

Report on memory writing performance tests

Tuan Mate Nguyen

20/11/2015

Differences between physmem and PDA

- PDA: a library for programming microdrivers
 - http://compeng.uni-frankfurt.de/fileadmin/Images/pda/eschweiler_lindenstruth_pda_rtlws.pdf
- Move large parts of a device driver code to the user space without loss of speed
- The device driver code is compatible with at least 28 kernel releases, easier to maintain

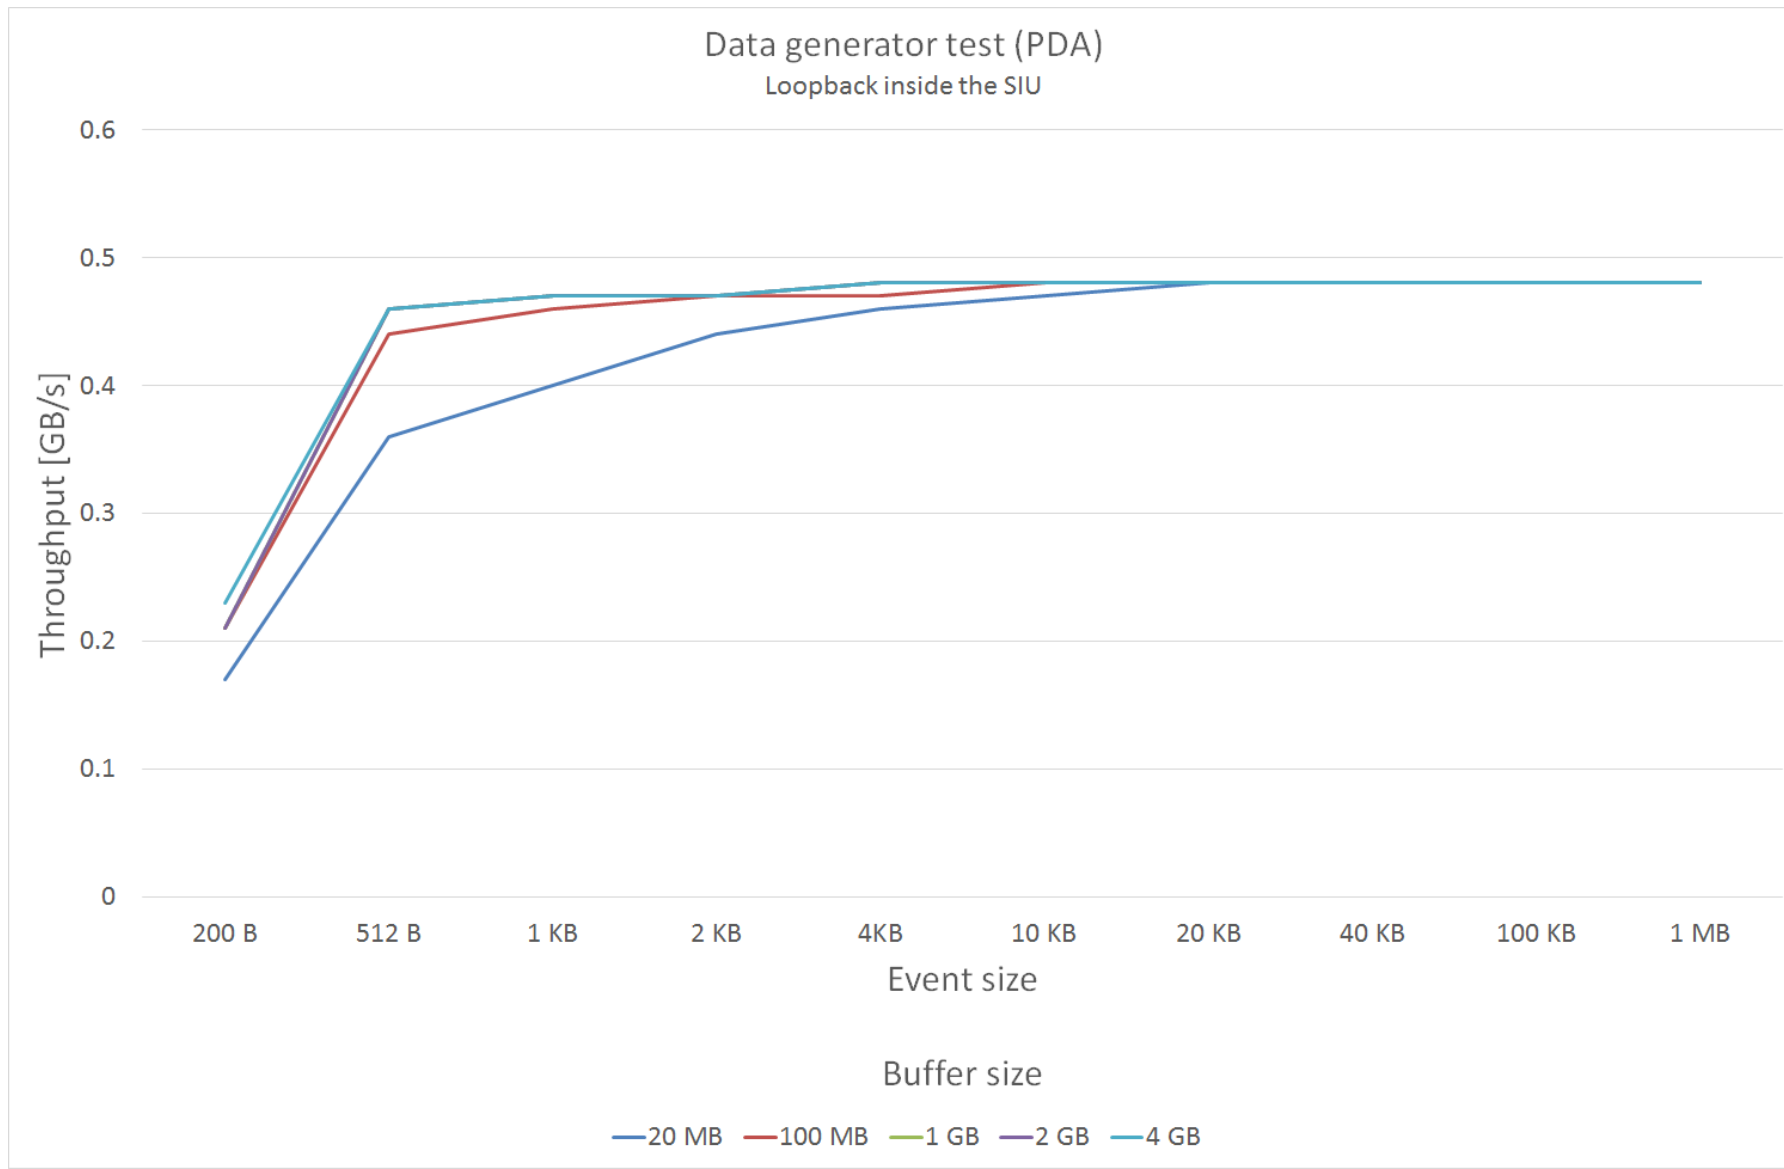
- Memory allocation: opening and writing to a sysfs file (*request*)
 - This triggers a callback function which allocates the requested memory buffer
- Memory buffer:
 - allocation from kernel space
 - More robust
 - allocation from user space
 - User space owns the memory but the device can access it even after deallocation
 - Translates virtual memory addresses into physical addresses
 - Increments reference counter on each physical page
- Usually memory is not continuous->scatter/gather list
 - Contrary to physmem's continuous buffer
- NUMA control

