

A Mid-Flight Synopsis of the BG Social Networking Benchmark

Shahram Ghandeharizadeh^(✉) and Sumita Barahmand

Database Laboratory, Computer Science Department,
USC, Los Angeles, CA 90089-0781, USA
{shahram, barahman}@usc.edu

Abstract. BG is a benchmark that rates the performance of a data store for processing interactive social networking actions such as view a member’s profile, invite a member to be friends, accept a friend request, and others. It is motivated by a proliferation of data stores from a variety of academic and industrial contributors including social networking companies, e.g., Voldemort by LinkedIn. BG is designed to provide a system architect with insights into alternative design principles such as the use of a weak consistency technique instead of a strong one, different physical data models such as relational and JSON, factors that impact vertical and horizontal scalability of a data store, the consistency versus availability tradeoff in the CAP theorem, among others. While BG is a recently introduced benchmark (less than a year old as of this writing), it combines elements of maturer benchmarks and extends them to simplify its use by the practitioners and experimentalists. This paper provides a synopsis of the BG benchmark by identifying its strengths and limitations in our daily use cases. The identified limitations shape our research activities and the obtained solutions shall be incorporated into future BG releases. Thus, this workshop paper is a mid-flight glimpse into our current research efforts with BG.

1 Introduction

In an article that appeared in the July 2012 issue of the Communications of the ACM, David Patterson observes when a discipline has good benchmarks, debates are settled and the discipline makes rapid progress [16]. Today, we have an abundance of architectures for data stores and services with only a handful of benchmarks to substantiate their many claims. Academia, cloud service providers such as Google and Amazon, social networking sites such as LinkedIn and Facebook, and computer industry continue to contribute systems and services with novel architectures and assumptions. In 2010, Rick Cattell surveyed 23 systems [7] and we are aware of 10 new¹ ones since that writing. In his survey, Cattell identified a “gaping hole” with a scarcity of benchmarks to substantiate

¹ Apache Jackrabbit and RavenDB, Titan, Oracle NoSQL, FoundationDB, STSdb, EJDB, FatDB, SAP HANA, CouchBase.

the claims made by the different systems. We have designed and implemented a social networking benchmark named BG [2] (visit <http://bgbenchmark.org>) to address certain aspects of the hole that is too large to address with just one benchmark.

BG's workload consists of actions that either read or write a small amount of data from big data, typically termed simple operations [10,21]. Today's BG is designed for high throughput data stores that provide interactive response times. A long term objective is to extend BG with complex analytics that require processing of a large amount of data using machine learning algorithms. This would make BG suitable to evaluate Hadoop and other implementation of the MapReduce [8] framework, see Sect. 7 for details.

We developed BG in 2012 and released a stable version of it in January 2013. Its conceptual schema and eleven actions are an abstraction of today's social networking sites such as Google+, Facebook and others. In [2], we provide a comprehensive list of the surveyed sites and a matrix that describes compatibility of BG's actions with those supported by a site. Figure 1 shows BG's conceptual schema. The concept of members with registered profiles befriending one another is at the core of this schema. Its implementation in a physical data store is dictated by an experimentalist. BG is data store agnostic and one may tailor both the physical schema and the implementation of actions to highlight the strengths of a data store. At the time of this writing, an implementation of BG's schema and actions is available for the following data stores:

- MySQL, Oracle, PostgreSQL, and VoltDB as relational data stores.
- MongoDB and CouchBase as document stores.
- Microsoft Azure as a cloud service provider.
- Hibernate as an Object Relational Mapping (ORM) framework.
- Cache augmented data stores with memcached, EhCache, Twemcache, and KOSAR.

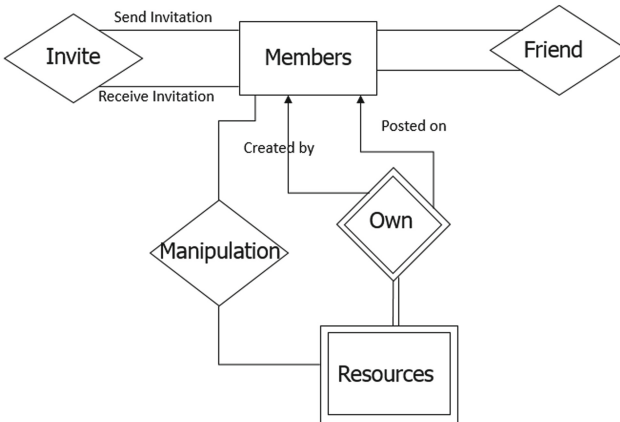


Fig. 1. Conceptual design of BG's database.

