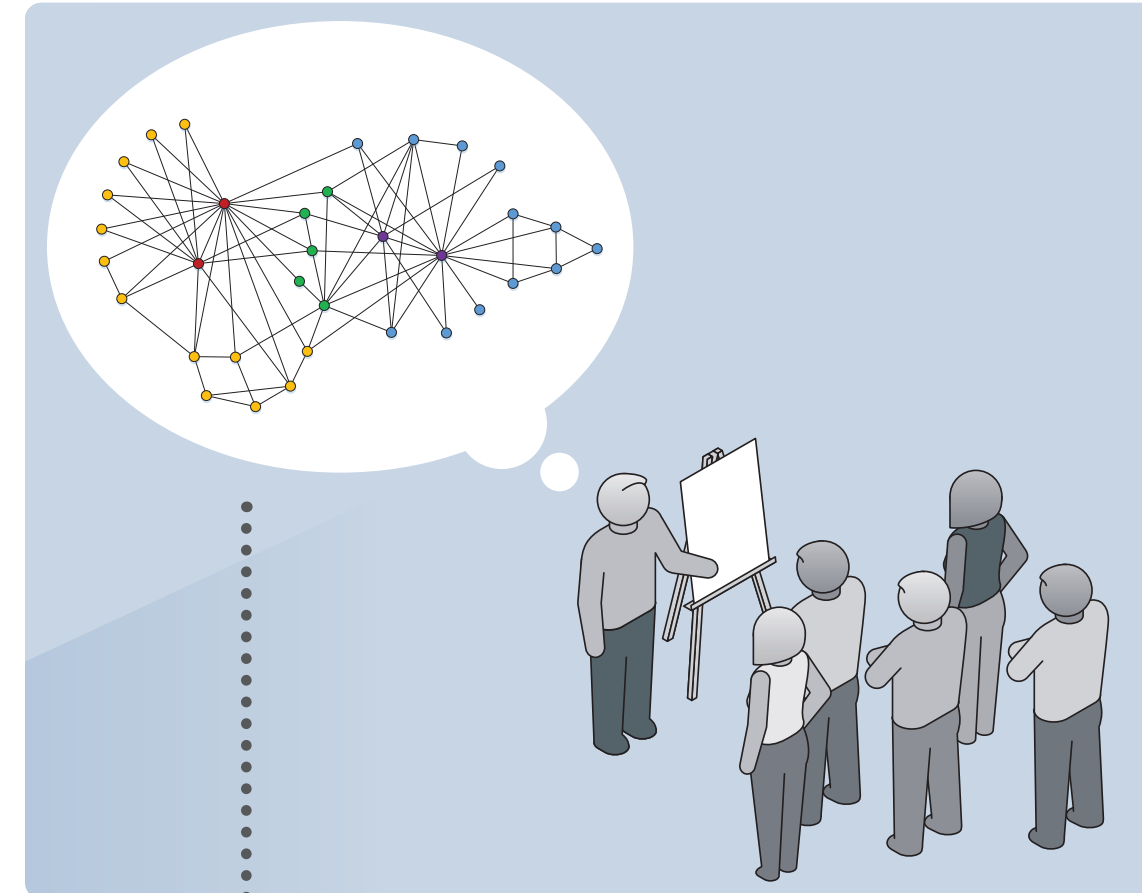


Automated Code Generation for High-Performance, Future-Compatible Graph Libraries


THE GRAPHBLAS APPLICATION PROGRAMMING INTERFACE (API) SPECIFICATION

released this year allows graph analytics experts and hardware experts to more easily come together to develop high performance graph algorithms. In the coming year we will be developing automated code generation tools to help hardware experts tune the graph primitives defined in the API for each new hardware architecture that is developed.

Algorithm Development



Hardware Architectures



GraphBLAS Application Programming Interface (API)

Goal: Write Once Run Everywhere

The GraphBLAS API Specification was released in May 2017. It defines a set of primitive operations that can be used to implement a wide variety of graph algorithms. It allows the separation of concerns where:

