

Toward Understanding Congestion in Tor

*DC-area Anonymity, Privacy, and
Security Seminar*


January 24th, 2014



Rob Jansen
U.S. Naval Research Laboratory

*Joint with John Geddes, Chris Wacek, Micah Sherr, Paul Syverson

Tor for ~~Awesomeness~~ Anonymity



Tor is Slow!!! Research*

- PCTCP: Per-Circuit TCP-over-IPsec Transport for Anonymous Communication Overlay Networks (CCS '13)
- Reducing Latency in Tor Circuits with Unordered Delivery (FOCI '13)
- How Low Can You Go: Balancing Performance with Anonymity in Tor (PETS '13)
- The Path Less Travelled: Overcoming Tor's Bottlenecks with Traffic Splitting (PETS '13)
- An Empirical Evaluation of Relay Selection in Tor (NDSS '13)
- LIRA: Lightweight Incentivized Routing for Anonymity (NDSS '13)
- Improving Performance and Anonymity in the Tor Network (IPCCC '12)
- Enhancing Tor's Performance using Real-time Traffic Classification (CCS '12)
- Torchestra: Reducing interactive traffic delays over Tor (WPES '12)
- Throttling Tor Bandwidth Parasites (USENIX Sec '12)
- LASTor: A Low-Latency AS-Aware Tor Client (Oakland '12)
- Congestion-aware Path Selection for Tor (FC '12)

*Not a comprehensive list

Tor is Slow!!! Research*

- PCTCP: Per-Circuit TCP-over-IPsec Transport for Anonymous Communication Overlay Networks (CCS '13)
- Reducing Latency in Tor Circuits with Unordered Delivery (FOCI '13)
- How Low Can You Go: Balancing Performance with Anonymity in Tor (PETS '13)
- The Path Less Travelled: Overcoming Tor's Bottlenecks with Traffic Splitting (PETS '13)
- An Empirical Evaluation of Relay Selection in Tor (NDSS '13)
- LIRA: Lightweight Incentivized Routing for Anonymity (NDSS '13)
- Improving Performance and Anonymity in the Tor Network (IPCCC '12)
- Enhancing Tor's Performance using Real-time Traffic Classification (CCS '12)
- Torchestra: Reducing interactive traffic delays over Tor (WPES '12)
- Throttling Tor Bandwidth Parasites (USENIX Sec '12)
- LASTor: A Low-Latency AS-Aware Tor Client (Oakland '12)
- Congestion-aware Path Selection for Tor (FC '12)

Where?

*Not a comprehensive list


Outline

- **Where** is Tor slow?
 - Understand Tor relay architecture
 - Measure and analyze relay congestion in realistic Tor networks
- Design **focused** solutions


Outline

- **Where** is Tor slow?
 - Understand Tor relay architecture
 - Measure and analyze relay congestion in realistic Tor networks
- Design focused solutions


