

Harnessing the Computing Continuum for Programming Our World

Pete Beckman* Jack Dongarra† Nicola Ferrier‡ Geoffrey Fox§ Terry Moore ¶
Dan Reed|| Micah Beck**

March 12, 2019

1 Introduction and Overview

The number of network-connected devices (sensors, actuators, instruments, computers, and data stores) now substantially exceeds the number of humans on this planet. This is a tipping point, and the societal and intellectual effects of this are not yet fully understood. Billions of things that sense, think, and act are now connected to a planet-spanning network of cloud and high-performance computing (HPC) centers that contain more computers than the entire Internet did just a few years ago. We are now critically dependent on this expanding network for our communications and social discourse; our food, health, and safety; our manufacturing, transportation, and logistics; and our creative and intellectual endeavors, including research and technical innovation. Despite our increasing dependence on this massive, interconnected system of systems in nearly every aspect of our social, political, economic, and cultural lives, we lack ways to analyze its emergent properties, specify its operating constraints, or coordinate its behavior.

Simply put, today we have the tools to instrument and embed intelligence in everything, and we are doing so at a prodigious pace. Although we are the globally distributed designers, builders, and users of this immense, multi-layered environment, we are not truly its masters. Each of us manages only some of the networks components, and we can neither predict its aggregate behavior nor easily specify our intensional goals in intuitive language. For all of us collectively, and each of us individually, this must change. Today, we program in the relatively small confines of a single node, defining individual device, instrument, and computing element behaviors, and we are regularly confounded by unanticipated outcomes and unexpected behavior that results once this individual node/device is exposed to the network collective. As consumers, we want our Internet-capable environmental devices (e.g., thermostats, lighting, and entertainment preferences) to adapt seamlessly to our changing roles and expectations, regardless of location. And yet, rather than specifying the ends we seek, we must specify detailed behaviors for home, office, car, and transient locale. In environmental health, we build and deploy arrays of wireless environmental sensors and edge devices when our goal may really be to “reprioritize edge resources to search for mosquitoes, given a statistically significant change in seasonal temperature and humidity across the nearby river basin”. In disaster planning, when satellites show hurricane formation, we manually redirect data streams and simulation software stacks, when our goal is really to “retarget advanced computing resources to predict storm surge levels along the eastern seaboard”. In science, when the Laser Interferometer Gravitational-Wave Observatory (LIGO) detects

* Argonne National Laboratory

† University of Tennessee and Oak Ridge National Laboratory

‡ Argonne National Laboratory

§ Indiana University

¶ University of Tennessee

|| University of Utah

**Dr. Beck is an Associate Professor at University of Tennessee, Knoxville. he is currently on detail to the National Science Foundation in the Office of Advanced Cyberinfrastructure. The work discussed herein was completed prior to his government service and does not reflect the views, conclusions, or opinions of the National Science Foundation or of the U.S. Government.

The Computing Continuum




Figure 1: The Computing Continuum: Cyberinfrastructure that spans every scale. Components vary from small, inexpensive devices with limited computer resources (IoT) to modest priced servers with mid-range resources to expensive high performance computers with extensive compute, storage and network capabilities. This range of capabilities, cost, and numbers forms a continuum.

a gravity wave, we scramble to reposition the global network of telescopes to capture multi-spectral data and launch simulations, when our true goal is to “identify and analyze correlated transient phenomena”.

Whatever the desires of consumers, companies, governments, and scientific researchers may be, we continue to build this increasingly digital world with only *ad hoc*, experiential, and intuitive expectations for the efficacy of alternative design choices. More perniciously, once these choices have been made, modifying or reversing them is often impossible. In large measure, this is because two constraints—resource capabilities and desired outcomes—are convolved, artificially and unnecessarily, on two time scales—design and deployment. The first of these is at the time of design and construction; the second is during outcome. At either time, the resource components may change (e.g., due to availability or failure) or expectations may shift (e.g., due to addition of new instruments or new questions). Moreover, the lifetime of many computations is not minutes or hours, but often days, months, or years. As Figure 1 shows, these components vary dramatically in capabilities and numbers but in aggregate define a complex collective that we call the “**computing continuum.**” While significant research and development exists at specific places along this continuum (i.e. focus on HPC, or cloud, or IoT), we seek to develop approaches that include the entire computing continuum as a collective whole. Just as early experimentalists who studied molecular behavior in isolation struggled for want of a predictive theory of gases and materials properties, so we struggle in the absence of a specification methodology and predictive understanding of this new computing continuum. We need a conceptual breakthrough for continuum programming that elevates specifications from components and behaviors to systems and objectives.

