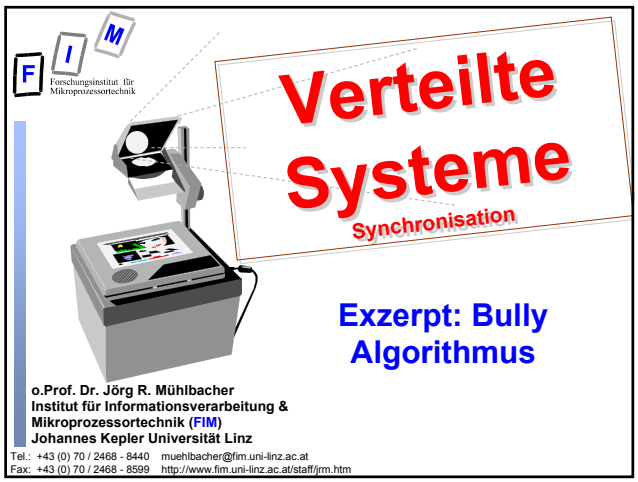

Election algorithms

- many distributed algorithms rely on the existence of a coordinator process
  - ➔ e.g.: mutual exclusion by centralized method
- fault tolerance?
  - ➔ if the coordinator crashes
    - » a new coordinator must be elected
    - » all the other process must agree on who the new coordinator is to be
- underlying assumptions:
  - ➔ processes do not know, which P_i are live or down
  - ➔ every process knows the process numbers of all the others

Election algorithms: The Bully Method

- bully = einschüchternd, Tyrann
  - ➔ Garcia - Molina, 1982
- most common and obvious method
- underlying principle:

  the biggest guy in town always wins
- other methods, not discussed here
  - ➔ring algorithm without a token

Election algorithms: The Bully Method

- a process P makes an election as follows
  - ➔ P sends an ELECTION message to all processes with higher process number
  - ➔ if no one responds, P wins and becomes coordinator, sending a "I am the COORDINATOR message" to all the processes
  - ➔ if one of the higher-ups answers, it takes over. P´s job is done
- if a process Q receives an election message from a P with lower number
  - ➔ it sends an OK indicating it is alive, and takes over: it makes an election
- if a P was down and comes back,
  - ➔ it makes an election

o. Univ.-Prof. Dr.Jörg R. Mühlbacher, Institut für Informationsverarbeitung und Mikroprozessortechnik (FIM),
Johannes Kepler Universität Linz, A-4040 Linz
Tel.: ++43 (0) 70 / 2468 - 8440, FAX: ++43 (0) 70 / 2468 - 8599
Email: muehlbacher@fim.uni-linz.ac.at, WWW: http://www.fim.uni-linz.ac.at/staff/jrm.htm