Energy Consumption Comparison of Interactive Cloud-Based and Local Applications

Arun Vishwanath, *Member, IEEE*, Fatemeh Jalali, *Student Member, IEEE*, Kerry Hinton,
Tansu Alpcan, *Senior Member, IEEE*, Robert W. A. Ayre, and Rodney S. Tucker, *Life Fellow, IEEE*

Abstract—**Interactive cloud computing and cloud-based applications are a rapidly growing sector of the expanding digital economy because they provide access to advanced computing and storage services via simple, compact personal devices. Recent studies have suggested that processing a task in the cloud is more energy-efficient than processing the same task locally. However, these studies have generally ignored the power consumption of the network and end-user devices when accessing the cloud. In this paper, we develop a power consumption model for interactive cloud applications that includes the power consumption of end-user devices and the influence of the applications on the power consumption of the various network elements along the path between the user and the cloud data centre. As examples, we apply our model to Google Drive and Microsoft Skydrive's Word processing, Presentation and Spreadsheet interactive applications. We demonstrate via extensive packet-level traffic measurements that the volume of traffic generated by a session of the application vastly exceeds the amount of data keyed in by the user. This has important implications on the overall power consumption of the service. We show that using the cloud to perform certain tasks consumes more power (by a Watt to 10 Watts depending on the scenario) than performing the same tasks locally on a low-power consuming computer and a tablet.**

Index Terms—**Interactive cloud-based applications, local processing, energy consumption.**

I. INTRODUCTION

Cloud computing and web-based cloud offerings are hailed as the new wave transforming the IT industry. Enterprise customers and home users are increasingly being offered the opportunity to move from running applications on stand-alone computers to using cloud-based services. As a result, these applications are expected to grow dramatically in the future as more businesses and consumers choose to access applications, documents and content remotely over the Internet [1]–[3].

There are three broad flavours to cloud computing – Infrastructure as a Service (Iaas), Platform as a Service (PaaS) and Software as a Service (SaaS) [4]. This paper focuses on SaaS because a large number of cloud service providers, such as Google, Microsoft and Amazon, promote SaaS products which have the same look-and-feel as desktop applications, to encourage users to make a transition to the cloud.

Cloud services offer numerous benefits in terms of cost, scalability, performance and maintenance. Several recent studies [5]–[7] have suggested that cloud offerings are "green" in the sense that they save energy relative to traditional desktop computing. The rationale for this is that data centres are generally optimised for energy efficiency, and migration of applications to the cloud permits replacing high-power desktop computers by low-power consuming computers such as netbooks and tablets. Further, the compute and storage resources in data centres are often shared by many users, in contrast to a single user running a dedicated desktop computer.

While intuitively reasonable, the above argument ignores two key factors: (1) energy required to transport data between the user and the cloud, and (2) power consumed by the end-user device when accessing the cloud. Although prior work advocates computation offloading [8]–[10], namely techniques to reduce the power consumption of end-user devices (e.g. tablets) when accessing the cloud, it largely ignores the energy consumed for *transporting* data from the end-user device to the cloud and back. Using a network-based model we have shown that as the data rate between the user and the cloud data centre increases, the transport energy becomes a dominant fraction of the total energy consumption of cloud computing, thus reducing the latter's energy efficiency [11].

Numerous interactive cloud-based applications have become available in recent years. Moreover, with the widespread deployment of high-bandwidth 3G/4G wireless networks, the number of mobile cloud users is expected to grow significantly [2], [3]. The large-scale migration to cloud computing makes it important to quantify the traffic and power consumption implications of using interactive cloud-based applications.

The work in this paper extends our earlier work in [11] by constructing a measurement based power consumption model for interactive cloud-based applications. This model includes all components of the interactive cloud service and the measurements expose the fact that the volume of traffic generated during an online session of the application can be as much as a 1000-times larger than the amount of data keyed in by the user. The model is then used to compare the power consumption of three scenarios: (i) Creating, editing and saving documents, presentations and spreadsheets in the cloud, (ii) Creating and editing the applications locally, and then saving the files in the cloud, and (iii) Performing the tasks locally (i.e. the cloud is absent). All the tasks are performed on the same low-power consuming end-user devices.