Table 1. Three mixes of social networking actions.

BG social actions	Type	Very low	Low	High
		(0.1 %) Write	(1 %) Write	(10 %) Write
View Profile (VP)	Read	40 %	40 %	35 %
List Friends (LF)	Read	5 %	5 %	5 %
View Friends Requests (VFR)	Read	5 %	5 %	5 %
Invite Friend (IF)	Write	0.04 %	0.4 %	4 %
Accept Friend Request (AFR)	Write	0.02 %	0.2 %	2 %
Reject Friend Request (RFR)	Write	0.02 %	0.2 %	2 %
Thaw Friendship (TF)	Write	0.02 %	0.2 %	2 %
View Top-K Resources (VTR)	Read	40 %	40 %	35 %
View Comments on Resource (VCR)	Read	9.9 %	9 %	10 %
Post Comment on a Resource (PCR)	Write	0 %	0 %	0 %
Delete Comment from a Resource (DCR)	Write	0 %	0 %	0 %

The first column of Table 1 shows the eleven actions that constitute BG. The name of the actions is self explanatory. All actions that reference members are binary consuming two member ids as input. For example, the two member ids specified with the View Profile action identify the member who is viewing a profile and the member whose profile is being viewed. Those actions that consume a resource id either read the resource and its comments or modify a comment on that unique resource. The different actions may either read or write data from a data store as highlighted by the second column of Table 1. We refer the interested reader to [2] for a detailed description of each action.

The last three columns of Table 1 show three different workloads corresponding to a very low, low, and a high percentage of write actions. (According to Facebook, more than 99 % of its workload consists of read actions [9].) Each column specifies a fixed percentage of occurrence for each action in a workload. We use the presented three workloads in our experiments on a daily basis. All three are symmetric workloads that cause an experiment to complete with approximately the same number of confirmed friendships and pending friendships as those in the beginning of the experiment. The number of these relationships is impacted by the frequency of the following actions: Invite Friend, Accept Friend Request, Reject Friend Request, and Thaw Friendship actions. For this number to remain unchanged, the rate at which BG generates friendships should equal the rate at which it thaws friendships. This is realized by satisfying the following two conditions: (1) percentage of Thaw Friendship and Accept Friend Request must be identical, and (2) percentage of Invite Friend must equal the sum of percentage of Reject Friend Request and Accept Friend Request. In Table 1, the frequency of Post Comment on a Resource (PCR) and Delete Comment from a Resource (DCR) are intentionally kept at zero to demonstrate that one may specify workloads consisting of either a single or a few actions. When specifying

frequencies of PCR and DCR, a symmetric workload should define the same frequency of occurrence for each action.

BG is a stateful benchmark that generates valid actions. For example, it extends a friendship from Member A to Member B only when they are not friends. It realizes this by maintaining a representation of the social graph in its memory. BG uses this representation to ensure its emulated simultaneous members and resources are unique at an instance in time.

A novel feature of BG is its ability to quantify the amount of unpredictable (stale, inconsistent, erroneous) data produced by a data store. BG evaluates a data store for a workload that specifies an SLA. An example SLA may require 95% of actions to be performed faster than 100 milliseconds with no more than 0.01% unpredictable data for $\Delta = 10$ min. BG includes a heuristic search technique to quantify the maximum throughput (actions per second) observed with a data store while satisfying the pre-specified SLA. This is termed the Social Action Rating, SoAR, of the data store.

