Detecting Smart, Self-Propagating Internet Worms

Jun Li (presenter)
Shad Stafford

at:
IEEE CONFERENCE ON COMMUNICATIONS AND NETWORK SECURITY
October 30, 2014

UNIVERSITY OF OREGON
Network & Security Research Laboratory

Outline

- Introduction
- SWORD — Self-propagating Worm Observation and Rapid Detection
- Experiment Methodology
- SWORD Performance Against Classic Worms
- SWORD Performance Against Smart Worms
- Concluding remarks

Network worm

- A network worm is a program that actively tries to copy itself to other hosts across network connections
- The environment for worms is getting more fertile
  - >700,000 Android devices activated per day
  - 100 million iCloud users
  - 25 **billion** iOS App Store downloads
- We need to be prepared for worm outbreaks

Worm defense

- Prevention
  - Design software and hardware with no vulnerabilities
  - Ensure all software is deployed with configuration that does not allow worm propagation
- Detection
  - Find worm outbreaks as they occur
- Disinfection
  - Isolate infected hosts or remove worm code from them

Worm detector taxonomy

Worm Detection

- **Host-based**
  - Buffer Overflow Detection
  - Input Correlation
  - System Calls
- **Honeypot-based**
  - Honeypot
- **Content-based**
  - Static Signature
  - Dynamic Signature
  - Advanced Signature
  - Protocol Field Length
- **Behavior-based**
  - Network Telescope
  - Connection Failures
  - Destination Addresses
  - Causation

Behavior-based detectors

TRW
RBS
TRWRBS
MRW
DSC
PGD

- Deployable at a single point (gateway)
- Does not require access to connection payload
  - content independent

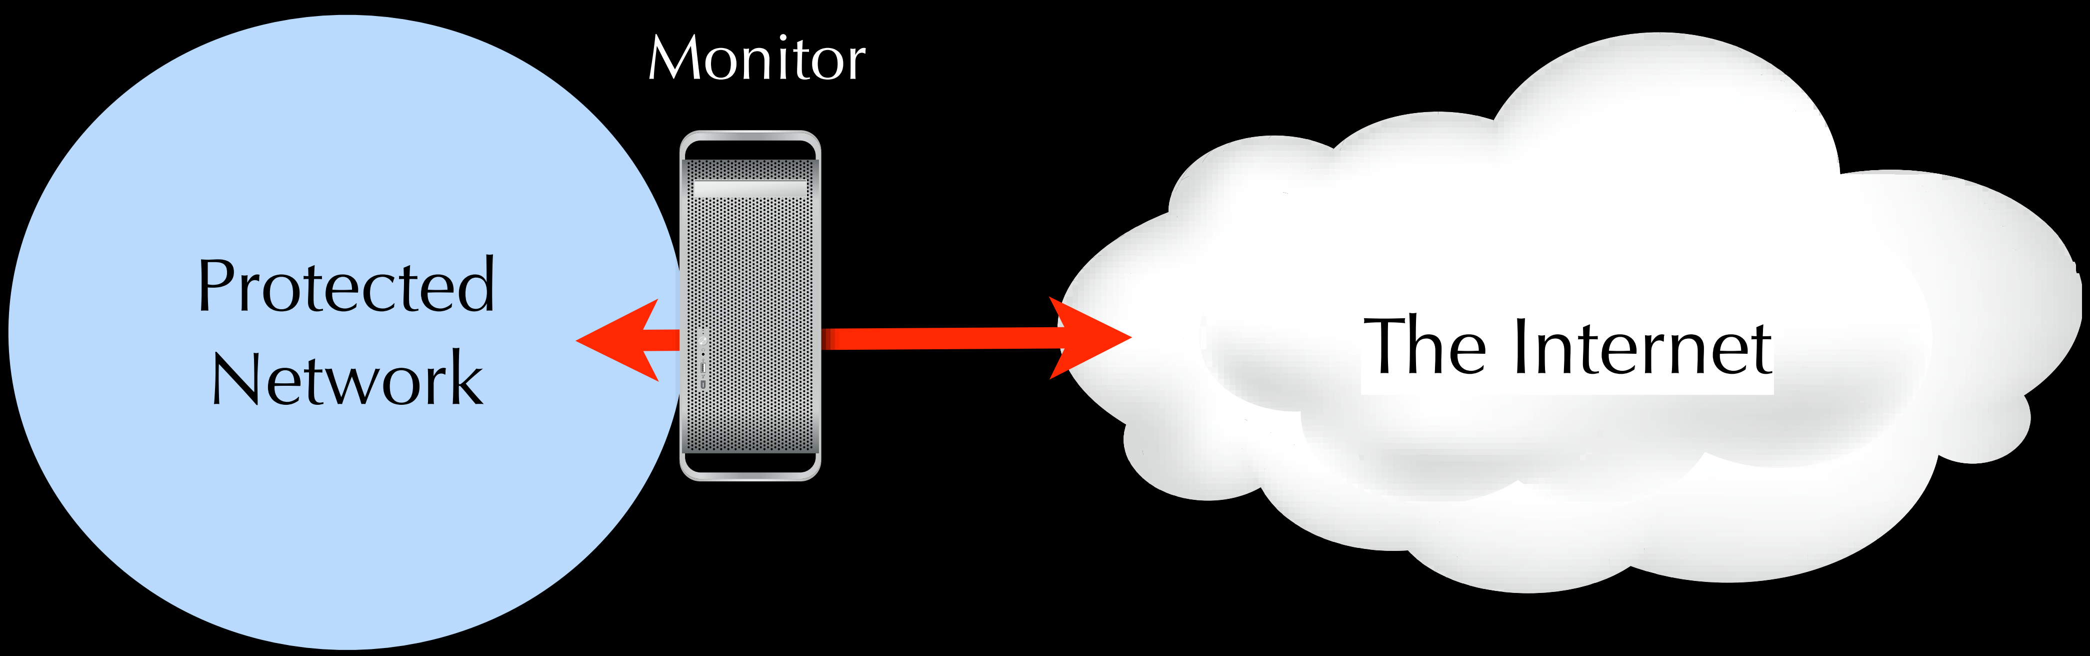

Smart, evasive Worm

- These detectors all assumed naive worm
  - more evaluation results later
- But worm may be able to avoid expressing certain behavior traits
  - worm authors not likely to be so simplistic
- Worm detectors must be effective even against smart, evasive worms

Outline

- Introduction
- SWORD — Self-propagating Worm Observation and Rapid Detection
- Experiment Methodology
- SWORD Performance Against Classic Worms
- SWORD Performance Against Smart Worms
- Concluding remarks

Design principles

- Effective detectors must target *essential behaviors* of worms
- Recall a self-propagating worm is defined as code that scans the network to find and infect new hosts
- We believe the only truly essential behavior of worms is that of connecting to new destinations

SWORD detector

- Two main modules: a Burst Duration Detector (BDD) and a Quiescent Period Detector (QPD)
- Two modules complement each other: if a worm does not violate one module, it will violate the other

Burst-duration detector (BDD)

- Self-propagating worms must contact new destinations
  - cannot be avoided as a worm seeks new victims
- BDD determines if a host is making first-contact connections faster than usual
- Every different size of burst has a threshold: the minimum duration learned during the training phase

Burst-duration detector (BDD)

- Self-propagating worms must contact new destinations
  - cannot be avoided as a worm seeks new victims
- BDD determines if a host is making first-contact connections faster than usual
- Every different size of burst has a threshold: the minimum duration learned during the training phase

Burst-duration detector (BDD)

- Self-propagating worms must contact new destinations
  - cannot be avoided as a worm seeks new victims
- BDD determines if a host is making first-contact connections faster than usual
- Every different size of burst has a threshold: the minimum duration learned during the training phase

Burst-duration detector (BDD)

- Self-propagating worms must contact new destinations
  - cannot be avoided as a worm seeks new victims
- BDD determines if a host is making first-contact connections faster than usual
- Every different size of burst has a threshold: the minimum duration learned during the training phase

Burst-duration detector (BDD)

- Self-propagating worms must contact new destinations
  - cannot be avoided as a worm seeks new victims
- BDD determines if a host is making first-contact connections faster than usual
- Every different size of burst has a threshold: the minimum duration learned during the training phase

Burst-duration detector (BDD)

- Self-propagating worms must contact new destinations
  - cannot be avoided as a worm seeks new victims
- BDD determines if a host is making first-contact connections faster than usual
- Every different size of burst has a threshold: the minimum duration learned during the training phase

Burst-duration detector (BDD)

- Self-propagating worms must contact new destinations
  - cannot be avoided as a worm seeks new victims
- BDD determines if a host is making first-contact connections faster than usual
- Every different size of burst has a threshold: the minimum duration learned during the training phase

Burst-duration detector (BDD)

- Self-propagating worms must contact new destinations
  - cannot be avoided as a worm seeks new victims
- BDD determines if a host is making first-contact connections faster than usual
- Every different size of burst has a threshold: the minimum duration learned during the training phase

Burst-duration detector (BDD)

- Self-propagating worms must contact new destinations
  - cannot be avoided as a worm seeks new victims
- BDD determines if a host is making first-contact connections faster than usual
- Every different size of burst has a threshold: the minimum duration learned during the training phase

Burst-duration detector (BDD)

- Self-propagating worms must contact new destinations
  - cannot be avoided as a worm seeks new victims
- BDD determines if a host is making first-contact connections faster than usual
- Every different size of burst has a threshold: the minimum duration learned during the training phase

Burst-duration detector (BDD)

- Self-propagating worms must contact new destinations
  - cannot be avoided as a worm seeks new victims
- BDD determines if a host is making first-contact connections faster than usual
- Every different size of burst has a threshold: the minimum duration learned during the training phase

Burst-duration detector (BDD)

- Self-propagating worms must contact new destinations
  - cannot be avoided as a worm seeks new victims
- BDD determines if a host is making first-contact connections faster than usual
- Every different size of burst has a threshold: the minimum duration learned during the training phase

Burst-duration detector (BDD)

- Self-propagating worms must contact new destinations
  - cannot be avoided as a worm seeks new victims
- BDD determines if a host is making first-contact connections faster than usual
- Every different size of burst has a threshold: the minimum duration learned during the training phase

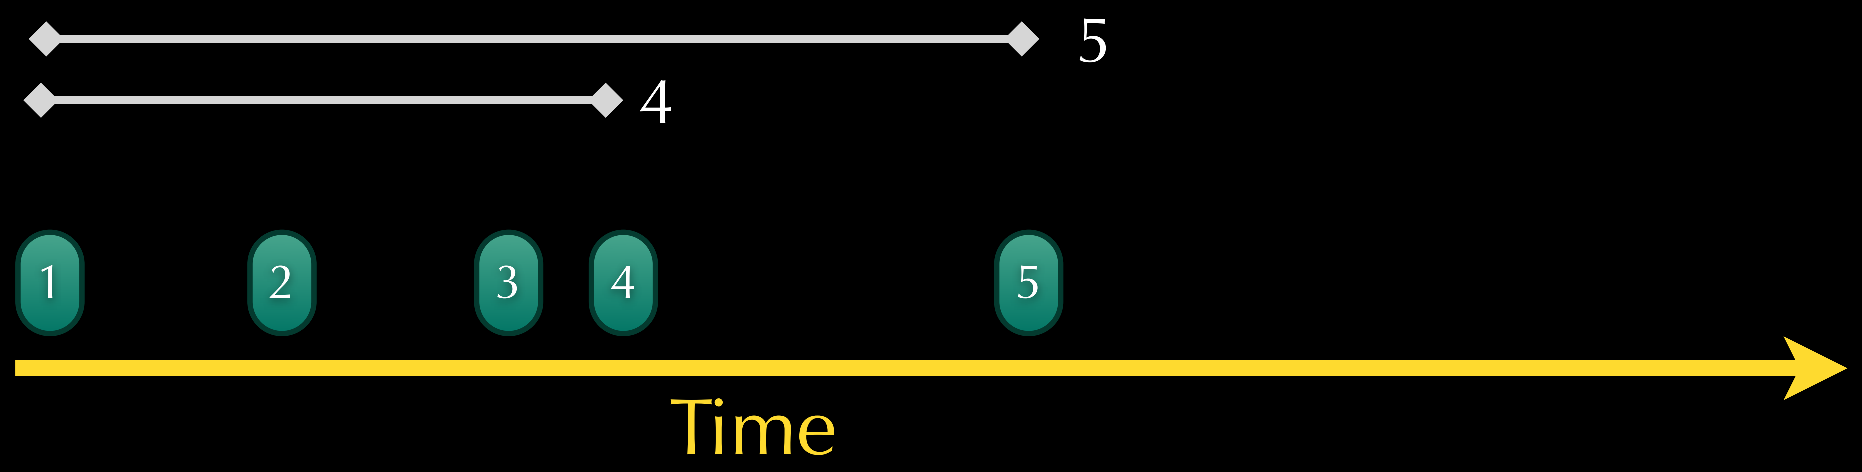

Burst-duration detector (BDD)

- Self-propagating worms must contact new destinations
  - cannot be avoided as a worm seeks new victims
- BDD determines if a host is making first-contact connections faster than usual
- Every different size of burst has a threshold: the minimum duration learned during the training phase

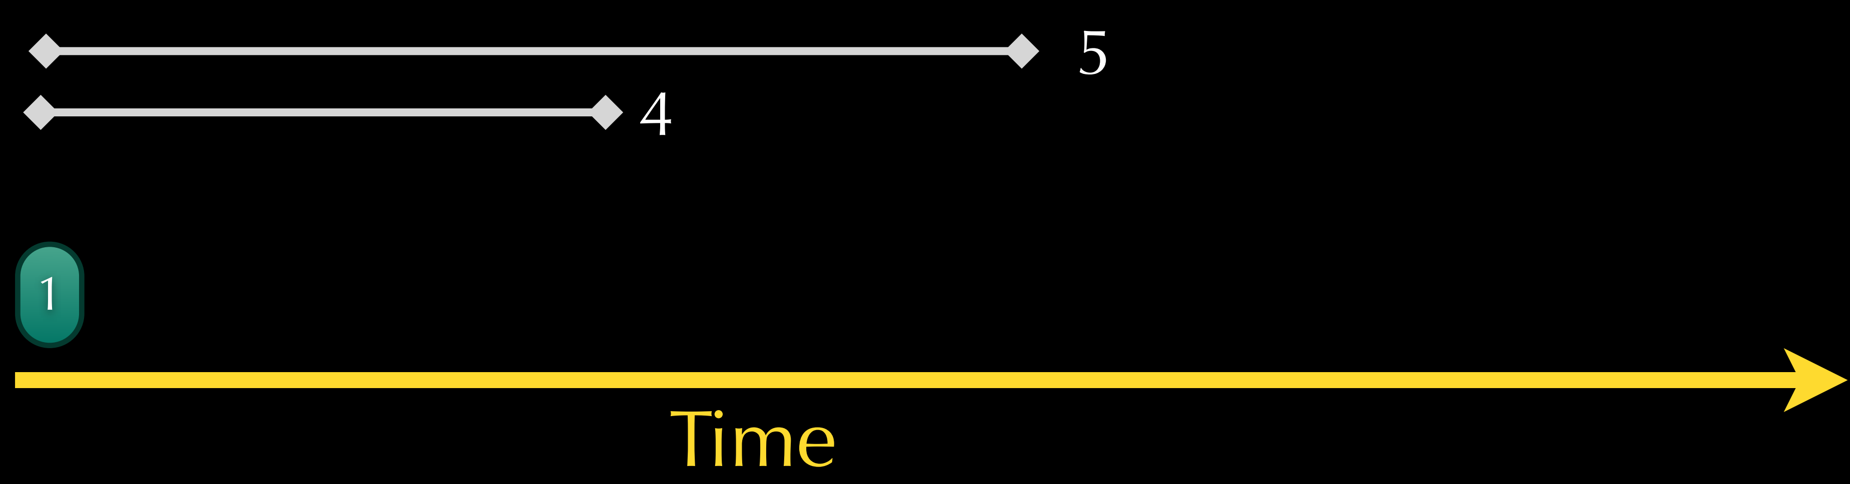

Burst-duration detector (BDD)

- Self-propagating worms must contact new destinations
  - cannot be avoided as a worm seeks new victims
- BDD determines if a host is making first-contact connections faster than usual
- Every different size of burst has a threshold: the minimum duration learned during the training phase