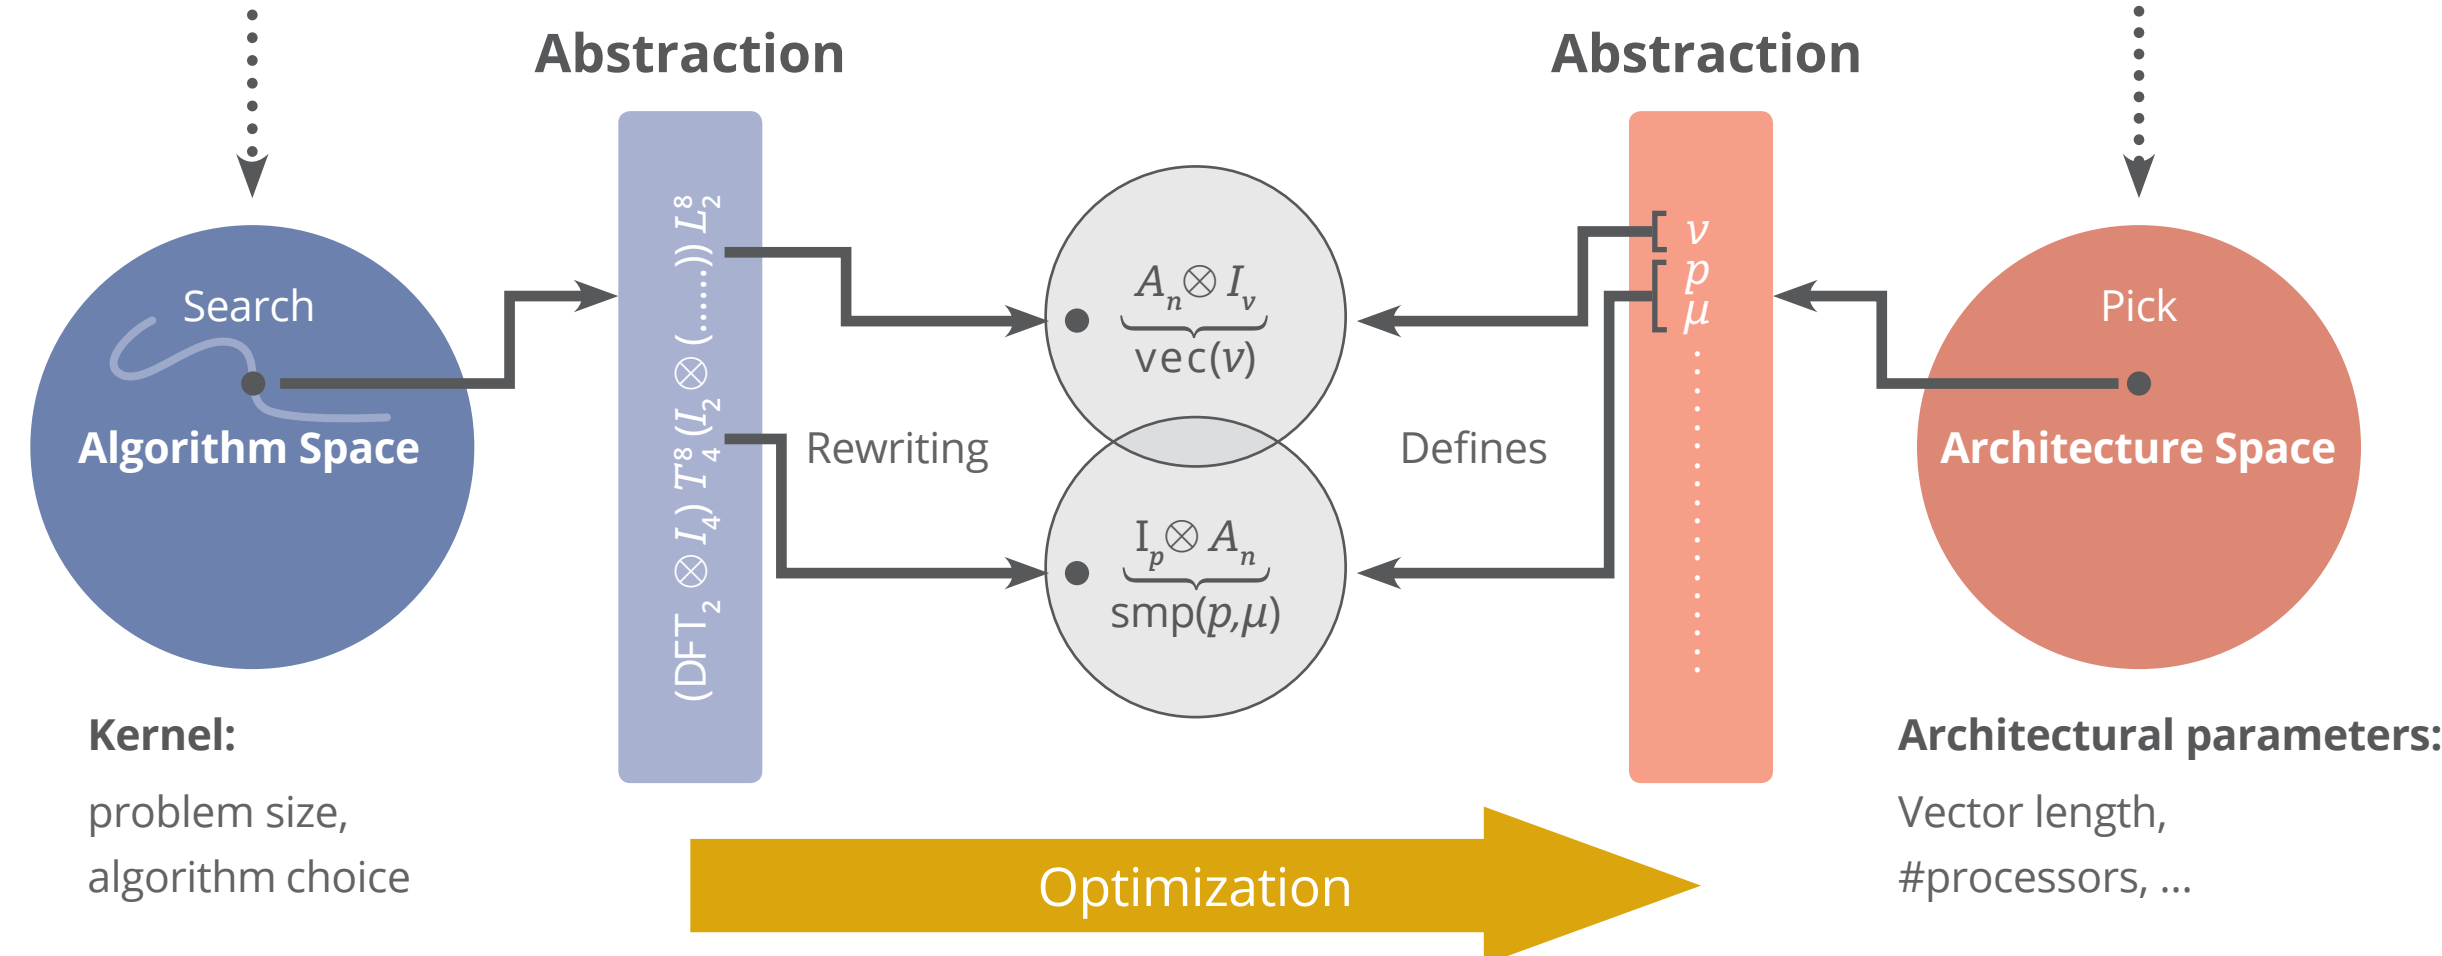
- Graph experts can more easily develop algorithms at a high level using the primitives defined by the API.
- Hardware experts can finely tune the primitives for their present and future hardware architectures

