

# Adaptive auto-tuning in HPX using APEX

**Mohammad Alaul Haque Monil**

Department of Computer and  
Information Science,  
University of Oregon  
Eugene, OR, USA  
mmonil@cs.uoregon.edu

**Bibek Wagle, Hartmut Kaiser**

School of Electrical Engineering and  
Computer Science,  
Louisiana State University  
Baton Rouge, LA, USA  
bwagle3@lsu.edu, hkaiser@cct.lsu.edu

**Kevin Huck**

Performance Research Laboratory,  
University of Oregon  
Eugene, OR, USA  
khuck@cs.uoregon.edu

## ABSTRACT

Finding an optimal value for a parameter that impacts application performance is a hard problem and often requires repetitive execution of the application. This auto-tuning effort by running the application many times causes wastage of HPC resources. In this research, we provide a preliminary study which demonstrates parameter value searching at runtime while ensuring better performance. We use APEX performance measurement library to implement an adaptive auto-tuning policy to tune parcel coalescing parameters based on a sampled counter of HPX runtime system. Our study shows APEX policy can converge on parcel coalescing parameters while reducing the network overhead and hence ensures reduction of execution time.

## KEYWORDS

Task-Based Runtime, APEX, HPX, auto-tuning, performance measurement and analysis.

## 1 EXTENDED ABSTRACT

Being on the verge of exascale era, researchers are trying to come up with novel programming model and runtimes to better utilize the HPC resources. Task-based programming model and runtime is one of the happening areas in the field High-Performance Computing(HPC). During the last decade, a good number of such programming model and runtimes are made available by the HPC research community. Some of the runtime systems provide asynchronous execution which is very different from tradition MPI programming model. This difference opens the door for new research to understand and look at application performance from a new perspective.

In this research, we worked on APEX [1] (Autonomic Performance Environment for Exascale), a performance measurement library for distributed, asynchronous tasking models/runtimes. i.e. HPX. It provides lightweight measurement (task < 1ms) and high concurrency. To support performance measurement in asynchronous task-based runtime systems, it depends on dependency chain rather call

stack. HPX is such an asynchronous task-based runtime system which builds on C++ runtime system based on the ParalleX model [3]. The HPX threading system employs lightweight tasks, known as HPX threads, that are scheduled on top of operating system threads. In a distributed environment, a locality in HPX is an abstraction for a physical node. The Active Global Address Space (AGAS) system in HPX provides a mechanism for addressing any HPX object globally.