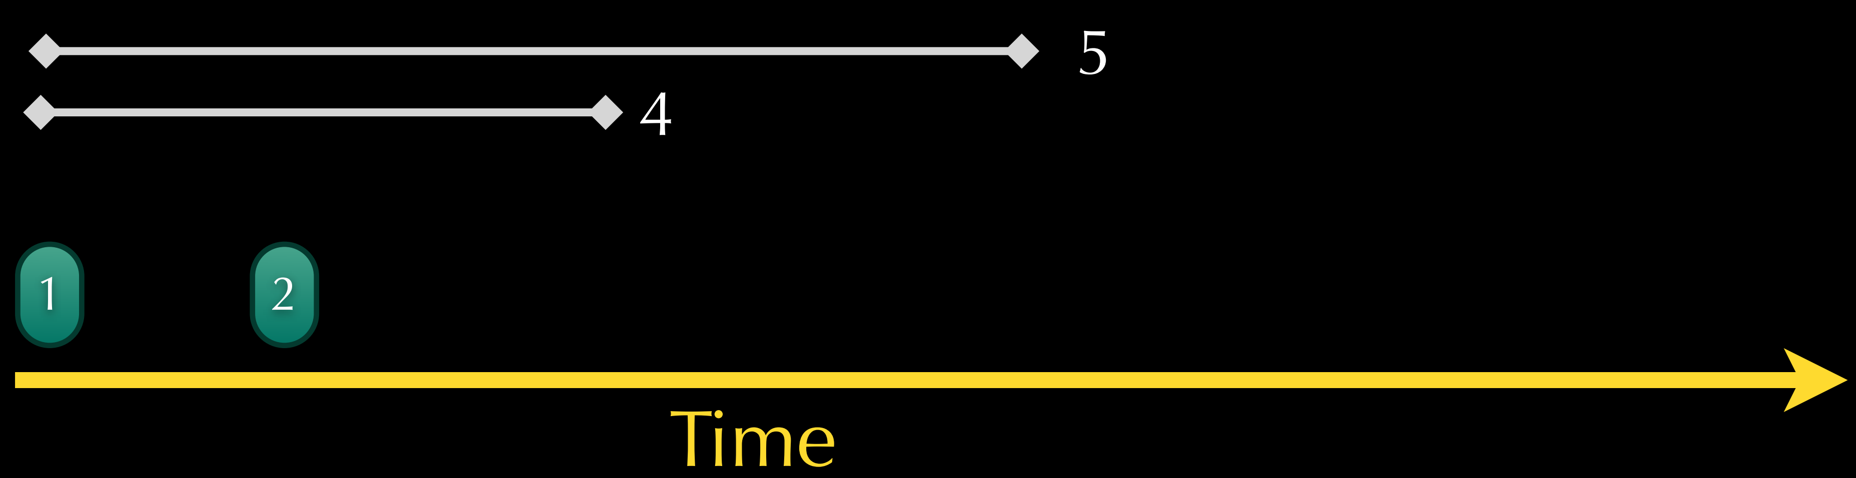

Burst-duration detector (BDD)

- Self-propagating worms must contact new destinations
  - cannot be avoided as a worm seeks new victims
- BDD determines if a host is making first-contact connections faster than usual
- Every different size of burst has a threshold: the minimum duration learned during the training phase

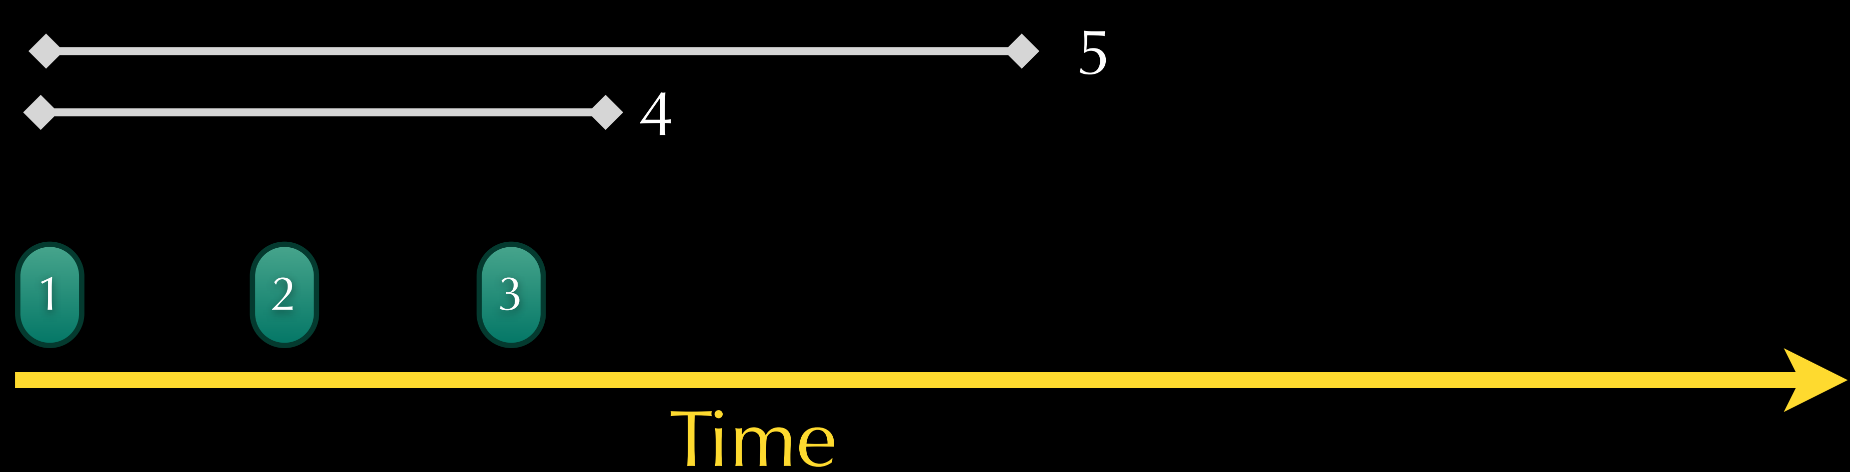

Burst-duration detector (BDD)

- Self-propagating worms must contact new destinations
  - cannot be avoided as a worm seeks new victims
- BDD determines if a host is making first-contact connections faster than usual
- Every different size of burst has a threshold: the minimum duration learned during the training phase

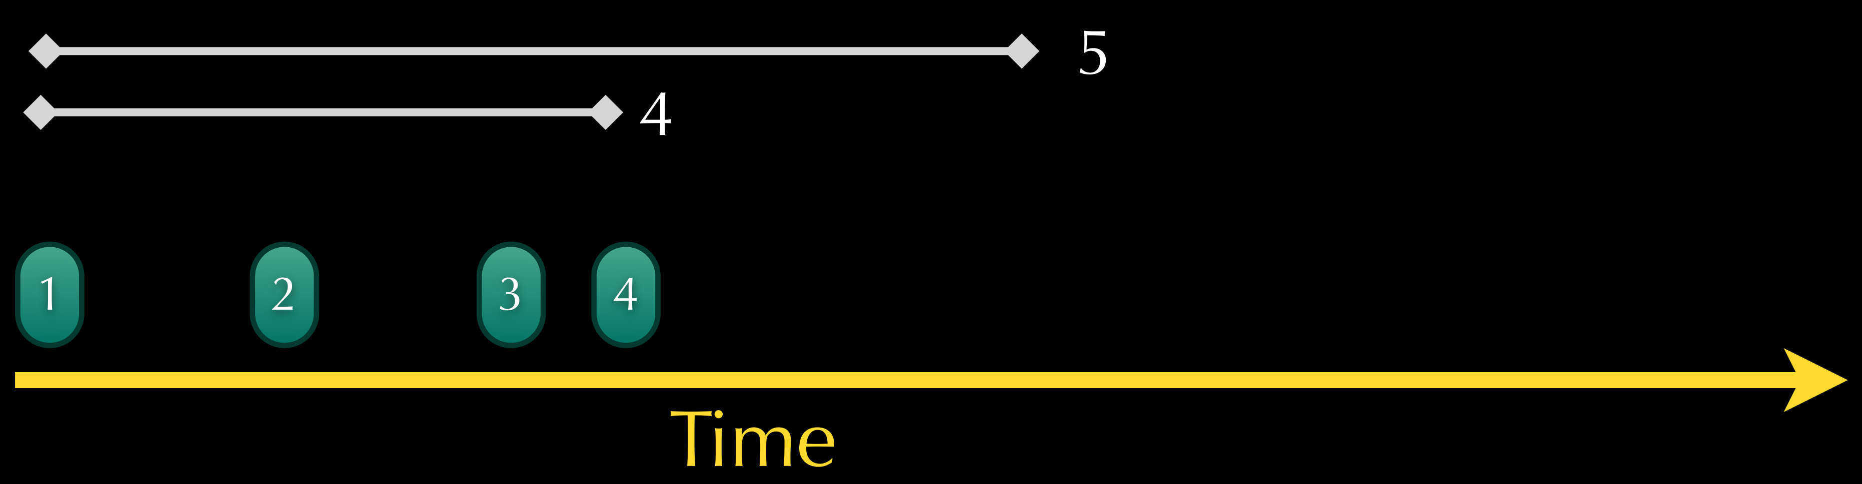

Burst-duration detector (BDD)

- Self-propagating worms must contact new destinations
  - cannot be avoided as a worm seeks new victims
- BDD determines if a host is making first-contact connections faster than usual
- Every different size of burst has a threshold: the minimum duration learned during the training phase

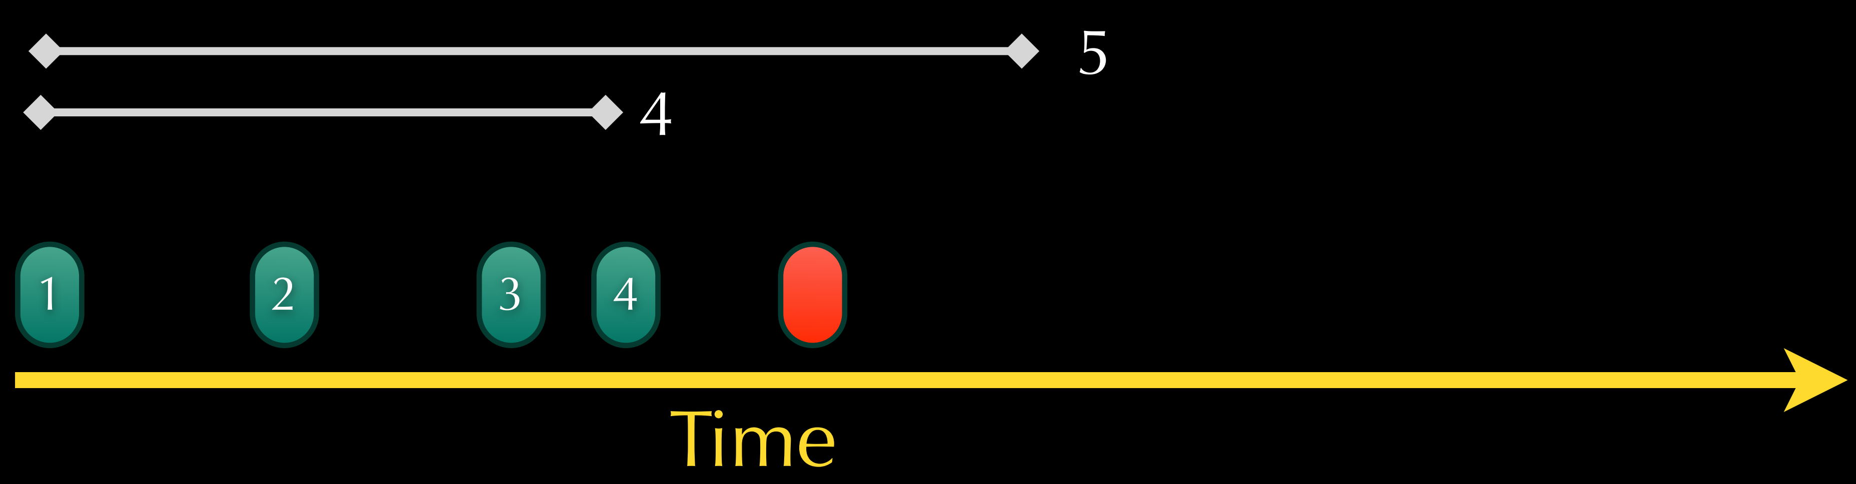

Burst-duration detector (BDD)

- Self-propagating worms must contact new destinations
  - cannot be avoided as a worm seeks new victims
- BDD determines if a host is making first-contact connections faster than usual
- Every different size of burst has a threshold: the minimum duration learned during the training phase

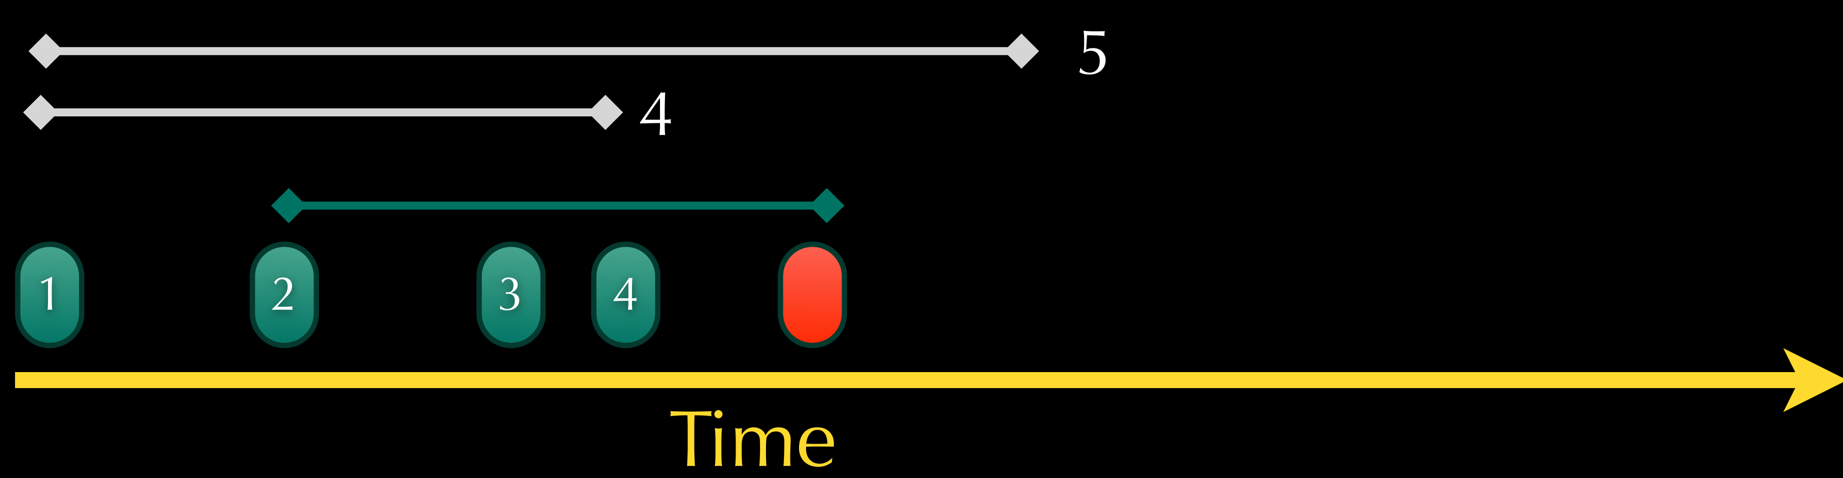

Burst-duration detector (BDD)

- Self-propagating worms must contact new destinations
  - cannot be avoided as a worm seeks new victims
- BDD determines if a host is making first-contact connections faster than usual
- Every different size of burst has a threshold: the minimum duration learned during the training phase

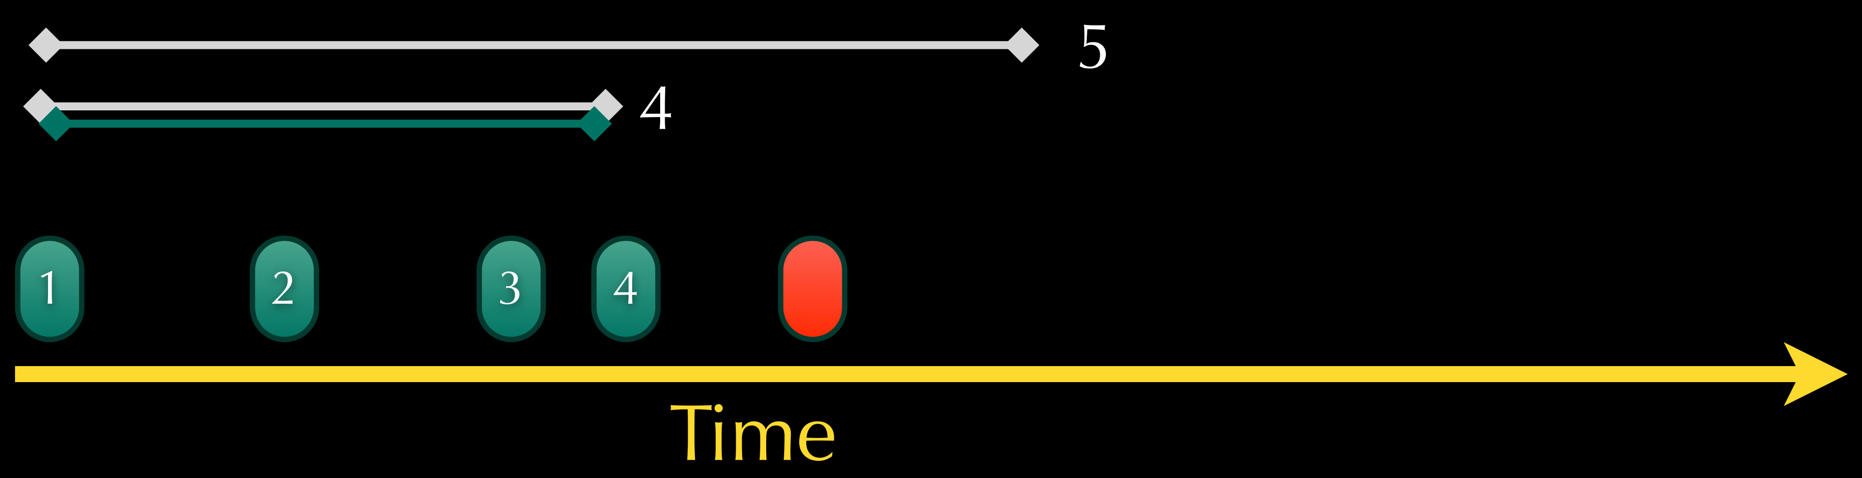

Burst-duration detector (BDD)

