

GDist-RIA Crawler: A Greedy Distributed Crawler for Rich Internet Applications

Seyed M. Mirtaheri¹(✉), Gregor von Bochmann¹, Guy-Vincent Jourdan¹,
and Iosif Viorel Onut²

¹ School of Electrical Engineering and Computer Science,
University of Ottawa, Ottawa, Ontario, Canada

staheri@uottawa.ca, {gvj,bochmann}@eecs.uottawa.ca

² Security AppScan[®] Enterprise, IBM, 770 Palladium Dr, Ottawa, Ontario, Canada
vioonut@ca.ibm.com

Abstract. Crawling web applications is important for indexing, accessibility and security assessment. Crawling traditional web applications is an old problem, for which good and efficient solution are known. Crawling Rich Internet Applications (RIA) quickly and efficiently, however, is an open problem. Technologies such as AJAX and partial Document Object Model (DOM) updates only make the problem of crawling RIA more time consuming to the web crawler. One way to reduce the time to crawl a RIA is to crawl a RIA in parallel with multiple computers. Previously published *Dist-RIA Crawler* presents a distributed breath-first search algorithm to crawl RIAs. This paper expands Dist-RIA Crawler in two ways. First, it introduces an adaptive load-balancing algorithm that enables the crawler to learn about the speed of the nodes and adapt to changes, thus better utilize the resources. Second, it present a distributed greedy algorithm to crawl a RIA in parallel, called *GDist-RIA Crawler*. The GDist-RIA Crawler uses a server-client architecture where the server dispatched crawling jobs to the crawling clients. This paper illustrates a prototype implementation of the GDist-RIA Crawler, explains some of the techniques used to implement the prototype and inspects empirical performance measurements.

Keywords: Web crawling · Rich internet application · Greedy algorithm · Load-balancing

1 Introduction

Crawling is the process of exploring and discovering states of a web application automatically. This problem has a long and interesting history. Throughout the history of web-crawling, the chief focus of web-crawlers has been on crawling traditional web applications. In these applications there is a one to one correspondence between the state of the web application and its URL. The new generation of web applications, called *Rich Internet Applications* (RIAs), take

advantage of availability of powerful client-side web-browsers and shift some part of application logic to the client. This shift often breaks the assumption of one-to-one correspondance between the URL and the state of the application. Thus, unlike a traditional web application, in crawling a RIA it is not sufficient to discover all application URLs, and it involves discovering all application states.

In a RIA, a client-side page, associated with a single URL, often contains executable code that may change the state of the page as seen by the user. This state is stored within the browser, and is called the *Document Object Model* (DOM). Its structure is encoded in HTML and includes the program fragments executed in response to user input. Code execution is normally triggered by events invoked by the user, such as *mouse over* or *clicking* events. To ensure that a crawler finds all application content it must execute every events from every reachable application states. Thus, under the assumption that a RIA is deterministic, the problem of crawling is reduced to the problem of executing all events in the application across all reachable DOMs.

One can reduce the time it takes to crawl a RIA by executing the crawl in parallel on multiple computational units. By considering each state of the application on the client side (henceforth simply referred to as *state*) as a vertex and each JavaScript event as an edge, the problem of the parallel crawling a RIA is mapped to the problem of parallel exploration of a directed graph.

Dist-RIA Crawler [27] introduced a distributed crawler for RIAs that achieves parallelism by having all the crawlers go to each application state, however, each crawler only explores a specific subset of the events in that vertex. The union of all these events covers all of the events in the state. In Dist-RIA Crawler, each crawler node implements a breath-first search algorithm in its own scope.

Dist-RIA Crawler assigns equal number of events to each node. The underlying assumption is that all nodes have equal processing power, and thus equal workload is to be assigned to the nodes. To enhance Dist-RIA Crawler to take advantage of heterogeneous set of nodes available, this paper introduces a mechanism to adapt to the perceived speed and processing power of the nodes. This algorithm is explained in Sect. 3.

In the context of RIA crawling, crawling strategy refers to the strategy the crawler follows to decide the next event to execute. Dincturk et al. [5, 12, 14] studied several crawling strategies to optimize the crawl in two dimensions: reducing the total time of the crawl, and finding new application states as soon as possible in the crawl. Among the strategies studied, the greedy algorithm [28] scores well in the majority of cases, and it is much better than breath-first and depth-first search strategies. This algorithm always chooses the closest application state with an un-executed event, goes to the state and execute the event. This paper studies distribution of the greedy algorithm.

In Dist-RIA Crawler, the nodes only broadcast the knowledge of application states, and no single node had the entire knowledge of the transitions between the states. This restriction does not allow a Dist-RIA Crawler to run the greedy algorithm: knowledge of application transitions is a prerequisite for the greedy

algorithm. At the same time, broadcasting all transitions to the entire group of workers can make the network a bottleneck.