An important finding of this work is that although migration to the cloud offers significant benefits, performing tasks in the cloud may not always be the most energy efficient way to undertake those tasks. The relative merits of using a cloud service, from the perspective of power consumption,

---

A. Vishwanath is with IBM Research-Australia. The other authors are with the Centre for Energy-Efficient Telecommunications, The University of Melbourne, Australia. Emails: (arunv@ieee.org, {jalalif, k.hinton, t.alpcan, rayre, r.tucker}@unimelb.edu.au).

This submission is an extended version of our paper presented at the IEEE ICC 2013 conference [1].

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/JSAC.2015.2393431, IEEE Journal on Selected Areas in Communications

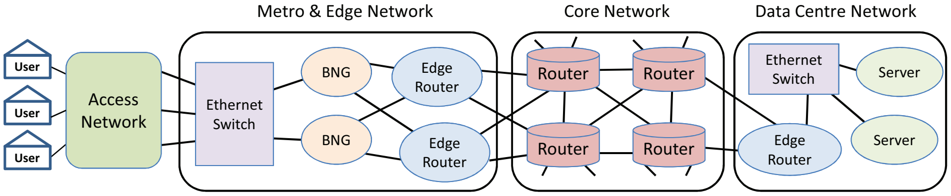

Fig. 1: Topology of the network between a end-user and the cloud data centre.

depends on factors such as the power consumption of the end-user device, access network technology used, computational complexity of the task to be performed, volume of traffic exchanged between the user and the cloud, and factors such as the number of users sharing a compute resource in the cloud.

The rest of this paper is organised as follows. In Section II, we develop a model for quantifying the power consumption per user incurred when using interactive cloud-based applications. In Section III, we report measurements of traffic, in particular the overhead multiplier. We present estimates of power consumption for various network elements in Section IV, and use this to estimate the power consumption per user in Section V. We conclude the paper in Section VI.

II. POWER CONSUMPTION MODEL

We consider a user accessing the cloud via the network topology shown in Fig. 1. The access network includes ADSL Ethernet, WiFi, or in the case of wireless, a 3G/4G (LTE) connection. The metro Ethernet switch aggregates traffic from several users, broadband network gateways (BNGs) regulate access and usage, and edge routers represent the gateway to the global Internet, which consists of many large core routers. Similar architectures have been used in previous studies (e.g. [11], [12]). The data centre network comprises an edge router connecting the data centre to the Internet, aggregation switches and application servers.

The power consumption per user, P_I, of using an interactive cloud-based application is a function of the bit-rate of the application, and the energy per bit incurred by the various network elements shown in Fig. 1, required to deliver the service to the user. This power can be expressed as follows:

$$P_I = P_u + E_a B + (N_c E_c + N_e E_e + E_{bng} + E_{sw})B + E_d B + P_d \tag{1}$$

where P_u is the power consumed by the end-user device to access the interactive cloud application, B is the bit-rate of the application, N_c (N_e) are the number of core (edge) routers along the path between the user and the application server in the data centre, E_c, E_e, E_{bng}, E_{sw} and E_d denote respectively the energy per bit of the core router, edge router, BNG, Ethernet and data centre switches, E_a is the energy per bit of the access network, and P_d is the power consumption per user of the server in the data centre. The power consumption of a server is a function of its CPU utilisation, which is related to the number of processes running on it. This in turn relates to the number of users assigned to that server. We have thus used power per user to model the server power consumption. For network equipment, power consumption is a function of the load [13], i.e. bits per second flowing through it, and is modelled using energy per bit, as described next.

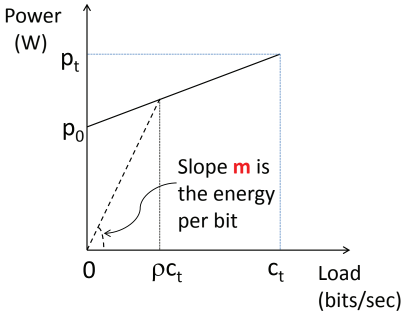

Fig. 2: Power consumption trend of routers and switches.

A. Energy per Bit Modelling

Fig. 2 shows the power consumption of routers and switches in the network (such as in Fig. 1) as a function of the load on the element. This dependence can accurately be modelled using a linear trend [13] as shown by the solid line in Fig. 2. The parameters p_0 and p_t denote a router or switch's idle and maximum power consumption, while c_t denotes the maximum capacity, measured in bits per second. Let ρ denote the average utilisation. Then the energy per bit of that network element, m, is given by $(p_0(1-\rho)+\rho p_t)/\rho c_t$, which is the slope of the dashed line shown in Fig. 2. We estimate this slope for all the routers and switches in Fig. 1 assuming a realistic $\rho = 30\%$ [14], and then apply (1) to estimate the power consumption due to the traffic generated when accessing the cloud application.

