A Social Content Delivery Network for Scientific Cooperation: Vision, Design, and Architecture

Kyle Chard*, Simon Caton†, Omer Rana‡, and Daniel S. Katz*
*Computation Institute, University of Chicago & Argonne National Laboratory, IL
{kyle, dsk}@ci.uchicago.edu
†Karlsruhe Service Research Institute, Karlsruhe Institute of Technology, Germany
simon.caton@kit.edu
‡School of Computer Science and Informatics, Cardiff University, UK
o.f.rana@cs.cardiff.ac.uk

***Abstract*—Data volumes have increased so significantly that we need to carefully consider how we interact with, share, and analyze data to avoid bottlenecks. In contexts such as eScience and scientific computing, a large emphasis is placed on collaboration, resulting in many well-known challenges in ensuring that data is in the right place at the right time and accessible by the right users. Yet these simple requirements create substantial challenges for the distribution, analysis, storage, and replication of potentially "large" datasets. Additional complexity is added through constraints such as budget, data locality, usage, and available local storage. In this paper, we propose a "socially driven" approach to address some of the challenges within (academic) research contexts by defining a Social Data Cloud and underpinning Content Delivery Network: a Social CDN (S-CDN). Our approach leverages digitally encoded social constructs via social network platforms that we use to represent (virtual) research communities. Ultimately, the S-CDN builds upon the intrinsic incentives of members of a given scientific community to address their data challenges collaboratively and in proven trusted settings. We define the design and architecture of a S-CDN and investigate its feasibility via a coauthorship case study as first steps to illustrate its usefulness.**

I. Introduction

Today, the methodologies for data delivery, distribution, and persistence are dramatically changing. Data are often no longer stored locally, but rather in data warehouses and storage clouds where storage capacity is limited only by a user's willingness to pay. However, the ease-of-use of network- or cloud-attached storage is compromised when data sizes run up against available network bandwidth, or must be accessed from large geographic distances for on-demand processing. This can be particularly challenging in contexts such as collaborative eScience and scientific computing, where large data sets also need to be readily available for analysis by international teams of researchers. To avoid research bottlenecks, data must be available in the right place at the right time with the right access permissions. Data must therefore be carefully distributed to support high performance, reliable, and trustworthy access that is optimized to meet the needs of collaborative research communities.

In other domains, such as the distribution of software or media, data distribution challenges are typically addressed through the architecting of sophisticated Content Delivery Networks (CDNs), or in simpler cases, server replication. However, when the data in question are (sensitive) research data, the access controls and policy management techniques that we know from Grid computing can obstruct or sometimes even block the use of such architectures. Even with approaches like GlobusOnline [1], it is still not straightforward to establish data distribution networks amongst teams of researchers.

In this paper, we propose a possible solution to these problems by leveraging social networks as a means to identify locations within a scientific community for trustworthy data storage, caching, data provenance management, access control, and accountability. Here, members of a scientific community contribute storage resources to act as nodes within a content delivery network for caching, temporary, as well as persistent storage. Each user allocates a portion of their hard disk or storage server (e.g., a directory or folder) that is used both as an interface to the content delivery network and for data storage by the content delivery system itself. The social network underlying the content delivery network provides the opportunity to implement unique distribution algorithms that take into account a node's trustworthiness and its "social proximity" to other users.

To capture a scientific community, the relationships that exist within its basic social fabrics, and the necessary notions of trust, we build upon the concept of a Social Cloud as a resource and service sharing framework that utilizes relationships established between members of a social network [2]. In this paper we propose the notion of a Social Content Delivery Network (S-CDN) for collaboration in scientific computing and eScience. We specifically focus on how the data needed to facilitate collaboration can be stored and accessed by collaborators within their communities. In other words, we do not yet focus on how these data are consumed and processed, but rather how a Social Data Cloud can be established within a given community of researchers using a CDN-like model.

This paper is structured as follows. In Section II, we outline the concept of a Social Cloud for Scientific Content Delivery and Sharing. In Section III, we discuss the basic premises of trust and participation needed for a S-CDN. In Section IV, we present medical image processing as a motivating use case to support our argument for leveraging social relationships as a means to support collaborative tasks potentially with BigData-like challenges. In Section V, we present our proposed architecture for a S-CDN and discuss high level performance indicators that permit an extensive evaluation of both the content delivery and social aspects of our approach, which act as anchors for future investigations and provide a basic premise for our initial case study. In Section VI, we present our initial case study using DBLP authorship networks to evaluate different replica placement algorithms, taking into account different trust thresholds and prior scientific collaborations in the form of publications. In Section VII, we discuss and evaluate existing work. Finally, in Section VIII, we summarize the paper and present our future work.

II. A Social Data Cloud for Scientific Computing

The concept of a Social Cloud [2], [3] builds upon two distinct strands of modern computer science. First is the ever increasing pervasiveness of social network platforms, such as Facebook and Twitter, that have profoundly changed how people communicate and interact. These systems allow us to interact within virtual platforms, establish virtual communities, and represent, document, and explore inter-personal relationships in a digital manner. Second is the basic premise of computing as a service encapsulated within the Cloud paradigm.

We defined the concept of a Social Cloud as a dynamic environment through which (new) Cloud-like provisioning scenarios can be established based upon the implicit levels of trust that transcend inter-personal relationships digitally encoded within social network platforms. The vision of a Social Cloud is motivated by the need of individuals, groups, or communities to access specific artifacts that they do not possess, but that can be made available by connected peers and peer groups. In essence, Social Clouds use social networks as mechanisms for collaboration and resource sharing, and social incentives to motivate sharing and non-malicious behavior, as users leverage their existing networks to share their (excess) capabilities and resources. This may initially seem unrealistic; however, if we consider that today an average Facebook user has 190 friends [4], and that each friend has at least one device to connect to Facebook, it is easy to conceive that large resource endowments exist within a given user's one-hop social network. We also know, from 30 years of research in volunteer computing, that such users are willing to contribute to "good" causes where they receive little to no personal benefit, and that their personal resource endowments are 60-95% idle [5]–[7]. Therefore, we argue that given social ties as a stimulus for collaborative and/or co-operative actions, social networks provide a more than adequate platform for accessing computational resources given appropriate middleware.

