Performance Analysis of Large Receive Offload in a Xen Virtualized System

Hitoshi Oi and Fumio Nakajima

The University of Aizu

January 23, 2009

International Conference on Computer Engineering and Technology 2009

Outline

- Performance Interference in Consolidated System
- Objectives and Methodologies
- Large Receive Offload (LRO)
- Xen Internal Network Architecture
- Experimental Results
- Conclusions and Future Work

Performance Interference due to Consolidation

- Xen provides multiple computing environments (*domains*).
- Each domain looks as a logically independent machine: (runs different instances of OS, with different policies, etc).
- Domains share system resources and better utilization is expected by consolidation.
- Conflicting resource demands from domains are controlled by Xen, in a better way than a shared server, but not perfectly – leading to performance interference among domains.

Objectives and Methodologies

- Run two domains on a Xen virtualized system and analyze the performance interference between domains.
- One domain runs a CPU-intensive workload (SPECjbb), a Java-implementation of TPC-C with varying CPU utilization.
- Another domain runs `netserver`, the receiver part of the `netperf` network instrumentation suite. We measure how workloads of these domains interfere each other.
- Large receive offload, an optimization for network throughput, is ported to Xen. We see how it affects the standalone and consolidated performance.

Large Receive Offload (LRO)

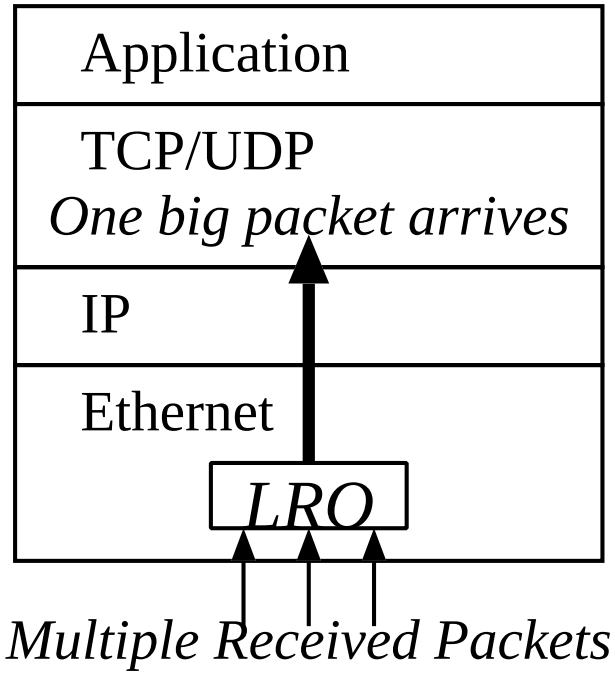

- Aggregate multiple receive packets into a single large packet.
- Reduce the overhead in packet handling operations leading to lower CPU time and higher throughput.
- First implemented in HW (in Neterion NIC) and then in SW (as a patch to Linux device driver).

Xen Internal Network Architecture

LRO has been ported to Physical and Virtual NICs in Xen (pLRO and vLRO).