B. Power Consumption Measurement

The power consumption of end-user devices when interacting with the cloud (e.g. Google Drive and Microsoft Skydrive) is measured directly using a power meter. In the measurements, we noted that the power consumption of a desktop PC or a high-end laptop was virtually unchanged when interacting with these cloud applications. In order to accurately isolate the power consumption of a end-user device, we used a MSI Wind U100 netbook computer [15] running Windows XP on a 1.6 GHz Intel Atom processor with 2 GB memory. This netbook computer is representative of cloud-ready low-power consuming user devices such as Google Chromebook, which consumes 11 W when awake [16] (similar to the netbook). We also performed measurements using a Samsung tablet [17]. A PowerMate power meter [18] (resolution of 10 mW) was used to record the power consumption of the netbook computer with the battery pack removed at intervals of 1 sec during each session. This enabled us to accurately determine the netbook computer's average power consumption. A custom-built power meter was used to record the power consumption of the tablet.

III. MEASURING CLOUD APPLICATION TRAFFIC

We used the setup shown in Fig. 3 to measure the volume of traffic generated by a session of a cloud application. A

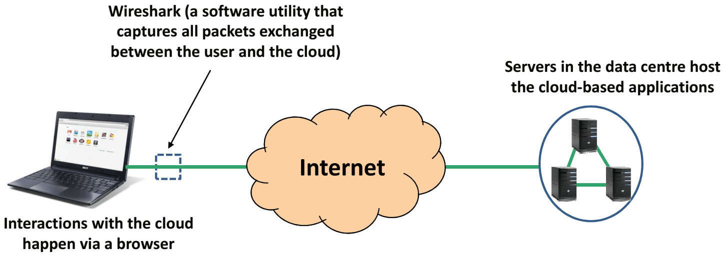

Fig. 3: Measurement setup to capture the volume of traffic generated when accessing cloud-based applications.

packet sniffer software utility (Wireshark [19]), running on the netbook computer captures statistics of all packets exchanged with the cloud server during each session. The file size and the number of key strokes when using the cloud applications were also measured. The applications used for the measurements were office-based applications, owing to their ease of use. The number of characters typed into each application varied from 50 to 500 in steps of 50 characters (equivalently the number of bytes entered varied from 50 to 500 in steps of 50 Bytes). Each session was repeated 10 times to obtain confidence in the results. We automated the typing process using Robosoft record-and-playback software [20]. This enabled us to repeat the experiments consistently across the different applications, ensuring that the typing speed was the same each time; ≈ 57 words per minute (speed of a professional typist).

Traffic measurements for two scenarios are considered, corresponding to how the cloud is used.
(i) Composing and editing Word documents, Presentations and Spreadsheets *online* in Google Drive and Microsoft Skydrive using a web browser (Edit online, Save in the cloud).
(ii) Composing and editing Word documents, Presentations and Spreadsheets *offline* (i.e. locally on the netbook computer), then saving the files in the Google Drive folder on the netbook, and finally synchronizing the folder with the cloud (Edit offline, Save in the cloud).

A. Online Interactive Word processing and Presentation Applications (Edit online, Save in the cloud)

Figures 4 and 5 show the total volume of data traffic (in Bytes) exchanged between the user and the cloud for the online interactive Word processing and Presentation applications from Google and Microsoft. The figures also show the traffic volumes in both the upstream and downstream directions. This data was generated after postprocessing the Wireshark logs. It can be observed from Fig. 4 and Fig. 5 that the total volume of data traffic is substantially larger than the amount of data typed into the application by the user. The overhead multiplier (in terms of the number of bytes) when using Google for both applications is more than a 1000-fold while the overhead multiplier when using Microsoft is 280-fold for Word processing, and 171-fold for Presentation.

B. Online Interactive Spreadsheet Applications (Edit online, Save in the cloud)

The volume of traffic generated by the Spreadsheet application from Google and Microsoft is shown in Figures 6(a) and 6(b). The former generates an overhead multiplier of 650, which is smaller than that of the other two applications, while the latter incurs a substantial overhead; in excess of 9000.

