Extension-Based Proofs

Faith Ellen
University of Toronto

THEOREM There is no wait-free algorithm to solve consensus among $n \geq 2$ processes in an asynchronous system where processes communicate using registers.

[Chor, Israeli & Li 1987, Loui & Abu Amara 1987, Abrahamson 1988] consensus

every process p_i has an input value x_i and, if it doesn't crash, must output a value y_i such that the following properties hold:

validity: $y_i \in \{x_1, \ldots, x_n\}$ and

agreement: all output values are the same.

wait-free = every process terminates within a finite number of steps, even if other processes crash

THEOREM There is no wait-free algorithm to solve consensus among $n \geq 2$ processes in an asynchronous system where processes communicate using registers.

LEMMA 1 Every consensus algorithm has a bivalent initial configuration.

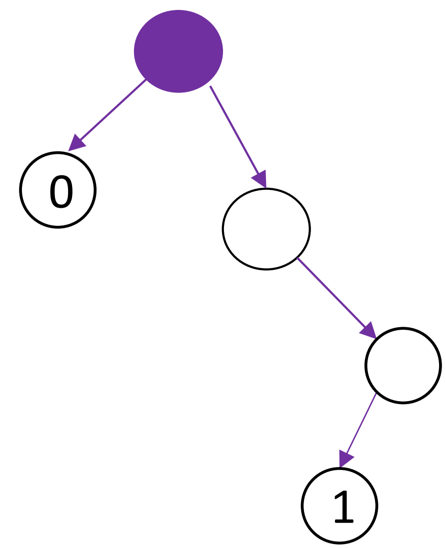

THEOREM There is no wait-free algorithm to solve consensus among $n \geq 2$ processes in an asynchronous system where processes communicate using registers.

LEMMA 1 Every consensus algorithm has a bivalent initial configuration.

LEMMA 2 From every bivalent configuration, there is a step that leads to a bivalent configuration.

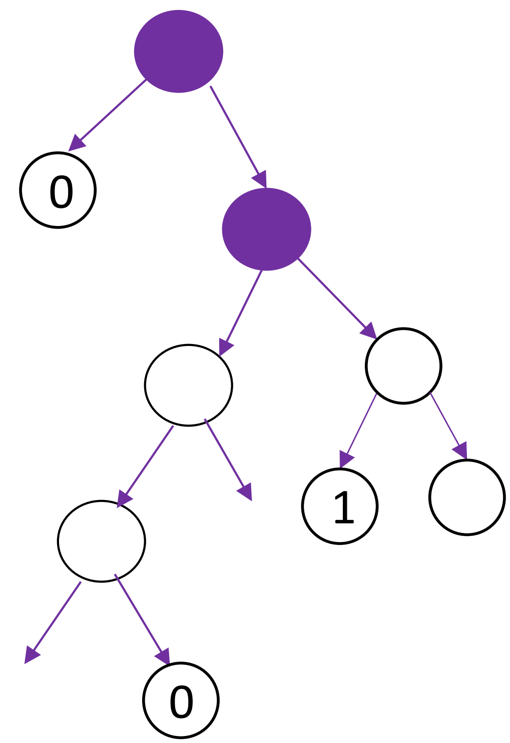

THEOREM There is no wait-free algorithm to solve consensus among $n \geq 2$ processes in an asynchronous system where processes communicate using registers.

LEMMA 1 Every consensus algorithm has a bivalent initial configuration.

LEMMA 2 From every bivalent configuration, there is a step that leads to a bivalent configuration.

