beta, U) = \max_{u \in U} \varrho(u, \beta(U)) = \max_{u \in U} \min_k \frac{u_k}{\beta(U)_k} \quad (2)$$

Therefore $\beta(U)$ is Pareto inefficient as soon as $I_{SDF}(\beta, U) > 1$ and the larger $I_{SDF}(\beta, U)$, the more inefficient the allocation.

Lemma 1. Let us define $\mathbb{1}_{\leq} = \{x \in \mathbb{R}^n | \exists k : x_k \leq 0\}$. We denote by $a\mathbb{1}_{\leq} = \{x \in \mathbb{R}^n | \exists k : x_k \leq a_k\}$.

$$\log(I_{SDF}(\beta, U)) \leq \varepsilon \Leftrightarrow \log(U) \subseteq (\log(\beta(U)) + \varepsilon)\mathbb{1}_{\leq}$$

Proof.

$$\begin{aligned} \log(I_{SDF}(\beta, U)) \leq \varepsilon &\Leftrightarrow \max_{u \in U} \min_k \log\left(\frac{u_k}{\beta(U)_k}\right) \leq \varepsilon \\ &\Leftrightarrow \forall u \in U, \exists k, \log(u_k) - \log(\beta(U)_k) \leq \varepsilon \\ &\Leftrightarrow \forall u \in U, \exists k, \log(u_k) \leq \log(\beta(U)_k) + \varepsilon \\ &\Leftrightarrow \forall u \in U, \log(u) \in (\log(\beta(U)) + \varepsilon)\mathbb{1}_{\leq} \\ &\Leftrightarrow \log(U) \subseteq (\log(\beta(U)) + \varepsilon)\mathbb{1}_{\leq} \quad \square \end{aligned}$$

Figure 8 depicts a graphical interpretation of this inefficiency measure. As illustrated by the previous lemma, this inefficiency seems to measure how much $\beta(U)$ should be increased so that it is not dominated by any other points in U . Therefore, $\log(I_{SDF}(\beta, U))$ somehow measures the distance in the log-space from $\beta(U)$ to the Pareto set. However, as we will see in the next section, this definition holds only because of the very specific shape of the set U used in this example.

Anyway, the *selfishness degradation factor* can, as in section 4.1.2, be defined from this inefficiency measure:

$$\phi = \sup_{U \in \mathcal{U}} I_{SDF}(\beta, U) = \sup_{U \in \mathcal{U}} \max_{u \in U} \min_k \frac{u_k}{\beta(U)_k}.$$

A system (e.g., queuing network, transportation network, load-balancing, ...) that would be such that the Nash equilibria are always Pareto optimal would have a selfishness degradation factor equal to one. The selfishness degradation factor may however be unbounded on systems where non-cooperative equilibria are particularly inefficient. The relevance of this definition is corroborated by the fact that ε -approximations of Pareto-sets defined by Yannakakis and Papadimitriou [8] have a degradation factor of $\exp(\varepsilon) \simeq 1 + \varepsilon$.

4.2 A Topological Point of View

In this section, we go back to the inefficiency measure introduced in the previous section and show that such a measure can be properly defined only when referring to the whole Pareto set. Indeed, what we are interested in is in fact some kind of distance of a point to the Pareto set. As researchers are used to look at factors when evaluating the performance of an algorithm, this distance to the Pareto set should be measured in the log space. As we have seen in the previous section, the inefficiency measure for the selfishness degradation factor is closely related to the distance to the Pareto set. More precisely, we prove that being close to the Pareto set implies a small measure of inefficiency. However, the converse is true only when the utility set has some particular properties.

The distance from $\beta(U)$ to the closure of the Pareto set $\overline{\mathcal{P}}(U)$ in the log-space is equal to:

$$d_\infty(\log(\beta(U)), \log(\overline{\mathcal{P}}(U))) = \min_{u \in \overline{\mathcal{P}}(u)} \max_k |\log(\beta(U)_k) - \log(u_k)|$$

Therefore, we can define

$$\begin{aligned} \tilde{I}_\infty(\beta, U) &= \exp(d_\infty(\log(\beta(U)), \log(\overline{\mathcal{P}}(U)))) \\ &= \min_{u \in \overline{\mathcal{P}}(u)} \max_k \max\left(\frac{\beta(U)_k}{u_k}, \frac{u_k}{\beta(U)_k}\right) \quad (3) \end{aligned}$$

Let us recall the classical expansion definition:

$$X \oplus a = \{y | d(x, y) \leq a, \text{ for some } x \in X\}$$

This definition can be easily expanded as:

$$\begin{aligned} X \otimes a &= \exp(\log(X) \oplus \log(a)) \\ &= \{y | \exp(d(\log(x), \log(y))) \leq a \text{ for some } x \in X\} \end{aligned}$$

Definition 9 (ε -approximation). [8] defines an ε -approximation of $\overline{\mathcal{P}}(U)$ as a set of points S such that for all $u \in U$ there exists some $s \in S$ such that $\forall k : u_k \leq (1 + \varepsilon)s_k$.