By separating the concerns between algorithm development and hardware architectures, the GraphBLAS API (graphblas.org) frees graph experts to develop algorithms while allowing hardware experts to fine tune primitives for existing and new hardware architectures.

Goal: Automated Code Generation for High-Performance

It still takes a large team of developers with knowledge of the complex hardware architectures to finely tune the GraphBLAS primitives. With the rapid development of new hardware architectures, the burden on these teams to develop high-performance implementations for each architecture is high. Using formal specifications of hardware capabilities, CMU's Spiral code generation technology can already automatically generate high performance signal processing codes. We are currently augmenting Spiral to support the GraphBLAS primitives so it can generate high performance code for graph algorithms.

GraphBLAS Math
 $C(L,z) = (L \oplus \cdot \otimes L^T)$
 $count = \oplus_{ij} C(i,j)$



Achieving high-performance from today's complex hardware architectures traditionally requires large teams of developers with hardware expertise. This year, we are building on automated code generation technology called Spiral (spiral.net) to use the mathematical formalization of the GraphBLAS algorithms to automatically generate the high-performance code for targeted hardware platforms.

By combining the mathematical descriptions of the graph primitives specified by GraphBLAS and the formal specifications of hardware used by Spiral, our technology will be able to automatically generate high-performance graph library code for today's hardware platforms as well as future architectures still being designed.

Copyright 2017 Carnegie Mellon University. All Rights Reserved.

This material is based upon work funded and supported by the Department of Defense under Contract No. FA8702-15-D-0002 with Carnegie Mellon University for the operation of the Software Engineering Institute, a federally funded research and development center.

The view, opinions, and/or findings contained in this material are those of the author(s) and should not be construed as an official Government position, policy, or decision, unless designated by other documentation.

NO WARRANTY. THIS CARNEGIE MELLON UNIVERSITY AND SOFTWARE ENGINEERING INSTITUTE MATERIAL IS FURNISHED ON AN "AS-IS" BASIS. CARNEGIE MELLON UNIVERSITY MAKES NO WARRANTIES OF ANY KIND, EITHER EXPRESSED OR IMPLIED, AS TO ANY MATTER INCLUDING, BUT NOT LIMITED TO, WARRANTY OF FITNESS FOR PURPOSE OR MERCHANTABILITY, EXCLUSIVITY, OR RESULTS OBTAINED FROM USE OF THE MATERIAL. CARNEGIE MELLON UNIVERSITY DOES NOT MAKE ANY WARRANTY OF ANY KIND WITH RESPECT TO FREEDOM FROM PATENT, TRADEMARK, OR COPYRIGHT INFRINGEMENT.

[DISTRIBUTION STATEMENT A] This material has been approved for public release and unlimited distribution. Please see Copyright notice for non-US Government use and distribution.

Internal use:* Permission to reproduce this material and to prepare derivative works from this material for internal use is granted, provided the copyright and "No Warranty" statements are included with all reproductions and derivative works.

External use:* This material may be reproduced in its entirety, without modification, and freely distributed in written or electronic form without requesting formal permission. Permission is required for any other external and/or commercial use. Requests for permission should be directed to the Software Engineering Institute at permission@sei.cmu.edu.

* These restrictions do not apply to U.S. government entities.

Carnegie Mellon® is registered in the U.S. Patent and Trademark Office by Carnegie Mellon University.

DM17-0743

Auto Code Generation for High-Performance Future-Compatible Graph Libraries