We have employed BG to investigate design and implementation of novel architectures for data intensive applications, e.g., to compare a relational representation of a social graph with its JSON representation [5], quantify the trade-offs associated with alternative consistency techniques for a cache augmented relational data store [12], and others. In these use cases, we have identified several limitations with BG’s design. These shape our research efforts to extend BG to maintain it as a state of the art benchmark. Most are in the context of the novel features of BG that make it unique. Below, we describe these in turn, detailing BG’s scalable request generation in Sect. 2, its closed emulation of socialites and an alternative open emulator in Sect. 3, its rating mechanism in Sect. 4, its validation phase in Sect. 5, and additional actions in Sect. 6. Section 7 provides our long term future research.

2 Scalability

BG employs a shared-nothing architecture and scales to a large number of nodes, preventing either the CPU, network, or memory resources of a single node from limiting its request generation rate. Its software architecture consists of one coordinator and N clients, termed BGCoord and BGClient, respectively. In our experiments with an 8 core CPU, a multi-threaded BGClient is able to utilize all cores fully as long as the client component of a data store does not suffer from the convoy phenomena [6] and the data store is able to process requests at the rate generated by BG. When the client component of a data store limits vertical scalability, as long as there is a sufficient amount of memory, one may execute multiple instances of BGClients on a single node to utilize all cores. BG scales horizontally by executing multiple BGClients across different nodes. BGCoord is responsible for initiating the BGClients, monitoring their progress, gathering their results at the end of an experiment, and aggregating the obtained results to compute the SoAR of a data store.

Once the BGClient instances are started, they generate requests independently with no synchronization. This is made possible using the following two concepts.

First, a BGClient implements a decentralized partitioning strategy that declusters a benchmark social graph into N disjoint sub-graphs where N is the number of BGClients. A BGClient is assigned a sub-graph to generate requests referencing members of its assigned sub-graph only. While the data store is not aware of this partitioning, the data generated and stored in the data store does correspond to the N disjoint graphs. One may conceptualize each sub-graph as a province whose citizens may perform BG’s actions with one another only. This means citizens of different provinces may not view one another’s profile or become friends with one another.

Second, BG employs a novel decentralized implementation of the Zipfian distribution, named D-Zipfian [3,24], that ensures the distribution of requests to the different members is independent of N . Thus, the distribution of access with one node is the same as that with several nodes. D-Zipfian in combination with partitioning of the social graph enables BG to utilize N nodes to generate requests without requiring coordination until the end of the experiment, see [2–4] for details.

While BG scales to a large numbers of nodes, its two concepts may fail to evaluate some data stores objectively. As an example, consider the architecture of Fig. 2 where an application is extended with a cache such as KOSAR or EhCache [12]. This caching framework consists of a KOSAR coordinator that maintains which application server has cached a copy of a data item in its KOSAR JDBC wrapper, KOSAR-Client for short. When one application server updates a copy of the data item, its KOSAR-Client informs the KOSAR coordinator of the impacted data item. In turn, the KOSAR coordinator invalidates a copy of this data item that resides in the KOSAR-Client of other application servers. With a skewed pattern of access to members and a workload that exhibits a low read to write ratio, a centralized KOSAR coordinator may become the bottleneck and dictate the overall system performance. The aforementioned

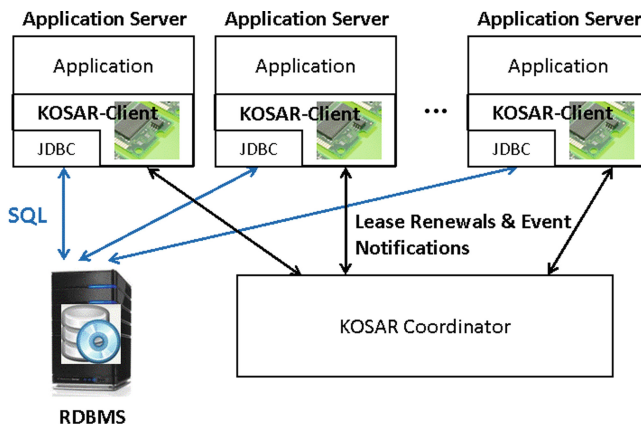


Fig. 2. A data intensive architecture using KOSAR.

two concepts employed by BG fail to cause the formation of such a bottleneck. To elaborate, each application server references data items that are unique to itself since its assigned sub-graph is unique and independent of the other sub-graphs. Hence, once an application server updates a cached data item, BG does not exercise the KOSAR coordinator informing KOSAR-Client of another application server.