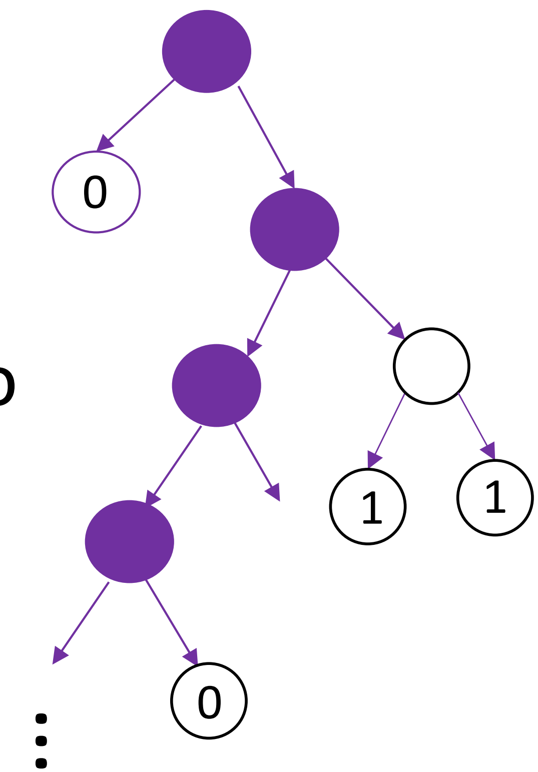

This implies there is an infinite execution, consisting of only bivalent configurations, violating wait-freedom.

k-set agreement

every process p_i has an input value x_i and,

if it doesn't crash, must output a value y_i

such that the following properties hold:

validity: $y_i \in \{x_1, \ldots, x_n\}$ and

agreement: at most k different values are output.

1-set agreement = consensus

THEOREM There is no wait-free algorithm to solve k-set agreement among $n > k \geq 2$ processes in an asynchronous system where processes communicate using registers.

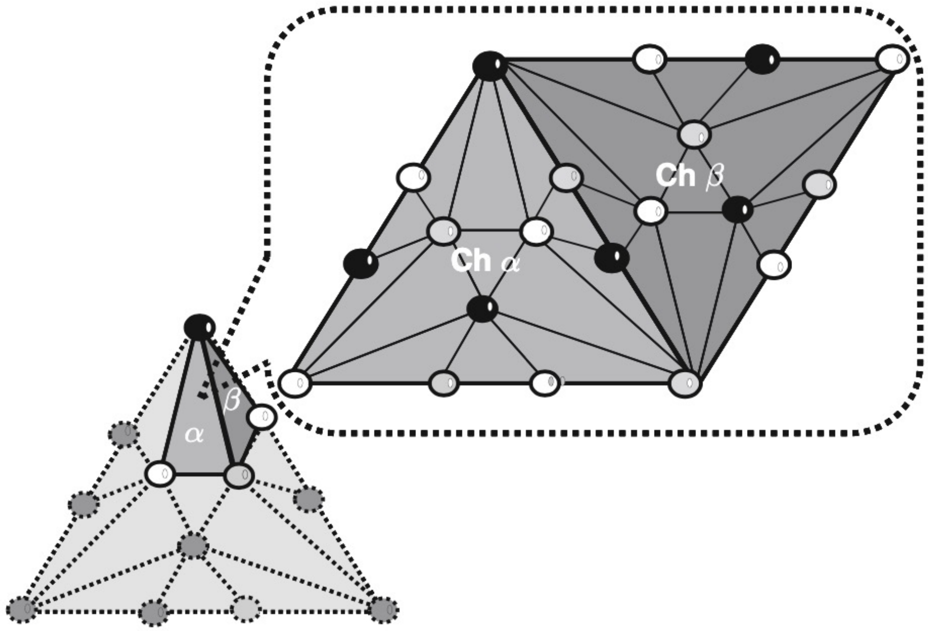

[Borowsky & Gafni, Herlihy & Shavit, Saks & Zaharoglu, 1993]

Alistarh, Aspnes, Ellen, Gelashvili, Zhu
STOC 2019, PODC 2020, SICOMP 2023

- Definition of extension-based proof
- There is no extension-based proof of the impossibility of a wait-free algorithm to solve k-set agreement among $n > k \geq 2$ processes in an asynchronous system.

