Data Networks
State – Part 2

Prof. Anja Feldmann, Ph.D.

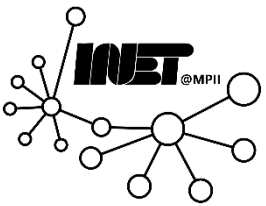

Design Principles

Goals:

- Identify, study common architectural components, protocol mechanisms, approaches do we find in network architectures?
- Synthesis: Big picture

Design Principles:

- Separation of data, control
- Hard state versus soft state
- Randomization
- Indirection
- Network virtualization / Overlays
- Resource sharing
- Design for scale

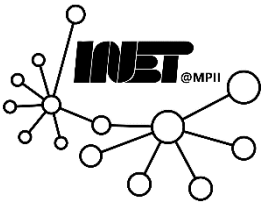

2: Maintaining network state

State: Information *stored* in network nodes by network protocols

- Hard-state
  - State installed by receiver via setup msg from sender
  - State removed by receiver via teardown msg from sender
- Soft-state
  - State installed by receiver via setup (trigger) msg from sender (typically, an endpoint)

    Sender sends periodic refresh msg indicating receiver should continue to maintain state
  - State removed by receiver via timeout (absence of refresh msg from sender)
  - Default assumption: State becomes invalid unless refreshed

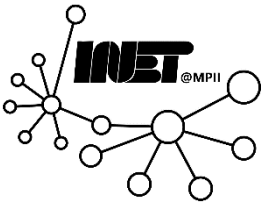

How do we model/evaluate?

Metrics

- Inconsistency ratio – fraction time participating nodes disagree
- Signaling overhead – average # of messages during session lifetime

- Measures of robustness?
  - Convergence time
  - Complexity?

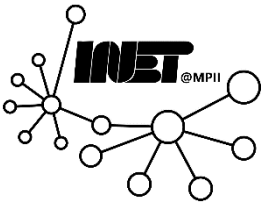

Single hop model

- Sender, receiver
- Single state variable
- State lifetime – exp(m)
- Updates – Poisson(l)
- Timers – exponentially distributed
  - Refresh - 1/T
  - State expiration – 1/X
- Link: Delay exp(1/D), loss prob. p

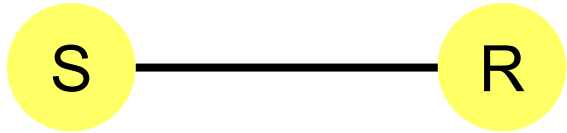

Transient Markov model

Model for SS

(Ji, Ge, Kurose and Towsley. A Comparison of Hard-state and Soft-state Signaling Protocols. SIGCOMM 2003)