This paper outlines a vision for how best to harness the continuum of interconnected sensors, actuators, instruments, and computing systems, from small numbers of very large devices to large numbers of very small devices. Our hypothesis is that only via a continuum perspective can we intentionally specify desired continuum actions and effectively manage outcomes and systemic properties—adaptability and homeostasis, temporal constraints and deadlines—and elevate the discourse from device programming to intellectual goals and outcomes.

2 Research Philosophy

As the deployment of intelligent network-connected devices accelerates, so does the urgency with which this research area must be addressed. Industry analysts have begun predicting that “the edge will eat the cloud” [1]. The challenge, however, is not that one form will supplant another, but that we lack a programming and execution model that is inclusive and capable of harnessing the entire computing continuum to program our new intelligent world. Thus, development of a framework for harnessing the computing continuum would catalyze new consumer services, business processes, social services, and scientific discovery.

Table 1: Exemplar Continuum Computing Science Applications

Project	Description
Array of Things Urban Science Instrument	Environmental sensors, computer vision, deep learning inference, triggered weather, and traffic computations
Atmospheric Radiation Measurement Climate Research Facility	Software-controlled radar, user-provided sensors, data archive, climate models
Large Synoptic Survey Telescope	Transient phenomena detection and multispectral image correlation
National Ecological Observatory Network	Field-deployed instrumented towers and sensor arrays, sentinel measurements, specimen collection protocols, remote sensing capabilities, natural history archives
Precision Weather Forecasting and Sustainable Agriculture	Inexpensive environmental sensors and citizen science drive customized simulation models for microclimate weather forecasts and aquifer depletion reduction
Intelligent Manufacturing	Sensor measurement, modeling (digital twin), analysis, and control

We believe programming the continuum is not only possible, but realizable, though breakthroughs in the concepts and abstractions are needed for its coordination and management. Our common models of computation assume enumerated resources and predictable computing capabilities. However, in the continuum, the capabilities and numbers of components change dramatically over time, and aggregate functions can be long running—lasting months or years. In such contexts, intelligent, adaptive, closed-loop operations are *de rigueur*.

To be adopted, any broadly applicable continuum framework must be domain independent and exploit and interconnect extant, lower-level programming models and software systems. It is neither possible nor practical to obviate extant tools and techniques. Thus, the challenge lies in developing new annotations and composable abstractions for continuum computation and data movement, where none currently exist.

Table 1 highlights a range of scientific examples where the fabric of computing spans many orders of magnitude, and complex, open, and closed-loop behavior is required [2]. For example, ecologists should be able to link ecosystem monitoring with cloud-based simulations. Edge analysis routines, written to process telescope data, could be connected with HPC work flows that retrain machine learning algorithms automatically; the results of this machine learning could then be used to reposition observational assets based on the detection of new phenomenon. Similarly, materials scientists could write programs linking live instrument analysis at the edge with predictive models running on large-scale computing systems to detect experimental configuration errors.

Quite clearly, however, the merit of advances in continuum programming is not limited to science but will be a catalyst for exploration and advancement in almost every human endeavor.

3 A Goal-Oriented Approach to Programming the Computing Continuum

As Figure 1 shows, along this continuum the product of device count and device size is roughly constant. At either extreme, the scale is large, the resources are geographically distributed, their availability varies over time, and they frequently span multiple control domains. Thus, the computing continuum future should integrate two primary activities with multiple subsidiary goals: (1) developing a goal-oriented annotation and high-level programming model that specifies desired outcomes for the aggregate, rather than a collection, of component behaviors; (2) building mapping tools, a run-time system, and an execution model for managing continuum resources as an abstract machine that also monitors behavior and triggers re-mappings when necessary; and

In one sense, this approach maps loosely to a classic view of computation: a program, a run-time system, and an abstract machine. However, this traditional motif for computation evolved from our experiences directly specifying the actions of hardware under our control. Even the word “program” evokes the notion that we provide commands via a sequence of steps.

However, in the computing continuum, we cannot uniformly command all components to do our bidding—the building blocks are too diverse, the scale is too large, and the component owners and operators are sometimes unknown. Hence, we must begin by describing the goal.