With the previous notations, it is easy to see that:

Theorem 7. $S \subseteq U$ is an ε -approximation of $\overline{\mathcal{P}}(U)$ iff $\overline{\mathcal{P}}(U) \subseteq S \otimes \exp(\varepsilon)$.

Figure 9(d) depicts the expansion of $\log(\overline{\mathcal{P}}(U))$ by ε so that it contains $\log(\beta(U))$. It is easy to show that:

Lemma 2. $\tilde{I}_\infty(\beta, U) \leq \exp(\varepsilon) \Leftrightarrow \beta(U) \in \overline{\mathcal{P}}(U) \otimes \exp(\varepsilon)$.

In other words, $\tilde{I}_\infty(\beta, U) \leq \exp(\varepsilon)$ iff $\beta(U)$ is no farther than ε from $\overline{\mathcal{P}}(U)$ in the log space.

When comparing the definitions of I_Σ , I_{SDF} and \tilde{I}_∞ , the latest may seem harder to compute as it relies on $\overline{\mathcal{P}}(U)$. However, what we are interested in is measuring the distance to the Pareto set and no index-based inefficiency measure can reflect this distance. Then can only reflect a particular property of the allocation such as fairness. Note that in mono-criteria situations, it is natural to compare a solution to an intractable optimal solution, generally using approximations or lower bounds. Therefore, similar approaches should be used in multi-criteria settings to compute \tilde{I}_∞ . This inefficiency measure is thus a natural extension of the classical mono-criteria performance ratio.

The previous definition should thus be used in the general case, even though in a some particular situations, the SDF definition is sufficient.

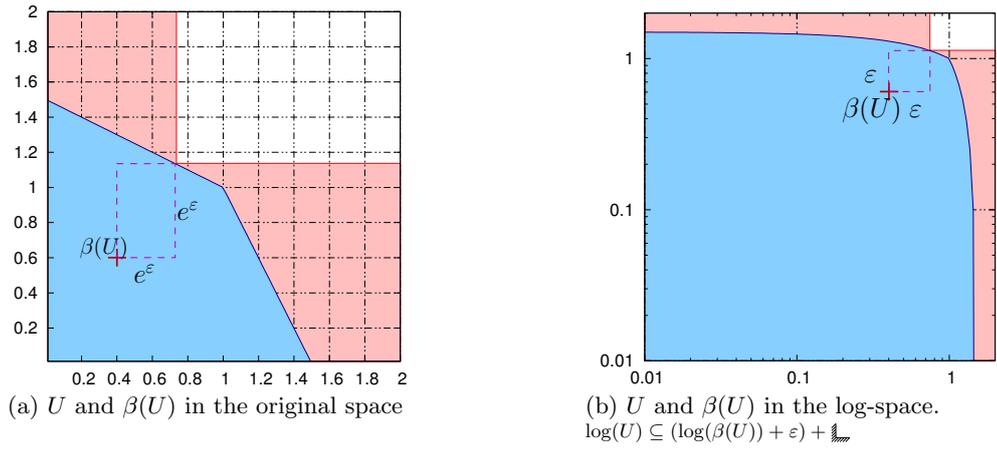


Figure 8: Inefficiency for the selfishness degradation factor: $\log(U) \subset (\log(\beta(U)) + \varepsilon)$

