

Emergency-Driven Assured Information Sharing in Secure Online Social Networks: A Position Paper

Bhavani Thuraisingham, Murat Kantarcioglu,
Latifur Khan
The University of Texas at Dallas

Barbara Carminati, Elena Ferrari, Leila Bahri
University of Insubria, Italy



Abstract—The United States and its Allied Forces have had tremendous success in combat operations. This includes combat in Germany, Japan and more recently in Iraq and Afghanistan. However not all of our stabilization and reconstruction operations (SARO) have been as successful. Several studies have been carried out on SARO by organizations such as the National Defense University and the Naval Post Graduate School. These studies have shown that security as well as power and jobs are key ingredients for success during SARO. One of the major conclusions is that we need to plan for SARO while we are planning for combat. These studies have also analyzed the various technologies that are needed for successfully carrying out SARO which include sensors, robotics and information management. As stated in the work by the Naval Postgraduate School, we need to determine the social, political and economic relationships between the local communities as well as determine who the important people are. One particular task during SARO is managing emergencies and crisis situations as well as carrying out humanitarian operations. In such situations, it is critical that the information between various individuals as well as agencies be shared. Furthermore, social relationships between the different parties need to be analyzed. To address the key technical challenges for emergency management and crisis response during SARO and related humanitarian efforts we are exploring assured information sharing and secure social networking to develop novel technologies to support emergency management and crisis response.

I. INTRODUCTION

Online Social Networks (OSN) are online applications that allow their users to connect by means of various link types. Since its inception in the mid 90s, social networks have provided a way for users to interact, reflecting of social networks or social relations among people, e.g., those who share interests and/or activities. Initially, critics regarded social media as a fad, a temporary fashion. But today, social media has demonstrated exponential growth, making it the most popular activity on the world wide web. As early as 2010, Twitter has recorded a 1,500 percent increase in the number of registered users, and over 70,000 applications have been developed according to the company. From 500,000 messages per quarter in 2007, the company has grown to 4 billion messages per quarter in 2010 (and many more billions in 2014). As the increase in popularity of social networking is on a constant rise, new uses for the technology are continuously being observed. Two major

technologies that have been developed for OSNs are (i) big data analytics technologies for analyzing these networks and extracting useful information such as location, demographics and sentiments of the participants of the network and (ii) security and privacy technologies that ensure the privacy of the participants of the network as well as provide controlled access to the information posted and exchanged by the participants.

Today, the importance of social media in emergency management is well-recognized. For instance, the Federal Emergency Management Agency (FEMA) wrote in its 2013 National Preparedness report that during and immediately following Hurricane Sandy, “users sent more than 20 million Sandy-related Twitter posts, or “tweets,” despite the loss of cell phone service during the peak of the storm.” Many different forms of social media can be utilized during emergencies such as the instant transmission of messages directed to followers, to raise the public’s situation awareness, to assist responders in gaining a more accurate picture of the situation, to help coordinate response efforts, to support rumor management and to correct misinformation, connect the community after an emergency to aid in the recovery effort, and provide a way for the community to express and offer support to its members.

While social media technologies have been extremely beneficial to emergency management applications, protecting the security of the social media data and the privacy of individuals is critical for these applications. For example, communication messages between the emergency responders as well as say the twitter users should not get into the hands of terrorists and criminals. Malicious individuals could divert the resources and exacerbate the problems faced by those trapped in hurricanes, earthquakes as well as other disasters. Furthermore, while location and other demographic information can be critical for crisis response, the users may want sensitive information about them to be protected. While security and privacy for social media data has received some attention, there is still a lot of work to be done to extract useful information as a result of social media analytics while at the same time ensuring the security of the data and the privacy of individuals.

To address the above challenges, we are examining security and privacy solutions for social media data as well as developing techniques for analyzing such data. We will

discuss our prior research in Section 2. Our approach will be discussed in Section 3. The paper is concluded in Section 4.

II. OUR PRIOR WORK

Our prior research in analyzing and securing social media as well as cloud-based assured information sharing has been reported in [1], [2], [5], [7], and [8]. These include extracting demographics data from social media, developing and enforcing security policies for social media, developing privacy-aware social analytic techniques, and developing a cloud for assured information sharing. Our research is building on our prior research and developing a scalable secure social media system for emergency management and crisis response.