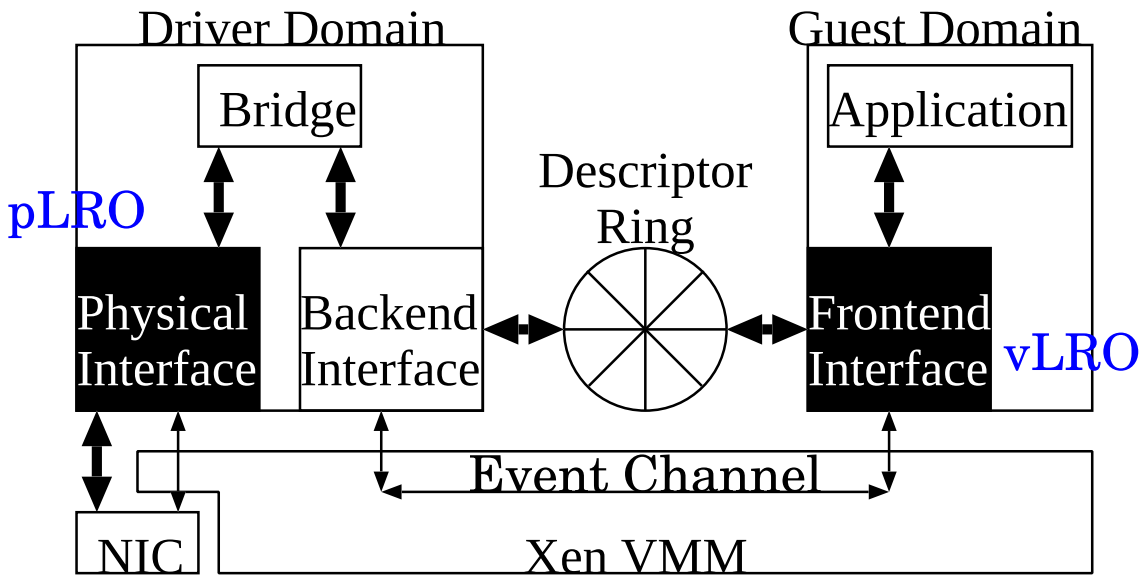

Benchmarking Environments

Component	Description
CPU	Xeon 3GHz
Memory	2GB
NIC	1Gbps
Operating System	Linux 2.6.18
VMM	Xen 3.1.1
Network Measurement	Netperf 2.4.4
Java workload	SPECjbb2001 v1.04
Guest Domains	
vCPU	1/domain
Memory	512MB/domain

Throughput (MTU = 1500B)