The Tor Network




Relay Overview




Relay Overview



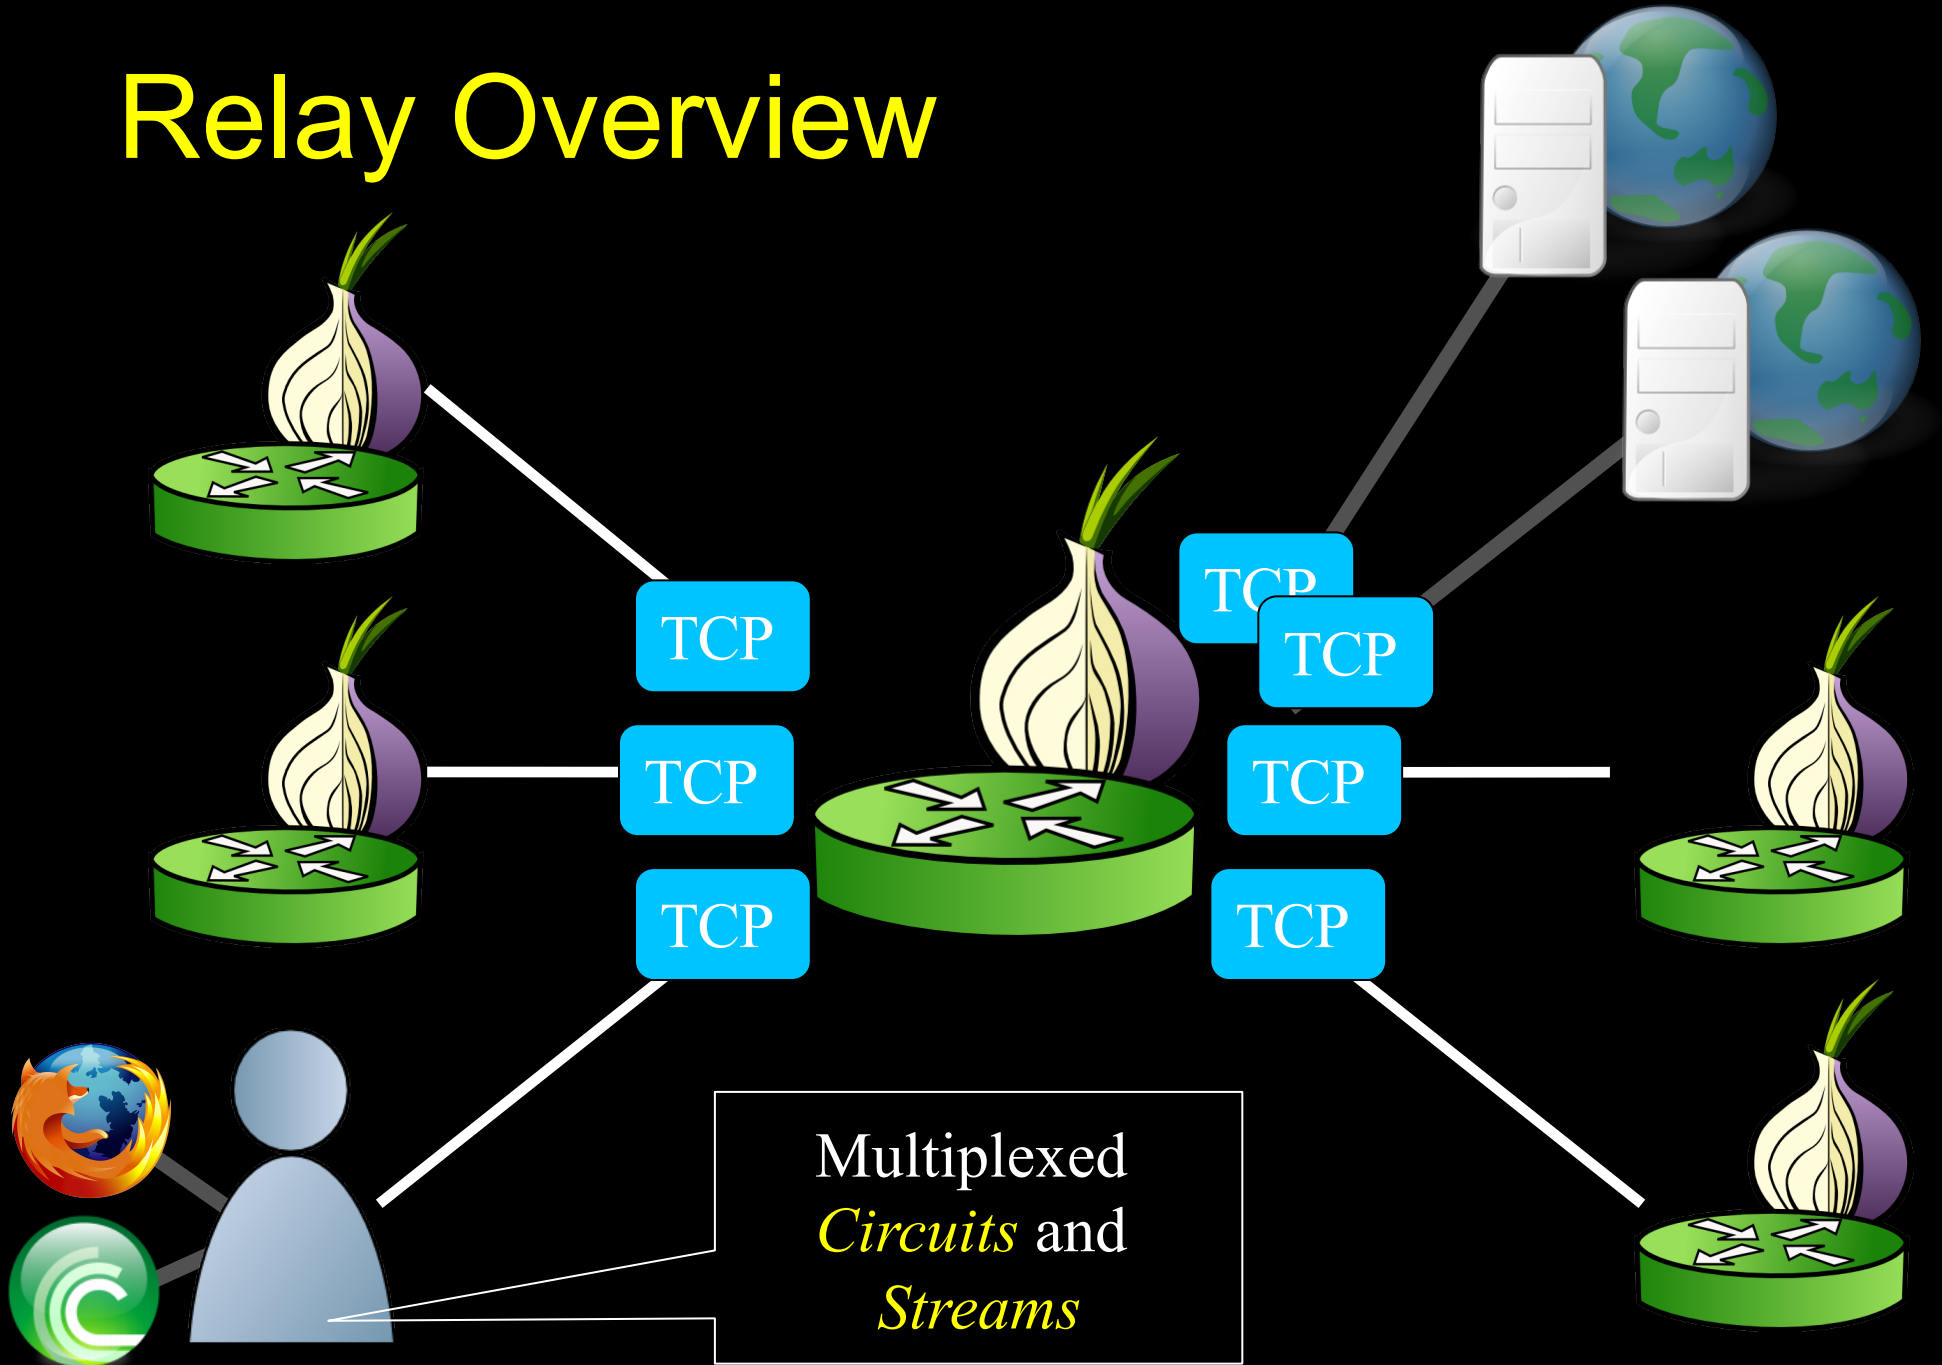
Relay Overview



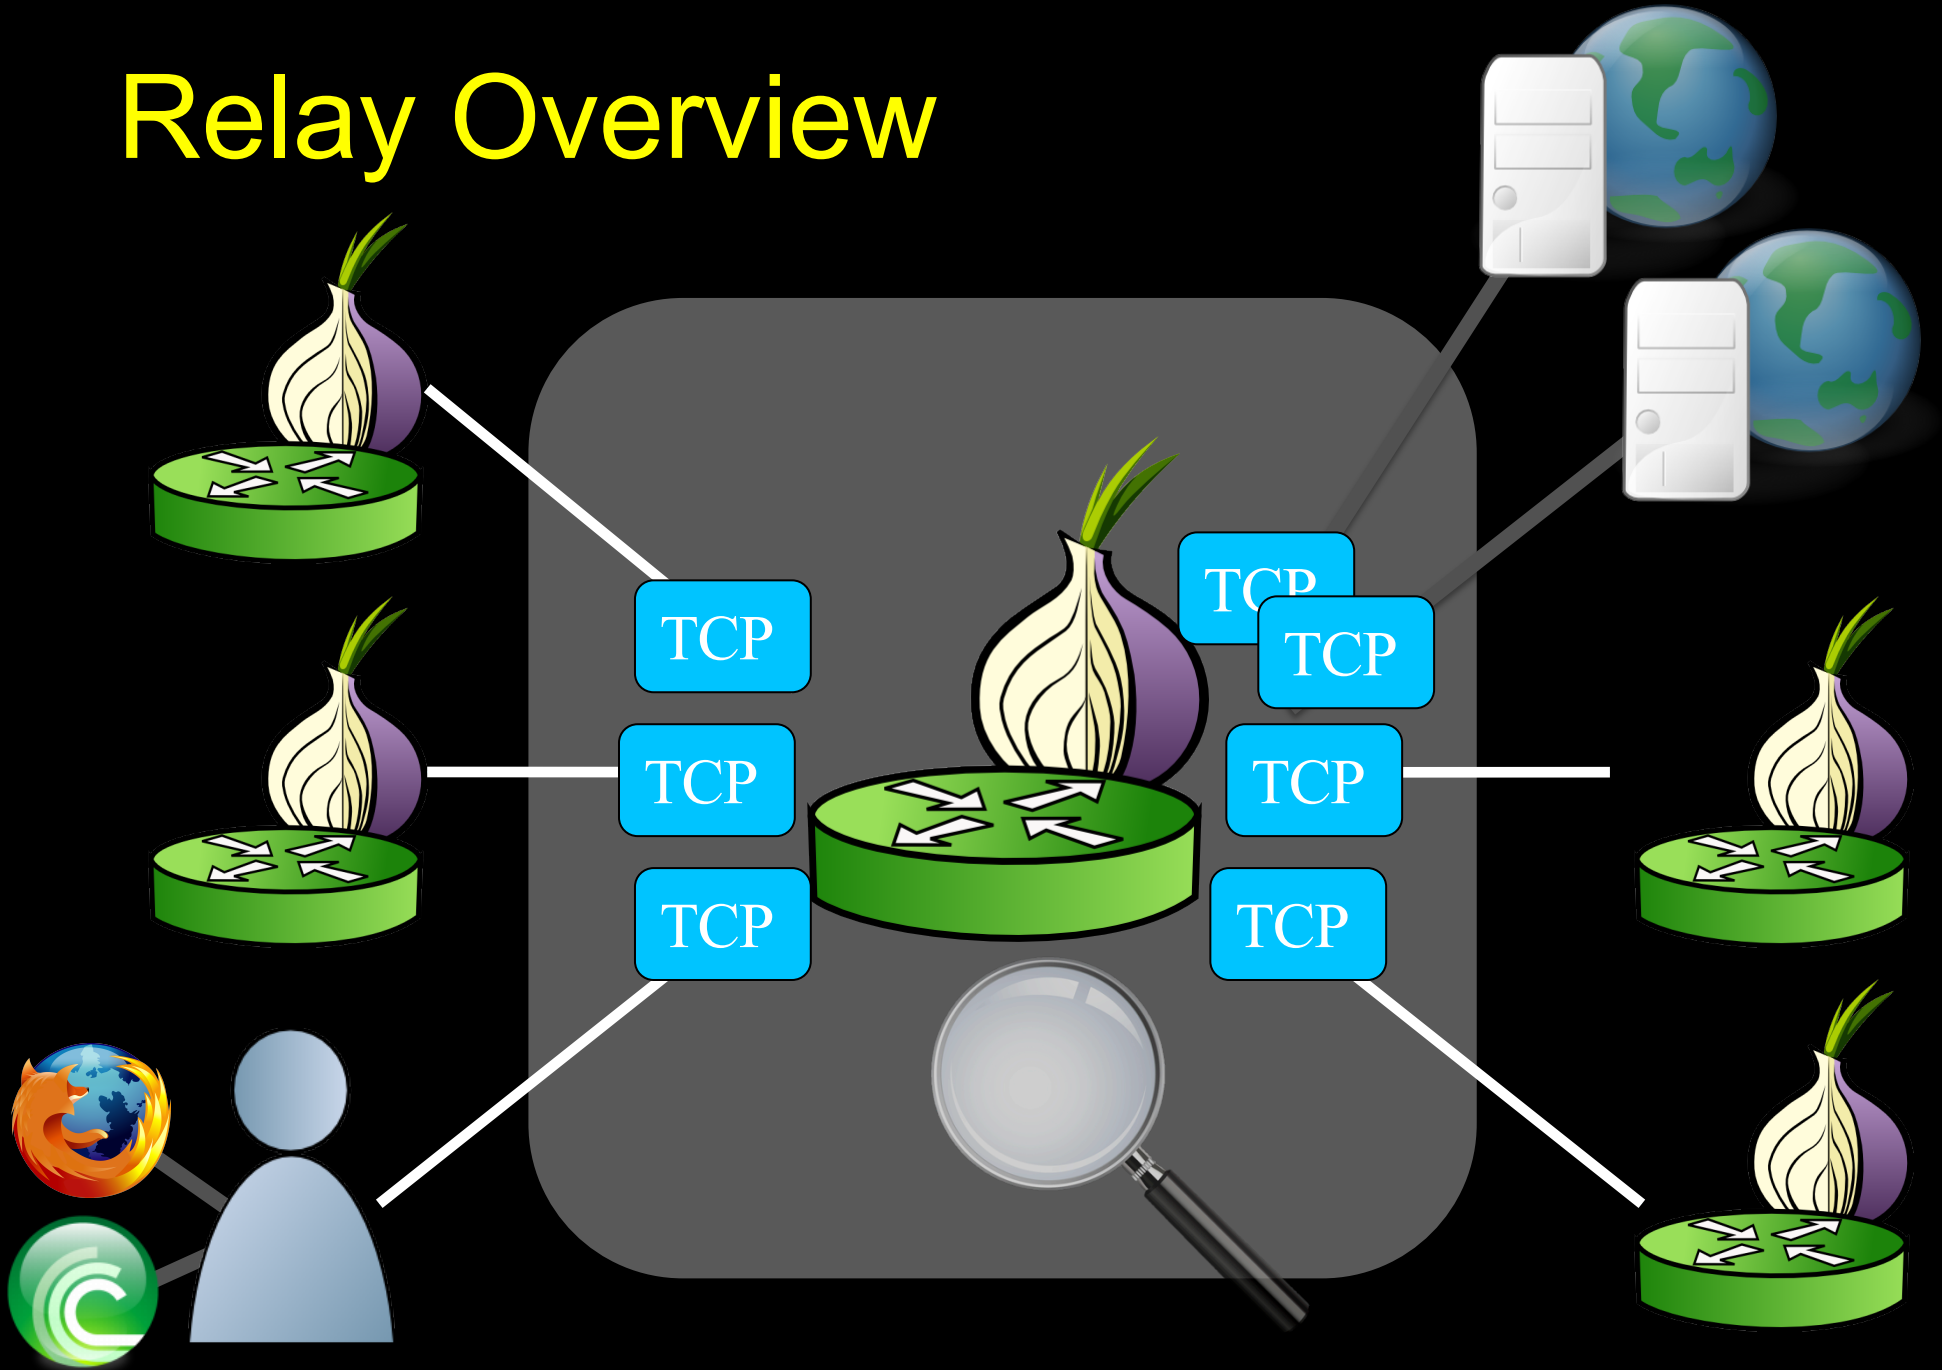
Relay Overview



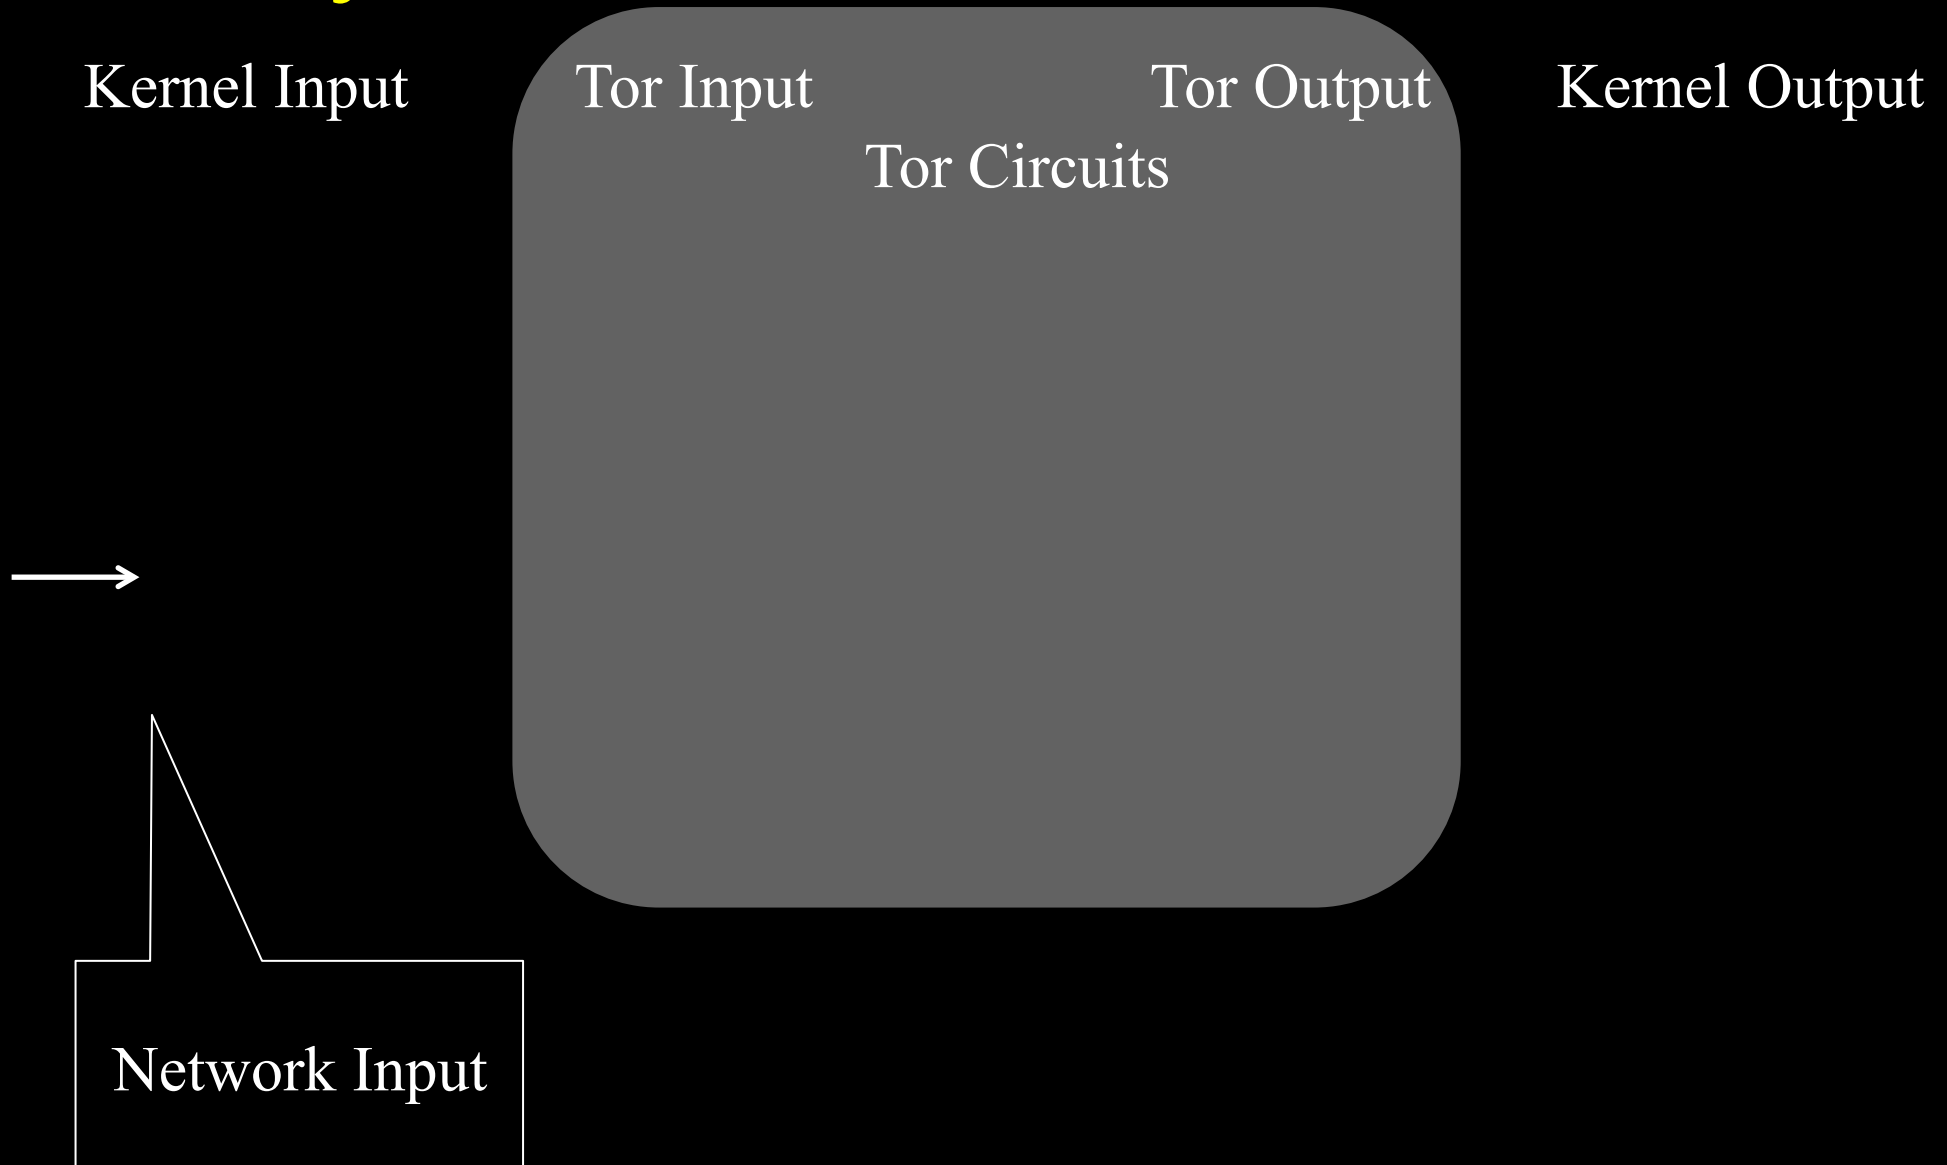
Relay Overview



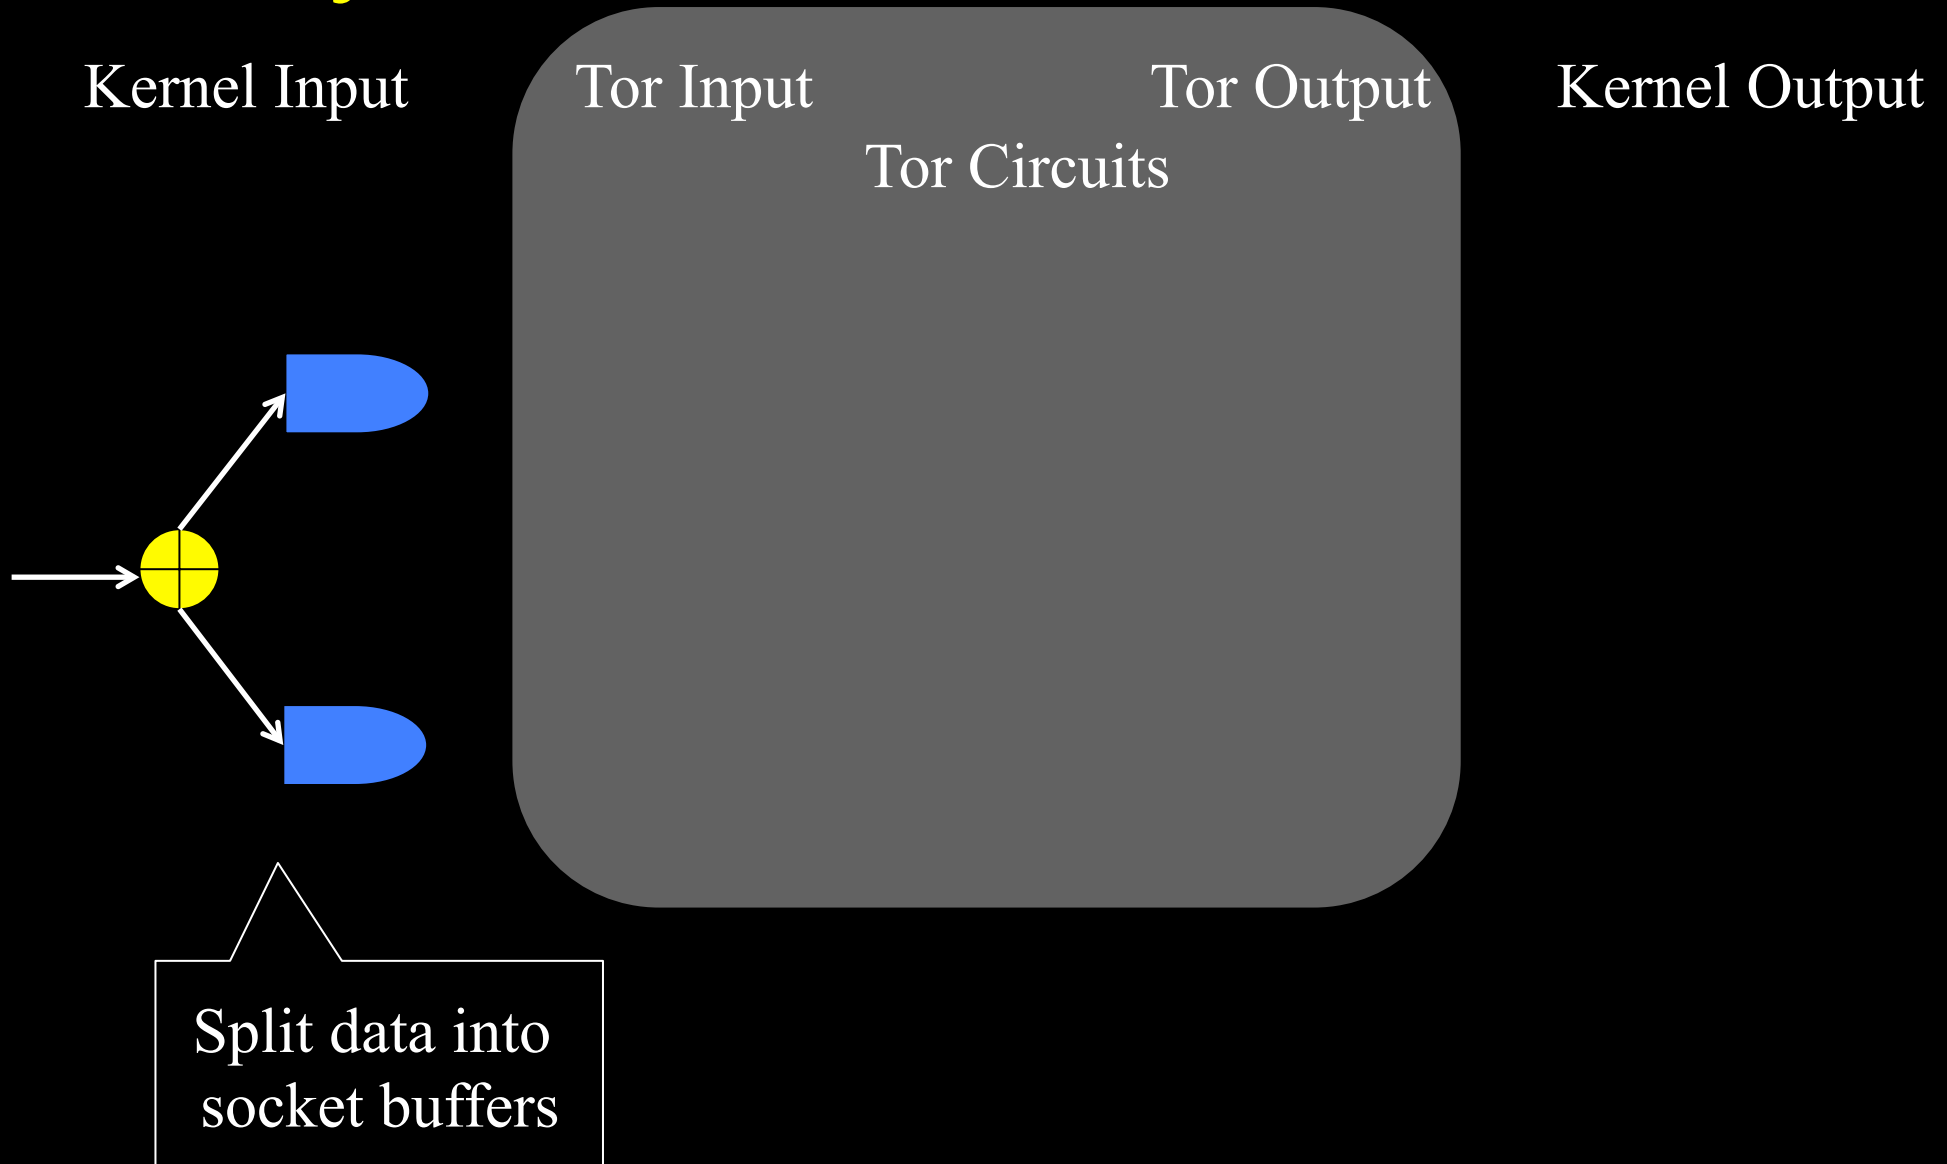