This paper introduces GDist-RIA Crawler, a client-server architecture to integrate the greedy algorithm into the architecture of the Dist-RIA Crawler. The GDist-RIA Crawler runs the greedy algorithm on the server and runs the crawling jobs to the client nodes. The server node is henceforth referred to as the *coordinator* and the client nodes responsible to crawl the website are henceforth referred to as the *nodes*. Nodes ask the coordinator for tasks to do, the coordinator runs the greedy algorithm on the application graph and responds them with a set of events to execute. Nodes execute the assigned tasks and inform the coordinator about the transition they discovered. The coordinator is the only computer that keeps the knowledge of application graph.

The greedy nature of the algorithm makes the GDist-RIA Crawler superior to the Dist-RIA Crawler (which runs breath-first search) by reducing the total number of events executed to crawl an application. The GDist-RIA Crawler is also superior to the centralized greedy algorithm in that it harnesses the power of multiple nodes to reduce the time it takes to crawl the target application. Further, it does not require the load-balancing algorithm introduced in Sect. 3 that is required by the breath-first search strategy, since only idle nodes ask for work from the coordinator, and thus no node becomes a bottleneck.

This paper contributes to the body of crawling literature by enhancing the previously presented Dist-RIA Crawler in two ways. First by introducing an adaptive load-balancing strategy to harness availability of heterogenous nodes. Second by introducing a client-server architecture to concurrently crawl RIAs. We share our empirical experience with the introduced model and some of the challenges we faced in capturing client-side events.

The rest of this paper is organized as follows: In Sect. 3 we introduce a new adaptive load-balancing algorithm. In Sect. 4 we give an overview of the GDist-RIA Crawler. In Sect. 5 we describe some of the technical aspects of implementing the GDist-RIA crawler. In Sect. 6 we evaluate various performance aspects of the GDist-RIA Crawler. In Sect. 2 we give an overview of the related works. Finally, in Sect. 7 we conclude this paper.

2 Related Works

This work is not the first of its kind in addressing the issue of RIA model construction and model checking. Duda et al. [15,18,23] uses Breadth-First search crawling strategy to crawl RIAs. Crawljax [24,25] leans toward Depth-First search strategy. Other works aim at constructing the FSM model of the application [1–3,22].

Model-based crawling is another area of research that gained momentum in recent years. Benjamin et al. [5,13] present they hypercube model that assumes the target application is a hypercube. Choudhary et al. [9,10] introduce Menu model that assumes events reach the same target state, irrelevant of the source state. Greedy strategy was explored by Peng et al. [28]; and Milani Fard and

Mesbah [26] in a tool called FeedEx. An empirical comparison of different crawling strategies is done by Dincturk et al. [12, 13].

Parallel crawling of traditional web applications has been explored extensively in the literature [6–8, 16, 17, 19, 20, 29, 30]. Parallel crawling of RIAs however is a new field and the only work we know of is Dist-RIA Crawler [27]. Dist-RIA Crawler performs a breath-first search over multiple independent nodes. This paper adds a load-balancing algorithm to the breath-first search. It also works on the superior and more efficient greedy algorithm.

A close topic to Model-based crawling is DOM equivalency. Duda et al. [15, 18, 23] used equality of DOMs to measure their equivalency. Crawljax [24, 25] uses edit distance to do so. Amalfitano et al. [2] compares the two DOMs based on the elements in them. *Imagen* [21] takes into account JavaScript functions closure, event listeners and HTML5 elements as well in identifying the state of the application. In this paper an DOM equality, the most strict form of DOM equivalency, is used.

3 Load-Balancing

The following notations are used in this section and the rest of the paper:

- s : Refers to an application state.
- e : Refers to an event.
- S : The total number of application states.
- E_s : The number of events in the application state s .
- E : Sum of the number of events in all application states.
- N : Number of crawler nodes.
- i : A unique identification number of a node, where $1 \leq i \leq N$.

As described earlier, in each state, Dist-RIA Crawler assigns equal shares of work to the nodes. The load-balancing algorithm presented in this section, referred to as *adaptive* approach, adjusts the portion of events assigned to each node as the crawling proceeds. The manipulation of the portion assigned to the nodes is used as a tool to reduce the workload of the overloaded nodes, and increase the workload of the idle nodes. One of the nodes, called *coordinator*, calculates the portion of the events to be assigned to each node at the time of state discovery. Tasks are not assigned equally, but assigned based on the perceived computational speed of the node and its current workload.