We believe a Social Data Cloud can facilitate collaboration within scientific computing because: 1) Issues pertaining to the trustworthiness of users as well as data provenance and accountability may embrace concepts of trust as an attribute of social relationships rather than the possession of a dehumanized signed certificate. The latter is arguably a reinterpretation of an underlying web of interpersonal relationships. 2) The concept of sharing and collaboration is inherent to the domain of eScience and scientific computing, and this facet of the domain is implicit within the constructs and context of a social network. This permits sharing via induced social capital – where the sum of even small contributions can result in a powerful resource infrastructure. 3) Reduced infrastructure costs, as basic infrastructure can be instantiated and provided by a community in small and easily manageable parts. (See [8], where we introduce the concept of a co-op infrastructure for the establishment of a Social Cloud.) 4) Data (i.e., replicas and caches) location can be optimized based upon their needed spatial and temporal requirements, which allows socially-tuned data aware scheduling. In other words, we can observe where data segments are most required and reliable, and then update data distribution in the Social Cloud.

III. Trust and Participation

In a Social Cloud, the notion of trust plays a critical role. Therefore in this section we briefly describe the premises of trust and its relevance to a S-CDN. Put simply, in the context of a Social Cloud trust entails two core aspects: 1) trust in the infrastructure via appropriate security and authentication mechanisms as well as policies, which we do not focus on here, and 2) inter-personal trust as an enabler of social collaboration.

We define the latter as: *a positive expectation or assumption on future outcomes that results from proven contextualized personal interaction-histories*[...] [9]. In other words, trust, when leveraged as a basis for the facilitation of exchange, cannot be observed or modeled as a simple binary entity. Trust is an intrinsic (subjective) attribute of an inter-personal social relation; it is the trust in the competence of an individual to be able to deliver a given resource or capability, i.e., a belief in the self-awareness and personal re-evaluation of self-efficacy undertaken by an individual as a precursor to partaking in a Social Cloud; and trust in an individual to deliver. Finally, the concept of proven trust relates to the occurrence of previous interactions with a positive outcome. In the context of scientific computing, this can be observed via publications or previous projects where elements of trust were crucial for a successful conclusion of the collaborative undertaking.

Upon an understanding of trust, we can leverage a social network as a mechanism for choosing suitable interaction partners, where interaction here relates to contribution and participation in a S-CDN. The Social Cloud concept is used as a basis to observe and establish trust relationships based on prior, and proven, interactions. Upon this basis, we can develop "trust models" validated through transactions over time to aid CDN algorithms with notions of reliability, availability, etc. in the management of the CDN overlay network. Through the observation of "successful" exchanges, we can also assume that future collaborative actions will follow. Although a simple and perhaps even naïve assumption, as project driven collaborations can dissipate when funding ends, we can tune classic

CDN algorithms with a social and historical perspective to help optimize the location of replicas, their contents, as well as the levels of redundancy they receive.

IV. Motivating Use Case: Medical Image Analysis

We present medical image analysis as an example of how a S-CDN could satisfy a complex array of constraints as well as approaching BigData-like challenges. We have intentionally selected a use case in a medical domain, as it is potentially one of the most challenging domains for collaborative science. This is due largely to the sensitive nature of data, i.e., a trusted setting is essential, and to data quantities that merit advanced techniques in data management and delivery. Various national standards already exist to protect the privacy of individuals when their health information (such as medical records) is shared. For instance, according to the Health Insurance Portability and Accountability Act (HIPAA) rules from the US Dept. of Health & Human Services, "... such rules require appropriate safeguards to protect the privacy of personal health information, and sets limits and conditions on the uses and disclosures that may be made of such information without patient authorization". In our approach, we assume that such privacy protection is already in place and the multi-center trials are being conducted between institutions that are already part of a "trusted" domain (i.e. that they all adhere to HIPAA rules and have been pre-approved for inclusion within the trial).

We therefore primarily focus on issues associated with data management and access (assuming all participants using or providing storage have agreements in place to be part of such a grouping). To motivate our approach, consider the data management options that currently exist in multi-center trials: 1) data are housed on a local system and transferred to other researchers when required; 2) data are stored in a central repository managed by a single institution; or 3) data are stored in a federated data store. In the first two cases, transfer mechanisms are often ad hoc, and it is not uncommon to observe external hard drives, DVDs, etc. sent via postal services, which makes it difficult for collaborators to share datasets or results quickly and securely. The third case represents an ongoing research challenge, to design and utilize efficient federated data stores, which is made even more difficult due to the heterogeneity of storage systems used in medical imaging.

Medical image processing represents both a computationally and data intensive application. A key example is magnetic resonance imagery (MRI). Research projects in this field tend to have experts from many domains, including medical practitioners, radiologists, and computer scientists. Therefore several collaborators and stakeholders exist, each of whom need to be able to obtain, analyze, and visualize datasets. Consequently many collections of raw and processed image-based datasets exist with various states of accessibility: from being available only locally to a single researcher or a co-located team of researchers, to a project repository or public repository, for example National Biomedical Image Archive (NBIA) or Picture Archiving and Communication Systems (PACS). Due to the scope of such projects, there is also a wide variety of image processing algorithms that need to be executed, for example to ensure data quality: collaborator sites undergo periodic site qualification tests to ensure the quality and anonymity of images. Once uploaded, images are processed with automated and manual analysis workflows. Examples from neurological studies include: brain extraction (removing the skull from the image), image registration (aligning the brain to a common atlas), region of interest annotation (to select particular regions to study), and Fractional Anisotropy (FA) calculations (determining the isotropy of diffusion). Such processes create multiple versions of a dataset, at potentially multiple sites, and with constraints on their accessibility due to their sensitive nature.