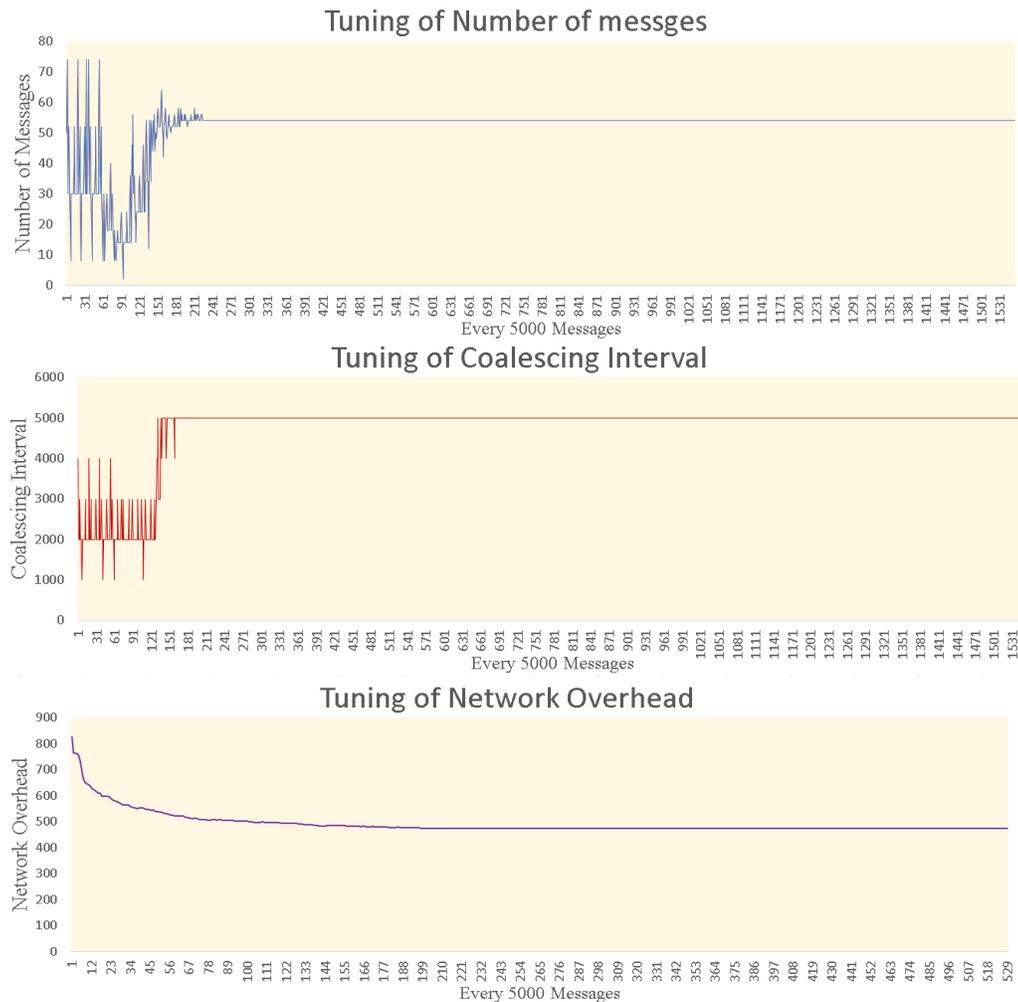
Finding an optimal value for a parameter that impacts application performance is a hard problem and often requires repetitive execution and hence incurs wastage of resources. In this research, we provide a preliminary study which demonstrates parameter value searching at runtime for better performance. We implement an adaptive auto-tuning policy to tune parcel coalescing parameters based on a sampled counter of HPX runtime system. Following previous research, where a direct correlation between network overhead and application execution time was shown provided a solid ground to develop an adaptive policy in APEX policy engine [2]. The main objective of this policy is to find out optimal value for two parameters: interval between the coalesced message sent and the number message to have coalesced. The APEX custom policy triggers the policy based on specified interval or based on the specified number of send events. When the policy is triggered, it invokes active harmony auto tuning library [4]. Active harmony observes the sampled network overhead counter and makes an adjustment to the values of the parameters. If the overhead is reduced in the next iteration, more change is made in that direction and if no improvement is found the parameters are changed in the opposite direction. For every policy triggering event, the callback function ensures the parameter values go closer to better performance. In this way, the policy search for parameter values for which the application demonstrates better performance based on the selected counter from the runtime system. Finding the counter that correlates with application performance and when to trigger the policy warrant research effort. While in this poster we demonstrate the counter correlation we still need to find when is the optimal time to trigger the policy. Figure 1. demonstrates the preliminary

results based on a toy application [2] that sends millions of messages between two nodes. We have conducted these experiments on Talapas, University of Oregon cluster (Compute node has dual E5-2690v4 processors). For this application, parcel coalescing policy has been triggered for every 5000 messages. First two graphs show how optimal parameter is searched. With Every change in the number of message and interval, the network overhead is sent to Active harmony to change the parameter values in a particular direction. By observing the overhead counter, it reaches to a point where no more improvement is possible. The third graph shows how network overhead reduced and reached to convergence and provide the proof of concept of this research.

We plan to continue this study by testing this policy for a couple of real applications in large scale. Moreover, we want to find out an adaptive approach to trigger this policy based on application characteristics

**REFERENCES**

- [1] K. A. Huck, A. Porterfield, N. Chaimov, H. Kaiser, A. D. Malony, T. Sterling, and R. Fowler. An autonomic performance environment for exascale. *Supercomputing frontiers and innovations*, 2(3), pp.49-66, 2015.
- [2] B. Wagle, S. Kellar, A. Serio and H. Kaiser. Active Message Coalescing in Task Based Runtime Systems. *International Workshop on Automatic Performance Tuning. IWAPT-2018.* (accepted)
- [3] H. Kaiser, M. Brodowicz, and T. Sterling, "Parallex an advanced parallel execution model for scaling-impaired applications," in *Proceedings of the 2009 International Conference on Parallel Processing Workshops, ICPPW '09.* Washington, DC, USA, 2009, pp. 394-401.
- [4] C. Tapus, I. Chung, and J. K. Hollingsworth. Active harmony: Towards automated performance tuning. In *2002 ACM/IEEE Conference on Supercomputing, SC '02*, pages 1-11, Los Alamitos, CA, USA, 2002.



**Figure 1: Tuning the number of message in a parcel and coalescing interval and its impact on network overhead**