The purpose of the assignment is to drive all nodes to finish together. The portion of events in state s that belong to node i is represented by $P_{s,i}$ where $P_{s,i} \in [0, 1]$. The coordinator uses the assignment of tasks to different nodes as a means to increase the chance of all nodes to finish together, and no node becomes a bottleneck. To achieve this goal, for every node i , the coordinator uses the number of events executed so far by the node (called ET_i) to calculate the execution speed of the node. This execution speed is used to forecast the execution rate of the node in the future. Based on the calculated speed for all nodes, and the given remaining workload of each node, the coordinator decides the portion of the tasks that are assigned to each node.

3.1 Adaptive Load-Balancing Algorithm

Assume that a new state s is discovered at time t . The coordinator calculates v_i , the speed of node i , as:

$$v_i = ET_i/t \quad (1)$$

where ET_i is the number of events executed by node i so far. The remaining workload of node i can be calculated as the difference between the number of assigned events (called AT_i) and the number of executed events ET_i . Based on the calculated speed v_i , the coordinator calculates the time it takes for node i to finish execution of remaining events assigned to it. This time to completion is represented by TC_i and is calculated as follow:

$$TC_i = \frac{AT_i - ET_i}{v_i} \quad (2)$$

After the coordinator distributes the new events of a newly discovered state s among the nodes, the time to complete all events will change. Assuming node i will continue executing events at rate v_i , the new estimation for the time to finish, called TC'_i , is:

$$TC'_i = TC_i + \frac{P_{s,i} \times E_s}{v_i} \quad (3)$$

To drive all nodes to finish together, the coordinator seeks to make TC' equal for all nodes. That is, it seeks to make the following equation valid:

$$TC'_1 = TC'_2 = \dots = TC'_N \quad (4)$$

Equation 4 can be re-written using Eq. 3:

$$TC_1 + \frac{P_{s,1} \times E_s}{v_1} = TC_2 + \frac{P_{s,2} \times E_s}{v_2} = \dots = TC_N + \frac{P_{s,N} \times E_s}{v_N} \quad (5)$$

Let us take the first two expressions and re-write then:

$$TC_1 + \frac{P_{s,1} \times E_s}{v_1} = TC_2 + \frac{P_{s,2} \times E_s}{v_2} \quad (6a)$$

$$\Rightarrow \frac{(TC_1 + \frac{P_{s,1} \times E_s}{v_1} - TC_2) \times v_2}{E_s} = P_{s,2} \quad (6b)$$

Similarly $P_{s,2}$, $P_{s,3}$, ... and $P_{s,N}$ can all be expressed as follow:

$$\forall i : 2 \leq i \leq N : P_{s,i} = \frac{(TC_1 + \frac{(P_{s,1} \times E_s)}{v_1}) - TC_i}{E_s} \times v_i \quad (7)$$

The coordinator intends to assign all of the events in the newly discovered states to the nodes. Thus the sum of all P_s for state s is 1. Therefore:

$$1 = \sum_{i=1}^N P_{s,i} \quad (8)$$

By expanding $P_{s,2}$, $P_{s,3}$, \dots and $P_{s,N}$ in Eq. 8, using Eq. 7, we get:

$$1 = P_{s,1} + \sum_{i=2}^N \frac{(TC_1 + \frac{(P_{s,1} \times E_s)}{v_1}) - TC_i}{E_s} \times v_i \quad (9a)$$

$$\Rightarrow P_{s,1} = \frac{1 - \sum_{i=2}^N \frac{(TC_1 - TC_i) \times v_i}{E_s}}{1 + \frac{E_s}{v_1 \times E_s} \times \sum_{i=2}^N v_i} \quad (9b)$$

Given the value of $P_{s,1}$ using Eq. 9, the value of $P_{s,2}$, $P_{s,3}$, \dots and $P_{s,N}$ can easily be calculated using Eq. 7.

The adaptive approach does not guarantee that all nodes finish together. The assignment eliminates bottlenecks only if there are enough events in a newly discovered state s to rescue every bottlenecked node. In other words, if there are not enough events in s , and the workload gap between the nodes is large, the adaptive approach fails to assign enough jobs to all idle nodes and make them busy so that all nodes finish together.

4 Overview of the GDist-RIA Crawler

This section describes the crawling algorithm that the GDist-RIA crawler uses.

4.1 Design Assumptions

The GDist-RIA Crawler makes the following assumptions:

- **Reliability:** Reliability of nodes and communication channels is assumed. It is also assumed that each node has a reachable IP address.
- **Network Bandwidth:** It is assumed that the crawling nodes and the coordinator can communicate at a high speed. This makes the network delay intangible. Note that there is no assumption made about the network delay between the server or servers hosting the target application and the crawling nodes.
- **Target RIA:** The GDist-RIA Crawler only targets deterministic finite RIAs. More formally, the GDist-RIA Crawler assumes that visiting a URL always leads to the same STATE; and from a given STATE, execution of a specific JavaScript event always leads to the same target STATE.