Continuing the example of neurological imaging, datasets from individual sessions are of varied size, depending on the number of imaging gradients, the image resolution and the type of MRI scan performed. As a guideline, such datasets are often in the region of 100s of MBs. However, for scientific analysis and clinical trials, overall datasets typically include many 10s or even 100s of subjects, each with multiple sessions. Moreover, at each stage in processing workflows, similar sized datasets are created. For example, a Diffusion Tensor Imaging (DTI) FA calculation workflow at the University of Chicago generates approximately 1.4 GB from a single raw session (of 100 MB). For a single study, it is therefore not unusual to have upward of 1 TB of data, and when considering multi-center trials, with multiple analysis workflows, the total data size can easily exceed 10s of TBs. With datasets of this magnitude and rigid constraints on access rights, it is crucial to have adequate mechanisms in place that facilitate scientific collaboration and have a specific focus on the (pre-)existence of trust.

The analysis algorithms used to process images differ greatly in computational requirements. In collaborative settings, the challenge that researchers face is twofold: first, the datasets must be accessible for processing when required by individuals, and second, the datasets must be moved into computation centers (which may not be owned by researchers interested in carrying out the analysis). We focus only on the first of these challenges and propose that S-CDN is a mechanism that allows better collaborative action by enabling processed datasets to be stored and shared with collaborators for further processing, analysis, etc.

We argue that S-CDN can act as a means to improve collaboration potential between researchers in the following ways: 1) The leverage of relationships encoded within a real world social/collaboration network provides a simple premise upon which to identify *trustworthiness* via previous scientific interactions or institutional affiliations. Although participants within such a social network would already have gained credentials (based on preserving privacy and confidentiality of data, for instance), it is still necessary to determine which center/participant is likely to possess suitable resources for carrying out any subsequent data analysis. 2) Building upon notions of interpersonal trust, an S-CDN can be used to ensure that data is available to those who are permitted to view it, and those who need it. We can also create mechanisms for redundancy without crossing the trusted boundaries within a given community. 3) Many approaches in eScience for patient related studies rely on arrays of (complex) middleware to deal with aspects such as authentication, the submission of jobs, data staging, etc. By employing a Social Middleware, we can aggregate much of this capability and reduce barriers to collaboration between participants. 4) S-CDN cannot replace the methods and processes used to determine when data is made available and in what manner. It can however, derive specific properties of the social graph as well as include new properties and constraints that can be used in access control, and data placement algorithms.

V. Social Content Delivery Networks: an Architecture

The typical usage of a CDN is to replicate data across many geographically distributed web servers that reside towards the edge of the network and therefore closer to end users. The major purpose of these architectures is to help web sites meet the demands of peak usage by improving scalability, availability and performance. Geographically distributed data caches have the effect of maximizing bandwidth while improving accessibility and latency. In a typical CDN configuration, when end users request content from a central server, the server redirects access to specific CDN nodes that serve web content to end users.

In our model, we use a CDN as a means of replicating scientific data at appropriate edge nodes (which in this case are researcher repositories) so that data storage is scalable, geographically distributed, and highly available. This in turn provides improved access to research data, making it faster to download, process, and share for researchers. The S-CDN is built upon a social networking fabric to represent the topology of a collaboration and to provide a trust overlay that ensures that data stays within the bounds of a particular project and on the nodes accessible by project members. This approach is similar to existing CDNs (e.g., Akamai, CDNs in Rackspace). In our approach, each user allocates a folder on their hard disk or storage server that can be used by a content distribution/management system for permanent and/or temporary data storage. We also require CDN folders to have associated properties of data integrity, uptime and availability, etc. Such QoS metrics can be used to select which participant is likely to be more trustworthy/reliable, and can be factored into replication/allocation strategies. Hence, the social network is used to chose participants who will be hosting various CDN folders.

Our vision of a S-CDN captures four core components: a *Social Network Platform*, *Allocation Servers*, *Individual Storage Repositories*, and a *Social Middleware*, which are illustrated in Fig. 1. To contribute resources to the CDN, researchers must join the Cloud through their social network platform and contribute one or more storage services (similar

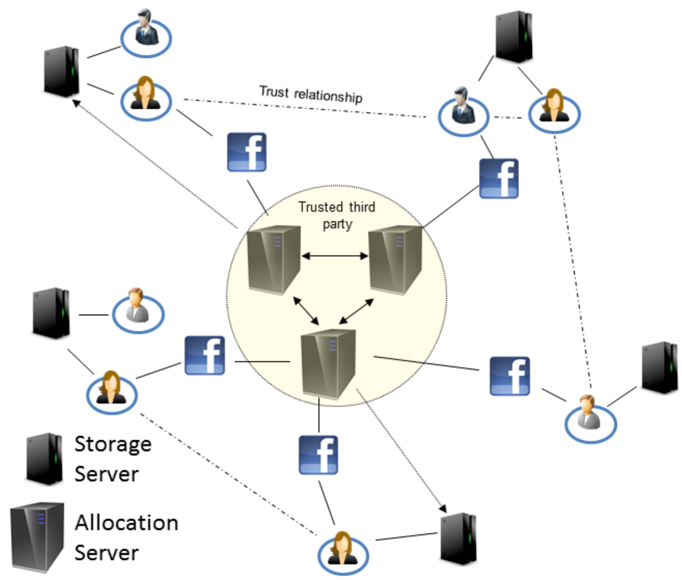

Fig. 1. Architecture of S-CDN

to the prototype presented in [2]). Here aspects such as authentication, social communication, and interaction within the scope of the project can take place. The *storage repository* acts as a CDN edge node on which research datasets (or fragments thereof) reside. One or more *allocation servers* act as catalogs for global datasets (for a particular Social Cloud); together they maintain a list of current replicas and place, move, update, and maintain replicas. The *social middleware* adds a layer of abstraction between users and the S-CDN (allocation server(s) and storage services) and provides authentication and authorization for the platform.

