

1 1. INTRODUCTION

2 Mobile devices have become ubiquitous over the last decade. In late 2011, there were an
3 estimated 5.9 billion cellular subscriptions worldwide with a global penetration rate of 87% [1].
4 In the United States, as of December 2012 there were 326.4 million cellular subscriptions, a
5 penetration rate of 102.2% (indicating that some users have multiple subscriptions) [2]. As a
6 result, the population's reliance on such devices for completing everyday tasks has increased.
7 According to a 2009 survey, 82% of Americans never leave their house without their phone, and
8 42% stated "they cannot live without their phone" [3].

9 Since mobile devices aren't restricted to a single location, they are natural tools that
10 travelers can use to access real-time information during a commute. Research has shown that
11 transit riders with access to real-time information wait around 2 minutes less for a bus than those
12 without real-time information [4]. In the same study, riders with real-time information perceived
13 their wait time to be around 30% shorter than riders without real-time information, indicating
14 that real and perceived benefits both exist [4].

15 As is common in the early years of technology growth, mobile device and mobile app
16 evolution has been organic, with a variety of companies and individuals filling the need for
17 various transportation services as they emerge, often using a variety of proprietary data formats.
18 Eventually, however, the issue of interoperability between systems must be addressed. For
19 example, many of today's mobile transit apps developed in one city cannot easily function in
20 another city due to differing real-time data formats. Standards, or formats agreed upon by a
21 group of stakeholders, can potentially benefit the industry by establishing a consensus on
22 communication protocols. This would, for example, enable one city to establish a real-time data
23 feed that is identical to another city, so that mobile apps for first city would be transferrable to
24 the second.

25 When examining data formats in the context of mobile device communication, key
26 differences from desktop computers have been observed in the areas of processing power,
27 communications, and energy constraints [5, 6]. Due to smaller form-factors, mobile devices
28 typically have lower-powered processors that generate less heat, resulting in lower processing
29 performance when compared to a desktop computer. Additionally, mobile devices have a
30 constrained wireless communication channel, both in terms of speed, which is slower than a
31 typical wired broadband connection, and in terms of the amount of allowed data transfer, which
32 is typically paid for by the wireless subscriber based on a tiered data plan (i.e., transferring more
33 data costs more money). Finally, mobile devices have significant energy limitations – they are
34 battery-powered, and all application actions (e.g., program execution, wireless transmissions)
35 consume energy.

36 New data standards should recognize the importance, and limitations, of mobile devices.
37 As of August 2013, the European Committee for Standardization (CEN) is in the process of
38 ratifying the new v2.0 version of the Service Interface for Real Time Information (SIRI)
39 (CEN/TS 15531) standard [7]. SIRI v2.0 includes a new "mobile-friendly" version of this
40 standard, or "SIRI Lite", based on the experience of the Metropolitan Transit Authority (MTA)
41 when implementing the MTA Bus Time real-time transit information system in New York [8].

42 With new "mobile friendly" data formats becoming available, and with the processing
43 capabilities of new smart phones on the market approaching 2.0 GHz and eight-core CPUs, one
44 may think that the burden on app developers for implementing smart and efficient apps is lifted.
45 However, this is far from true.

1 This paper presents a performance evaluation of device-to-server communication using
2 the SIRI data format on a modern smart phone. The results indicate that not only is performance
3 still an issue, but that app developers should carefully consider certain software design choices to
4 avoid exposing mobile users to extensive wait times (e.g., for real-time transit arrival
5 information). The results also demonstrate that information service providers should always
6 offer mobile-friendly interfaces when possible. Optimizations to reduce user wait times are also
7 presented.

8 Fast processing of data from a server is important because it has a direct effect on the
9 amount of time a user has to wait before they can see results fetched from the internet. For
10 example, when a user removes a phone from their pocket, opens the device, and requests updated
11 transit arrival information for their current location, the device must retrieve this information
12 from a server (e.g., a SIRI API). The user will be shown a “Please wait...” screen until the
13 device receives and completely parses a server response. Slower parsing also typically results in
14 longer CPU, screen, and radio usage, which all negatively affect battery life.

15 The benchmarking software used for these experiments is made available as an open-
16 source application so that others can perform their own experiments, and so that app developers
17 can use this library as a foundation for building new applications using the SIRI format [9].

18 **2. DATA COMMUNICATIONS FOR MOBILE DEVICES**

19 Fast processing of data from a server is important because it has a direct effect on the amount of
20 time a user has to wait before they can see results fetched from the internet. For example, when
21 a user removes a phone from their pocket, opens the device, and requests updated transit arrival
22 information for their current location, the device must retrieve this information from a server
23 (e.g., a SIRI API). The user will be shown a “Please wait...” screen until the device receives and
24 completely parses a server response. Slower parsing also typically results in longer CPU, screen,
25 and radio usage, which all negatively affect battery life.

26 When discussing device-to-server communication and performance, there are two typical
27 high-level design choices for a web service, or API:

- 28 1. The protocol used to transfer data
- 29 2. The format of the data being transferred

30 **Protocols**

31 The Hypertext Transfer Protocol (HTTP) underlies most internet communication, and therefore
32 typically plays a role in most API designs [10]. However, a second protocol, SOAP, emerged in
33 the early 2000s to support advanced enterprise server-to-server communication [11]. SOAP uses
34 the Extensible Markup Language (XML) to define protocol properties.

35 While SOAP has proven useful for advanced server-to-server communication, it is not
36 well suited for mobile devices. The extensive use of XML results in a large overhead being
37 added to each message.