C. Insights into the Traffic Overhead for Online Interactive Applications (Edit online, Save in the cloud)

The Word processing, Presentation and Spreadsheet applications from Google and Microsoft are essentially client-server applications, the browser is the client and the server is accessed via the cloud. Moreover, their look-and-feel, responsiveness and user experience are very similar to that of local stand-alone applications. To support these features, a considerable amount of communication occurs in the background between the browser and server (a brief overview from Google's applications appears in [21]). We noted from the Wireshark logs and while performing the measurements that changes made to the applications were *automatically* saved in the cloud server, thereby ensuring no data loss. Although this provides high service reliability, it incurs a significant traffic overhead.

1) Word processing and Presentation Applications: In the case of Google's Word processing and Presentation applications, logs of the traffic between the user and the data centre show that every key stroke triggers an application synchronisation event between the user and the server. Fig. 7 shows a log excerpt from Wireshark for the Word processing application from Google. A single key pressed at the traffic log time 20.63384 sec is sent as a 1314 Byte TCP packet to the server. This is followed by three (relatively small) packets. The packets are transported using HTTPS making it difficult to decipher their content. The traffic logs indicate that the browser could communicate the key that was typed or deleted (for auto-saving), and the position of the cursor in the browser window to the server as part of every synchronisation event. This occurs whether the event is an insert or delete operation. The synchronisation process ends at time 22.8355 sec at which point the client and server "see" the same document. The next key press event starts at time 25.64 sec and the process repeats.

The behaviour of Microsoft's Document and Presentation applications is similar to that of Google's. However, these applications generate less overhead because the latter typically synchronises with the cloud following every key stroke (as described above), while the former synchronises only when the user pauses or stops typing, as in between words. This results in a smaller volume of traffic exchanged between the user and the cloud server, reducing the traffic overhead.

2) Spreadsheet Applications: The Google Spreadsheet synchronises with the cloud only when the cursor (i.e. focus) shifts from one "cell" in the Spreadsheet to the next. This reduces the frequency of updates, and explains why the overhead (of 650) incurred by Spreadsheet is smaller than that of the other two applications. In the case of Microsoft Skydrive's Spreadsheet application however, we note that the overhead is significantly

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/JSAC.2015.2393431, IEEE Journal on Selected Areas in Communications

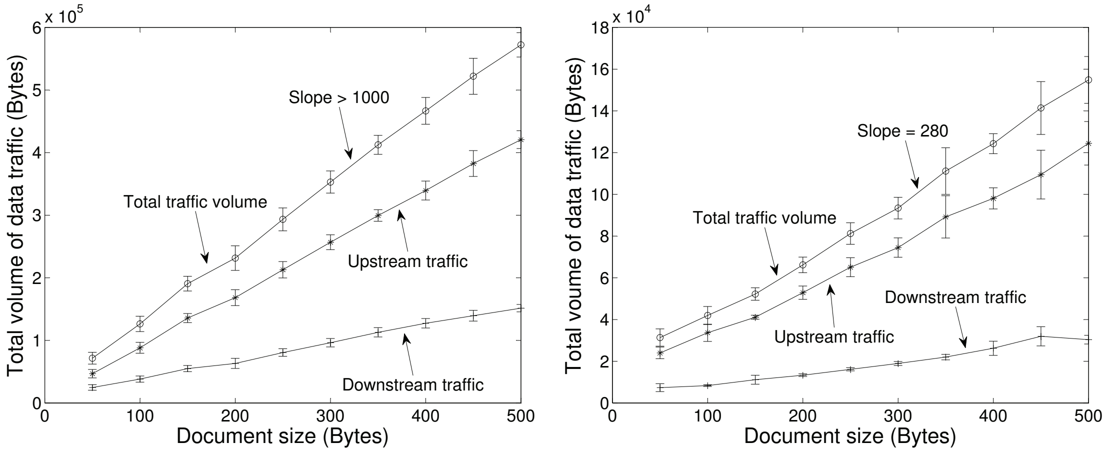

(a) Traffic volume from Google Docs. (b) Traffic volume from Microsoft Docs.

Fig. 4: Volume of traffic generated vs the size of the document for (a) Google Drive and (b) Microsoft Skydrive word processing applications.