Relay Overview



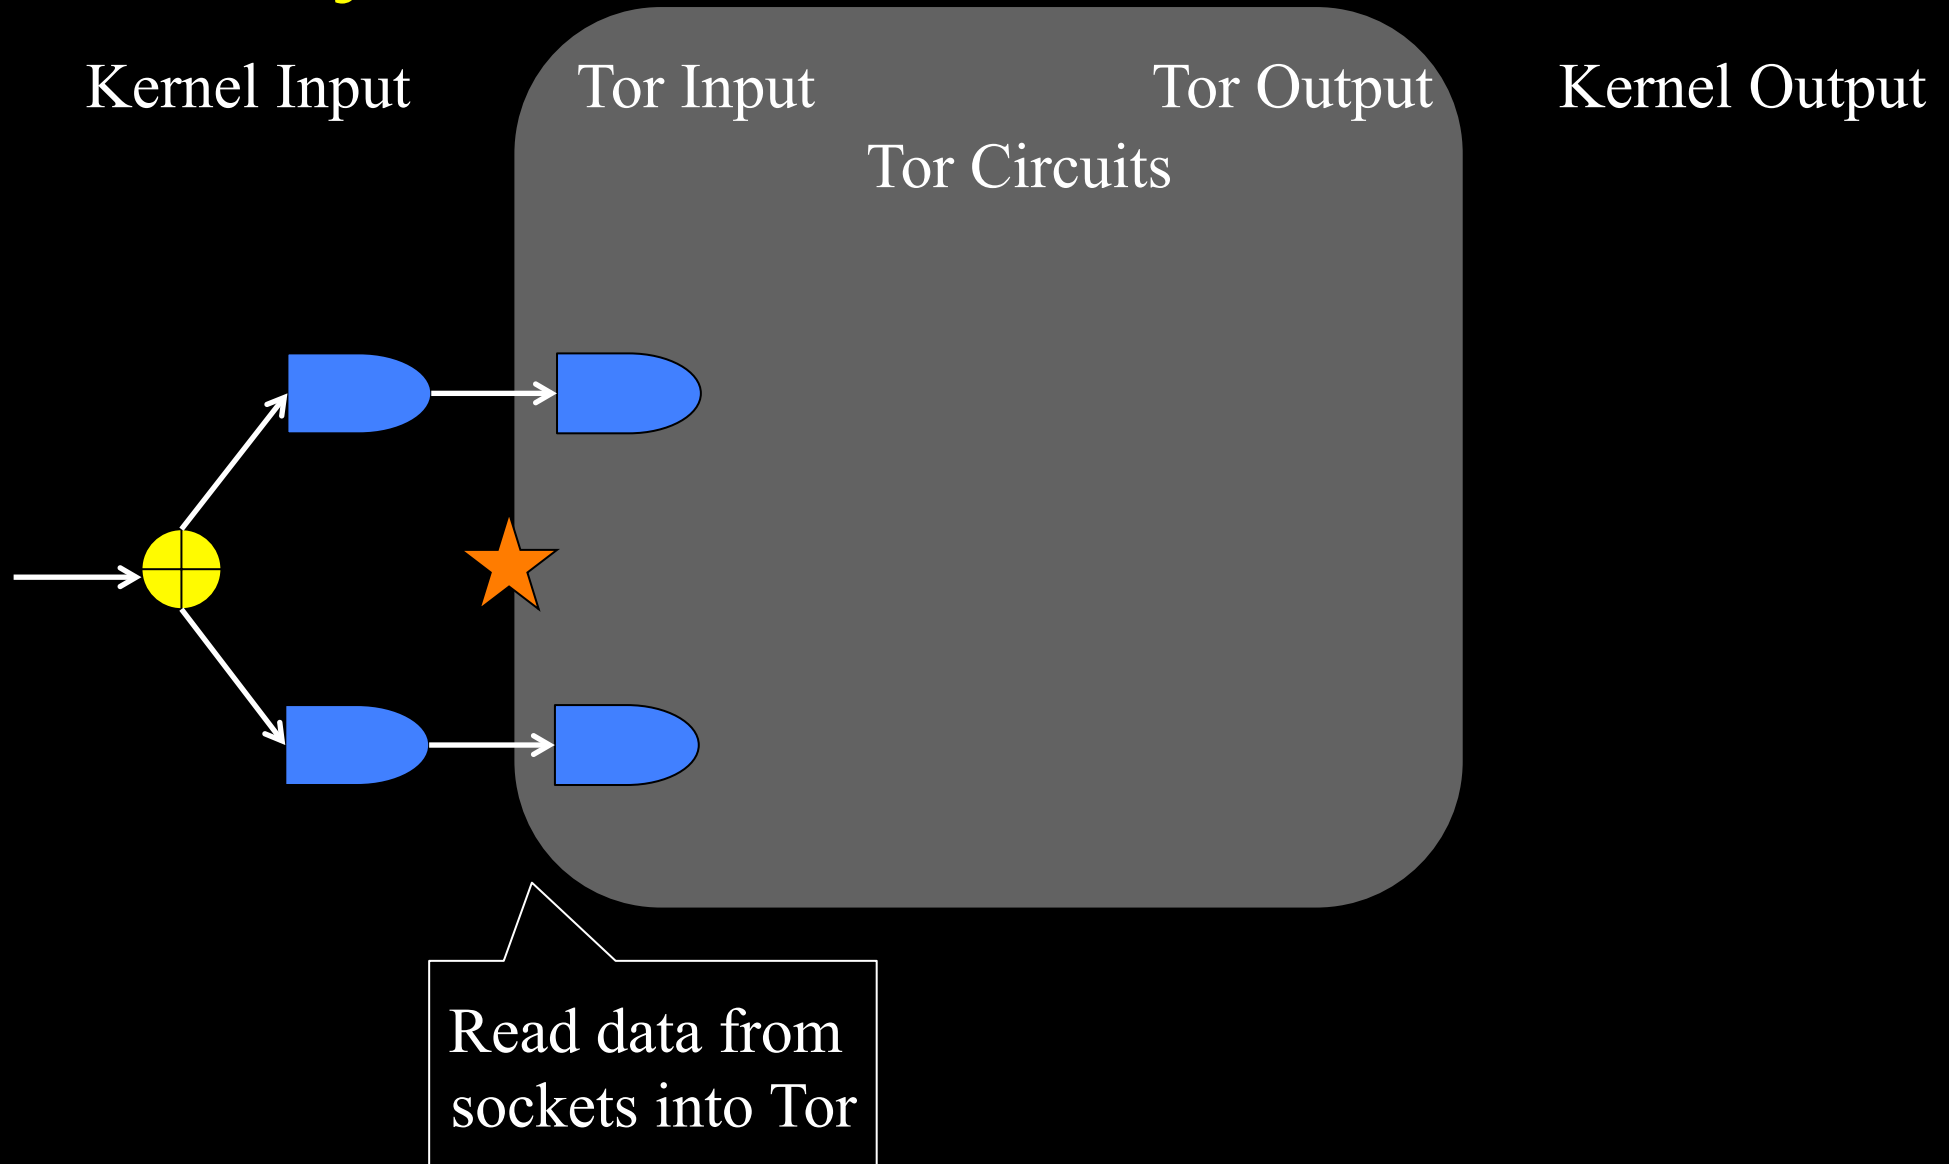
Relay Internals



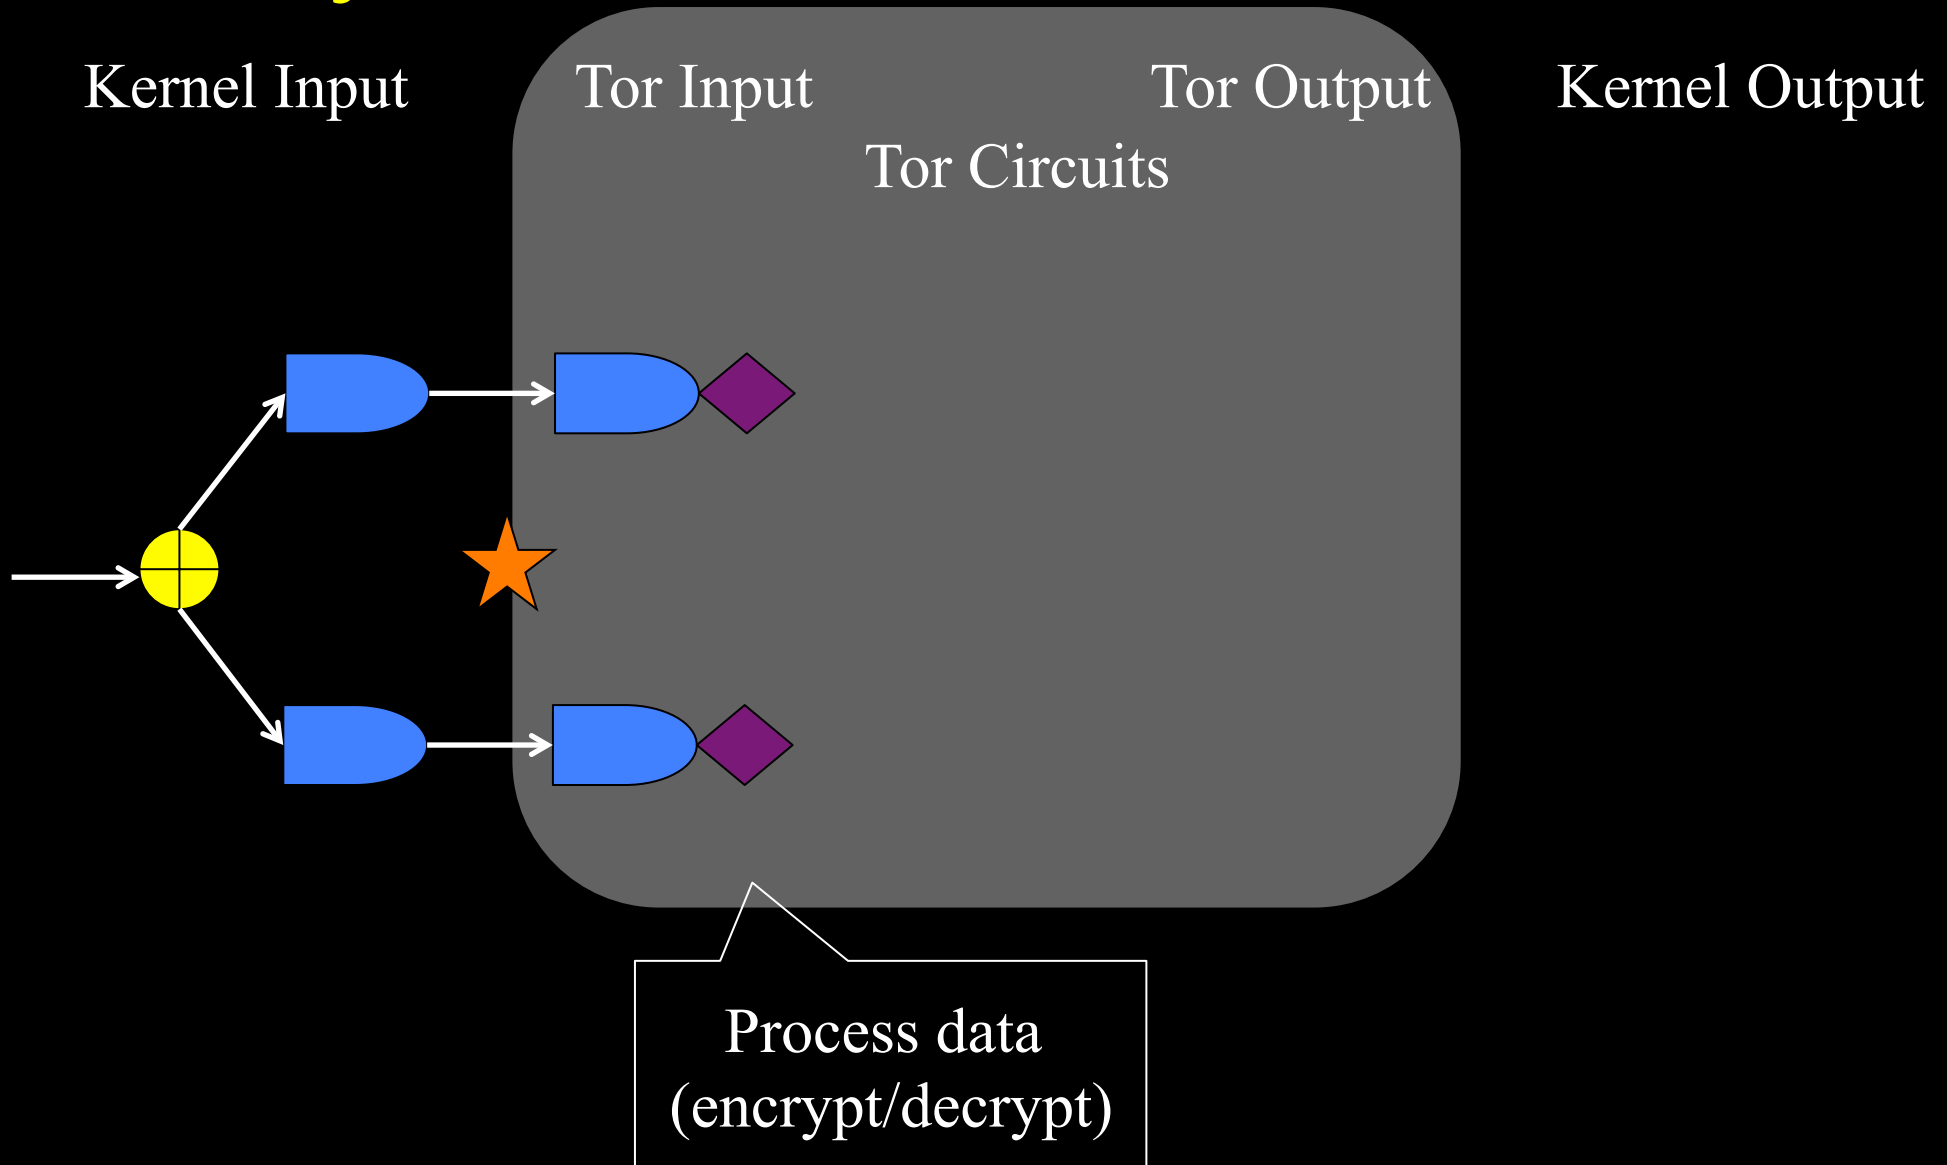
Relay Internals



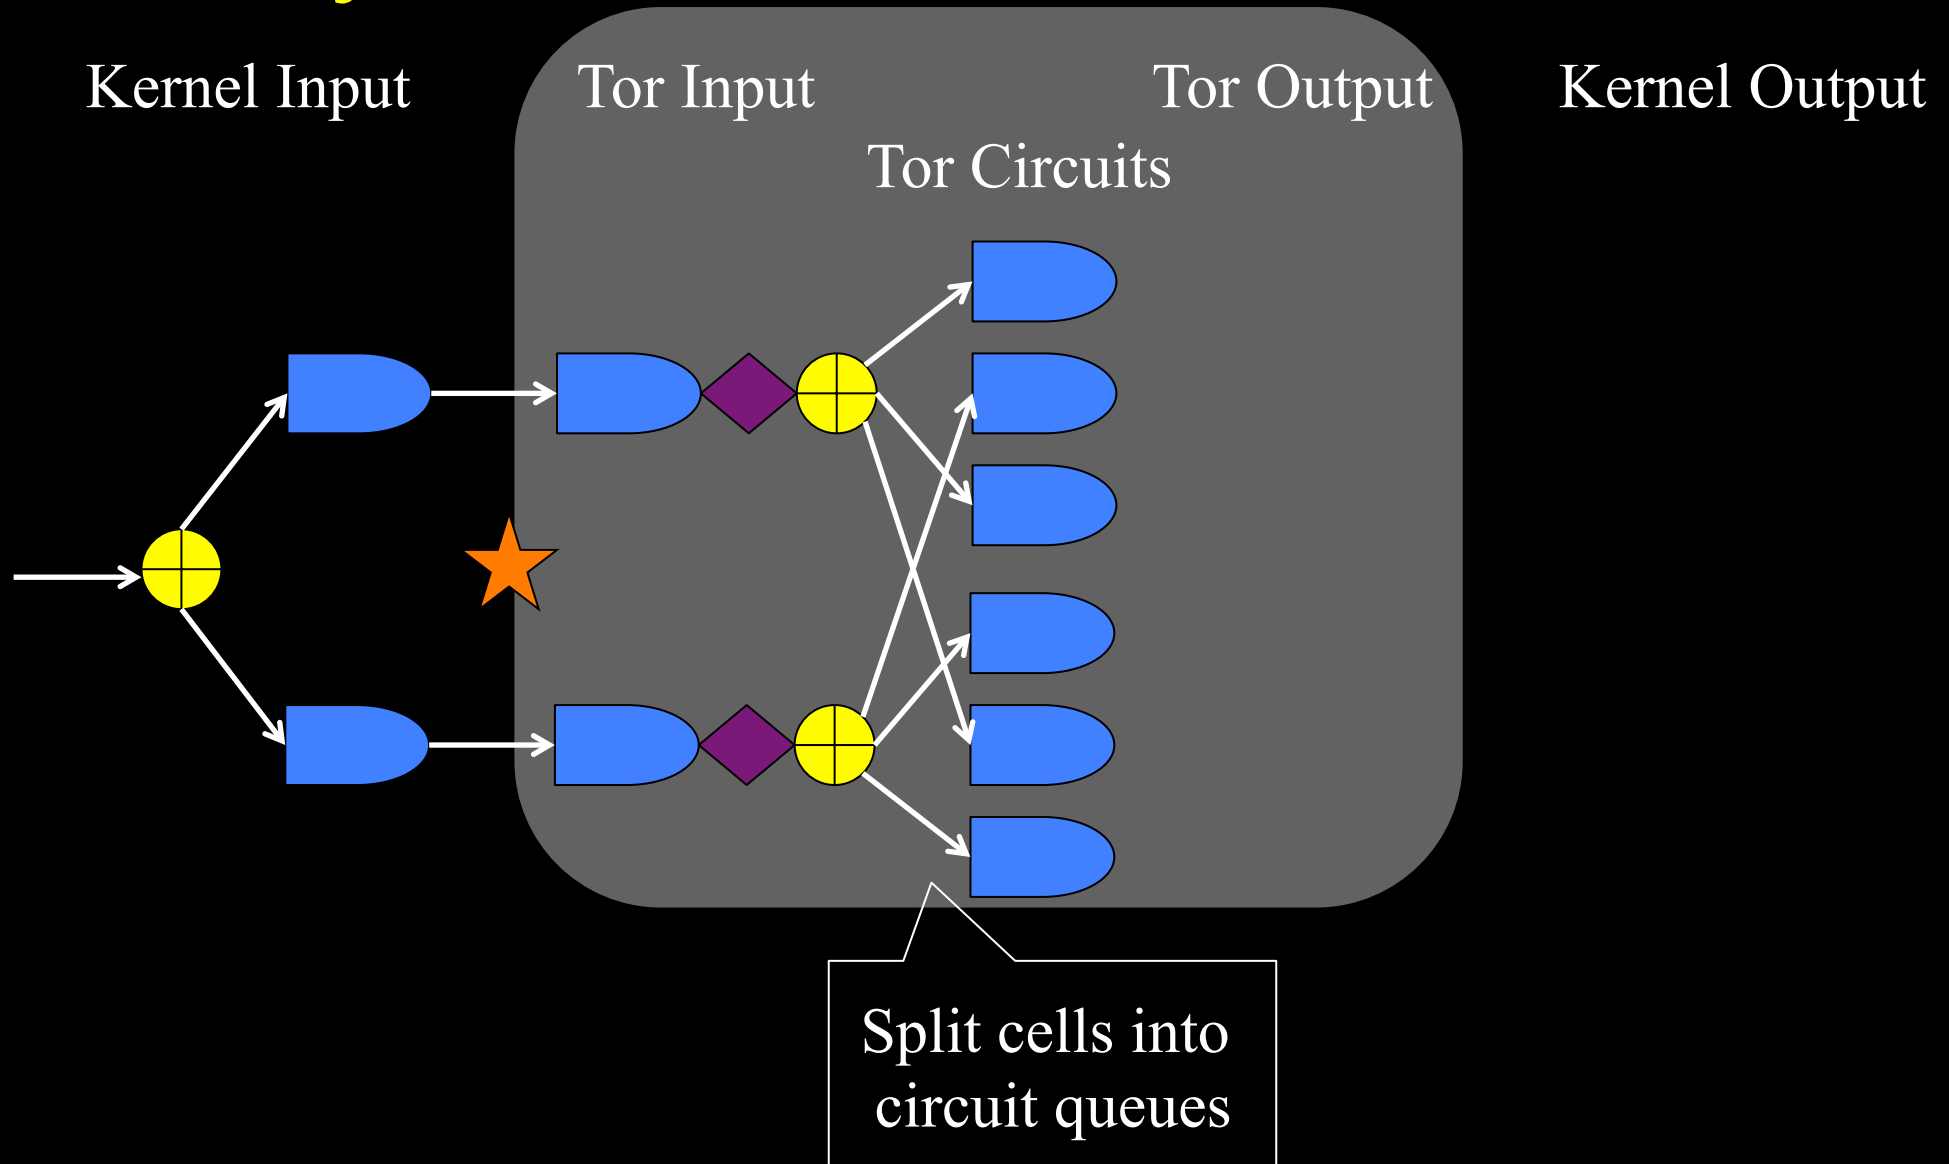
Relay Internals



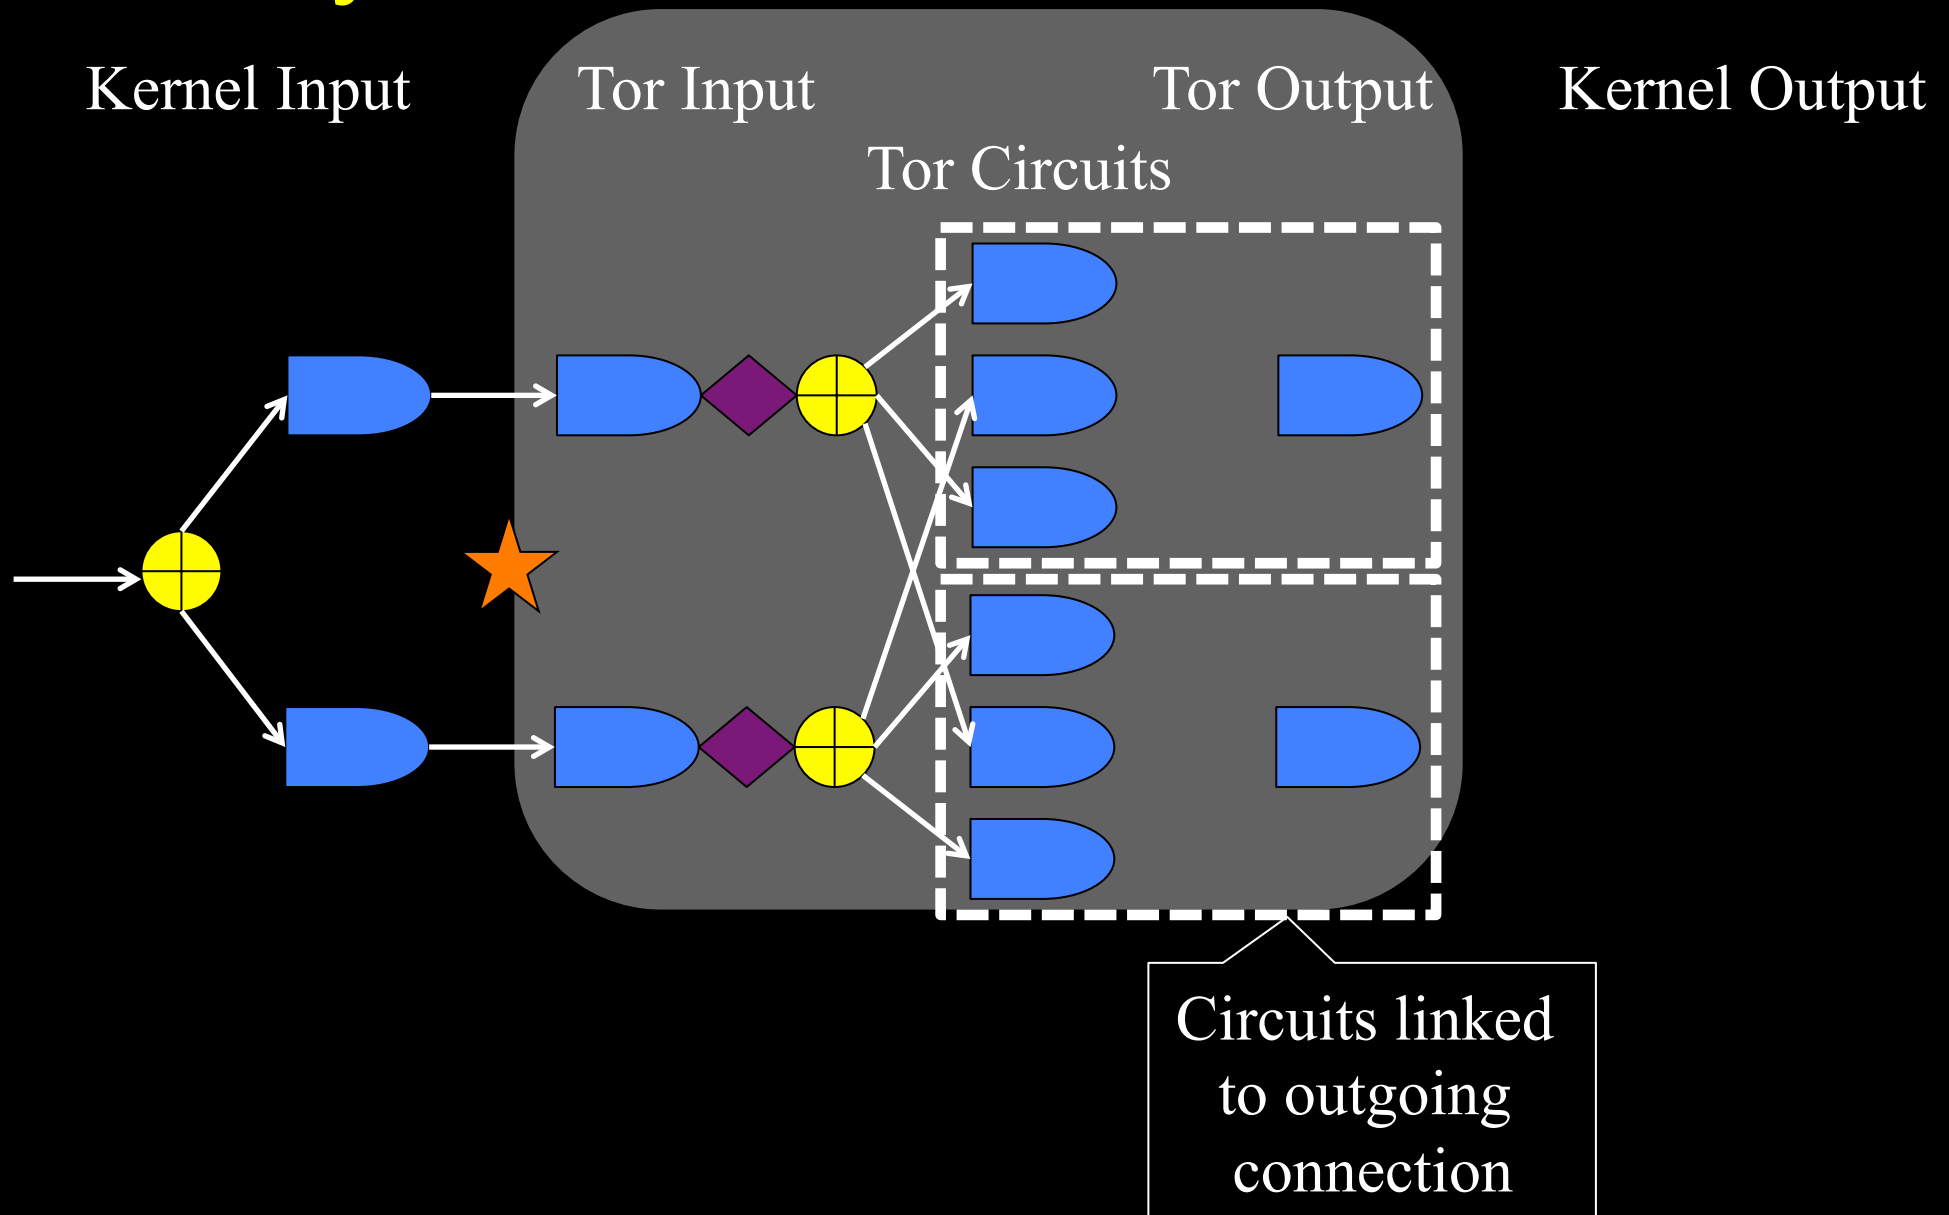
Relay Internals