38 An alternate, simpler protocol for web services, termed “RESTful” web services, is to
39 rely only on HTTP in a state-less design, which does not require advanced record-keeping on the
40 mobile device or server during communication and does not use XML in the protocol itself.
41 Previous research has shown that RESTful web services have a significantly less impact on
42 mobile devices resources [5, 6] than SOAP-based web services. One of the improvements to the
43 SIRI v2.0 “SIRI Lite” format was the inclusion RESTful query support. Since the benefits of


```
1         <OperatorRef>MTA NYCT</OperatorRef>
2         <OriginRef>MTA NYCT_200001</OriginRef>
3     </MonitoredVehicleJourney>
4 </VehicleActivity>
5 </VehicleMonitoringDelivery>
6 <ServiceDelivery>
7 <Siri>
```

8
9 One of the improvements to the SIRI v2.0 “SIRI Lite” format was the inclusion of non-XML
10 encoding, including JSON.

11 3. METHODOLOGY

12 The Jackson JSON and XML processor [13] was chosen for these experiments because previous
13 tests have shown that it outperforms other parsers [14, 15], it supports both JSON and XML in a
14 single parser, and it is open-source so that the internal functionality can be examined. It should
15 be noted that certain Java libraries had to be modified so that Jackson would function properly on
16 Android, due to differences in Java for Android versus desktop computers [16]. While the
17 results discussed in this paper are specific to tests performed using an Android application on an
18 Android device, the same general constraints and design considerations should also apply to
19 other mobile platforms such as Apple iOS and Windows Phone.

20 The SiriRestClientUI app [17] was used to perform benchmarks of JSON vs. XML
21 parsing using the SiriRestClient library [9] on a Samsung Galaxy S3 SPH-L710 with Android
22 4.1.1, 1.5 GHz dual core processor, 2GB RAM (Power saving mode off). Jackson 2.1.2 with
23 Aalto 0.9.8 was used. The Jackson Internal HTTP connection was used, as well as the
24 ObjectReader for JSON parsing. These tests were performed on the University of South Florida
25 (USF) WiFi network. Results from SpeedTest.NET Android app at time of test were 51,353 kbps
26 down, 49,554 kbps up, and ping of 16ms.
27


1
2

FIGURE 1 - XML vs. JSON Parsing - Elapsed Time from Request to Parsed Response for 50 Requests

3 50 requests were performed back-to-back using the MTA BusTime SIRI StopMonitoring
 4 API. A timestamp recording feature in the SiriRestClientUI app and SiriRestClient library
 5 captured how long a request took, from when the request was issued to when a Siri object
 6 became available from Jackson.

7 4. RESULTS

8 The elapsed time from request to parsed response for 50 sequential requests is shown in FIGURE
 9 1. There is a substantial difference between the "cold start" times (i.e., the first request) for both
 10 JSON and XML. The time for the XML cold start response is almost 18 seconds, over 4 times as
 11 long as the JSON cold start response (approximately 4 seconds). After the cold start, the
 12 differences between the response times for JSON and XML are much smaller.

13 FIGURE 2 shows the summary statistics for XML and JSON parsing time for this test.
 14 JSON outperforms XML in average response time of 401ms, vs. the XML average response time
 15 of 625ms. 95th percentile of elapsed times is closer, with JSON having a 95th percentile of
 16 501ms and XML having a 95th percentile of 626ms. The increase in standard deviation for XML
 17 response time reflects the initial large cold start value that is substantially larger than the
 18 following warm starts. The size of the JSON response was approximately 4KB, with the XML
 19 response being approximately 5KB.

20




FIGURE 2 - XML vs. JSON Parsing - Summary Statistics for the Elapsed Time of 50 Requests

5. DISCUSSION

The first execution of a request to the server (i.e., a cold start) typically takes much longer than subsequent requests (i.e., warm starts). This is because Jackson will dynamically construct Class model from Java class definitions the first time deserializers are needed (typically when the application initiates the request and the `readValue()` method is called). Here, JSON yields significant better performance than XML - JSON performance is over 4 times faster than XML with a time difference of 14 seconds. This is likely due to the additional annotation accesses required for XML, an area in which Android currently performs poorly [18]. However, poor XML performance when compared to JSON is not unique to Android. Other performance evaluations for desktop Java virtual machines have found similar relationships, with JSON processing outperforming XML processing for the same encoded content [19]. This likely indicates that XML processing for cold starts will typically perform worse than JSON processing on any Java virtual machine, and possibly on virtual machines for other programming languages as well.

After this initial cold start, Jackson will typically parse subsequent responses more quickly for both JSON and XML. As stated above, JSON has a slight performance advantage for warm starts too - an average of 224ms faster than XML. Since recent human-computer interaction studies have indicated that users can perceive time differences of 100ms when waiting for a response [20], JSON still yields a noticeable performance increase for warm starts from the user's perspective.

On Android, the warm start state can persist even if the app is closed by the user and re-opened. Android devices will typically keep recently closed apps in memory as a cached background process to enhance future startup performance. When the user "starts" the app, it is actually loading the application from the cached process in memory, which loads all the Jackson data structures needed to perform quick parsing on warm starts without needing to re-initialize them from a cold start state. As a result, when comparing tests, it's important to note that the first

1 execution of the app will typically show significantly worse performance than subsequent warm
 2 starts.