To address the above limitation, we are extending BG to employ N BGClients with one social graph. The key concept is to hash partition members and resources across the N BGClients. Each BGClient is aware of the hash function and employs the original Zipfian distribution (instead of D-Zipfian) to generate member ids. When a BGClient BGC_i references a data item that does not belong to its assigned partition, it contacts the BGClient that owns the referenced data (say BGC_j) to lock that data item for exclusive use by BGC_i and to determine if its intended action is possible. BGC_j grants the lock request if there is no existing lock on the referenced data item and the action is possible, enabling BGC_i to proceed to generate a request with the identified data item to the data store. Once the request is serviced, BGC_i contacts BGC_j to release the exclusive lock on the referenced data item to make it available for use by other BGClients. This design raises the following interesting questions:

- When BGC_j fails to grant an exclusive lock² to the referenced data item due to an existing lock, how should the framework handle the conflict? Three possibilities are as follows. First, it may block BGC_i until the referenced data item becomes available. Second, it may return an error to BGC_i to generate a different member/resource id and try again. Third, it may simply abort this action and generate a new action all together. We intend to quantify the tradeoff associated with these three possibilities and their impact on both the distribution of requests and the benchmarking framework.
- What is the scalability characteristic of the proposed technique? The proposed request generation technique requires different BGClients to exchange messages to lock and unlock data items and to determine the feasibility of actions. We plan to quantify this overhead and its impact on the scalability of this request generation technique. This intuition should enable us to propose refinements to enhance scalability.
- How different are the obtained results with N disjoint social graphs (current version of BG) and one social graph (the proposed change)? This question applies to those systems that may use the current version of BG. We intend to repeat our published experiments such as those reported in [5] to quantify differences if any.

An investigation of these questions shapes our short term research direction.

² BGC_j may return an error code when the action is not possible. For example, Thaw Friendship using Member A may not be feasible because A has no friends. In these cases, BGC_i may either abort the action or may reference a new member for the same action.

3 Closed Versus Open

BGclients generate requests using a fixed number of threads T . Each thread emulates a random member of a social networking site performing one of the eleven actions. The randomly selected member is conditioned using the D-Zipfian distribution. This is termed a closed emulation model because a thread does not emulate a new member generating a new action until its emulation of a current member completes. This model may include a think time between emulation of different members issuing actions. Historically, this is a model³ of a financial institution with a fixed number of tellers (ATM machines) with T concurrent customers (threads) performing financial transactions simultaneously [13].

An open emulator is a more realistic model of a social networking site [17] (and web sites in general). With this model, the emulator generates requests based on a pre-specified arrival rate, λ . This model is depicted in Fig. 3 where a factory generates members who issue a social networking action independently. (A member who is performing an action is termed a *socialite*.) The factory does not wait for the data store to service a request issued by a socialite. Instead, it generates λ socialites issuing requests per unit of time using a distribution such as random, uniform, or Poisson. A Poisson distribution results in a pattern of requests that is bursty. This means λ is an average and the number of simultaneous socialites at an instance in time might be higher than λ .

While the open emulator is more realistic, its design and implementation requires a careful study. This is because today's data stores service requests at such a high rate that the emulator must support λ values in the order of a million without exhausting its CPU resources. In addition to the scalability discussions of Sect. 2, the emulator must generate requests in a burst consistent with the Poisson distribution. At the time of this writing, we are evaluating the feasibility of such an open emulator and its implementation in BG.

4 Rating Mechanism

BG *rates* a data store to compute its Social Action Rating (SoAR) and Socialites rating for processing a workload. A workload consists of a mix of the eleven actions (see Table 1 and its discussion in Sect. 1), an exponent for the D-Zipfian distribution to control its degree of skew when referencing members, and a pre-specified SLA, see Sect. 1 for an example SLA. The SoAR of a data store is the highest throughput provided by that data store for the specified workload. The Socialite rating of a data store is the maximum number of simultaneous socialites (threads) that may generate requests corresponding to the specified workload. The pre-specified SLA imposes constraints on the acceptable response times and the amount of unpredictable data to constrain the SoAR and Socialite ratings of a data store. Given several data stores, the data store with the highest SoAR and Socialite rating is the superior one.