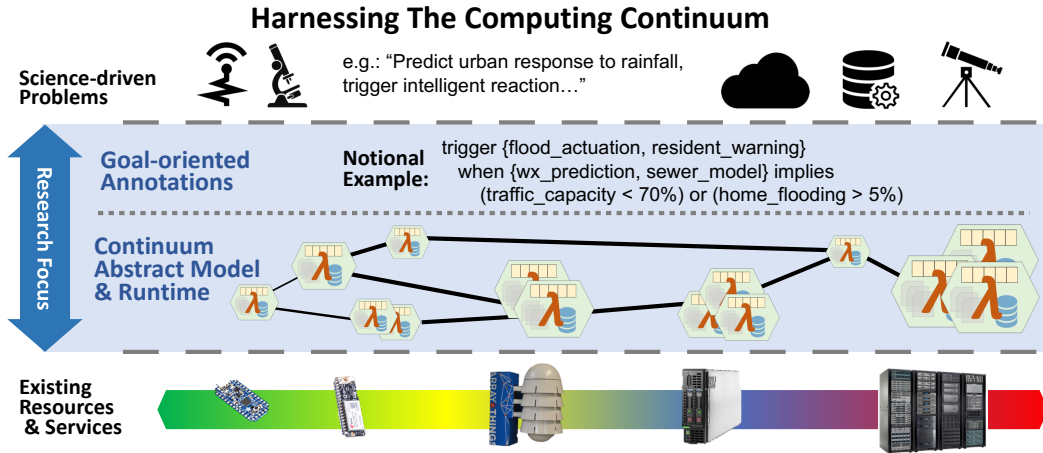


Figure 2: Continuum Computing Research Areas: A pictorial depiction of the computer science areas that require research to successfully program the computing continuum. To address problems, such as the sciences examples given in Table 1, using existing resources and services, we need an abstract programming model with goal oriented annotations, along with a run-time system and an execution model.

3.1 A Motivating Continuum Example

To highlight the breakthrough we envision, consider a realistic scenario from the south Chicago neighborhood of Chatham, which suffers from the highest level of home flooding and associated insurance claims in the city. Chicago’s infrastructure is a “combined system,” where storm water and raw sewage share the same underground structures, and when rainwater exceeds capacity, sewage enters homes, endangering health in addition to ruining personal property. Local residents are eager to see sensors, which we actually have deployed in Chatham, linked with weather models to provide warnings for impending danger.

The science/public policy problem can be succinctly stated as: “Eight hours before anticipated rainfall [...] predict underground infrastructure responses and trigger intelligent reactions and warnings if greater than 5% of homes flood or if traffic capacity falls below 70% of normal; then monitor and dynamically adapt urban controls to reduce harm.” Decomposing this example, we can sketch the components and control flow, as below.

- “Eight hours before anticipated rainfall” \implies On an HPC system, periodically run focused weather forecasts. As the risk of rainfall increases, run the forecast more frequently and with a finer resolution.
- “Predict underground infrastructure responses” \implies Three cloud/HPC models are coupled to predict the response of the storm infrastructure: (1) a regional floodplain model predicts waterway inputs to the neighborhood; (2) a computational hydrology model predicts the absorption of incoming water by soil and infrastructure; and (3) a model of the underground infrastructure predicts when capacity is exceeded. Inputs for the linked models come from other parts of the continuum: (a) live measurements from soil moisture monitors, sump pumps, basement water, and sewer levels provide immediate hyper-local data; and (b) Array of Things (AoT) nodes use intelligent libraries and edge computing resources to analyze camera images of rising flood water in rivers, on streets, and in rain gardens and detention ponds [3].
- “Trigger intelligent reactions” \implies Based on a computational sewer model and learned responses to historical actuator settings (e.g., the behavior of pumps and valves under stress) reconfigure local fog/edge components to autonomously respond resiliently, adjusting flow rates and motors within the disrupted network environment.
- “Trigger public warnings if greater than 5% of homes flood or if traffic capacity falls below 70% of normal” \implies Using cloud/HPC models, calculate the impact of the potential flooding and alert city officials and residents as needed. Furthermore, as new flood data become available from home and citizen science sensors, use image data from consumer cameras and sensors to push computer vision algorithms

along optimized intelligent libraries and new learning models to edge resources to identify and then predict when bridges, underpasses, or quickly flowing water will disable vehicles.

- “Monitor and dynamically adapt urban controls to reduce harm” \implies As the flood danger increases, re-prioritize available edge and fog resources to detect, predict, and report on events like rising waters, stalled vehicles, and water erupting from manholes and drains.