3 While the warm start state provides a significant advantage in response time
 4 performance, it unfortunately cannot be relied upon for consistent performance increases after an
 5 app is started on the device. Android may remove a cached process from memory if the platform
 6 is running low on memory. Given the multitasking that typically occurs on most cell phones,
 7 especially in a scenario where the user is waiting for a bus to arrive, performing tasks such as
 8 checking email, internet browsing, or using social networking apps may result in the real-time
 9 transit app being removed from the app cache. At this point, the app will revert to the cold start
 10 state and there may be a significant delay in retrieving transit data. The size of the
 11 SiriRestClientUI cached process was observed to typically be between 24-31MB, which is large
 12 for a frequently used app. For comparison, on a Samsung Galaxy S3 the following apps had the
 13 following cached sizes - Calendar app = 7.2MB, Clock = 16MB, Google+ = 17MB, Maps =
 14 13MB. Since the largest non-system processes are typically the first targets for process cache
 15 eviction, it is likely that this app would frequently be restored to a cold start state.

16 The above observations lead to the development of a manual caching strategy for Jackson
 17 objects on Android in an attempt to consistently reduce the cold start penalty for response times.
 18 The follow section discusses the results of these Jackson object manual caching optimizations.

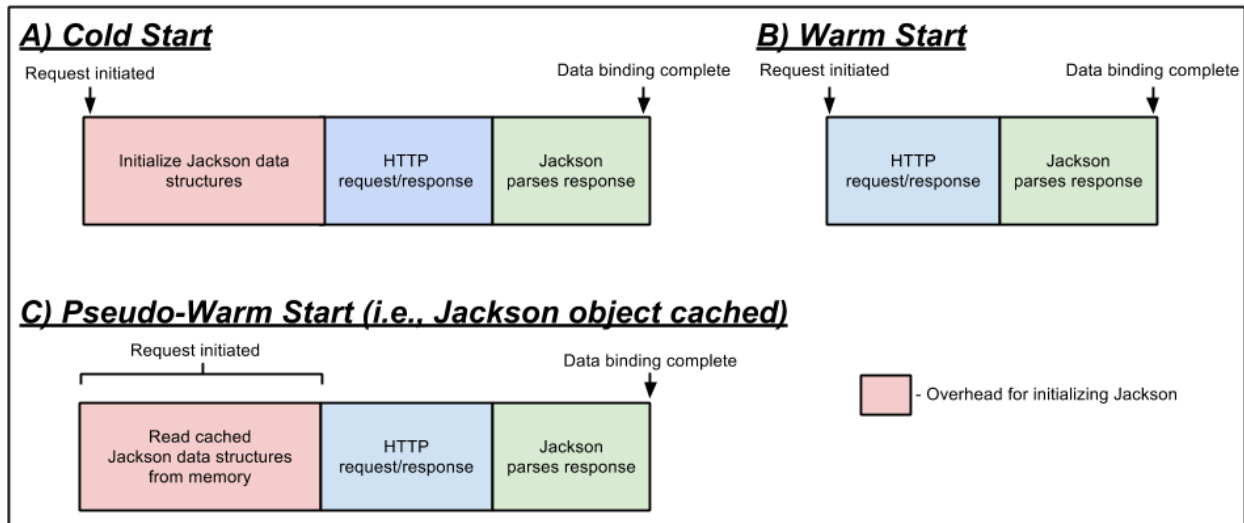
19 **6. OPTIMIZATIONS USING PARSER OBJECT CACHING**

20 To improve cold start performance of retrieving and parsing the SIRI responses using Jackson,
 21 further control over the caching process must be achieved. A strategy called "Pseudo-warm
 22 starts" was developed, which is defined as follows:

- 24 • Pseudo-warm start = An artificial warm-start by the app, where the Jackson object is
 25 manually cached to persistent memory by the app and read from memory during app
 26 startup (in what would otherwise be a cold start).

27 FIGURE 3 shows the stages of application execution for A) cold starts, B) warm starts, and C)
 28 pseudo-warm starts.

29



30
 31

FIGURE 3 - The Stages of App Execution for Cold, Warm, and Pseudo-warm Starts

1 The red-shaded blocks (i.e., the first block in FIGURE 3A and FIGURE 3C) indicate the
2 overhead required to initialize Jackson (note that warm starts, FIGURE 3B, do not have this
3 overhead). In FIGURE 3C, the normal Jackson initialization process is replaced with reading the
4 cached Jackson objects that were initialized in a previous application execution and then written
5 to a persistent cache (i.e., the pseudo-warm start). A note-worthy property of pseudo-warm starts
6 is that the cache read can be initiated on application start-up without requiring user initiation.
7 This flexibility is indicated in FIGURE 3 via the "Request initiated" bracket over the entire cache
8 read period, as the request may be initiated by the user at any point during the cache read. The
9 flexibility of when the cache read is initialized is important in later discussions of this
10 optimization.

11 The goal of the next set of experiments was to evaluate whether the cache read used in
12 pseudo-warm starts was faster than the normal Jackson initialization process in cold starts.