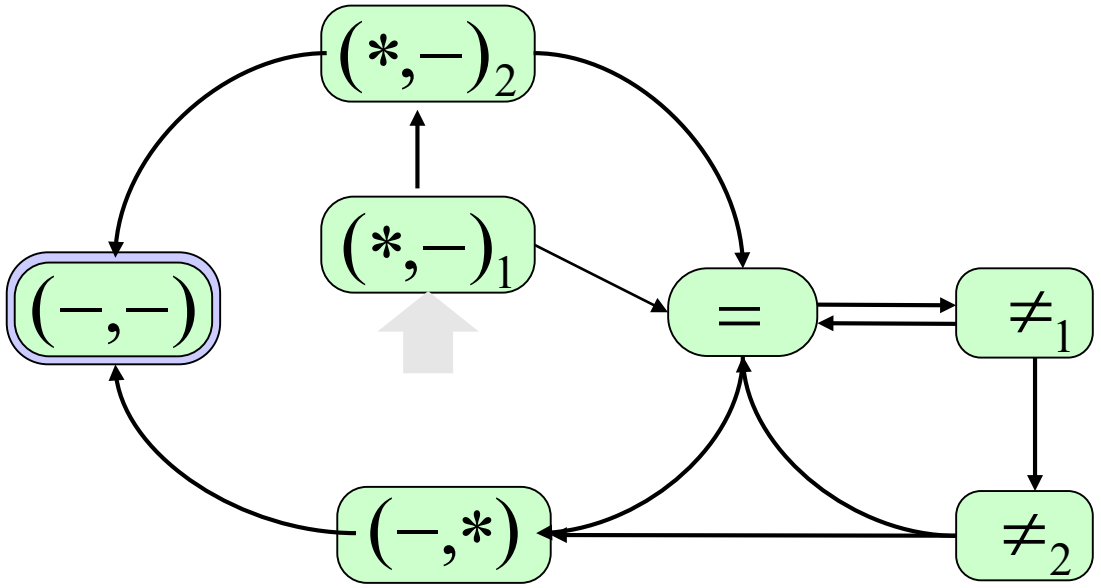

Model for SS

$(*,-)$ Signaling state generated at sdr, not installed at rcvr

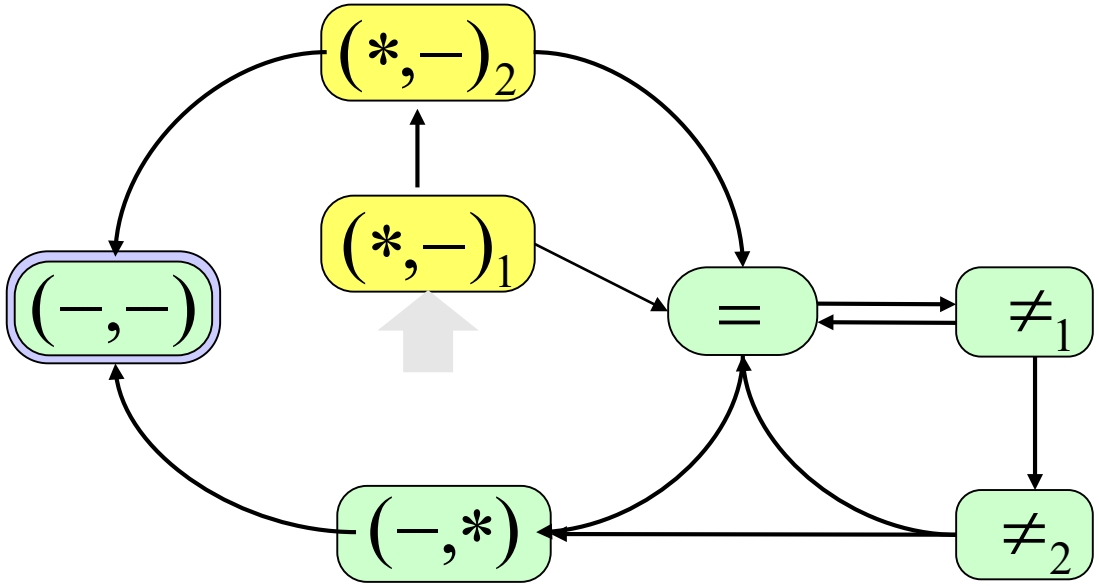

Model for SS

$(*,-)$ Signaling state generated at sdr, not installed at rcvr

$=$: Signaling state consistent at sdr/rcvr

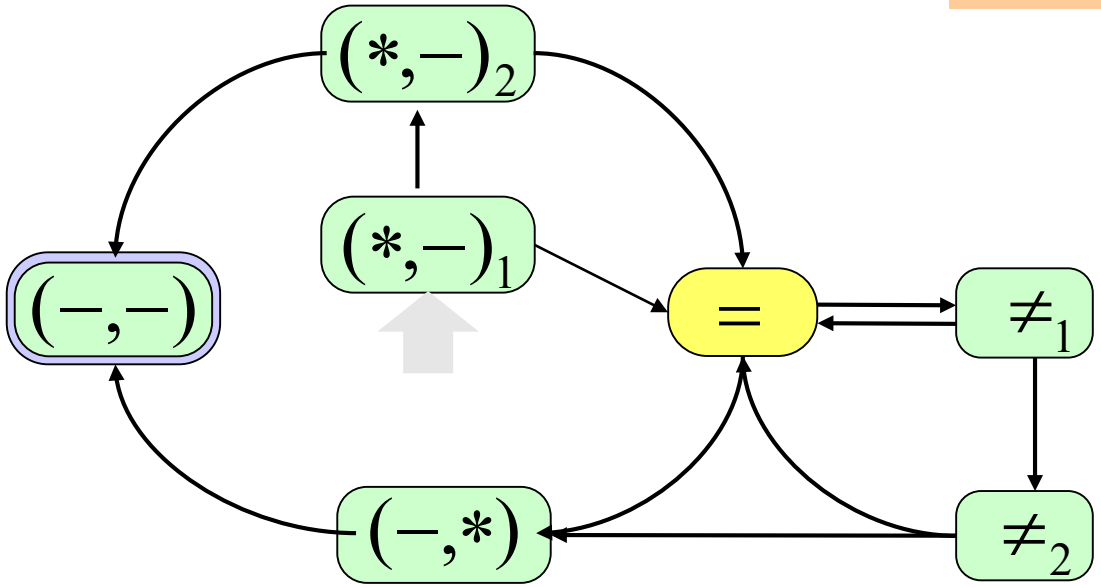

Model for SS

$(*,-)$ Signaling state generated at sdr, not installed at rcvr

$=$: Signaling state consistent at sdr/rcvr