Pitassi, Beame, Impagliazzo Comput. Complex. 1993

- There is no proof of the pigeon-hole principle using relativized bounded arithmetic.

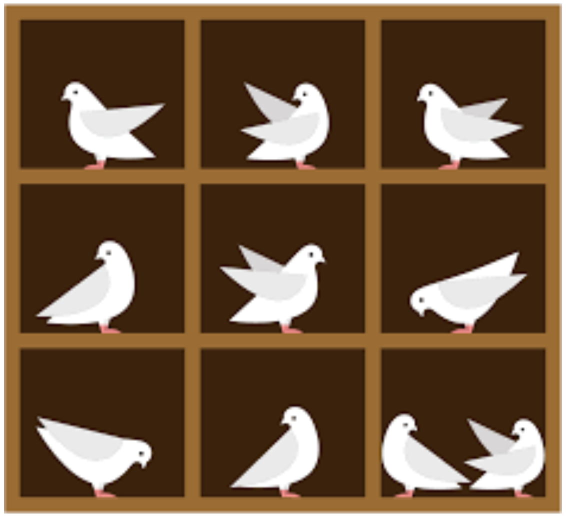

Extension-Based Proof

A sequence of interactions between a prover and an algorithm, divided into phases.

Initially, the prover has reached the initial configurations of the algorithm.

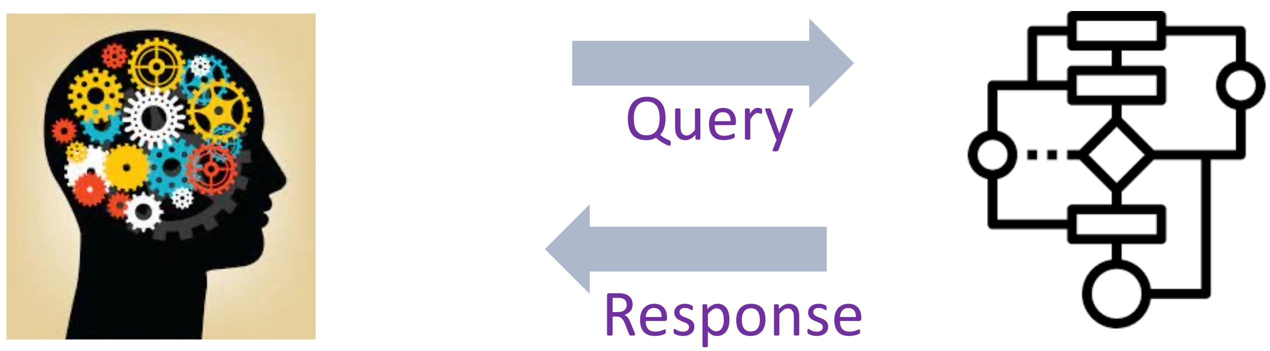

Extension-based proof

- The prover may ask a single-step query by choosing a configuration C it has reached and a process p that hasn't terminated in C.
- The algorithm responds with the configuration C' resulting from p taking one step from C. Now the prover has reached C'.

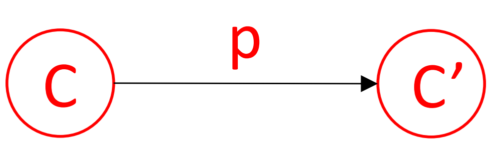

Extension-based proof

The prover wins (shows that the algorithm is incorrect) if the algorithm responds with a configuration in which the outputs of the processes violate the specifications.