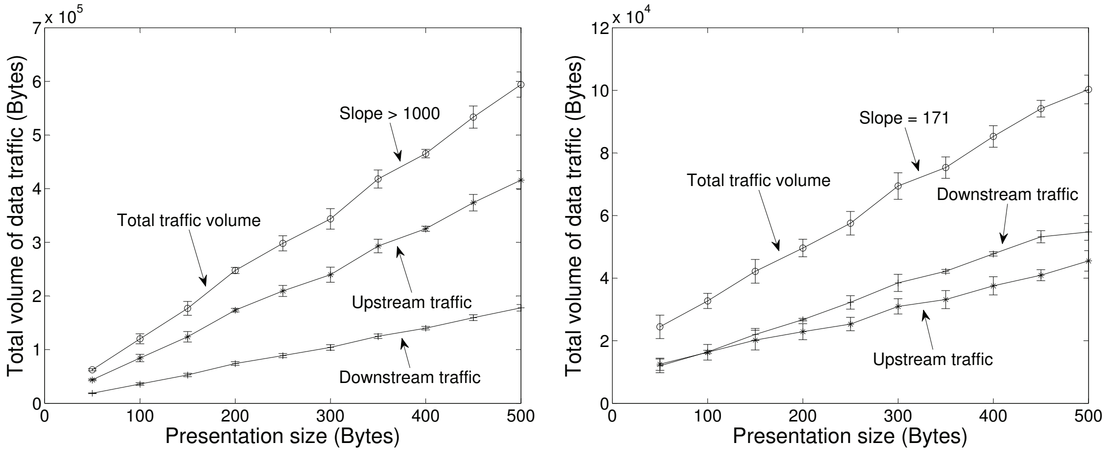

(a) Traffic volume from Google Presentation. (b) Traffic volume from Microsoft Presentation.

Fig. 5: Volume of traffic generated vs the size of the presentation for (a) Google Drive and (b) Microsoft Skydrive presentation applications.

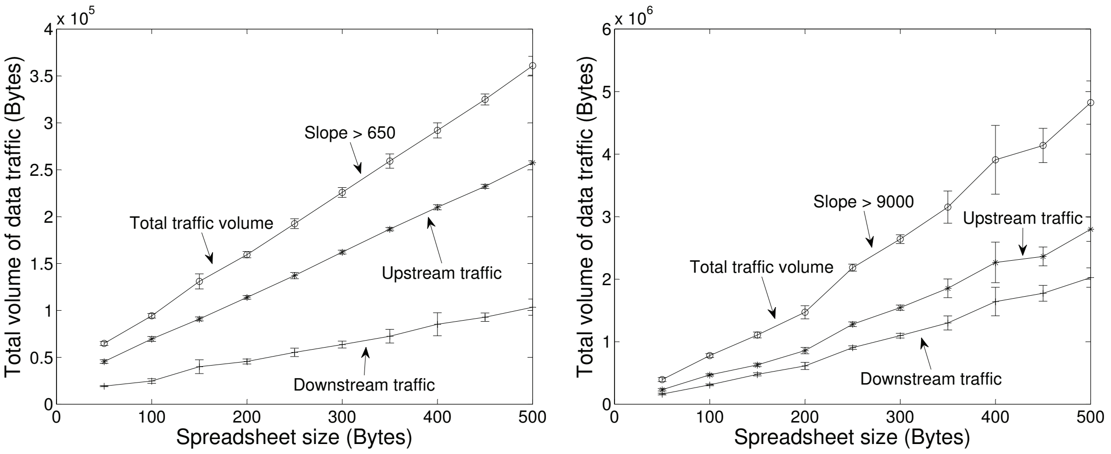

(a) Traffic volume from Google Spreadsheet. (b) Traffic volume from Microsoft Spreadsheet

Fig. 6: Volume of traffic generated vs the size of the spreadsheet for (a) Google Drive and (b) Microsoft Skydrive spreadsheet applications.

larger, as shown in Fig. 6(b). Postprocessing the Wireshark logs revealed that this application generates a large number of TCP sessions and a vast majority of these TCP sessions lasts only a few sec. These sessions handle synchronisation of content with the cloud. For example, it took about 30 sec to enter 50 characters in the Spreadsheet. During this time, there were 20 TCP connections, each lasting on average 4.5 sec. The number of TCP sessions established grew rapidly with the size of the Spreadsheet. Entering 500 characters took 331 sec resulting in 174 TCP sessions, each lasting on average 6.3 sec. We were unable to elicit the content of the sessions because they were encrypted and transported using HTTPS. The traffic logs indicate that the large traffic overhead is associated with establishing/tearing down TCP sessions very frequently and