- Self-propagating worms must contact new destinations
  - cannot be avoided as a worm seeks new victims
- BDD determines if a host is making first-contact connections faster than usual
- Every different size of burst has a threshold: the minimum duration learned during the training phase

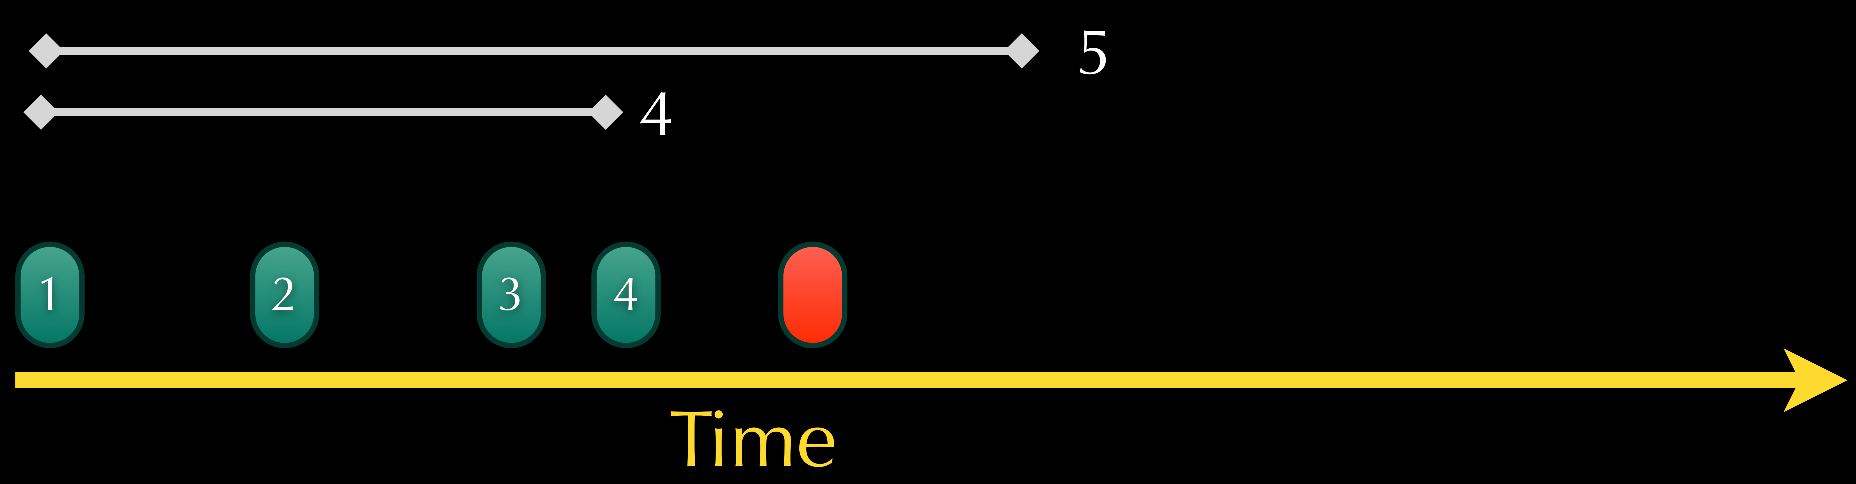

Burst-duration detector (BDD)

- Self-propagating worms must contact new destinations
  - cannot be avoided as a worm seeks new victims
- BDD determines if a host is making first-contact connections faster than usual
- Every different size of burst has a threshold: the minimum duration learned during the training phase

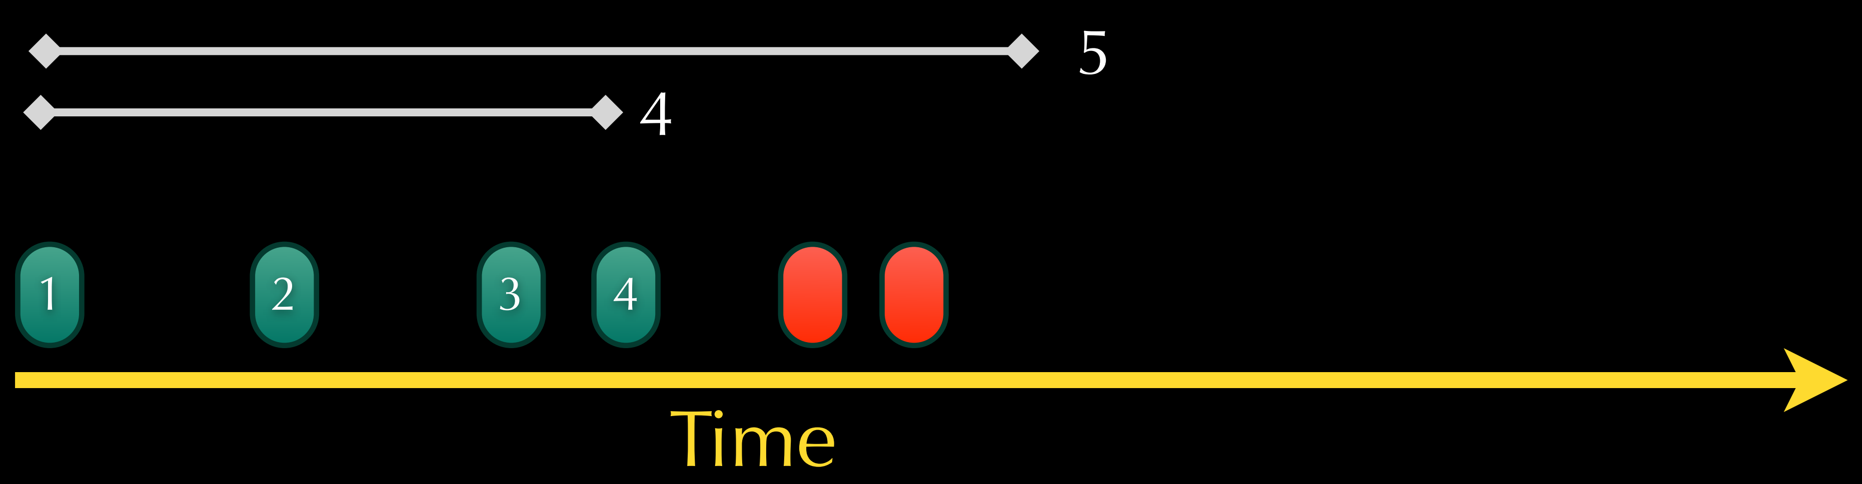

Burst-duration detector (BDD)

- Self-propagating worms must contact new destinations
  - cannot be avoided as a worm seeks new victims
- BDD determines if a host is making first-contact connections faster than usual
- Every different size of burst has a threshold: the minimum duration learned during the training phase

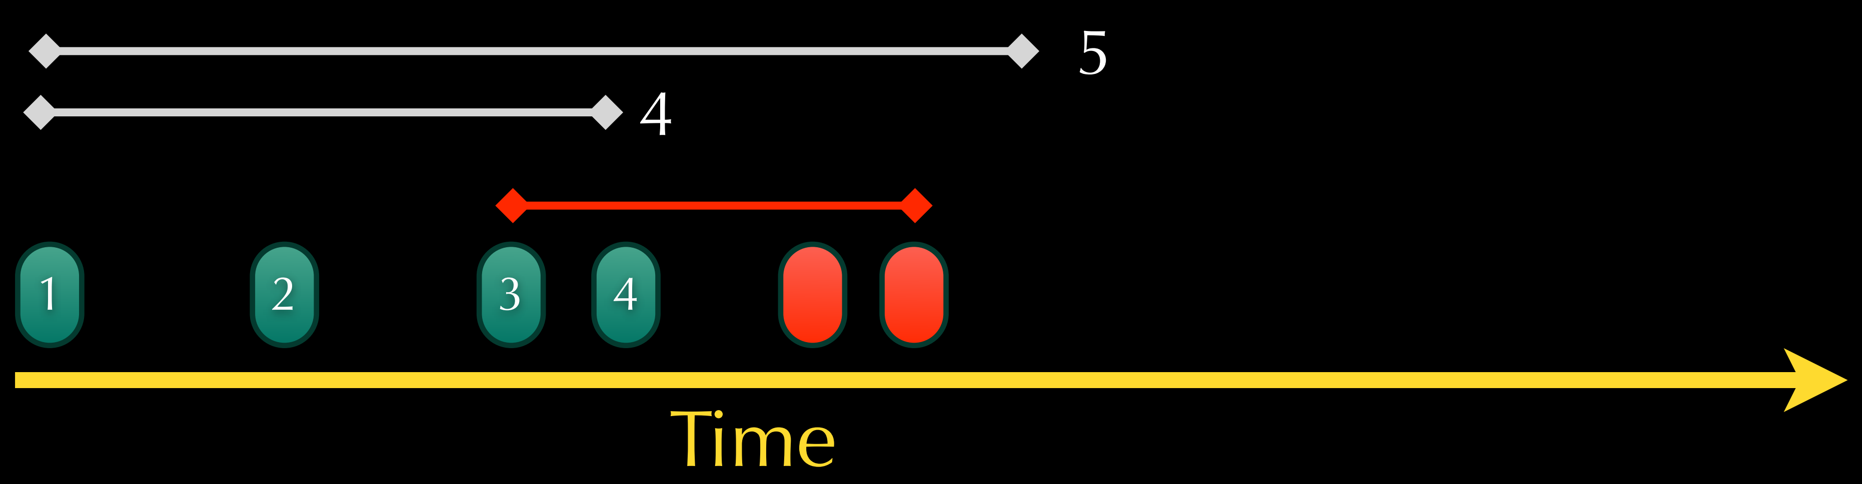

Burst-duration detector (BDD)

- Self-propagating worms must contact new destinations
  - cannot be avoided as a worm seeks new victims
- BDD determines if a host is making first-contact connections faster than usual
- Every different size of burst has a threshold: the minimum duration learned during the training phase

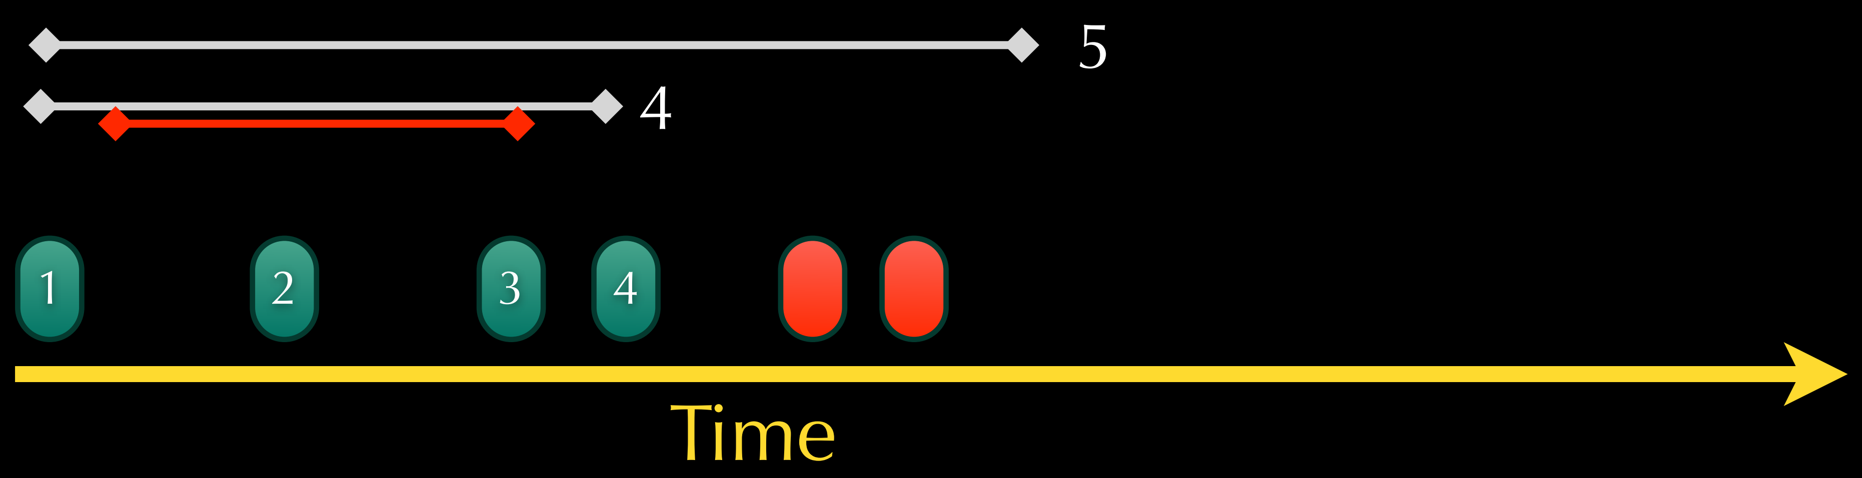

Quiescent period detector (QPD)

- Hosts are not constantly active and smart worms can use those quiet periods to propagate (which may also escape BDD)
- For each quiet period of a certain length or longer, QPD learns the threshold as the maximum duration for an active period
- QPD raises the alarm if the threshold is exceeded

Quiescent period detector (QPD)

- Hosts are not constantly active and smart worms can use those quiet periods to propagate (which may also escape BDD)
- For each quiet period of a certain length or longer, QPD learns the threshold as the maximum duration for an active period
- QPD raises the alarm if the threshold is exceeded

Quiescent period detector (QPD)

- Hosts are not constantly active and smart worms can use those quiet periods to propagate (which may also escape BDD)
- For each quiet period of a certain length or longer, QPD learns the threshold as the maximum duration for an active period
- QPD raises the alarm if the threshold is exceeded

Quiescent period detector (QPD)

- Hosts are not constantly active and smart worms can use those quiet periods to propagate (which may also escape BDD)
- For each quiet period of a certain length or longer, QPD learns the threshold as the maximum duration for an active period
- QPD raises the alarm if the threshold is exceeded

Quiescent period detector (QPD)

- Hosts are not constantly active and smart worms can use those quiet periods to propagate (which may also escape BDD)
- For each quiet period of a certain length or longer, QPD learns the threshold as the maximum duration for an active period
- QPD raises the alarm if the threshold is exceeded

Quiescent period detector (QPD)

- Hosts are not constantly active and smart worms can use those quiet periods to propagate (which may also escape BDD)
- For each quiet period of a certain length or longer, QPD learns the threshold as the maximum duration for an active period
- QPD raises the alarm if the threshold is exceeded

Quiescent period detector (QPD)

- Hosts are not constantly active and smart worms can use those quiet periods to propagate (which may also escape BDD)
- For each quiet period of a certain length or longer, QPD learns the threshold as the maximum duration for an active period
- QPD raises the alarm if the threshold is exceeded

Quiescent period detector (QPD)

- Hosts are not constantly active and smart worms can use those quiet periods to propagate (which may also escape BDD)
- For each quiet period of a certain length or longer, QPD learns the threshold as the maximum duration for an active period
- QPD raises the alarm if the threshold is exceeded

Quiescent period detector (QPD)

- Hosts are not constantly active and smart worms can use those quiet periods to propagate (which may also escape BDD)
- For each quiet period of a certain length or longer, QPD learns the threshold as the maximum duration for an active period
- QPD raises the alarm if the threshold is exceeded

Quiescent period detector (QPD)

- Hosts are not constantly active and smart worms can use those quiet periods to propagate (which may also escape BDD)
- For each quiet period of a certain length or longer, QPD learns the threshold as the maximum duration for an active period
- QPD raises the alarm if the threshold is exceeded

Quiescent period detector (QPD)

- Hosts are not constantly active and smart worms can use those quiet periods to propagate (which may also escape BDD)
- For each quiet period of a certain length or longer, QPD learns the threshold as the maximum duration for an active period
- QPD raises the alarm if the threshold is exceeded

Quiescent period detector (QPD)

- Hosts are not constantly active and smart worms can use those quiet periods to propagate (which may also escape BDD)
- For each quiet period of a certain length or longer, QPD learns the threshold as the maximum duration for an active period
- QPD raises the alarm if the threshold is exceeded

Quiescent period detector (QPD)

- Hosts are not constantly active and smart worms can use those quiet periods to propagate (which may also escape BDD)
- For each quiet period of a certain length or longer, QPD learns the threshold as the maximum duration for an active period
- QPD raises the alarm if the threshold is exceeded

Quiescent period detector (QPD)

- Hosts are not constantly active and smart worms can use those quiet periods to propagate (which may also escape BDD)
- For each quiet period of a certain length or longer, QPD learns the threshold as the maximum duration for an active period
- QPD raises the alarm if the threshold is exceeded

Quiescent period detector (QPD)

- Hosts are not constantly active and smart worms can use those quiet periods to propagate (which may also escape BDD)
- For each quiet period of a certain length or longer, QPD learns the threshold as the maximum duration for an active period
- QPD raises the alarm if the threshold is exceeded

Quiescent period detector (QPD)

