Brief Announcement: Compact Topology of Shared-Memory Adversaries*

Petr Kuznetsov[1], Thibault Rieutord[2], and Yuan He[3]

1 **LTCI, Télécom ParisTech, Université Paris-Saclay, Paris, France**
`petr.kuznetsov@telecom-paristech.fr`
2 **LTCI, Télécom ParisTech, Université Paris-Saclay, Paris, France**
`thibault.rieutord@telecom-paristech.fr`
2 **UCLA, Los Angeles, USA**
`yuan.he@cs.ucla.edu`

Abstract

The paper proposes a simple topological characterization of a large class of *adversarial* distributed-computing models via *affine tasks*: sub-complexes of the second iteration of the standard chromatic subdivision. We show that the task computability of a model in the class is precisely captured by iterations of the corresponding affine task. While an adversary is in general defined as a *non-compact* set of infinite runs, its affine task is just a finite subset of runs of the 2-round iterated immediate snapshot (IIS) model. Our results generalize and improve all previously derived topological characterizations of distributed-computing models.

1998 ACM Subject Classification C.2.4 Distributed Systems, F.1.1 Models of Computation.

Keywords and phrases Adversarial models, Affine tasks, Topological characterization.

Digital Object Identifier 10.4230/LIPIcs.DISC.2017.56

1 Introduction

Distributed computing is a jungle of models, parameterized by types of failures, synchrony assumptions, and employed communication primitives. Determining relative computability power of these models ("is model A more powerful than model B") is an intriguing and important problem.

The task computability of the *wait-free* model of computation, which makes no assumptions about the number of failures that can occur, was characterized by Herlihy and Shavit [7] through the existence of a specific continuous map from a subdivision of the input complex of a task $\mathcal{I}$ to its output complex $\mathcal{O}$. (The reader is referred to [6] for a thorough discussion of the use of combinatorial topology in distributed computability.) In particular, the characterization can consider the *iterated* standard chromatic subdivision ($\mathrm{Chr}\,\mathbf{s}$ depicted in Figure 1a) and, thus, derive that a task is wait-free solvable if and only if it can be solved in the IIS model.

The aim of this paper is to generalize this topological characterization to models beyond the wait-free model using the formalism of *affine tasks* [5]. An affine task is defined through a pure subcomplex of an iterated *standard chromatic subdivision* or, equivalently, a subset of finite runs of the IIS model. Iterations of such affine tasks provide a failure-less *compact* model (according to the "longest-prefix" metric [1]).

Given that many fundamental tasks are not wait-free solvable, the prominent *adversarial* failure model [2] has been introduced to strengthen the wait-free model. An adversary

* A full version of the paper is available at `https://hal.archives-ouvertes.fr/hal-01572257`.

31st International Symposium on Distributed Computing (DISC 2017).
Editor: Andréa W. Richa; Article No. 56; pp. 56:1–56:4
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

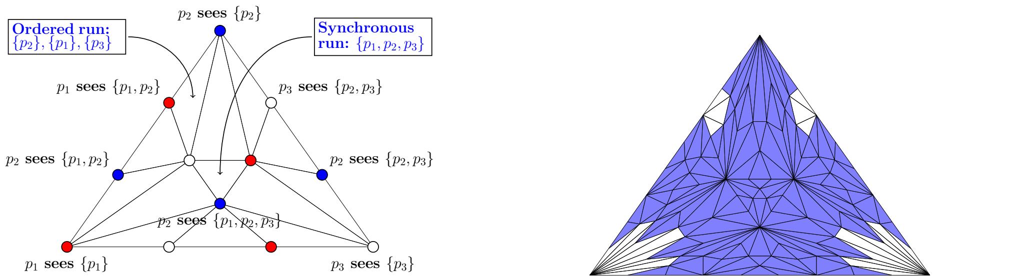

(a) $\mathrm{Chr}^1(\mathbf{s})$, equivalent to the output complex of the immediate snapshot (IS) task.

(b) The affine task of the adversary consisting of $\{p_1\}$, $\{p_2, p_3\}$ with all super-sets (in blue).

Figure 1 Standard chromatic subdivision and an affine task example for $n = 3$.

$\mathcal{A}$ is defined through a collection of process subsets, called *live sets*. In every run of the corresponding *adversarial $\mathcal{A}$-model*, the set of processes taking infinitely many steps must be a live set. The sub-class of *fair* adversaries [8] does not, intuitively, allow a *subset* of processes participating in a computation to achieve a better set consensus than the whole set of participants (processes taking at least one step). The class of *fair* adversaries is pretty large, as it includes the existing sub-classes of *superset-closed* and *symmetric* adversaries.

We show that a specific affine task $\mathcal{R}_{\mathcal{A}}$, defined as a subcomplex of the second iteration of the standard chromatic subdivision, captures the *task computability* of any fair adversary $\mathcal{A}$. A task is solved in the $\mathcal{A}$-model if and only it is solvable in the set of IIS runs resulting from iterations of $\mathcal{R}_{\mathcal{A}}$ (denoted $\mathcal{R}_{\mathcal{A}}^*$).