4.2 Algorithm

The GDist-RIA Crawler consists of multiple nodes. The nodes do not share memory and work independently of each other. Nodes communicate with the coordinator using a client-server architecture. Nodes start by contacting the coordinator for the seed URL. After loading the seed URL (i.e. the URL to reach the

starting state of the RIA), and after executing any path of events, a node sends the hash of the serialized DOM (henceforth referred to as the ID of s), as well as E_s to the coordinator.

In response, the coordinator who has the knowledge of the application graph calculates the closest application state to s with an unexecuted event and sends a chain of events that lead to that state back to the probing node. This path may start with a reset order by visiting the seed URL. In addition, the coordinator sends the index of the un-executed event in the target state to the probing node.

The probing node executes the assigned event and sends the transition to the coordinator. The coordinator again runs the greedy search algorithm and responds to the client with a new chain of events. This process continues until all the events in all the application states are executed. If at any point the coordinator realizes that there is no path from the state of the probing node to a state with unexecuted events, it orders the node to reset. In effect, by resetting the node jumps back to the seed URL. Since all application states are reachable from the seed URL, the node will find events to execute after the reset.

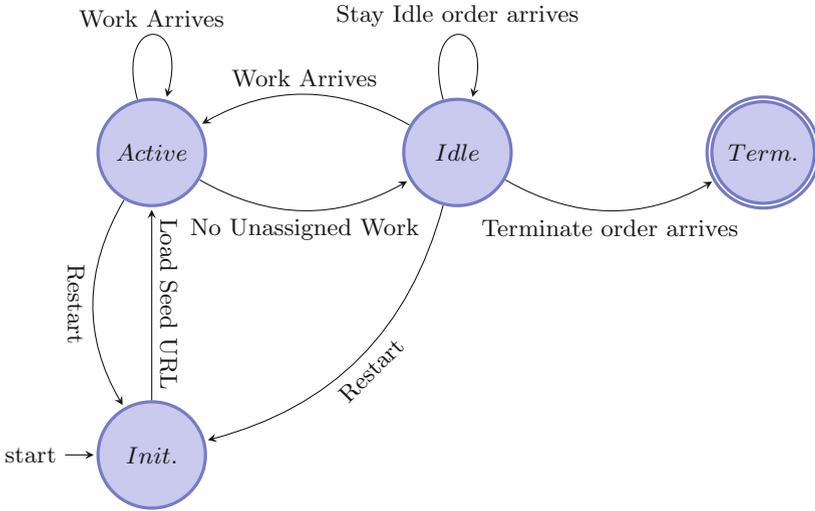


Fig. 1. The node status state diagram.

Figure 1 shows the node state diagram of a crawler node. The crawler starts in the *Initial* state. In this state, the crawler starts up a headless browser process. It then loads the seed URL in the headless browser and goes into the *Active* state. Crawling work happens in the *Active* state. After finishing the assigned task, the node goes to the *Idle* state. The node stays in the *Idle* state until either more work becomes available or a termination order from the coordinator marks the end of the crawl. During *Active* and *Idle* states, the coordinator may order the node to restart so it can reach states that are unreachable from the current state of the node.

4.3 Termination

When the following two conditions are met the coordinator initiates the termination protocol by sending all nodes a *Terminate* order:

- All nodes are all in Idle state.
- There is no Unassigned work in the coordinator i.e. all events in the discovered states are assigned to the nodes.

5 Implementation

To ensure that the proposed algorithm is practical a prototype of the system was implemented. This section explains some of the technical challenges in implementing the prototype of the GDist-RIA crawler.

5.1 Running a Virtual Headless Browser

The GDist-RIA Crawler uses an engine, called *JS-Engine*, to handle web client events¹. The primary task of JS-Engine is to execute JavaScript events and it uses *PhantomJS*², an open source headless WebKit, to emulate a browser with the capability to do so.

Due to the asynchronous nature of the JavaScript, the crawler can not simply trigger an event and consider the execution finished when the call returns. Executing an event in JavaScript may trigger an asynchronous call to the server, or schedule an event to happen in the future. When these events happen the state of the application may change. More formally, two main types of the events that may have dormant ramifications include: *Asynchronous calls* and *Clock events*.

Upon triggering an event on the target application, the JS-Engine waits until the event and all its ramifications are over. For this to happen successfully, the JS-Engine requires a mechanism to keep track of all asynchronous calls in progress and wait for their completion before continuing. Unfortunately, JavaScript does not offer a method to keep track of AJAX calls in progress. Thus the JS-Engine redefines *send* and *onreadystatechange* methods of *XMLHttpRequest* object, the native JavaScript object responsible for performing asynchronous requests, such that the target web application notifies the crawler application automatically upon start and finish of every asynchronous call (Listing 1.1)^{3,4}.

¹ This paper only focuses on JavaScript events and leaves other client side events such as Flash events to the future studies.