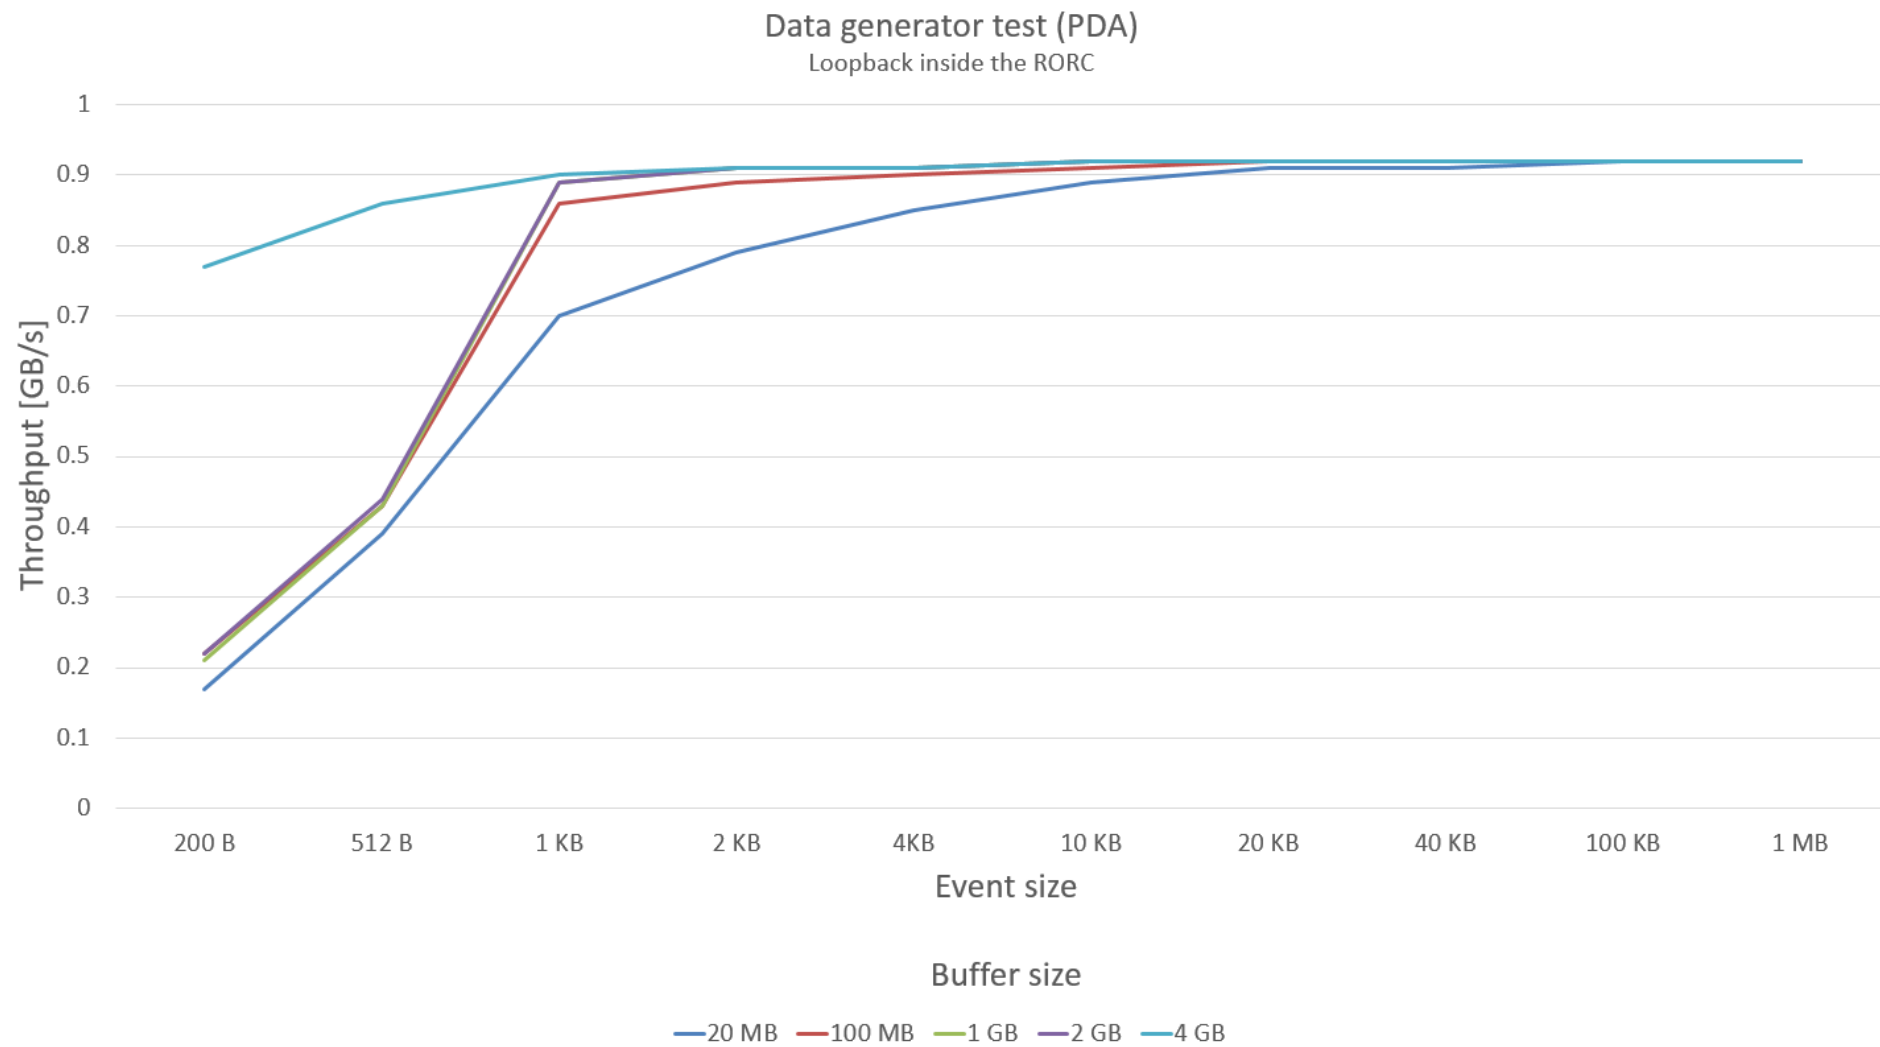
Import RORC library functions to use with PDA library

