THE UNIVERSITY of EDINBURGH
informatics

Bankrupt Covert Channel:

Turning Network Predictability into Vulnerability

Dmitrii Ustiugov, Plamen Petrov, Siavash Katebzadeh, Boris Grot
University of Edinburgh

This work is supported by ARM Center of Excellence at University of Edinburgh

Data Breaches Never Been More Relevant

2020 is on Track to Hit a New Data Breach Record

Hackers Attack Every 39 Seconds

COVID-19 News: FBI Reports 300% Increase in Reported Cybercrimes

Containing Data Breaches in Public Cloud

Data breaches happen ☹

- Spyware, side channels, …

Cloud vendors strive to contain stolen info

- Firewalls, authentication, …

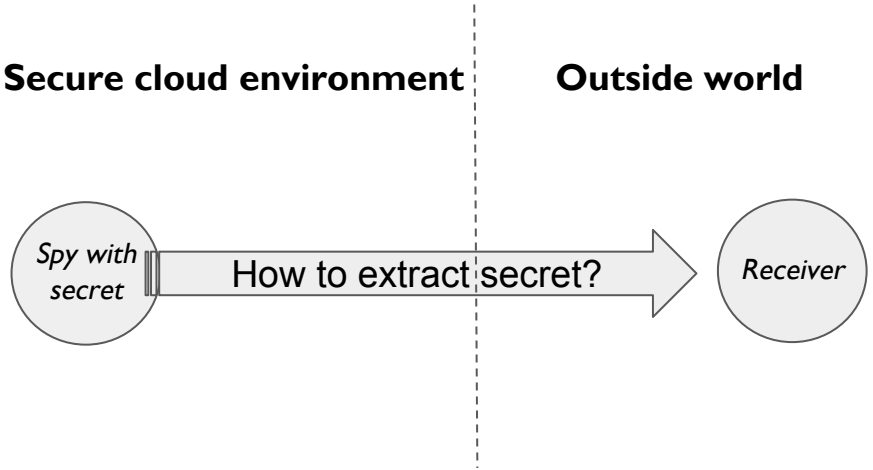

Containing Data Breaches in Public Cloud

Data breaches happen ☹

- Spyware, side channels, …

Cloud vendors strive to contain stolen info

- Firewalls, authentication, …

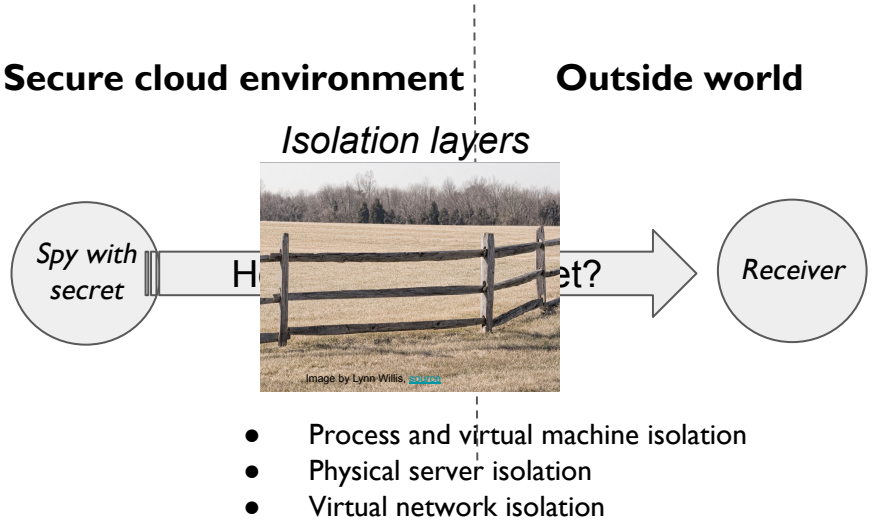

Containing Data Breaches in Public Cloud

Data breaches happen ☹

- Spyware, side channels, …

Cloud vendors strive to contain stolen info

- Firewalls, authentication, …

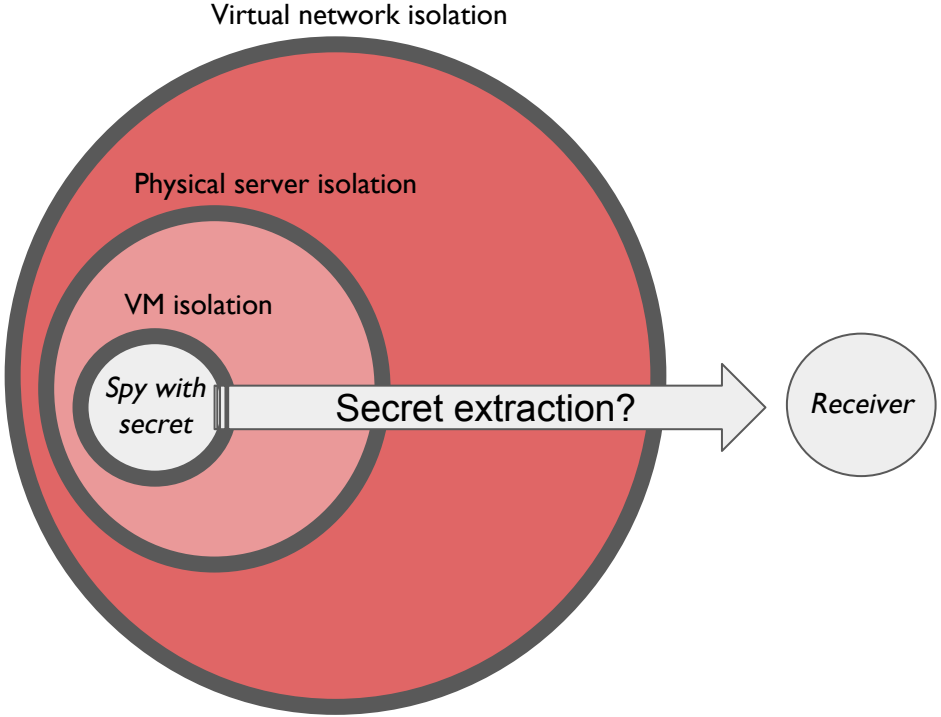

Are secrets safe now?