² <http://phantomjs.org/>

³ *XMLHttpRequest* is the module responsible for asynchronous calls in many popular browsers such as Firefox and Chrome. Microsoft Internet Explorer however does not use module, and instead it uses *ActiveXObject*.

⁴ Due to space limitation rest of code snippets in this section are omitted.

Listing 1.1. Hijacking Asynchronous Calls

```
XMLHttpRequest.prototype.sendOriginal = XMLHttpRequest.prototype.send;
XMLHttpRequest.prototype.send = function (x){
    cvar onreadystatechangeOriginal = this.onreadystatechange;
    this.onreadystatechange = function(){
        onreadystatechangeOriginal(this);
        parent.ajaxFinishNotification();
    }
    parent.ajaxStartNotification();
    this.sendOrig(x);
};
```

The second source of asynchronous behaviour of a RIA with respect to the time comes from executing clock functions, such as *setTimeout*. This methods is used to trigger an event in the future. In many cases, such events can help animating the website, and adding fade-in fade-out effects. Knowledge of the existence of such dormant functions may be necessary to the JS-Engine. Similar to the asynchronous events, JavaScript does not offer a method to keep track of the time events. Thus JS-Engine re-defines *setTimeout* to hijack time events.

JS-Engine needs to identify the user interface events (i.e. the events that can be triggered by the user interacting with the interface) in the page. Events that leave a footprint in the DOM are easy to detect: Traversing the DOM and inspecting each element can find these events. Attached events using *addEventListener*, however, do not reflect themselves on the DOM. The final challenge faced by JS-Engine is to detect these client-side events attached through event listeners.

These events are added through a call made to *addEventListener* and are removed through a call made to *removeEventListener*. To handle event listeners, JS-Engine redefines *addEventListener* and *removeEventListener* methods such that whenever a call is made to *addEventListener* an entry is added to a global object, and when a call is made to *removeEventListener* the corresponding element is removed. Hence at any given point, JS-Engine can simply check the contents of this object to get elements with attached events.

6 Evaluation

The coordinator prototype is implemented in PHP 5.3.10 and MySQL 14.14. The coordinator contacts the node using SSH channel. The nodes are implemented using PhantomJS 1.9.2, and they contact the coordinator through HTTP. The coordinator as well as the nodes are hosted on a Linux[®] Kernel 3.8.0 operating system with an Intel[®] Intel[®] Core(TM)2 Duo CPU E8400 @ 3.00 GHz and 3 GB of RAM. The communication happens over a 10 Gbps network.

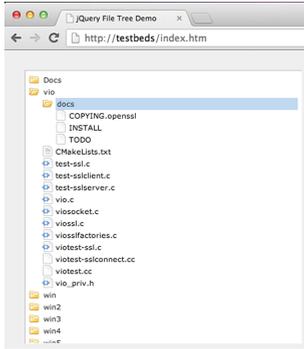


Fig. 2. File tree browser RIA screen-shot

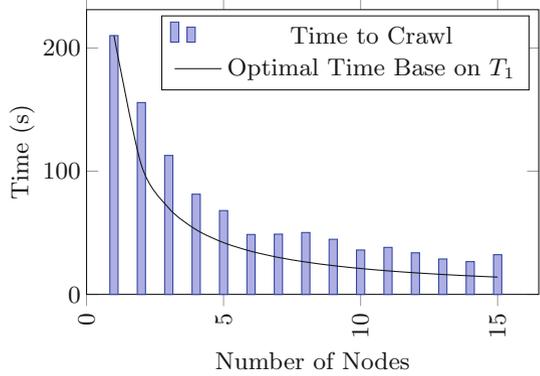


Fig. 3. The total time to crawl the target RIA with multiple nodes.

6.1 Testbed

To measure the performance of the crawler in practice a jQuery based RIA called *jQuery file tree*⁵ was chosen. This open source library creates a web interface that allows the user to brows a set of files and directories through a browser. Similar to most file browsers, directories can be expanded and collapsed, leading to the new client side states. Expanding a directory triggers an asynchronous call to the server to retrieve the contents of that directory. Figure 2 shows a picture of a jQuery file tree application.

6.2 Results

To capture the performance of the algorithm as the number of nodes increases, we crawled the target RIA with different number of nodes, from 1 node to 15 nodes.

Figure 3 shows the total time it takes to crawl the RIA as the number of nodes increase (the bar chart) and compares it with the theoretical optimal time to crawl the RIA with multiple nodes (the line chart). The theoretical optimal time it calculated by taking the time it takes to crawl the RIA with one node (T_1), and divide the number by the number of nodes used by the crawler. This theoretical number serves as a base line to measure the efficiency of the crawler. A the figure shows, a good speedup is achieved as the number of nodes increases. The best performance is achieved with 14 nodes.