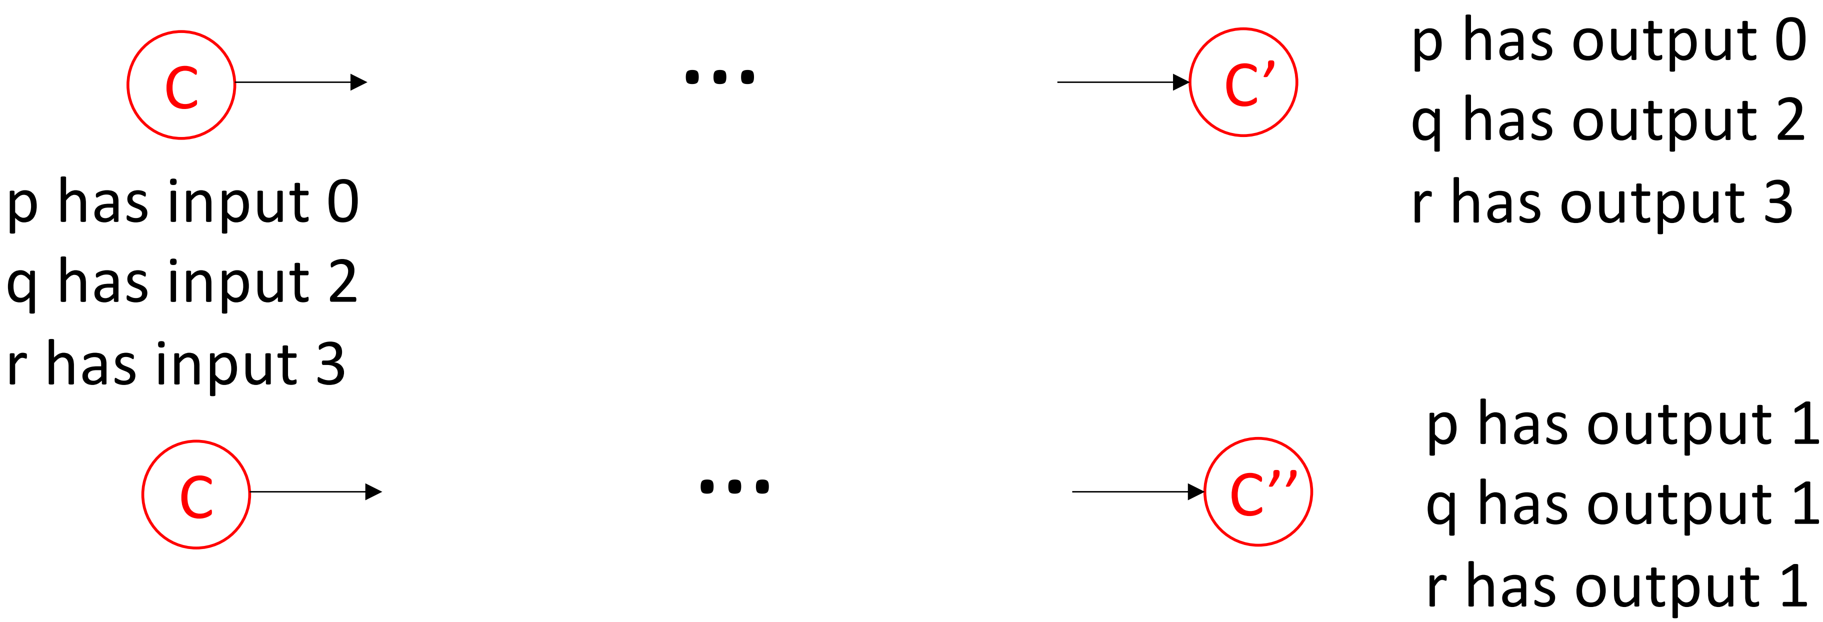

Extension-based proof

A chain of queries is a finite or infinite sequence of single-step queries

$(C_0, p_0), (C_1, p_1), \ldots,$

where C_{i+1} is the configuration that results when p_i takes 1 step from C_i, for each $i \geq 0$.