Covert Channels

Definition: Communication without using legitimate data transfer mechanisms

- Usually via resource sharing (e.g., CPU cache)
- Example: Timing channel via access latency modulation
  - High latency for transmitting "1", low for "0"

Covert channels allow bypassing isolation layers

Network Covert Channels

Allow communication **across** cluster/datacenter

Breach **many** isolation layers **at once**

Stereotypical thinking: Networks are **noisy**

⇨ low accuracy and low throughput channels

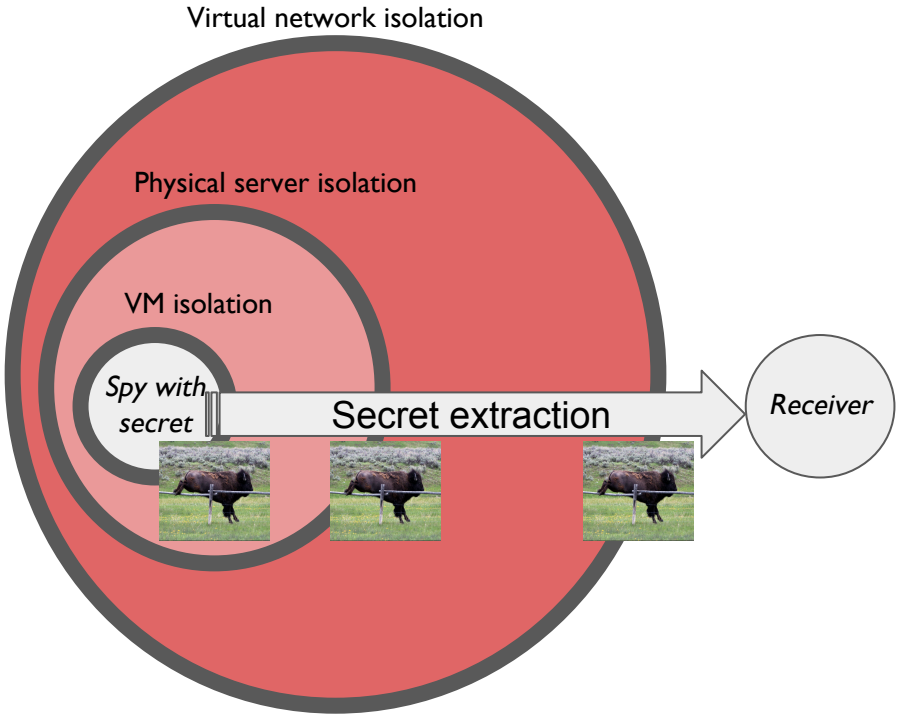

But… Are modern networks noisy?

Emerging Networks in Public Cloud

Remote Direct Memory Access (RDMA)

- Today most cloud providers offer RDMA networks
  - AWS, Azure, Alibaba, Oracle, …

RDMA network packets **bypass** destination CPU

- Low round-trip latency: **2-4µsec**
- High BW with commodity NICs: **100+Gb/s**

Nodes use **one-sided** reads/writes to their **private** data in remote node's memory

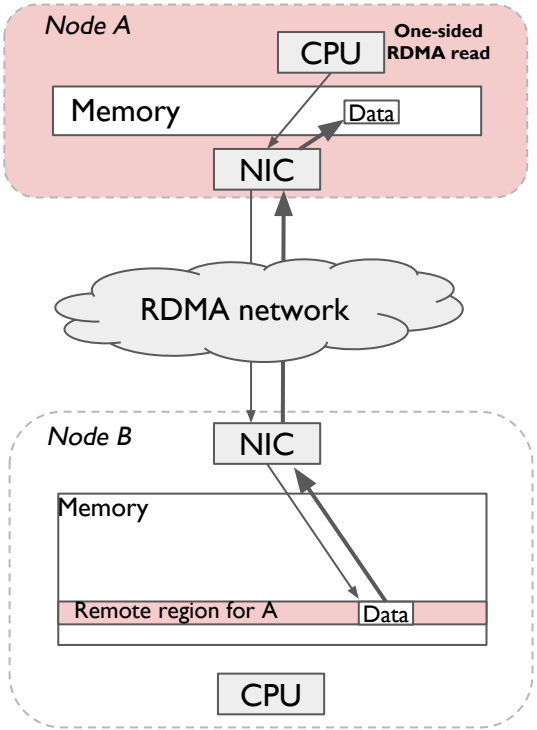

Network BW vs. Memory BW Discrepancy

First glance at bandwidth in modern servers

- RDMA NICs offer 100-200Gb/s
- Memory delivers >100GB/s (=800Gb/s)

Expectation: Memory BW always **much larger**

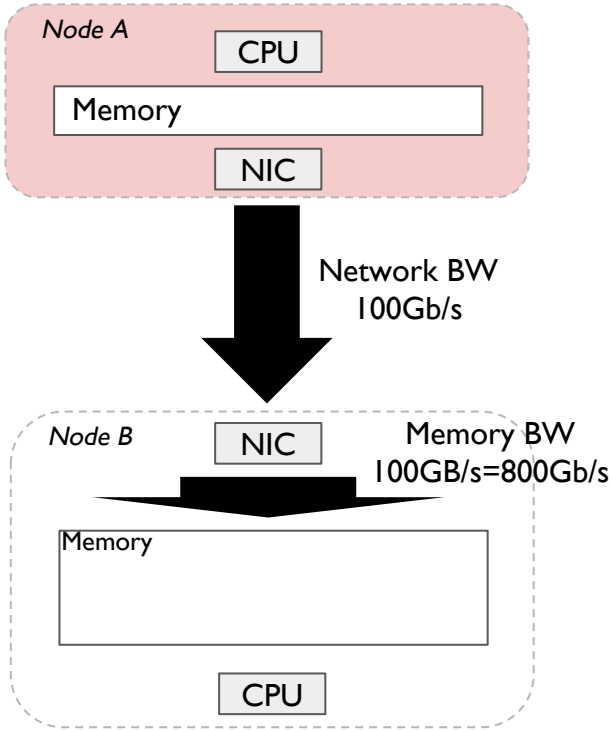

Network BW vs. Memory BW Discrepancy

First glance at bandwidth in modern servers

- RDMA NICs offer 100-200Gb/s
- Memory delivers >100GB/s (=800Gb/s)

Expectation: Memory BW always **much larger**

Wrong!

