Impact of remote memory and network performance on execution performance of disaggregated micro data centers

Akishige IKOMA, Yuichi OHSITA, Masayuki MURATA
Graduate School of Information Science and Technology, Osaka university

Micro Data Center (µDC)

- Cloud data center problems
  - Large latency between cloud data center and end device
  - Difficult to provide time-sensitive applications
- Micro data center (µDC)
  - Small data centers deployed near users
    - Smaller latency than the cloud
  - More limited amount of resources compared with large DC

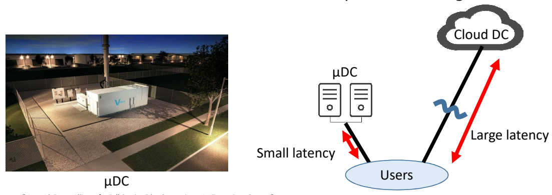

Source ："Vapor IO Forms Alliance for Full-Service Edge Computing at Its Tower Data Centers" https://www.datacenterknowledge.com/vapor-io/vapor-io-forms-alliance-full-service-edge-computing-its-tower-data-centers, last accessed on 2021-11-28.

Disaggregated µDC (µDDC)

- µDC constructed of resources connected by a network
  - Achieve efficient resource utilization
    - Optimization per resource
  - More easily Resource upgrade
- Assigned to each resource for each process
  - Improving resource utilization

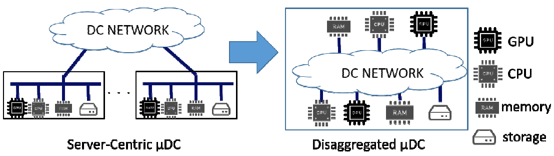

Differences between a server-centric data center and a disaggregated micro data center

Problem of disaggregation

- Performance degradation by a network latency
  - Latency due to communication between resources
- Communication delay between the CPU and remote memory has large impact[5]
  - Direct cause of performance degradation
  - Need to work to reduce the impact of delays

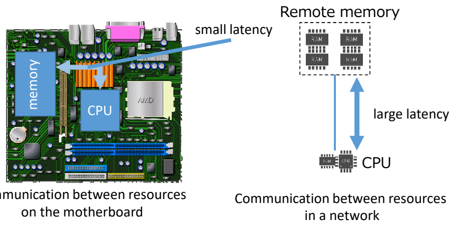

[5] P. X. Gao, A. Narayan, S. Karandikar, J. Carreira, S. Han, R. Agarwal, S. Ratnasamy, and S. Shenker, "Network requirements for resource disaggregation," in 12th USENIX Symposium on Operating Systems Design and Implementation (OSDI 16), pp. 249– 264, Nov. 2016.

Approach

- Research Objectives
  - Investigating the impact of network performance of µDDC on execution performance
- Approach
  - Add communication delays when accessing remote memory to emulate applications on a µDDC
  - Measure the performance degradation rate with execution time for image classification processes
    - Assumed to run at the edge
  - Measure the ratio of the communication time between CPU and memory to the execution time

Evaluation environment

- CPU mounted with cache (local memory)
  - Low latency between CPU and cache
  - Local memory size：200MB
- Remote memory communicates via the local memory
  - Communicate when a page fault occurs
- Communicating uses a paging technique
  - Page size：4KB

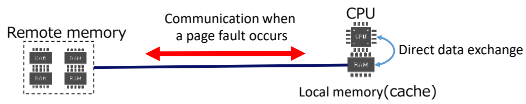

Communication between remote memory and local memory

Emulation

- Insert delay when communicating between local memory and remote memory
- Split the original memory into local memory and remote memory
- Insertion delay is based on latency and bandwidth
  - Latency and bandwidth are parameters

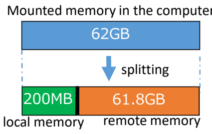

Create method of local memory and remote memory

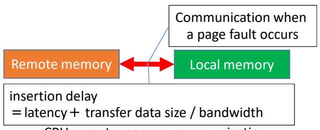

CPU-remote memory communication

CPU	: Intel(R) Xeon(R) CPU E5-2687W 0 @ 3.10GHz
memory	: DDR3 SDRAM
OS	: CentOS7.7

CPU, memory, and OS used in the experiment

Evaluation

- Evaluate performance by running machine learning image classification
  - Run 4 machine learning models with Tensorf
- Evaluate performance by performance degradation rate
  - Perforamance degradation rate = $\frac{T^{disaggregated} - T^{traditional}}{T^{traditional}}$
    - $T^{disaggregated}$: μDDC's process time
    - $T^{traditional}$: traditional computer's process time
- Changing latency and bandwidth to study the impact on performance degradation rate

	latency	bandwidth
evaluation1	Fixed 8μs	40Gbps,100Gbps
evaluation2	0.2μs,2μs,8 μs	Fixed 100Gbps