13 **Methodology**

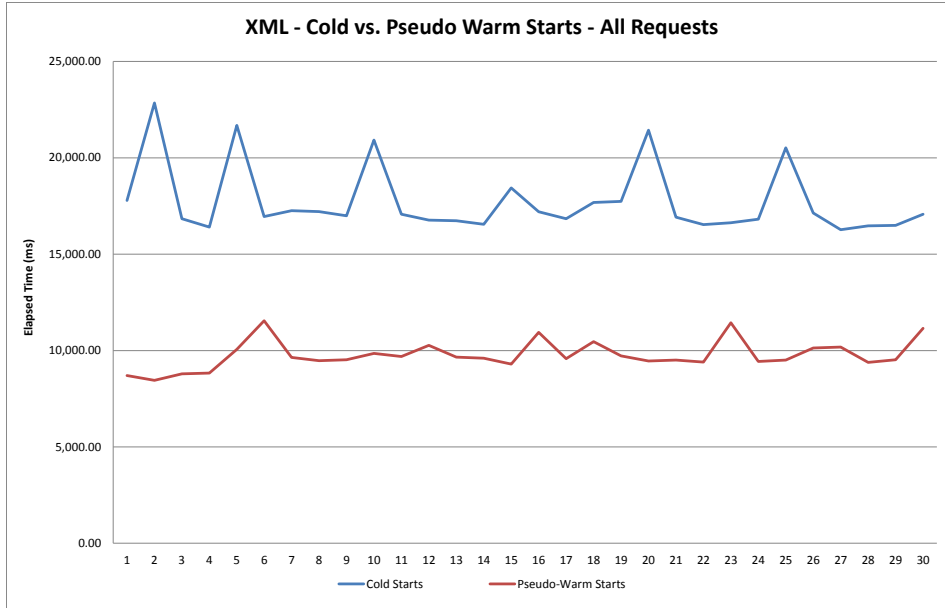
14 The SiriRestClientUI app with the SiriRestClient library was again used to perform benchmarks
15 of JSON vs. XML parsing on the same Samsung Galaxy S3 device and with the same settings as
16 the previous tests. Results from SpeedTest.NET Android app at time of test were 52,176 kbps
17 down, 78,721 kbps up, and ping of 15 ms.

18
19 30 requests were performed using the MTA BusTime SIRI StopMonitoring API. The
20 following steps were repeatedly taken to reset to a cold start state between each cold start test,
21 and to the pseudo-warm start state for the pseudo warm start tests:

- 22 A. A cold start response time was measured,
- 23 B. Caching was turned on in the app and another request was made so the app caches the
24 Jackson object used to make requests to persistent memory - cache write time and cached
25 Jackson object size were measured,
- 26 C. The Android cached process was manually removed from memory via the Application
27 Manager to reset to a cold start state,
- 28 D. A pseudo-warm start response time was measured (= Time to read cached object +
29 elapsed response time).
- 30 E. Caching was turned off in the app, and the Android cached process was manually
31 removed from memory via the Application Manager to reset to a cold start state

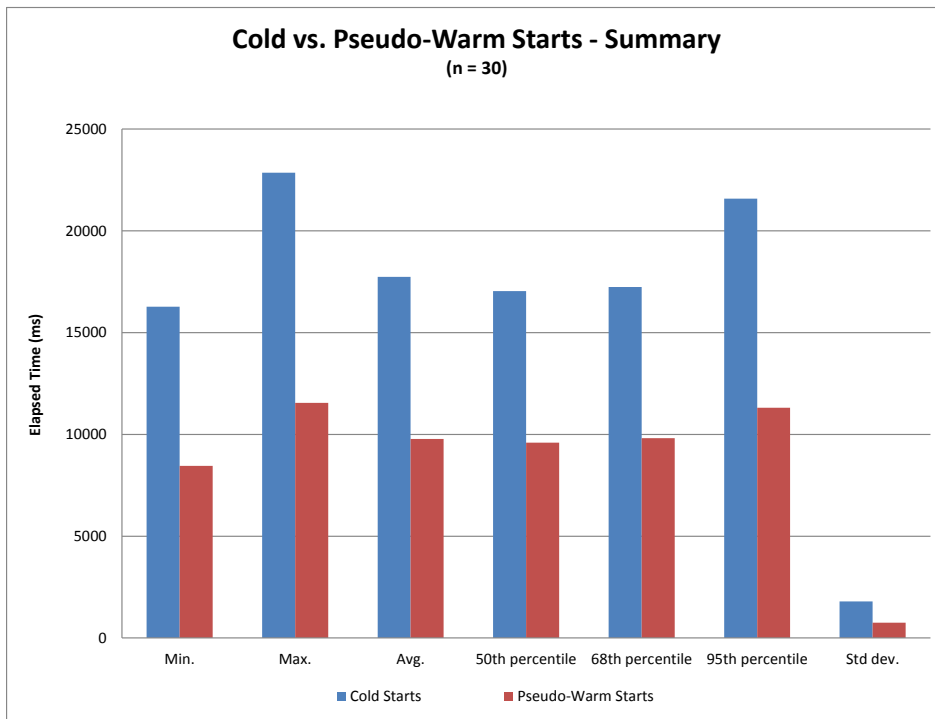
32 **Results**

33 First, the performance of cold vs. pseudo-warm starts when using XML to format the server
34 response was examined. FIGURE 4 shows the elapsed time of 30 cold and pseudo-warm start
35 tests, and FIGURE 5 shows the summary of the results from these tests.
36



1
2

FIGURE 4 - Cold vs. Pseudo-Warm Start performance for XML responses from 30 requests




3
4


FIGURE 5 - Cold vs. Pseudo-Warm Start performance for XML responses - Summary

5 Because of the large cold start penalties, improve the performance of XML parsing via pseudo-
6 warm starts and caching is substantially improved. Cold starts had an average elapsed time of
7 17.7 seconds, and pseudo-warm starts had an average of only 9.7 seconds, an 8 second
8 improvement. Overall, pseudo-warm starts thereby yield an average of a 44% performance
9 increase over cold starts, a dramatic improvement.

1 Next, the performance of cold vs. pseudo-warm starts when using JSON to format the
2 server response was examined. FIGURE 6 shows the elapsed time of 30 cold and pseudo-warm
3 start tests, and FIGURE 7 shows the summary of the results from these tests.
4



5
6 **FIGURE 6 - Cold vs. Pseudo-Warm Start performance for JSON responses from 30 requests**



7
8 **FIGURE 7 - Cold vs. Pseudo-Warm Start performance for JSON responses - Summary Results**

9

1 Pseudo-warm starts perform slightly better than cold starts, with an average elapsed time of
 2 3,785ms vs. the cold start average of 3,963ms - a difference of 178ms on average. This
 3 difference may still be noticeable to the user [20], but is not nearly as dramatic as the
 4 improvement of the XML pseudo-warm start.

