

Enabling Wide Area Data Analytics with Collaborative Distributed Processing Pipelines (CDPPs)

Anja Feldmann
TU Berlin
Berlin

Email: anja@inet.tu-berlin.de

Manfred Hauswirth
Fraunhofer FOKUS
TU Berlin, Berlin

Email: manfred.hauswirth@tu-berlin.de

Volker Markl
TU Berlin
Berlin

Email: volker.markl@tu-berlin.de

I. MOTIVATION

Life without the Internet is no longer possible nor thinkable. Consider the effect a prolonged Internet outage: In the least impactful way, most of our kids and peers just would no longer be able to interact with their peers. They might severely miss out on the quality of their leisure time activities which increasingly rely on, e.g., social networks, online games, YouTube, and other online entertainment offers. This may be a nuisance but still tolerable. More seriously and economically relevant, manufacturing and trade would be no longer work as all interactions in and among companies rely on a working Internet. Indeed, “just-in-time” ordering mechanisms and “Internet of Things (IoT)”-enhanced production chains within the Industry 4.0 framework would no longer operational as old-style communication means such as phone and faxes have completely been replaced. Indeed, neither of these alternative mechanisms, faxes, phone, and also messaging, would not be available either as they also rely on Internet technology. Even worse, the control of critical infrastructures would also be affected severely as they increasingly rely on the Internet for gathering input data and propagating control information. Moreover, all big data analytic applications, including financial transactions, would fail as they can no longer gather and process their input data. Even worse, the fact that there is “no communication without energy” nowadays also means that the reciprocal statement applies that there is no “energy without communication.”

In 10 to 20 years, staggering amounts of data will be available and include proprietary as well as public sources. Moreover, analytic, and processing capabilities will have advanced further and will include novel intelligent machine learning mechanisms. In addition, we will take ubiquitous access to information from everywhere for granted. We will not be bound to specific devices, rather, the most appropriate device that is in our immediate vicinity will be chosen based on our preferences. This allows us to interact with the information when and wherever we want. Additionally, interactions will no longer require keyboards—they will be much more “natural” that today and will be based on speech, gestures,

emotion recognition, or even brain computer interfaces (BCI).

Thus, novel concepts to organize the distribution and processing of information over the Internet in a secure and safe way which is scalable are needed. An important building block for achieving this goal are mobile edge clouds. A mobile edge cloud is provides micro data centers which can move in the network according to load. This is another form of virtualisation which is transforming the the infrastructure everywhere, including network components, end-user devices, and eventually even sensors to transparently use any kind of resource. To work efficiently and scalably, mobile edge clouds need to be complemented by an appropriate world-wide network as well as sufficient backend computing power and storage. Combining all of this, we will see a seamless integration of network, storage, and computing in the future.

Services may then offer data availability “warranties” and possibly even an “insurance” against data loss and tampering. While putting an emphasis on reliability we also have to realize that problems and system failures cannot be avoided. Due to the scale of this infrastructure, individual components or even subsystems will fail. Thus, we must include scalable methods that enable us to debug this complex Internet infrastructure and its services. We have to be able to trace service problems, failures and possibly Byzantine behavior to the responsible system components. We also need means to determine which services are affected if a system component fails. With regards to data analysis we will need to answer questions such as “Based on which data elements was this information derived?”, “Did anyone tamper with this data?”, etc:

II. TRENDS

In the process of becoming such a hugely successful infrastructure the usage of the Internet and, thus, the structure of the Internet has undergone continues changes. Initially, the Internet was mainly about communication. Nowadays, communication, computation, and storage have become so intertwined that they cannot be separated and looked at individually anymore. Cloud providers, e.g., Amazon, offer new virtualized infrastructures providing virtually infinite computation and

storage resources. But this is not the end of the development. Through collaboration with Equinix, Level 3, and various IXPs, etc. they provide “on-demand” connectivity from the user to the cloud, since data in the cloud is of limited value if access to it is suboptimal. ISPs are also following this trend. They are in the process of deploying fog networks, i.e., micro-datacenters co-located with network aggregation points. This is an attractive way to scale services without the need to deploy service-specific hardware around the globe while offering good performance to the users. Thus, we observe two opposing trends: On the one hand companies tend to build isolated platforms to limit the impact of competitors and bind as many customers as possible. On the other hand ISPs, IXPs, and CDNs build infrastructures that can be combined to host flexible ecosystems on the same physical infrastructure.