$\neq$: Signaling state inconsistent at sdr/rcvr

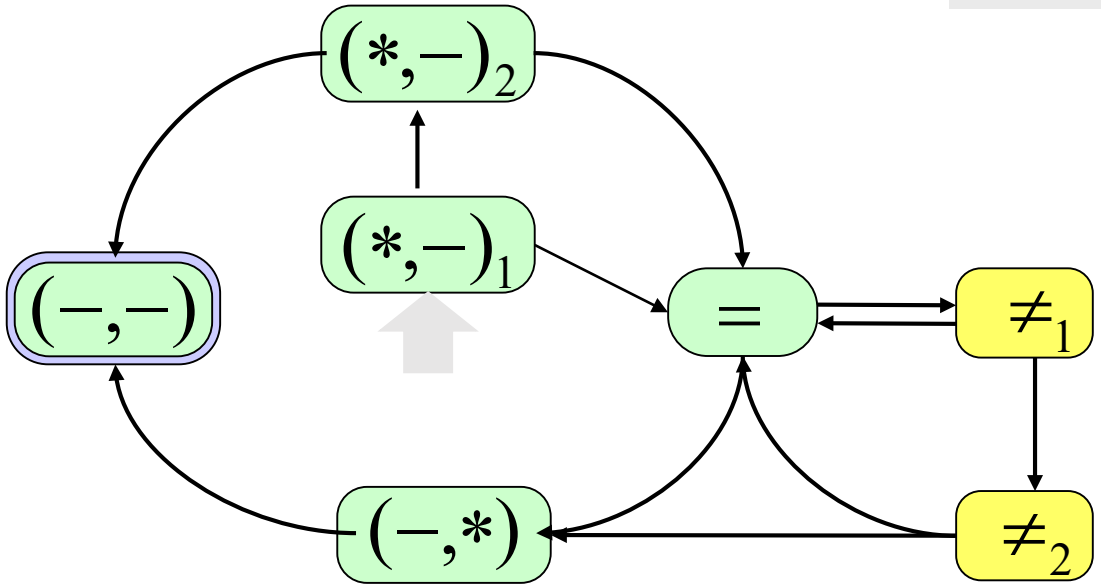

Model for SS

$(*,-)$ Signaling state generated at sdr, not installed at rcvr

$=$: Signaling state consistent at sdr/rcvr

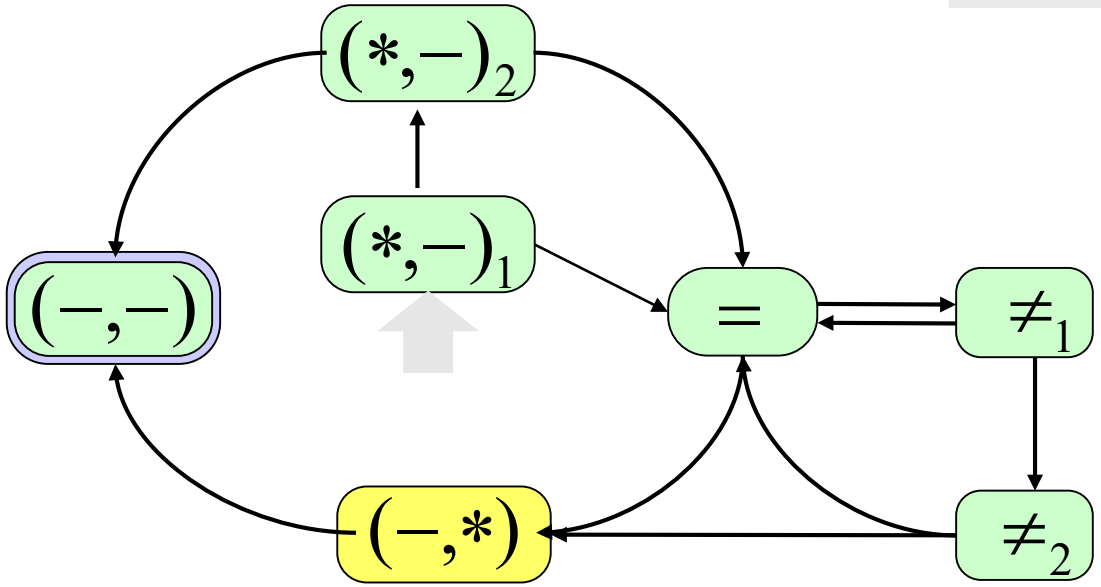

$\neq$: Signaling state inconsistent at sdr/rcvr

$(-,*)$ Signaling state removed at sender, present at receiver

Model for SS