A. Storage Repositories

The S-CDN model is built on a collaborative pool of storage repositories hosted by participating researchers. Each storage repository contains some part of the shared dataset, which is made available to all other researchers in the Social Cloud. Following a DropBox-like model, storage is both accessed through and contributed to the CDN through a shared file structure on researchers' resources. This is similar to the use of CloudFiles in OpenStack/RackSpace [10], which makes use of the Akamai content delivery mechanism to distribute data content. A key difference between this approach and the proposed model is that the CDN is composed of user supplied servers, rather than use of servers offered by an infrastructure provider (as in CloudFiles). Depending on how such infrastructure is provisioned (i.e. it can include folders that are hosted on a server operated by an individual, or by an institution) one is likely to see a much lower availability within such an infrastructure compared to an Akamai-supported CDN (for instance). Availability in this context would also be influenced by the use of NATs and firewalls available at participating sites.

In addition to a storage repository, a CDN client and file transfer client for managing the repository and moving data respectively, are made available by the contributor. The CDN client is a lightweight server that is configured with the user's social network credentials to interact with the CDN. It also manages the contributed storage repository and monitors system statistics such as availability and performance. System and usage statistics are sent to allocation servers to identify the location and number of replicas needed. The client also acts as a proxy to the contributed repository to perform tasks such as initiating data transfers between replicas. The transfer client is responsible for moving data between users and replicas, and while this could be based on any transfer mechanism (e.g., FTP), we have designed the system to use GlobusTransfer [11], as it provides a high performance, secure, and reliable third-party transfer mechanism.

When a shared folder is first registered in the CDN, it is partitioned for transparent usage as a replica and also as general storage for the user. Data stored in the replica partition are accessible as a read-only volume by the user; they are therefore not able to be deleted as the volume is managed by the CDN. When users attempt to access data that are not currently in the replica partition, the client makes a call to an allocation server to discover the location of an available and suitable replica. The client then initiates a transfer between the chosen replica and the user's shared folder to retrieve the data. These data are kept under the user's control; they may, however, also be copied to the replica partition if so instructed by an allocation server.

B. Allocation Servers

Rather than relying on a completely decentralized Peer-to-Peer (P2P) architecture, we initially use a centralized group of allocation servers to manage the CDN, to enable more efficient discovery of replicas. The allocation servers may be hosted by trusted third parties – in the case of scientific collaborations, these may be national laboratories or universities. The allocation servers have three major tasks: 1) selection of replicas and data allocation, 2) data discovery and transfer management, and 3) general CDN management. Selection of replicas and allocation of data segments to replicas relies on information obtained from users' clients. A mapping between data sets and replicas is maintained by each allocation server, which is used to resolve requests. In terms of general management, allocation servers are responsible for ensuring availability by increasing the number of replicas needed (and selecting their locations) based on demand and migrating replicas when required. Access to allocation servers can only take place after users have been authenticated through their social network, as illustrated in Fig. 1.

C. Social Middleware

The social middleware leverages two key aspects of the social network. First, the S-CDN authenticates users and ensures data access is limited to those that should have access through the social network's authentication and authorization mechanisms. In other words, it uses the credentials of the social network platform, as well as the existence of social relationships and the notions of trust that they entail as a premise upon which to represent collaborative scientific projects. Second, the social network provides the network and key user properties (such as research interests or current location) that assist CDN algorithms in the selection and use of storage repositories – essentially creating the CDN overlay.

D. Replica Selection and Data Allocation Algorithms

Underpinning the performance and value of the S-CDN are the algorithms used for replica allocation. Various factors can be used to select appropriate replicas. We consider two broad categories in this work: traditional data location metrics (geographic location, availability, data usage patterns) and social metrics (relationships with other researchers, centrality and betweenness values derived from the social connectivity graph of the researcher). In designing appropriate selection algorithms we build on the research area of Distributed Online Social Networks (DSONs). DSONs, such as Diaspora, do not have a centralized management infrastructure; they are decentralized social networks hosted by the users that participate in the network.

There are two distinct stages of allocation that are used in S-CDN. First, the CDN aims to create a highly available storage fabric that selects the replicas accessible to members of the network while also utilizing replication algorithms to ensure that data are highly available. Second, data partitioning algorithms are used to assign data segments to replicas based on usage records and social information. The first aim can be accomplished through a combination of different aspects, for instance, by using socially based algorithms to determine appropriate base replica locations, for example determining important, well connected individuals, and combining geographic information. Novel availability graphs, as used in My3 [12], can then be used to select additional replicas required to create a highly available and high performance network. In the first case, graph theory metrics such as centrality, clustering coefficient, and node betweenness can be used to determine nodes that are important within a network. In the second case, a graph can be constructed that has edges between nodes if the availability of two nodes overlaps, and a "distance" weighting assigned to each edge that describes the transfer characteristics of the connection. When allocating replicas, we can then select a subset of nodes that cover the entire graph with the lowest-cost edges.

The second aim of optimizing dataset partitioning leverages traditional data partitioning models along with information inferred from the social connections between peers. Traditionally data partitioning, as is common in databases and P2P networks, partitions data across nodes with little regard for access patterns, and those that do, do so based only on individual users and data access patterns. In the S-CDN we aim to build upon this model to incorporate social information to group similar users based on their social connections, information obtained from the social network, and data access patterns.

E. Measuring the Success of a S-CDN

To measure the success of S-CDN we identify two types of metrics: 1) those that assess the quality of the CDN itself, and 2) those that assess social performance of collaboration through the construct of a S-CDN.

To measure the performance of a CDN the following metrics are typically observed: availability, scalability, reliability, redundancy, response time, stability. These metrics can enable observations on whether a CDN functions adequately. Through the specification of policies and goals that consider one or more of these metrics we can investigate the appropriateness of constructing a scientific CDN by leveraging social networks. These metrics can also aid in the CDN's core management algorithms to ensure that the quality of service provided by S-CDN is sufficient. These metrics can also act as placeholders for analysis on different levels of quality – for example whether the current level(s) of redundancy and replication are necessary or insufficient.