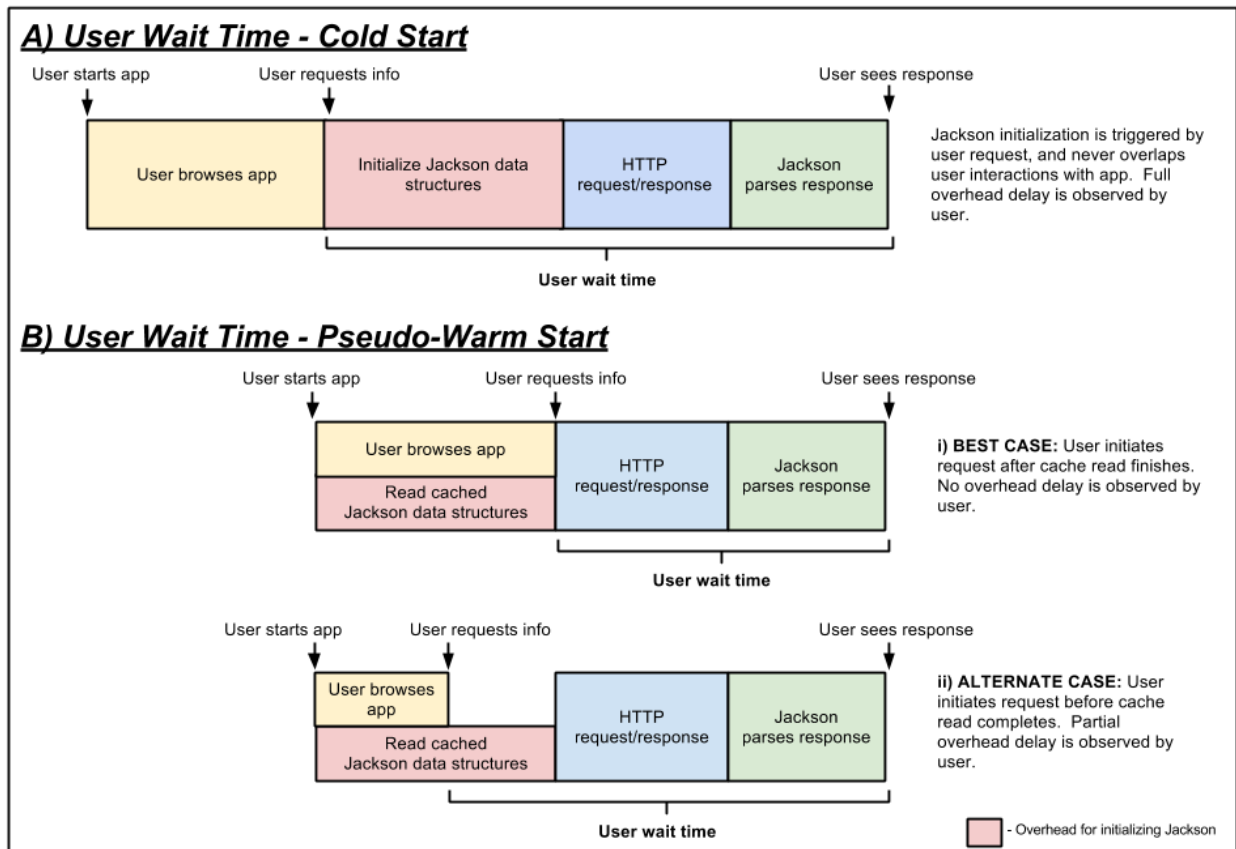
5 Overall, for JSON the pseudo-warm start amounted to an average 3.96% performance
 6 increase over the cold start.

7 **Discussion**

8 The dramatic improvement in XML parsing performance using pseudo-warm starts indicates that
 9 the time required to read a cached Jackson object from mobile device persistent memory is far
 10 less than the time required to newly instantiate that same object. Therefore, a huge 44%
 11 performance increase on average can be produced by using pseudo-warm starts.

12 JSON pseudo-warm start performance improvements (3.96% on average), however, are
 13 not nearly as impressive as their XML counterparts. This difference is partially due to the fact
 14 that XML cold starts are significantly larger than JSON cold starts to begin with (by a factor of
 15 4), which gives the XML pseudo-warm start a much greater margin of potential improvement.

16 After reviewing these results, one may think that a pseudo-warm start implementation for
 17 JSON parsing isn't worth the effort. However, a significant advantage of pseudo-warm starts
 18 over cold starts hasn't yet been discussed - the potential to hide part of the delay from the user by
 19 beginning the cache read when the app is first started.
 20