- Memory has 100s of internal devices (banks)
- Each bank delivers just **~10Gb/s**
  - E.g., same for both Micron DDR2, DDR4
  - Bank behaves as FIFO: ~50ns fixed service time

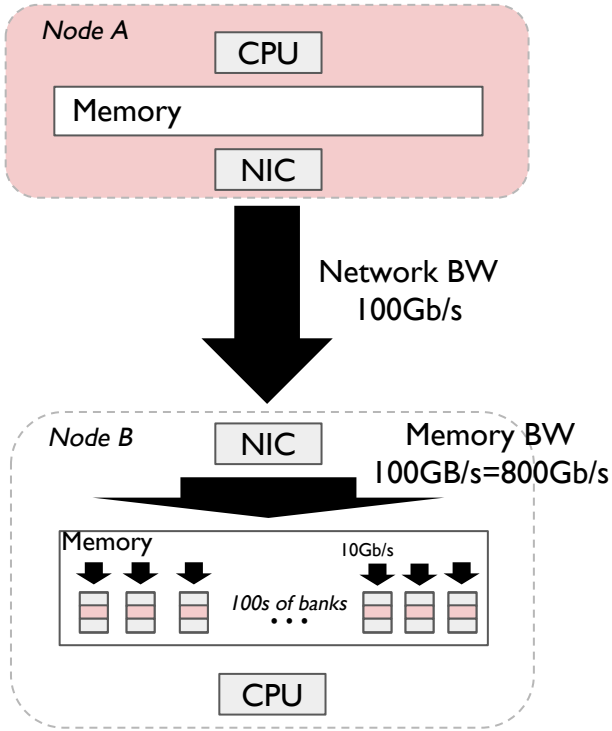

Network traffic can easily congest one memory bank

Bankrupt: RDMA Intra-Cluster Covert Channel

Key features

- No direct communication between Sender and Receiver
- Extremely stealthy!

Basic idea

- Sender transmits the secret by modulating the latency of one memory bank on an **Intermediary** node
- Receiver probes the bank latency and decodes the message
- Intermediary is unrelated innocuous node
  - **No shared memory** between Sender and Receiver

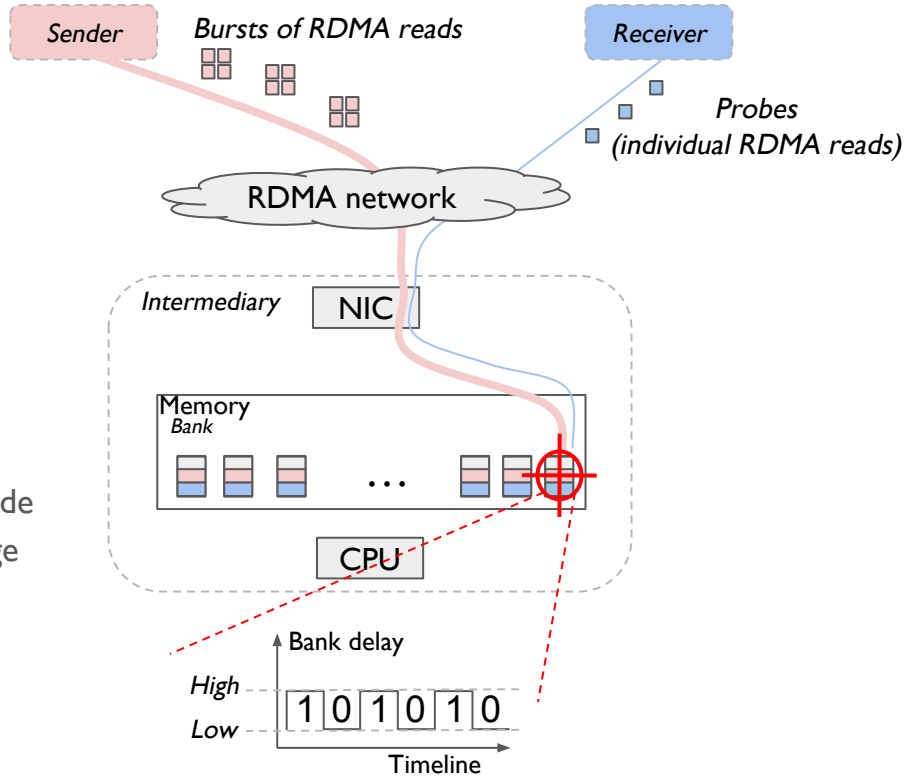

Constructing Bankrupt

1. **Search for addresses that map to target bank**
   - Challenge: CPU hashes addresses to determine the bank
   - Sender and Receiver search addresses independently
   - Addresses different for Sender and Receiver
     - Recall: No memory sharing!

2. **Determining communication parameters**
   - Sender side
     - How many RDMA reads per burst?
     - Transmission frequency?
   - Receiver side
     - Receiving (probes) frequency?