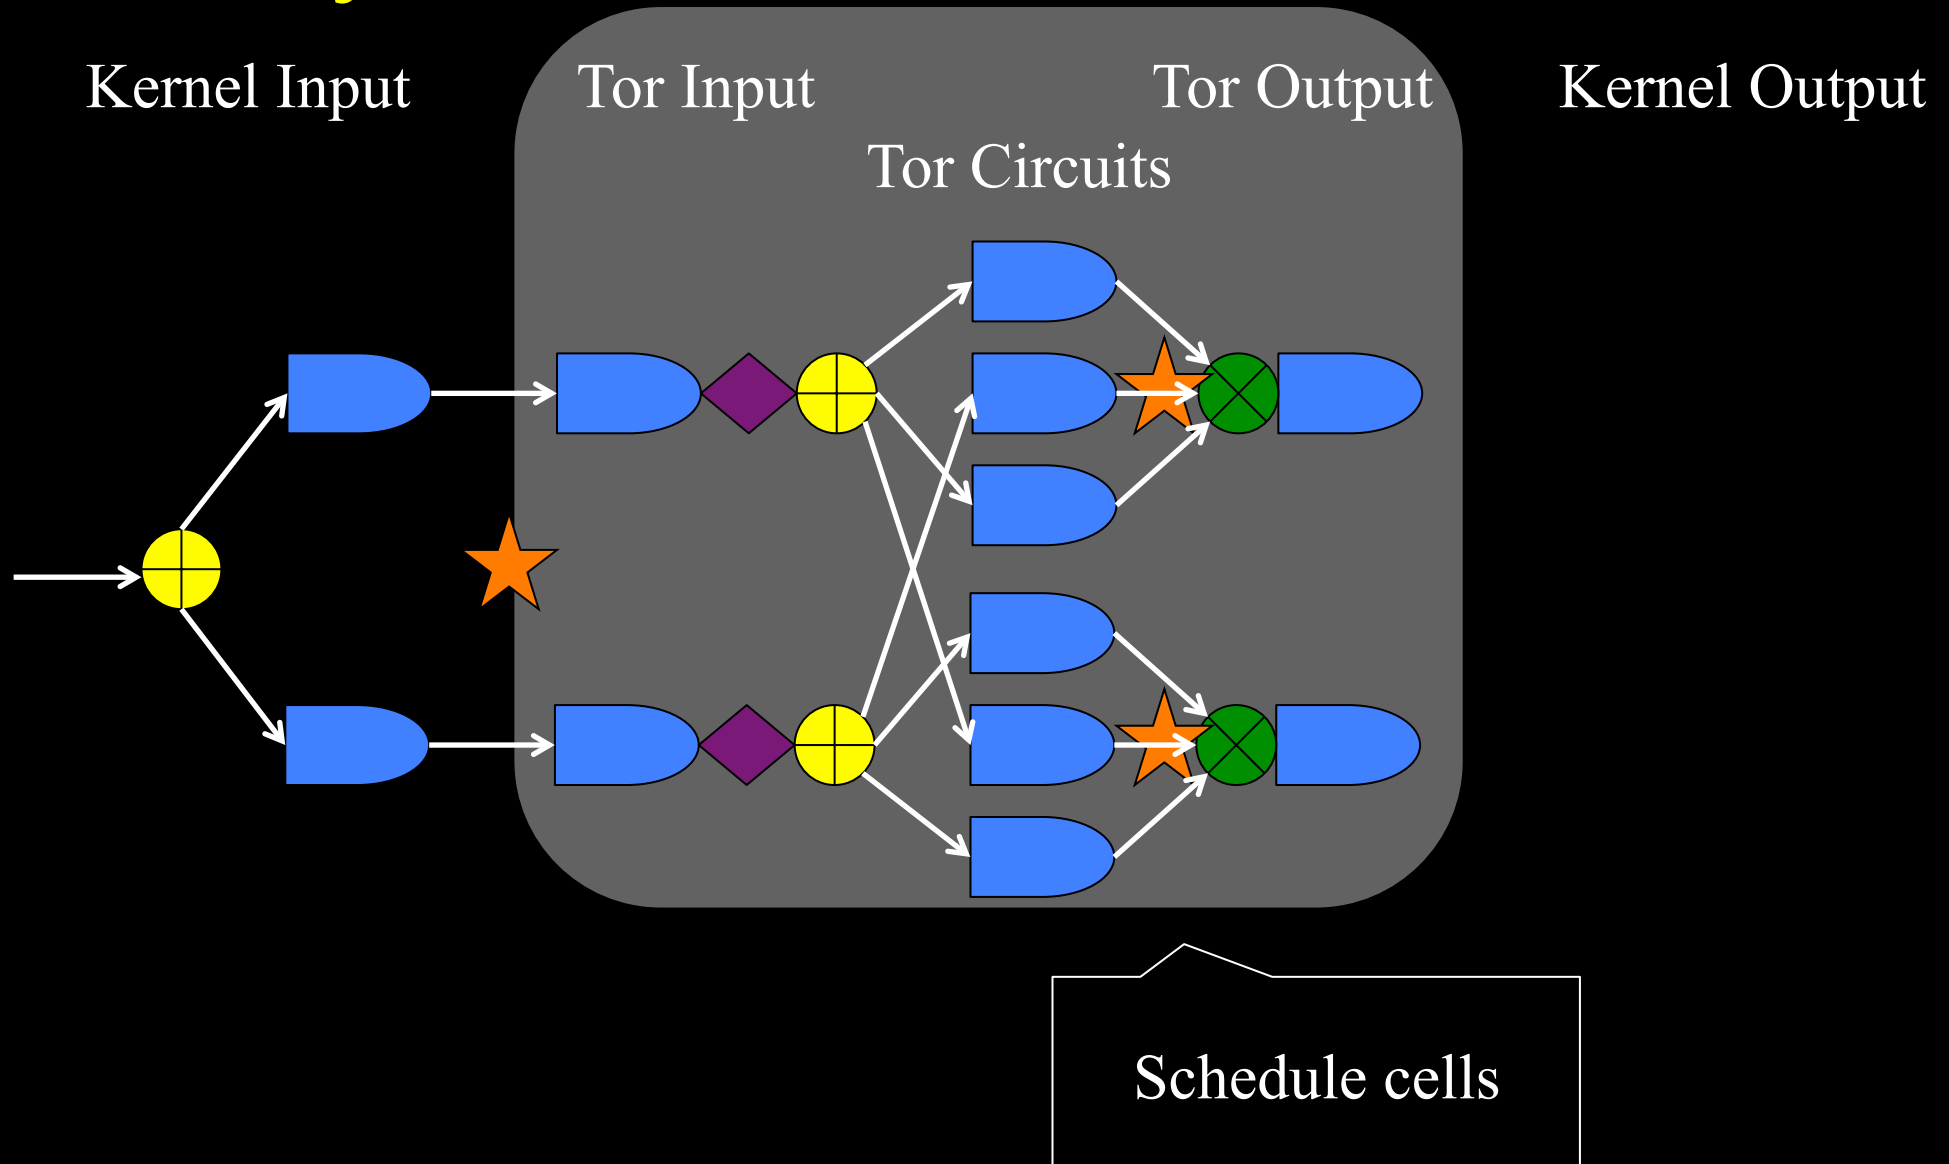
Relay Internals



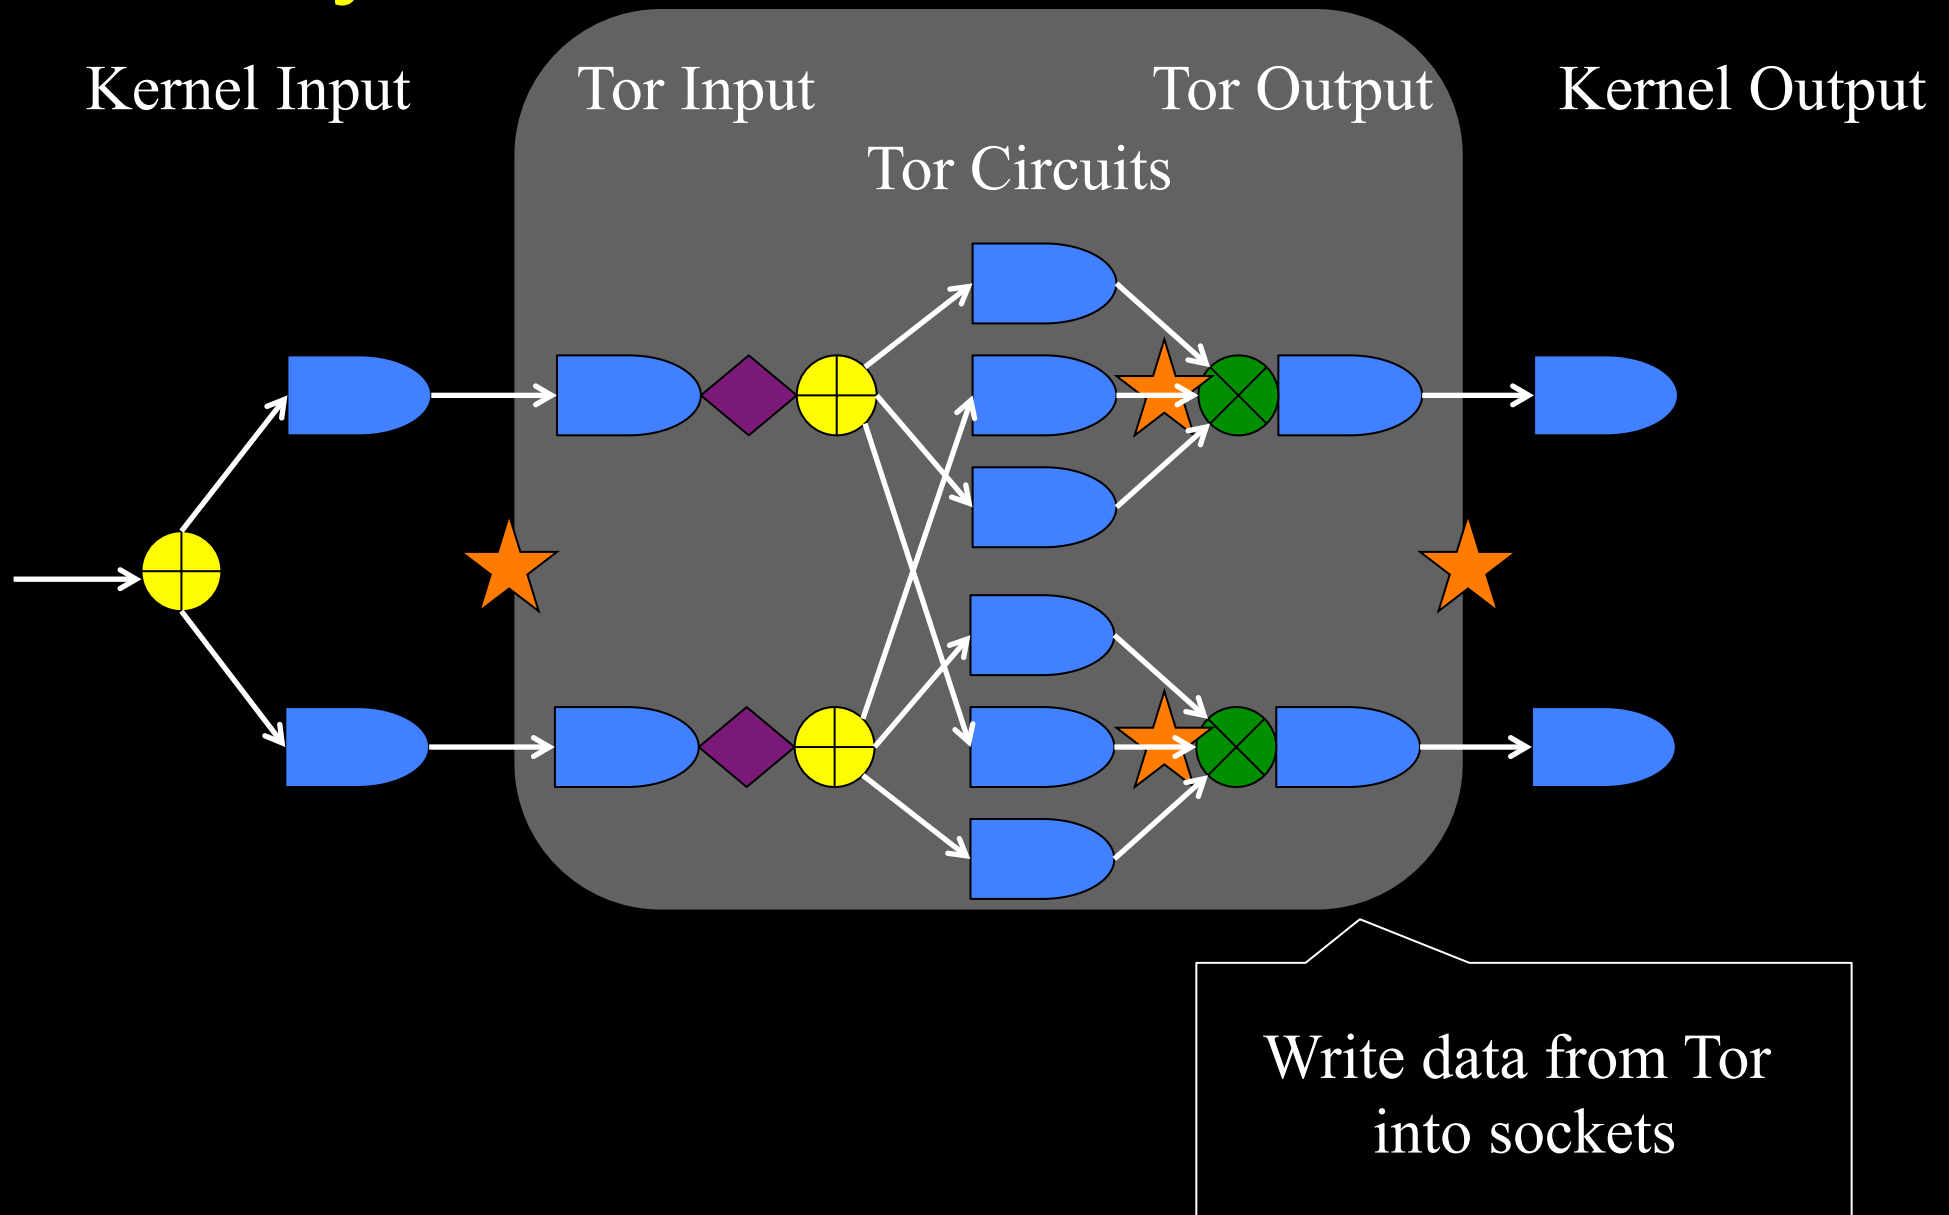
Relay Internals



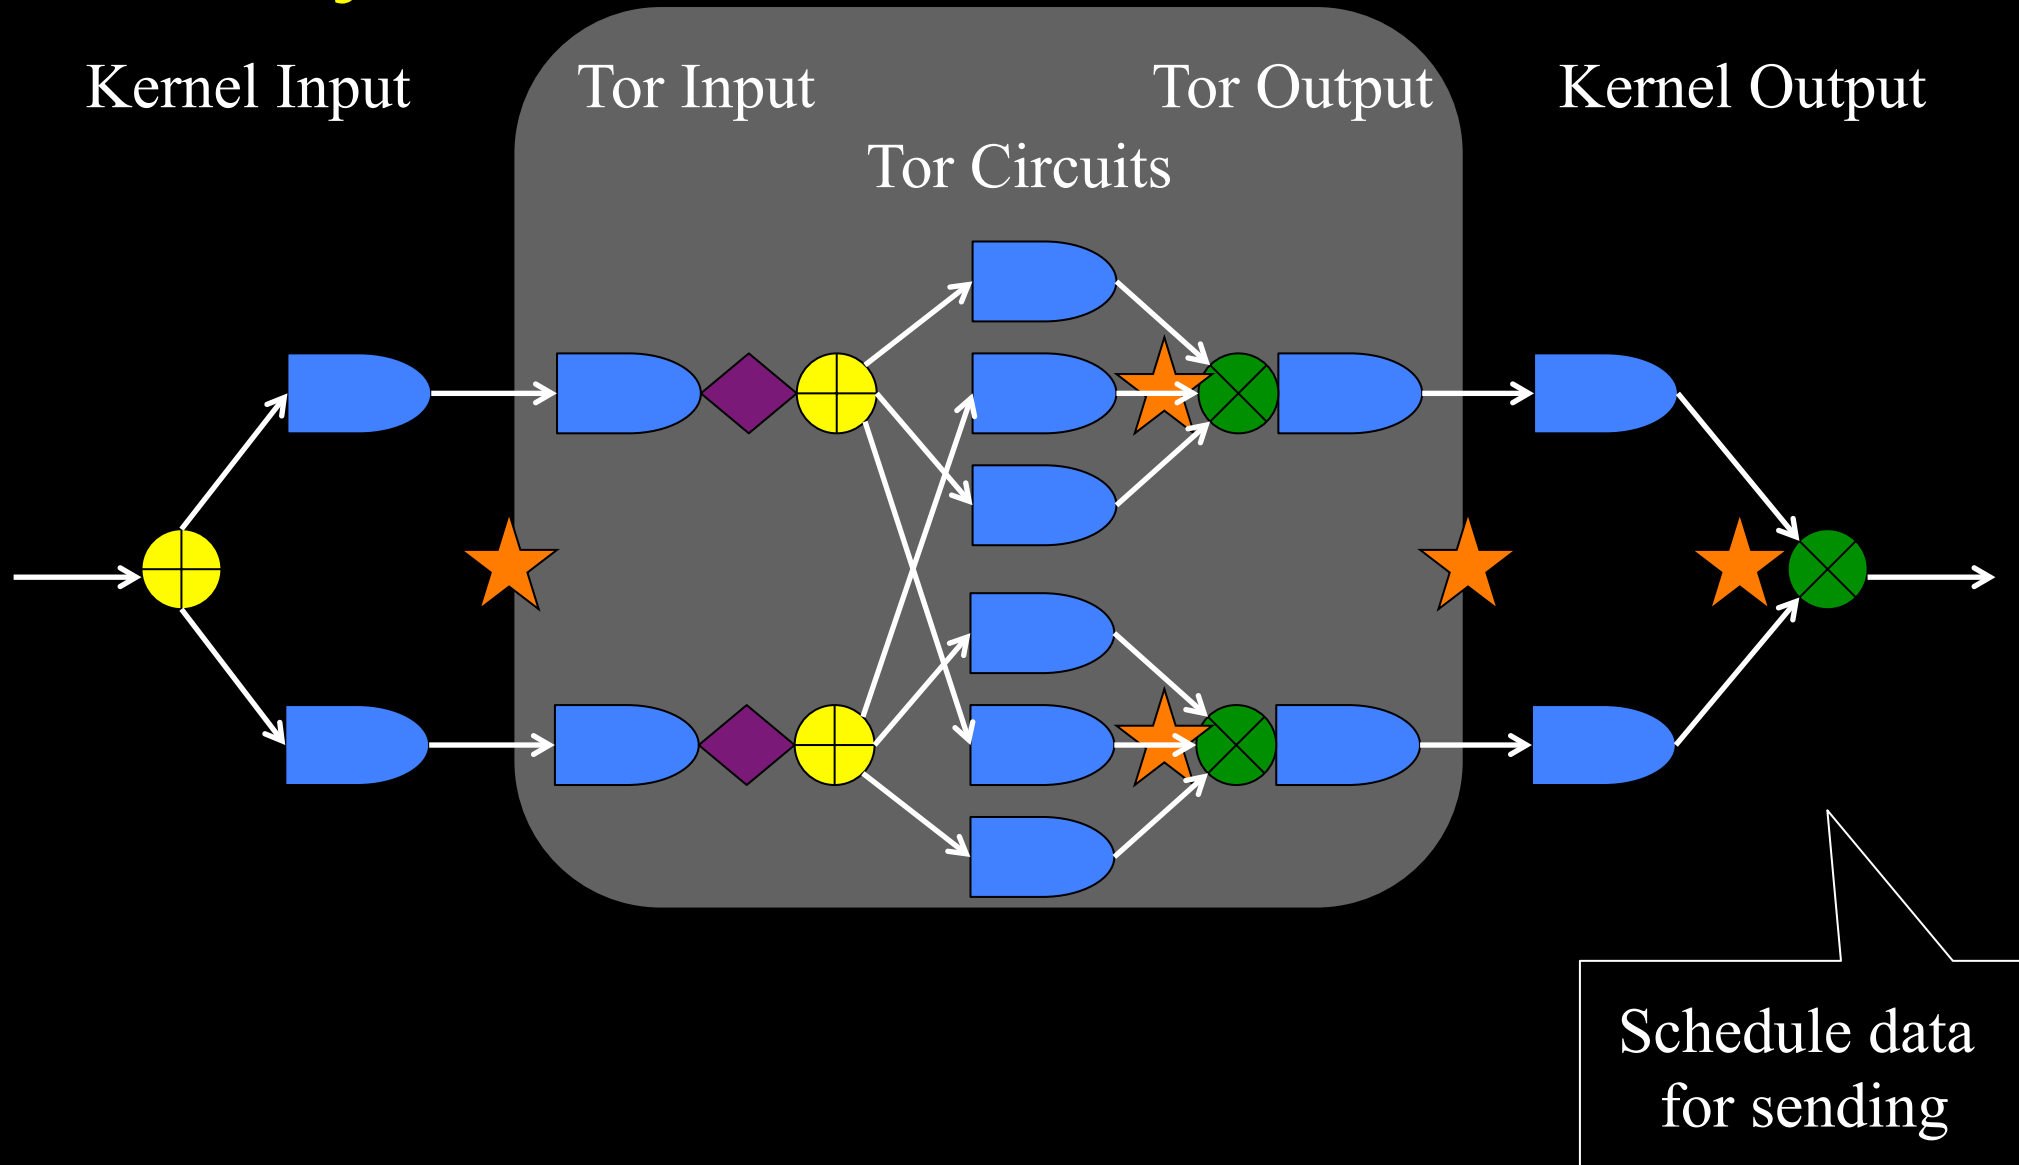
Relay Internals



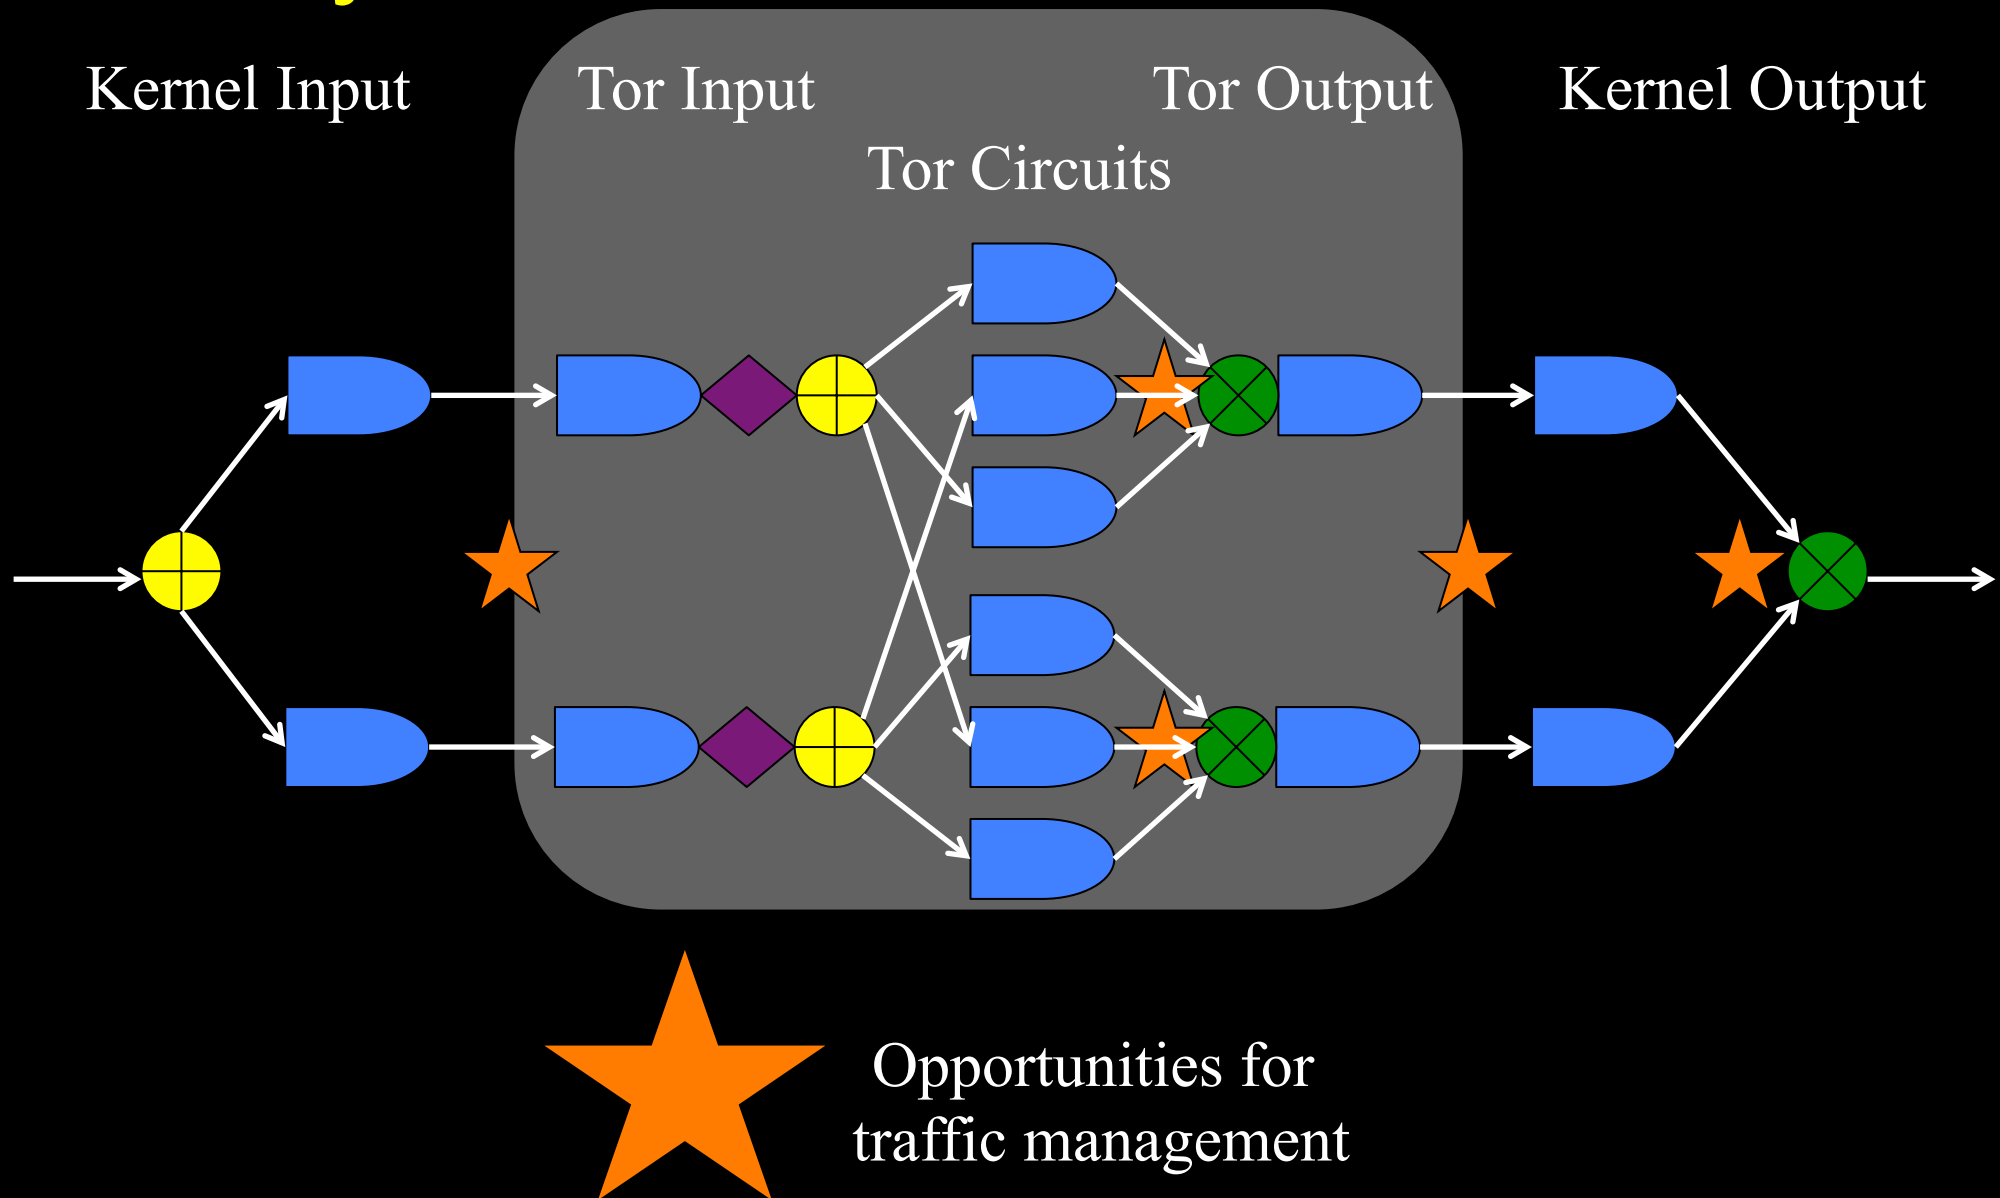
Relay Internals



Relay Internals