(*,–) Signaling state generated at sdr, not installed at rcvr

= : Signaling state consistent at sdr/rcvr

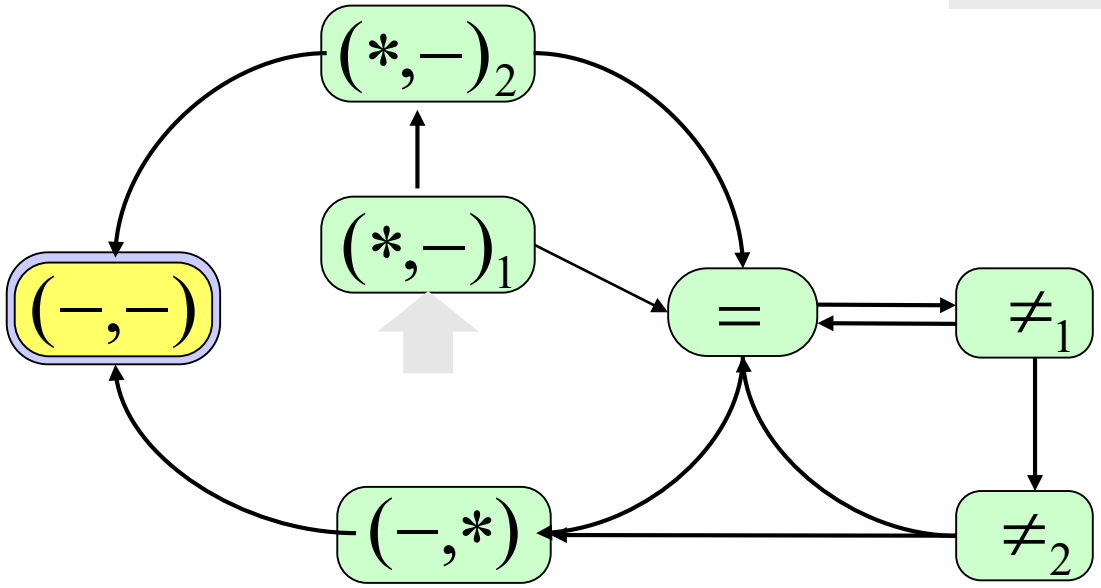

≠ : Signaling state inconsistent at sdr/rcvr

(–,–) : Signaling state removed at sdr/rcvr

(–,*) Signaling state removed at sender, present at receiver

Model for SS

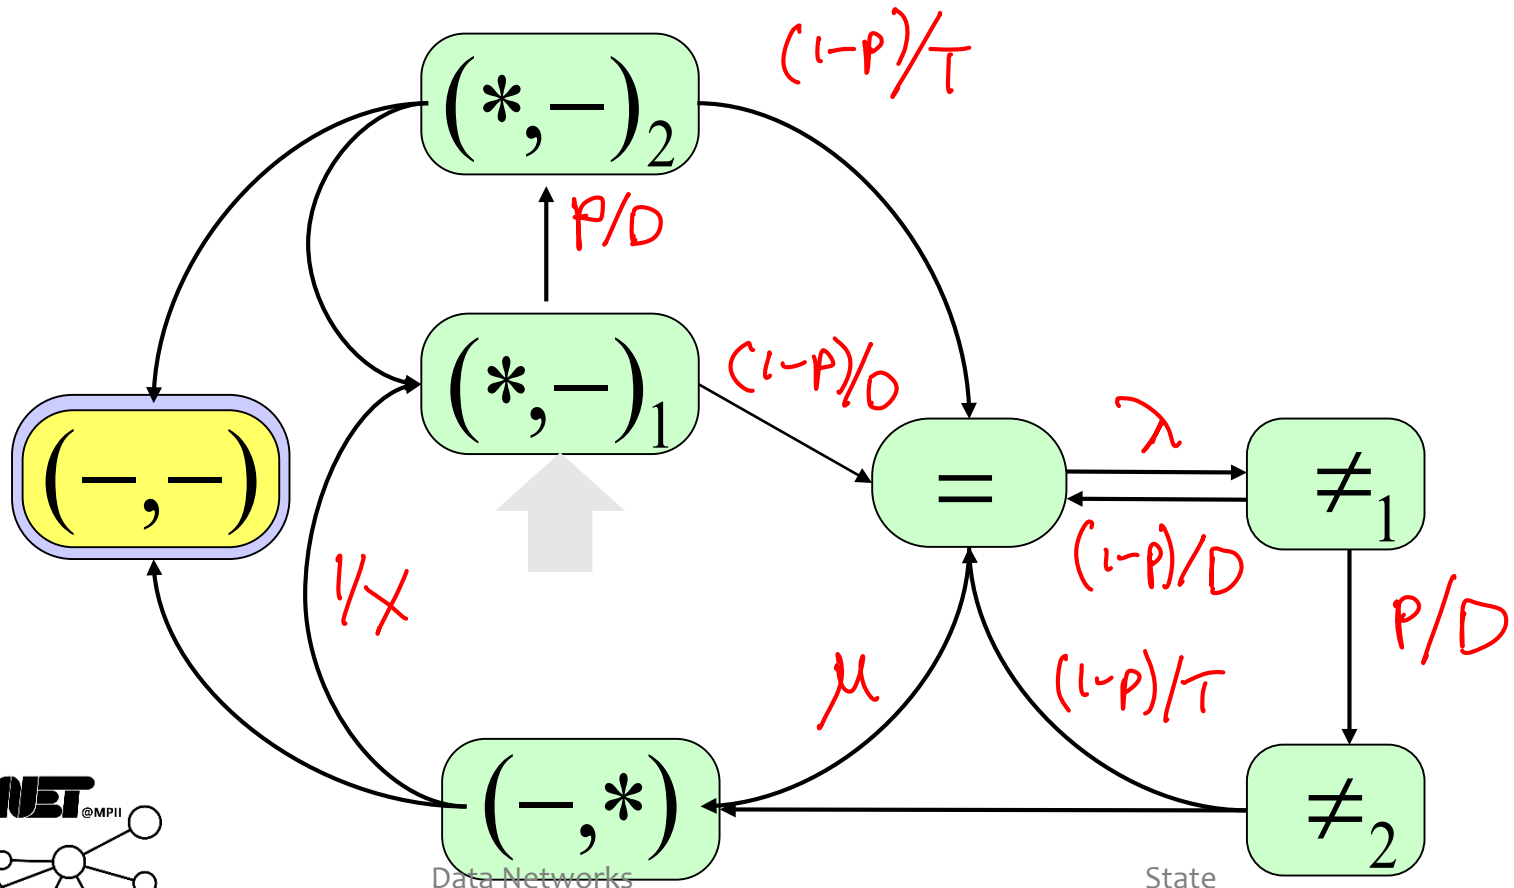

(*,–): Signaling state generated at sdr, not installed at rcvr

=: Signaling state consistent at sdr/rcvr

≠: Signaling state inconsistent at sdr/rcvr

(–,*): Signaling state removed at sender, present at receiver

(–,–): Signaling state removed at sdr/rcvr

Performance metrics (SS)

- Inconsistency ratio:

$$\delta = 1 - \pi =$$

- Signaling overhead

$$\Gamma = (1 + l + 1/T) / m$$

Parameter settings

- Mean lifetime: 30 min
- Refresh timer: T = 5sec
- State timer: X = 15 sec
- Update rate: 1/20sec
- Signal loss rate: 2 %