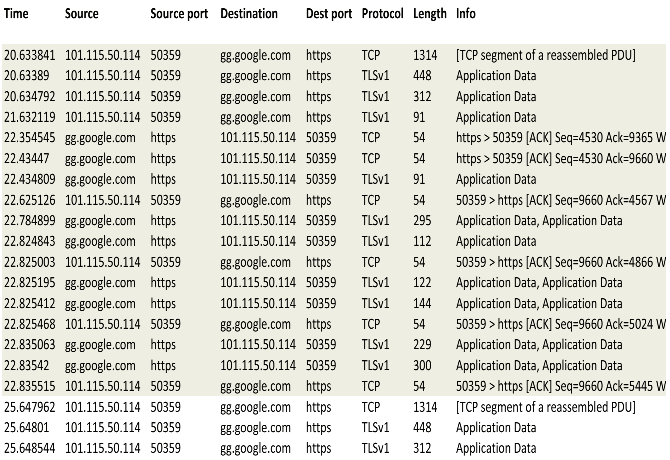

```
Time       Source          Source port  Destination      Dest port  Protocol  Length  Info

20.633841  101.115.50.114  50359        gg.google.com    https      TCP       1314    [TCP segment of a reassembled PDU]
20.63389   101.115.50.114  50359        gg.google.com    https      TLSv1     448     Application Data
20.634792  101.115.50.114  50359        gg.google.com    https      TLSv1     312     Application Data
21.632119  101.115.50.114  50359        gg.google.com    https      TLSv1     91      Application Data
22.354545  gg.google.com   https        101.115.50.114   50359      TCP       54      https > 50359 [ACK] Seq=4530 Ack=9365 W
22.43447   gg.google.com   https        101.115.50.114   50359      TCP       54      https > 50359 [ACK] Seq=4530 Ack=9660 W
22.434809  gg.google.com   https        101.115.50.114   50359      TLSv1     91      Application Data
22.625126  101.115.50.114  50359        gg.google.com    https      TCP       54      50359 > https [ACK] Seq=9660 Ack=4567 W
22.784899  gg.google.com   https        101.115.50.114   50359      TLSv1     295     Application Data, Application Data
22.824843  gg.google.com   https        101.115.50.114   50359      TLSv1     112     Application Data
22.825003  101.115.50.114  50359        gg.google.com    https      TCP       54      50359 > https [ACK] Seq=9660 Ack=4866 W
22.825195  gg.google.com   https        101.115.50.114   50359      TLSv1     122     Application Data, Application Data
22.825412  gg.google.com   https        101.115.50.114   50359      TLSv1     144     Application Data, Application Data
22.825468  101.115.50.114  50359        gg.google.com    https      TCP       54      50359 > https [ACK] Seq=9660 Ack=5024 W
22.835063  gg.google.com   https        101.115.50.114   50359      TLSv1     229     Application Data, Application Data
22.83542   gg.google.com   https        101.115.50.114   50359      TLSv1     300     Application Data, Application Data
22.835515  101.115.50.114  50359        gg.google.com    https      TCP       54      50359 > https [ACK] Seq=9660 Ack=5445 W
25.647962  101.115.50.114  50359        gg.google.com    https      TCP       1314    [TCP segment of a reassembled PDU]
25.64801   101.115.50.114  50359        gg.google.com    https      TLSv1     448     Application Data
25.648544  101.115.50.114  50359        gg.google.com    https      TLSv1     312     Application Data
```

Fig. 7: Wireshark trace following a single key being pressed in Google's interactive cloud-based Word processing application.

the volume of data transported to and from the user per session (tens to hundreds of Kilobytes). This behaviour was not observed with Google Spreadsheet.

The qualitative explanations above are based on observed traffic measurements. A more precise explanation would require an accurate understanding of the way these applications are designed, which remains proprietary. It is evident that the underlying protocols used by the applications to provide a secure and rich user experience involve frequent and encrypted communication of data between the browser and cloud server, giving rise to the large traffic overheads.