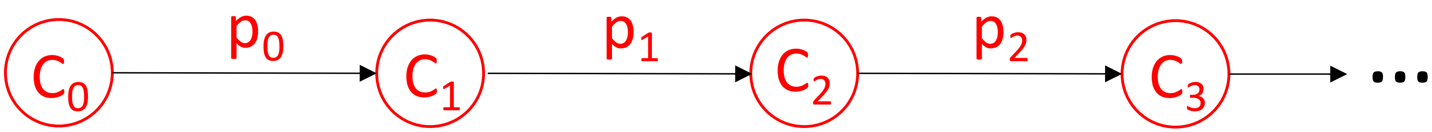

If the prover constructs an infinite chain of queries, it wins, since the algorithm is not wait-free.

Extension-based proof

The prover may make an output query (C, Q, y), where C is a configuration it has reached, Q is a set of processes, and y is a possible output value.

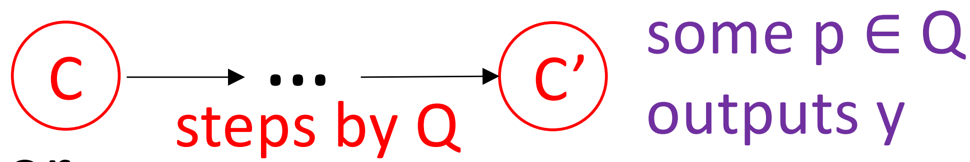

Then the algorithm must either

- respond with a finite sequence of steps by processes in Q such that, starting from C, one of them outputs the value y or
- say that no such sequence exists.

Extension-based proof

After making finitely many output queries and chains of queries in a phase without winning, the prover must

- choose a configuration C it first reached during this phase and
- start the next phase

In the next phase, the prover can only ask queries about configurations that are reachable from C.

Extension-based proof

The prover loses if all processes have terminated in the configuration chosen at the end of some phase.

Extension-based proof

The prover wins if:

- it asks an infinite chain of queries or
- there are an infinite number of phases

because it has demonstrated that the algorithm is not wait-free.

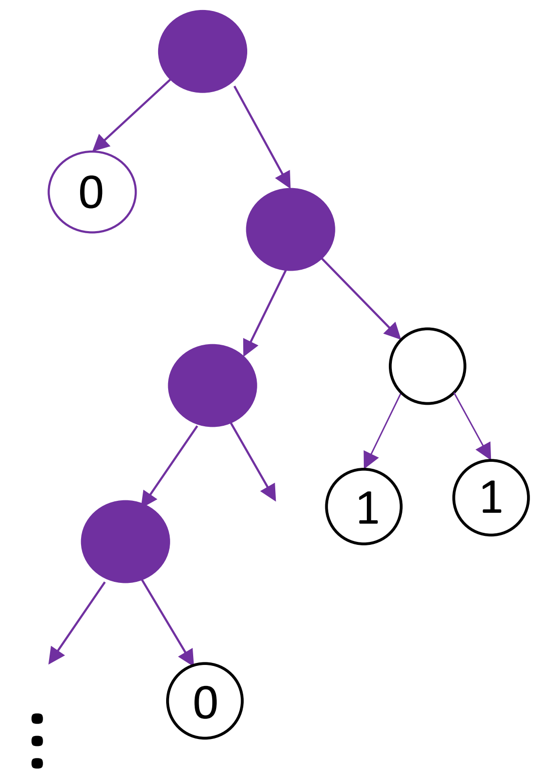

Alistarh, Aspnes, Ellen, Gelashvili, Zhu
STOC 2019, PODC 2020, SICOMP 2023

- Definition of extension-based proof
- There is no extension-based proof of the impossibility of a wait-free algorithm to solve k-set agreement among $n > k \geq 2$ processes in an asynchronous system.

Alistarh, Ellen, Rybicki
SIROCCO 2021, SICOMP 2023

- There is no extension-based proof of the impossibility of a wait-free algorithm to solve approximate agreement among $n > 2$ processes on a cycle of length 4 in an asynchronous system.