21
 22
 23

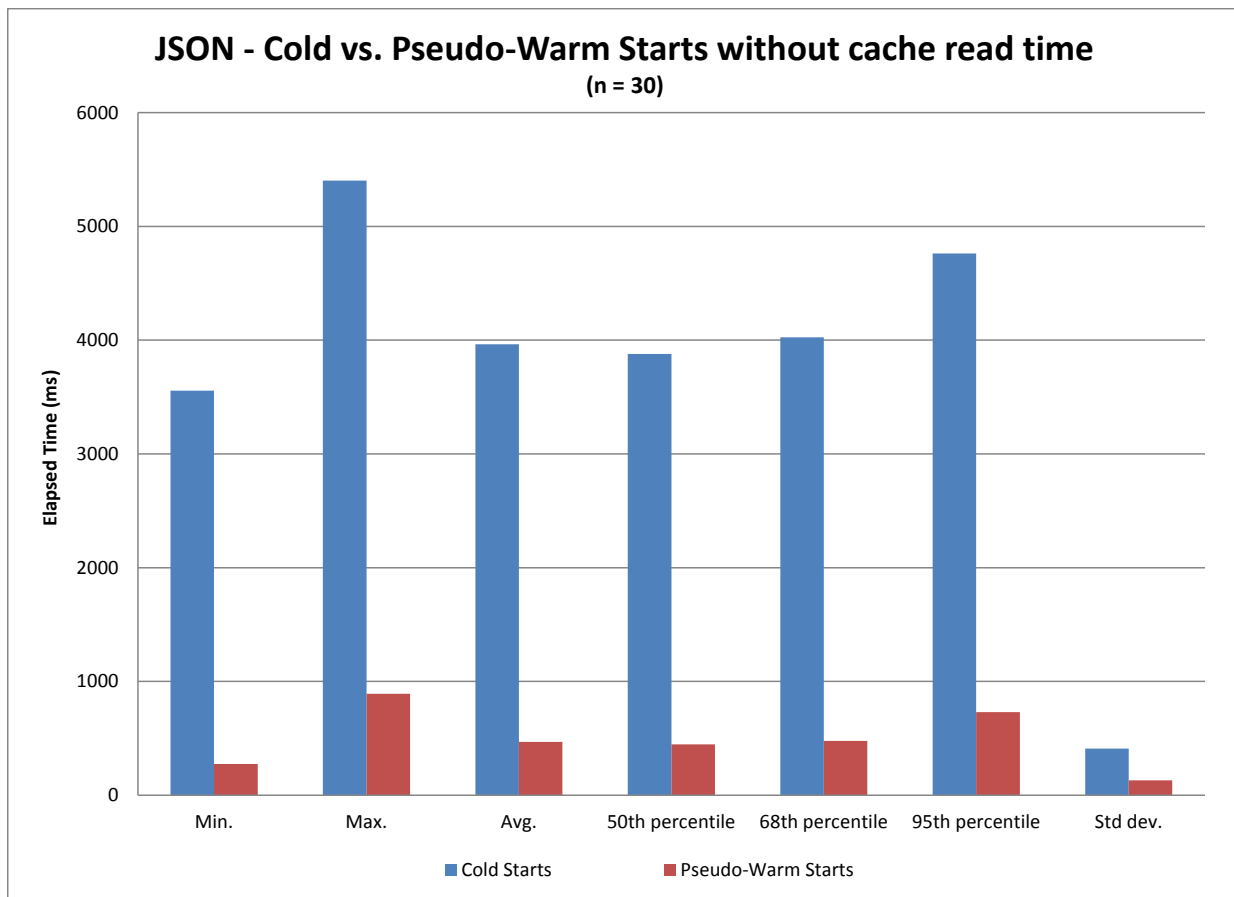
FIGURE 8 - With A) cold starts, the user always observes the fully overhead delay, but with B) pseudo-warm starts the overhead delay can be hidden from the user

1 FIGURE 8 shows how user interactions (shaded in yellow) occur in context of cold starts
 2 (FIGURE 8A) compared with pseudo-warm starts (FIGURE 8B).

3 As mentioned earlier, the cache read used in pseudo-warm starts can be initiated
 4 immediately upon application start up and can execute in the background while the user is
 5 browsing through the application. Therefore, using pseudo-warm starts, a considerable amount of
 6 the overhead time may pass before the user actually triggers a request to the server (e.g., for real-
 7 time transit arrival information). In the best-case scenario, shown in FIGURE 8Bi, the entire
 8 cache read finishes prior to the user initiating a request, and the user does not observe any delay.
 9 In this case, pseudo-warm starts are equivalent to warm starts in performance. Cold starts, on the
 10 other hand, cannot hide any of this overhead wait time from the user - Jackson initialization must
 11 always start when the user initiates a server request, as shown in FIGURE 8A.

12 An alternate pseudo-warm start scenario is shown in FIGURE 8Bii. Here, the user
 13 browses the app while the cache read starts, but initiates the request to the server before the
 14 cache read can complete. In this situation, the user would observe a partial delay, depending on
 15 how much time has elapsed and how long the cache read will take to complete.

16 When considering that the cache read time can now be partially or completely hidden
 17 from the user, the seemingly small improvements in JSON pseudo-warm start performance
 18 suddenly become quite large. FIGURE 9 shows the differences between pseudo-warm starts and
 19 cold starts if the user spends the entire cache read time browsing through the app.
 20



21
 22 **FIGURE 9 - If user interactions with the app hide the cache read latency, JSON pseudo-warm starts perform**
 23 **significantly better than cold starts**

1 Assuming that this cache read time is hidden from the user, the average JSON pseudo-warm start
2 elapsed time is 469ms, compared to nearly 4 seconds of cold start time. The average cache read
3 time in this test was 3.3 seconds, meaning the user would need to spend 3.3 seconds in the app
4 before requesting real-time transit information for the cache read to remain completely hidden.
5 Given the initial app start-up process and other activity of the user before they may retrieve real-
6 time transit information, the expectation that at least some of the cache read will be hidden seems
7 reasonable.

8 Revisiting XML pseudo-warm starts while considering the potential to hide the cache
9 read yields even greater performance increases over cold starts. Average pseudo-warm starts
10 improve to 473ms, versus average cold starts of 17.7 seconds. However, one must also consider
11 that to fully hide the cache read for XML parsing, the user would need to spend an average of 9.3
12 seconds in the app before triggering a request to the server. This is a much longer amount of time
13 than that required for JSON, and therefore less of the XML cache read time will be hidden from
14 the user on average. The longer cache read time for XML when compared to JSON may be
15 explained by the larger cache file size for the XML (1,200kB) vs. JSON (260kB).