³ This model is a representative of a web server configured with a maximum number of threads.

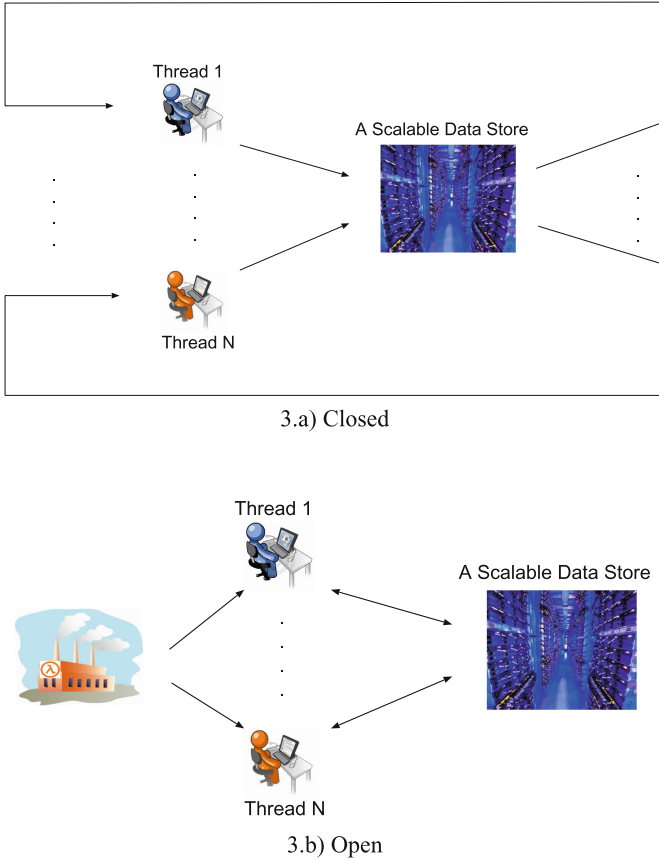


Fig. 3. Closed and open emulation of socialites issuing actions to a data store.

BG’s rating process and its assumptions are detailed in [2]. Briefly, the rating process consists of a heuristic search that conducts multiple experiments. Each experiment uses the same workload specified by an experimentalist. With those workloads that impact the state of the database, the rating process might be required to re-load the database prior to each experiment. The repeat loading of a benchmark database may constitute a significant portion of the rating process. For example, the time to generate a one million member database with a data store requires 11 days [4]. If the rating process conducts ten experiments each 30 min in duration, the time to generate the database each time would require almost 4 months.

Three alternatives that re-generate the database expeditiously are detailed in [4]. One is the Database Image Loading (*DBIL*) technique that maintains the original image of a database and copies it as the current database in advance of each experiment to reduce the load time. The time to copy the one million member database is 30 min [4], reducing the time to conduct ten experiments to

11 days and 10 h. The 11 days incurred to create the database for the first time is a one time overhead. By maintaining the image and re-using it, the duration of rating of the data store with different workloads is reduced dramatically. For example, the time to conduct ten experiments is now 10 h. We refer the interested reader to [4] for a detailed description of DBIL and two other techniques.

A key research question is the duration of each conducted experiment. A possible answer is to employ the duration specified by the SLA, Δ . However, if an experimentalist select a high value for this input parameter, then the rating process may consume more time than necessary. It is desirable to run experiments for shorter durations than Δ to identify a region in the search space that establishes the true SoAR of a data store. Ideally, the duration of an experiment should be the smallest possible value, δ , that reflects the behavior of a data store as if the experiment was running for Δ time units. The ideal δ is both data store and workload dependent and can be analyzed in a pre-processing step, prior to the rating process. This step involves multiple experiments issuing the given workload against the data store to select the smallest duration that results in a *steady* system behavior defined as one whose resource utilization and observed throughput do not change in time. For such a system the recently observed behavior will continue to hold into the future.