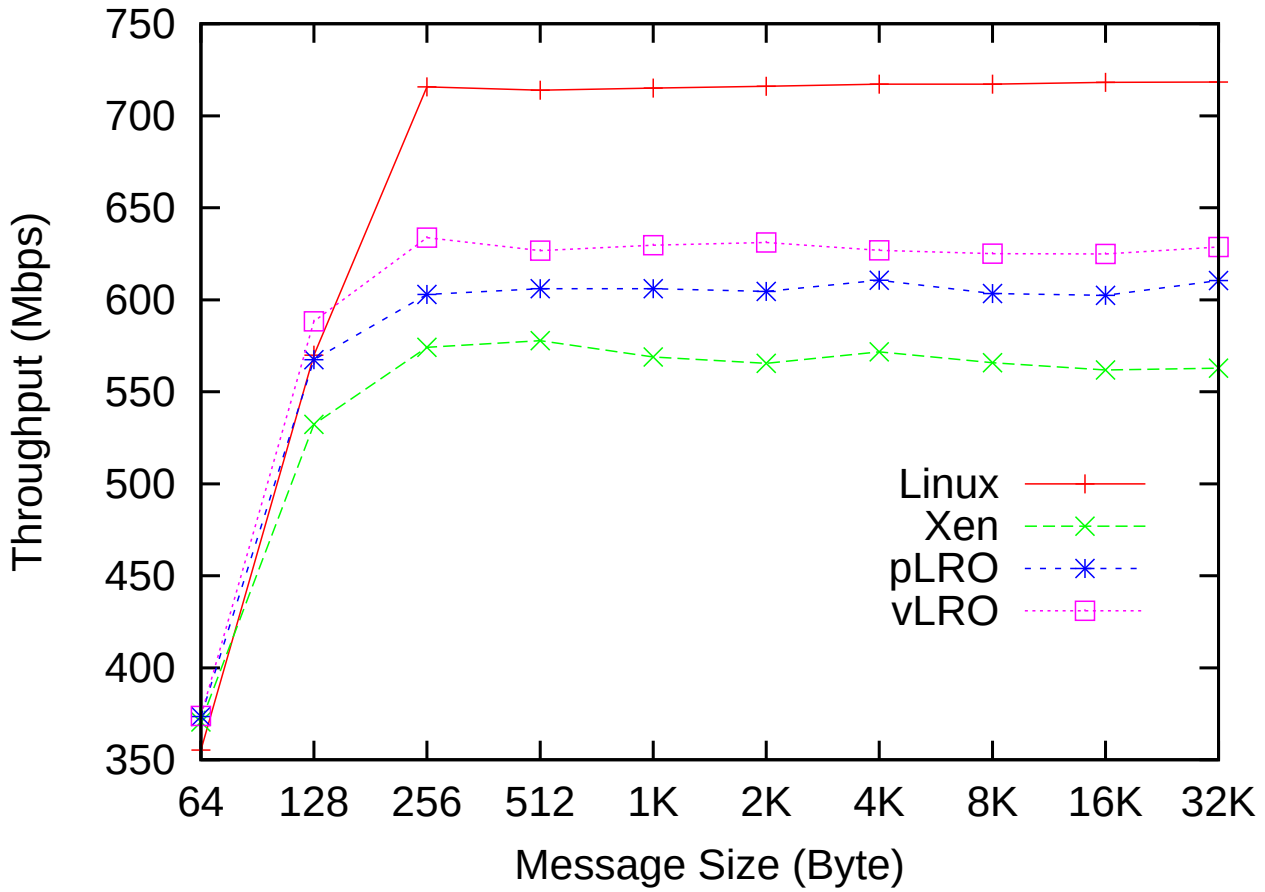

Throughput (MTU = 4000B)

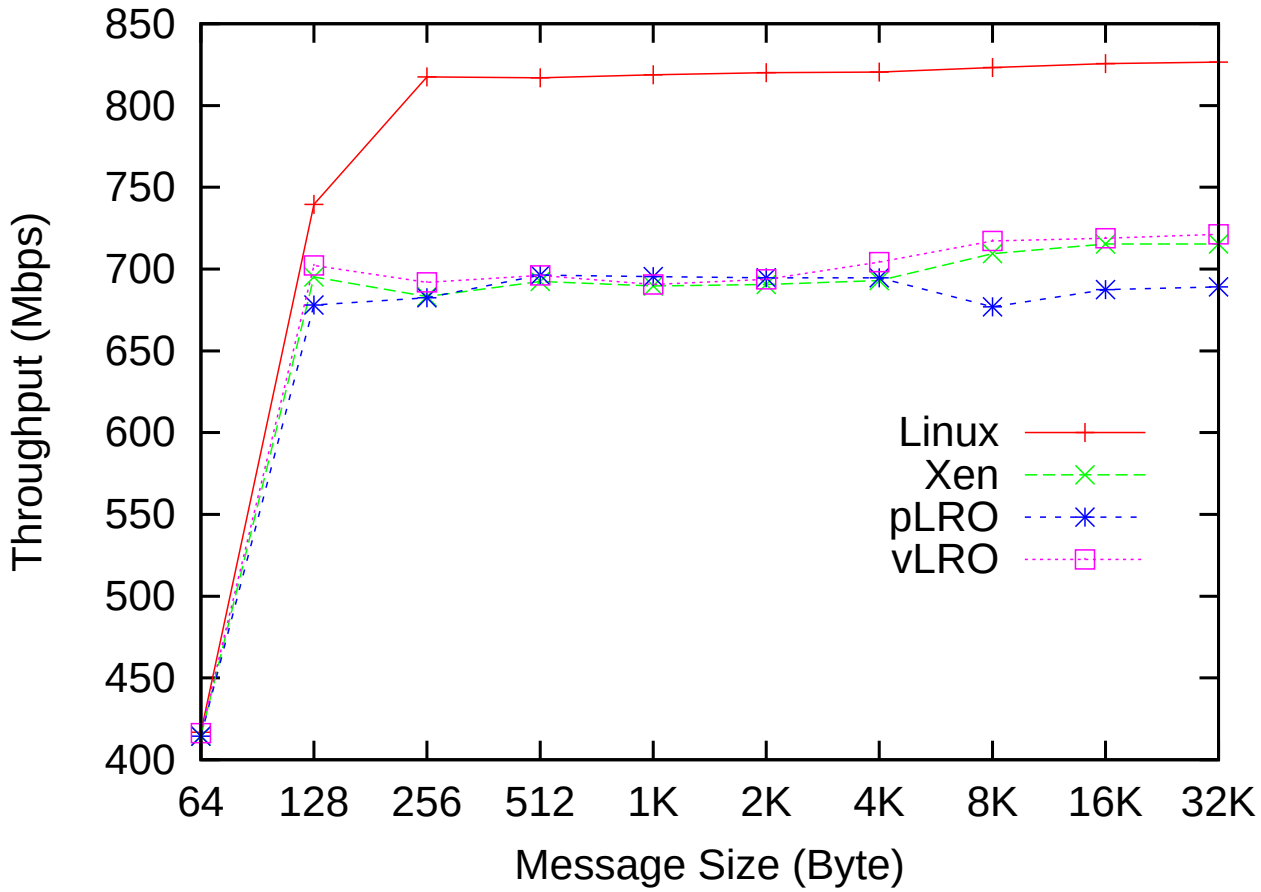

Throughput (MTU = 4500B)