D. *Word processing, Presentation and Spreadsheet Applications (Edit offline, Save in the cloud)*

The total volume of data traffic exchanged (in Bytes) between the user and the cloud for editing the Google and Microsoft Word, Presentation and Spreadsheet applications locally and then saving them to the cloud is only marginally greater than the size of the file stored in the hard disk. The observed extra traffic is only due to the added bytes for secure transmission through the Internet, and the number of key strokes used to compose the file does not impact the traffic generated during the upload, i.e. the overhead multiplier, as described above, is absent in this scenario.

IV. Power Consumption of Various Components

In this section we determine values of the various parameters in (1) needed to estimate the power consumption per user, P_I.

A. *Bit-rate Measurements for Interactive Cloud-Based Word processing Applications*

We used the setup shown in Fig. 3 to compose a 2-page document on the cloud. This experiment is representative of a typical instance where a user accesses the cloud to perform a word processing task. The experiment consisted of typing 649 words (4224 characters), inserting a picture, as well as a table comprising 4 rows and 3 columns. Each session on Google and Microsoft lasted on average 12 mins (± 1 sec), and 11 mins and 50 sec (± 10 sec), providing us sufficient data to quantify the bit-rate of the applications. We ran a total of 30 sessions for each application.

As explained previously, we used Wireshark to capture all packets generated during each session. We noted from the logs that the bit-rate – i.e. B in (1) – for the online interactive Word processing application varied between 45 Kbps and 60 Kbps for Google, and between 10 Kbps and 12 Kbps for Microsoft. The bit-rates are not a constant because the applications use TCP, and the performance of TCP varies depending on factors such as link congestion, delay and packet loss.

Identical measurements were conducted to determine the bit-rate of Word processing with Google Drive when the files are edited locally (offline) and then saved to the Google cloud. The bit-rate varied between 1.1 Kbps and 1.5 Kbps.

B. *Bit-rate Measurements for Interactive Cloud-Based Presentation Applications*

Using the automated setup described above, we composed 5 slides each on the two Presentation applications. The experiment consisted of typing 127 words (735 characters), inserting a picture and a table comprising 4 rows and 4 columns. Each session on Google and Microsoft lasted 4 mins and 50 sec (± 2 sec), and 4 mins 57 sec (± 16 sec), respectively. A total of 30 sessions for each application was performed. From the Wireshark logs we noted that the bit-rate B for the Presentation application varied between 37 Kbps and 40 Kbps for the Google application, and between 25 Kbps and 30 Kbps for the Microsoft application.

Again, identical measurements were conducted to determine the bit-rate of Presentation with Google Drive for the case when the files are edited locally and then saved to the Google cloud. The bit-rate varied between 2.5 Kbps and 2.7 Kbps.

C. *Bit-rate Measurements for Interactive Cloud-Based Spreadsheet Applications*

We composed a Spreadsheet by entering numbers along 200 rows and 2 columns. The total number of characters (i.e. digits) was 700. We then performed basic numerical operations such as determining the min, max, mean, median and mode of the numbers. Subsequently, we plotted a (x, y) graph, and noted that the graph was updated dynamically as we sorted the numbers in each of the two columns. We repeated this measurement 30 times for each application. Each session on Google lasted 7 mins and 34 sec (± 2 sec), and each session on Microsoft lasted 9 mins and 8 sec (± 5 sec). The bit-rate B, obtained after postprocessing the Wireshark logs, of Google Spreadsheet varied between 25 Kbps and 30 Kbps, while for Microsoft it varied between 110 Kbps and 150 Kbps.

These measurements were also repeated to quantify the bit-rate of Spreadsheet when the files are edited locally and then saved to the Google Drive cloud. The bit-rate varied between 0.3 Kbps and 0.6 Kbps.

Table I summarises the bit-rates of the different applications as obtained from our measurements. The substantial differences in the bit-rate between edit online and edit offline scenarios is due to the cost of incremental updates of file segments that occurs with the edit online scenario.

D. *Average Power Consumption P_u of the Netbook Computer*

The idle power consumed by the netbook computer with all network interfaces disabled was 10.8 W. We performed

	Application	Bit-rate
Google Drive **Edit online, Save in the cloud**	Word processing	45-60 Kbps
	Presentation	37-40 Kbps
	Spreadsheet	25-30 Kbps
Microsoft Skydrive **Edit online, Save in the cloud**	Word processing	10-12 Kbps
	Presentation	25-30 Kbps
	Spreadsheet	110-150 Kbps
Google Drive **Edit offline, Save in the cloud**	Word processing	1.1-1.5 Kbps
	Presentation	2.5-2.7 Kbps
	Spreadsheet	0.3-0.6 Kbps

TABLE I: Summary of bit-rates for Google and Microsoft Skdyrive's Word processing, Presentation and Spreadsheet applications.

