



Delft University of Technology

Document Version

Final published version

Citation (APA)

Si Salem, T., Iosifidis, G., & Neglia, G. (2023). Enabling Long-term Fairness in Dynamic Resource Allocation. *Performance Evaluation Review*, 51(1), 31-32. <https://doi.org/10.1145/3606376.3593541>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

In case the licence states "Dutch Copyright Act (Article 25fa)", this publication was made available Green Open Access via the TU Delft Institutional Repository pursuant to Dutch Copyright Act (Article 25fa, the Taverne amendment). This provision does not affect copyright ownership.

Unless copyright is transferred by contract or statute, it remains with the copyright holder.

Sharing and reuse

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

This work is downloaded from Delft University of Technology.

Green Open Access added to TU Delft Institutional Repository

'You share, we take care!' - Taverne project

<https://www.openaccess.nl/en/you-share-we-take-care>

Otherwise as indicated in the copyright section: the publisher is the copyright holder of this work and the author uses the Dutch legislation to make this work public.



Enabling Long-term Fairness in Dynamic Resource Allocation

Tareq Si Salem
tareq.si-salem@inria.fr
Inria
Sophia Antipolis, France

George Iosifidis
G.Iosifidis@tudelft.nl
TU Delft
Delft, The Netherlands

Giovanni Neglia
giovanni.neglia@inria.fr
Inria
Sophia Antipolis, France

ABSTRACT

We study the fairness of dynamic resource allocation problem under the α -fairness criterion. We recognize two different fairness objectives that naturally arise in this problem: the well-understood slot-fairness objective that aims to ensure fairness at every timeslot, and the less explored horizon-fairness objective that aims to ensure fairness across utilities accumulated over a time horizon. We argue that horizon-fairness comes at a lower price in terms of social welfare. We study horizon-fairness with the regret as a performance metric and show that vanishing regret cannot be achieved in presence of an unrestricted adversary. We propose restrictions on the adversary's capabilities corresponding to realistic scenarios and an online policy that indeed guarantees vanishing regret under these restrictions.

CCS CONCEPTS

- Applied computing → Multi-criterion optimization and decision-making;
- Theory of computation → Online learning algorithms; Algorithmic game theory.

KEYWORDS

Online Learning, Multi-timescale Fairness, Axiomatic Bargaining, Dynamic Resource Allocation

ACM Reference Format:

Tareq Si Salem, George Iosifidis, and Giovanni Neglia. 2023. Enabling Long-term Fairness in Dynamic Resource Allocation. In *Abstract Proceedings of the 2023 ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Systems (SIGMETRICS '23 Abstracts)*, June 19–23, 2023, Orlando, FL, USA. ACM, New York, NY, USA, 2 pages. <https://doi.org/10.1145/3578338.3593541>

1 INTRODUCTION

Achieving fairness when allocating resources in communication and computing systems has been a subject of extensive research, and has been successfully applied in numerous practical problems. Fairness is leveraged to perform congestion control in the Internet [6, 13], to select transmission power in multi-user wireless networks [17, 22], and to allocate multidimensional resources in cloud computing platforms [3, 20, 21]. Depending on the problem at hand, the criterion of fairness can be expressed in terms of how the service performance is distributed across the end-users, or in

Permission to make digital or hard copies of part or all 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. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

SIGMETRICS '23 Abstracts, June 19–23, 2023, Orlando, FL, USA

© 2023 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-0074-3/23/06.

<https://doi.org/10.1145/3578338.3593541>

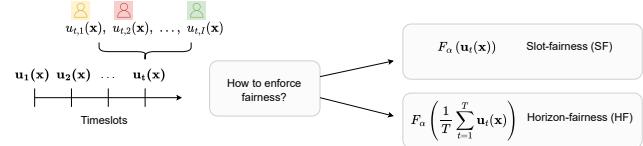


Figure 1: The different ways fairness can be enforced in a time-slotted dynamic system. The decision maker can either consider the α -fairness objective F_α at every timeslot (slot-fairness), or at the end of the time horizon T (horizon-fairness).

terms of how the costs are balanced across the servicing nodes. A prevalent fairness metric is α -fairness, which encompasses the utilitarian principle (Bentham-Edgeworth solution [5]), proportional fairness (Nash bargaining solution [15]), max-min fairness (Kalai-Smorodinsky bargaining solution [10]). All these fairness metrics have been used in different cases for the design of resource management mechanisms [14, 16].

A common limitation of the above works is that they consider *static* environments. That is, the resources to be allocated and, importantly, the users' utility functions, are fixed and known to the decision maker. This assumption is very often unrealistic for today's communication and computing systems.