16 One drawback to the caching strategy of pseudo-warm starts over cold starts is the
17 persistent memory consumed by the cache file. However, the observed sizes of the cache files
18 should be negligible for most smart phones.

19 **7. RELATED WORK**

20 The strategy of pseudo-warm starts is just one potential avenue for improving user experience in
21 mobile applications that use real-time information from a server. This paper focuses on reducing
22 the amount of time necessary to retrieve new real-time information from a server when an update
23 needs to be retrieved. Response caches, available on Android 4.0 and up, can be used to avoid
24 retrieving a full response to the mobile device when no change in the data has occurred since the
25 last request. Compression can also be used to reduce the size of HTTP data transfers (although
26 the computational costs of device-side decompression must also be taken into account). The
27 refresh interval on the device can also be adjusted to avoid querying the server too frequently. As
28 previously mentioned, the protocol used to transfer information also affects performance. Past
29 research has shown that RESTful web services are heavily preferred to SOAP-based web
30 services on mobile devices [5, 6].

31 **8. FUTURE WORK**

32 Another strategy for hiding the latency for cold starts from the user is to initiate a "dummy" read
33 of a small bit of data on startup, which will force initialization of many of the internal Jackson
34 data structures used for deserialization. While the dummy read isn't expected to increase the
35 general initialization performance, it does offer the advantage of being able to hide some of the
36 latency from the user vs. a normal cold start. Future work could compare the performance of
37 dummy reads against cache read times to quantify the differences.

38 The number of cores in mobile processors continues to increase, with eight core
39 processors just around the corner for mobile phones. Further experimentation could also
40 examine the potential for parallelizing computations on Android to speed up both JSON and
41 XML processing. However, such speedups are not expected to effect the relative relationship
42 between the time required for JSON and XML processing when the same general processing
43 model is used for both (which is the case in this paper).

44 Additional future benchmarking could also be performed in a multi-tasking environment
45 where multiple applications are being context-switched as they are brought to the foreground and

1 returned to the background during JSON and XML process. However, the authors do not expect
2 the general performance results presented in this paper to differ drastically in this situation. The
3 parsing and generation of data formats like JSON and XML is a relatively simple linear CPU-
4 intensive (and CPU-bound) process, and as such concurrency issues are rarely problematic when
5 benchmarking this type of performance. As a result, linear single-threaded tests are usually
6 sufficient when performing such benchmarking. Additionally, in many real-world time-sensitive
7 application scenarios (e.g., a user checking the estimated arrival time while waiting at a bus
8 stop), the user is waiting for the result to be processed while the application is in the foreground,
9 in a very similar design to the tests in this paper.

10 Different size responses could also be analyzed in future work, to determine if this would
11 affect the relative relationship between JSON and XML processing time (e.g., if JSON is better
12 for smaller datasets, while XML is better for larger datasets). The authors of this paper believe
13 that the general processing relationships of JSON being faster than XML would continue to hold,
14 no matter the dataset size, as the processing models for JSON and XML are very similar at
15 conceptual and implementation level.

16 Finally, implementation-specific future work could also be examined. Tests presented in
17 this paper used Jackson v2.1.2. Newer versions of Jackson are now available and should be used
18 in future benchmarking tests. Further investigation into the time required to access annotations
19 on Android revealed an issue on the Android platform [18]. It appears that this issue is now
20 fixed in the Android Open-Source Project, but has not yet been included in any Android releases
21 (4.3 and lower) for Android devices. Future benchmarking should be performed on releases of
22 Android with this improvement to see if performance increases – however, as noted earlier, other
23 experiments on desktop computer using Java have yielded similar performance benefits of JSON
24 over XML, and therefore improvements in Android are not expected to drastically affect the
25 overall results of the experiments presented in this paper.

26 9. CONCLUSIONS

27 When using real-time information services, users must often wait for their phone to retrieve the
28 latest real-time information from a server. This paper presented an evaluation of the effect of
29 data formats on the time required for a mobile device to retrieve updated information from a
30 server, in the context of real-time estimated arrival information for public transportation. The
31 results indicate that app developers should carefully consider certain software design choices to
32 avoid exposing mobile users to extensive wait times (e.g., for real-time transit arrival
33 information). The results also demonstrate that information service providers should always
34 offer mobile-friendly data (i.e., RESTful web services with JSON encoding) when possible.

35 JSON was shown to be a preferred data transfer format over XML for mobile devices.
36 Average performance for cold starts (i.e., when the user first starts the mobile app) was over 4
37 times faster for JSON than XML, with an average time difference of 14 seconds. JSON also had
38 a noticeable performance advantage in warm starts, being an average of 224ms faster than XML.

39 An optimization strategy, pseudo-warm starts, was also presented that aims at reducing
40 the large performance penalty of cold starts. Pseudo-warm starts, which use cached Jackson
41 objects instead of re-initializing the objects on app startup, provide a dramatic increase in
42 performance for XML parsing, reducing elapsed parsing time from an average of 17.7 seconds to
43 an average of only 9.7 seconds, an 8 second improvement. Pseudo-warm start improvements for
44 JSON were more modest, with an average elapsed time of 3,785ms vs. the cold start average of
45 3,963ms - a difference of 178ms on average.