The performance of the crawler in Fig. 3 is better described by breaking down the time into most time consuming operations. Box plots in Figs. 4, 5, 6 and 7 show this break down:

- Figure 4: This plot shows the time it takes to load the seed URL into JS-Engine. This plot is interesting in that, this operation is the only operation

⁵ <http://www.abeautifulsite.net/blog/2008/03/jquery-file-tree/>

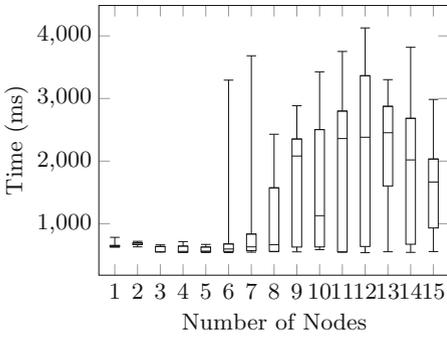


Fig. 4. Time to load the seed URL into JS-Engine.

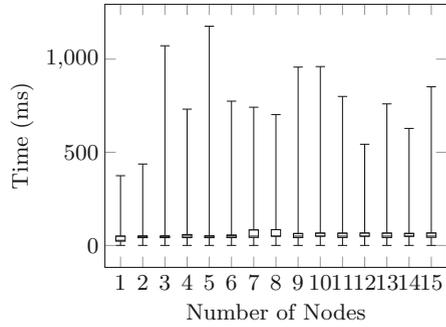


Fig. 5. Time to update application graph.

that gets more expensive as the number of crawlers increase. Compared to normal asynchronous calls, the seed URL contains large files and libraries. As the number of crawling nodes increase, the host server disk operation becomes a bottleneck and a jump is observed around node 6.

- **Figure 5:** This plot shows the time it takes for the coordinator to maintain and update the application graph. This includes adding new states and transitions to the application graph stored in the MySQL database. As expected, this operation is impacted by the number of crawlers.
- **Figure 6:** This plot shows the time it takes for the coordinator to calculate the closest state from the state of the probing node with un-executed events in it. The time to do this calculation does not vary much and it is often close to 50 ms. The calculation itself is rather fast, and the majority of the 50 ms is spent on retrieving the application graph from the database and constructing the auxiliary structures in the memory. As expected, the figure shows that the measured values are independent of the number of crawlers and are not impacted by it.
- **Figure 7:** Finally this plot shows the time it takes to execute a single JavaScript event. Based on our calculations, executing JavaScript events is fairly fast when there is no asynchronous call to the server. Asynchronous calls make event execution time substantially longer, and often increase the execution time by two orders of magnitude. At the scale we ran the experiments, the application server is not bottlenecked by executing JavaScript events. Eventually as the number of nodes increases, the application server will become a bottleneck and the time it takes to execute asynchronous requests rises.

6.3 Discussion

From the presented break down, it is obvious that the most time consuming operation is loading the seed URL into the JS-Engine. The second most time consuming operation that happens frequently is executing JavaScript events.

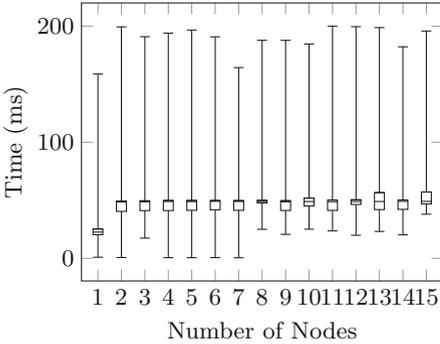


Fig. 6. Time to calculate the next task using the greedy algorithm.

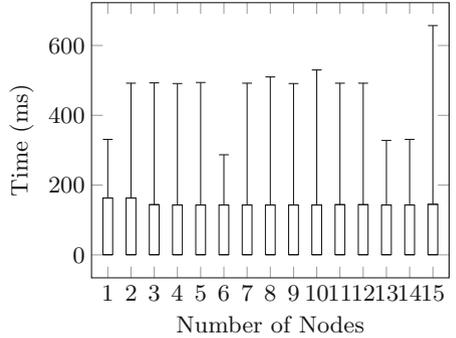


Fig. 7. Time to execute JavaScript events.

Executing a JavaScript event can be particularly time consuming if it involves an asynchronous call to the server.

The design decision of performing the greedy algorithm in a centralized location is inspired by the large discrepancy in the time it takes to find the path greedily and the time it takes to execute the path. As the experiments presented suggests, executing a single asynchronous event can take an order of magnitude longer than calculating the entire shortest path.