2 MAIN RESULTS

This paper makes the next step towards enabling long-term fairness in dynamic systems. We consider a system that serves a set of agents \mathcal{I} , where a controller selects at each timeslot $t \in \mathbb{N}$ a resource allocation profile \mathbf{x}_t from a set of eligible allocations \mathcal{X} based on past agents' utility functions $\mathbf{u}_t : \mathcal{X} \rightarrow \mathbb{R}^{\mathcal{I}}$ for $t' < t$ and of α -fairness function $F_\alpha : \mathbb{R}_{\geq 0}^{\mathcal{I}} \rightarrow \mathbb{R}$. The utilities might change due to unknown, unpredictable, and (possibly) non-stationary perturbations that are revealed to the controller only after it decides \mathbf{x}_t . We employ the terms *horizon-fairness* (HF) and *slot-fairness* (SF) to distinguish the different ways fairness can be enforced in a such time-slotted dynamic system (see the illustration in Fig. 1). Under horizon-fairness, the controller enforces fairness on the aggregate utilities for a given time horizon T , whereas under slot-fairness, it enforces fairness on the utilities at each timeslot separately. Both metrics have been studied in previous work, e.g., see [7, 9, 11, 19]. Our focus is on horizon-fairness, which raises novel technical challenges and subsumes slot-fairness as a special case.

We design the *online horizon-fair* (OHF) policy by leveraging *online convex optimization* (OCO) [8], to handle this reduced information setting under a powerful *adversarial* perturbation model. Adversarial analysis is a modeling technique to characterize a system's

performance under unknown and hard-to-characterized exogenous parameters. In our context, the performance of a resource allocation policy \mathcal{A} is evaluated by the *fairness regret*, which is defined as the difference between the α -fairness, over the time-averaged utilities, achieved by a static optimum-in-hindsight (*benchmark*) and the one achieved by the policy:

$$\mathfrak{R}_T(F_\alpha, \mathcal{A}) \triangleq \sup_{\{\mathbf{u}_t\}_{t=1}^T \in \mathcal{U}^T} \left\{ \max_{\mathbf{x} \in \mathcal{X}} F_\alpha \left(\frac{1}{T} \sum_{t \in \mathcal{T}} \mathbf{u}_t(\mathbf{x}) \right) - F_\alpha \left(\frac{1}{T} \sum_{t \in \mathcal{T}} \mathbf{u}_t(\mathbf{x}_t) \right) \right\}. \quad (1)$$

If the fairness regret vanishes over time (i.e., $\lim_{T \rightarrow \infty} \mathfrak{R}_T(F_\alpha, \mathcal{A}) = 0$), policy \mathcal{A} will attain the same fairness value as the static benchmark under any possible sequence of utility functions. A policy that achieves sublinear regret under these adversarial conditions, can also succeed in more benign conditions where the perturbations are not adversarial, or the utility functions are revealed at the beginning of each slot.

The fairness regret metric (1) departs from the template of OCO. In particular, the scalarization of the vector-valued utilities, through the α -fairness function, is not applied at every timeslot to allow for the controller to easily adapt its allocations, instead is only applied at the end of the time horizon T . Our first result characterizes the challenges in tackling this learning problem. Namely, we prove that when utility perturbations are only subject to four mild technical conditions, such as in standard OCO, it is impossible to achieve vanishing fairness-regret. Similar negative results were obtained under different setups of primal-dual learning and online saddle point learning [1, 12, 18], but they have been devised for specific problem structures (e.g., online matrix games) and thus do not apply to our setting.

In light of this negative result, we introduce additional *necessary* conditions on the adversary to obtain a vanishing regret guarantee. Namely, the adversary can only induce perturbations to the time-averaged utilities we call budgeted-severity or partitioned-severity constrained. These conditions capture several practical utility patterns, such as non-stationary corruptions, ergodic and periodic inputs [2, 4, 11, 23]. We proceed to propose the OHF policy which adapts dynamically the allocation decisions and provably achieves $\mathfrak{R}_T(F_\alpha, \mathcal{A}) = o(1)$.

The OHF policy employs a novel learning approach that operates concurrently, and in a synchronized fashion, in a primal and a dual (conjugate) space. Intuitively, OHF learns the weighted time-varying utilities in a primal space, and learns the weights accounting for the global fairness metric in some dual space. To achieve this, we develop novel techniques through a convex conjugate approach.

Finally, we apply our fairness framework to a representative resource management problem in virtualized caching systems where different caches cooperate by serving jointly the received content requests. We evaluate the performance of OHF with its slot-fairness counterpart policy through numerical examples. We evaluate the price of fairness of OHF, which quantifies the efficiency loss due to fairness, across different network topologies and participating agents. Lastly, we apply OHF to a Nash bargaining scenario, a concept that has been widely used in resource allocation to distribute to a set of agents the utility of their cooperation.

3 CONCLUSION AND FUTURE WORK

In this work, we proposed a novel OHF policy that achieves horizon-fairness in dynamic resource allocation problems. Our work paves the road for several interesting next steps. A future research direction is to consider decentralized versions of the policy under which each agent selects an allocation with limited information exchange across agents. Another important future research direction is to bridge the horizon-fairness and slot-fairness criteria to target applications where the agents are interested in ensuring fairness within a target time window. A final interesting research direction is to consider a limited feedback scenario where only part of the utility is revealed to the agents (e.g., bandit feedback). Our policy could be extended to this setting through gradient estimation techniques [8].