1 Finally, the potential impact of pseudo-warm starts vs. cold starts in the context of user-
2 observable performance was discussed. Assuming that the entire cache read time is hidden from
3 the user, the average JSON pseudo-warm parsing start elapsed time improves to 469ms,
4 compared to nearly 4 seconds of cold start time. Similarly, assuming a fully hidden cache read,
5 average XML pseudo-warm starts further improve to 473ms, versus average cold starts of 17.7
6 seconds. However, given that XML cache reads take an average of 9.3 seconds, versus average
7 JSON cache read of 3.3 seconds, it may not be realistic to completely hide XML cache reads
8 from the user.

9 10. ACKNOWLEDGMENTS

10 The research described in this paper was funded by the National Center for Transit Research at
11 the University of South Florida.

12 11. REFERENCES

- 13 [1] International Telecommunications Union (2011). "ICT Facts and Figures - The World in
14 2011." Available at [http://www.itu.int/ITU-](http://www.itu.int/ITU-D/ict/facts/2011/material/ICTFactsFigures2011.pdf)
15 [D/ict/facts/2011/material/ICTFactsFigures2011.pdf](http://www.itu.int/ITU-D/ict/facts/2011/material/ICTFactsFigures2011.pdf)
- 16 [2] CTIA - The Wireless Association. "Wireless Quick Facts - Year-End Figures." Accessed
17 July 31, 2013 from <http://www.ctia.org/advocacy/research/index.cfm/aid/10323>
- 18 [3] Synovate. "Synovate mobile phones survey." Accessed January 31, 2010 from
19 <http://www.synovate.com/insights/infact/issues/200909/>
- 20 [4] Kari Edison Watkins, Brian Ferris, Alan Borning, G. Scott Rutherford, and David Layton
21 (2011), "Where Is My Bus? Impact of mobile real-time information on the perceived and
22 actual wait time of transit riders," *Transportation Research Part A: Policy and Practice*,
23 Vol. 45 pp. 839-848.
- 24 [5] Sean J. Barbeau, Nevine L. Georggi, and Philip L. Winters (2010), "Global Positioning
25 System Integrated with Personalized Real-Time Transit Information from Automatic
26 Vehicle Location," *Transportation Research Record: Journal of the Transportation*
27 *Research Board*, pp. 168-176.
- 28 [6] Sean Barbeau, Rafael Perez, Miguel Labrador, Alfredo Perez, Philip Winters, and Nevine
29 Georggi (2011), "A Location-Aware Framework for Intelligent Real-Time Mobile
30 Applications," *Pervasive Computing, IEEE*, Vol. 10 pp. 58-67.
- 31 [7] Michael Frumin. "CEN TC278 WG3 SG7: SIRI Version 2.0." Accessed July 31, 2013
32 from https://groups.google.com/forum/#!topic/siri-developers/apb20g8_0EI
- 33 [8] Metropolitan Transportation Authority. "MTA Bus Time(R)." Accessed August 1, 2012
34 from <http://bustime.mta.info/>
- 35 [9] University of South Florida. "SIRI REST Client." Accessed July 31, 2013 from
36 <https://github.com/CUTR-at-USF/SiriRestClient/wiki>
- 37 [10] Internet Engineering Task Force, "Request for Comments (RFC) 2616 - Hypertext
38 Transfer Protocol -- HTTP/1.1," ed, 1999.
- 39 [11] World Wide Web Consortium (W3C), "SOAP Version 1.2 Part 1: Messaging Framework
40 (Second Edition)," ed, 2007.
- 41 [12] Ecma International, "Standard ECMA-262 3rd Edition," ed. Geneva, Switzerland: Ecma
42 International,, 2011.
- 43 [13] LLC FasterXML, . "Jackson JSON and XML Processor Wiki." Accessed July 31, 2013
44 from <http://wiki.fasterxml.com/JacksonHome>

- 1 [14] Martin Adamek. "JSON Parsers Performance on Android (With Warmup and Multiple
2 Iterations)." Accessed July 31, 2013 from [http://martinadamek.com/2011/02/04/json-](http://martinadamek.com/2011/02/04/json-parsers-performance-on-android-with-warmup-and-multiple-iterations/)
3 [parsers-performance-on-android-with-warmup-and-multiple-iterations/](http://martinadamek.com/2011/02/04/json-parsers-performance-on-android-with-warmup-and-multiple-iterations/)
- 4 [15] Tatu Saloranta. "JSON data binding performance (again!): Jackson / Google-gson / JSON
5 Tools... and FlexJSON too." Accessed July 31, 2013 from
6 http://www.cowtowncoder.com/blog/archives/2009/12/entry_345.html
- 7 [16] S. Barbeau. "Parsing JSON and XML on Android." Accessed July 31, 2013 from
8 [https://github.com/CUTR-at-USF/SiriRestClient/wiki/Parsing-JSON-and-XML-on-](https://github.com/CUTR-at-USF/SiriRestClient/wiki/Parsing-JSON-and-XML-on-Android)
9 [Android](https://github.com/CUTR-at-USF/SiriRestClient/wiki/Parsing-JSON-and-XML-on-Android)
- 10 [17] University of South Florida. "SIRI Rest Client UI." Accessed July 31, 2013 from
11 <https://github.com/CUTR-at-USF/SiriRestClientUI/wiki>
- 12 [18] Android Open Source Project. "Issue 43827: AnnotationFactory's cache has major
13 performance problems." Accessed July 31, 2013 from
14 <https://code.google.com/p/android/issues/detail?id=43827>
- 15 [19] Eishay Smith. "JVM Serializers." Accessed November 15, 2013 from
16 <https://github.com/eishay/jvm-serializers/wiki>
- 17 [20] R. Jota, A. Ng, P. Dietz, and D. Wigdor, "How Fast is Fast Enough? A Study of the
18 Effects of Latency in Direct-Touch Pointing Tasks," presented at the CHI 2013, Paris,
19 France, 2013.
20
21