5 Validation

A novel feature of BG is its ability to quantify the amount of unpredictable data (stale, inconsistent, erroneous) produced by a data store. A data store may produce unpredictable data for a variety of reasons. Examples include use of a weak consistency technique such as eventual [20, 23] and use of a cache [12] in a manner that results in dirty reads [15] and inconsistent cache states [11].

BG is able to measure the amount of unpredictable data because it is a stateful benchmark that is aware of the initial state of a data item in the database and its updates to the database. It maintains the start and end of each action to enumerate the finite number of ways a read may overlap multiple concurrent write actions that reference the same data item. BG enumerates these to compute a range of possible values that should be observed by the read operation. If a data store produces a different value then it has produced unpredictable data. This process is named *validation*.

BG decouples generation of requests to quantify the performance of a data store from the validation phase, performing validation in an off-line manner. When generating requests, each BG thread generates a log record for each of its actions: read log records for read actions and write log records for write actions. These log records are written to separate files. One file for the read log records and a second file for the write log records.

The validation phase assumes in-memory data structures to compute the amount of unpredictable data as follows. First, it maintains an interval tree for (1) each member and write actions that impact her friendships, (2) each member and write actions that impact her pending friend invitations, and (3) each

resource that is annotated with a write action. It constructs interval trees for a member/resource on demand as it reads the write log records in memory. The start and end time stamp of the write log records are indexed by an interval tree. Once the write log records are staged in memory, the validation phase retrieves the read log records. It employs the member id (resource id) and the action to identify the interval tree with the relevant write log records. Next, it uses the start and end time stamp of each read log record to enumerate the number of ways it overlaps with the different write actions, computing a range of valid values that should be observed by the read action. If the value observed by the read action (recorded in the log record) does not match one of the valid values then this read action has observed an unpredictable value.

The current implementation of the validation phase is fast when there is a sufficient amount of memory to stage the write log records in interval trees. This enables the validation phase to read the log files once to process both read and write log records in one pass. This multi-threaded process may utilize multiple cores fully because (1) write log records are inserted into different interval trees based on their referenced member/resource id, and (2) different read log records may share and read the same interval tree simultaneously.

By performing validation in an off-line manner, BG reduces the number of nodes required to generate request to evaluate a data store. A key limitation of the current validation technique is that it may exhaust the available physical memory, causing the operating system to exhibit a thrashing behavior that results in an unacceptably long validation process. This is especially true with high throughput multi-node data stores and cache augmented data stores such as KOSAR that process requests in the order of millions of actions per second. We intend to extend the validation phase to analyze the size of files produced during an experiment to estimate the amount of required memory to ensure it does not exhaust the physical memory. In passing, it is interesting to note that the log records pertaining to each unique member can be processed independently to identify unpredictable data. Hence, a MapReduce [8] framework such as Hadoop is effective in implementing the validation phase.

6 Additional Actions

The eleven actions of BG are a good start to evaluate a data store. However, there are other actions that are common to many social networking sites that we intend to abstract. An implementation of these would extend BG with new actions. Here, we focus on two new actions to be released soon, namely, Share Resource (SR) and Retrieve Feed (RF). Both are in support of feed following [1, 18, 19]. This action is supported by sites such as Google+, Twitter, Facebook, My Yahoo and others. It enables users to create personalized feed by selecting one or more event streams they wish to follow.

Figure 4 shows the high level (incomplete) ER diagram for feed following. While the conceptual model looks complex, it is based on the concept of aggregation that establishes the many-to-many relationship between producers and

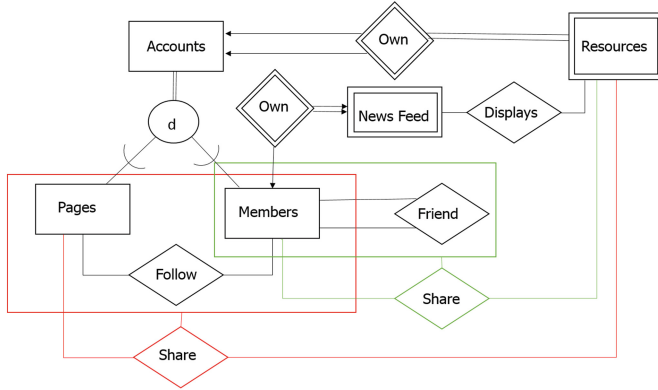


Fig. 4. Conceptual design of BG’s feed following (Color figure online).