We claim that future services are most likely realized within service specific *CloudNets*: virtual networks that combine clouds with networking and offer virtual resources: storage, computation, and communication. By decoupling from the physical constraints of the underlying substrate network, virtual networks and CloudNets—much like cloud resources—can grow, shrink, and/or be moved dynamically. CloudNets can be application specific and, thus, can rely on tailored networking protocols (beyond IP). Thus, CloudNets can overcome current limitations of IP and allow to introduce new functionality in the network core. Among the future challenges in this context are (i) algorithm for efficiently scale, shrink, and place CloudNets, (ii) what tailored protocol to use for what application, (iii) how to design and utilize mechanism for interactions between the infrastructure and the application.

In a few years staggering volumes of data will be continuously generated almost everywhere. Moreover, this data will grow exponentially. At the same time our analytic and processing capabilities will have further advanced and, e.g., offer intelligent machine learning mechanisms. Moreover, everyone wants to be able to have ubiquitous access to information from everywhere at any time. Hereby, the number of users and possible combinations of data and information is rapidly expanding. Thus, data streams will have to be processed and distributed in a coordinated manner in real-time. Thus, our current approach to data analytics: bring all data to a single central location and then use tools such as Apache Spark [[spark.website](#)], Apache Flink [[carbone2015apache](#), [Stratosphere14](#)], Hadoop [[hadoop.website](#), [white2012hadoop](#)], etc. will no longer scale. Rather, we will need a distributed processing platform where processing and data can move around freely and securely in an optimal fashion enabling fast reaction time and minimal resource consumption. In the process data provenance, quality criteria, and time constraints, both varying per customer, will have to be taken into account.

III. COLLABORATIVE DATA PROCESSING PIPELINES

This requires the integration of information processing and networking into a single paradigm. We envision that data will

flow along various *Collaborative Data Processing Pipelines (CDPP)*. CDPPs will provide the means to process data distributively, aggregate or sample it, add meta-data, and store it. CDPPs will allow joins and splits of multiple data sources as well as branches to multiple consumers. This will enable collaboration and sharing of the data. CDPPs will have to: (1) exercise “global” control with multiple distributed controllers; (2) scale under the constraints of the digital infrastructure and its computing and network resource limits, e.g., by trading-off local processing and data movement; (3) offer a holistic view of the information that is created incrementally at different data aggregation levels for different consumers. Among the future challenges in this context are both the control as well as data plane of such overlay networks which can be build and designed on top of CloudNets.

A. State of the art

Historically, big-data was invented during absolutism between the 16th and 19th century as public administrative data. Today, big-data is defined as “Information assets characterized by such a High Volume, Velocity, and Variety to require specific Technology and Analytical Methods” [[de2015big](#)]. Big-data started in 2004 when Google introduced MapReduce where data is stored in a *local* distributed file system and is processed according to “map” and “reduce” operators. In an industry context the term “Big” is also often seen as referring to the capabilities enabled by Big Data. “Big Data” is often seen as the driver/enabler of digitization.