In this example, realizing end-to-end data capture, modeling and analysis, and timely response (closed loop) requires systemic coordination between sensors (edge), behavior detection (fog), models (cloud/HPC), and actuators (edge). Notice that it would be necessary to push new computational elements from the cloud to edge/fog resources during the storm—the continuum is bi-directional, and code is dynamic. Traditional computational work-flow systems like Pegasus [4] and Kepler [5] provide mechanisms for triggering data movement and computation. However, our goal is much broader, as we intend to *program* the continuum so that new algorithms and deep learning models can be pushed to appropriate locations (i.e., edge, fog, cloud, and/or HPC computing resources) using a simple lambda (function as a service) [6, 7] abstraction. Ideally, both the annotation and the high-level mapping would be portable to other contexts, spanning cities—Portland to Barcelona—and rural areas, even when the edge hardware and cloud providers differ significantly. For rural and agricultural areas, the challenges shift to dam spillway management, agricultural runoff and water supply quality assessment, and community evacuation.

Realizing and implementing this continuum programming model requires balancing conflicting constraints and translating the high-level specification into a form suitable for execution on a unifying abstract machine model. In turn, the abstract machine must implement the mapping of specification demands to end-to-end resources.

3.2 Goal-Oriented Annotations for Intensional Specification

This approach to programming the continuum stems from the fluid nature of the underlying resources, especially those at the edge. If every program component and its behavior were static, an imperative behavioral specification could be mapped directly to resources, assuming the resources themselves were static. For static cases, and when developers wish to implement all dynamic management, one could expose the resource demand graph specifications and control mechanisms for direct use. However, this is rarely possible or practical, as the following examples illustrate:

- When weather models predict roadways will degrade, adjust traffic signaling and preferred routing based on local conditions to optimize safety for both pedestrians and vehicles.
- When waterway sensors detect increased phosphates, use edge device sensor data and satellite image analysis as inputs to simulations that can predict harmful algae blooms.
- After significant seismic activity, re-prioritize edge computation to detect smoke, distress calls, and natural gas leaks. Based on air quality and local weather, predict location of the source(s).

Simply put, the programming equivalent to source routing, where the packet originator completely determines the route, is rarely possible in the continuum.

For dynamic cases—the majority—the community must devise an approach that is more declarative and constraint based. It must succinctly describe the *aim* and enable efficient mapping (and remapping) to disjoint, heterogeneous, shifting resources that behave more like independent agents than a single, cohesive machine. This is analogous to packet routing, which describes *where* the data should be sent but not how it should be sent.

Attacking this problem requires a two-pronged approach. First, one must devise languages for describing the resources of the continuum, including intelligent libraries, sensors, instruments, and cloud services. Tightly coupled with this, our abstract machine and runtime system will keep historical metrics of performance, interconnection bandwidth, and computational capacity that can be used for building execution graphs. Such work could build on the World Wide Web Consortium’s (W3C’s) specifications for the Semantic Web [8], including a Resource Description Framework (RDF), Web Ontology Language (OWL), and a query language for RDF (SPARQL). Although the Semantic Web has been slow to evolve, several well-developed

technologies for describing resources, data, and ontologies have been deployed.

This first prong is insufficient to fully harness the continuum, for one must also specify goals and desired outcomes. Thus, we must also link the exciting and intense resurgence in autonomous agent research, fueled by advances in machine learning, to build goal-based specifications that can be mapped to resources and computation. The foundations for this field were built decades ago [9–11]. Today, we see successful work in a wide range of fields—from goal-based, human-machine teaming to flocks of autonomous drones.

By fusing the research from these domains with our novel work on edge computing, intelligent libraries, data logistics, and HPC, we believe the community could revolutionize the computing continuum.

3.3 A Mapping and Run-time System for the Computing Continuum

Realizing the continuum programming model requires translation of high-level specifications into a form suitable for execution on underlying resources. Cleanly separating the intensional specification from the execution strategies is key to managing temporally varying application demands and shifting resource availability and capability, which is a defining element of the continuum. To accomplish this, we would need to translate goal-oriented application specifications into an annotated resource demand graph and a set of constraints. In turn, continuum resources will be represented by an annotated resource capabilities graph with its own set of constraints (an abstract machine). This abstract machine must instantiate specification demands on continuum resources. An intelligent run-time system is then responsible for mapping and adaptively remapping the resource demands to continuum capabilities. As Figure 2 illustrates, these elements define the programming model for the computing continuum.

With this backdrop, consider our motivating example once again: “Eight hours before anticipated rainfall, predict underground infrastructure responses and trigger intelligent reactions and public health warnings if highway traffic capacity is significantly impacted or if more than 5% of homes flood, then monitor and adapt controls during the storm.” This high-level specification results in a resource demand graph of (a) an HPC-based mesoscale weather simulation, (b) water level sensors, (c) roadway flooding detection via image analysis, and (d) public health models and warning and evacuation models. Constraints include computationally intense models, geographic sensor dispersion, and real-time adaptation and actuator control.