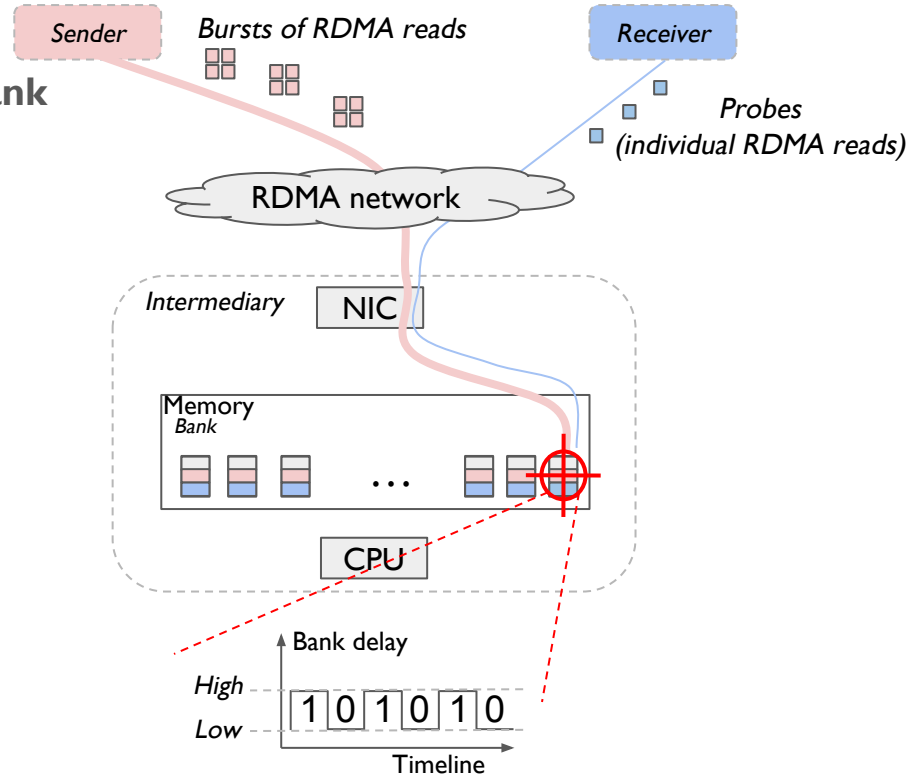

Finding Addresses in Same Bank

Virtual Memory Addressing

Virtual addresses translated to physical upon access

- Translation at **page** granularity
- Same mechanism for local and remote (over RDMA) accesses

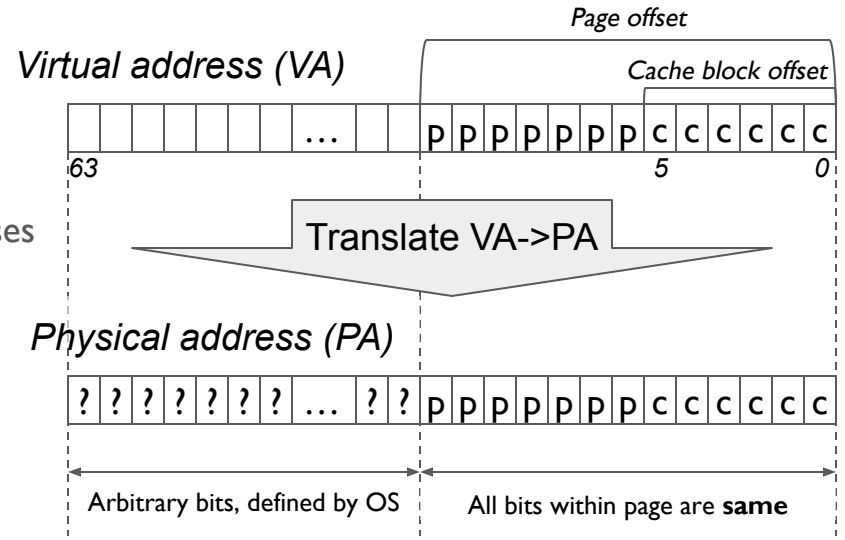

Within a page, physical address bits same as in virtual address

Bank Location

Some physical address bits, "**bank bits**", define bank

- **Low-order** bits to maximize bank-level parallelism

How to find addresses in same bank?

- These addresses have **same** bank bits

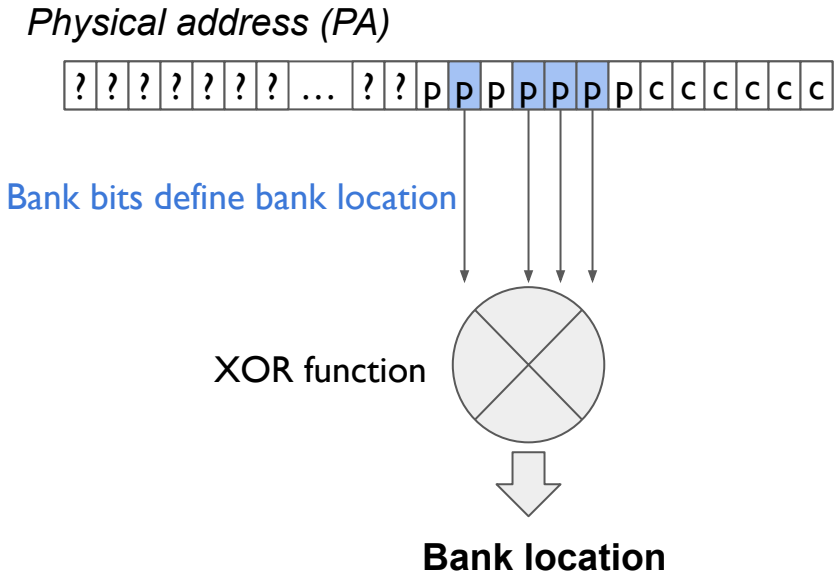

Need to find exact bank bits positions

Same-Bank Addresses Search: Iteration 1

Attacker (Sender and Receiver independently):

1. Chooses arbitrary addresses in remote memory
   - Reads to same cache blocks (64B) coalesced
     ⇨ set **{5:0}** bits to 0

2. Issues RDMA reads to chosen addresses and measures network throughput
   - Network BW = number of serving banks x10Gb/s