Motivated by Kazaa

Impact of state lifetime

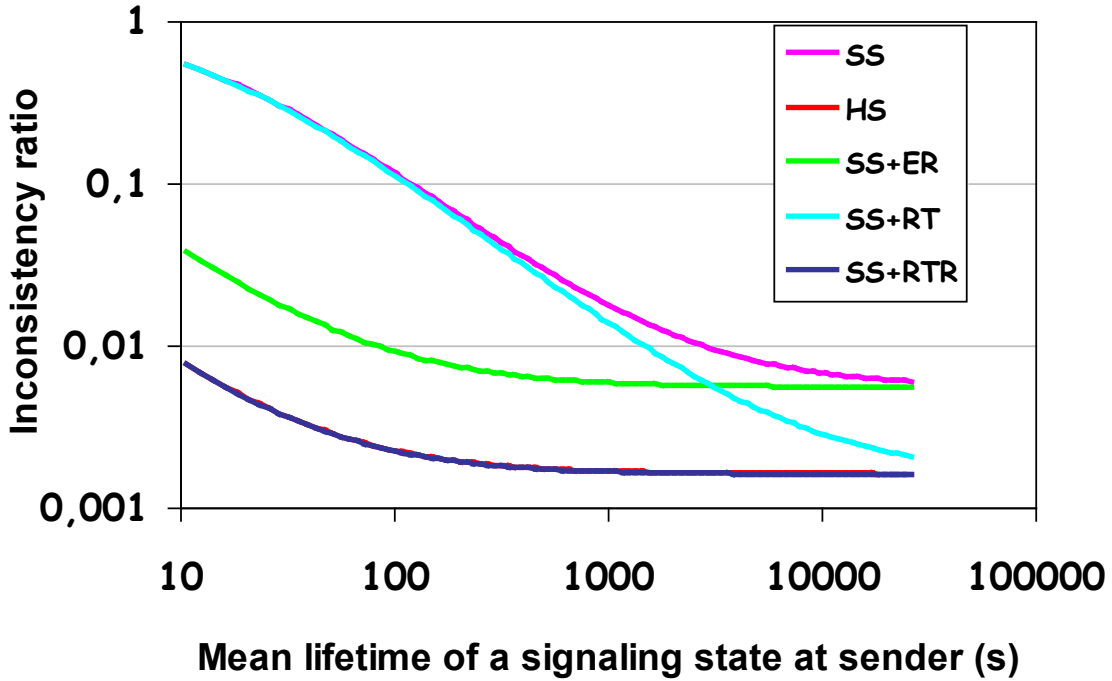

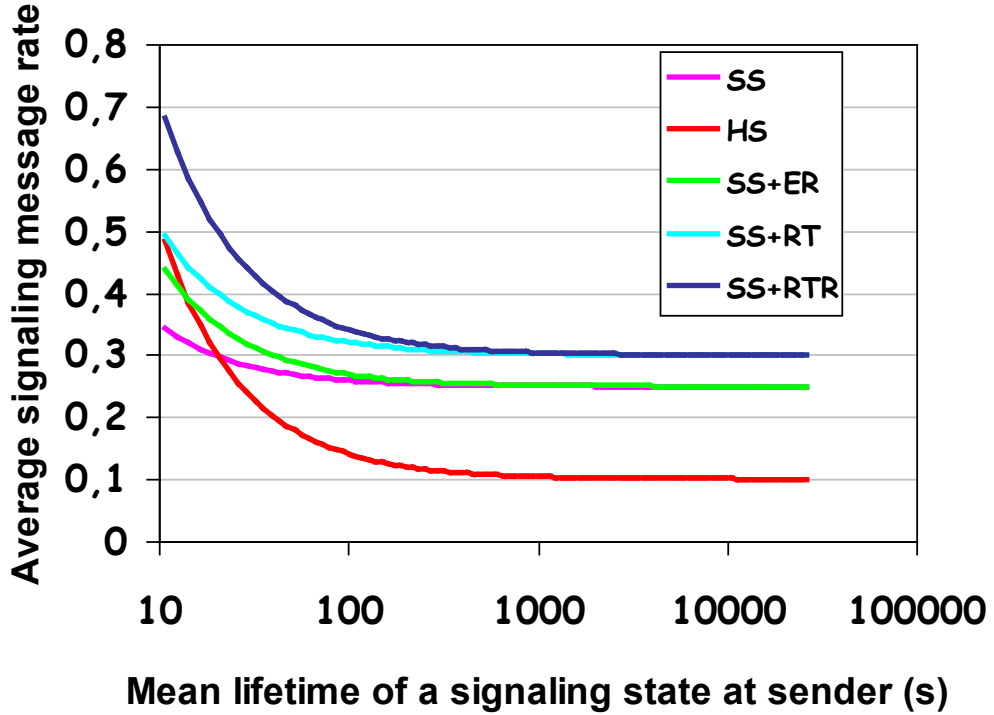

- Inconsistency, overhead decrease as state life-time increases
- Explicit removal improves consistency with little additional overhead

Soft-state: Setting timer values (Recall)

Q: How to set refresh/timeout timers

- State-timeout interval = n * refresh-interval-timeout
  - What value of n to choose?
- Will determine amount of signaling traffic, responsiveness to change
  - Small timers: Fast response to changes, more signaling
  - Long timers: Slow response to changes, less signaling
- Ultimately: Consequence of slow/fast response, msg loss probability will dictate appropriate timer values