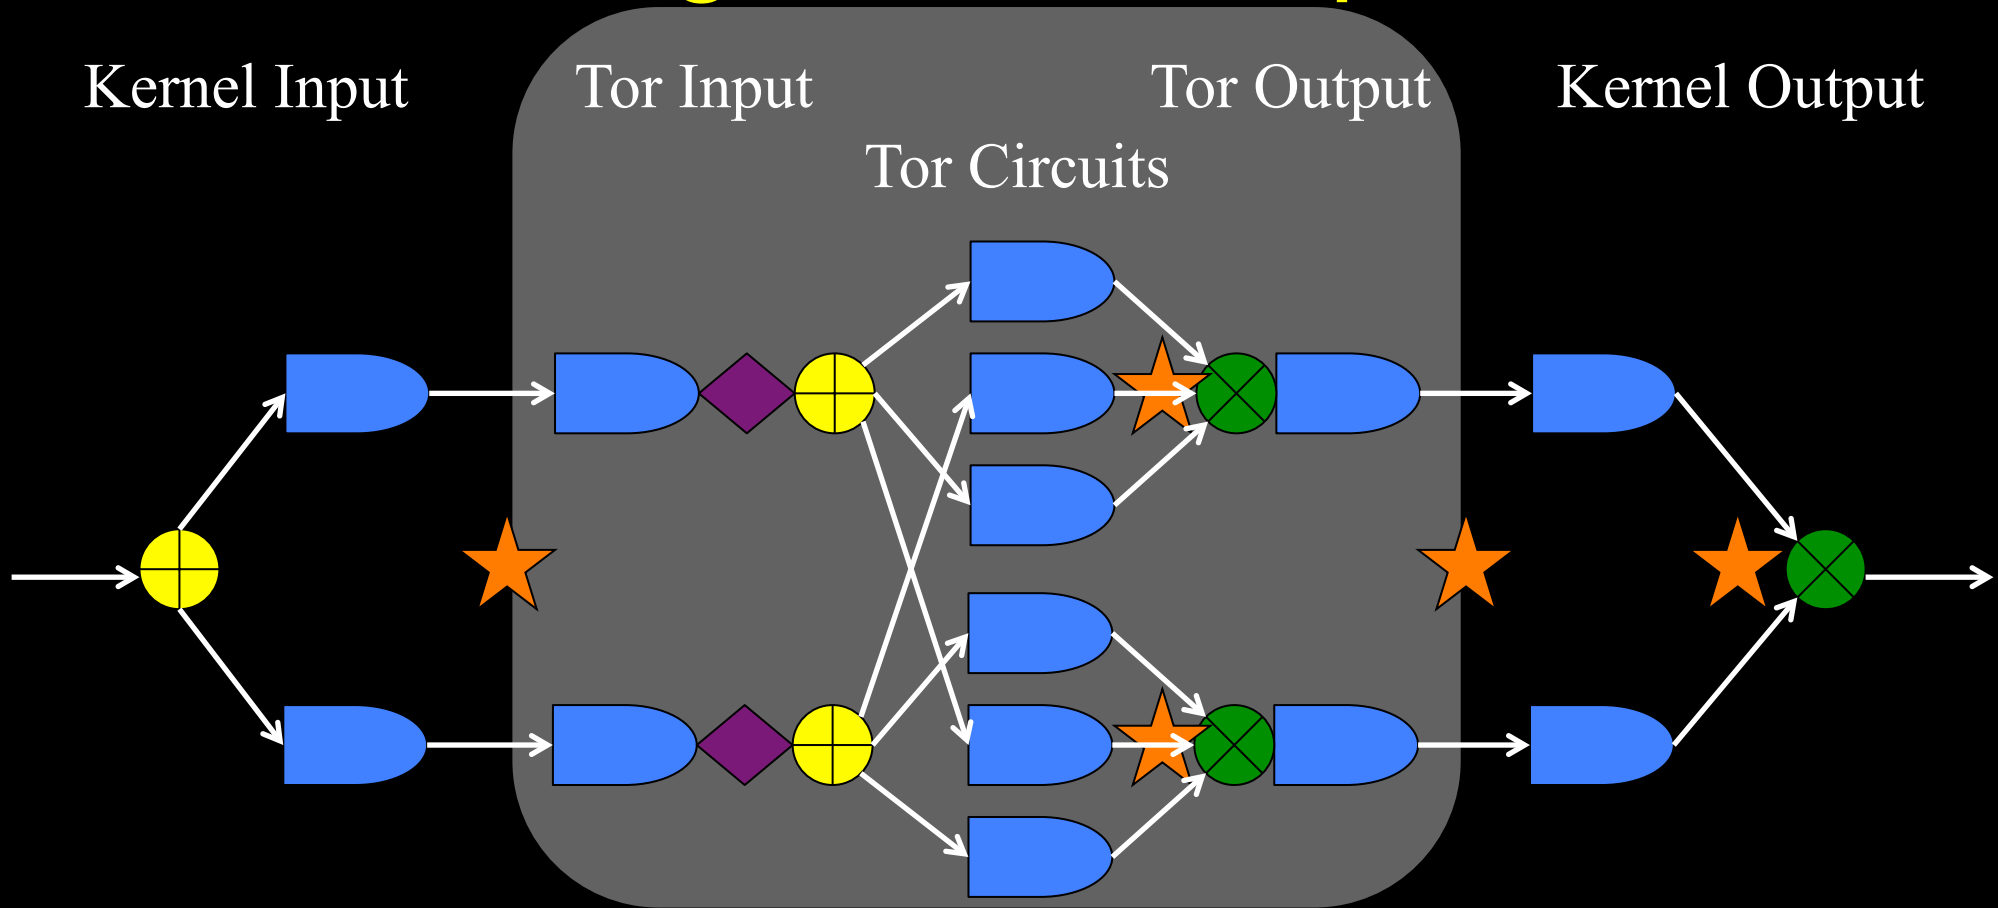
Relay Internals



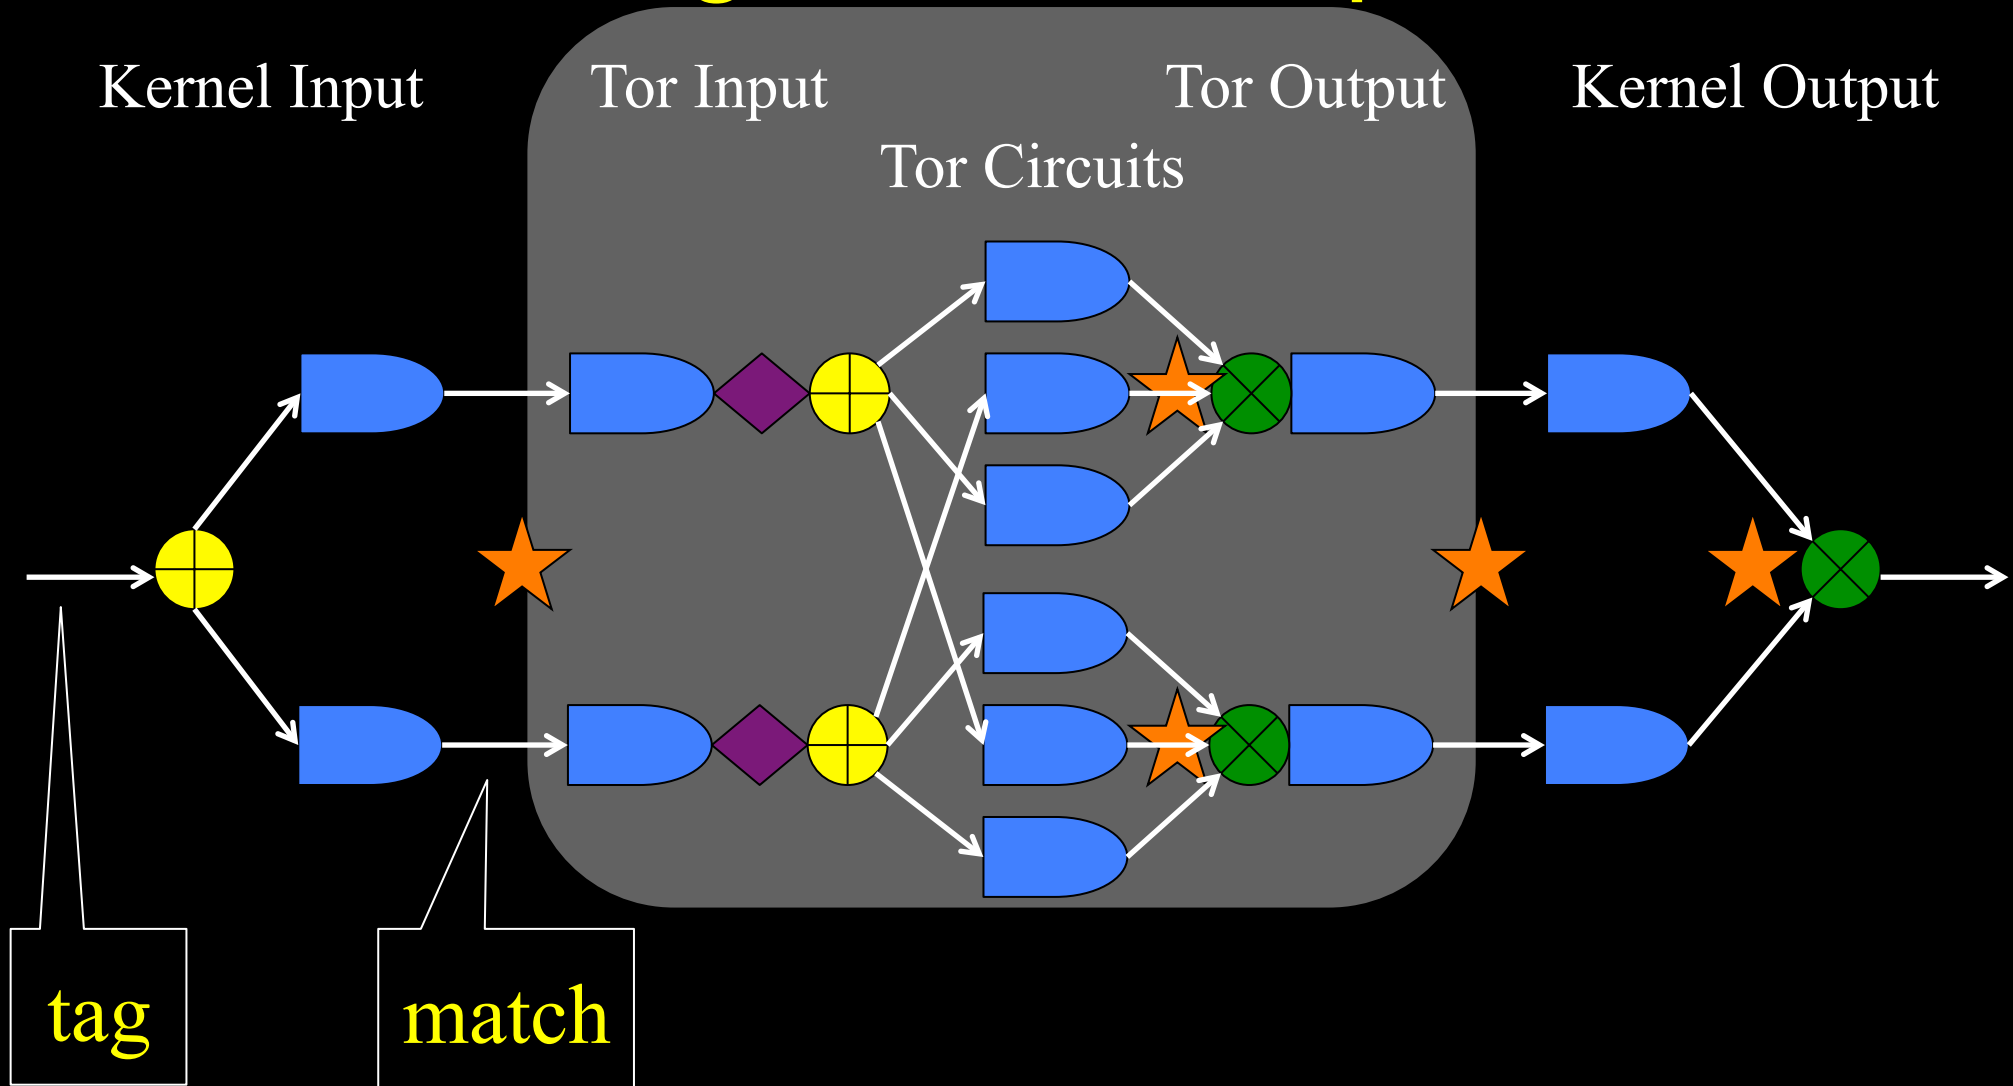
Outline

- **Where** is Tor slow?
 - Understand Tor relay architecture
 - Measure and analyze relay congestion in realistic Tor networks
- Design focused solutions