- Hosts are not constantly active and smart worms can use those quiet periods to propagate (which may also escape BDD)
- For each quiet period of a certain length or longer, QPD learns the threshold as the maximum duration for an active period
- QPD raises the alarm if the threshold is exceeded

Quiescent period detector (QPD)

- Hosts are not constantly active and smart worms can use those quiet periods to propagate (which may also escape BDD)
- For each quiet period of a certain length or longer, QPD learns the threshold as the maximum duration for an active period
- QPD raises the alarm if the threshold is exceeded

Quiescent period detector (QPD)

- Hosts are not constantly active and smart worms can use those quiet periods to propagate (which may also escape BDD)
- For each quiet period of a certain length or longer, QPD learns the threshold as the maximum duration for an active period
- QPD raises the alarm if the threshold is exceeded

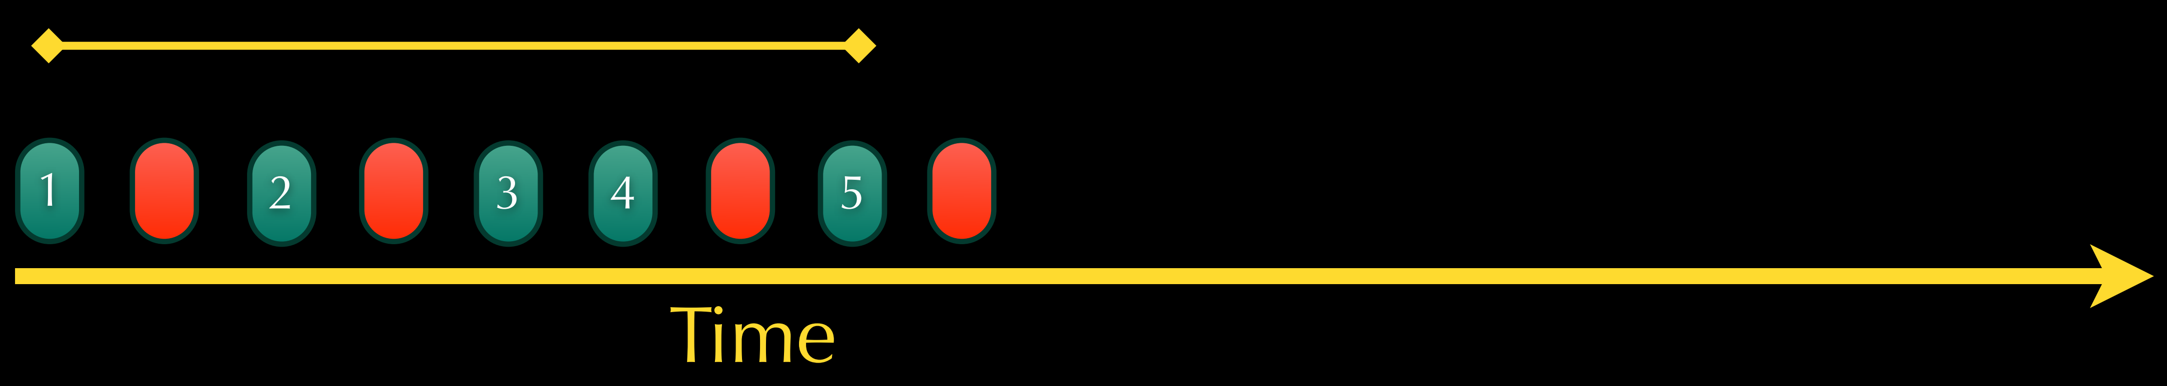

Quiescent period detector (QPD)

- Hosts are not constantly active and smart worms can use those quiet periods to propagate (which may also escape BDD)
- For each quiet period of a certain length or longer, QPD learns the threshold as the maximum duration for an active period
- QPD raises the alarm if the threshold is exceeded

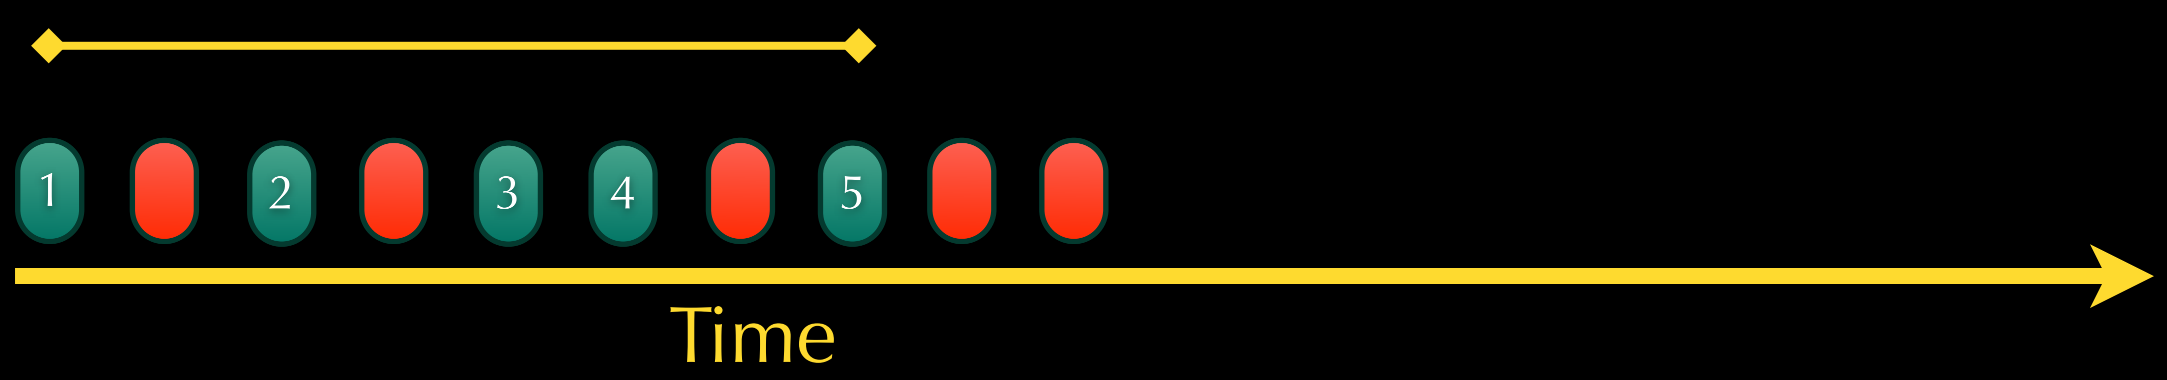

Quiescent period detector (QPD)

- Hosts are not constantly active and smart worms can use those quiet periods to propagate (which may also escape BDD)
- For each quiet period of a certain length or longer, QPD learns the threshold as the maximum duration for an active period
- QPD raises the alarm if the threshold is exceeded

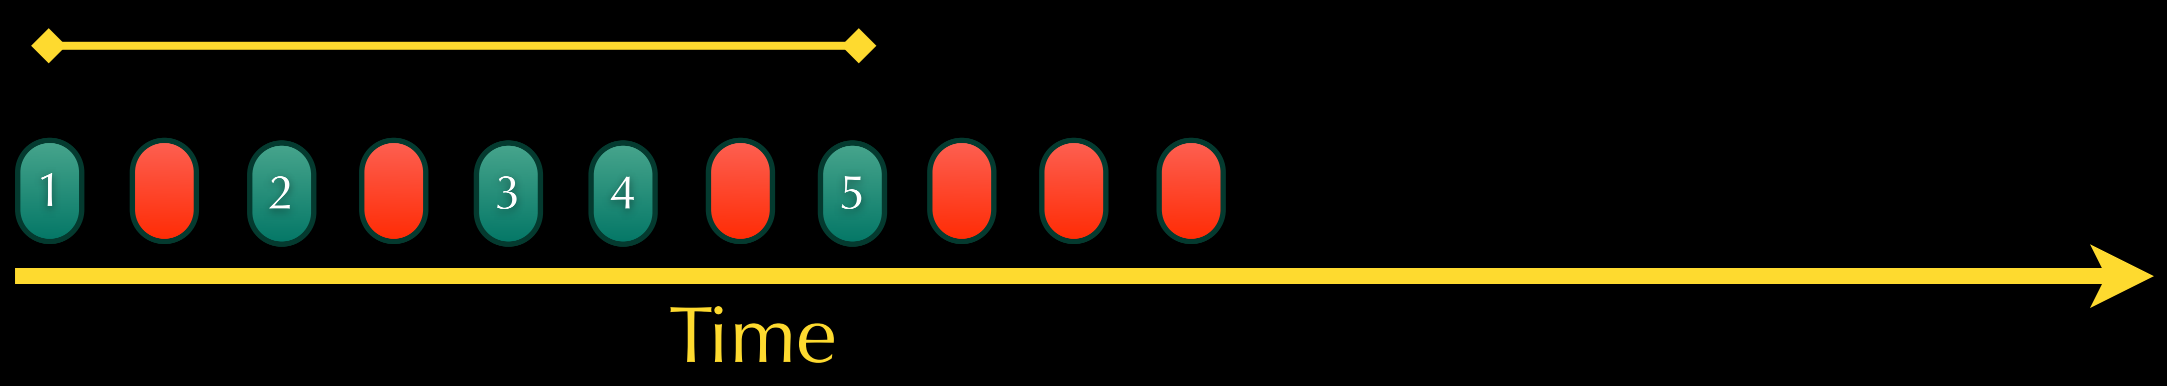

Quiescent period detector (QPD)

- Hosts are not constantly active and smart worms can use those quiet periods to propagate (which may also escape BDD)
- For each quiet period of a certain length or longer, QPD learns the threshold as the maximum duration for an active period
- QPD raises the alarm if the threshold is exceeded

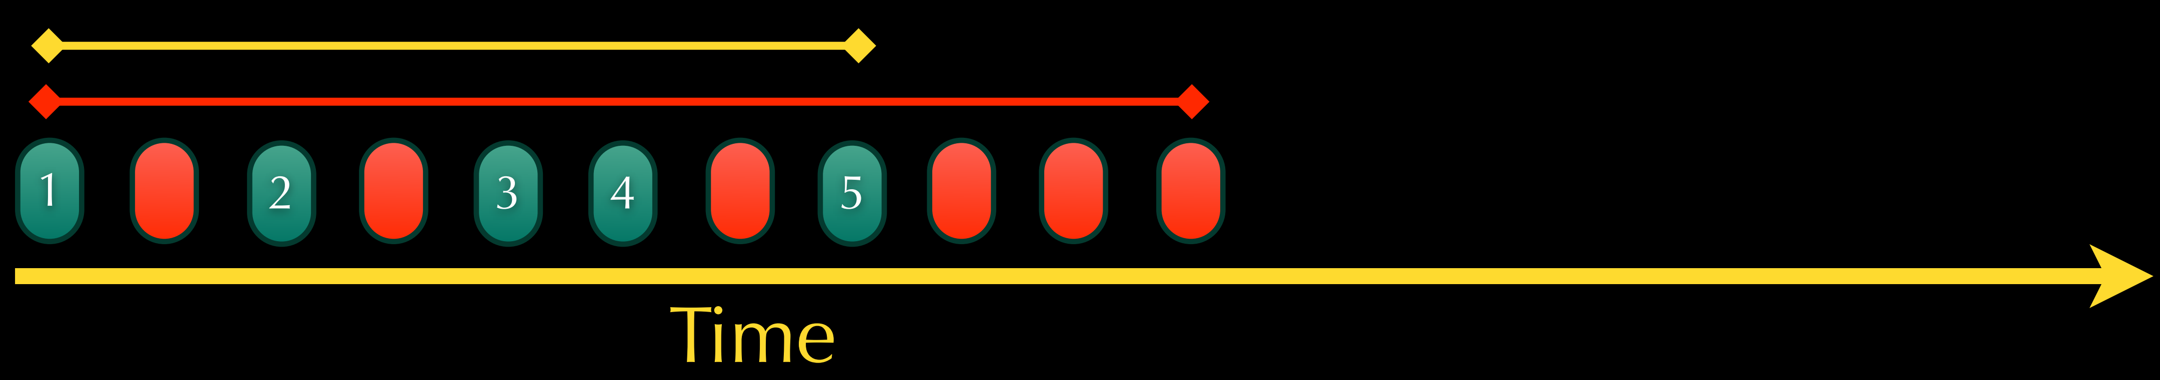

Quiescent period detector (QPD)

- Hosts are not constantly active and smart worms can use those quiet periods to propagate (which may also escape BDD)
- For each quiet period of a certain length or longer, QPD learns the threshold as the maximum duration for an active period
- QPD raises the alarm if the threshold is exceeded

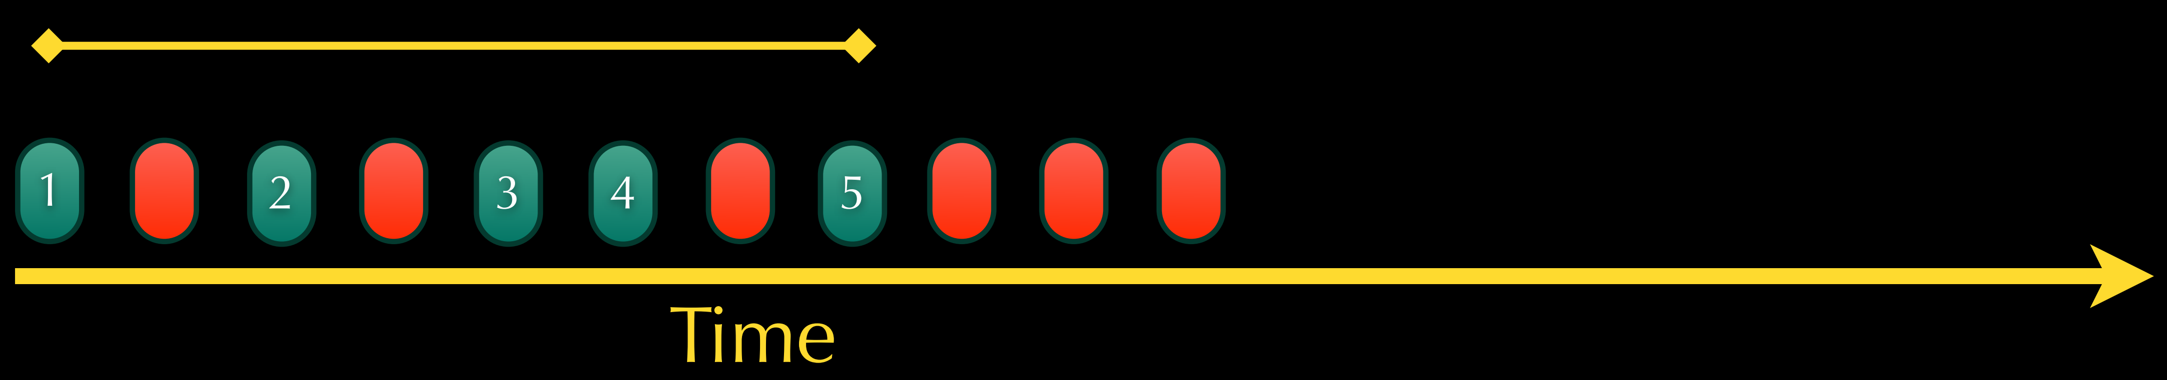

Combine BDD and QPD

- SWORD declares a host to be infected with a worm when either BDD or QPD raises an alarm
- If a worm wishes to escape BDD, it cannot shorten the duration of a burst of any size, and will have to lengthen active periods, thus caught by QPD
- If a worm wishes to escape QPD, it has to ensure quiescent periods, and will have to insert its connections to active periods, which makes certain bursts to be shorter than permitted, thus caught by BDD

Outline

- Introduction
- SWORD — Self-propagating Worm Observation and Rapid Detection
- Experiment Methodology
- SWORD Performance Against Classic Worms
- SWORD Performance Against Smart Worms
- Concluding remarks

Evaluation framework

- Support multiple standardized network environments
- Allow pluggable worm implementations that can support advanced worm types
- Make it easy to implement a variety of detectors, and include popular detectors as benchmarks

Metrics

- False positive rate
  - By host: the number of false alarms raised during a time period (limited to one alarm per host)
  - By time: Percentage of minutes during time period τ when a false alarm is triggered
- False negative rate
  - Percent of instances where a worm infection occurs but is not detected in time period
- Detection latency
  - The number of outbound worm connections prior to detection

Outline

- Introduction
- SWORD — Self-propagating Worm Observation and Rapid Detection
- Experiment Methodology
- SWORD Performance Against Classic Worms
- SWORD Performance Against Smart Worms
- Concluding remarks