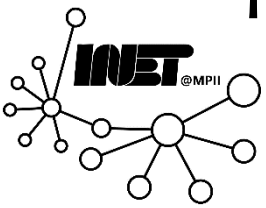

Impact of state timeout timer

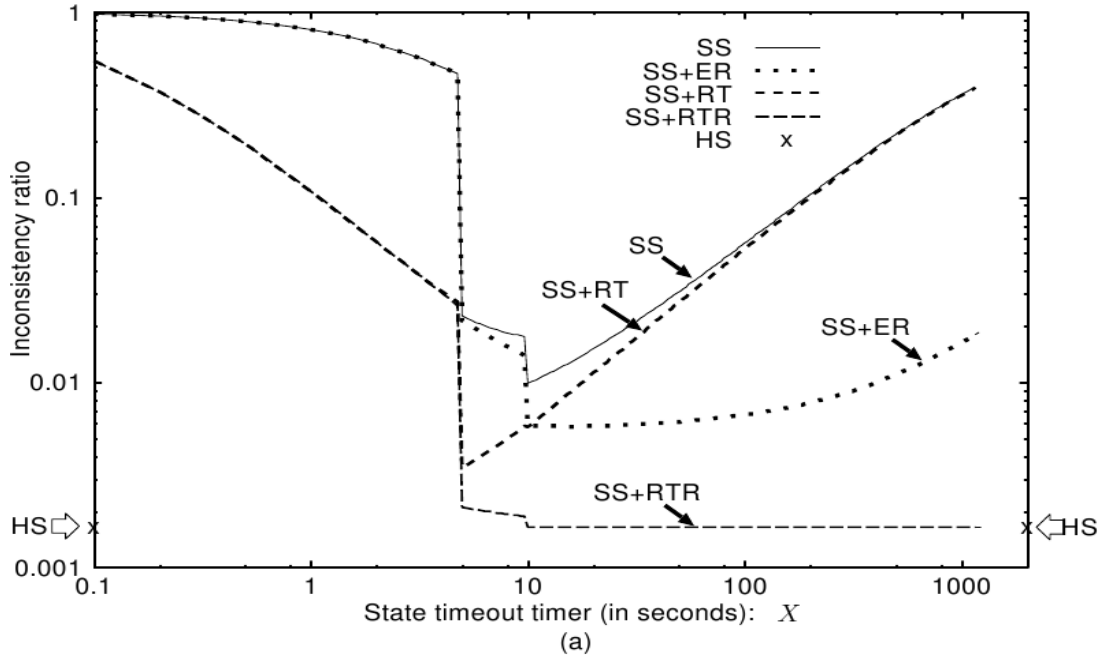

Recall: T=5s

- $X < T$: Inconsistency high (premature state removal)
- $X > 2T$: Increasing $X \Rightarrow$ increasing inconsistency for SS, SS+ER, SS+RT (due to orphan state)
- $X = 2T$: Sweet spot

Hard-state versus soft-state: Discussion

Q: Which is preferable and why?

Hard state:

- Better if message OH really high
- Potentially greater consistency
- System wide coupling -> difficult to analyze

Soft state:

- Robustness, shorter convergence times
- Implicit reliability
- Easier error recovery
- Easily decomposed -> simpler analysis