Similarly, the resource capabilities graph for the sensor set, edge devices, actuators, and cloud or HPC resources would include annotations that specify: (1) performance characteristics such as node interconnection bandwidth and connectivity, storage capacity, and computation speed; (2) programmability (i.e., fixed function device or “over the air” programmable); (3) multiplicity (i.e., an estimate of the number of instances, recognizing these vary over time); (4) control span (i.e., single or shared function and ownership); and (5) domain-specific constraints (e.g., geographic location, power limitations, or maximum usage frequency).

The mapping function would then instantiate environmental monitoring by tapping data streams from a statistical sample of the available water sensors, re-prioritize flood image analysis on AoT sensors and fog devices, and then launch a weather model parameterized by a terrain model with a real-time constraint on prediction cycles. Because these resource demands may conflict with or sometimes exceed resource availability, and the resources themselves may shift over time (e.g., due to sensor loss or replacement), any mapping is necessarily imprecise. Limited cloud or HPC availability might force a reduction in forecast accuracy and redeployment of alternative library versions to meet deadline constraints. Thus, the execution system must monitor the efficacy of each mapping and adapt accordingly.

It is also necessary to explore several techniques for mapping annotated resource demand graphs, mapping constraints to resource capability graphs, and learning and adaptively remapping these elements as demands and resources shift. As shown in Figure 3, these techniques include, but need not be limited to: (1) the intensional, goal-oriented program specification, translated to an annotated resource demand graph via autonomous agents and machine learning; (2) a resource registry that contains a time-varying list of available resources, attributes, and constraints; (3) optimized, multi-version libraries suitable for deployment on continuum resources, spanning computation-limited sensors, sophisticated edge devices, HPC systems, and cloud resources; (4) temporal “fuzzy logic” [12, 13] for qualitative constraint specification—an approach we used successfully in real-time adaptive control of parallel scientific applications [14, 15]; (5) machine




Figure 3: Continuum mapping and execution: Research is needed to explore techniques for mapping goal-oriented specifications to the available resources. This graphic shows a possible technique that translates the goal specification into an annotated resource demand graph that is used by the mapping engine, along with resource capability, availability and constraint information, multi-version libraries and behavioral history.

learning [16] and auto-tuning that exploit behavioral data and temporal fuzzy logic constraint specifications to map and remap resource demands to extant resources; and (6) very low overhead, distributed behavioral monitoring tools [17], a behavioral repository, and data sharing via Message Queuing Telemetry Transport (MQTT) [18, 19].

3.4 Building Blocks and Enabling Technologies

There are many potential building blocks and enabling technologies for building the continuum. Several of these have been developed by the authors.

The Array of Things (AoT). AoT [3] is an NSF-sponsored project at the University of Chicago, Northwestern University, and Argonne National Laboratory focused on supporting urban sciences. AoT is deploying 500 sensor nodes in Chicago, with more than 100 already operational. The core AoT hardware/software platform, Waggle [20, 21], is also being used by a diverse community of scientists for things like exploring the hydrology of a pristine prairie owned by the Nature Conservancy and for studying pollen and asthma in Chattanooga, Tennessee. These sophisticated sensors as well as their lightweight, battery-powered, thumb-sized cousins, support edge computation – the ability to directly analyze the data stream *in situ*. In addition to air quality sensors, the nodes support edge computing for in-situ computer vision, machine learning, and audio analysis. For example, the number of pedestrians using a crosswalk can be calculated by GPU-based computer vision algorithms using a combination of Kalman filters and deep learning. The computed results are pushed into the cloud.

Iowa Quantified (IQ). A complementary project, when Reed was at the University of Iowa, has been designing inexpensive battery and solar-powered wireless sensors with multiple wireless protocols (WiFi, LoRaWAN [22]) that are integrated with cloud and data analytics services. Configurable sensors span a wide range of environmental data, including temperature, humidity, gases, and particulate matter. This IQ project has deployed multi-depth soil sensors to enable understanding of moisture dynamics and agricultural productivity [23], and other projects are being planned around microscale weather forecasting and aquifer depletion minimization for precision agriculture.