REFERENCES

- [1] Daron Anderson, George Iosifidis, and Douglas J Leith. 2022. Lazy Lagrangians with Predictions for Online Learning. *arXiv preprint arXiv:2201.02890* (2022).
- [2] Santiago R Balseiro, Haihao Lu, and Vahab Mirrokni. 2022. The Best of many worlds: Dual mirror descent for online allocation problems. *Operations Research* (2022).
- [3] C. Joe-Wong, S. Sen, T. Lan, and M. Chiang. 2013. Multiresource Allocation: Fairness-Efficiency Tradeoffs in a Unifying Framework. *IEEE/ACM Trans. on Networking* 21, 6 (2013), 1785–1798.
- [4] John C Duchi, Alekh Agarwal, Mikael Johansson, and Michael I Jordan. 2012. Ergodic mirror descent. *SIAM Journal on Optimization* 22, 4 (2012), 1549–1578.
- [5] Francis Ysidro Edgeworth et al. 1881. Mathematical Psychics. *History of Economic Thought Books* (1881).
- [6] F.P. Kelly, A. Maulloo, and D. Tan. 1998. Rate Control for Communication Networks: Shadow Prices, Proportional Fairness, and Stability. *J. Oper. Res. Soc.* 49, 3 (1998), 237–252.
- [7] Swati Gupta and Vijay Kamble. 2021. Individual Fairness in Hindsight. *J. Mach. Learn. Res.* 22, 144 (2021), 1–35.
- [8] Elad Hazan. 2016. Introduction to Online Convex Optimization. *Foundations and Trends® in Optimization* 2, 3–4 (Aug. 2016), 157–325.
- [9] Devansh Jalota and Yinyu Ye. 2022. Online Learning in Fisher Markets with Unknown Agent Preferences. *arXiv preprint arXiv:2205.00825* (2022).
- [10] Ehud Kalai, Meir Smorodinsky, et al. 1975. Other Solutions to Nash's Bargaining Problem. *Econometrica* 43, 3 (1975), 513–518.
- [11] Luofeng Liao, Yuan Gao, and Christian Kroer. 2022. Nonstationary Dual Averaging and Online Fair Allocation. *ArXiv e-prints* (Feb. 2022).
- [12] Shie Mannor, John N Tsitsiklis, and Jie Yuan Yu. 2009. Online Learning with Sample Path Constraints. *Journal of Machine Learning Research* 10, 3 (2009).
- [13] Jeonghoon Mo and Jean Walrand. 2000. Fair End-to-End Window-Based Congestion Control. *IEEE/ACM Transactions on networking* 8, 5 (2000), 556–567.
- [14] Dritan Nace and Michal Pioro. 2008. Max-min fairness and its applications to routing and load-balancing in communication networks: a tutorial. *IEEE Communications Surveys & Tutorials* 10, 4 (2008), 5–17.
- [15] John F. Nash. 1950. The Bargaining Problem. *Econometrica* 18, 2 (1950), 155–162.
- [16] Bozidar Radunovic and Jean-Yves Le Boudec. 2007. A unified framework for max-min and min-max fairness with applications. *IEEE/ACM Transactions on networking* 15, 5 (2007), 1073–1083.
- [17] Resource Allocation and Cross Layer Control in Wireless Networks. 2006. L. Georgiadis, M. J. Neely, and L. Tassiulas. *Found. Trends Netw.* 1, 1 (2006), 1–143.
- [18] Adrian Rivera, He Wang, and Huan Xu. 2018. The Online Saddle Point Problem and Online Convex Optimization with Knapsacks. *preprint arXiv:1806.08301* (2018).
- [19] Sean R. Sinclair, Siddhartha Banerjee, and Christina Lee Yu. 2022. Sequential Fair Allocation: Achieving the Optimal Envy-Efficiency Tradeoff Curve. *SIGMETRICS Perform. Eval. Rev.* 50, 1 (jun 2022), 95–96.
- [20] T. Bonald, and J. W. Roberts. 2015. Multi-Resource Fairness: Objectives, Algorithms and Performance. In *ACM Sigmetrics*.
- [21] W. Wang, B. Li, and B. Liang. 2014. Dominant Resource Fairness in Cloud Computing Systems with Heterogeneous Servers. In *IEEE INFOCOM*.
- [22] X. Lin, N. B. Shroff, and R. Srikant. 2006. A Tutorial on Cross-Layer Optimization in Wireless Networks. *IEEE J. Sel. Areas Commun.* 24, 8 (2006), 1452–1463.
- [23] Yu-Hang Zhou, Chen Liang, Nan Li, Cheng Yang, Shenghua Zhu, and Rong Jin. 2019. Robust online matching with user arrival distribution drift. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 33. 459–466.