In order to measure the performance of the social and collaborative aspects of S-CDN, we propose the following metrics: *request acceptance rate (%)* – how often requests from the CDN's overlay management algorithms are accepted by storage participants; the *number of data exchanges* undertaken; *immediacy of allocation* – how fast (on average) are participants at accepting requests from the CDN; *ratio of successful to unsuccessful exchanges*; *ratio of freeriders or strategic players to producers/consumers* – this metric is self-evident, but it is an important and interesting aspect to observe within the scope of social collaboration *transaction volume* – network usage; *ratio of allocated to unallocated resources* – resource abundance; and *ratio of scare to abundant resource locations* – a metric to determine whether resource provisions are are well geographically distributed amongst participants.

By also observing metrics relating to the social and collaborative aspects of S-CDN, we can identify whether the "social" aspect in the CDN has a positive or negative influence on the performance of the CDN in general. The challenge with these metrics, however, is that they are difficult to observe outside of real implementations of S-CDN.

VI. Case Study: Replica Placement Based on Successful Science

To evaluate the feasibility of a S-CDN, we present a case study into replica placement using authorship networks from DBLP as a real world context for modeling scientific collaboration.

A. Evaluation Methodology

We construct our case study as an explorative investigation into the availability of data through replica placement within a Social Graph. In other words, we focus on the ability of a user to access data within their Social Cloud for (new) collaborative undertakings. We extrapolate collaborative research from the publication history of a scientist. That is, in this study we assume data access to be analogous to authorship. To support this approach, we extract the publications history of one author (Kyle Chard) from DBLP for the time span of 2009 – 2011 and explode his authorship network to a maximum social distance of 3 hops (coauthors of Kyle's coauthors' coauthors). Note that in our evaluation, we consider publications from the entire network, and not just from the graph seed.

Having extracted the authorship graph, we use the years 2009 and 2010 as a training set to identify locations for CDN replica placement using various algorithms. Here, a replica is a shared dataset for a given subgraph of the coauthorship network. It contains a complete repository of data needed for past as well as future scientific undertakings. A location relates to a scientist's locally available data resources, e.g., a NAS, SVN, Samba server, etc. Note that we don't yet model data segmentation, clustering, or redundancy; these are considerations for future work. Instead, we focus on how different replica placement algorithms affect the performance (in this case availability) of data within the Social Cloud using three subgraphs composed from different trust heuristics. Once a distribution of replicas has been assigned, we then use publications from 2011 of any author in the subgraph to determine how available datasets are, given new collaborations as well as cases of new collaborators.

To perform replica placement we use the following algorithms:

1) **Random:** Replicas are randomly assigned to nodes irrespective of any other factors.
2) **Node Degree:** Replicas are assigned to nodes with the highest degree (number of coauthors).
3) **Community Node Degree:** Replicas are assigned to a node within a community (direct neighbors) with the highest degree. That is, replicas are not placed as direct neighbors to one another.
4) **Clustering Coefficient:** Replicas are assigned to nodes with the highest clustering coefficient. Clustering coefficient is defined as the likelihood that nodes (b, c) that are directly connected to a single node (a) are also connected to one another (i.e. there is a coauthorship edge between b and c) .

In order to capture different notions of trust within the graph we prune the complete graph prior to replica placement using different heuristics. At this stage we only focus on 3 different trust graphs: 1) the initial subgraph with no trust threshold, 2) a subgraph composed of nodes and edges where authors have coauthored more than 1 publication together in the time period considered, and 3) a subgraph that includes only publications with fewer than 6 authors. The rationale behind this initial exploration is that we believe that multiple authorship between authors can be indicative of a closer working relationship and therefore a better predictor of future collaboration and conversely, publications with many coauthors are less useful for predicting collaborative relationships as there is not typically a strong link between all authors. The resulting subgraphs are summarized in TABLE I and depicted in Fig. 2. While the graph topologies are increasingly sparse with fewer nodes and edges, it is important to note that the maximum span is

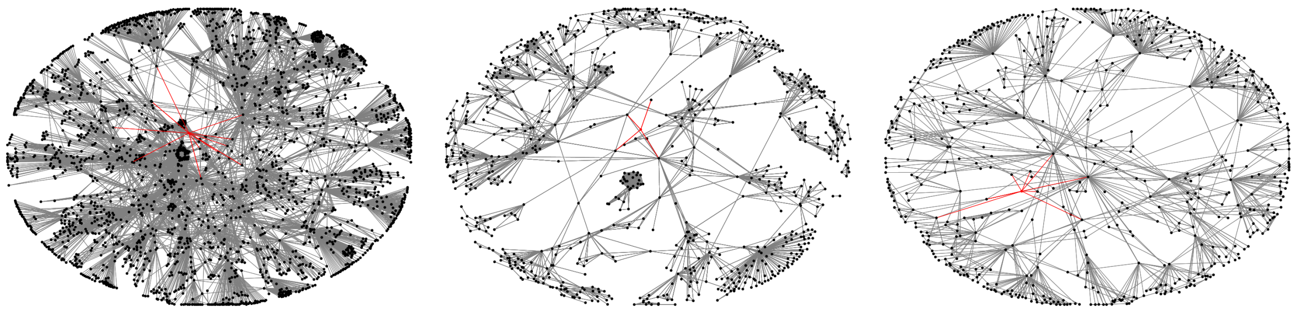

(a) Baseline Graph (b) Double Coauthorship (c) Number of Authors

Fig. 2. Subgraph topologies for the baseline, double coauthorship and number of authors graphs. Authors are depicted as nodes in the graph and edges represent coauthorship on one or more publications in the time period considered. The node marked in red is the initial seed node for the graph and the red edges indicate the first degree relationships from this seed node.

still 6 hops between nodes (authors). Another observation is that, unlike the other topologies, Fig. 2(b) includes isolated islands formed due to the pruning algorithm requiring at least two coauthorships between nodes, this can have a significant affect on the allocation of replicas as they will be disconnected from the rest of the network however it also serves to identify communities of trusted researchers.