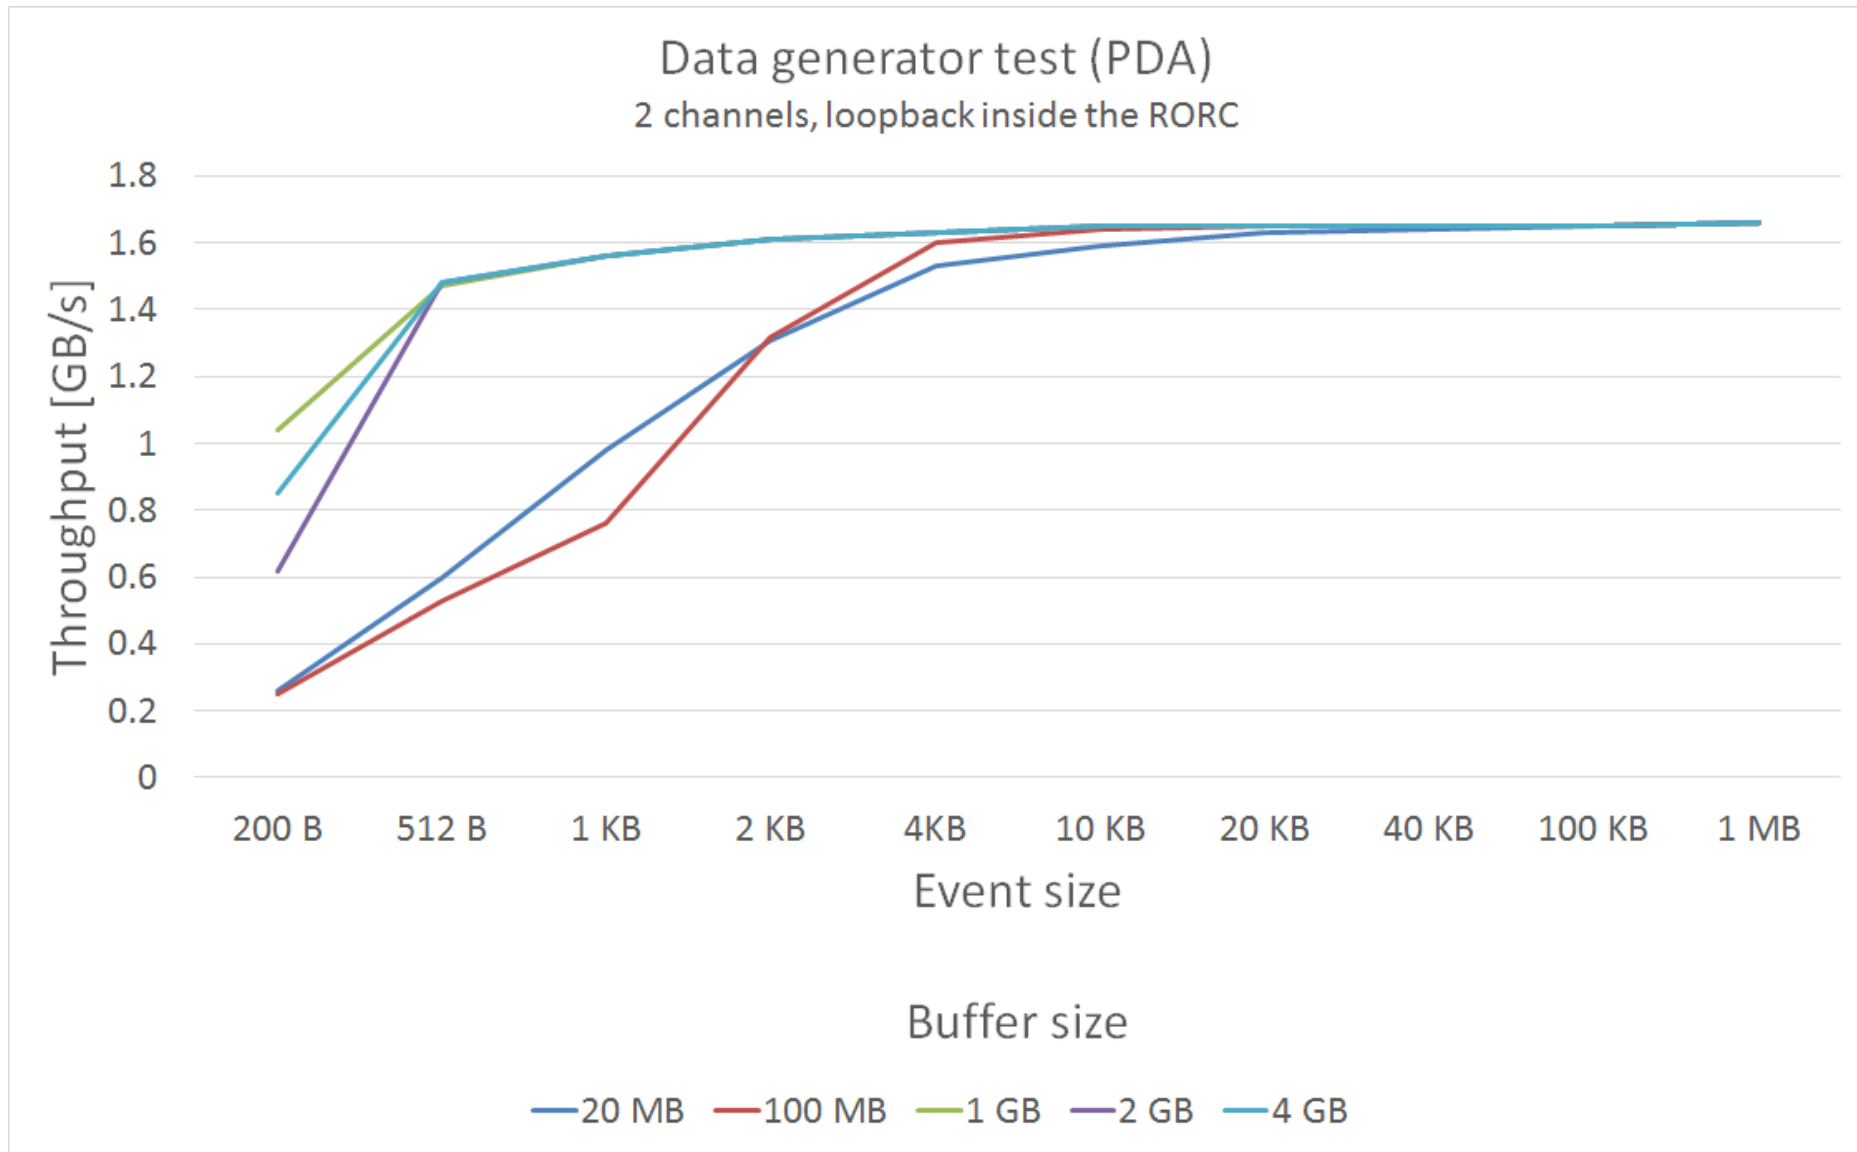
- PDA library provides functions to access a PCI device's configuration space
- Getting the buffer registers' addresses basically allows one to control the rorc by simply reading from/writing to the registers
- Imported the necessary functions from rorc_lib, rorc_ddl and rorc_receive for DMA
- Slight modifications were needed in order to be compatible with PDA

Two performance tests with the RORC's internal data generator

- Event sizes from 200 B to 1 MB, buffer sizes from 20 MB to 4 GB
- Page size constant, 2 MB
- Average of 30 measurements
- Loopback inside the SIU
 - Good for validation
 - The data link means a physical limitation of 0.5 GB/s
 - The achieved max throughput is 0.48 GB/s
- Loopback inside the RORC
 - The achieved max throughput is 0.92 GB/s
- Faster for bigger event sizes – as expected
 - More data is written, the overhead is the same
- Faster for bigger buffer sizes – as expected
 - Scatter/gather lists of bigger allocated buffers contains bigger continuous memory blocks







Ram performance test

- Same setup but CPU writes blocks of data (event size) to the memory
- Dependency on event size
 - Maximum throughput at around 4 KB which is the system's page size
 - Max ~18 GB/s
 - After that every 4 KB means a new page: additional overhead – a drop in performance
 - The bottleneck is not the RAM
- Speed difference:
 - Faster for smaller buffer size: 20, 100 MB
 - possible explanation:
 - Size of L3 cache is 15 MB, this could make a difference in case of small buffers
- Same results with physmem