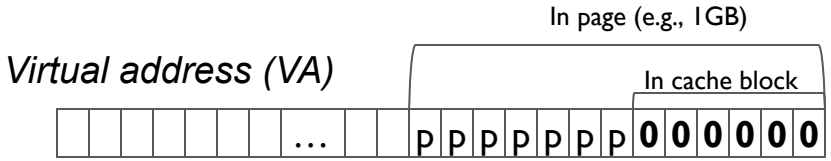

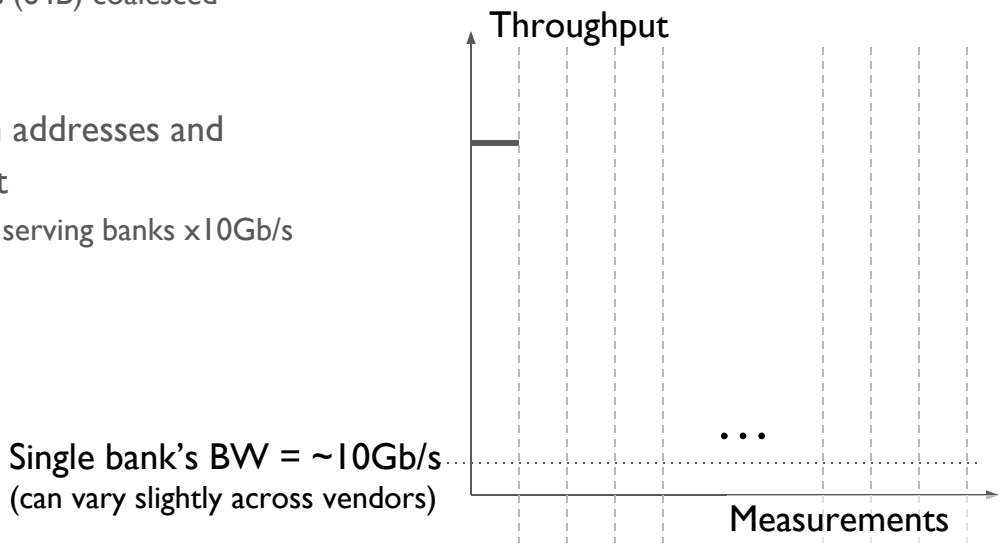

Same-Bank Addresses Search: Iteration 2

Attacker (Sender and Receiver independently)

1. Reduces subset of addresses
   a. Set **{6:0}** bits to 0
2. Issues RDMA reads & measures throughput

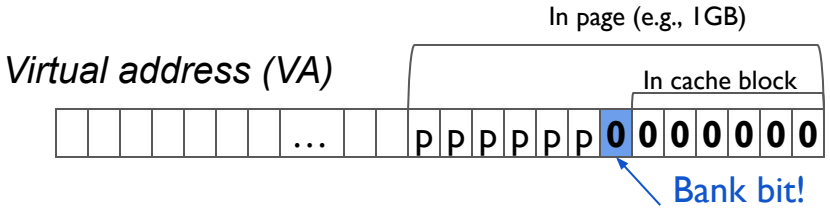

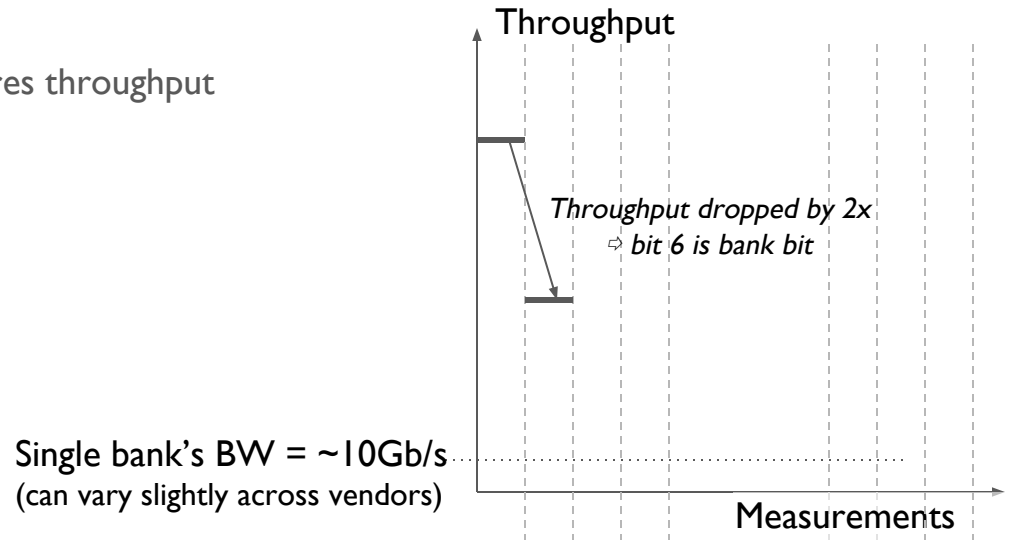

Same-Bank Addresses Search: Iteration 3

Attacker (Sender and Receiver independently)