A. Analyzing Social Networks

With respect to analyzing OSNs, our work has focused on mining for location and other attributes of individuals. At the forefront of emerging trends in social networking sites is the concept of "real time" and "location-based". So what makes location-based social media services so important? We briefly describe three situations where knowing the location of a user has a direct impact.

1) *Privacy and safety*: To change the default, adjust the Gowalla is a location-based social networking website which allows users to 'check-in' at places in their local vicinity, either through a dedicated mobile application or through the mobile website. As a reward, users will sometimes receive items from check-ins. This is achieved through the use of dedicated applications available on smartphones or by logging on to Gowalla's mobile website. Posting such an update publishes your current location to the user, and can prove to be an invitation to get your house robbed. In addition to this, knowing the location of the user makes it easier for spammers to attack the user in a more personalized manner.

2) *Trustworthiness*: Social media may be a source but is it a reliable one? Time Magazine in partnership with CNN discussed the impact of Twitter on the coverage of developments after the Iran elections in 2009. On June 12th, Iran held its presidential elections between incumbent Ahmadinejad and rival Mousavi. The result, a landslide for Ahmadinejad, led to violent riots across Iran, charges of voting fraud, and protests worldwide. Even as the government of that country was evidently restricting access to opposition websites and text-messaging, on Twitter, a separate uprising took place, as tweets marked with the hash tag #cnnfail began tearing into the cable-news network for devoting too few resources to the controversy in Iran. US State Department officials reached out to Twitter and asked them to delay a network upgrade that was scheduled for the night of June 15th. This was done to protect the interests of Iranians and assess the sentiment of the public about the Iran elections. In such cases, it was important for the US State Department to be able to trust the location of the user.

Another major area where trustworthiness of the user location is important is when companies use social media to obtain location-based opinions of products. Companies all over the world, from Burger King to Ford Fiesta, are using social media to target customers

3) *Advertising and marketing*: Here, we talk about two broad aspects on how location-based social media can affect companies in doing business in a better manner. OSNs connect people at low cost; this can be beneficial for entrepreneurs and small businesses looking to expand their contact bases. These networks often act as a customer relationship management tool for companies selling products and services. Companies can also use OSNs for advertising in the form of banners and text ads. Since businesses operate globally, OSNs can make it easier to keep in touch with contacts around the world. Having highlighted the importance of the location of the user in social media, it is important to understand that it is not provided explicitly by the users for a number of reasons. Some of the users are concerned about their privacy and security, and others do not find any incentive in sharing the location. The end result is the impeded growth of location-based services in the present world scenario. Social networking companies have shifted their focus from building relationships to identifying temporal and spatial patterns in messages. Twitter has very recently introduced a new feature whereby users can find location-based trending topics. But this is still not sufficient as firstly, it only covers the users who mention their locations explicitly and secondly, the topic of search is limited to the trending topics.

But to mine the location of the user is not an easy task in itself. In our work we have shown that traditional text-based location extraction techniques do not perform well in the domain of OSNs. The presence of multiple locations mentioned in the text makes it difficult to identify a single location which serves the purpose of the page focus for the messages. Additionally, we observe a lack of relationship between the location of the user and the location mentioned in the text. Secondly, the high running time associated with such mining techniques makes them unusable in practical scenarios. We have developed a system called TweetHood that outperforms the typical gazetteer-based approach in both accuracy and running time.

B. Securing Social Networks

The second part of our research is on the security and privacy of OSNs. As part of their offerings, many of the OSNs such as Facebook allow people to list details about themselves that are relevant to the nature of the network. For instance, Facebook is a general-use social network, so individual users list their favorite activities, books, and movies. Conversely, LinkedIn is a professional network; because of this, users specify details which are related to their professional life (i.e. reference letters, previous

employment, etc.) This personal information allows social network application providers a unique opportunity: direct use of this information could be useful to advertisers for direct marketing. However, in practice, privacy concerns can prevent these efforts. This conflict between desired use of data and individual privacy presents an opportunity for social network data mining – that is, the discovery of information and relationships from social network data. The privacy concerns of individuals in a social network can be classified into one of two categories: privacy after data release, and private information leakage. Privacy after data release has to do with the identification of specific individuals in a dataset subsequent to its release to the general public or to paying customers for specific usage. Perhaps the most illustrative example of this type of privacy breach (and the repercussions thereof) is the AOL search data scandal. In 2006, AOL released the search results from 650,000 users for research purposes. However, these results had a significant number of “vanity” searches – searches on an individual’s name, social security number, or address – that could then be tied back to a specific individual.