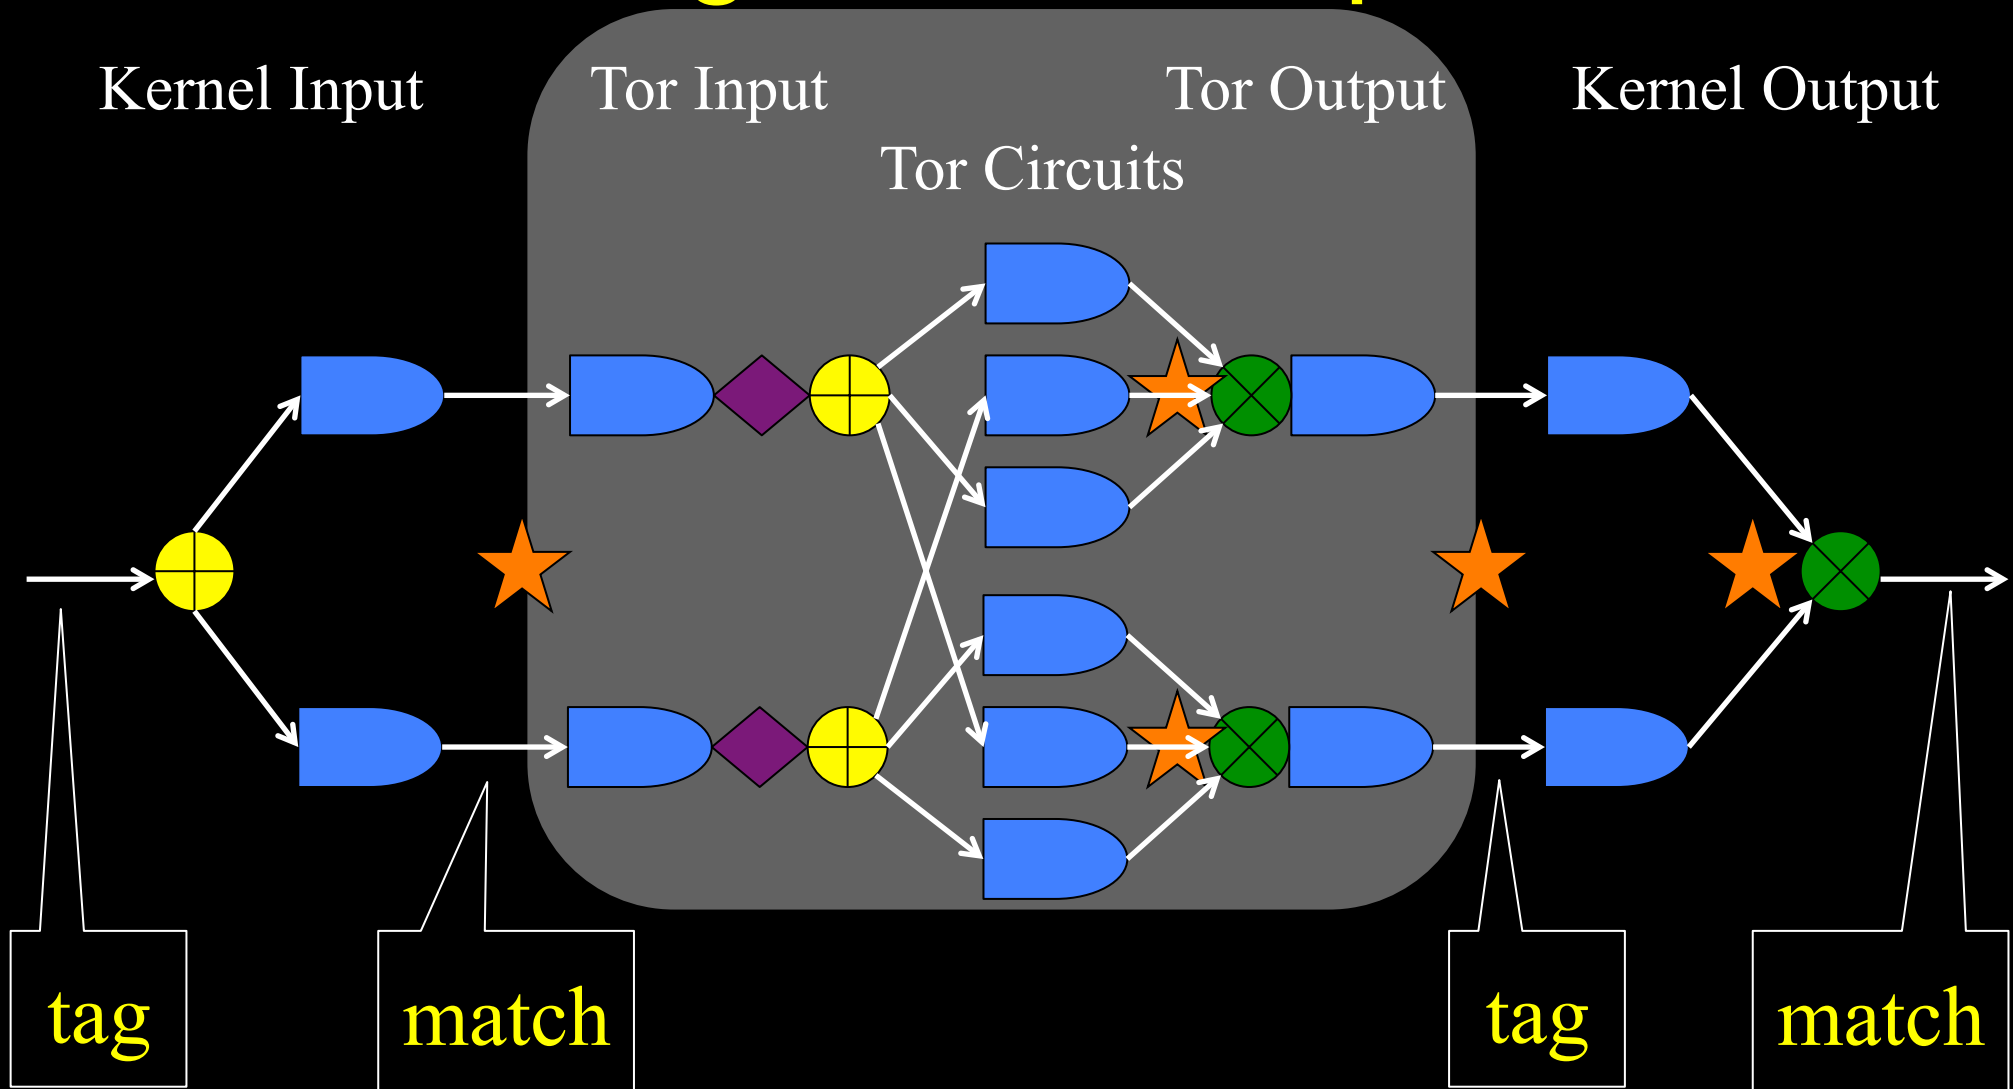
Kernel Congestion: libkqtime



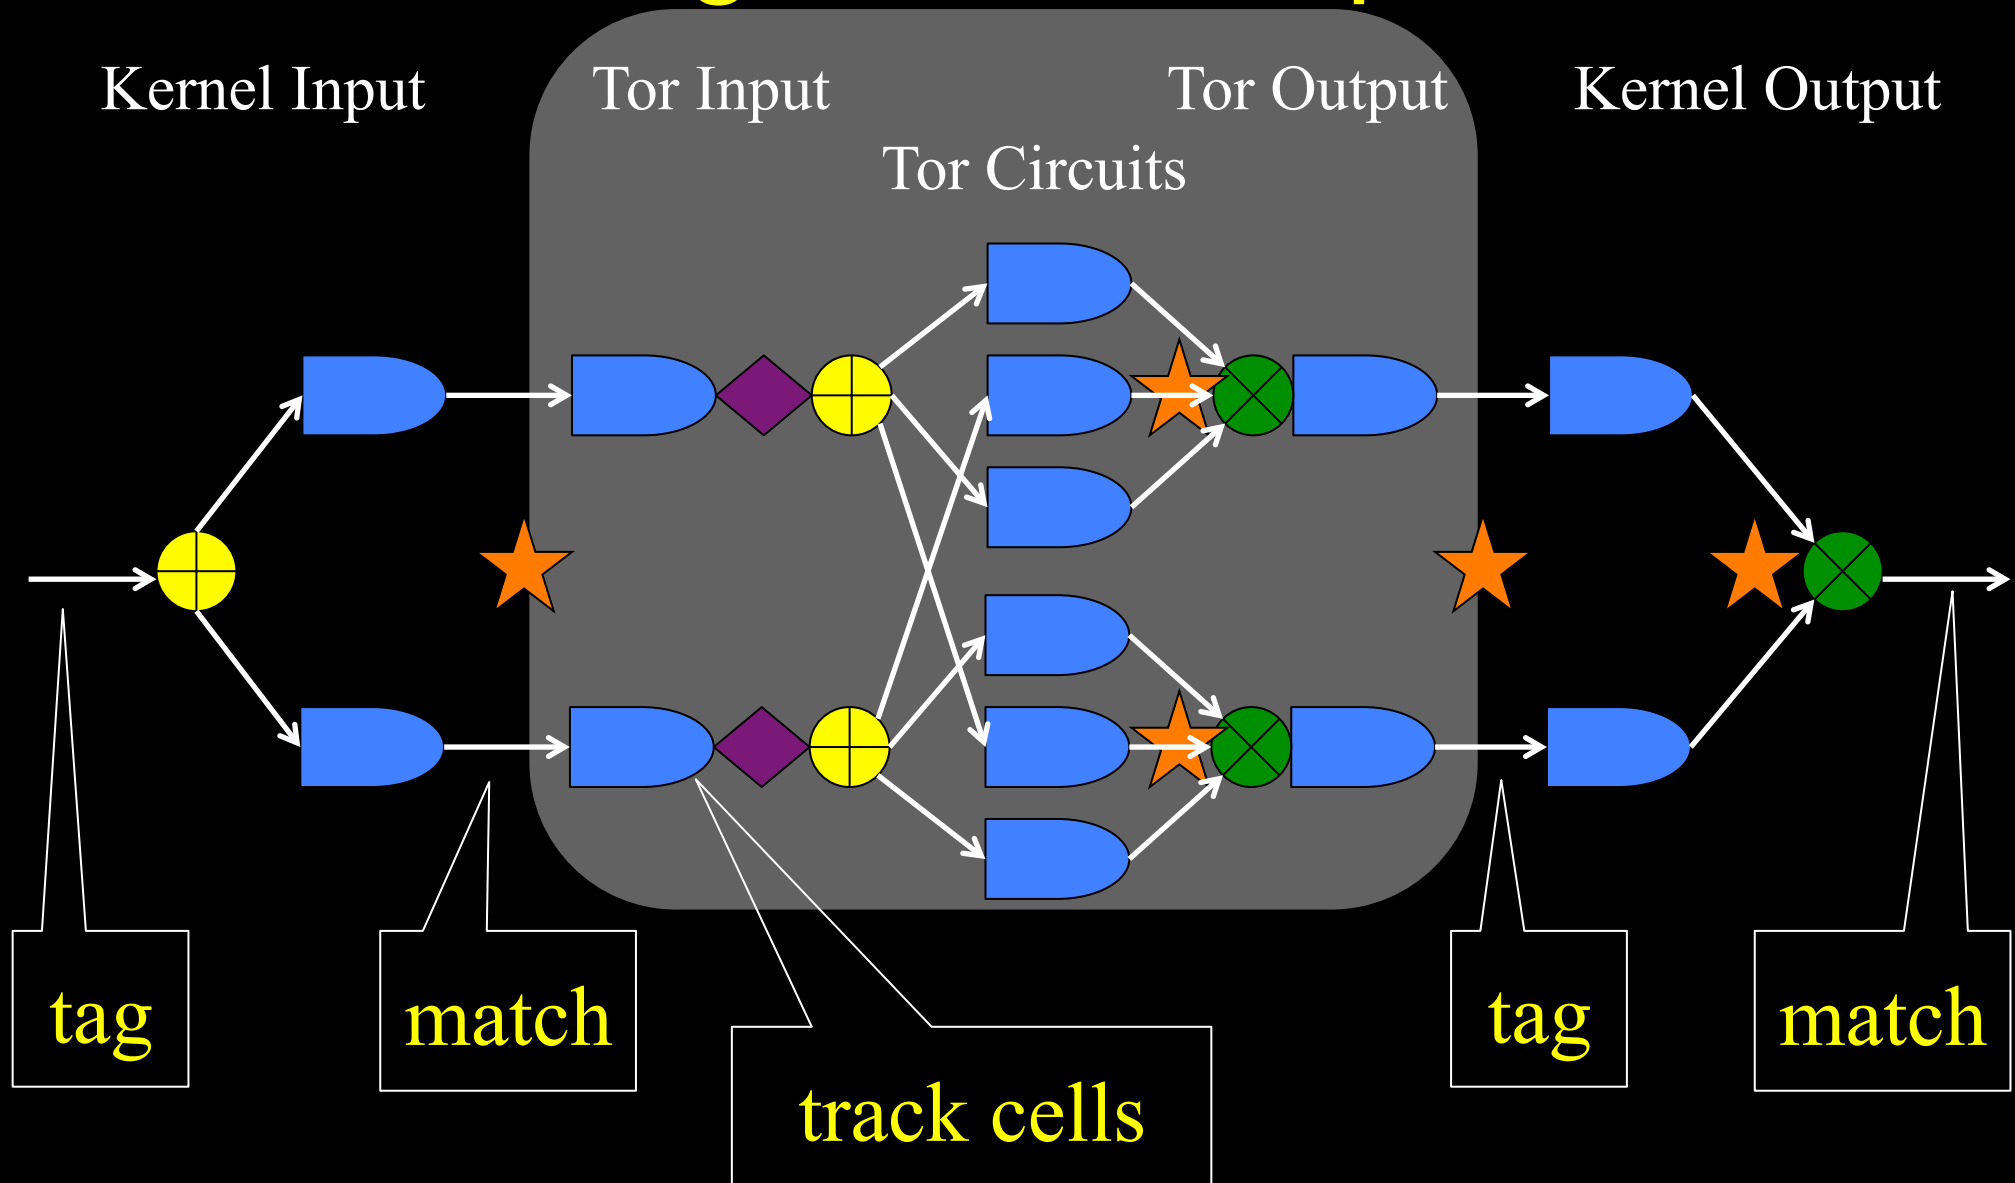
Kernel Congestion: libkqtime



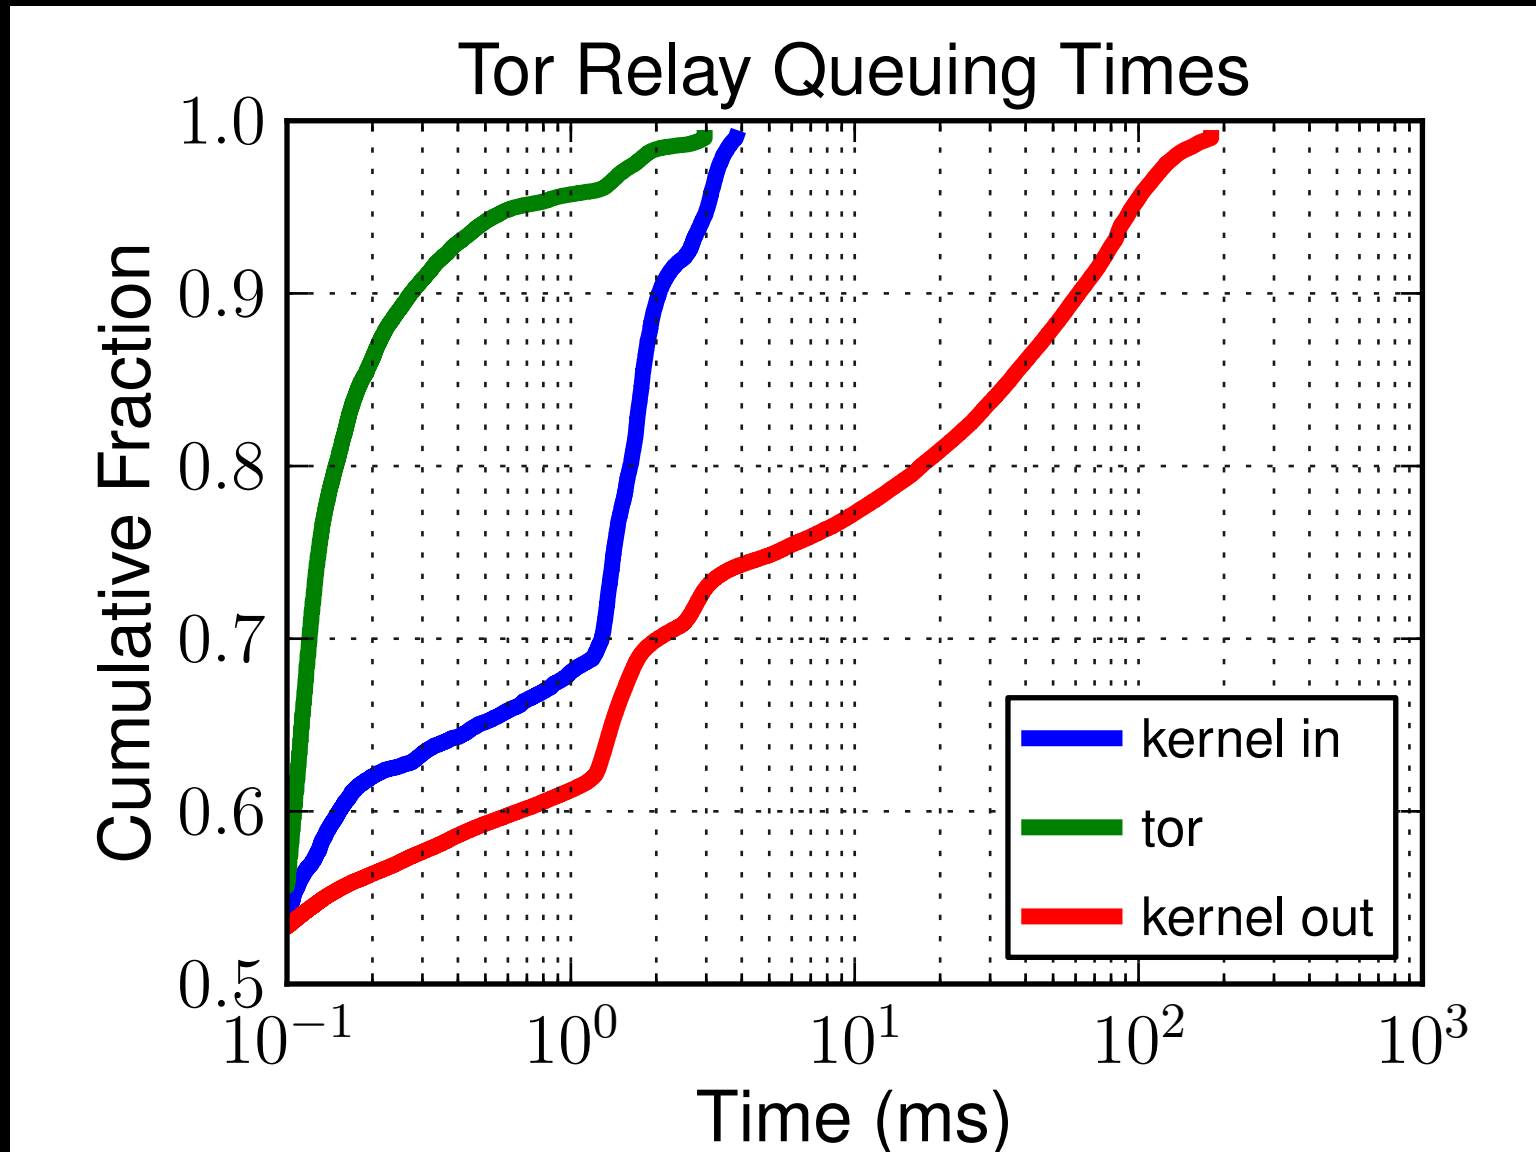
Kernel Congestion: libkqtime



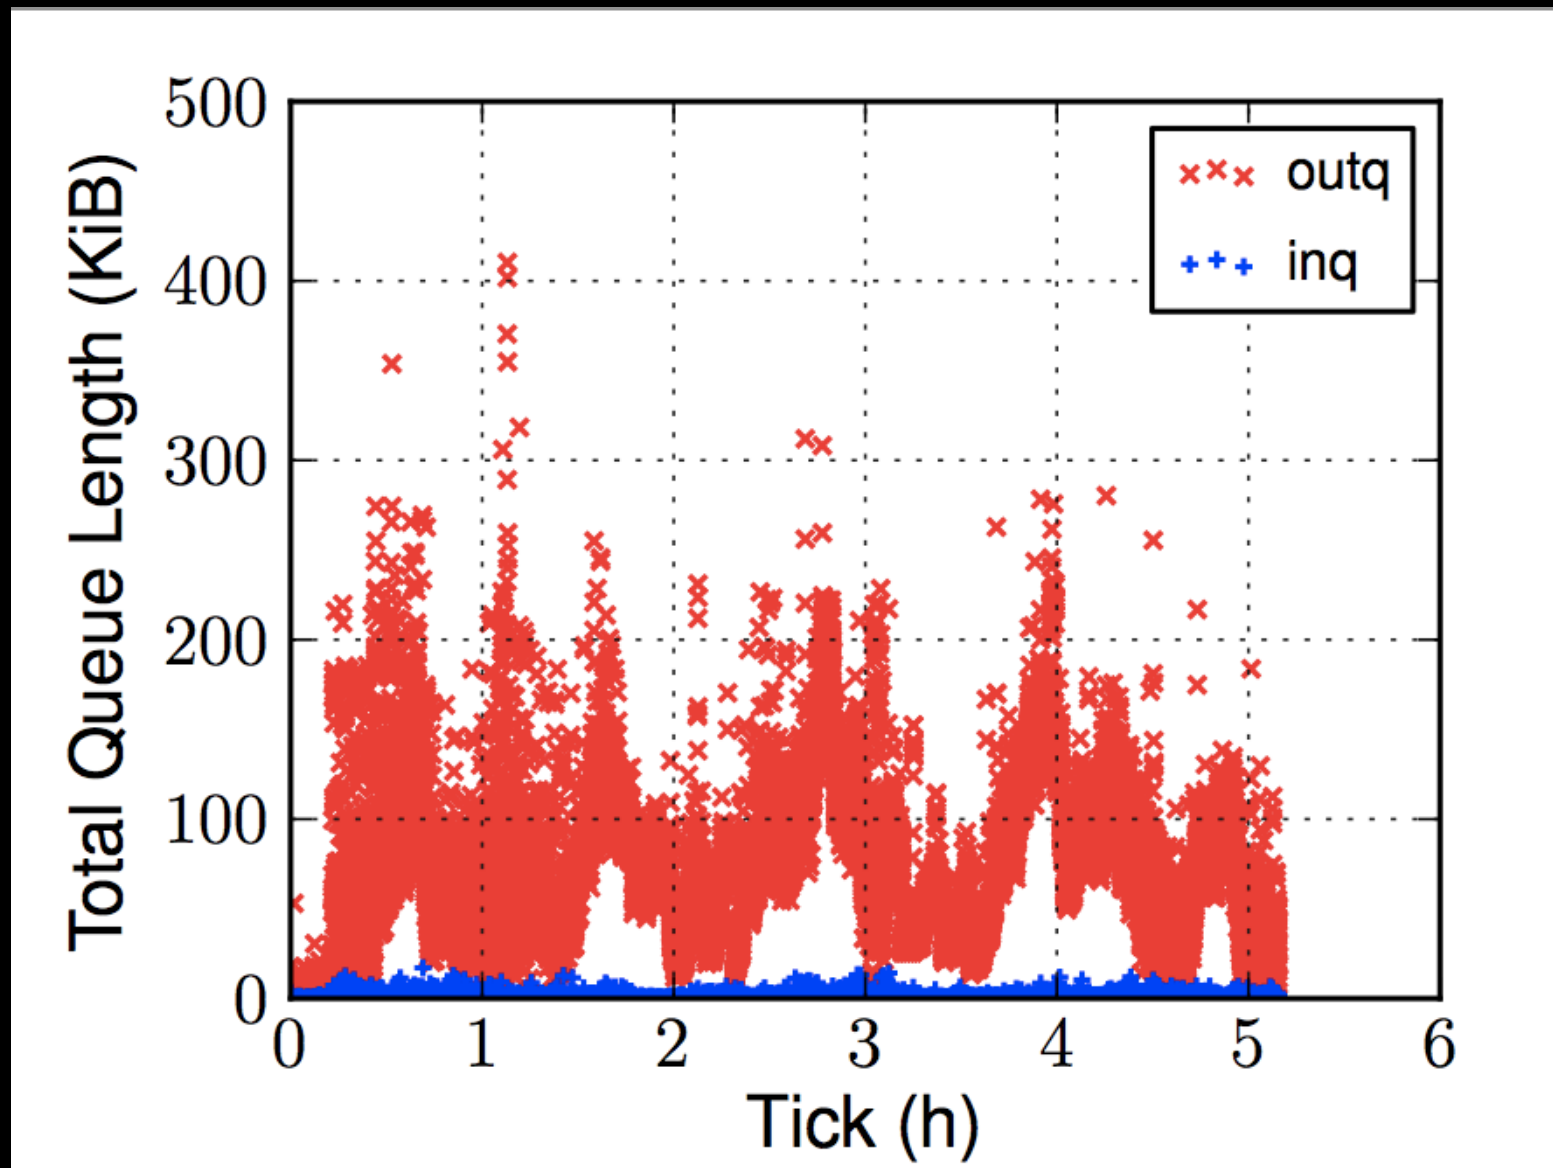
Kernel Congestion: libkqtime




Congestion Analysis



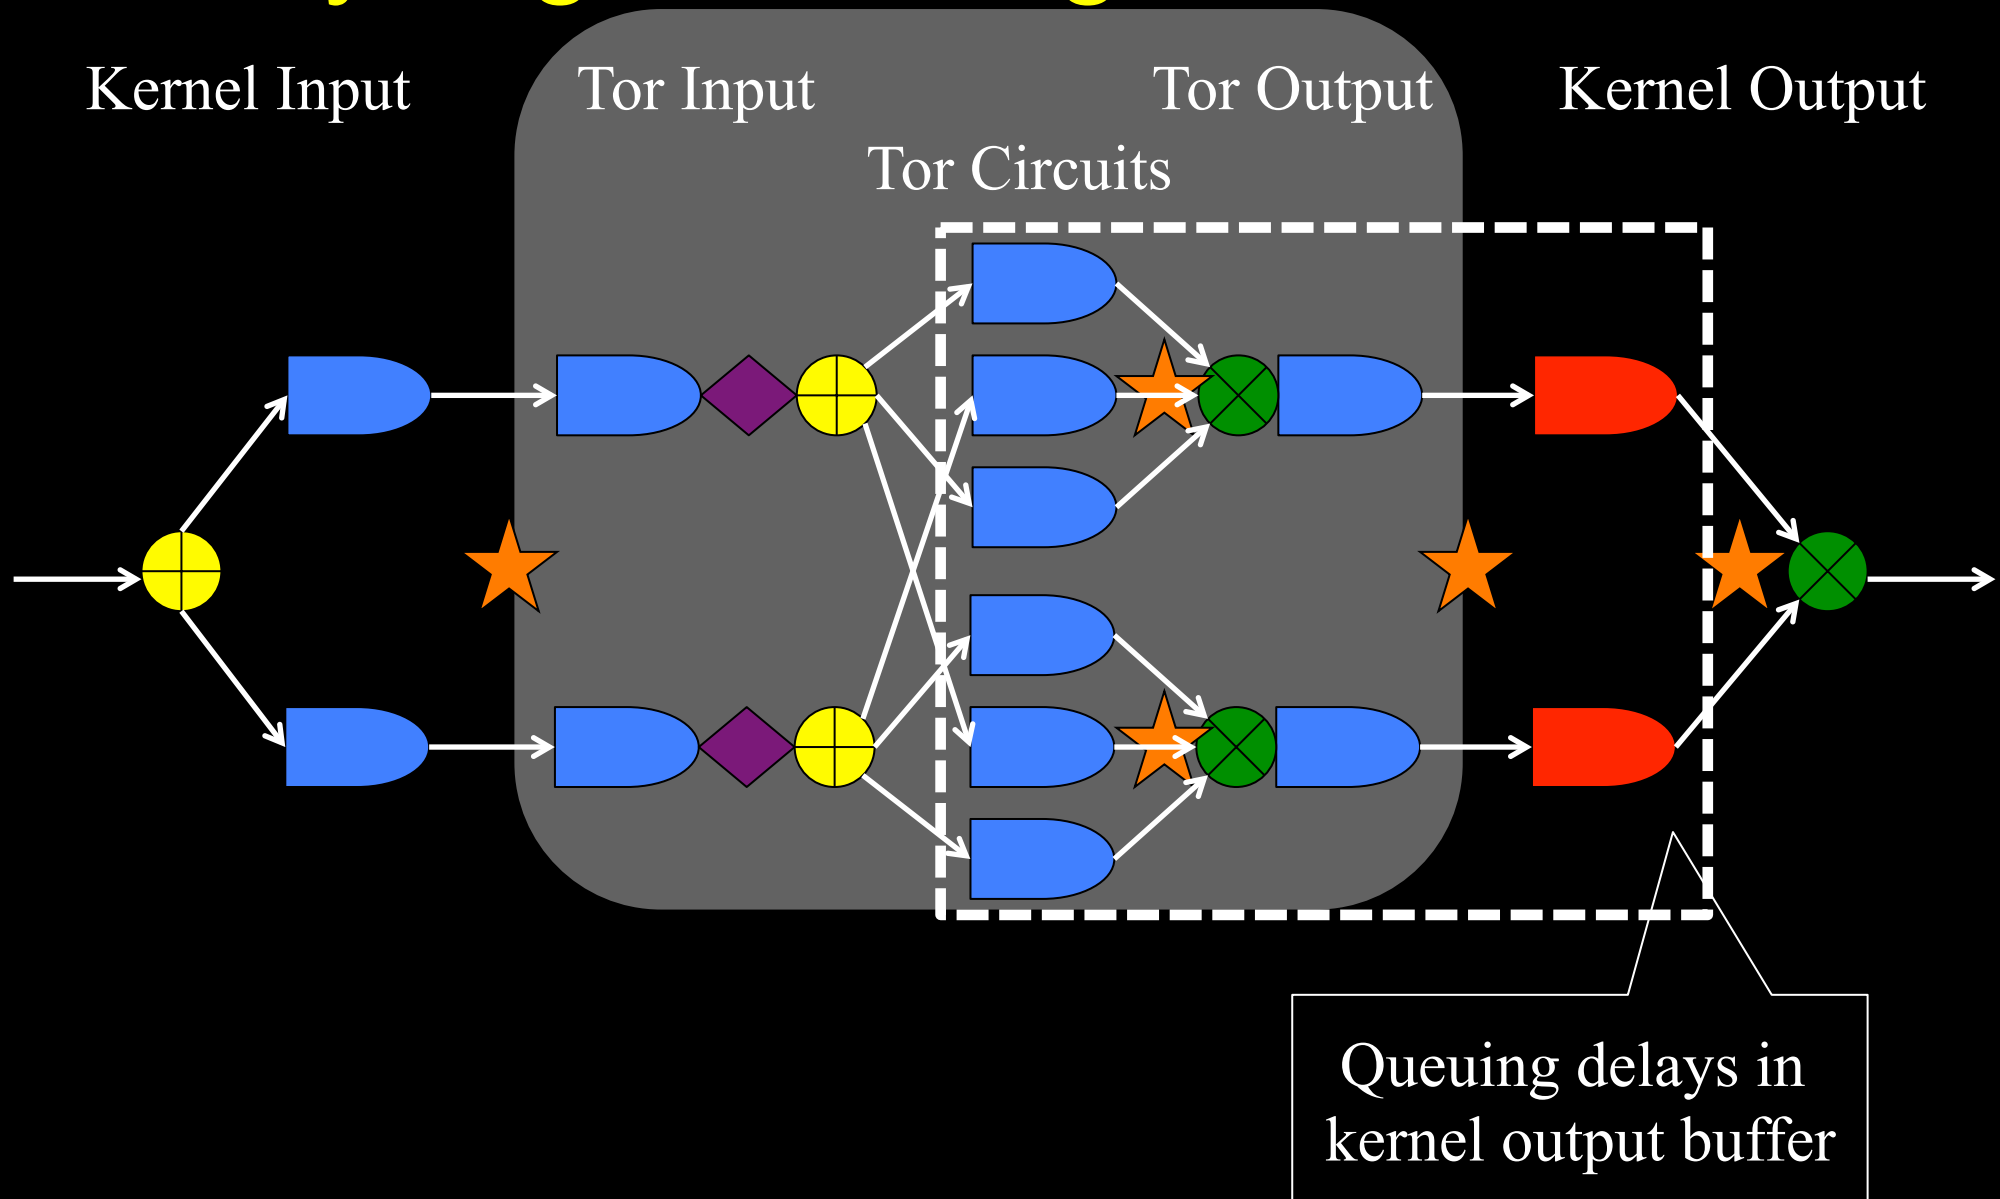