Value of the parameter for each evaluation

Result

- ResNet50 had minimal impact on performance
  - Because the model is small and page faults are infrequent
- Small difference in performance degradation with bandwidth
  - Bandwidth has only a small impact on the performance
- As latency increases, performance drops significantly
  - Latency has a large impact on the performance

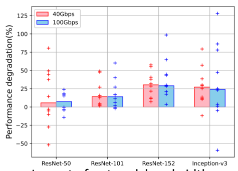

Impact of network bandwidth on performance

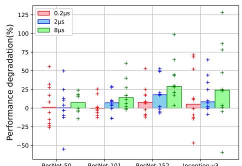

Impact of latency on performance

Result

- ResNet50 had minimal impact on performance
  - Because the model is small and page faults are infrequent
- Small difference in performance degradation with bandwidth
  - Bandwidth has only a small impact on the performance
- As latency increases, performance drops significantly
  - Latency has a large impact on the performance

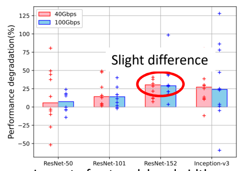

Impact of network bandwidth on performance

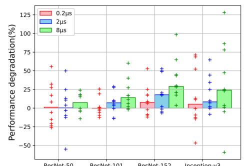

Impact of latency on performance

Result

- ResNet50 had minimal impact on performance
  - Because the model is small and page faults are infrequent
- Small difference in performance degradation with bandwidth
  - Bandwidth has only a small impact on the performance
- As latency increases, performance drops significantly
  - Latency has a large impact on the performance

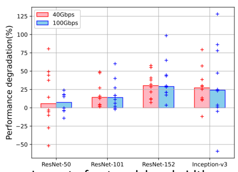

Impact of network bandwidth on performance

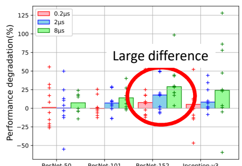

Impact of latency on performance

Remote memory access time

- More than double the time is required to obtain the data when the latency is 8 μs , 0.2 μs
- When the latency is 8μs, the access time occupies about 1/5 to 1/3 of the total time.
- Latency has a significant impact.
- Overall, CPU processing time than remote memory access time is larger

	0.2μs	2μs	8μs
ResNet50	7.94%	10.15%	18.44%
ResNet101	10.96%	14.94%	27.73%
ResNet152	12.50%	16.16%	27.31%
Incerion-v3	9.93%	12.31%	22.24%

Percentage of remote memory access time in the sum of remote memory access time and CPU processing time

Remote memory access time

- More than double the time is required to obtain the data when the latency is 8 μs , 0.2 μs
- When the latency is 8μs, the access time occupies about 1/5 to 1/3 of the total time.

- Latency has a significant impact.
- Overall, CPU processing time than remote memory access time is larger

More than twice the difference

	0.2μs	2μs	8μs
ResNet50	7.94%	10.15%	18.44%
ResNet101	10.96%	14.94%	27.73%
ResNet152	12.50%	16.16%	27.31%
Inceprion-v3	9.93%	12.31%	22.24%

Percentage of remote memory access time
in the sum of remote memory access time and CPU processing time

16

Remote memory access time

- More than double the time is required to obtain the data when the latency is 8 μs , 0.2 μs
- When the latency is 8μs, the access time occupies about 1/5 to 1/3 of the total time.

- Latency has a significant impact.
- Overall, CPU processing time than remote memory access time is larger

The remaining 72.69% is CPU processing time

	0.2μs	2μs	8μs
ResNet50	7.94%	10.15%	18.44%
ResNet101	10.96%	14.94%	27.73%
ResNet152	12.50%	16.16%	27.31%
Inceprion-v3	9.93%	12.31%	22.24%

Percentage of remote memory access time
in the sum of remote memory access time and CPU processing time

17

Discussion

- Bandwidth has only a small impact on the performance
- Effect of using a paging technique
  - CPU obtains only 4 KB of data every time a page fault occurs
  - 4 KB data is very small relative to the bandwidth
- Much of the processing time is consumed by the CPU
- CPU allocation is important
  - Compensation for communication delays
- Need to address both communication delay reduction and appropriate resource allocation methods

19

Conclusion

- Investigating the impact of network performance of μDDC on execution performance
  - Latency has a large impact on the performance
  - Much of the processing time is consumed by the CPU.
  - Communication delay reduction and resource allocation methods is needed
- Future work
  - Establish a resource allocation method specific to disaggregation
    - Considering communication delay and resource performance
    - Consider potential future resource demands.

20