Today, various so called “big data” frameworks exist. For an overview see, e.g., [[ranjan2014streaming](#)]. These frameworks range from (a) data application programming frameworks to large scale data mining frameworks, e.g., GraphLab, FlexGP, Apache Mahout, to distributed message queuing frameworks, e.g., Amazon Kinesis and Apache Kafka, to NoSQL databases, e.g. MongoDB, Cassandra and Amazon Dynamo. Examples of data application programming frameworks include Apache Hadoop, which consists of a distributed file system (HDFS [[shvachko2010hadoop](#)]) and the MapReduce [[dean2008mapreduce](#)] programming model, and Apache Storm and Flink for stream processing which support DAGs of computation over the data. These frameworks enable users to rapidly develop applications that can process massive amounts of data within a data center.

Among the major drawbacks of all of above systems is that they presume that the input data is centralized meaning that it is available in a single HFS within one cloud data center. Due to (a) the pure size of the data, (b) regulatory and privacy concerns which forbid centralized data processing [[mohan2012gupt](#)], and (c) organizational boundaries this will not be feasible in the future. Initial steps towards Wide-Area Big Data analysis have been proposed by Vulimiri et al. [[vulimiri2015wanalytics](#), [vulimiri2015global](#)] with Geode which builds upon Hive [[hive](#)] and minimizes bandwidth usage. Iridium [[pu2015low](#)] offers a system which focuses on minimizing response times for queries that happen across the WAN. Rabkin et al. [[rabkin2014aggregation](#)] with JetStream

which uses data cubes. JetStream incorporates storage into the processing system and, thus, facilitate aggregation. Moreover, it adds support for degradation in order to trade-off which query results are needed in real-time with which accuracy while meeting real-time deadlines.

At the same time social factors influence data quality[seysen09], e.g., the organizational model underlying the data collecting process or the specific sets of attitudes and behavior of the person initiating the data collection. Different social contexts develop so-called quality conventions of interaction[Diaz15], i.e., different social contexts differ regarding their beliefs, e.g., about the goals of data collection, adding to the difficulty of defining data meta-structures.

IV. CHALLENGES

The main challenge is to realize a distributed computing platform, where data and processing are optimally distributed to address the volume, variety, velocity, and veracity aspects of big data while minimizing resource consumption in an environment of wide-area distributed devices with heterogeneous capabilities. Additional challenges are the locality and time-scale spread, legal, privacy, and security constraints regarding data movement, highly volatile sets of data providers and consumers, limited and unpredictable WANs, as well as creating a timely and consistent centralized view across this decentrally organized large-scale infrastructure.

Thus, we have to develop scalable mechanisms to support collaborative data processing pipelines. Here, we need to tackle both: their **control** as well as their **data processing** planes. Among the challenges that have to be addressed by both **control planes** are the following:

- What are the right abstractions of the infrastructure for the CDPP control plane?
- What language do we use to describe the requirements and the infrastructure? How do we map between them?
- What functionality is hosted where inside the network?
- What service is mapped to which network function?
- How to specify what service guaranties are needed from the Infrastructure? How to map these to the needed resources? How to adapt the requirements over time and scale the service?
- How to automate CloudNet deployment?
- How to enable debuggability and traceability?
- How to avoid/detect misconfigurations?
- What is the price/value of infrastructure components?
- Can we do revenue sharing between the service, content, and the infrastructure provider?

We also have to address questions of the **CDPP data plane**:

- How to find data that enables us to answer a question?
- How to combine data in intelligent manner to derive information via novel machine learning mechanisms?
- How and where to sample data/information?
- Where to store which data/information?
- Which access control mechanism should be used?
- How to prevent tampering with the data/information?

- How to keep the data from being misused?
- How to trace the data processing pipeline?
- How to provide long term scalable secure backup?
- How to enable information sharing?
- How do we price information? How do we price data?
What are the resulting pricing models?

To realize ubiquitous access we have to tackle the questions of decentralized control. Due to the need for quick decisions not all decisions can be made centrally. Thus, we need control architectures which delegate some control decisions to decentralized controllers. This opens many tradeoffs and new interaction opportunities and threats. Indeed, we need this functionality across multiple operators and service providers. At the same time it needs to be invisible to the end-user.