

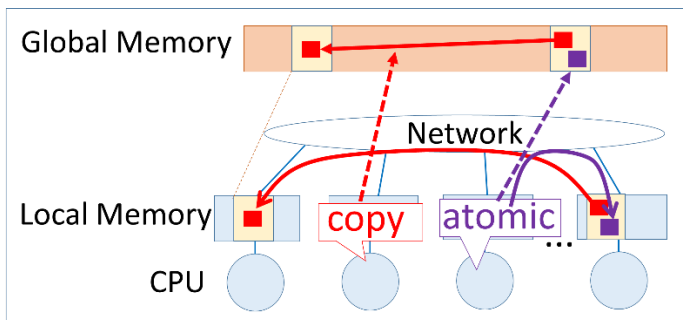
PGAS-based Scalable Communication Library ACP, and its Applications

Overview

Memory-efficiency of communication libraries is becoming important issue in large scale supercomputers where the number of processes is expected to be tens of millions. To achieve both high performance and low memory-consumption, this project has developed a PGAS-based communication library, ACP, and its applications.

ACP (Advanced Communication Primitives): a PGAS-based Communication Library

ACP supports a communication model that transfers data via "Global Memory", that is a virtual memory space managed by the library. Each process can expose a region of its local memory to other processes by registering that region to the global memory.



Basic Layer: Global Memory Access Operations of ACP

● Copy operation


Copy data from one place in the Global Memory to another. Internally, the library performs appropriate communication operations for each copy operation so that the data is transferred between corresponding registered regions of local memories.

● Atomic operation


Permutation, increment/decrement and logarithmic operations, on the value stored in the specified place of the Global Memory.

Fundamental Performance of ACP:

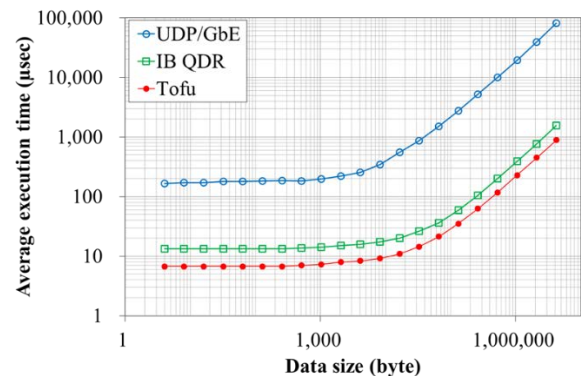
● Local to Remote:



● Remote to Local:



● Remote to Remote



Project Web Site

<http://ace-project.kyushu-u.ac.jp>

ACP Library

Objective: PGAS-based communication library.

<https://github.com/project-ace/ACP>

NSIM-ACE

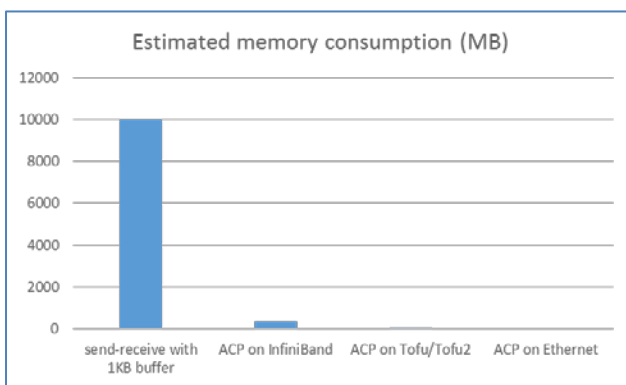
Objective: Network simulator for one-sided communication.

http://ace-project.kyushu-u.ac.jp/main/en/05_download

Estimated Memory Consumption of ACP:

- Estimation with 1,000,000 processes:


Memory	send-receive with 1KB buffer	ACP		
		InfiniBand	Tofu/Tofu2	Ethernet
	10GB	370MB	67MB	34MB



Middle Layer: Additional Operations of ACP


- **Communication Library**

Explicit creation / free of channels among processes that support one-directional send / receive.



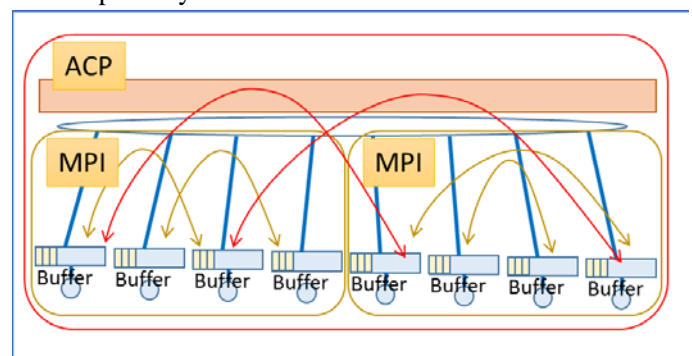
- **Data Library**

Create / modify / destruct data structures, such as vector, list, deque map and set, on the Global Memory.



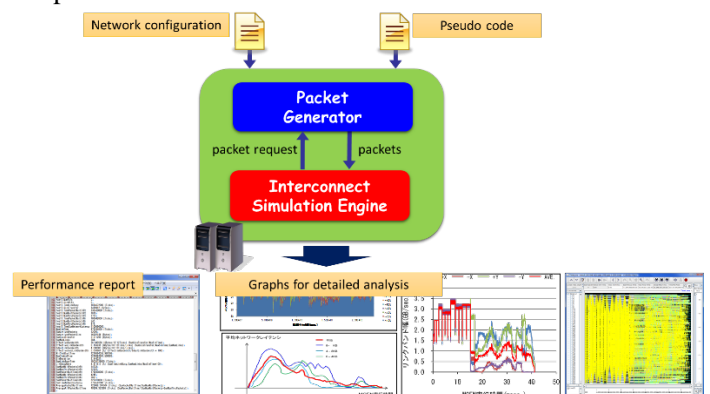
ACP + MPI:

Existing MPI programs can call ACP. Also, ACP can be used to connect multiple MPI programs to reduce memory consumption by communicators of MPI.



Network Simulator NSIM-ACE:

Study behavior of one-sided communications performed with ACP. Packets are generated automatically, and simulated according to the configuration of networks and the pseudo code.



Project information:

Project leader: Takeshi Nanri (Kyushu Univ.)

Project member: Hiroaki Honda, Ryutarou Susukita, Taizo Kobayashi, Yoshiyuki Morie (Kyushu U.), Shinji Sumimoto, Yuichiro Ajima, Naoyuki Shida, Kazushige Saga, Takafumi Nose (Fujitsu Ltd.), Hidetomo Shibamura, Takeshi Soga (ISIT Kyushu), Keiichiro Fukazawa (Kyoto U.), Toshiya Takami (Oita U.)

Contact: takeshi.nanri.995@m.kyushu-u.ac.jp