At the scale presented in this paper, the coordinator is far from being a bottleneck. As the number of crawling nodes increases, however, the coordinator is bound to become one. In Dist-RIA Crawler [27] nodes uses a deterministic algorithm to autonomously partition the search space and execute JavaScript events in the application. As a future improvement, similar techniques can be used to improve the GDist-RIA crawler by allowing the crawling nodes to autonomously decide (at least partly) the events to execute.

7 Conclusion and Future Improvements

This paper studies distributed crawling of RIAs using a greedy algorithm. A new client-server architecture to dispatch crawling jobs among the crawling nodes, called GDist-RIA Crawler, is introduced. Upon finishing a task, nodes ask the coordinator for the next tasks to do. The coordinator runs the greedy algorithm to assign new task to the probing node, and responds the node with the task. A prototype of the algorithm is implemented and experimental results are provided.

The GDist-RIA Crawler achieves a satisfactory speed up while running the system with up to 15 crawling nodes. This speedup is a result of the low cost of running the greedy search in the application graph at the coordinator, compared to executing the found path by a crawler node. The GDist-RIA Crawler can be improved in many directions, including: Multiple Coordinators to scale better, a peer-to-peer architecture is to shift the greedy algorithm from the coordinator to the crawling nodes, parallelizing other Model-based Crawling strategies (such

as probabilistic model or menu model) [4, 5, 11, 14], and Cloud Computing to be more elastic with respect to the resources available and disappearing resources.

Acknowledgments. This work is largely supported by the IBM[®] Center for Advanced Studies, the IBM Ottawa Lab and the Natural Sciences and Engineering Research Council of Canada (NSERC). A special thank to Sara Baghbanzadeh.

Trademarks

IBM, the IBM logo, ibm.com and AppScan are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at “Copyright and trademark information” at www.ibm.com/legal/copytrade.shtml. Intel, and Intel Xeon are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries. Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both. Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

References

1. Amalfitano, D., Fasolino, A.R., Tramontana, P.: Reverse engineering finite state machines from rich internet applications. In: Proceedings of the 2008 15th Working Conference on Reverse Engineering, WCRE '08, pp. 69–73. IEEE Computer Society, Washington, DC (2008)
2. Amalfitano, D., Fasolino, A.R., Tramontana, P.: Experimenting a reverse engineering technique for modelling the behaviour of rich internet applications. In: IEEE International Conference on Software Maintenance, ICSM 2009, pp. 571–574, September 2009
3. Amalfitano, D., Fasolino, A.R., Tramontana, P.: Rich internet application testing using execution trace data. In: Proceedings of the 2010 Third International Conference on Software Testing, Verification, and Validation Workshops, ICSTW '10, pp. 274–283. IEEE Computer Society, Washington, DC (2010)
4. Benjamin, K., von Bochmann, G., Jourdan, G.-V., Onut, I.-V.: Some modeling challenges when testing rich internet applications for security. In: Proceedings of the 2010 Third International Conference on Software Testing, Verification, and Validation Workshops, ICSTW '10, pp. 403–409. IEEE Computer Society, Washington, DC (2010)
5. Benjamin, K., von Bochmann, G., Dincturk, M.E., Jourdan, G.-V., Onut, I.V.: A strategy for efficient crawling of rich internet applications. In: Auer, S., Díaz, O., Papadopoulos, G.A. (eds.) ICWE 2011. LNCS, vol. 6757, pp. 74–89. Springer, Heidelberg (2011)
6. Boldi, P., Codenotti, B., Santini, M., Vigna, S.: Ubicrawler: a scalable fully distributed web crawler. Proc. Aust. World Wide Web Conf. **34**(8), 711–26 (2002)