Private information leakage, conversely, is related to details about an individual that are not explicitly stated, but, rather, are inferred through other details released and/or relationships to individuals who may express that trait. A trivial example of this type of information leakage is a scenario where a user, say John, does not enter his political affiliation because of privacy concerns. However, it is publicly available that he is a member of the College Democrats. Using this publicly available information regarding a general group membership, it is easily guessable what John’s political affiliation is. We note that this is an issue both in live data (i.e. currently on the server) and in any released data. Our research has focused on the problem of private information leakage for individuals as a direct result of their actions as being part of an online social network. We model an attack scenario as follows: Suppose Facebook wishes to release data to Electronic Arts for their use in advertising games to interested people. However, once Electronic Arts has this data, they want to identify the political affiliation of users in their data for lobbying efforts. This would obviously be a privacy violation of hidden details. We explore how the online social network data could be used to predict some individual private trait that a user is not willing to disclose (e.g. political or religious affiliation) and explore the effect of possible data sanitization approaches on preventing such private information leakage, while allowing the recipient of the sanitized data to do inference on non-private traits. In order to protect privacy, we sanitize both details and link details. That is, we delete some information from a user’s profile and also remove links between friends. We then study the effect this has on combating possible inference attacks. Additionally, we present a modification of the Naïve Bayes classification algorithm that will use details about a node, as well as the node’s link structure, to predict private details.

We then compare the accuracy of this new learning method against the accuracy of the traditional Naïve Bayes classifier.

While mining OSNs and at the same time protecting the privacy of individuals is becoming critical for developing useful OSNs, controlling access to OSNs is also becoming a major concern. However, most of the current OSNs implement very basic access control systems, by simply making a user able to decide which personal information is accessible by other members by marking a given item as public, private, or accessible by their direct contacts. In order to give more flexibility, some OSNs enforce variants of these settings, but the principle is the same. For instance, besides the basic settings, Bebo, Facebook, and Multiply support the option “selected friends”; Last.fm the option “neighbors” (i.e., the set of users having musical preferences and tastes similar to mine); Facebook, Friendster and Orkut the option “friends of friends”; Xing the options “contacts of my contacts” (2nd degree contacts), and “3rd” and “4th degree contacts”. It is important to note that all these approaches have the advantage of being easy to implement, but they lack flexibility. In fact, the available protection settings do not allow users to easily specify their access control requirements, in that they are either too restrictive or too loose. Furthermore, existing solutions are platform-specific and they are hard to implement for various different OSNs.

To address some of these limitations, we have designed and developed an extensible, fine-grained OSN access control model based on semantic web technologies. Our main idea is to encode social network-related information by means of an ontology. In particular, we suggest modeling the following five important aspects of OSNs using semantic web ontologies: (1) users’ profiles, (2) relationships among users (e.g., Bob is Alice’s close friend), (3) resources (e.g., online photo albums), (4) relationships between users and resources (e.g., Bob is the owner of the photo album), (5) actions (e.g., post a message on someone’s wall). By constructing such an ontology, we model the Social Network Knowledge Base (SNKB). The main advantage for using an ontology for modeling OSN data is that relationships among many different social network concepts can be naturally represented using OWL (Web Ontology Language). Furthermore, by using reasoning, many inferences about such relationships could be done automatically. Our access control enforcement mechanism is then implemented by exploiting this knowledge. In particular, the idea is to define security policies as rules, whose antecedents state conditions on SNKB, and the consequents specify the authorized actions. In particular, we will encode the authorizations implied by security policies by means of an ontology, obtaining the Security Authorization Knowledge Base (SAKB). Thus, security policies have to be translated as rules whose antecedents and consequents are expressed on the ontology. To achieve this goal, we use the Semantic Web Rule Language (SWRL). As

a consequence, the access control policies can be enforced by simply querying the authorizations, that is, the SAKB. The query can be easily directly implemented by the ontology reasoner by means of instance checking operations, or can be performed by a SPARQL query, if the ontology is serialized in RDF. We have developed a model for fine-grained social network access control system using semantic web technologies. We also assume that a centralized reference monitor hosted by the social network manager will enforce the required policies. Since our approach depends on extensible ontologies, it could be easily adapted to various OSNs by modifying the ontologies in our SNKB. Furthermore, semantic web tools allow us to define more fine-grained access control policies than the ones provided by current OSNs. We have also designed and implemented an access control model for a semantic web-based social network system.