Intelligent, Multi-Version Libraries. Across the continuum, the efficient application of scientific computing techniques requires specialized knowledge in numerical analysis, architecture, and programming models that many working researchers do not have the time, energy, or inclination to acquire. With good reason scientists expect their computing tools to serve them and not the other way around. Unfortunately, the highly interdisciplinary problems of programming the continuum, using more and more realistic simulations on increasingly complex computing platforms, will only exacerbate the problem. To address this, we believe a two-pronged strategy is needed that leveraged expertise in developing optimized, automatically tuned

libraries that can be deployed dynamically based on the goal-based annotation and mapping described above. At a low level, the community needs operations that closely mimic the building blocks of the Basic Linear Algebra Subroutines (BLAS) [24] but perform well across the continuum. Studying different approaches to optimization (e.g., mixed-precision algorithms and empirical auto-tuning for different platforms on the continuum) will enable us to provide the intelligent libraries needed to meet the performance constraints of the target application scenarios. At a high level, defining a common name space for different types of operations and services (e.g., “compress,” “merge,” and “SVD”) will enable us to move the locus of computation, in a seamless fashion, between the cloud and the edge. To integrate the bi-directional workflow across the variety of actors, the run-time system should match the resource requirements of the named components of a prototype library with the resources that are actually available when the work flow executes.

Data Flow Execution for Big Data. Currently, we are exploring the Twister2 environment [25, 26], which is an early prototype of some of the ideas expressed above. Twister2 has a modular implementation with separate components summarized in Table 2, which include the key features of existing “big data” programming environments and have been designed as a toolkit so that one can pick and choose capabilities from different sources. Further, all components have been redesigned as necessary to obtain high performance. Twister2 contains three separate communication packages: (1) one aimed at classic parallel programming, (2) Twister:Net [27] aimed at distributed execution, and (3) Publish-subscribe messaging as seen in Apache Storm and Heron to connect edge to cloud.

The first of these communication environments addresses problems in parallel programming that are well known from MPI, and implements Bulk Synchronous Processing. It extends MPI by adding some collectives required for solving big data problems including topic modelling and graph analytics [28–30]. However, the leading programming environments Spark, Storm, Heron, Flink that target big data offer as default, a different communication model built around data *flow*. Twister2 has implemented this as a separate communication system, Twister:Net, which offers a user friendly data API rather than a messaging API. Twister:Net automatically breaks data bundles into messages and, if necessary, uses disk storage for large volume transfers. A second and more critical feature is that the communication system supports distinct source and target tasks; therefore, Twister:Net can be used in a fully distributed environment, as seen in edge-cloud environments. Twister:Net also implements a rich set of collectives such as keyed reduction [31] where all the famous MPI and MapReduce operations (gather, scatter, shuffle, merge, join etc.) can be considered as particular versions of “keyed collectives”. Here we define a collective operation as a combined communication and computation (as in reduce) operation that involves some variant of all to all linkage of the source and target tasks. The final communication sub-system in Twister2, also supports different source and target tasks but uses publish-subscribe messaging. It is more suitable than Twister:Net for unreliable links as are found in edge to cloud communication.

Twister2 also has a carefully designed data flow execution model that supports linking intelligent nodes; this feature enables support for fault tolerance (i.e., as in the creation of Resilient Distributed Datasets [RDDs] used in Apache Spark) and allows wrapping nodes with rich tools such as learning systems. Further, Twister2 recognizes that some coarse-grain data flow nodes (e.g., those seen in job work flows) are not performance sensitive, but other finer-grain nodes need low overhead implementations with, in particular, use of streaming data without the overhead associated with reading and writing of intermediate files. Note currently if you use Apache Storm or Spark to link distributed sub-systems - data center to data center or edge to fog to data center – different jobs in each sub-system must be linked. Twister2 offers the possibility of invoking a single data flow across all parts of a distributed system.

4 Summary

Modern society is now critically dependent on a global and pervasive network of intelligent devices, both large and small, that are themselves connected to a planet-spanning network of cloud and HPC centers. Despite this dependence, analyzing this burgeoning network’s emergent properties and coordinating its integrated behavior is not straightforward. This paper outlines an ambitious plan to elevate programming,

