



A Hybrid Attributed Generic Graph Library Environment

Andrew Lumsdaine (PI), John Feo (PM), Giovanni Castellana (PNNL)

Architectures Thrust: Hierarchical Identify Verify Exploit (HIVE)

Improve analyst productive while maximizing performance on purpose-built systems

Break the owner-compute, superstep programming model

Shared memory, hybrid data structures; dynamic data/task scheduling; latency tolerant data movement

High-productivity, high-performance, SCALABILITY

Principal features

Extensible and generic SDK with three levels of graph primitives

Support for static, persistent, dynamic, and streaming data

Hybrid data model --- relational tables, property graphs, edge matrices

Extended Abstract Graph Machine to reason about data flow, data locality, and task scheduling

Control data flow intermediate representation for code transformation, optimization, and scheduling with introspection

Abstract runtime model supporting a variety of memory and execution models

ARTS


Efficient, a/synchronous resource aware task scheduling

Global address space to facilitate data movement

Dynamic dependency based synchronization

Light-weight multi-threading

- Active messages
- Asynchronous put/get/AMOs



Flexible distributed coherency protocols to trade-off data duplication with data movement

Introspection framework to support performance analysis and adaptive execution

HAGGLE architecture

HAGGLE is organized in layers.

SHAD, the Scalable High Performance Algorithms and Data structures Library, provides flexible, high-performance data structures and methods

The Extended Abstract Graph Machine (EAGM) applies code transformations to optimize data and task mappings


The Abstract Runtime System (ARTS) implements data movements, task scheduling and control operations for the unique hardware features of target systems

Layer	Abstraction Layer	Algorithms	Data Structures / Objects
Application	What move will I see		
Level 2	BFS, GC		Unstructured: Table
Level 1	For each		Structured: List, Array
Level 0	Set, get		Vertex, Edge, Data element
ARTS	Put, get, spawn		Interface mapping
Runtime system	Put, get, spawn		Tasks, locality, Blocks, distribution
Node-level OS	Spawn, send, receive		Threads, active messages, MAT, Message, global calls, SPI
HIVE Node	Scalable gather		Scalable gather, Local, store

SHAD

SHAD is a C++ library of data structures and algorithms that hides the complexity of programming a distributed system.

SHAD implements HAGGLE's Hybrid Data Model and supports dynamic insertion and removal of elements from the data structures.




HAGGLE HARDWARE ASKS

Fast context switching: latency tolerance for unpredictable data accesses

Fine-grained global address space: no partitioning, simplifies code development


Efficient synchronization: reduces hotspots with synchronization intensive workloads



EAGM

AGM model decomposes graph algorithms into processing functions and strict weak ordering of work items → algorithm taxonomy

EAGM model describes hybrid hierarchical algorithms using different orderings at different algorithmic levels (to match hardware performance)



This research was developed with funding from the Defense Advanced Research Projects Agency (DARPA).

The views, opinions and/or findings expressed are those of the author and should not be interpreted as representing the official views or policies of the Department of Defense or the U.S. Government.

Distribution Statement A - Approved for Public Release, Distribution Unlimited



THE ELECTRONICS RESURGENCE INITIATIVE