C. Cloud-based Assured Information Sharing

The cloud computing paradigm enables the sharing of large amounts of data securely and efficiently. Furthermore, the advent of cloud computing and the continuing movement toward software as a service (SaaS) paradigms have posed an increasing need for assured information sharing (AIS) as a service in the cloud. In order to satisfy the cloud-centric AIS needs of a coalition organization, there is a critical need to develop an AIS framework that operates in the cloud.

We have developed a secure cloud computing framework as well as multiple secure cloud query processing systems. Our framework uses Hadoop to store and retrieve large numbers of RDF triples by exploiting the cloud computing paradigm and we have developed a scheme to store RDF data in a Hadoop Distributed File System. We implemented XACML-based policy management and integrated it with our query processing strategies. For secure query processing with relational data, we utilized the HIVE framework. More recently, we have developed strategies for secure storage and query processing in a hybrid cloud. In particular, we have developed algorithms for query processing wherein a user's local computing capability is exploited alongside public cloud services to deliver an efficient and secure data management solution. We have also developed techniques for secure virtualization using the XEN hypervisor to host our cloud data managers as well as an RDF-based policy engine hosted on our cloud computing framework. We have also developed demonstration systems (The University of Texas at Dallas, University of Insubria, Italy and Kings College, University of London) for AIS in the cloud.

III. SYSTEM DESIGN

A. Objectives and Motivational Scenario

1) *Objectives:* Our main objective is to design and develop a system based on sound theoretical foundations for assured cloud-based information sharing for emergency management and crisis response in a social media setting. We will address questions such as how we can combine the data mining tools with graph theoretic analysis to extract

detailed information to enhance the design of OSNs? How can we extract accurate demographic information such as locations of the emergency responders and/or those who are trapped in emergency situations such as a tsunami or terrorism/war? What are the appropriate access control models for OSNs? How can the privacy-preserving data analytics algorithms be applied on OSNs?

2) *The Process:* We assume that users' profiles are validated. For this purpose, a building block of the entire process is a community-based identity validation process. Then, once an emergency occurs, the process has to build the emergency-related community. This has to be seen as the merge of portions of social networks (e.g., Facebook, g+, Twitter) selected based on the emergency location and other demographic information. For this purpose, work from UTD on location inference will be exploited and possibly extended. Responder selection will be performed over this community. The selection process will exploit validated users' profiles, reputation, and any other useful information mined from the social media. The selected responders will be temporarily authorized to access only personalized views of information needed for the emergency response. We plan to investigate this process according to two different approaches, as described in the following:

a) *Centralized Emergency Manager:* In this architecture we assume the presence of a trusted entity, the emergency manager, which is in charge of coordinating the community-based identity validation, listening for emergency situations, eventually building the corresponding community, and selecting the proper responders. Requests for help/collaboration will be sent directly by the emergency manager to responders, along with the corresponding authorized view of information. As such, the emergency manager enforces the emergency policies and retrieves the needed information by interacting directly with data owners.

b) *Decentralized enforcement:* As the name says, this approach does not rely on an emergency manager. In contrast, the emergency is detected directly by users, and requests for help/collaboration are propagated along the network by exploiting the social network graphs as well as the location-based networks (e.g., MANET network). We assume the presence of social apps devoted to handle this decentralized collaboration among users. Also, a repository of well-known emergency situations (i.e., emergency plans and corresponding emergency policies) should exist such that as soon as an emergency is detected, the request for collaboration propagates in the community together with the corresponding emergency policies. The policy enforcement is carried out on the users' side (e.g., mobile device) by matching the (validated) user profile with policy requirements. In case of authorization, users can retrieve the authorized view directly from the data owner. Both the architectures bring benefits and drawbacks, as well as, new research issues to be investigated. The centralized