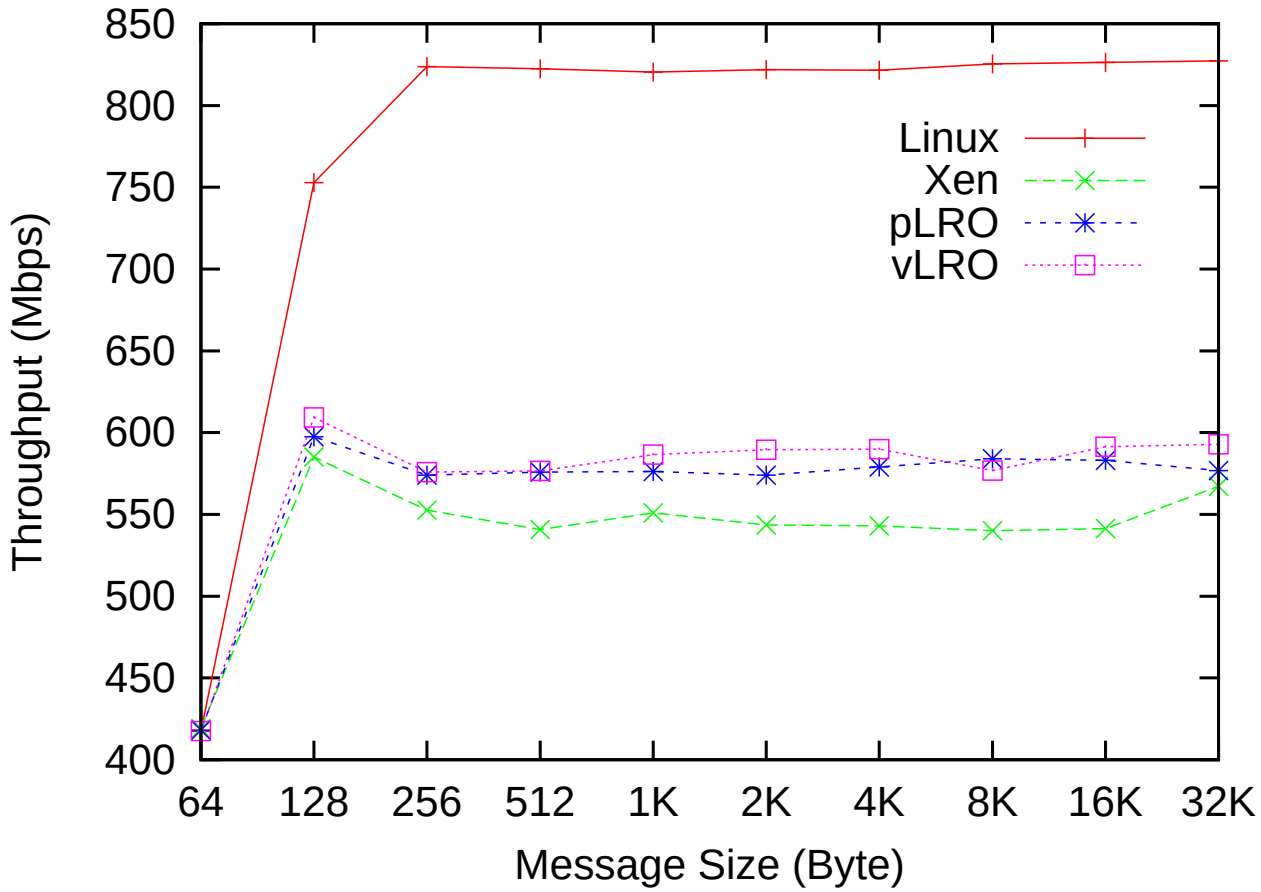

LRO Rate

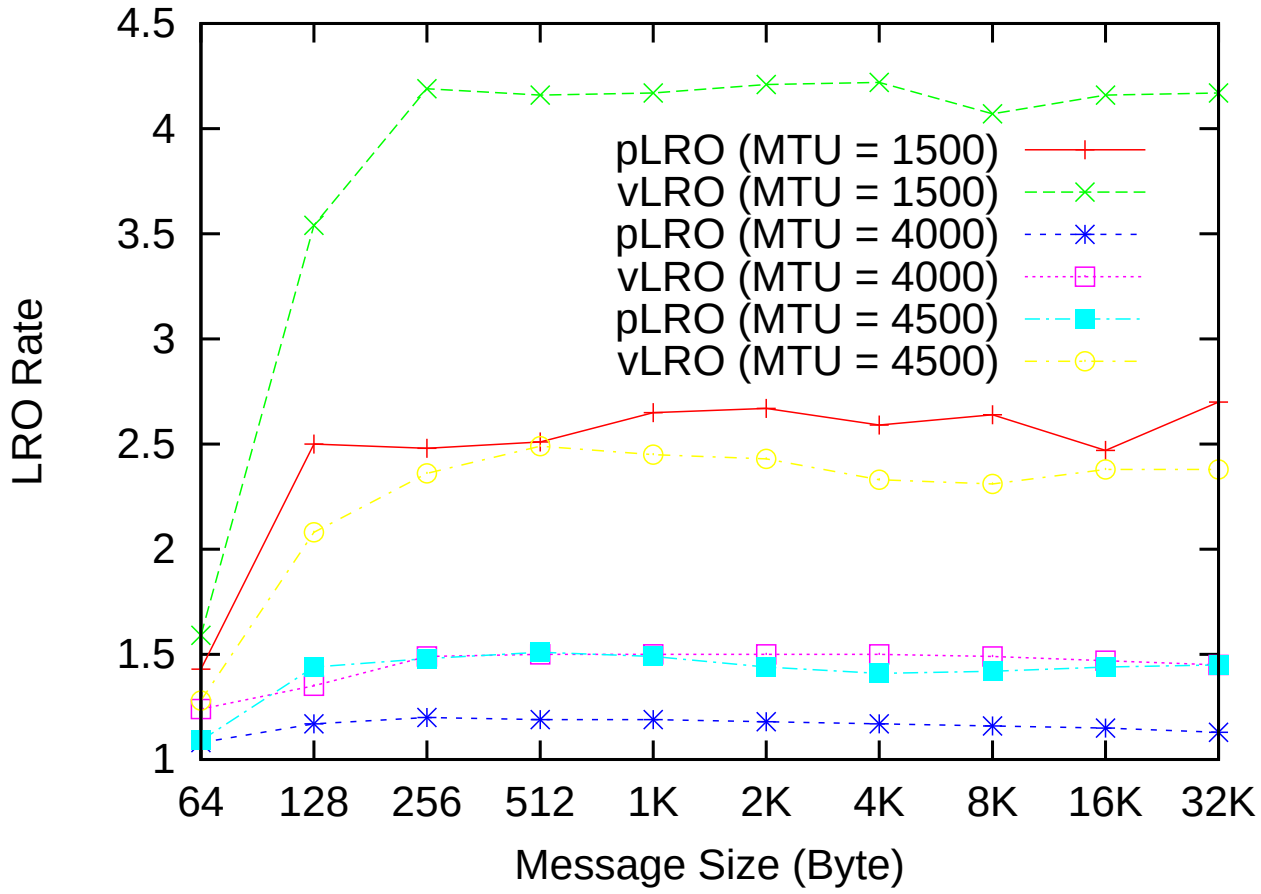
LRO Rate
4.5
4
3.5
3
2.5
2
1.5
1
64
128
256
512
1K
2K
4K
8K
16K
32K
Message Size (Byte)
pLRO (MTU = 1500)
vLRO (MTU = 1500)
pLRO (MTU = 4000)
vLRO (MTU = 4000)
pLRO (MTU = 4500)
vLRO (MTU = 4500)