Table 2: Twister2 components and status

Component	Area	Current Implementation	Future Implementation
Connected DataFlow	Internal fine-grain DataFlow or external DataFlow (workflow)	Dynamic Internal DataFlows	Ongoing and add Coarse Grain external DataFlow
High Level API's	Distributed Data Set, SQL, Python, Scala, Graph	Tsets (Twister2 implementation of RDD and Streamlets), Java	Dataflow optimizations, SQL, Python, Scala, Graph
Task System	Task migration	Not started	Streaming job task migrations
	Streaming	Streaming execution of task graph	FaaS Function as a Service
	Task Execution	Process, Threads	More executors
	Task Scheduling	Dynamic Scheduling, Static Scheduling; Pluggable Scheduling Algorithms	More algorithms
	Task Graph	Static Graph, Dynamic Graph Generation	Cyclic graphs for iteration as in Timely DataFlow
Communication	Dataflow Communication DFW	Twister:Net. MPI Based or TCP. Batch and Streaming Operations	Integrate to other big data systems, Integrate with RDMA
	BSP Communication	Conventional MPI, Harp with extra collectives	Native MPI Integration
Job Submission	Job Submission (Dynamic/Static) Resource Allocation	Plugins for Slurm, Mesos, Kubernetes, Aurora, Nomad	Yarn, Marathon
Data Access	Static (Batch) Data	File Systems including HDFS	NoSQL, SQL
	Streaming Data	Kafka Connector for Pub-Sub Communication, Storm API	RabbitMQ, ActiveMQ

coordination, and control of the continuum from *ad hoc* component assembly to intensional specification and intelligent, homeostatic control.

References

- [1] Maverick* Research: The Edge Will Eat the Cloud. Online, 2017. <https://www.gartner.com/doc/3806165/maverick-research-edge-eat-cloud>.
- [2] Stuart Anderson, Ewa Deelman, Manish Parashar, Valerio Pascucci, Donald Petravick, and Ellen M. Rathje. Report from the NSF Large Facilities Cyberinfrastructure Workshop. September 2017. <http://facilitiesci.org>.
- [3] *Array of Things*, March 2019. <http://www.ArrayOfThings.org>.
- [4] Ewa Deelman, Karan Vahi, Mats Rynge, Gideon Juve, Rajiv Mayani, and Rafael Ferreira da Silva. Pegasus in the cloud: Science automation through workflow technologies. *IEEE Internet Computing*, 20(1):70–76, 2016.
- [5] Ilkay Altintas, Chad Berkley, Efrat Jaeger, Matthew Jones, Bertram Ludascher, and Steve Mock. Kepler: an extensible system for design and execution of scientific workflows. In *Proc. of 16th Int’l Conference on Scientific and Statistical Database Management*, pages 423–424. IEEE, 2004.
- [6] *Apache OpenWhisk*, 2018 (accessed April 2018). <https://openwhisk.apache.org>.
- [7] *Amazon Web Services Lambda Features*, 2018 (accessed April 2018). <https://aws.amazon.com/lambda/features/>.
- [8] Tim Berners-Lee, James Hendler, and Ora Lassila. The semantic web. *Scientific American*, 284(5):34–43, May 2001.
- [9] J. M. Bradshaw, M. Johnson, A. Uszok, J. Dalton, A. Tate, S. Aitken, and R. Jeffers. Kaos policy management for semantic web services. *IEEE Intelligent Systems*, 19:32–41, 07 2004.
- [10] L. Kagal, T. Finin, and Anupam Joshi. A policy language for a pervasive computing environment. In *Proceedings POLICY 2003. IEEE 4th International Workshop on Policies for Distributed Systems and Networks*, pages 63–74, June 2003.
- [11] H. Hexmoor, C. Castelfranchi, and R. Falcone. *Agent Autonomy. Multiagent Systems, Artificial Societies, and Simulated Organizations*. Springer US, 2003.
- [12] L. A. Zadeh. Commonsense reasoning based on fuzzy logic. In *Proceedings of the 18th Conference on Winter Simulation, WSC ’86*, pages 445–447, New York, NY, USA, 1986. ACM.
- [13] Achille Frigeri, Liliana Pasquale, and Paola Spoletini. Fuzzy time in linear temporal logic. *ACM Trans. Comput. Logic*, 15(4):30:1–30:22, August 2014.
- [14] Kelvin K. Droegemeier, Dennis Gannon, Daniel Reed, Beth Plale, Jay Alameda, Tom Baltzer, Keith Brewster, Richard Clark, Ben Domenico, Sara Graves, Everette Joseph, Donald Murray, Rahul Ramachandran, Mohan Ramamurthy, Lavanya Ramakrishnan, John A. Rushing, Daniel Weber, Robert Wilhelmson, Anne Wilson, Ming Xue, and Sepideh Yalda. Service-oriented environments for dynamically interacting with mesoscale weather. *Computing in Science and Engg.*, 7(6):12–29, November 2005.
- [15] Lavanya Ramakrishnan and Daniel A. Reed. Performability modeling for scheduling and fault tolerance strategies for scientific workflows. In *Proceedings of the 17th International Symposium on High Performance Distributed Computing, HPDC ’08*, pages 23–34, New York, NY, USA, 2008. ACM.
- [16] Alexander Collins, Christian Fensch, and Hugh Leather. Masif: Machine learning guided auto-tuning of parallel skeletons. In *Proceedings of the 21st International Conference on Parallel Architectures and Compilation Techniques, PACT ’12*, pages 437–438, New York, NY, USA, 2012. ACM.