architecture has the advantage of being easier to implement in that it relies on a central trusted entity in charge of managing the whole process. Still it brings some open research questions such as for instance, how community-based identity validation can be securely and efficiently performed? Which access control models should be used during emergency situations? Which policy-enforcing techniques should be applied (since responses are in real-time, it is a paramount requirement in many emergency situations)? However, one of the main drawbacks of the centralized solution is that the centralized emergency manager represents a single point of failure and can become a bottleneck of the whole architecture. Additionally, this architecture works under the assumption that the emergency manager is trusted; it knows all the user data and related emergency policies, which may bring serious concerns about the privacy of the involved data. This is overcome by a decentralized solution that does not rely on a trusted central entity and therefore gives network nodes more control over their data. But decentralization brings much more complexity. We need to solve issues related to decentralized data management, identity validation, responder selection, and the private and secure decentralized enforcement of emergency policies through collaboration among the network nodes. Efficiency and privacy/security issues related to all these processes are more complex than in a centralized setting because we cannot rely on a single emergency manager.

3) *Motivational Scenarios*: Our approach will be guided by motivational scenarios. Let us consider the scenario of a flood emergency. In particular, suppose that due to severe weather, the local river suddenly floods different zones of a target city. By using social media (e.g., tweet, photos, post) citizens bring to the public's attention the current emergency situations. As a consequence, social networks are fed emergency-related information, e.g., information on exact flood locations, news of interrupted roads, possible alternative directions, call for missing people, etc. Let us also assume that some houses of the target city have been flooded. The response plan implies to immediately evacuate the flooded house/apartments. This activity requires the help of several responders to be sought in the emergency-related community. In particular, there is the need for generic volunteers willing to offer their accommodations and/or equipment, as well as volunteers with given skills. An example of the latter are geologists, which could help in identifying flood-safe zones/houses in the target city. In this case, if responders with this skill are available in the emergency-related community, they have to be provided with detailed city cartography, so as to make them able to estimate the flood hazards.

B. Responders Selection

Our main focus will be on one of the most strategic one, that is, the responder selection. Note that this selection is not about responders from government offices whose participation is driven by laws and well-defined emergency plans. In contrast, we are interested in exploiting social media for the selection of responders in the community, that is normal persons with capabilities and/or skills to be of support to the emergency plans and who are reasonably close to the emergency location. Our approach will build on the work discussed in [2], [3], [6], and [9].

We believe that a key functionality that should be provided for an effective emergency management system through social media is not only detecting the responders but empowering them with the necessary information for emergency management. Indeed, during an emergency, it is not only important to be able to quickly retrieve the best persons able to respond to an emergency but also to give them the ability to access resources that are not allowed during normal situations and this should be done in an efficient and timely manner, while at the same time not sacrificing the security of the whole system. Such information sharing should be emergency-driven, personalized, controlled and temporary (that is only for the time required to respond to the emergency). Therefore, our goal is to design a system that is able to:

- Leverage social network connections/profile information and, more in general, social media data, to find the best possible responders for a given emergency. To do this, the system should consider a set of dimensions, such as the skills of the users, their proximity to the emergency location, their willingness to help, the size of the emergency, etc.
- Deliver information to the responders in a personalized, timely and secure manner. This will be obtained through the enforcement of emergency policies. In our view, emergency policies are used to represent the additional information requirements that arise during the management of a given emergency and therefore they basically state which user (i.e., emergency responder) can access which information and under which emergency situation.

Both responder identification and event detection will be based on the results of the research activity described in Section 3.1. In particular, responder selection will be aided by the output of two modules:

1) *Reputation Module*: We believe that a key component for responder selection is a reputation-based module that allows assessing the trustworthiness of a user with regard to emergency management. Just as when we buy a product online we look for feedback from other users, we can use the same rationale to help responder selections. For instance, if a person has received negative feedback on the management of previous similar emergencies, his/her reputation score will be lowered and this will be taken into account during responder

selection. Many reputation management techniques have been proposed for different application scenarios, such as for instance recommender systems or e-commerce which can be extended and tailored to the emergency management scenario.