Consolidated Performance

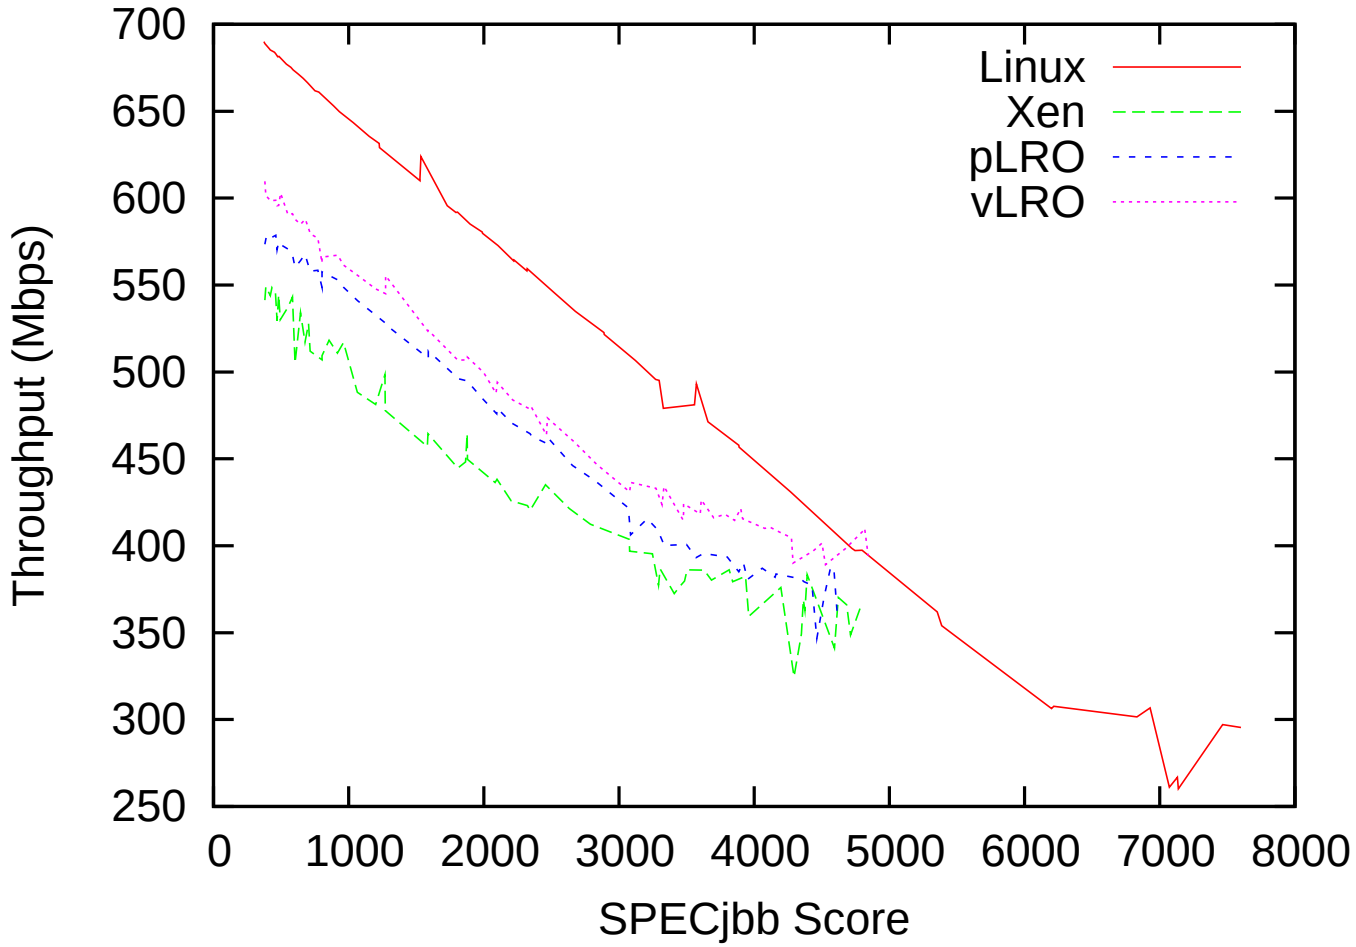

Conclusions and Future Work

- Ported LRO into NICs of Xen and evaluated its effectiveness under varying MTU and message length.
- Run two guest domains on Xen, one with SPECjbb and the other with network receiver.
- Evaluated performance interference between two domains and effectiveness of LRO in this environment.
- Topics of future work include: (1) LRO optimization, (2) further investigation and modeling of performance interference.

Thanks for You Attention

Any question ?

(Akabeko – a local handicraft of Aizu)