1. Reduces subset of addresses
   a. Set **{7:0}** bits to 0
2. Issues RDMA reads & measures throughput

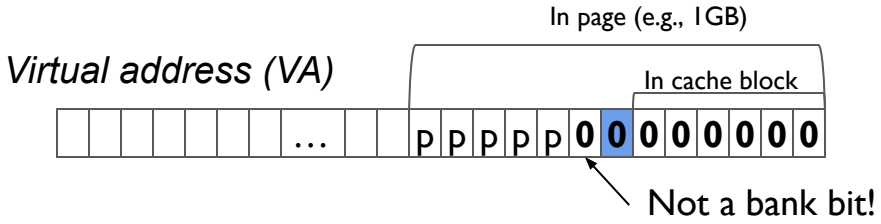

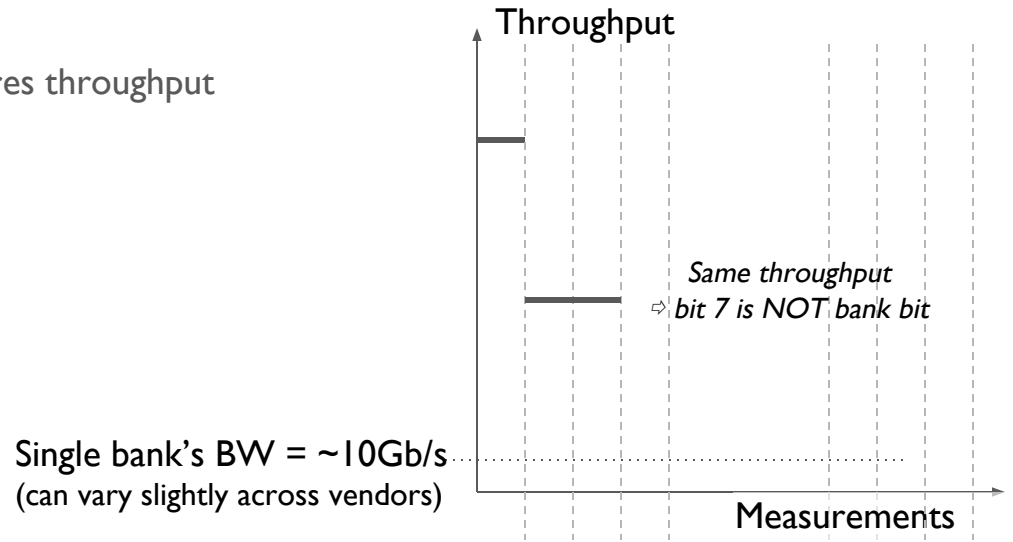

Same-Bank Addresses Search: Iteration 4

Attacker (Sender and Receiver independently)

1. Reduces subset of addresses
   a. Set **{8:0}** bits to 0
2. Issues RDMA reads & measures throughput

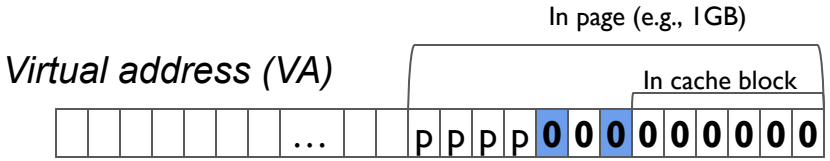

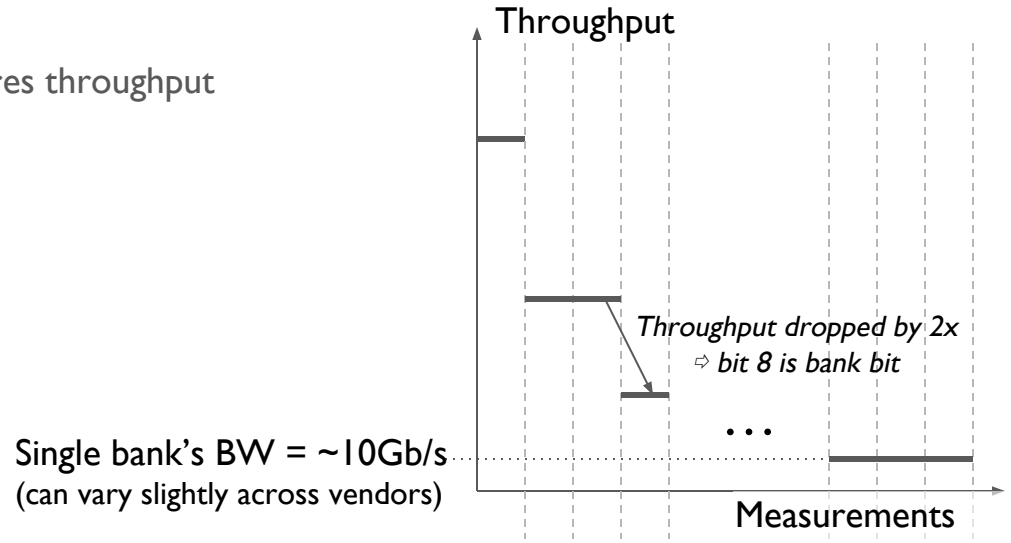

Same-Bank Addresses Search: Iteration N

Attacker (Sender and Receiver independently)