consumers of news feeds. The producers are a specialization of the Accounts entity set, identifying Pages and Members as two different categories because pages may have a significantly larger number of followers⁴ and do not perform some of the actions performed by members. Examples of pages are celebrity fan pages and company and brand pages among the others. Resources such as images and tweets posted by a page are shared with their followers. This is represented by the red rectangle that aggregates the follow relationship between members and pages as an entity that participates in the “Share” relationship (red line) with the Resources entity set. Similarly, resources posted by the members are shared with their circle of friends which are also members. This is denoted by the green aggregation that participates in the “Share” relationship with the Resources entity set. Each member owns a “News feed” that contains resources shared by those whom they are following.

In the near future, we intend to investigate alternative techniques to implement feed following; see [19] for a pull, push, and a hybrid of these two techniques. We are also exploring techniques that compute an approximate feed. We intend to use BG to quantify both the scalability of these techniques and the performance gain of the approximate technique relative to its amount of unpredictable data.

7 Conclusions and Future Research

Social networks are emerging in diverse applications that strive to provide a sense of community for their users. These diverse applications range from financial web sites such as online trading system to academic institutions. BG is a benchmark to evaluate the performance of a data store for processing interactive social networking actions such as viewing a member’s profile, extending a friendship request to a member, accepting a friendship request, and others as shown

⁴ As of November 12, 2013, Katy Perry had more than forty million Twitter followers.

in Table 1. We use BG on a daily basis to evaluate the performance of novel architectures that enable high throughput, low latency data stores. Obtained insights enable us to introduce novel designs and implementations. We identified key features of BG and how we are refining these to ensure BG remains a state of the art benchmark. Each section focuses on one feature of BG. However, a change in one aspect of BG impacts the other features of BG. For example, the two design principles that enable BG to generate requests in a scalable manner (see Sect. 2) also enables multiple BGClients to perform the validation phase (see Sect. 5) independently and in parallel with one another. The proposed modifications of Sect. 2 on how BG generates requests requires a revisit of how BG performs its validation phase.

We intend to extend today’s BG beyond simple analytics to include additional Web 2.0 workloads including complex analytics that require the use of machine learning algorithms. An example is a recommender system that analyzes the social graph and attribute values to suggest friends for a member. A challenge here is to generate the social graph in a manner where the output of the recommendation system is known, extending BG with metrics such as precision and recall. The graph database may include annotations on the edges of the graph with attribute values. For example, the friendship relationship between two members A and B might be tagged with a family relationships (such as daughter), number of likes given by Member A to postings by B’s friends, number of comments posted by B’s friends on A’s resources, and others. We may change how the value of textual attributes for a comment are generated in a manner that is more realistic using techniques such as those utilized by the TREC benchmark. This may evaluate alternative (1) natural language processing and (2) information retrieval techniques in the context a recommender system for a social graph.

Finally, we are investigating the viability of a *Benchmark Generator*, BG+, that inputs an abstraction of an application, its actions and their dependencies, metrics of interest to be quantified, and a control parameter. Its output is a benchmark specific to that application. This is appropriate for those applications with diverse use cases such as data sciences [14,22]. In essence BG+ is an extensible toolkit that is configured with its input to enable a data scientist to develop a benchmark for their activities rapidly. Similar to BG and YCSB, it would expose the implementation of the database schema and the abstracted actions to be implemented by the data scientists. This enables BG+ to support diverse types of data models such as structured (relational), un-structured, JSON, extensible, images, audio and video as input. The control parameter may manipulate factors such as noise in the data (similar to today’s system load controlled by parameter T) to analyze the behavior of an algorithm. This would enable BG+ to output a rating mechanism that computes a single value (a maxima such as SoAR) by manipulating the value of the control parameter.