Graph	Nodes	Publications	Edges
Baseline	2335	1163	17973
Double-Author	811	881	5123
Number of Authors	604	435	1988

TABLE I
THE NUMBER OF NODES AND EDGES IN EACH OF THE SUBGRAPHS.

B. Results and Discussion

Fig. 3 shows the percentage of replica “hits” based on 2011 publications coauthored by at least one author in the subgraphs presented above. In this analysis we define a “hit” as an author with a direct link to a replica (hop=1), and a “miss” as an author without a direct link to a replica. We report misses only when the author exists in the subgraph; misses for authors that are not in the subgraph are constant across algorithms and therefore do not affect the results except to reduce the overall hit ratio. Each of the experiments presented has been run 100 times to account for randomness.

Fig. 3(a), 3(b), and 3(c) show the hit percentage for each replica placement algorithm on the baseline, double-author, and number of authors subgraphs respectively. The graphs show an increase in overall hit rate for each subgraph. While it is not surprising that smaller graphs have a higher hit ratio, this is not the only factor in the improvement as each graph still has six hops between edge nodes. The increased hit rate in trusted networks is perhaps as a result of being a better indicator of future coauthorship.

In all cases using a community “elected” replica outperforms the other replica placement algorithms because it ensures replicas are distributed across the network rather than grouped together. This shows the advantage of using social principles to allocate replicas to avoid clustering of replicas too close together. In Fig. 3(c), the hit ratio of community election and node degree are similar. This is because the nodes with higher degree are more evenly spread across the network in this graph. In general, clustering coefficient is shown to be a bad metric for determining ”important” nodes in the network; in many cases the nodes with high clustering coefficient are those with few coauthors who are equally connected in a tight cluster. However, clustering coefficient can provide a good basis for determining trust in subgroups and could therefore provide a mechanism to establish groups of users with similar data access requirements.

Fig. 3(a) shows a near flat increase in hit rate for the node degree algorithm with more than two replicas. In fact, the rate increases by only 0.17% when 10 replicas are used. Investigation shows this is caused by a group of authors extracted from a single publication: [13]. This publication has 86 authors, which has the effect of creating an artificially high node degree for many of these edge authors, the result of which, when allocating replicas, is that the subsequent replicas added are also authors in this cluster, which only minimally increases the hit rate as these nodes are already, at most, one hop away from a replica. This is an interesting observation, as in some domains, it is not uncommon to have large coauthorship numbers. However, this also serves to support our assumption that publications with a large number of authors is not indicative of a close relationship between authors.

C. Relationship to Use Case

In the medical image analysis use case presented in Section IV, we assume that a set of researchers embark on a multi center trial led by a small number of lead institutions with a larger number of collaborative institutions. The lead institution is responsible for assembling a group of allocation servers and selecting, through a social network, the set of users and groups that are included in the collaboration. In our architecture we focus on large scale social networks, such as Facebook – however, it is also possible to use community specific tools (such as myExperiment) to establish the social network. Researchers in the collaboration join the S-CDN by configuring their CDN client and storage repository with their personal social network credentials and information about the

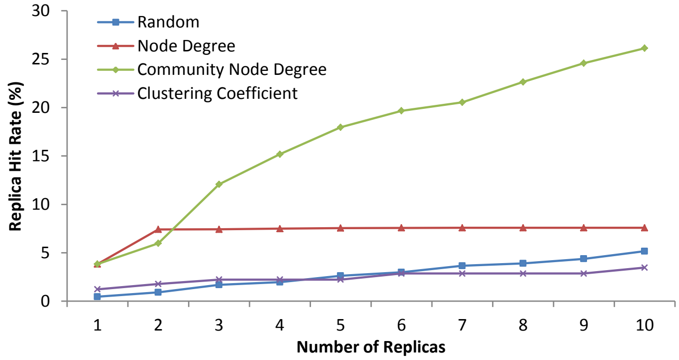

(a) Baseline Graph

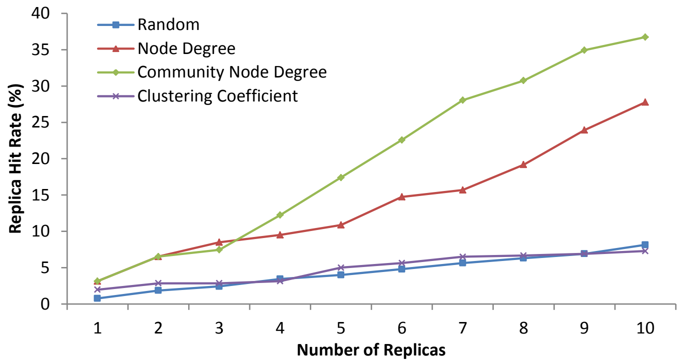

(b) Double Coauthorship

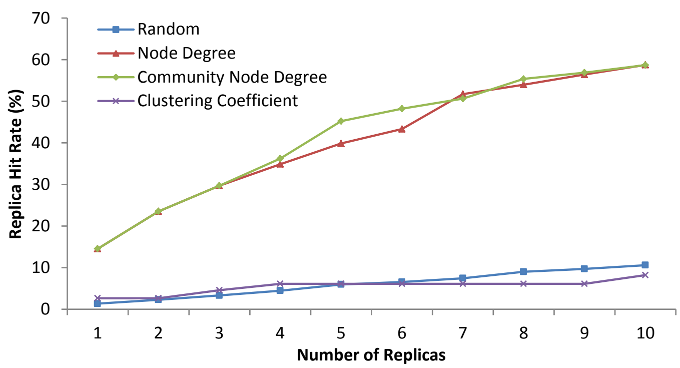

(c) Number of Authors

Fig. 3. Replica hit rate for each subgraph and algorithm combination.

lead institution's allocation servers.

