General Caching Is Hard: Even with Small Pages

Jiří Sgall

(in collaboration with Lukáš Folwarczný)

Computer Science Institute of Charles University, Faculty of Mathematics and Physics, Malostranské nám. 25, CZ-11800 Praha 1, Prague, Czech Republic

Caching (also known as *paging*) is a classical problem concerning page replacement policies in two-level memory systems. *General caching* is the variant with pages of different sizes and fault costs. The strong NP-hardness of its two important cases, the *fault model* (each page has unit fault cost) and the *bit model* (each page has the same fault cost as size) has been established in 2010 by Chrobak et al. [1], however, the reduction uses pages as large as half of the cache size. We prove that the strong NP-hardness holds already when page sizes are bounded by a small constant: The bit and fault models are strongly NP-complete even when page sizes are limited to $\{1, 2, 3\}$.

Considering only the decision versions of the problems, general caching is equivalent to the *unsplittable flow on a path problem* and therefore our results also improve the hardness results for this problem.

The results were presented in Proceedings of the 26th International Symposium on Algorithms and Computation (ISAAC 2015) [2].

References

[1] Chrobak, M., Woeginger, G.J., Makino, K., Xu, H.: Caching is hard – Even in the fault model. Algorithmica **63**(4), 781–794 (2012). A preliminary version appeared at ESA 2010.

[2] Folwarczný, L., Sgall, J.: General caching is hard: Even with small pages. In: K.M. Elbassioni, K. Makino (eds.) Algorithms and Computation - 26th International Symposium, ISAAC 2015, Nagoya, Japan, December 9-11, 2015, Proceedings, *Lecture Notes in Computer Science*, vol. 9472, pp. 116–126. Springer (2015).