2) *Identity validation module*: The system will also provide an innovative identity validation module. Indeed, identity management in online social networks is a challenging requirement and it is very crucial when social networks are used for emergency management. When registering on an OSN, almost all the existing commercialized solutions require users to create a profile that exposes personal information with a varying degree of detail. However, when creating a profile, most, if not all, of the provided information is not verifiable. A new joiner, for example, can claim to be a doctor, but the OSN has no one single way of validating the veracity of this information. On the other hand, relying on information which has not been validated may have very serious consequences for emergency management, but the OSN cannot decline access to a new user just because his/her information cannot be validated. The literature offers several proposals addressing issues related to identity breaches and/or identity related attacks on OSNs, but none of them provide a satisfactory solution at the issue of assessing users' reliability in terms of trustworthiness of their profile information. Therefore, we plan to build on the recent research results we have obtained on an identity validation process relying on OSN community feedback and active learning techniques to assign OSN users identity trustworthiness levels in a privacy-preserving and efficient way.

C. Secure Data Integration for Emergency Responders

We are expanding the framework discussed in Section 3.2 to enable integration of OSN-based data with traditional data sources to improve the response capabilities of the traditional government responders as well. For example, during an emergency, a first responder (e.g., firefighter) who is asked to respond to a fire may need automatic access to the location of any chemical plant near the incident site to determine potential problems and warn the general public about the incident. In such a scenario, the data coming from social network mining (e.g., tweets about red smoke seen by people near the incident location) could be integrated with the traditional data sources (e.g., state-wide database that contains the locations of the chemical plants with types of chemicals found in those locations). During such integration, the security of the information coming from sensitive resources needs to be protected while enhancing them with the OSN data. Furthermore, information may need to be shared among the first responders by leveraging OSNs. To enable this vision of securely integrating traditional first responders' data with OSN data, we need to address various challenges:

1) *Policy-based Location Aware Data Integration*: We plan to extend the policy framework discussed in Section 3.2 to include location-aware policies for the first responders. For example, the fireman responding to a fire in a location should automatically see the important information about the location such as the elderly people who live in the area, dangerous chemicals that can impact the fire, any tweets coming from the area that may be related to fire (e.g., tweets with certain keywords like smoke from location x). We plan to disclose such information based on the location and needs of the first responder. In some cases, the first responder may need to access previously unanticipated data that is not defined in the policy. To address such challenges, we plan to also explore integrating break-glass policies [3] with location-aware data integration tools.

2) *Rumor Detection and Response*: Although the reputation-based detection system may allow filtering of some false information and rumors, it may not be able to filter all the rumors. We may want to securely leverage the feedback of the first responders to further detect rumors and allow first responders to quickly make announcements to stop the spread of such rumors. To address this challenge, we plan to integrate the data mining algorithms discussed in Section 3.4 with active learning models so that the feedback of the first responders can be used to improve the data mining algorithms. To address this challenge, we need to capture the first responders' responses securely without overburdening them with time-consuming responses and explore various active learning algorithms for robustness against unreliable information.

3) *Information Sharing Among Responders: Secure Overlay Social Network for First Responders*: One of the most important needs during any emergency is to report to the public what is going on and share information among multiple different responders. For example, we may want to capture the pictures of the injured people and if possible, their identities while they are being sent to the hospital. Later on, such information may be shared with the families that are looking for their loved ones. In addition, first responders may want to make announcements and warnings to other first responders. Although, we do not anticipate any first responders to communicate with each other via twitter, the emergency response coordinators may want to document important information to be seen by all responders, but not the general public. Instead of building a separate OSN for first responders, we plan to leverage an existing OSN such as Facebook to securely share information among first responders.

D. Scalable Solutions Secure Social Media Systems for Emergency Management

The security and privacy policies discussed in the previous sections will be applied to the social network data

as well as to the data that is extracted from the social networks. Several efforts have been reported on extracting information from social network data. These include our work on extracting demographics from the social network data as well as sentiment analysis from such data. However, the scalability of the techniques is a major issue due to the volumes of data. Therefore we will apply big data management techniques to develop scalable solutions for social network integration and analytics. It is critical that the data extracted be accessed by the emergency responders. Therefore, a third component of our system is to develop scalable solutions for social media integration, extraction and security.

As a case study of OSN, we will focus on Twitter. Twitter allows its users to specify their geographical location as user information (Metadata). This location information is manually entered by the user or updated with a GPS-enabled device. The feature to update the user location with a GPS-enabled device has not been adopted by a significant number of users. Hence, this geographic location data for most users may be missing or incorrect. There are several drawbacks to relying on a user's manual update of location.