Approximate Agreement on a Graph G =(V,E)

Each process p_i has an input $x_i \in V$ and,
if it does not crash, must output $y_i \in V$
such that the following properties hold:
shortest path validity: every output y_i lies on
a shortest path between two inputs and
approximate agreement: the set of outputs are
the nodes of a clique in G

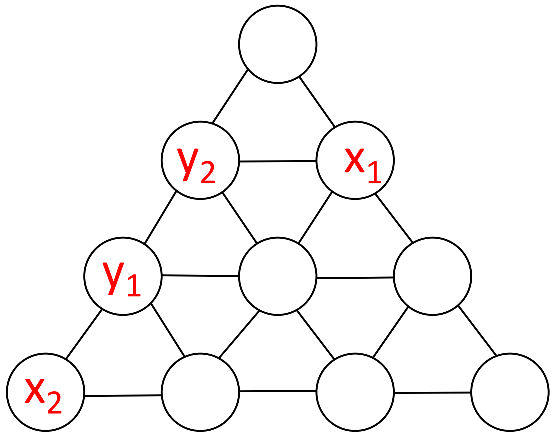

Liu
OPODIS 2022

- There is no extension-based proof of the impossibility of a wait-free algorithm to solve approximate agreement among $n > 2$ processes on any connected graph in an asynchronous system.

If problem $\mathcal{T}$ reduces to problem $\mathcal{S}$
and $\mathcal{T}$ is impossible to solve,
then $\mathcal{S}$ is impossible to solve.

If problem $\mathcal{T}$ reduces to problem $\mathcal{S}$
and $\mathcal{T}$ is impossible to solve,
then $\mathcal{S}$ is impossible to solve.

If problem $\mathcal{T}$ reduces to problem $\mathcal{S}$
and there is an extension-based proof that
$\mathcal{T}$ is impossible to solve,
then there is an extension-based proof that
$\mathcal{S}$ is impossible to solve.

Brusse, Ellen
PODC 2021

If problem $\mathcal{T}$ reduces* to problem $\mathcal{S}$
and there is an augmented extension-based proof that
that $\mathcal{T}$ is impossible to solve,
then there is an augmented extension-based proof
that $\mathcal{S}$ is impossible to solve.

* for a large, natural class of reductions

Our Class of Reductions

Given inputs $x_1,\ldots,x_n$

the n processes first solve the problem $\mathcal{R}_1$,

then solve the problem $\mathcal{S}$, and

finally solve the problem $\mathcal{R}_2$,

where $\mathcal{R}_1$ and $\mathcal{R}_2$ are solvable problems.

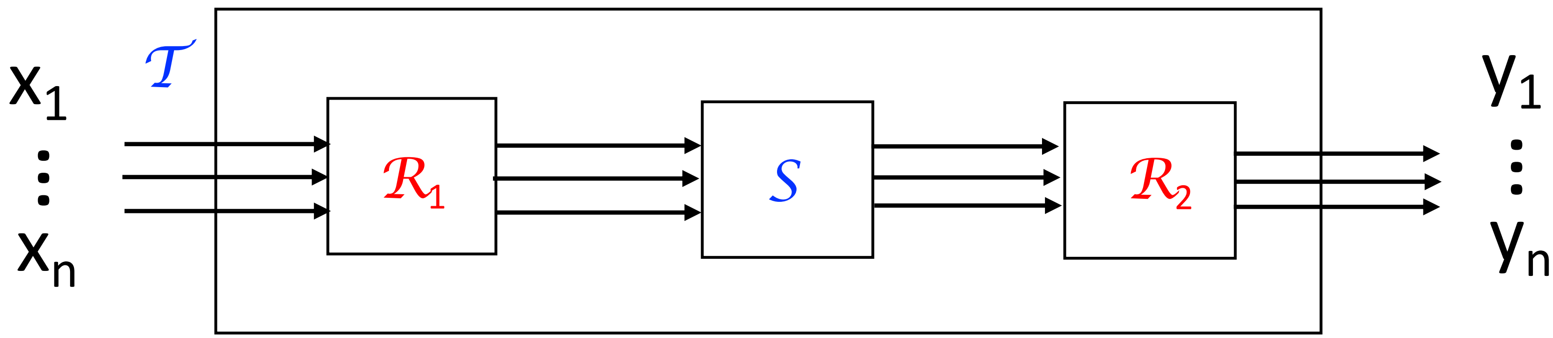

Augmented Extension-based proof

The prover may make an assignment query (C, Q, f) where C is a configuration it has reached,
Q is a set of processes, and
f is an assignment from $Q' \subseteq Q$ to possible output values.