SWORD vs other detectors

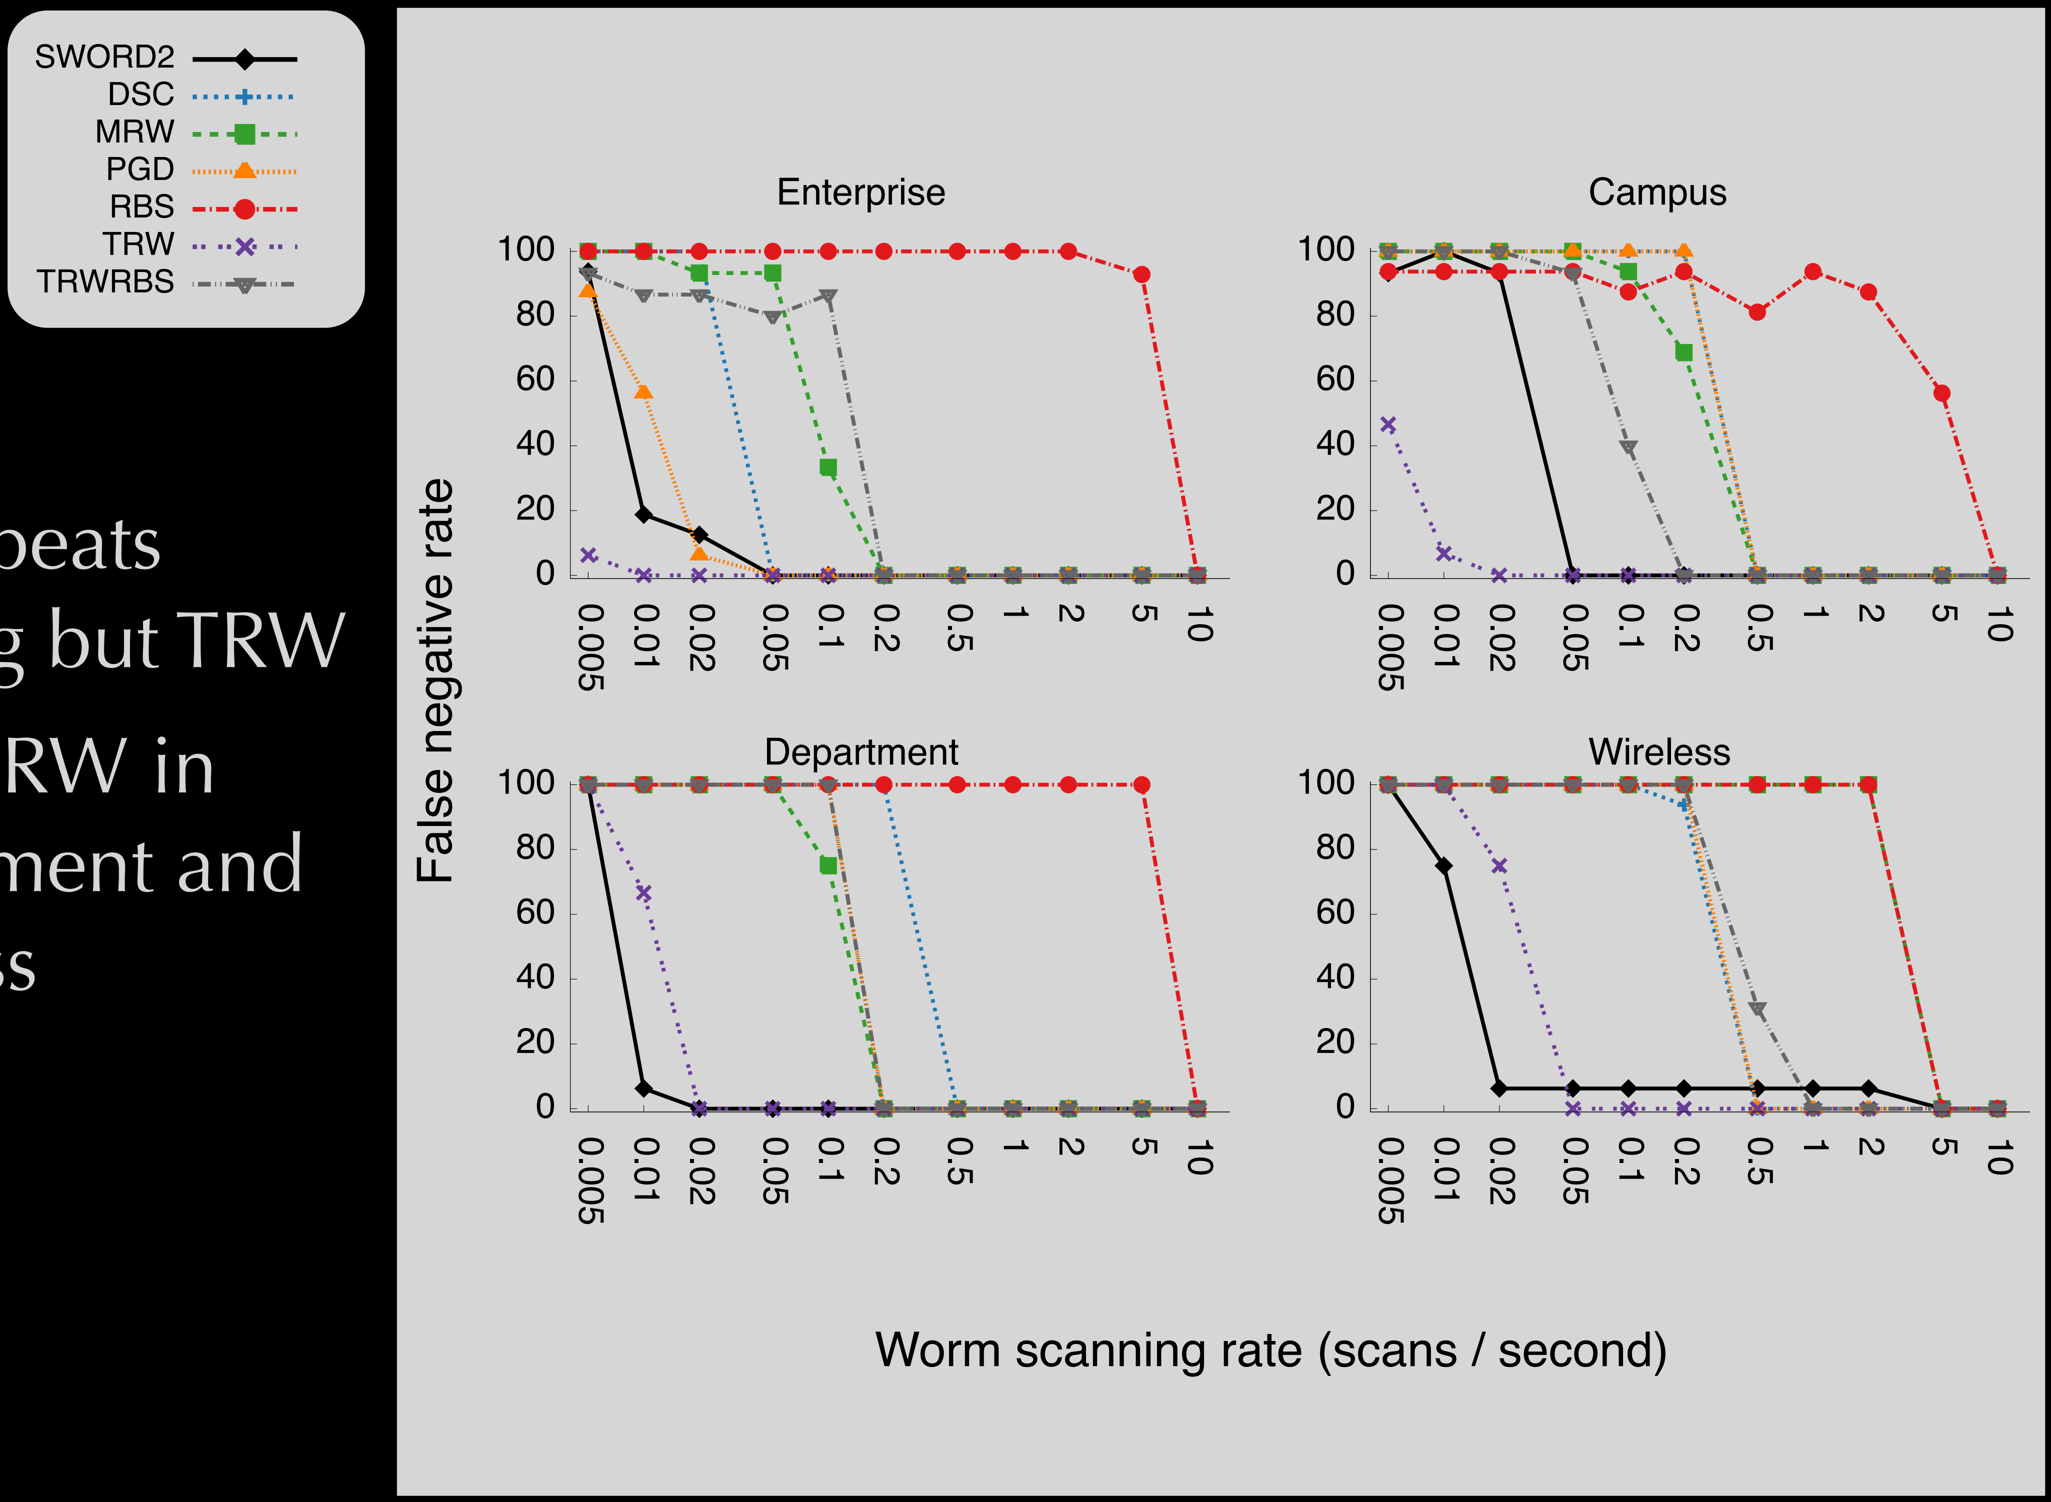

- SWORD beats everything but TRW
  - beats TRW in department and wireless

Latency

Detector	Campus	Enterprise	Department	Wireless
SWORD	21.73	24.97	22.99	264.94
DSC	2	22	19	15.93
MRW	28.88	51.7	43.64	1014.16
PGD	93.8	28.11	25.81	621.75
RBS	17.36	4.25	26.44	349.53
TRW	4.23	11.13	24.75	49.93
RWRBS	57.97	30.39	58.66	167.95

Outline

- Introduction
- SWORD — Self-propagating Worm Observation and Rapid Detection
- Experiment Methodology
- SWORD Performance Against Classic Worms
- SWORD Performance Against Smart Worms
- Concluding remarks

Goals

- Determine how fast a worm can scan while evading detection most of the time
- The lower this rate, the more time we have to intervene before the worm causes extensive damage

Capabilities

- What capabilities should an evasive worm have?

Knowledge of the network	Knowledge of the detector
blind: can't see any traffic	*speculative*: guesses at the detector configuration, but knows nothing about it
perceptive: *can see all* traffic on the network	*informed*: has detailed knowledge of the detector configuration

Methodology

- Evasive worms evaluated using the same framework and methods as the previous experiments
- Introduce a parameter: load factor

Load factor

- Imagine a rate-based detector that would raise an alarm when >20 connections per second are made
- Evasive worm models this detector internally
  - load factor of 1.0, would cause worm to issue 20 connections per second.
  - load factor of 0.5 would cause worm to issue 10 connections per second

Metrics

- Evasion Rate
  - for a given scenario, in what % of experiments a worm was able to evade detection
- Effective Scan Rate
  - how many worm scans a worm was able to make while avoiding detection
- **Maximum Effective Scan Rate**
  - the maximum effective scan rate achieved with an evasion rate of 90% or greater

Effective scanning rate

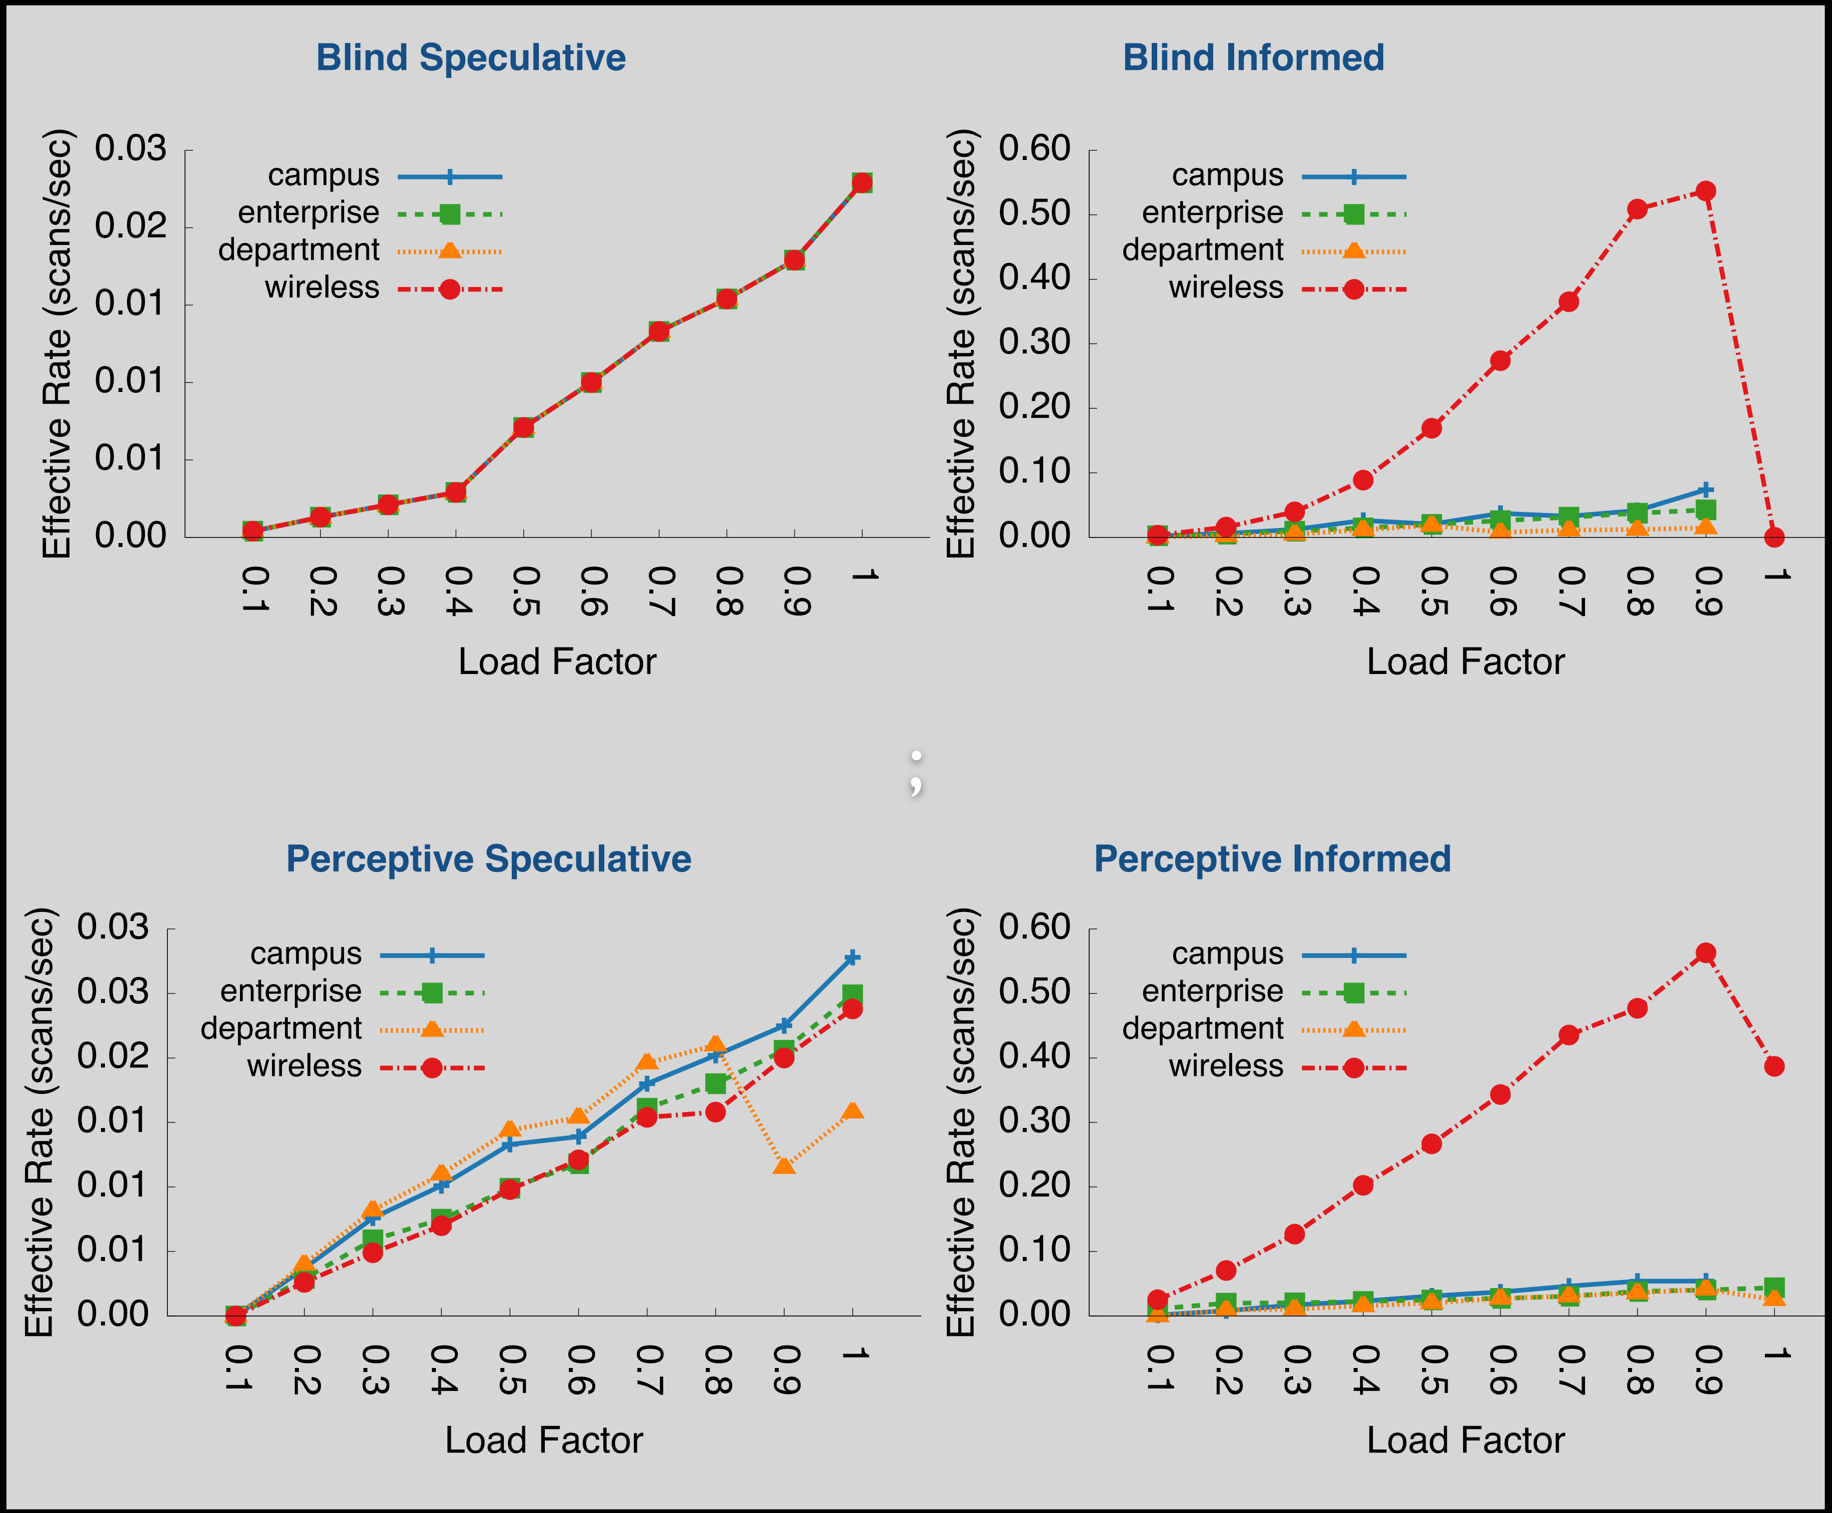

- Effective rate increases linearly with load factor
- Knowledge of environment can make a substantial difference