- Users may have incorrect geographic location data. For example, a Twitter user may enter his/her location as 'Krypton'. This may not be the name of a real geographic location.
- Users may not always have a city-level location. Users can input location names vaguely, such as the name of a state (E.g. Arizona) or the name of a country (E.g. U.S). These location names cannot be directly used in determining the city-level location of the user.
- Users may have multiple locations. If a user travels to different locations, he/she might mention more than one location in the metadata of his/her Twitter page. This makes it very difficult to determine their current, singular city-level geographic location.
- Users may have incomplete location data. A user may have specified an ambiguous name that may refer to different locations. For example, if a user specifies a location such as 'Washington', this name can be related to a state name or a city name (Washington D.C). These types of ambiguous names make it difficult to determine the exact user location.

Therefore, the reliability of such data for determining a city-level geographic location of a user is low. To overcome this problem of sparsely available geo-location information of users, we have evaluated the Twitter users' city-level geographic location based on tweet content and social graph. For content-based, we focus on tweet content along with the content of the related reply-tweet messages. We have used a probabilistic framework that considers a

distribution of terms used in the tweet messages of a certain conversation containing reply-tweet messages, initiated by the user [1].

With regard to graph-based, we predict the location of the user purely on the basis of his social network, using the strong theoretical framework of semi-supervised learning. In particular, we employ the label propagation algorithm. On the city locations returned by the algorithm, the system performs agglomerative clustering based on geospatial proximity and their individual scores to return cluster of locations with higher confidence. We perform extensive experiments to show the validity of our system in terms of both accuracy and running time [1].

Whatever approaches we have used as mentioned above, both suffer from scalability issues. In other words, when the number of users grows, the algorithm takes a significant amount of time. To speed up processing, we would like to utilize a distributed framework that runs on commodity hardware. In particular, we will exploit Hadoop MapReduce [7] or Spark [10]. Hadoop MapReduce is a software framework to store a large volume of data and process/retrieve relevant data in parallel on large clusters [6]. Those large clusters are built on a commodity hardware in a reliable fault-tolerant manner. Apache Spark [10] is a fast and general engine for large-scale data processing including stream data. It can run programs up to 100x faster than Hadoop MapReduce in memory or 10x faster on disk. For graph-based processing, we will exploit Giraph. Apache Giraph is an iterative graph processing system built for high scalability. For example, the Shortest path algorithm can be implemented efficiently using Giraph. In the future, we will implement our approaches in parallel to address scalability using Hadoop MapReduce, Spark and Giraph. We will measure the performance of various distributed frameworks and report timing and accuracy.

E. Cloud-based Assured Social Computing

To share the large amounts of social media data securely and efficiently, there clearly needs to be a seamless integration of the policy and the social media information managers in the cloud. Therefore in order to satisfy the cloud-centric assured information sharing needs of organization such as the DoD (Department of Defense) and the NSA (National Security Agency), we need (i) a policy manager running in the cloud that would enforce the information sharing policies across the social media and (ii) a social media data manager running in the cloud, for the storage and retrieval of data that would seamlessly integrate with the policy manager. To our knowledge, no such system currently exists and we plan to develop such a system.

Our plan is to first extend our proof-of-concept prototype of a system called CAISS (Cloud-centric Assured Information Sharing System) by incorporating the social

media data. The resulting system will be called CAIS-SM (Cloud-based Assured Information Sharing for Social Media). CAIS-SM will consist of two components: a cloud-centric policy manager that enforces policies developed as part of our research and a cloud-centric social media manager that will store and manage data. We will specify our policies in languages such as RDF and execute queries in SPARQL. This RDF data manager is essentially a SPARQL query engine where SPARQL is the language specified by the semantic web community to query RDF data. RDF is a semantic web language that is considerably more expressive than languages such as XACML for specifying and reasoning about policies. Furthermore, our policy manager and data manager will have seamless integration since they both manage RDF data. We have chosen this RDF-based approach for cloud-centric assured information sharing during Phase 1 because of the expressive power of RDF and the seamless integration of the RDF policy manager with the RDF query engine [4]. Having parts of the two critical components needed to build a useful cloud-centric assured information sharing system puts us in an excellent position to build the useful proof-of-concept demonstration system CAIS-SM. Specifically, we will enhance our RDF-based policy engine to operate on a cloud, extend our cloud-centric RDF data manager to integrate with the policy manager and build an integrated framework for CAIS-SM.