1. Reduces subset of addresses
   a. Set **{N-6:0}** bits to 0
2. Issues RDMA reads & measures throughput

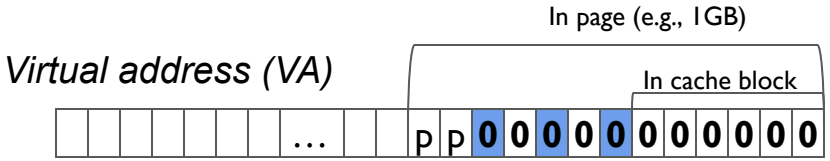

Knowing bank bits locations, choose arbitrary addresses with

- bank bits equal to 0
- cache block bits equal to 0

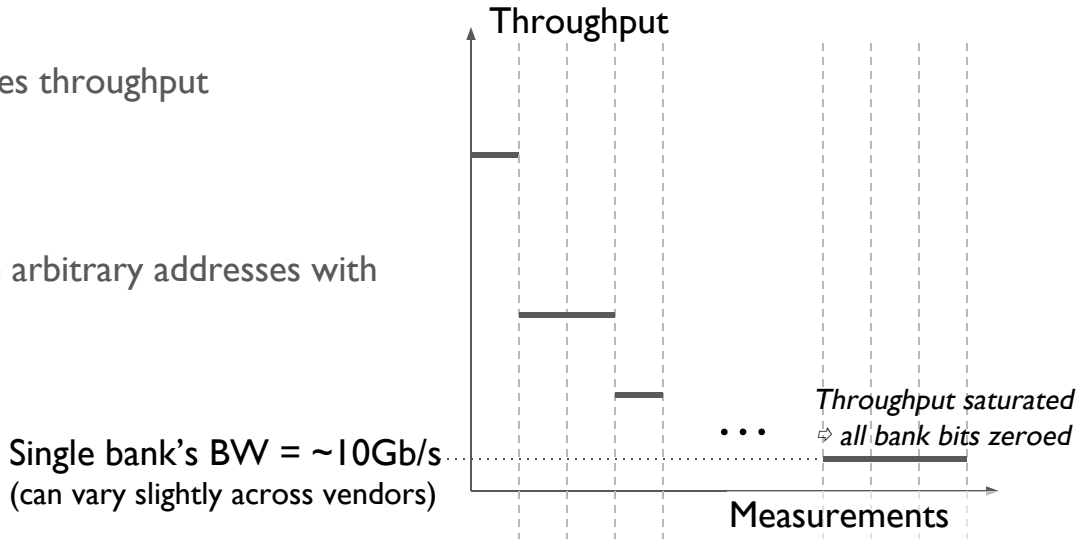

Trivial complexity: Remote attacker finds addresses in <1 second

Determining Communication Parameters

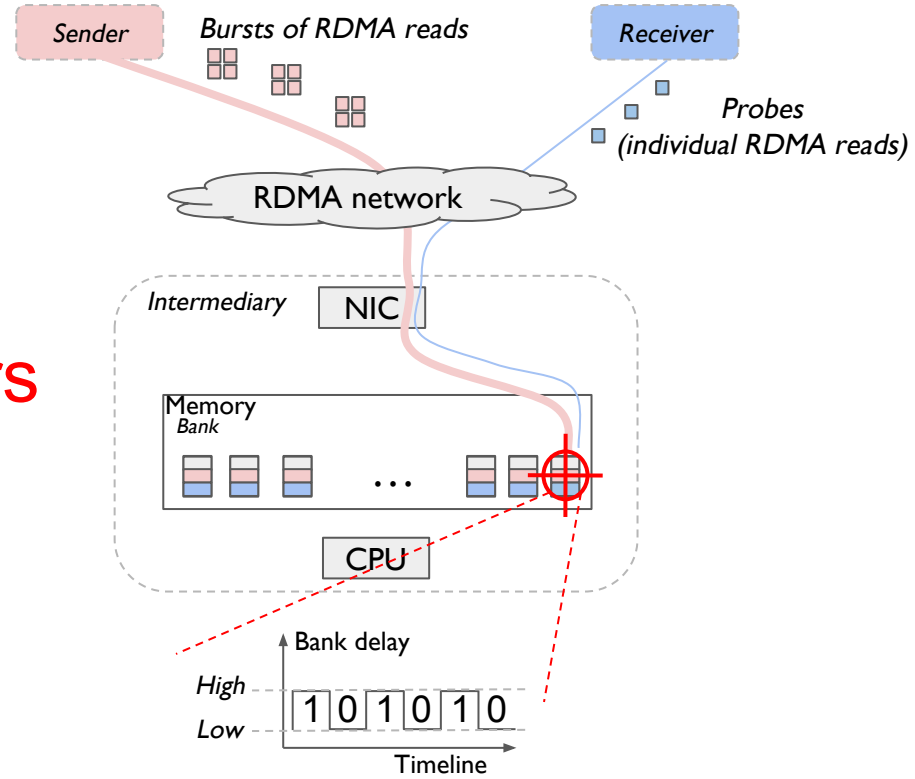

Sender Side

Key parameter: Sender's burst size

- Larger bursts more pronounced
  - ⇨ higher accuracy
    - Especially in noisy networks
- Smaller bursts drain quicker
  - ⇨ higher frequency

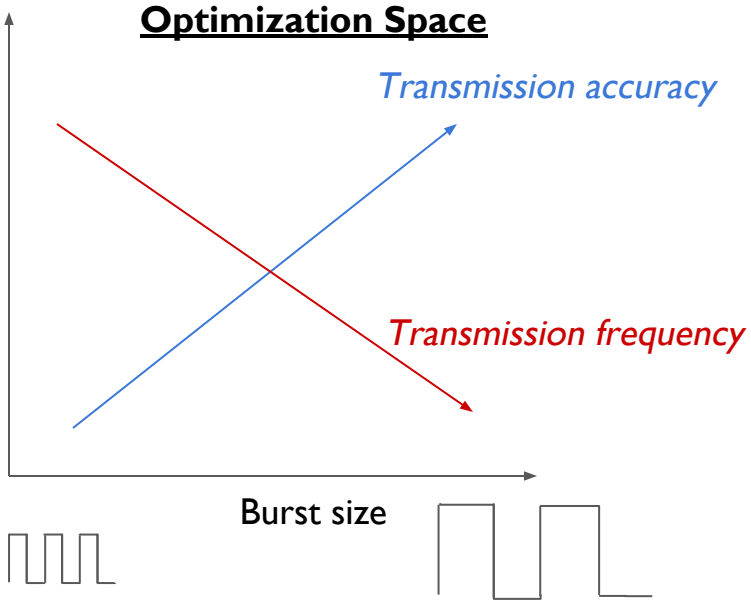

Receiver Side

Transmission period estimation

- Transmitted packets comprise **fixed-size** preamble and payload
  - Example: 32-bit preamble & 200-bit payload
- Receiver iteratively determines the transmission period by looking for **pre-agreed** preamble value

Key parameter: Probing frequency

- Several probes (measurements) per transmission period
- Found little sensitivity on decoding accuracy with probing frequency > 2MHz (1/0.5µseconds)

Probe Round-Trip Delay

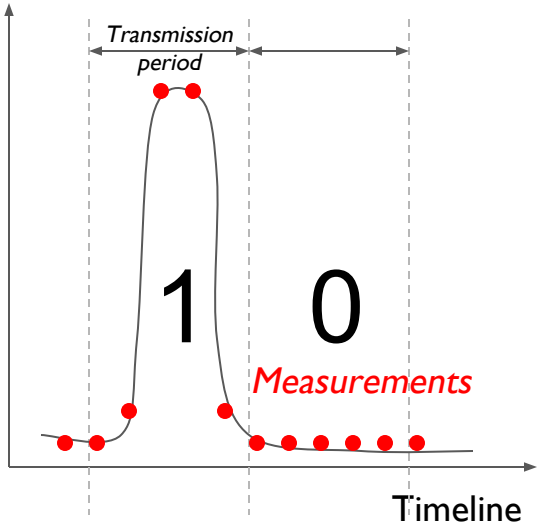

Evaluation Platforms

Private Cluster
(isolated and loaded network)

Cluster size: **6 nodes**
Infiniband
CPU: Xeon E5-2630v4 (Broadwell)
RAM: 64GB, DDR4-2400
NIC: Mellanox **CX-5, 56Gb/s**

Public Cloud: CloudLab Utah
(80% utilized during measurements)

Cluster size: **200 nodes**
Infiniband
CPU: Xeon E5-2640v4 (Broadwell)
RAM: 64GB, DDR4-2400
NIC: Mellanox **CX-4, 50Gb/s**