Evasion rate

- Evasion rate is always excellent at low load factors
- Results not meaningful until combined with effective rate

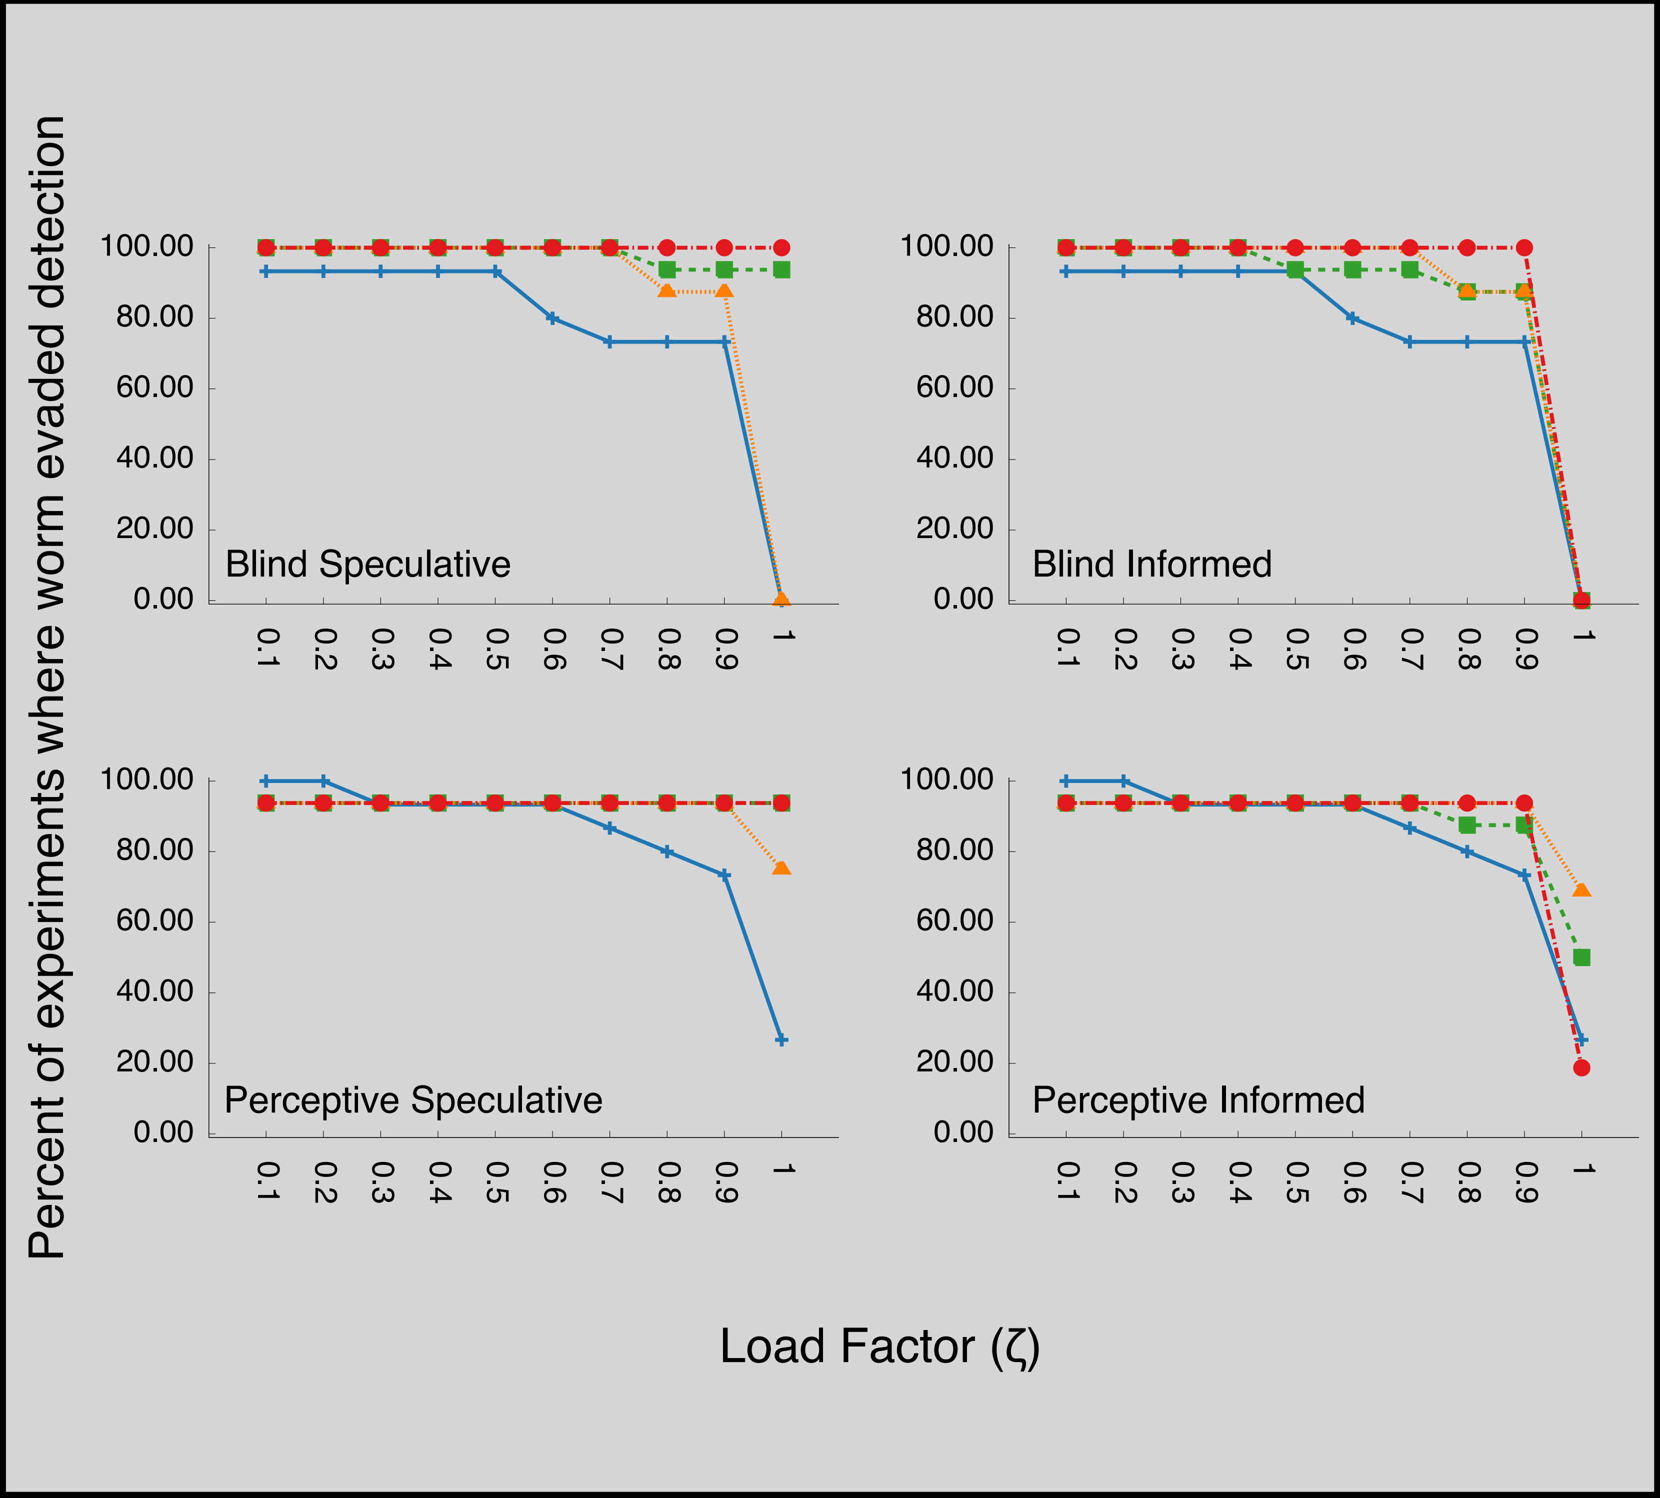

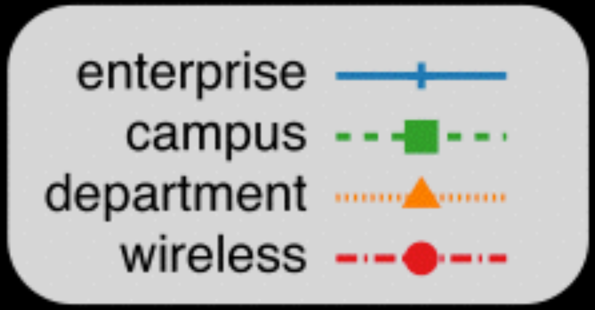

Max effective rate (campus)