While CAIS-SM, will be the first of its kind for cloud-centric assured information sharing, it will operate only on a single trusted cloud and does not have the capability for information sharing across multiple clouds. Second, CAIS does not address issues such as scalability and flexibility. Our goal is to develop a fully functional and robust assured information sharing system called CAIS-SM++ which will incorporate our research on security and privacy policy management and data integration for social media together with the scalable solutions discussed in the previous sections. We believe that building CAIS-SM is an important step towards CAIS-SM++. This is because CAIS-SM can be used as a baseline framework against which CAIS-SM++ can be compared for several performance factors such as efficiency of storage model and the effectiveness of OWL-based policies.

IV. CONCLUSION

Our work directly addresses problems of critical importance not only for defense applications but also for healthcare and other emergency preparedness applications. SARO is becoming very important for us especially due to the winding down of the wars in Iraq and Afghanistan. In addition, new challenges are emerging in the former Soviet Union as well as in South East Asia (e.g., Myanmar). Managing crisis situations is critical for successful SARO as well as related humanitarian operations. Our theoretical

results and practical languages and tools will improve our ability to build and manage open, heterogeneous information systems that will move us toward a "need to share" philosophy in a social media setting while maintaining appropriate protections for security and privacy. The system that we will be developing will advance the promise of computational policies by grounding them in sharable ontologies, incorporating new models of usage control, and linking them with utility models.

ACKNOWLEDGMENT

Our prior research has been supported in part by grants from the AFOSR and EOARD.

REFERENCES

- [1] Satyen Abrol, Latifur Khan: Tweethood: Agglomerative Clustering on Fuzzy k-Closest Friends with Variable Depth for Location Mining. *SocialCom/PASSAT 2010*: 153-160
- [2] Leila Bahri, B. Carminati, E. Ferrari. Community-based Identity Validation on Online Social Networks. In *proceedings of The 34th International Conference on Distributed Computing Systems (ICDCS2014)*, Madrid, Spain, 2014, 21-30.
- [3] Achim D. Brucker and Helmut Petritsch. 2009. Extending access control models with break-glass. In *Proceedings of the 14th ACM symposium on Access control models and technologies (SACMAT '09)*. ACM, New York, NY, USA, 197-206.
- [4] Tyrone Cadenhead, Murat Kantarcioglu, Vaibhav Khadilkar, Bhavani M. Thuraisingham: Design and Implementation of a Cloud-Based Assured Information Sharing System. *MMM-ACNS 2012*.
- [5] Barbara Carminati, Elena Ferrari, Raymond Heatherly, Murat Kantarcioglu, Bhavani M. Thuraisingham: A semantic web based framework for social network access control. *SACMAT 2009*: 177-186
- [6] Jeffrey Dean and Sanjay Ghemawat, MapReduce: Simplified Data Processing on Large Clusters, in *Proc. of OSDI'04: Sixth Symposium on Operating System Design and Implementation*, San Francisco, CA, December, 2004.
- [7] Haque, A.; Parker, B.; Khan, L., "Labeling Instances in Evolving Data Streams with MapReduce," *Big Data (BigData Congress)*, 2013 IEEE International Congress on Big Data, July 2013.
- [8] Raymond Heatherly, Murat Kantarcioglu, Bhavani Thuraisingham, Preventing Private Information Inference Attacks on Social Networks, *IEEE Transactions on Knowledge and Data Engineering*, 2013.
- [9] Josang, Audun, Ismail, Roslan and Boyd, Colin A. (2007) A survey of trust and reputation systems for online service provision. *Decision Support Systems* 43(2):pp. 618-644.
- [10] Matei Zaharia, Mosharaf Chowdhury, Tathagata Das, Ankur Dave, Justin Ma, Murphy McCauley, Michael J. Franklin, Scott Shenker, and Ion Stoica. 2012. Resilient distributed datasets: a fault-tolerant abstraction for in-memory cluster computing. In *Proceedings of the 9th USENIX conference on Networked Systems Design and Implementation (NSDI'12)*. USENIX Association, Berkeley, CA.