



HAL
open science

Brief Announcement: Tight Space Bounds for Memoryless Anonymous Consensus

Leqi Zhu

► **To cite this version:**

Leqi Zhu. Brief Announcement: Tight Space Bounds for Memoryless Anonymous Consensus. DISC 2015, Toshimitsu Masuzawa; Koichi Wada, Oct 2015, Tokyo, Japan. hal-01207887

HAL Id: hal-01207887

<https://hal.science/hal-01207887>

Submitted on 1 Oct 2015

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Brief Announcement: Tight Space Bounds for Memoryless Anonymous Consensus

Leqi Zhu*

University of Toronto, Toronto, ON M5S 3G4, Canada
leqzhu@cs.toronto.edu

Introduction. Tight $\Theta(n^2)$ bounds are known for the total step complexity of randomized algorithms for n -process consensus from registers [1]. However, there is a large gap between the best known space lower bound of $\Omega(\sqrt{n})$ registers [2] and the $\Theta(n)$ space complexity of the best existing algorithms. We prove matching upper and lower bounds of n for the space complexity of nondeterministic solo-terminating consensus in a restricted computational model. Specifically, we consider an asynchronous system with n anonymous processes, which communicate through an m -component multi-writer snapshot. Each process alternately performs SCAN and UPDATE. The location and value of each UPDATE can depend only on the result of the preceding SCAN by the same process. The only exception is the first UPDATE by each process, which can also depend on its input. We call algorithms designed for this model *memoryless*.

Lower Bound. Let C be any configuration of an n -process memoryless anonymous consensus algorithm. A process is *free* in C if it has already taken at least one step and is poised to perform a SCAN in C . Since processes are anonymous and memoryless, an adversary can cause any two free processes in C to behave identically if they both see the same result on their next SCAN. We say C is *solo v -deciding for free processes* if there is a solo execution by a free process starting from C that decides v . We say C is (P, q) -*bivalent* if P is a set of processes covering distinct components and $q \notin P$ is a process covering a component such that $C\beta_P$ is solo v -deciding for free processes and $C\beta_q$ is solo \bar{v} -deciding for free processes, for some $v \in \{0, 1\}$, where β_q is an UPDATE by q and β_P consists of one UPDATE by each process in P .

Given a $(Z \cup \{p\}, q)$ -bivalent configuration C , consider the longest prefix α' of q 's solo terminating execution α from C such that $C\alpha'\beta_{Z \cup \{p\}}$ is solo v -deciding for free processes. The next step δ by q in α after α' must be an UPDATE to a component not covered by $Z \cup \{p\}$. Running p from $C\alpha'\beta_P$ until it is about to perform an UPDATE yields a $(\{p\}, q)$ -bivalent configuration C' . Since C' is also $(\{q\}, p)$ -bivalent, the same argument implies that there is an execution from C' in which q takes at least two steps, such that the resulting configuration C'' is $(\{q\}, p)$ -bivalent and every process in Z is free. Note that at least $|Z|+2$ different components were updated in the execution from C to C'' (via δ and $\beta_{Z \cup \{p\}}$).

* I would like to thank my advisor, Dr. Faith Ellen, and David Solymosi. This work was supported by the Natural Sciences and Engineering Council of Canada.

We show that, if C is $(\{p\}, q)$ -bivalent and Z is a set of free processes in C , there is an execution γ from C in which p and q take at least two steps, such that $C\gamma$ is $(Z \cup \{p\}, q)$ -bivalent. The proof is by induction on $|Z|$. The base case, when $Z = \emptyset$, holds by the previous paragraph. Fix any $Z' \subset Z$ with $|Z'| = |Z| - 1$. By induction, there is an execution γ_1 in which p and q take at least two steps, such that $C\gamma_1$ is $(Z' \cup \{p\}, q)$ -bivalent. By the preceding paragraph, there is an execution γ_2 from $C\gamma_1$ in which at least $|Z'| + 2$ components have been updated and p and q have taken at least 2 steps, such that $C\gamma_1\gamma_2$ is $(\{p\}, q)$ -bivalent and each process in Z' is free in $C\gamma_1\gamma_2$. By induction, there is an execution γ_3 such that $C\gamma_1\gamma_2\gamma_3$ is $(Z' \cup \{p\}, q)$ -bivalent. Among the components updated in $\gamma_2\gamma_3$, there is at least one component j which is not covered by $Z' \cup \{p\}$ in $C\gamma_1\gamma_2\gamma_3$. Let $z' \in Z' \cup \{p, q\}$ be the last process to UPDATE component j prior to $C\gamma_1\gamma_2\gamma_3$, and let σ' be the SCAN by z' before this UPDATE. Note that every process in Z' is free immediately before σ' . Modifying the execution $\gamma_1\gamma_2\gamma_3$ to let the remaining free process in $Z - Z'$ perform its SCAN immediately after σ' gives a $(Z \cup \{p\}, q)$ -bivalent configuration.

To obtain the lower bound, construct a $(\{p\}, q)$ -bivalent configuration having a set Z of $n - 2$ free processes. Apply the previous argument to get a $(Z \cup \{p\}, q)$ -bivalent configuration. Running q until it is about to UPDATE a component not covered by $Z \cup \{p\}$ gives a configuration with n components covered. This proof method, which uses induction on the number of free processes to build larger coverings, seems applicable in the general case. In fact, we recently used this method to give a different, much simpler proof of the $\Omega(\sqrt{n})$ lower bound in [2].

Upper Bound. We describe an n -process memoryless anonymous obstruction-free consensus algorithm using an n -component multi-writer snapshot, matching the lower bound. The algorithm can also be made randomized wait-free by [3].

Intuitively, 0 and 1 are competing to complete *laps*. If v gets a substantial lead on \bar{v} , then v is decided. Initially, each component contains $(0, 0)$. If a process with input x sees this initial state in a SCAN, it updates component 1 with (\bar{x}, x) . Otherwise, it determines the laps, ℓ_0 and ℓ_1 , of 0 and 1 by finding the largest values in the first and second entries of the components returned by its SCAN. If some component is not (ℓ_0, ℓ_1) , then it updates the first such component with (ℓ_0, ℓ_1) . So, suppose all components are the same. If value v is ahead of value \bar{v} by at least 2 laps, for some $v \in \{0, 1\}$, then it decides v . If not, it increments the larger of ℓ_0 and ℓ_1 (breaking ties in favour of ℓ_0) and updates component 1 with (ℓ_0, ℓ_1) . This is repeated until the process decides.

References

1. Hagit Attiya and Keren Censor. Tight bounds for asynchronous randomized consensus. *J. ACM*, 55(5):20, 2008.
2. Faith Fich, Maurice Herlihy, and Nir Shavit. On the space complexity of randomized synchronization. *J. ACM*, 45(5):843–862, 1998.
3. George Giakkoupis, Maryam Helmi, Lisa Higham, and Philipp Woelfel. An $O(\sqrt{n})$ space bound for obstruction-free leader election. In *27th DISC*, pages 46–60, 2013.