The allocation servers create a CDN overlay over the registered storage repositories. The first step is forming a network across the registered users, this is done by combining the social network topology and a publication network extracted from the pubmed repository – the National Institutes of Health's publication repository. After assembling the social network, the allocation servers will attempt to create trusted networks using the principles outlined above, for example by disregarding publications with many authors and increasing the importance of multiple coauthorships between researchers. The trusted subgraphs are then parsed to identify groups of users with similar data usage requirements, for instance by using clustering coefficient to determine tightly connected groups and perhaps topic modeling to extract areas of interest. Finally replicas are selected based on metrics such as node degree (to identify a particular community of users). Over the lifetime of the collaboration, allocation servers will also include aspects such as data access profiling and node availability to customize the replicas appropriately.

VII. Related Work

There are many examples of commercial CDNs available, such as Akamai, Amazon Cloud Front, and Limelight. Commercial CDN providers have thousands of dedicated resources geographically distributed to provide high performance global access to (web) data. However, usage of these services is typically expensive and if they were used to store research-scale data the cost would be prohibitive. Open source CDN platforms, such as Coral CDN, provide an alternative means of creating personal P2P-based CDNs, though without significant investment in infrastructure, it is difficult to construct a scalable distributed network.

MetaCDN [14], [15] was proposed as a "poor man's CDN" capable of creating a CDN over different Storage Cloud providers (e.g., Amazon S3). In this work the use of existing storage clouds is shown to provide a low cost alternative that is able to deliver similar functionality to commercial options. MetaCDN provides an overlay network across this basic Internet-enabled file storage to provide replication, failover, geographical load redirection, and load balancing. In a somewhat tangential approach, Content-as-a-service (CoDaaS) [16] is a Cloud-based model designed to host and deliver user generated content through virtual content delivery services. CoDaaS focuses on user generated content, that is, content created by non-professional producers, which like research data has unique hosting requirements that current CDNs are unable to provide. In particular, aspects such as the data's long-tail nature (most content is of low interest to most users) is somewhat orthogonal to the goals of traditional CDN providers that aim to serve high profile content. CoDaaS constructs a dynamic content distribution tree, aiming to be cost- and QoS-effective, over a media cloud.

Distributed online social networks (DOSN) have been proposed as a solution to privacy concerns with online social networks. The premise of DOSNs is collaborative social content hosting through peers rather than a potentially untrusted centralized network. Examples of DOSNs include Diaspora [17], Peerson [18], My3 [12], and a DOSN Social CDN [19]. While the focus of our work differs in that we use a CDN to provide efficient, available and trustworthy access to data that is too big for traditional CDNs, many of these DOSNs include novel content distribution algorithms that are applicable in our approach.

My3 uses replication techniques across peers to ensure a user's content is only available to themselves and their friends even when they are offline. In this model, updates propagate amongst replicas until profiles are eventually consistent, replicas (or *trusted proxy sets*) are selected based on the availability and performance of a user's friends, and novel social algorithms using geographical location and analysis of when users and their friends are online. The Social CDN [19] is composed of social caches [20] which are used to communicate social updates amongst friends. Like our approach, the authors use the social network topology rather than traditional geographic location to assemble the CDN. The authors also propose several distributed selection algorithms to optimize cache selection and information dissemination in DOSNs. Of particular interest is the Social Score algorithm which takes into consideration social networking principles such as node centrality computed by each node to select social caches.

A social virtual private network (SocialVPN) [21], provides another useful approach for bridging the gap between social and overlay networking. A SocialVPN enables an automatic establishment of Peer-to-Peer links between participants that are connected through a social network. Establishment of such an overlay network involves the discovery of peers and the identification of cryptographic public certificates (e.g. X.509 certificates or X.509 certificate fingerprints) to associate with these peer identities. In the context of this work, a social networking infrastructure can range from Facebook messages to an encrypted Google talk chat session, and even a PGP-signed email exchange amongst peers. A prototype is demonstrated by the authors using the IPOP virtual network, across a variety of social networks (Facebook, the Drupal content management system, a PGP-signed email exchanges system). Experiments are carried out over 500 SocialVPN routers across five continents using the PlanetLab infrastructure, and over 100 SocialVPN virtual endpoints deployed dynamically over the Amazon Elastic Cloud (EC2) infrastructure. This approach can strongly complement the work described in this paper – as a SocialVPN can be used to establish secure links between participants involved in a S-CDN.

VIII. Summary and Future Work

The rapid growth of scientific data has led to a world in which collaborative researchers struggle to manage and organize big data. Often, it is difficult just to ensure that data is in the right place at the right time while also being accessible to those in a collaboration. In this paper we propose a unique approach to address these problems, by combining social principles with content delivery networks to create a Social CDN in which researchers are able to collaboratively host and access research scale datasets. To underly this approach, we present an analysis of trust from the perspective of collaborative researchers and also define a real-world data collaboration use case based on the analysis of medical images.

The S-CDN model builds upon a proven approach for delivering scalable access to large datasets, albeit in a vastly different domain. We extend the traditional CDN model by leveraging user-contributed storage resources to act as "edge nodes" in the system and aim to use novel replica placement and data partitioning algorithms to optimize data segmentation. Moreover, our architecture utilizes the social fabric of a social network both as a means of incorporating trust in the CDN and also as a basis for replica selection.

The work presented in this paper serves as an outline for our future research. Because trust is a primary motivator for the construction of a S-CDN, we aim to continue our analysis and formalization of trust as an enabler of collaboration and look into other mechanisms to extract trust from scientific networks. Before embarking on implementation, we will extend our analysis platform to simulate a more diverse range of attributes, such as data access algorithms, different research networks, and indicators of trust. We will use this platform to analyze new social algorithms and continue to explore different trust thresholds. Finally, we will implement a proof of concept S-CDN based on the architecture and network analysis presented in this paper.

References

[1] I. Foster. Globus online: Accelerating and democratizing science through cloud-based services. *IEEE Internet Computing*, 15(3):70–73, 2011.

[2] K. Chard, K. Bubendorfer, S. Caton, and O. Rana. Social cloud computing: A vision for socially motivated resource sharing. *IEEE Transactions on Services Computing*, 99(PrePrint):1, 2011.