7. Boldi, P., Marino, A., Santini, M., Vigna, S.: BUBiNG: massive crawling for the masses. In: WWW (Companion Volume), pp. 227–228 (2014). <http://doi.acm.org/10.1145/2567948.2577304>
8. Brin, S., Page, L.: The anatomy of a large-scale hypertextual web search engine. In: Proceedings of the Seventh International Conference on World Wide Web 7, WWW7, pp. 107–117. Elsevier Science Publishers B.V., Amsterdam (1998)
9. Choudhary, S.: M-crawler: crawling rich internet applications using menu meta-model. Master's thesis, EECS - University of Ottawa (2012). <http://ssrg.site.uottawa.ca/docs/Surya-Thesis.pdf>
10. Choudhary, S., Dincturk, M.E., Mirtaheri, S.M., Jourdan, G.-V., Bochmann, G., Onut, I.V.: Building rich internet applications models: example of a better strategy. In: Daniel, F., Dolog, P., Li, Q. (eds.) ICWE 2013. LNCS, vol. 7977, pp. 291–305. Springer, Heidelberg (2013)
11. Choudhary, S., Dincturk, M.E., von Bochmann, G., Jourdan, G.-V., Onut, I.-V., Ionescu, P.: Solving some modeling challenges when testing rich internet applications for security. In: ICST, pp. 850–857 (2012)
12. Choudhary, S., Dincturk, M.E., Mirtaheri, S.M., von Bochmann, G., Jourdan, G.-V., Onut, I.-V.: Crawling rich internet applications: the state of the art. In: Proceedings of the 2012 Conference of the Center for Advanced Studies on Collaborative Research, CASCON '12, IBM Corp., Riverton (2012)
13. Dincturk, M.E.: Model-based crawling - an approach to design efficient crawling strategies for rich internet applications. Master's thesis, EECS - University of Ottawa (2013). http://ssrg.eecs.uottawa.ca/docs/Dincturk_MustafaEmre_2013_thesis.pdf
14. Dincturk, M.E., Choudhary, S., von Bochmann, G., Jourdan, G.-V., Onut, I.V.: A statistical approach for efficient crawling of rich internet applications. In: Brambilla, M., Tokuda, T., Tolksdorf, R. (eds.) ICWE 2012. LNCS, vol. 7387, pp. 362–9. Springer, Heidelberg (2012)
15. Duda, C., Frey, G., Kossmann, D., Matter, R., Zhou, C.: Ajax crawl: making ajax applications searchable. In: Proceedings of the 2009 IEEE International Conference on Data Engineering, ICDE '09, pp. 78–89. IEEE Computer Society, Washington, DC (2009)
16. Edwards, J., McCurley, K., Tomlin, J.: An adaptive model for optimizing performance of an incremental web crawler (2001)
17. Expoto, J., Macedo, J., Pina, A., Alves, A., Rufino, J.: Geographical partition for distributed web crawling. In: Proceedings of the 2005 workshop on Geographic information retrieval, GIR '05, pp. 55–60. ACM, New York (2005)
18. Frey, G.: Indexing ajax web applications. Master's thesis, ETH Zurich (2007). <http://e-collection.library.ethz.ch/eserv/eth:30111/eth-30111-01.pdf>
19. Heydon, A., Najork, M.: Mercator: a scalable, extensible web crawler. World Wide Web **2**, 219–9 (1999)
20. Li, J., Loo, B., Hellerstein, J., Kaashoek, M., Karger, D., Morris, R.: On the feasibility of peer-to-peer web indexing and search. In: Kaashoek, M.F., Stoica, I. (eds.) IPTPS 2003. LNCS, vol. 2735, pp. 207–15. Springer, Heidelberg (2003)
21. Lo, J., Wohlstadter, E., Mesbah, A.: Imagen: runtime migration of browser sessions for javascript web applications. In: Proceedings of the International World Wide Web Conference (WWW), pp. 815–825. ACM (2013)
22. Marchetto, A., Tonella, P., Ricca, F.: State-based testing of ajax web applications. In: Proceedings of the 2008 International Conference on Software Testing, Verification, and Validation, ICST '08, pp. 121–130. IEEE Computer Society, Washington, DC (2008)

23. Matter, R.: Ajax crawl: making ajax applications searchable. Master's thesis, ETH Zurich (2008). <http://e-collection.library.ethz.ch/eserv/eth:30709/eth-30709-01.pdf>
24. Mesbah, A., Bozdag, E., van Deursen, A.: Crawling ajax by inferring user interface state changes. In: Proceedings of the 2008 Eighth International Conference on Web Engineering, ICWE '08, pages 122–134. IEEE Computer Society, Washington, DC (2008)
25. Mesbah, A., van Deursen, A., Lenselink, S.: Crawling ajax-based web applications through dynamic analysis of user interface state changes. *TWEB* **6**(1), 3 (2012)
26. Fard, A.M., Mesbah, A.: Feedback-directed exploration of web applications to derive test models. In: Proceedings of the 24th IEEE International Symposium on Software Reliability Engineering (ISSRE), p. 10. IEEE Computer Society (2013)
27. Mirtaheri, S.M., Zou, D., Bochmann, G.V., Jourdan, G.-V., Onut, I.V.: Dist-ria crawler: a distributed crawler for rich internet applications. In: Proceedings of the 8th International Conference on P2P, Parallel, Grid, Cloud and Internet Computing (2013)
28. Peng, Z., He, N., Jiang, C., Li, Z., Xu, L., Li, Y., Ren, Y.: Graph-based ajax crawl: Mining data from rich internet applications. In: 2012 International Conference on Computer Science and Electronics Engineering (ICCSEE), vol. 3, pp. 590–594 (2012)
29. Shkapenyuk, V., Suel, T.: Design and implementation of a high-performance distributed web crawler. In: Proceedings of the International Conference on Data, Engineering, pp. 357–368 (2002)
30. tsang Lee, H., Leonard, D., Wang, X., Loguinov, D.: Irlbot: Scaling to 6 billion pages and beyond (2008)