The notion of *agreement functions* [8] was instrumental for this result. (An agreement function α associates each set of processes P with the best level of set consensus solvable when only processes in P might participate.) Fair adversaries are characterized by their agreement function in the sense that they belong to the weakest equivalence class (in terms of task computability) of models with the same agreement function. Our characterization can then be put as a generalization of the celebrated Asynchronous Computability Theorem (ACT) [7]:

> A task $T = (\mathcal{I}, \mathcal{O}, \Delta)$, where $\mathcal{I}$ is the input complex, $\mathcal{O}$ is an output complex, and Δ is a map from $\mathcal{I}$ to sub-complexes of $\mathcal{O}$, is solvable in a fair adversarial $\mathcal{A}$-model if and only if there exists a natural number ℓ and a simplicial map $\phi : \mathcal{R}_{\mathcal{A}}^{\ell}(\mathcal{I}) \to \mathcal{O}$ carried by Δ (informally, respecting the task specification Δ).

2 Affine tasks for fair adversaries.

Two classes of affine tasks were recently defined. The class $\mathcal{R}_{t-res}$ was introduced in [10], with $\mathcal{R}_{t-res}^*$ equivalent to the t-resilient model. Similarly, the class $\mathcal{R}_k$ was introduced in [4], with $\mathcal{R}_k^*$ equivalent to the k-concurrent model. Interestingly, these models correspond to two "well-behaved" sub-classes of fair adversaries on opposite sides of the spectrum. In a sense, a fair adversary can be seen as a combination of *concurrency* and *resilience*, grasped using, resp., *contention* and *critical* simplices:

- **Contention simplices:** If processes are executed sequentially, they not only obtain distinct views out of IS, but also obtain the same view (inclusion) ordering out of multiple iterations. But to be combined with resiliency features, concurrency restrictions must be weaken to focus on "fully" conflicting processes. This is why we say that a simplex, or a group of processes, forms a 2-*contention simplex* if any two of its processes have distinct views in both IS iterations, ordered alternatively in each (see [9] for a formal definition).

- **Critical simplices:** For fair adversaries, concurrency may rise with participation irregularly. Critical simplices act as representatives of each increase of participation resulting in a concurrency increase. They are selected among processes with the smallest first IS output providing an observed participation corresponding to some non-nul level of set-consensus power. Moreover, in order to be identifiable as critical, they are selected as critical simplices only if grouped with sufficiently many other processes with the same IS output, so that if they are withdrawn from their own first IS observed participation the remaining participation is associated to a strictly smaller set-consensus power. Hence observing in the second IS all members of a critical simplex is enough to check that they together form a critical simplex.

Now we are ready to define the subcomplex $\mathcal{R}_A \subseteq \text{Chr}^2\,\mathbf{s}$. The idea is that a large 2-constention simplex may be allowed only if it terminates after a critical simplex associated with a large enough view, i.e., concurrency may rises only after sufficient ensured resilience. A $(n-1)$-dimensional simplex $\sigma \in \text{Chr}^2\,\mathbf{s}$ (composed of n vertices) belongs to $\mathcal{R}_A$ if and only if every sub-simplex of σ of size k which (1) is a 2-contention simplex; (2) does not include critical simplices members; (3) does not include processes observed by identifiable critical simplices (with a smaller second IS view); must observe a critical simplex with a first IS view associated to an agreeent power greater than or equal to k (see Figure 1b for an example).

The proof of equivalence between $\mathcal{R}^*_{\mathcal{A}}$ and the fair $\mathcal{A}$-adversary model is done using the equivalent α-model [8] which (1) allows for a simple resolution of $\mathcal{R}_{\mathcal{A}}$ in the α-model by simply executing two rounds of an IS algorithm, with a waiting phase in between the rounds (similarly to [10]); and (2) can be simulated easily as soon as α-adaptive set-consensus (see [8]) is solvable in the presence of read-write memory (similarly to [4]).

To summarize, this paper generalizes all previous topological characterizations of distributed computing models [7, 5, 4, 10]. We believe that the results can further be extended to all "practical" restrictions of the wait-free model of computations, beyond fair adversaries, which may potentially result in a complete computability theory for distributed computing [3].

References

1. Bowen Alpern and Fred B. Schneider. Defining liveness. *Information Processing Letters*, 21(4):181–185, October 1985.
2. Carole Delporte-Gallet, Hugues Fauconnier, Rachid Guerraoui, and Andreas Tielmann. The disagreement power of an adversary. *Distributed Computing*, 24(3-4):137–147, 2011.
3. Eli Gafni. Private communication. 2002.
4. Eli Gafni, Yuan He, Petr Kuznetsov, and Thibault Rieutord. Read-Write Memory and k-Set Consensus as an Affine Task. In *OPODIS 2016*, volume 70, pages 6:1–6:17, 2017. `doi:10.4230/LIPIcs.OPODIS.2016.6`.
5. Eli Gafni, Petr Kuznetsov, and Ciprian Manolescu. A generalized asynchronous computability theorem. In *PODC*, pages 222–231, 2014. `doi:10.1145/2611462.2611477`.
6. Maurice Herlihy, Dmitry N. Kozlov, and Sergio Rajsbaum. *Distributed Computing Through Combinatorial Topology*. Morgan Kaufmann, 2014.
7. Maurice Herlihy and Nir Shavit. The topological structure of asynchronous computability. *Journal of the ACM*, 46(2):858–923, 1999.
8. Petr Kuznetsov and Thibault Rieutord. Agreement functions for distributed computing models. In *NETYS*, pages 175–190, 2017. `doi:10.1007/978-3-319-59647-1_14`.
9. Petr Kuznetsov, Thibault Rieutord, and Yuan He. Compact topology of shared-memory adversaries. *HAL*, https://hal.archives-ouvertes.fr/hal-01572257, 2017.

10 Vikram Saraph, Maurice Herlihy, and Eli Gafni. Asynchronous computability theorems for t-resilient systems. In *DISC*, pages 428–441, 2016.