- [17] Todd Gamblin, Bronis R. de Supinski, Martin Schulz, Rob Fowler, and Daniel A. Reed. Clustering performance data efficiently at massive scales. In *Proceedings of the 24th ACM International Conference on Supercomputing*, ICS '10, pages 243–252, New York, NY, USA, 2010. ACM.
- [18] Andrew Banks, Ed Briggs, Ken Borgendale, and Rahul Gupta (eds). MQTT Version 5.0, OASIS Committee Specification 01. December 2017. <http://docs.oasisopen.org/mqtt/mqtt/v5.0/cs01/mqtt-v5.0-cs01.html>.
- [19] Pietro Manzoni, Enrique Hernández-Orallo, Carlos T. Calafate, and Juan-Carlos Cano. A proposal for a publish/subscribe, disruption tolerant content island for fog computing. In *Proceedings of the 3rd Workshop on Experiences with the Design and Implementation of Smart Objects*, SMARTOBJECTS '17, pages 47–52, New York, NY, USA, 2017. ACM.
- [20] *Waggle Project*, March 2019. <http://wa8.gl>.
- [21] Pete Beckman, Rajesh Sankaran, Charlie Catlett, Nicola Ferrier, Robert Jacob, and Michael Papka. Waggle: An open sensor platform for edge computing. In *2016 IEEE SENSORS*, pages 1–3. IEEE, 2016.
- [22] LoRaWAN Wireless Specification. 2018. <https://loro-alliance.org>.
- [23] Kang-Pyo Lee, Spencer J. Kuhl, Henry J. Bockholt, Benjamin P. Rogers, and Daniel A. Reed. A Cloud-based Scientific Gateway for Internet of Things Data Analytics. In *Practice and Experience in Advanced Research Computing*, to appear, 2018.
- [24] J. J. Dongarra, Jerney Du Cruz, Sven Hammarling, and I. S. Duff. Algorithm 679: A set of level 3 basic linear algebra subprograms: Model implementation and test programs. *ACM Trans. Math. Softw.*, 16(1):18–28, March 1990.
- [25] J.Ekanayake, H.Li, B.Zhang, T.Gunarathne, S.Bae, J.Qiu, and G.Fox. Twister: A runtime for iterative MapReduce. In *Proceedings of the First International Workshop on MapReduce and its Applications of ACM HPDC 2010 conference June 20-25, 2010*. ACM, 2010.
- [26] Kamburugamuve, Supun and Govindarajan, Kannan and Wickramasinghe, Pulasthi and Abeykoon, Vibhatha and Fox, Geoffrey. Twister2: Design of a big data toolkit. *Concurr. Comput.*, EXAMPI 2017 workshop at SC17 conference, 2019.
- [27] Supun Kamburugamuve, Pulasthi Wickramasinghe, Kannan Govindarajan, Ahmet Uyar, Gurhan Gunduz, Vibhatha Abeykoon, Geoffrey Fox. Twister:Net - communication library for big data processing in HPC and cloud environments. In *Proceedings of Cloud 2018 Conference*. IEEE.
- [28] Intel parallel universe issue 32 page 31: Judy qiu, Harp-DAAL for High-Performance big data computing. <https://software.intel.com/sites/default/files/parallel-universe-issue-32.pdf>, http://dsc.soic.indiana.edu/publications/Intel-Magazine-HarpDAAL_10.pdf. Accessed: 2018-9-30.
- [29] Harp-DAAL high-performance data analytics framework website. <https://github.com/DSC-SPIDAL/harp>, <https://dsc-spidal.github.io/harp/docs/harpdaal/harpdaal/>. Accessed: 2018-9-30.
- [30] Langshi Chen, Bo Peng, Bingjing Zhang, Tony Liu, Yiming Zou, Lei Jiang, Robert Henschel, Craig Stewart, Zhang Zhang, Emily McCallum, Zahniser Tom, Omer Jon, and Judy Qiu. Benchmarking Harp-DAAL: High performance hadoop on KNL clusters. In *IEEE Cloud 2017 Conference*. IEEE.
- [31] Geoffrey Fox. Big data overview for twister2 tutorial. In *BigDat2019 5th International Winter School on Big Data*.