Congestion Analysis



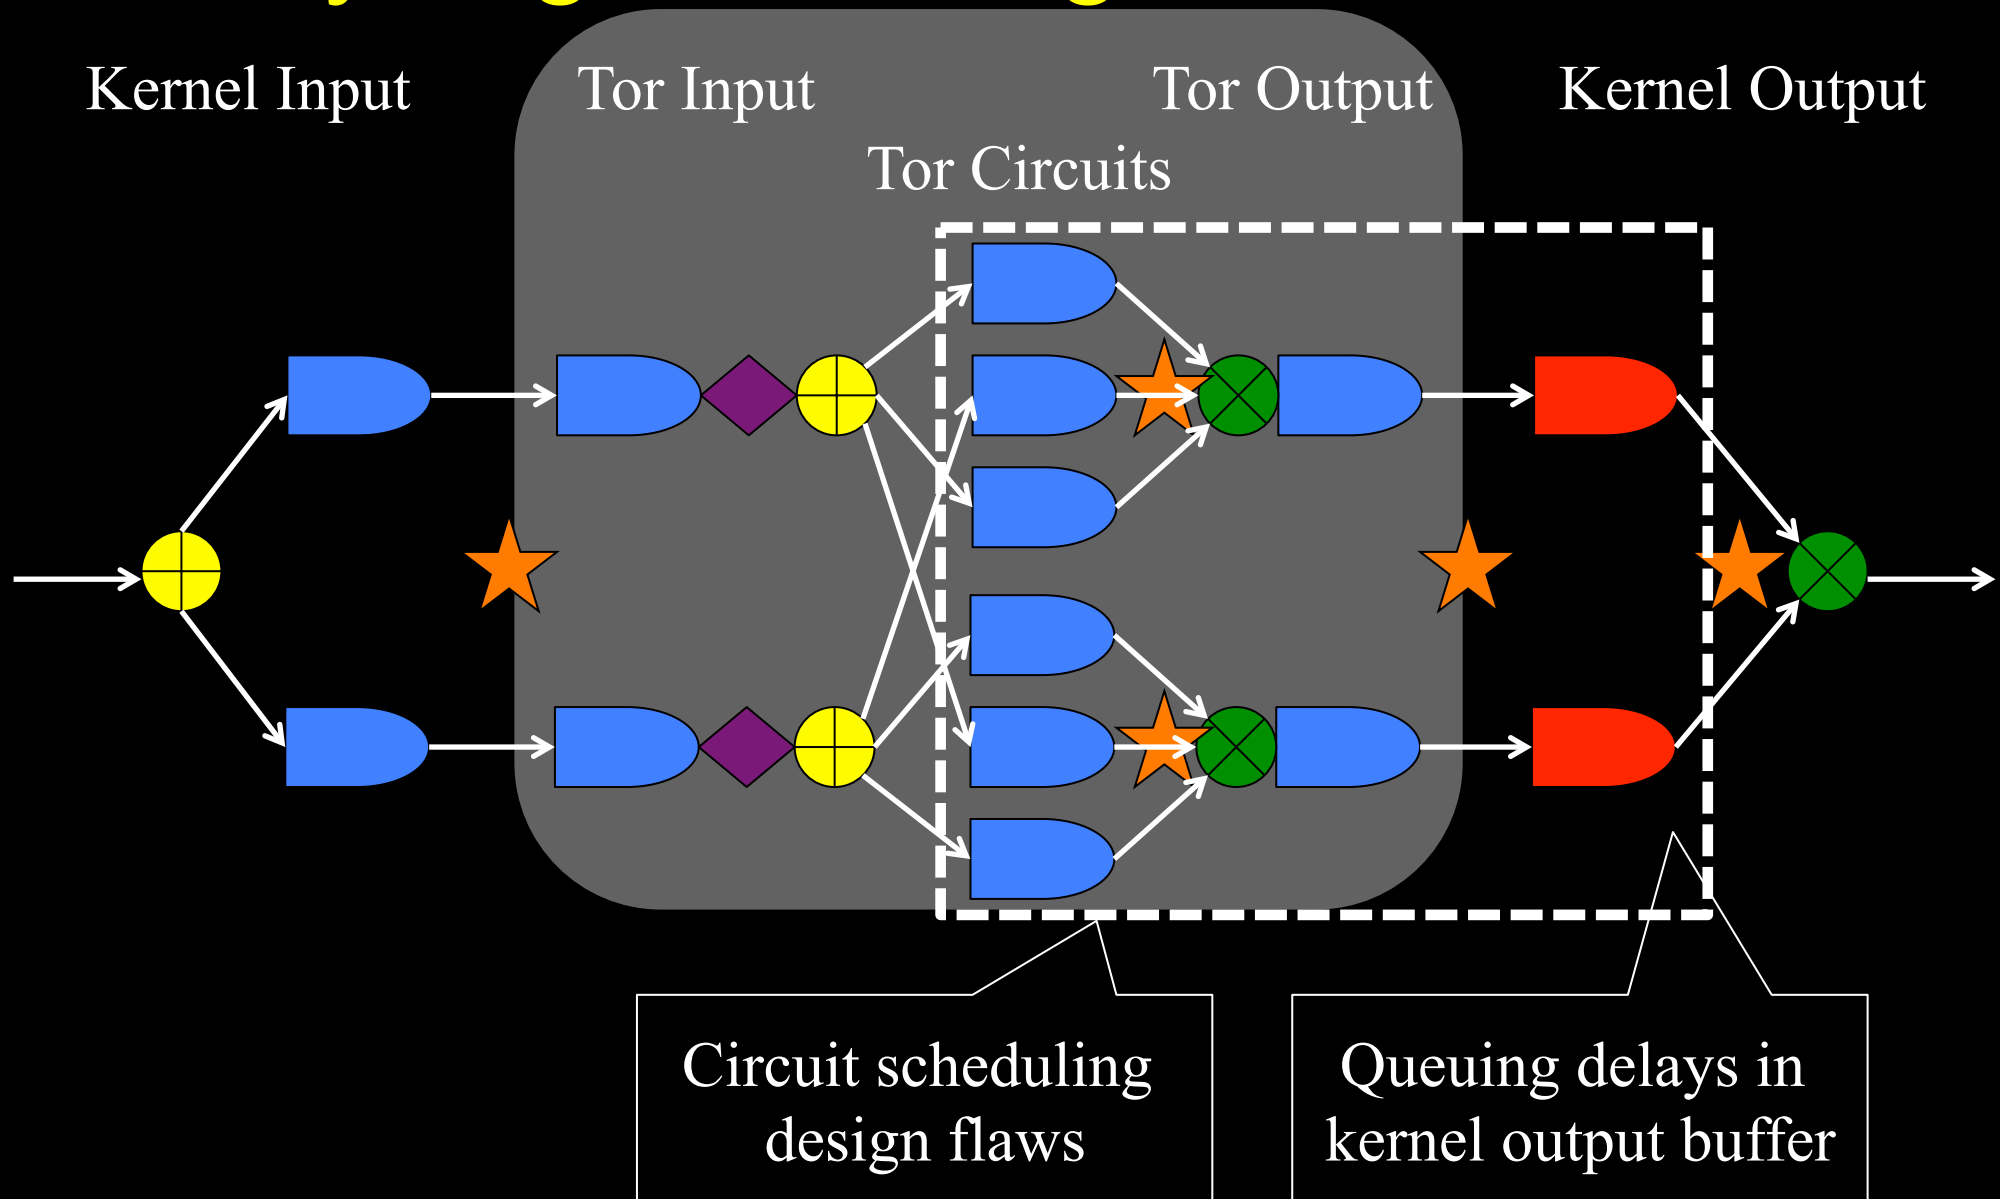
Analyzing the Design



Analyzing the Design





Analyzing the Design



Outline

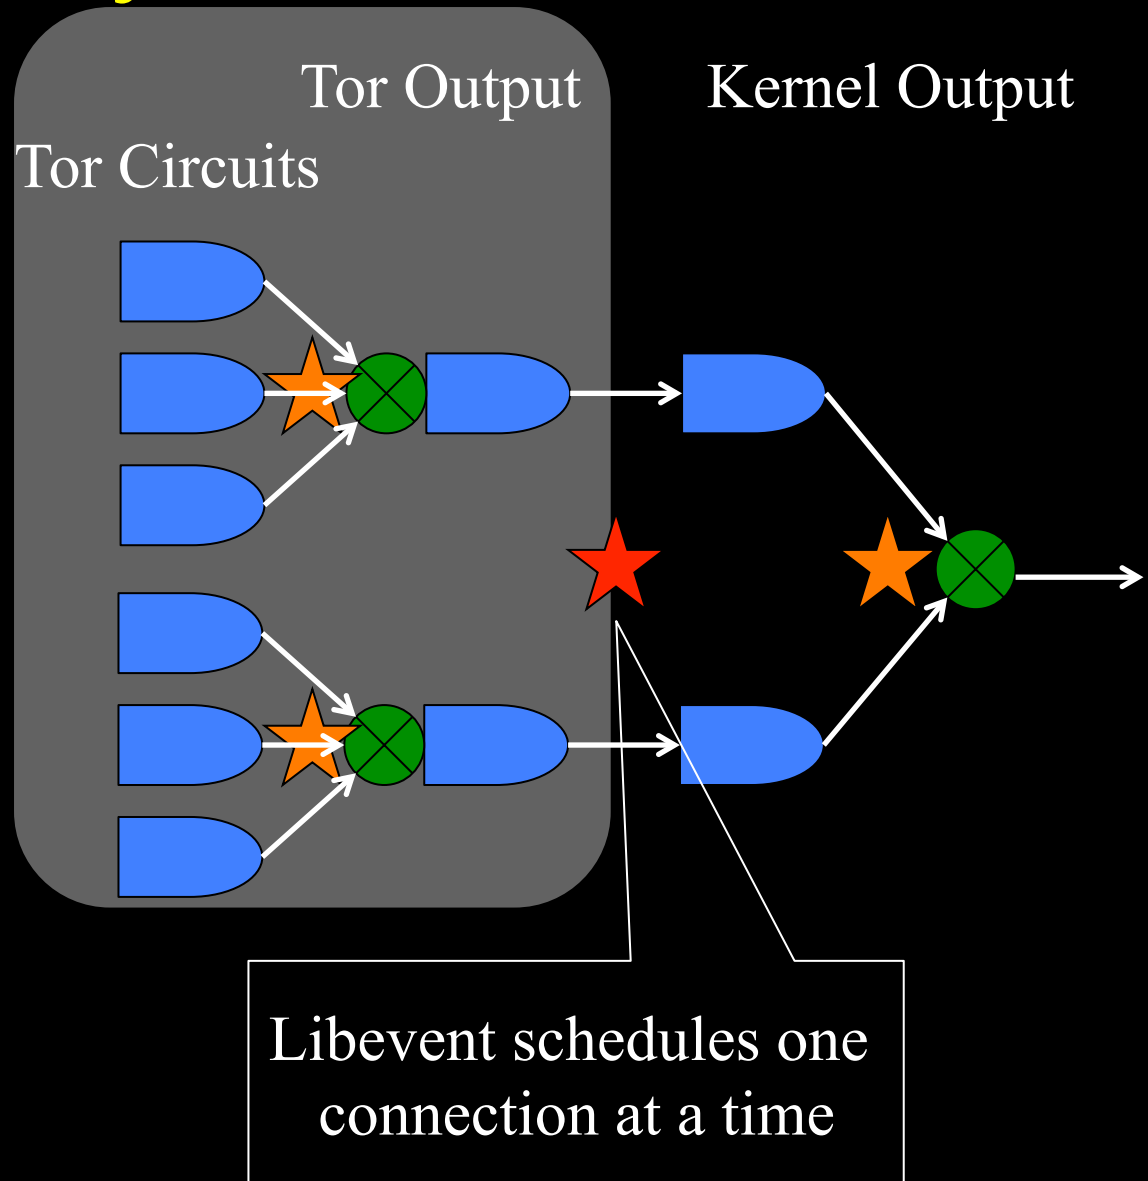
- Where is Tor slow?
 - Understand Tor relay architecture
 - Measure and analyze relay congestion in realistic Tor networks
- Design **focused** solutions