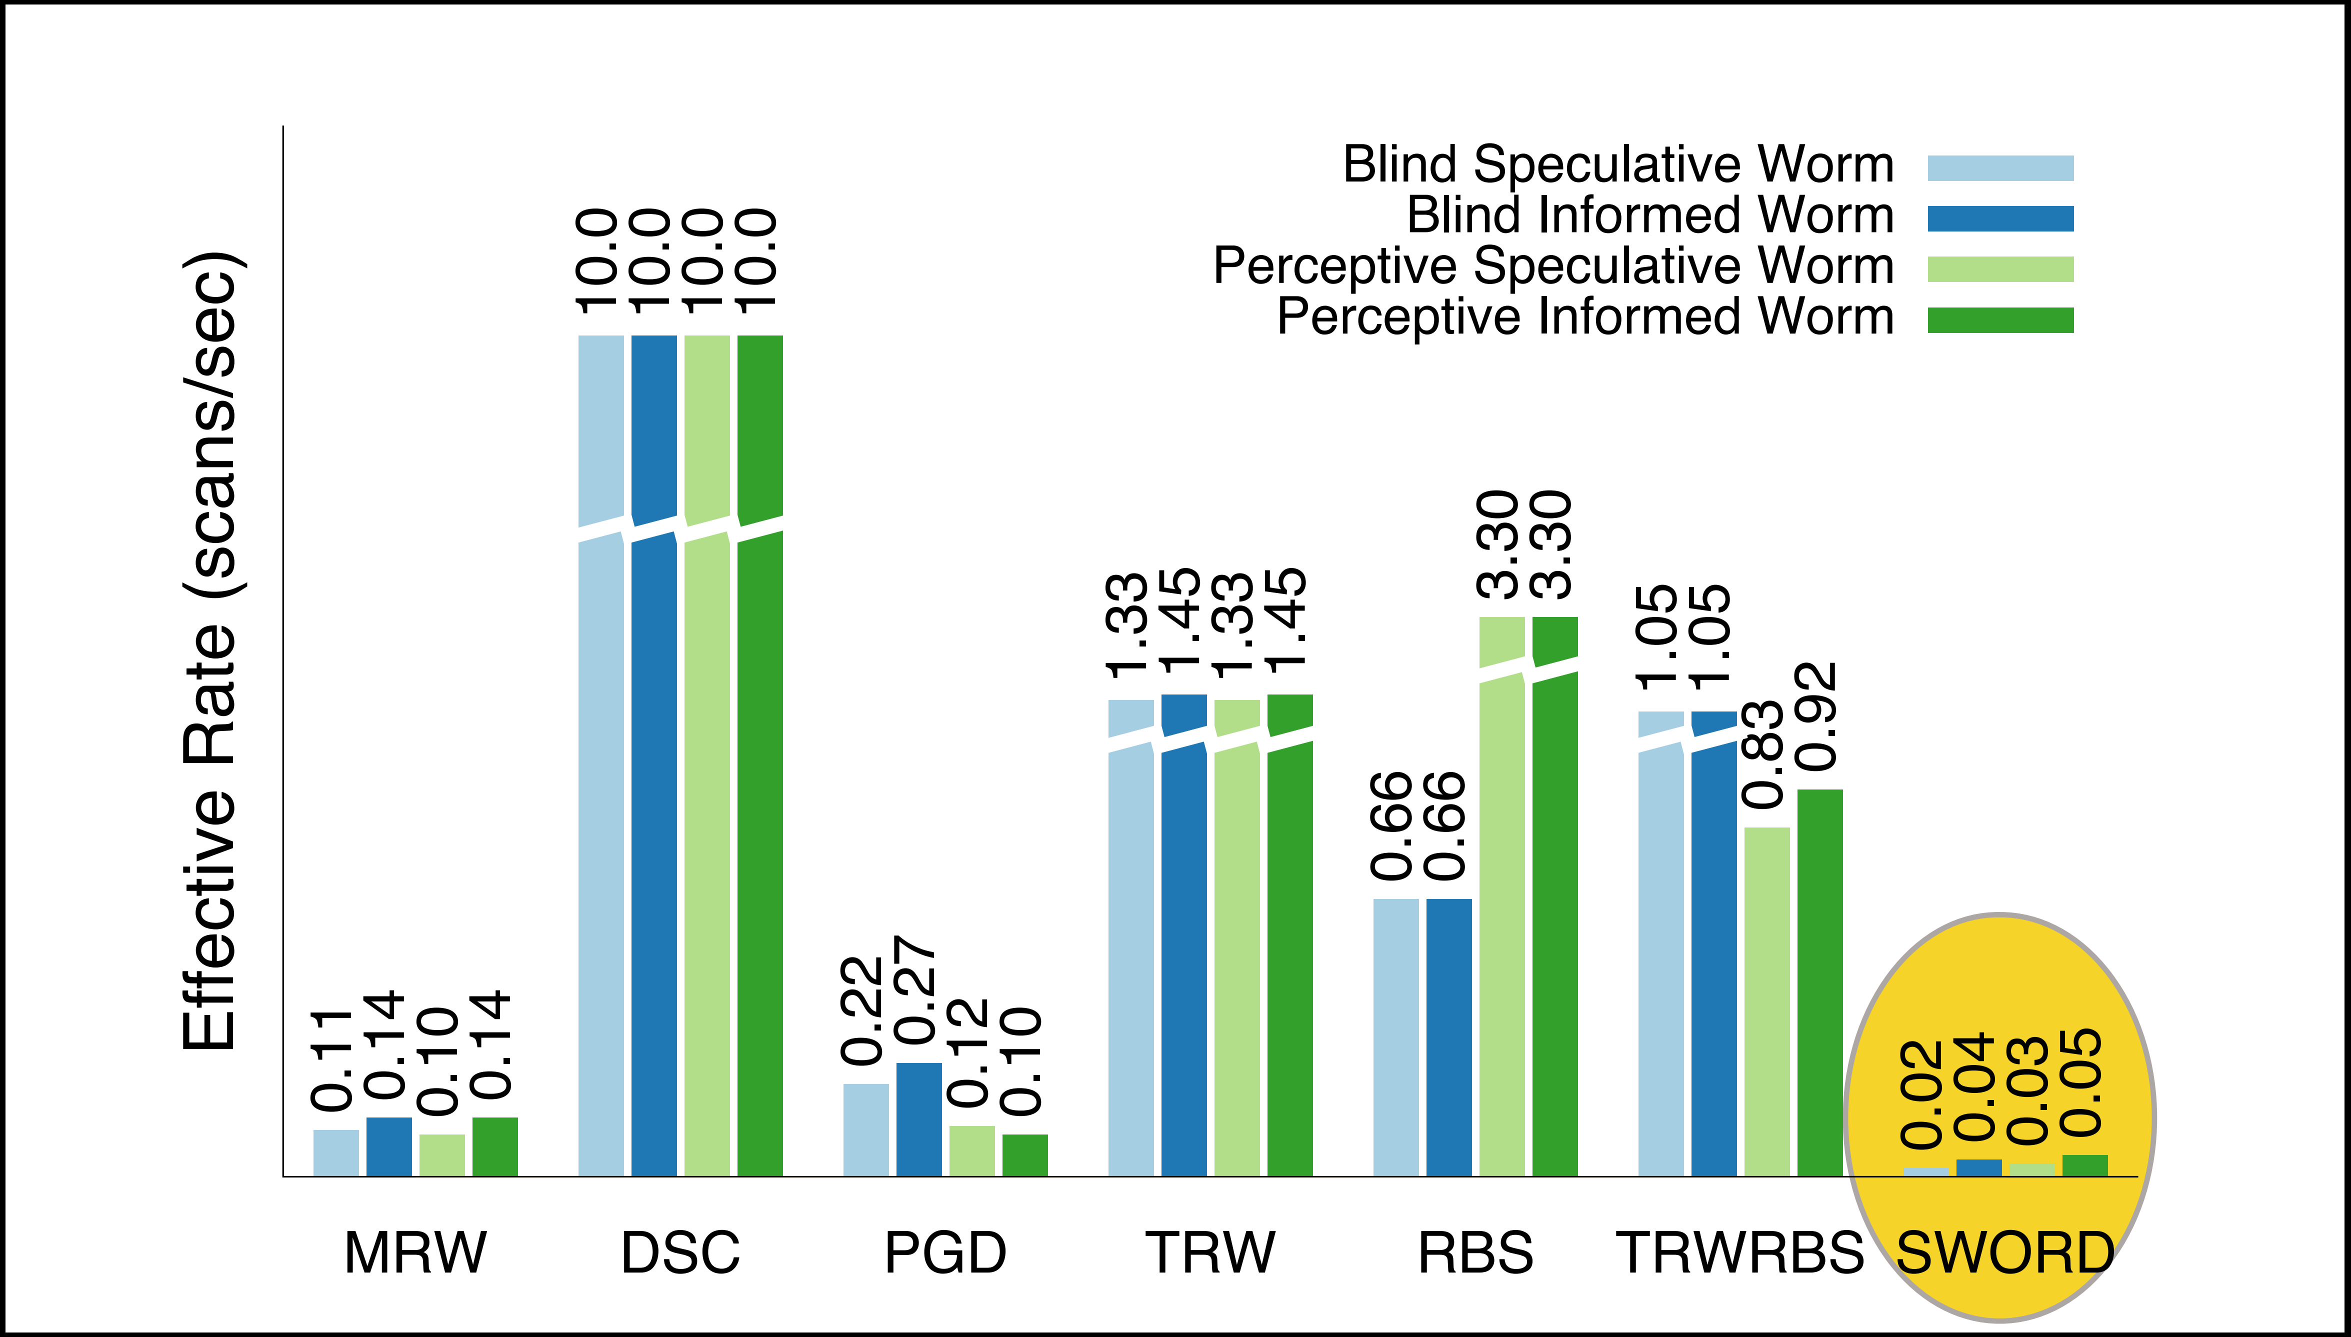

Max effective rate (enterprise)

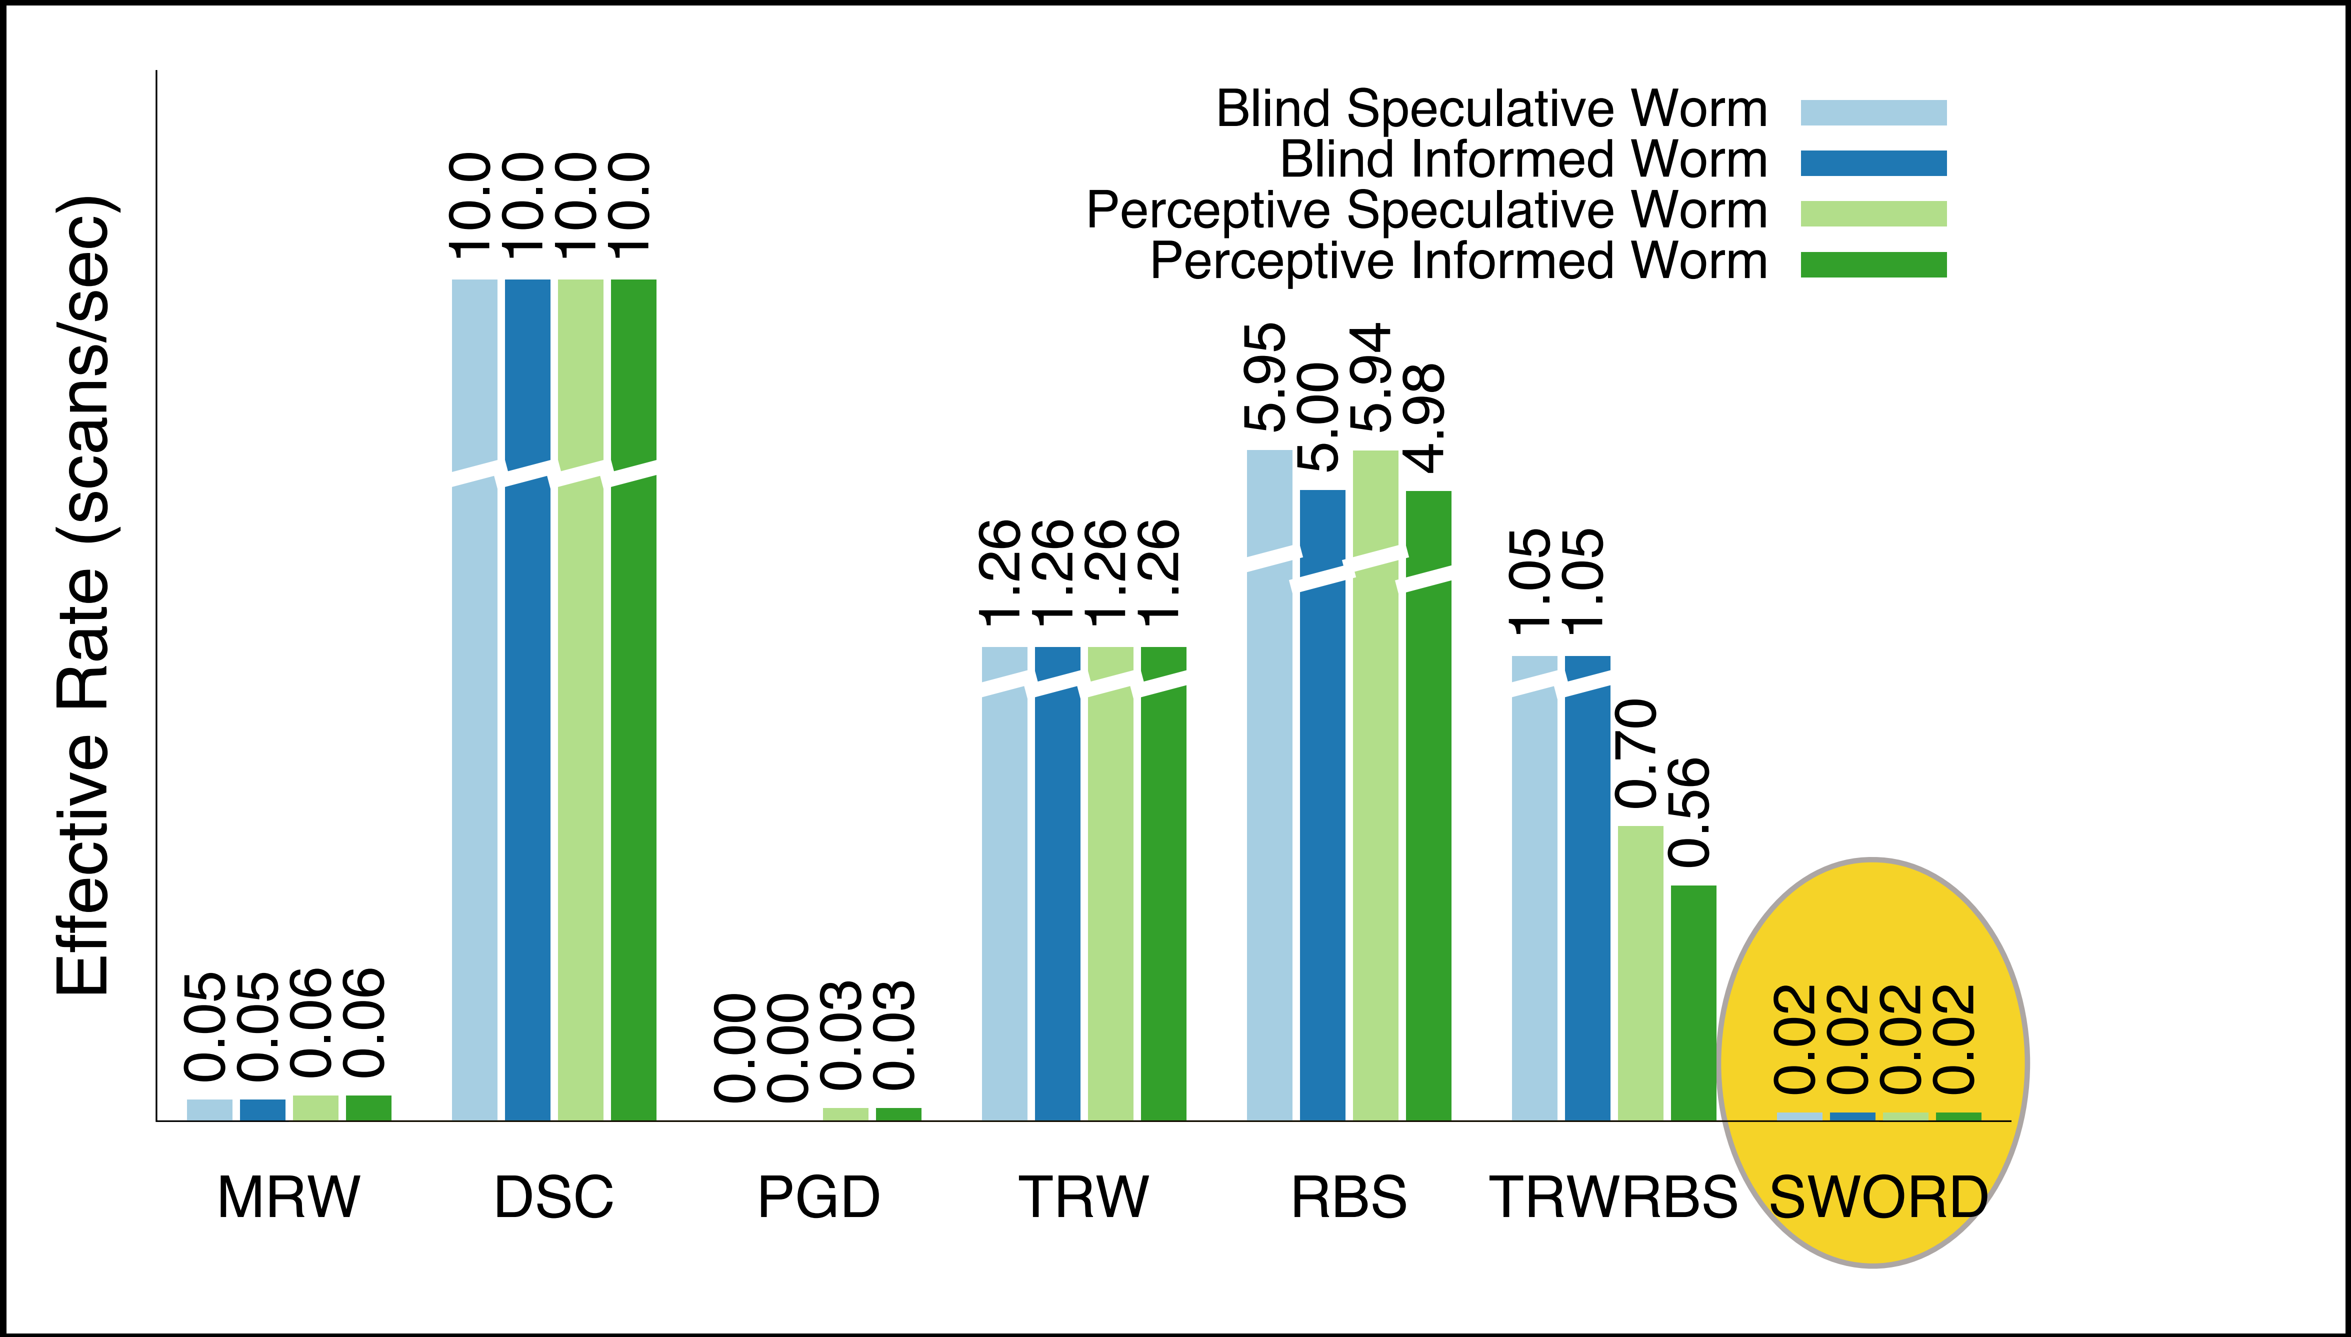

Max effective rate (department)