Private Cluster: Isolated Environment

Burst size:

32

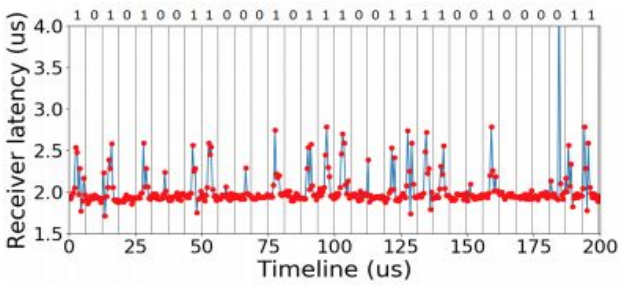

Burst size 32 (x 64 bytes) minimum required for reliable decoding

128

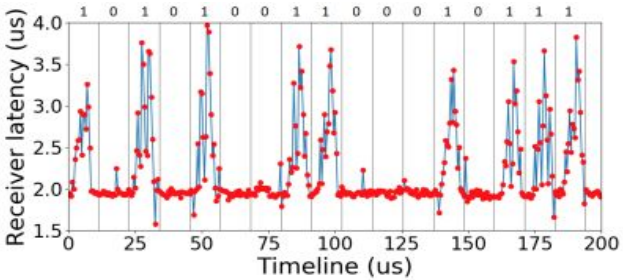

Larger burst sizes

- Decrease transmission frequency
- Make signal more pronounced
  - Larger gap between high and low delay

Private Cluster: Noisy Environment

Burst size:

32

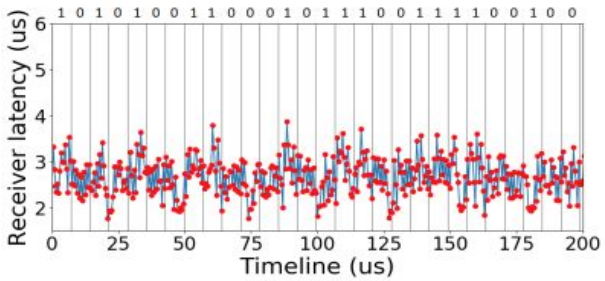

Network loading µbenchmark issues 40Gb/s (70% link BW) RDMA read traffic to Intermediary

Signal with burst size 32 indistinguishable

128

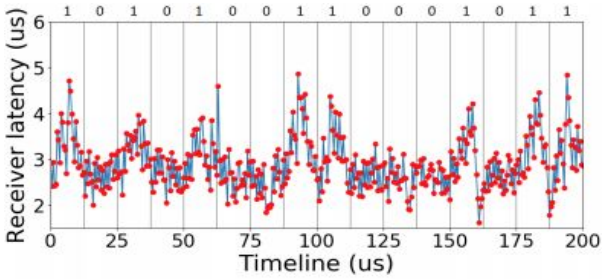

Signal clear with burst size of 128

Noise efficiently compensated for with larger burst size

Private Cluster: Stealthiness

1. CPU hardware counters
   - Memory bandwidth monitoring
     - Bankrupt loads only one bank, <1% of memory
     - CPU counters too coarse-grain

2. Measure local memory access time (LMAT) with software random-access µbenchmark
   - Using **`RDTSC`** timestamps
   - With Bankrupt, LMAT is affected only at **$99.9^{\text{-th}}$ percentile**

3. Network counters non-helpful: No network resources congested

Virtually undetectable with current HW and SW

Private Cluster vs. CloudLab: Throughput & Accuracy

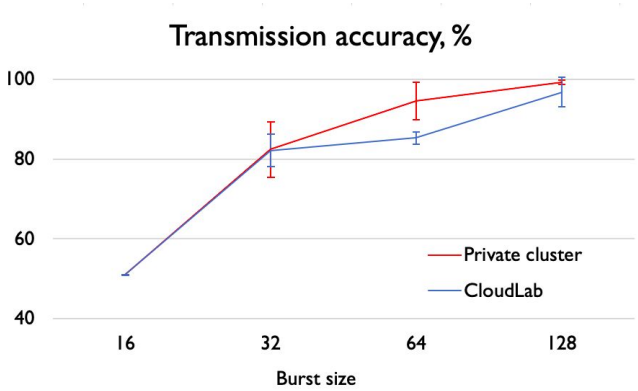

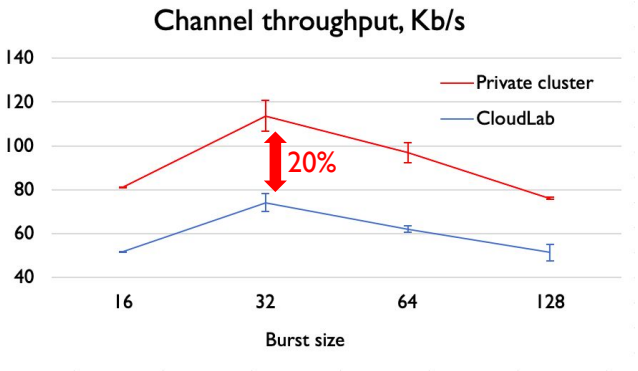

Accuracy as % of correctly transmitted bits

Throughput: True data rate without preambles and errors

Similar accuracy but 20% lower throughput in CloudLab

- Used larger preambles to improve accuracy

Optimal burst size of 32

Sweet spot between transmission frequency & accuracy

THE UNIVERSITY of EDINBURGH
informatics

Takeaways

Covert channels allow bypassing cloud isolation layers

We introduce **Bankrupt** covert channel

- No direct communication between Sender and Receiver
- Affects the timing of single memory bank on Intermediary node
- Delivers **74Kb**/s throughput & robust in **noisy** public cloud network

See paper for mitigation strategies and other details

THE UNIVERSITY of EDINBURGH
informatics

Thank you!

Source code available at: github.com/ease-lab/bankrupt

Contact details: dmitrii.ustiugov(at)ed.ac.uk

Authors thank **ARM Center of Excellence** at University of Edinburgh for their support