Ineffective Priority

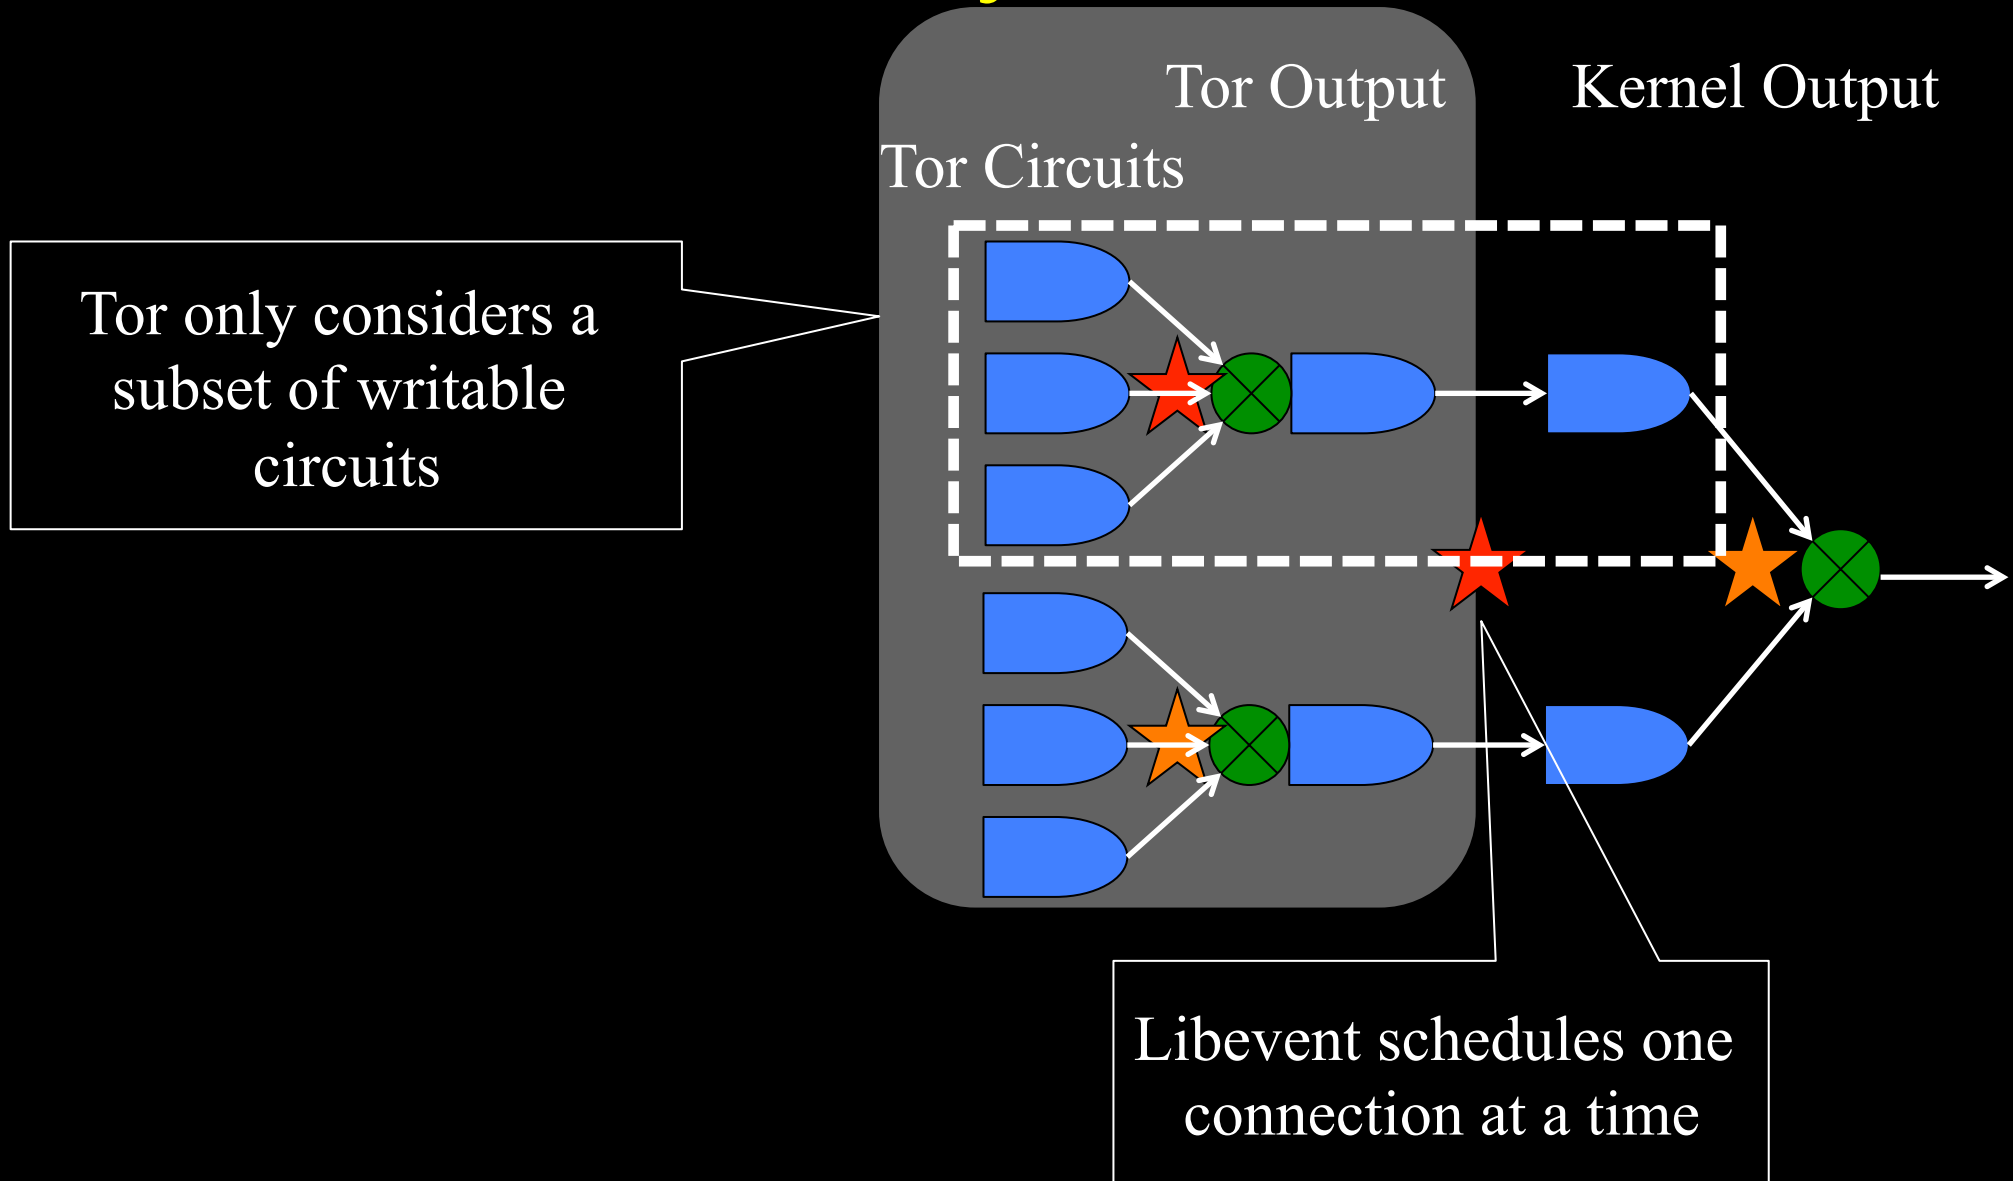


Circuit schedulers are ineffective at prioritization

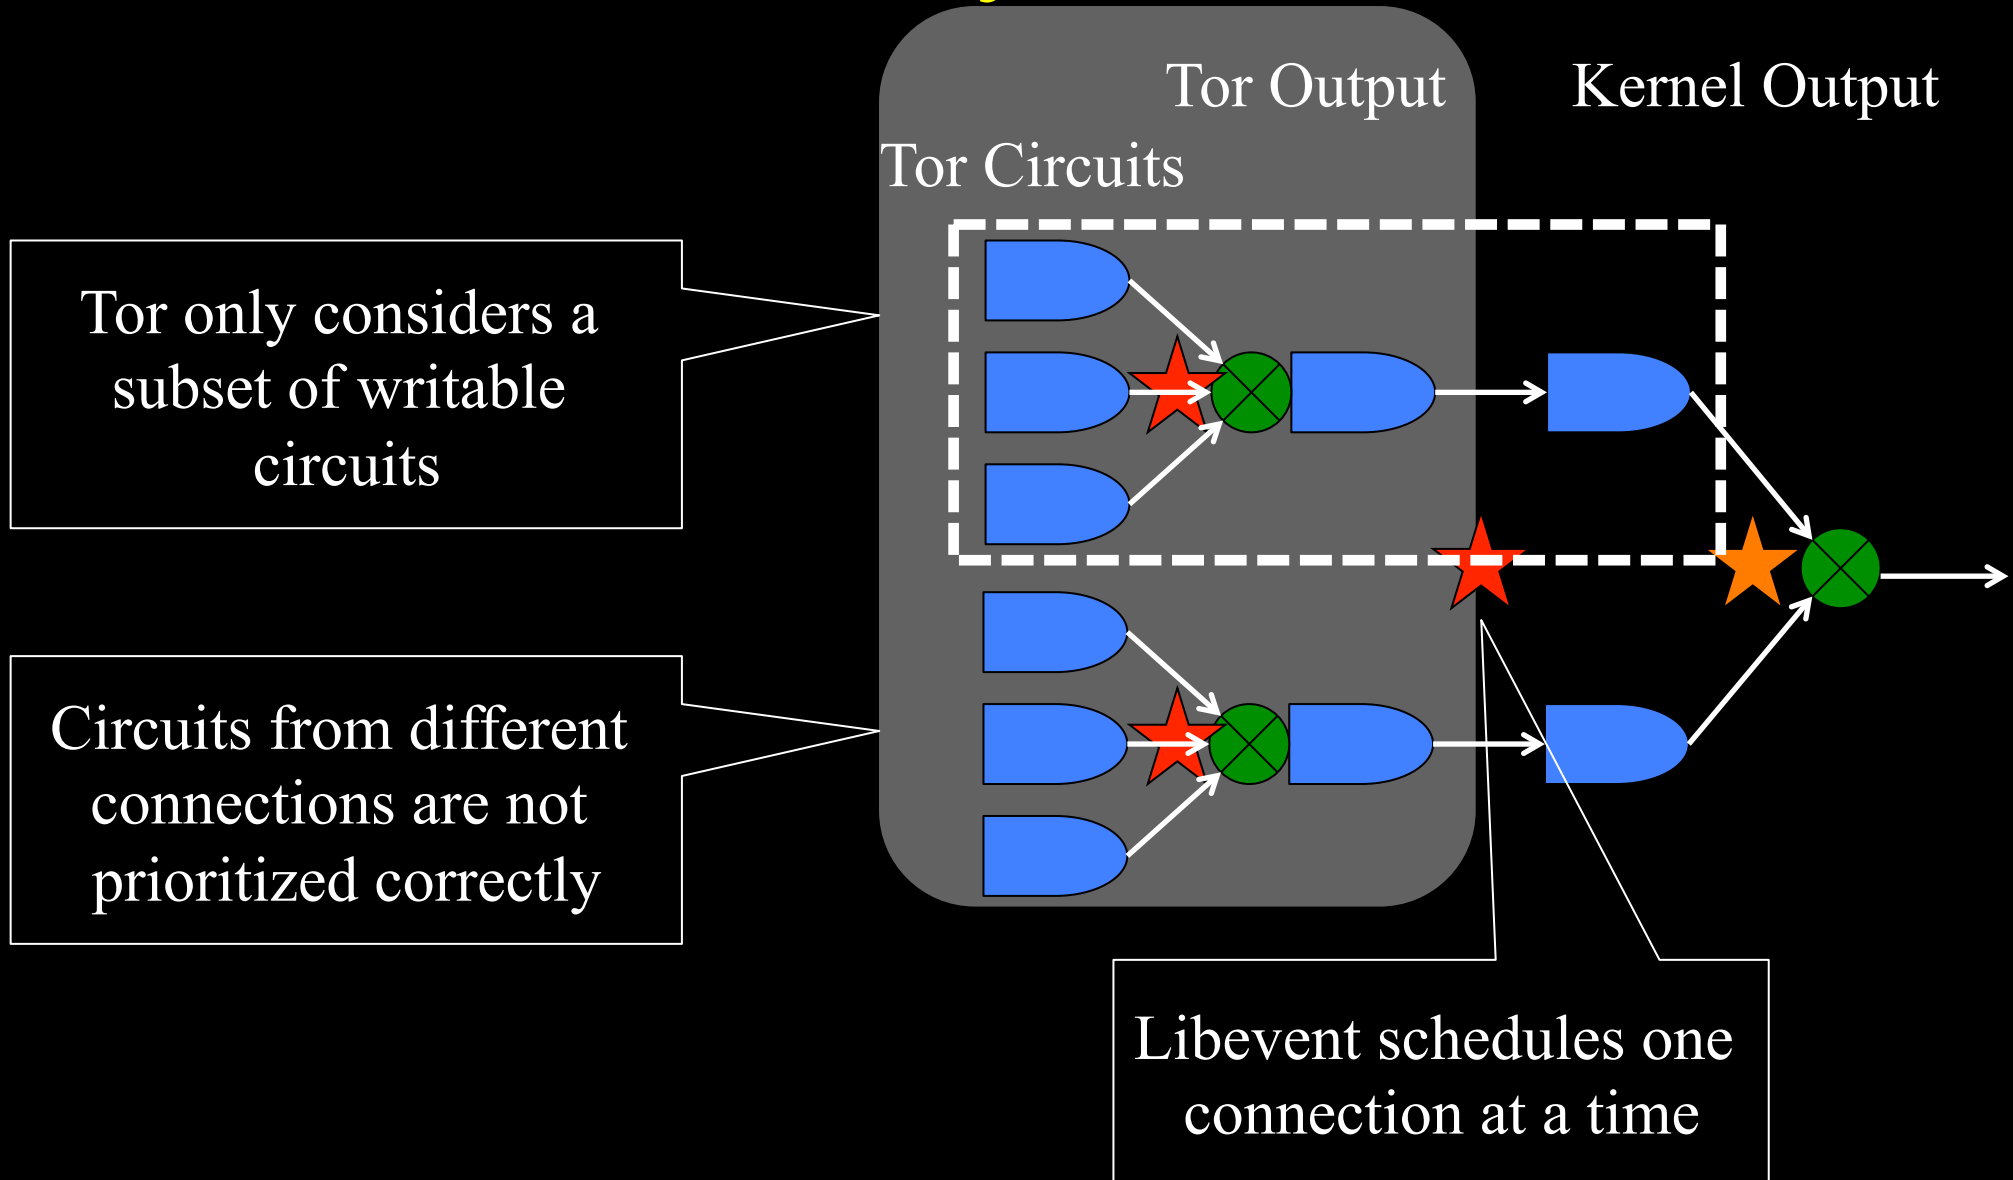
Ineffective Priority



Ineffective Priority




Ineffective Priority




Scheduling Problems

Scenario A



Shared Connection


Scenario B






No Shared Connection

Scheduling Problems

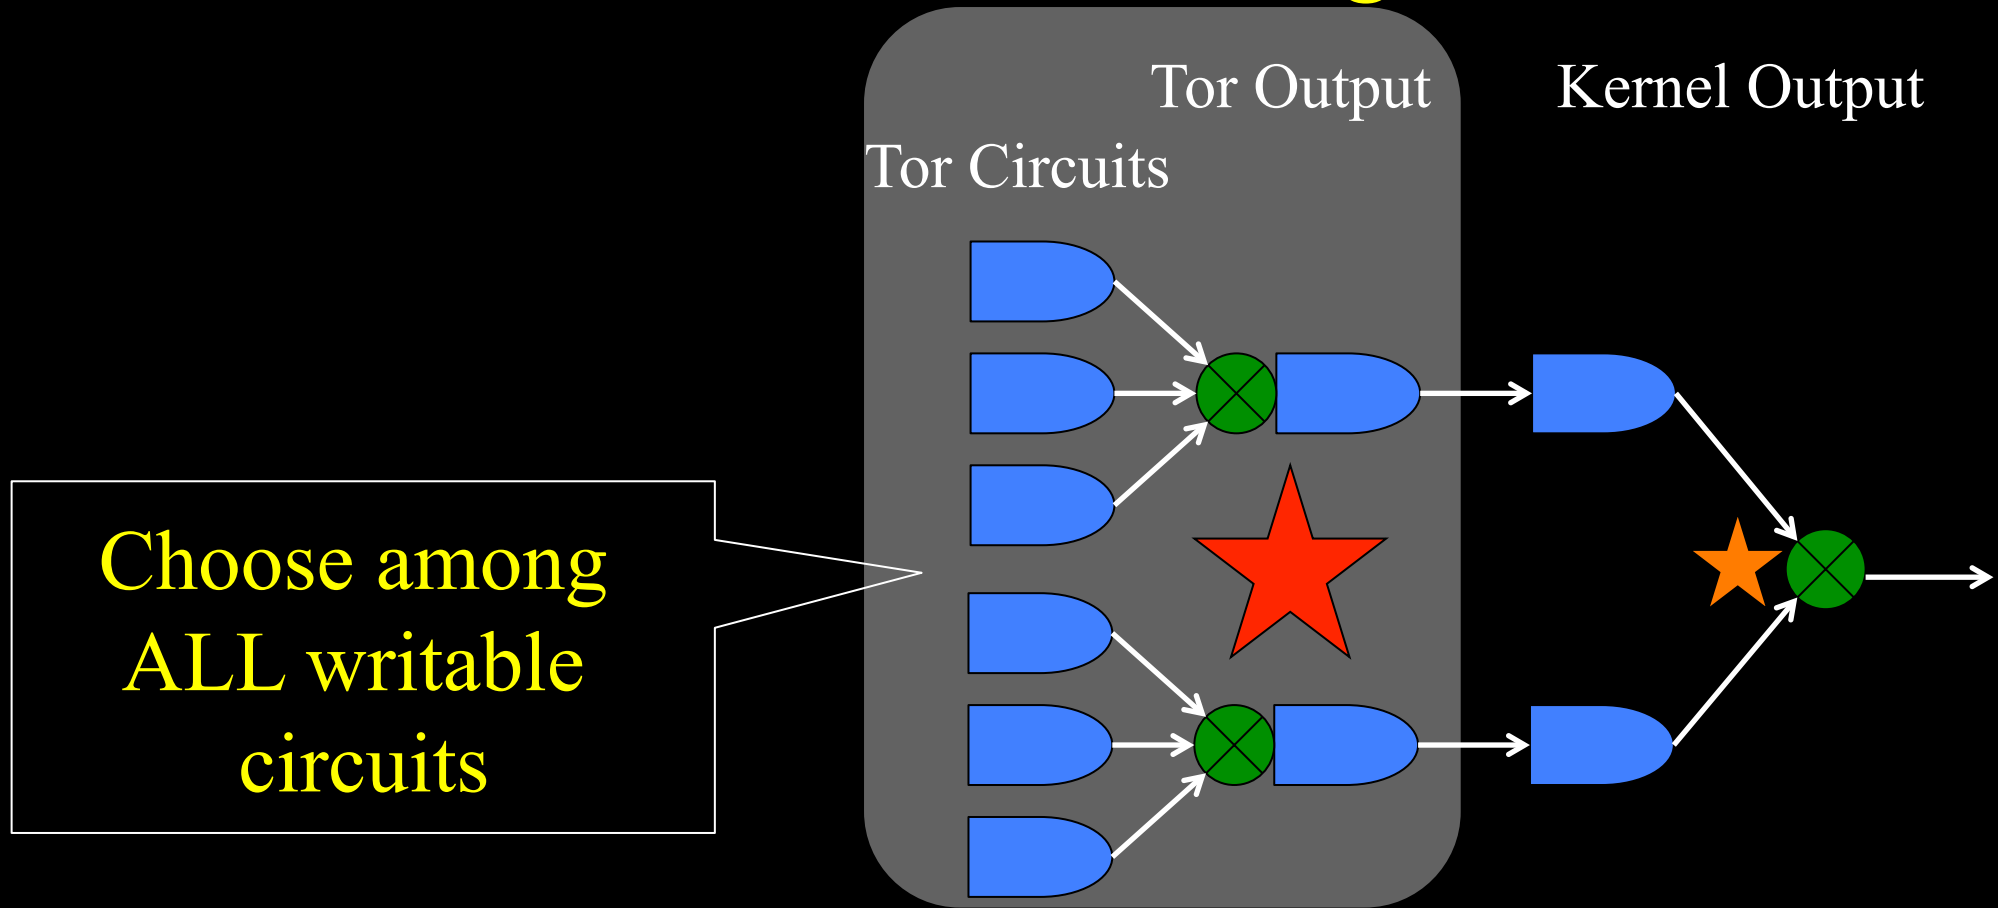
Scenario A




Scenario B



Global Circuit Scheduling




Kernel Buffer Bloat



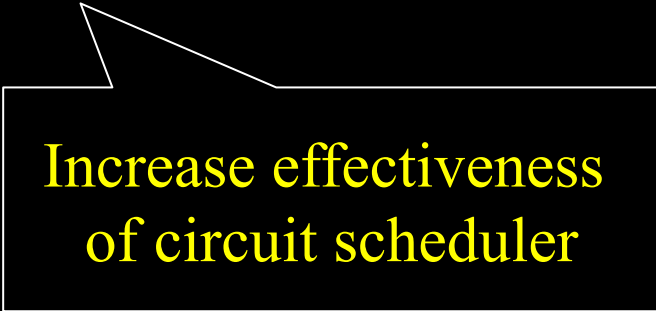
Kernel Buffer Bloat

- Too many large kernel queues
- More data in kernel than it can send
- Circuit scheduler timing issues



Tor Output Auto-tuning

- Don't write what the kernel can't send
- Smartly write to kernel using
 - Socket queue lengths and sizes
 - TCP windows
 - Node bandwidth capacity
- Check again before kernel starvation



Increase effectiveness
of circuit scheduler

Questions?

cs.umn.edu/~jansen
rob.g.jansen@nrl.navy.mil

think like an adversary



libkqtime