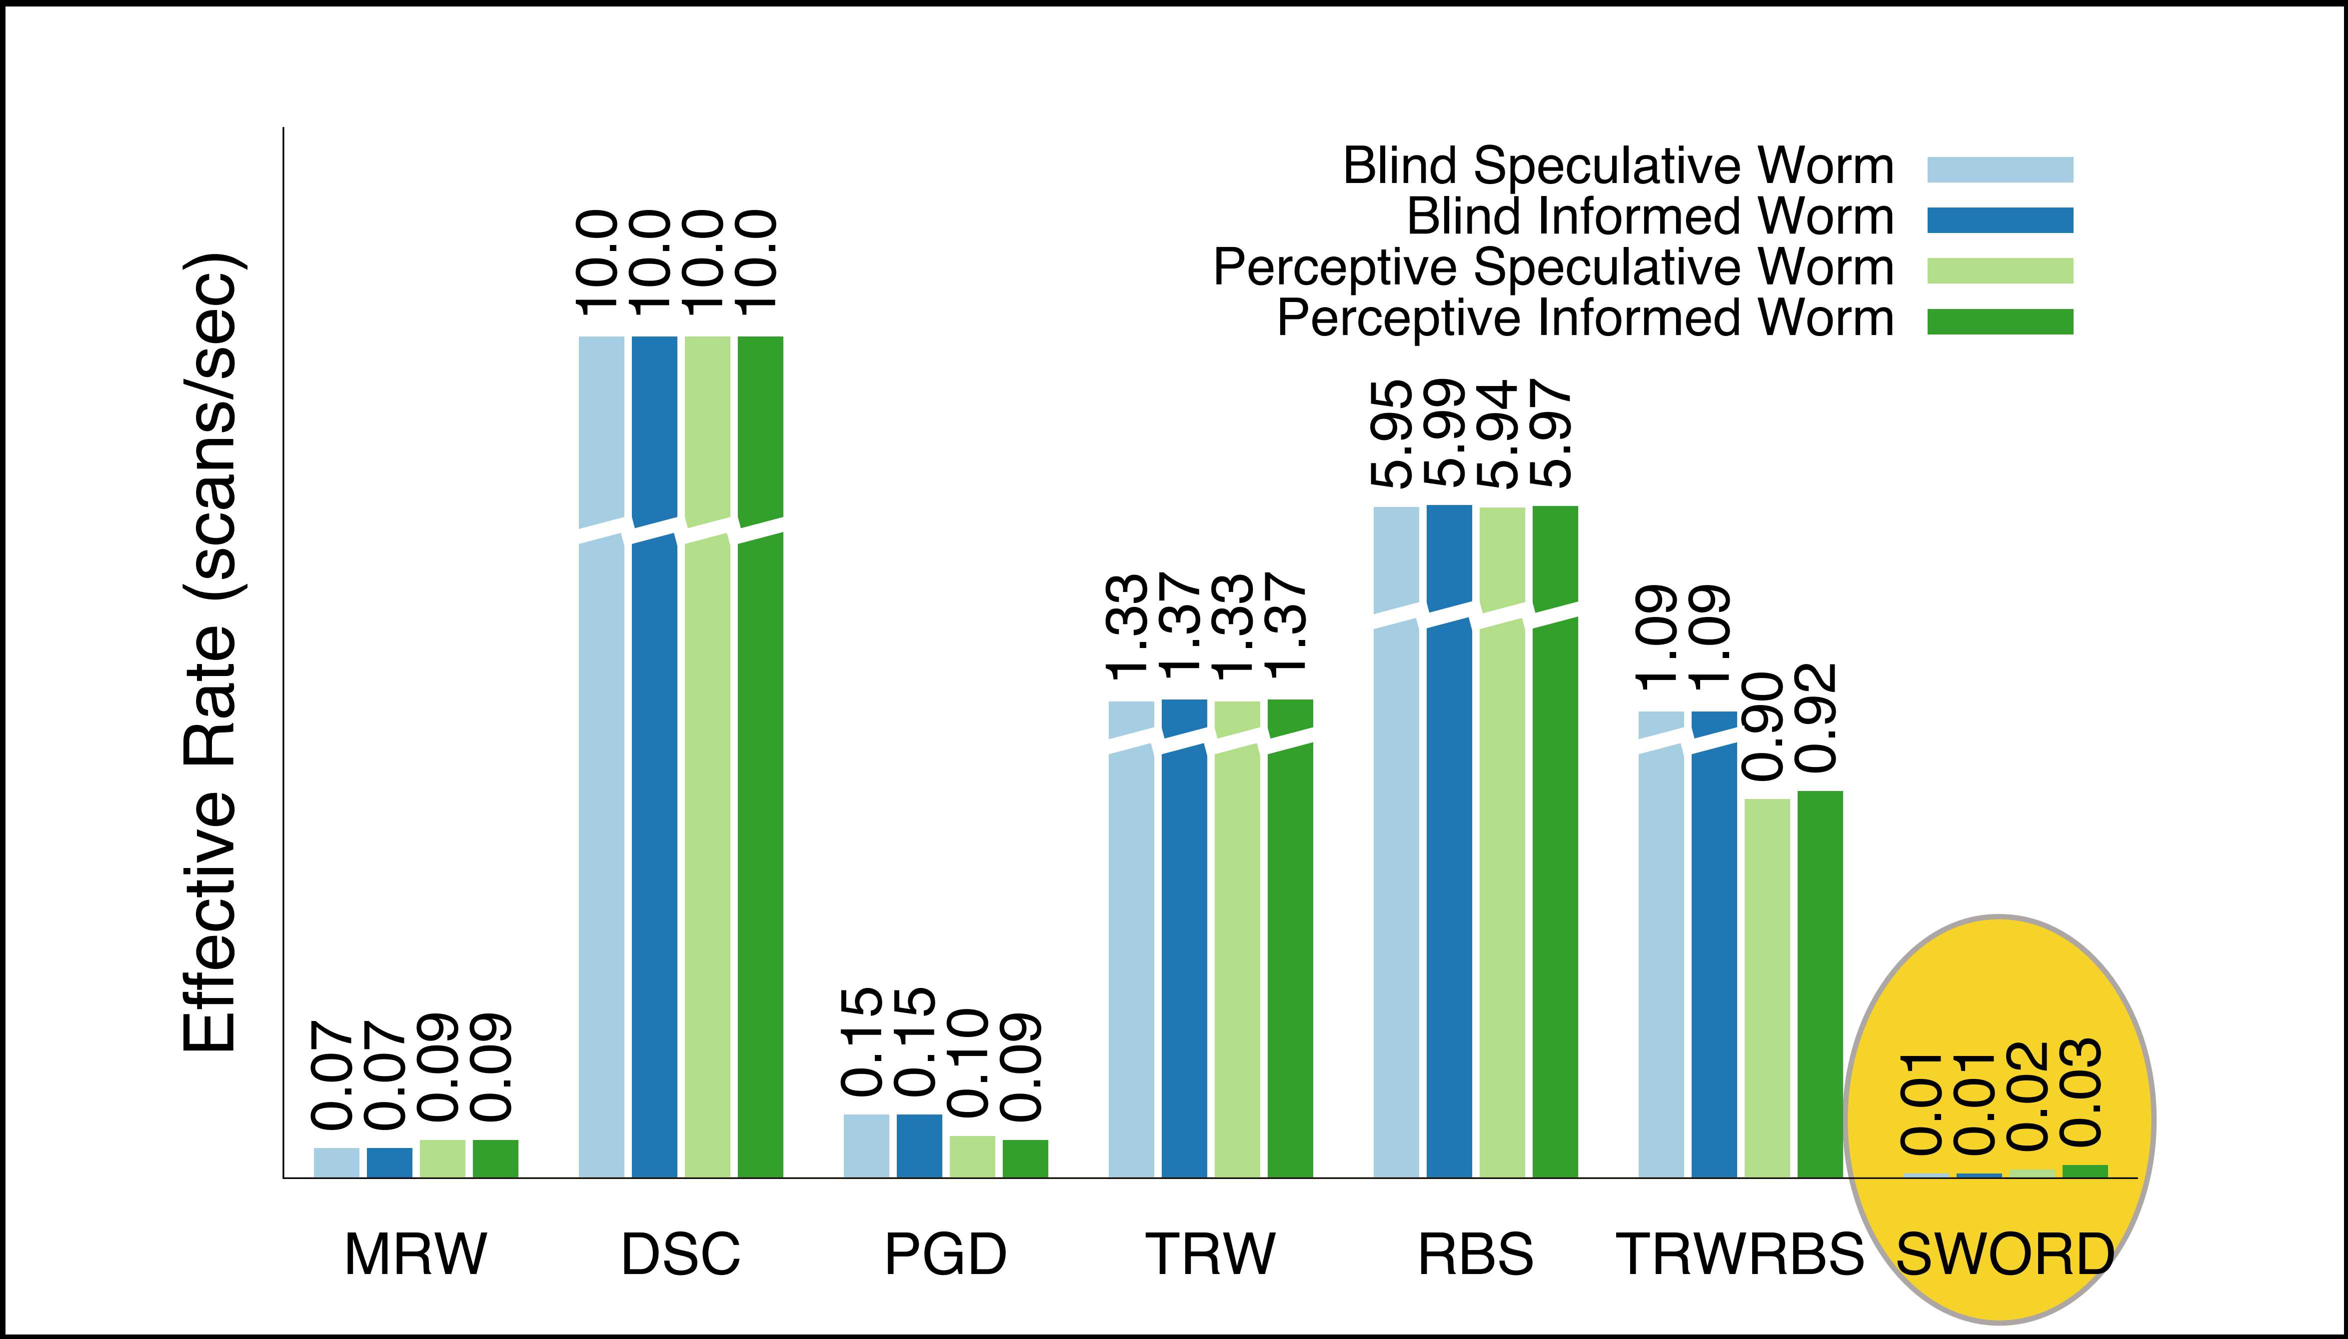

Max effective rate (wireless)

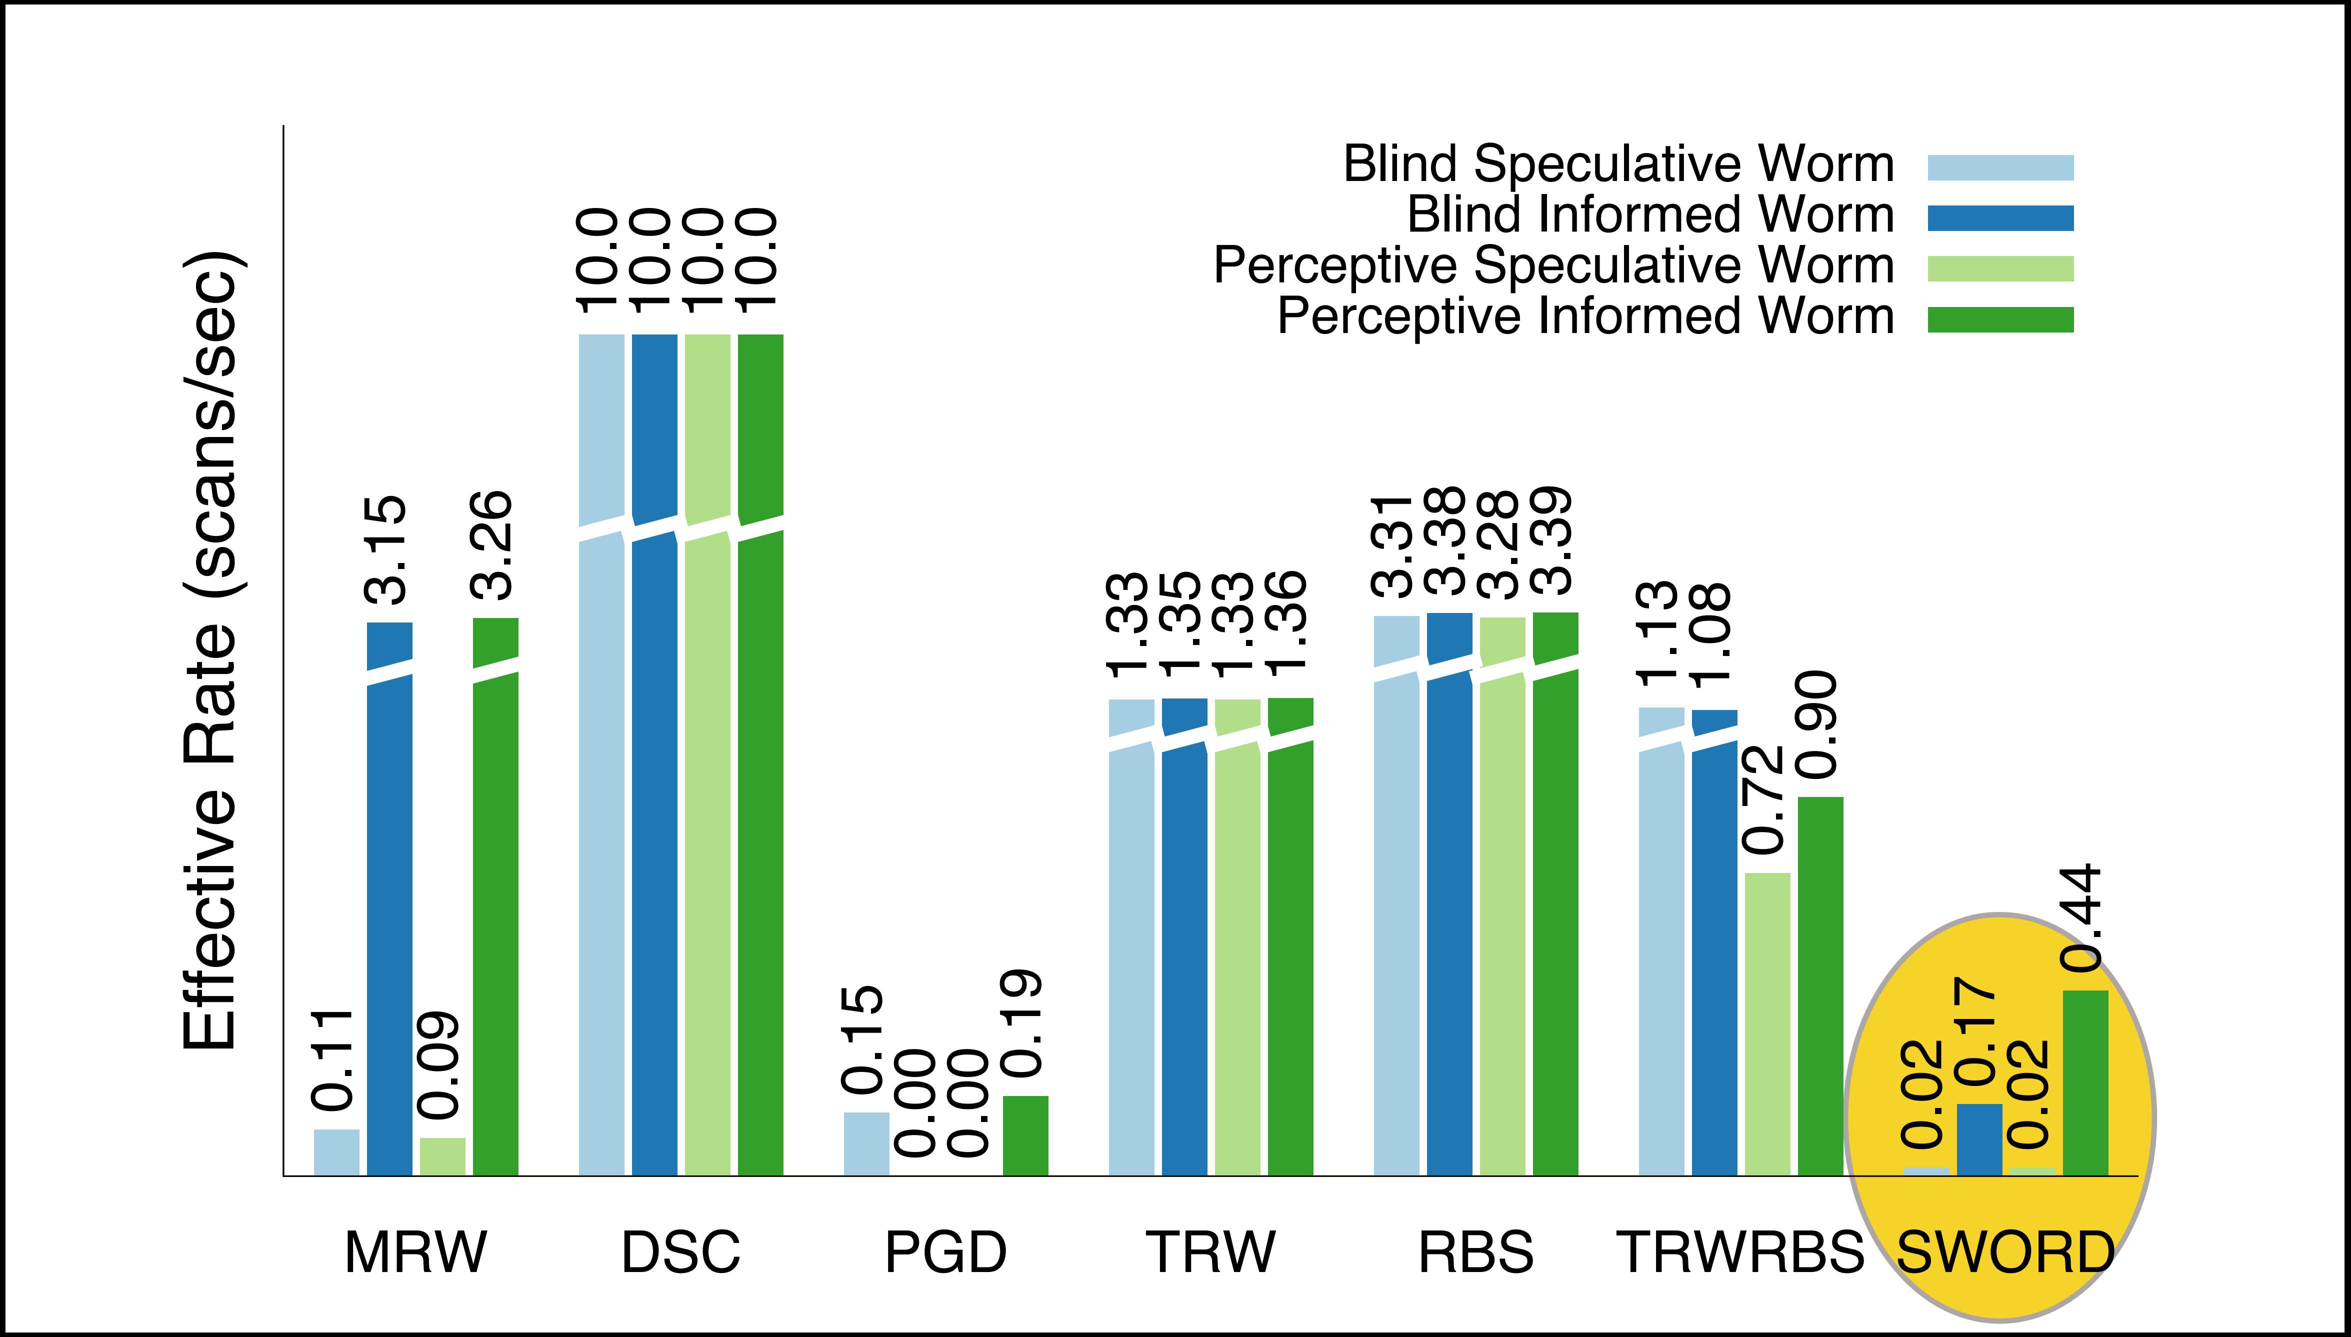

Findings

- Q) Can worms get smarter and evade detection?
- A) Against some worm detectors, YES!
  - DSC performed well against naive worms, but miserably against evasive worms
  - when TRW has known hosts available, it can scan rapidly while evading detection
  - detectors based on counting first-contact connections not easily evaded

Findings (2)

- SWORD outperforms existing detectors
  - against both classic and evasive worms
- PGD is the only detector that even comes close
  - SWORD outperforms PGD against classic worms with better sensitivity and latency in **all** environments
  - SWORD outperforms PGD in 11 of the 16 evasive worm scenarios

Outline

- Introduction
- SWORD — Self-propagating Worm Observation and Rapid Detection
- Experiment Methodology
- SWORD Performance Against Classic Worms
- SWORD Performance Against Smart Worms
- Concluding remarks

Conclusions

- Internet worms continue to be a significant risk and research is required to defend against them
- SWORD can detect worms effectively by focusing on the fundamental worm behaviors
- SWORD outperforms other behavior-based detectors against both classic and smart worms

This research is supported by the National Science Foundation (NSF) Award No. CNS-0644434. Any opinions, findings, and conclusions or recommendations expressed in this research are those of the author(s) and do not necessarily reflect the views of the NSF.

?

Jun Li, Ph.D.
Associate Professor, Computer and Information Science
Network & Security Research Laboratory
University of Oregon
Email: lijun@cs.uoregon.edu
http://www.cs.uoregon.edu/~lijun