References

1. Bai, X., Junqueira, F.P., Silberstein, A.: Cache refreshing for online social news feeds. In: CIKM, pp. 787–792 (2013)
2. Barahmand, S., Ghandeharizadeh, S.: BG: a benchmark to evaluate interactive social networking actions. In: CIDR, Jan 2013
3. Barahmand, S., Ghandeharizadeh, S.: D-Zipfian: a decentralized implementation of Zipfian. In: ACM SIGMOD DBTest Workshop (2013)
4. Barahmand, S., Ghandeharizadeh, S.: Expedited benchmarking of social networking actions with agile data load techniques. In: CIKM (2013)
5. Barahmand, S., Ghandeharizadeh, S., Yap, J.: A comparison of two physical data designs for interactive social networking actions. In: CIKM (2013)
6. Blasgen, M.W., Gray, J., Mitoma, M.F., Price, T.G.: The convoy phenomenon. *Oper. Syst. Rev.* **13**(2), 20–25 (1979)
7. Cattell, R.: Scalable SQL and NoSQL data stores. *SIGMOD Rec.* **39**, 12–27 (2011)
8. Dean, J., Ghemawat, S.: MapReduce: simplified data processing on large clusters. In: Symposium on Operating Systems Design and Implementation, vol. 6 (2004)
9. Nishtala, R., et al.: Scaling memcache at Facebook. In: NSDI (2013)
10. Floratou, A., Teletria, N., DeWitt, D.J., Patel, J.M., Zhang, D.: Can the elephants handle the NoSQL onslaught? In: VLDB (2012)
11. Ghandeharizadeh, S., Yap, J.: Gumball: a race condition prevention technique for cache augmented SQL database management systems. In: ACM SIGMOD DBSocial Workshop (2012)
12. Ghandeharizadeh, S., Yap, J.: Cache augmented database management systems. In: ACM SIGMOD DBSocial Workshop, June 2013
13. Gray, J., Reuter, A.: *Transaction Processing: Concepts and Techniques*, pp. 677–680. Morgan Kaufmann, San Francisco (1993)
14. Greenberg, C.: Overview of the NIST data science evaluation and metrology plans. In: Data Science Symposium, NIST, 4–5 Mar 2014
15. Gupta, P., Zeldovich, N., Madden, S.: A trigger-based middleware cache for ORMs. In: *Middleware* (2011)
16. Patterson, D.: For better or worse, benchmarks shape a field. *Commun. ACM* **55**, 104 (2012)
17. Schroeder, B., Wierman, A., Harchol-Balter, M.: Open versus closed: a cautionary tale. In: NSDI (2006)
18. Silberstein, A., Machanavajjhala, A., Ramakrishnan, R.: Feed following: the big data challenge in social applications. In: DBSocial, pp. 1–6 (2011)
19. Silberstein, A., Terrace, J., Cooper, B.F., Ramakrishnan, R.: Feeding frenzy: selectively materializing users event feeds. In: SIGMOD Conference, pp. 831–842 (2010)
20. Stonebraker, M.: Errors in database systems, eventual consistency, and the CAP theorem. *Commun. ACM. BLOG@ACM*, Apr 2010
21. Stonebraker, M., Cattell, R.: 10 rules for scalable performance in simple operation datastores. *Commun. ACM* **54**, 72–80 (2011)
22. Talukder, A.: Overview of the NIST data science program. In: Data Science Symposium, NIST, 4–5 Mar 2014
23. Vogels, W.: Eventually consistent. *Commun. ACM* **52**(1), 40–45 (2009)
24. Yap, J., Ghandeharizadeh, S., Barahmand, S.: An analysis of BGs implementation of the Zipfian distribution. USC DBLAB technical report 2013-02 (2013). <http://dmlab.usc.edu/Users/papers/zipf.pdf>



<http://www.springer.com/978-3-319-10595-6>

Advancing Big Data Benchmarks

Proceedings of the 2013 Workshop Series on Big Data
Benchmarking, WBDB.cn, Xi'an, China, July 16–17, 2013 and
WBDB.us, San José, CA, USA, October 9–10, 2013, Revised
Selected Papers

Rabl, T.; Raghunath, N.; Poess, M.; Bhandarkar, M.;
Jacobsen, H.-A.; Baru, C. (Eds.)

2014, XII, 203 p. 84 illus., Softcover

ISBN: 978-3-319-10595-6