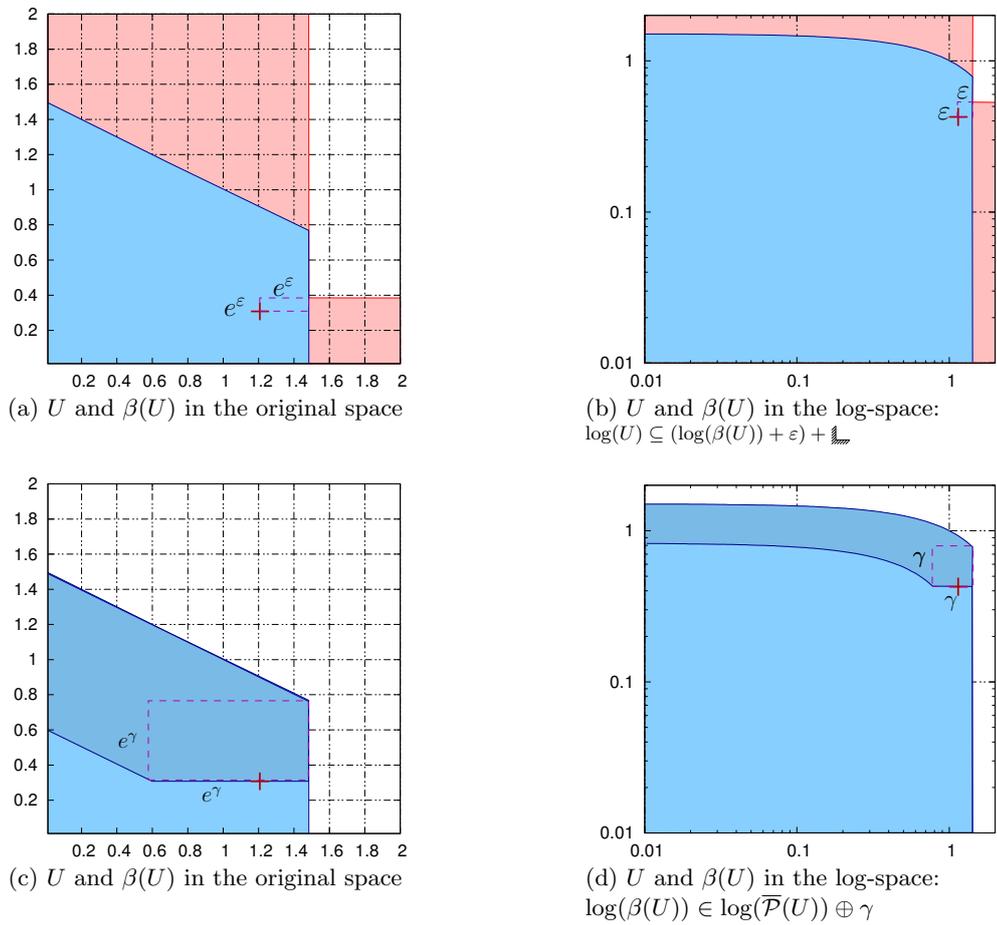


Figure 9: Distance to the Pareto set

5. CONCLUSION

In this paper, we have addressed the question of how to properly measure efficiency of allocations, may they be obtained as the result of some index-function optimization or some general policy. We have shown a number of results, both at qualitative and quantitative level. In particular, we have shown that:

- Monotonicity is the link between index-optimization and Pareto optimality.
- When utilities are continuous with the system's resources, solution allocations can be continuous in the resources only when the utility sets are convex.
- Even with Braess-free allocations, there always exists instances where resource increase is detrimental to at least one user.
- A policy optimizing a given index leads to erratic values for another index when utility sets grow.
- Both the Jain index and the price of anarchy have flaws as measures of the inefficiency of an equilibria.
- A correct general inefficiency measure can be defined and is based on the log space as the distance of a point to the Pareto border.

We believe that these results can serve as a general theoretical milestones to any researcher aiming at analyzing the performance of an allocation in a specific problem.

6. REFERENCES

- [1] M. Barnsley. *Fractals Everywhere*. Morgan Kaufmann, 1988.
- [2] M. Detyniecki. *Mathematical Aggregation Operators and their Application to Video Querying*. PhD thesis, University of Paris 6, France, 2000.
- [3] R. K. Jain, dah Ming W. Chiu, and W. R. Hawe. A quantitative measure of fairness and discrimination for resource allocation in shared computer systems. Technical Report DEC-TR-301, DEC Research Report, September 1984.
- [4] H. Kameda. Bounds on benefits and harms of adding connections to noncooperative networks. In N. Mitrou, K. Kontovasilis, G. N. Rouskas, I. Iliadis, and L. Merakos, editors, *NETWORKING 2004*, volume 3042 of *LNCS*, pages 405–417. Springer Verlag, 2006.
- [5] E. Koutsoupias and C. Papadimitriou. Worst-case equilibria. In *STACS*, 1998.
- [6] A. Legrand and C. Touati. Non-cooperative scheduling of multiple bag-of-task applications. In *Proceedings of the 25th Conference on Computer Communications (INFOCOM'07)*, Alaska, USA, May 2007.
- [7] J. Mo and J. Walrand. Fair end-to-end window-based congestion control. *IEEE/ACM Transactions on Networking*, 8(5):556–567, 2000.
- [8] C. H. Papadimitriou and M. Yannakakis. On the approximability of trade-offs and optimal access of web sources. In *FOCS '00: Proceedings of the 41st Annual Symposium on Foundations of Computer Science*, page 86, Washington, DC, USA, 2000. IEEE Computer Society.
- [9] C. Touati, E. Altman, and J. Galtier. Generalised Nash bargaining solution for bandwidth allocation. *Computer Networks*, 2006 (to appear).