

Walzer, Stefan:

**Cuckoo hashing with overlapping buckets**

**DOI:** [10.22032/dbt.42358](https://doi.org/10.22032/dbt.42358)  
**URN:** [urn:nbn:de:gbv:ilm1-2020200354](https://nbn-resolving.org/urn:nbn:de:gbv:ilm1-2020200354)

---

*Original published in:* Theory and applications of hashing : report from Dagstuhl Seminar 17181 / Dietzfelbinger, Martin; Mitzenmacher, Michael; Pagh, Rasmus; Woodruff, David P.; Aumüller, Martin. - Wadern : Schloss Dagstuhl. - 7 (2017), 5, p. 16.  
(Dagstuhl Reports : Dokumentationen zu Dagstuhl-Seminaren und Dagstuhl-Perspektiven-Workshops / Schloss Dagstuhl, Leibniz-Zentrum für Informatik. - Wadern : Schloss Dagstuhl. - 7 (2017), 5, p. 1-21. - DOI: [10.4230/DagRep.7.5.1](https://doi.org/10.4230/DagRep.7.5.1))

*ISSN:* 2192-5283  
*Original published:* 2017-12-19  
*[Visited:* 2020-03-12]



This work is licensed under a [Creative Commons Attribution 3.0 Unported](https://creativecommons.org/licenses/by/3.0/) license. To view a copy of this license, visit <https://creativecommons.org/licenses/by/3.0/>

### 3.21 Cuckoo Hashing with Overlapping Buckets

*Stefan Walzer (TU Ilmenau, DE)*

License  Creative Commons BY 3.0 Unported license  
© Stefan Walzer

Cuckoo hashing is the task of finding an injective assignment of  $m$  objects to  $n$  memory positions, where each object  $e$  is only allowed to reside in one of  $k \geq 2$  different positions that are chosen for  $e$  independently and uniformly at random. It is well understood how the number of choices  $k$  affects the critical density  $c^* \in (0, 1)$  around which the probability that a valid assignment exists jumps from almost 0 to almost 1.

In order to profit from memory caches when searching for an object, it is reasonable in practice to weaken the requirement that all choices be independent, and instead assign each object  $k$  contiguous memory regions, each of size  $\ell$  (totalling  $k \cdot \ell$  choices) that can be searched with  $k$  cache faults only. In previous work, the possible memory regions formed a partition of the total memory available. Dietzfelbinger and Weidling noticed that space utilisation can be improved if memory regions are allowed to overlap imperfectly, i.e. the choices for each object are the union of  $k$  intervals of length  $\ell$ , where the start of an interval is not necessarily a multiple of  $\ell$ . We confirm the experimental results by deriving corresponding thresholds rigorously.

Our main tool is a theorem due to Lelarge, Leconte and Massoulié which uses methods from statistical physics. It allows to derive the thresholds from properties of infinite trees that the underlying hypergraphs converge to.