[3] K. Chard, S. Caton, O. Rana, and K. Bubendorfer. Social cloud: Cloud computing in social networks. In *Proceedings of the 3rd IEEE International Conference on Cloud Computing (CLOUD)*, pages 99–106, 2010.

[4] The anatomy of facebook. http://tinyurl.com/fb-anatomy. Last accessed September 2012.

[5] F. Douglis and J. Ousterhout. Transparent process migration: Design alternatives and the sprite implementation. *Software Practice & Experience*, Vol. 21, No. 8:757 785, 1991.

[6] M. W. Mutka. Estimating capacity for sharing in a privately owned workstation environment. *IEEE Transactions on Software Engineering*, Vol. 18, No. 4:319 328, 1992.

[7] A. Acharya, G. Edjlali, and J. Saltz. The utility of exploiting idle workstations for parallel computation. In *Proceedings of ACM International Conference on Measurement and Modelling of Computer Systems (SIGMETRICS '97)*, pages 225–234, 1997.

[8] C. Haas, S. Caton, K. Chard, and C. Weinhardt. Co-operative infrastructures: An economic model for providing infrastructures for social cloud computing. In *Proceedings of the 46th Annual Hawaii International Conference on System Sciences (HICSS).*, 2013.

[9] S. Caton, C. Dukat, T. Grenz, C. Haas, M. Pfadenhauer, and C. Weinhardt. Foundations of trust: Contextualising trust in social clouds. In *Proceedings of the 2nd International Conference on Social Computing and Its Applications (SCA 2012)*, 2012.

[10] Cloudfiles from rackspace/openstack. http://www.rackspace.co.uk/cloud-files/. Last accessed October 2012.

[11] B. Allen, J. Bresnahan, L. Childers, I. Foster, G. Kandaswamy, R. Kettimuthu, J. Kordas, M. Link, S. Martin, K. Pickett, and S. Tuecke. Software as a service for data scientists. *Communications of the ACM*, 55:81–88, 2012.

[12] R. Narendula, T. G. Papaioannou, and K. Aberer. My3: A highly-available p2p-based online social network. In *Proceedings of the IEEE International Conference on Peer-to-Peer Computing (P2P)*, pages 166–167, 2011.

[13] M. Riedel, E. Laure, Th. Soddemann, L. Field, J. P. Navarro, J. Casey, M. Litmaath, J. Ph. Baud, B. Koblitz, C. Catlett, D. Skow, C. Zheng, P. M. Papadopoulos, M. Katz, N. Sharma, O. Smirnova, B. Kónya, P. Arzberger, F. Würthwein, A. S. Rana, T. Martin, M. Wan, V. Welch, T. Rimovsky, S. Newhouse, A. Vanni, Y. Tanaka, Y. Tanimura, T. Ikegami, D. Abramson, C. Enticott, G. Jenkins, R. Pordes, N. Sharma, S. Timm, N. Sharma, G. Moont, M. Aggarwal, D. Colling, O. van der Aa, A. Sim, V. Natarajan, A. Shoshani, J. Gu, S. Chen, G. Galang, R. Zappi, L. Magnoni, V. Ciaschini, M. Pace, V. Venturi, M. Marzolla, P. Andreetto, B. Cowles, S. Wang, Y. Saeki, H. Sato, S. Matsuoka, P. Uthayopas, S. Sriprayoonsakul, O. Koeroo, M. Viljoen, L. Pearlman, S. Pickles, David Wallom, G. Moloney, J. Lauret, J. Marsteller, P. Sheldon, S. Pathak, S. De Witt, J. Mencák, J. Jensen, M. Hodges, D. Ross, S. Phatanapherom, G. Netzer, A. R. Gregersen, M. Jones, S. Chen,

P. Kacsuk, A. Streit, D. Mallmann, F. Wolf, Th. Lippert, Th. Delaitre, E. Huedo, and N. Geddes. Interoperation of world-wide production e-science infrastructures. *Concurrency and Computation: Practice & Experience*, 21(8):961–990, 2009.

[14] J. Broberg and Z. Tari. Metacdn: Harnessing storage clouds for high performance content delivery. In *Proceedings of the 6th International Conference on Service-Oriented Computing (ICSOC '08)*, pages 730–731, Berlin, Heidelberg, 2008. Springer-Verlag.

[15] J. Broberg, R. Buyya, and Z. Tari. Metacdn: Harnessing 'storage clouds' for high performance content delivery. *Journal of Network and Computer Applications*, 32(5):1012–1022, 2009.

[16] Y. Jin, Y. Wen, G. Shi, G. Wang, and A. V. Vasilakos. Codaas: An experimental cloud-centric content delivery platform for user-generated contents. In *Proceedings of the International Conference on Computing, Networking and Communications (ICNC)*, pages 934–938, 2012.

[17] The diaspora project. http://diasporaproject.org/. Last accessed October 2012.

[18] S. Buchegger, D. Schiöberg, L. Vu, and A. Datta. Peerson: P2p social networking: early experiences and insights. In *Proceedings of the 2nd ACM EuroSys Workshop on Social Network Systems (SNS '09)*, pages 46–52, 2009.

[19] L. Han, M. Punceva, B. Nath, S. Muthukrishnan, and L. Iftode. Socialcdn: Caching techniques for distributed social networks. In *Proceedings of the 12th IEEE International Conference on Peer-to-Peer Computing (P2P)*, pages 191–202, 2012.

[20] L. Han, B. Nath, L. Iftode, and S. Muthukrishnan. Social butterfly: Social caches for distributed social networks. In *SocialCom*, pages 81–86, 2011.

[21] P. St. Juste, D. Wolinsky, P. Oscar Boykin, M. J. Covington, and R. J. Figueiredo. Socialvpn: Enabling wide-area collaboration with integrated social and overlay networks. *Computer Networks (Elsevier)*, 54(12):1926–1938, 2010.