Brief Announcement: The 1-2-3-Toolkit for Building Your Own Balls-into-Bins Algorithm

Pierre Bertrand[1] and Christoph Lenzen[2]

[1] Ecole Normale Suprieure Cachan
Avenue du prsident Wilson, 94230 Cachan
pierre.bertrand@ens-cachan.fr
[2] MPI for Informatics
Campus E1 4, 66123 Saarbrcken
clenzen@mpi-inf.mpg.de

Abstract. We examine a generic class of simple distributed balls-into-bins algorithms and compute accurate estimates of the remaining balls and the load distribution after each round. Each algorithm is classified by (i) the load that bins accept in a given round and (ii) the number of messages each ball sends in a given round. Our algorithms employ a novel ranking mechanism resulting in notable improvements. Simulations independently verify our results and their high accuracy.

1 Problem and Algorithm

Consider a distributed system of n anonymous balls and n anonymous bins, each having access to (perfect) randomization. Communication proceeds in synchronous rounds, each of which consists of the following steps.

1. Balls perform computations and send messages to bins.
2. Bins receive them, perform computations, and respond to received messages.
3. Each ball may commit to a bin, inform it, and terminate.

The main goals are to minimize the maximal number of balls committing to the same bin, the number of rounds, and the number of messages. This fundamental load balancing task has a wide range of applications, cf. [5].

Today, we understand the asymptotics of this problem very well [3,4,6]. However, lower and upper bounds have in common that they are not very precise. Arguably, with running time bounds like, e.g., $\Theta(\log\log n/\log\log\log n)$ or $\log^* n + \mathcal{O}(1)$, the involved constants are essential. In this work, we provide a simple, yet accurate analysis of a general class of algorithms. We introduce a novel *ranking* mechanism, resulting in superior performance.

Concretely, in each round $i \in \mathbb{N}$, the following steps are executed.

1. Each ball sends $M_i \in \mathbb{N}$ messages to uniformly independently random (u.i.r.) bins. These messages carry ranks $1, \dots, M_i$.
2. A bin of current load ℓ responds to (up to) $L_i - \ell$ balls, where smaller ranks are preferred. Ties are broken by choosing u.i.r.
3. Each ball that receives a response commits to the responding bin to which it sent the message of smallest rank.

2 Techniques and Results

Applying Chernoff's bound, it is not hard to show that the number of bins with a given load and the number of remaining balls are strongly concentrated around the expected values. With high probability, the error resulting from assuming that these expected values are matched exactly is hence negligible. Using this argument (and the union bound) repeatedly, we can infer that it suffices to compute expected values, approximating the true distribution by expected values. We complement the derived analytical results by simulations, confirming that the deviations are indeed very small. Moreover, we use the simulations to compare to other algorithms from the literature.

goal	rounds	max. load	messages	exp. fraction of balls left	L	M
small load	3	2	$< 5.5n$	$< 6 \cdot 10^{-7}$	$(2, 2, 2)$	$(2, 5, 5)$
few rounds	2	3	$< 5.5n$	$< 6 \cdot 10^{-10}$	$(2, 3)$	$(2, 5)$
few messages	3	3	$< 3.5n$	$< 5 \cdot 10^{-8}$	$(2, 3, 3)$	$(1, 2, 2)$
safe termination	3	3	$< 3.85n$	$< 6 \cdot 10^{-19}$	$(2, 2, 3)$	$(1, 4, 5)$

Table 1. Evaluated specific scenarios (analytical and simulation results match).

Our simulations also show that the proposed algorithms compare favorably with all previous ones from the literature. The full paper, comprising a discussion of related work, the derivation of the analytical bounds, and details on the simulation results, is available on arxiv [2]. The used code can be found online [1].
Acknowledgements. Christoph Lenzen has been supported by the Deutsche Forschungsgemeinschaft (DFG, reference number Le 3107/1-1).

References

1. Bertrand, P.: Python scripts used for simulations and computations (2014), http://people.mpi-inf.mpg.de/∼clenzen/babi/
2. Bertrand, P., Lenzen, C.: The 1-2-3-Toolkit for Building Your Own Balls-into-Bins Algorithm. Computing Research Repository abs/1407.8433 (2014)
3. Even, G., Medina, M.: Parallel Randomized Load Balancing: A Lower Bound for a More General Model. In: Proc. 36th Conference on Theory and Practice of Computer Science (SOFSEM). pp. 358–369 (2010)
4. Lenzen, C., Wattenhofer, R.: Tight Bounds for Parallel Randomized Load Balancing: Extended Abstract. In: Proc. 43rd Symposium on Theory of Computing (STOC). pp. 11–20 (2011)
5. Mitzenmacher, M., Richa, A., Sitaraman, R.: Handbook of Randomized Computing, vol. 1, chap. The Power of Two Random Choices: A Survey of the Techniques and Results, pp. 255–312. Kluwer Academic Publishers, Dordrecht (2001)
6. Stemann, V.: Parallel Balanced Allocations. In: Proc. 8th Symposium on Parallel Algorithms and Architectures (SPAA). pp. 261–269 (1996)