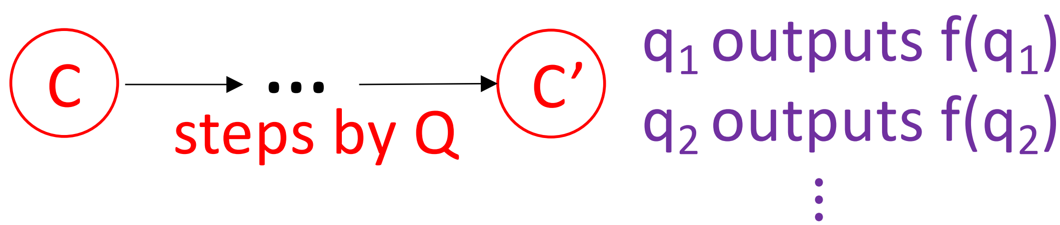

Then the algorithm must either

- respond with a finite sequence of steps by processes in Q such that, starting from C, every process $q_i \in Q'$ of outputs the value $f(q_i)$ or
- say that no such sequence exists.

Augmented Extension-based proof

An output query (C, Q, y) can be simulated by
|Q| assignment queries (C, Q, f_P), where
f_P:{p} → {y} assigns the output value y to process $p \in Q$.

Thus augmented extension-based proofs are at least as powerful as extension-based proofs.

THEOREM There is no extension-based proof of the impossibility of a wait-free algorithm to solve k-set agreement among $n > k \geq 2$ processes in an asynchronous system.

THEOREM There is no augmented extension-based proof of the impossibility of a wait-free algorithm to solve k-set agreement among $n > k \geq 2$ processes in an asynchronous system.

THEOREM **If** $\mathcal{R}_2 \circ \mathcal{S} \circ \mathcal{R}_1$ is a reduction from problem $\mathcal{T}$ to problem $\mathcal{S}$

and there is an augmented extension-based proof that
$\mathcal{T}$ is impossible to solve,

then there is an augmented extension-based proof that
$\mathcal{S}$ is impossible to solve.

There are reductions from k-leader election and k-test-and-set to k-set agreement.

[Borowsky & Gafni, 1993]

Hence, there are no augmented extension-based proofs of the impossibility of wait-free algorithms to solve k-leader election and k-test-and-set among $n > k \geq 2$ processes in an asynchronous system.

THEOREM There are no anonymous wait-free algorithms to solve weak symmetry breaking or (2n-2)-renaming among $n \geq 2$ processes in an asynchronous shared memory system where processes communicate using registers.

A algorithm is anonymous if the steps taken by a process do not depend on its identifier.

[Castaneda & Rajsbaum 2010]

THEOREM **If** $\mathcal{R}_2 \circ \mathcal{S} \circ \mathcal{R}_1$ is an anonymous reduction from problem $\mathcal{T}$ to problem $\mathcal{S}$

and there is an augmented extension-based proof that $\mathcal{T}$ is impossible to solve anonymously,

then there is an augmented extension-based proof that $\mathcal{S}$ is impossible to solve anonymously.

- There is no augmented extension-based proof that k-set agreement is impossible to solve anonymously.
- There are anonymous reductions from weak symmetry breaking and (2n-2)-renaming to (n-1)-set agreement.

Hence there are no augmented extension-based proofs of the impossibility of solving weak symmetry breaking and (2n-2)-renaming anonymously.

[Herlihy, Kozlov & Rajsbaum 2013]

Happy 60th Birthday, Toni