Application	Access network technology	Average power consumed by the Netbook computer
Google Drive Word Processing **Edit online, Save in the cloud**	Ethernet	13.6 W
	WiFi	14.0 W
	4G	16.1 W
Microsoft Skydrive Word Processing **Edit online, Save in the cloud**	Ethernet	14.4 W
	WiFi	14.5 W
	4G	16.7 W
Google Drive Word Processing **Edit offline, Save in the cloud**	Ethernet	13.2 W
	WiFi	13.7 W
	4G	15.1 W

TABLE II: Average power consumed by the netbook computer for using Google and Microsoft's Word processing applications.

Application	Access network technology	Average power consumed by the Netbook computer
Google Drive Presentation **Edit online, Save in the cloud**	Ethernet	14 .0 W
	WiFi	14.2 W
	4G	16.1 W
Microsoft Skydrive Presentation **Edit online, Save in the cloud**	Ethernet	12.8 W
	WiFi	13.0 W
	4G	15.8 W
Google Drive Presentation **Edit offline, Save in the cloud**	Ethernet	13.4 W
	WiFi	13.9 W
	4G	15.3 W

TABLE III: Average power consumed by the netbook computer for using Google and Microsoft's Presentation applications.

Application	Access network technology	Average power consumed by the Netbook computer
Google Drive Spreadsheet **Edit online, Save in the cloud**	Ethernet	16.1 W
	WiFi	16.6 W
	4G	17.8 W
Microsoft Skydrive Spreadsheet **Edit online, Save in the cloud**	Ethernet	14.3 W
	WiFi	14.7 W
	4G	16.2 W
Google Drive Spreadsheet **Edit offline, Save in the cloud**	Ethernet	13.4 W
	WiFi	14.3 W
	4G	15.2 W

TABLE IV: Average power consumed by the netbook computer for using Google and Microsoft's Spreadsheet applications.

Type	Model	Max capacity (c_t) (bidirectional)	Max power (p_t)	Idle power (p_0)	Energy per bit (slope *m*)
Core router	CRS-3	4480 Gbps	12300 W	11070 W	8.5 nJ/bit
Edge router	7609	560 Gbps	4550 W	4095 W	25.2 nJ/bit
BNG	ASR 9010	320 Gbps	1890 W	1701 W	18.3 nJ/bit
Ethernet Switch	Catalyst 6509	256 Gbps	1766 W	1589 W	21.4 nJ/bit
Data Centre Switch	Catalyst 6509	320 Gbps	2020 W	1818 W	19.6 nJ/bit

TABLE V: Energy per bit of equipment in the metro, edge, core and data centre networks of Fig. 1.

experiments at different times during the day (to address the issue of variability in the situations the user may experience) on the interactive cloud applications described in the previous section, and noted that the power consumption of the netbook computer was not sensitive to the time-of-day variation. Measurements were performed using three different access technologies available in the netbook, i.e. Ethernet, WiFi and 4G (via a USB dongle), and the power consumed in each of these cases was recorded.

1) P_u for Word processing Applications: Column three in Table II gives the average power consumed by the netbook, P_u, for composing the 2-pages using Google and Microsoft's Word processing applications. We can see that 13.6 W is consumed when accessing the interactive Word processing application from Google using Ethernet. This increases to 16.1 W when using 4G high-speed wireless technology. A similar trend is observed with the Microsoft application.

2) P_u for Presentation Applications: Table III shows the netbook's average power consumption to access the cloud when composing 5-slides in the Presentation applications. We note that the power consumed by the netbook in this scenario is similar to that for the Word processing applications described above.

3) P_u for Spreadsheet Applications: Table IV shows the power consumption when composing the Spreadsheet. We note that P_u of Google Spreadsheet is greater than 16 W regardless of the type of access technology.

4) Energy per Bit of Routers and Switches: Table. V lists the key network equipment (used in the metro, edge, core and data centre networks) corresponding to Fig. 1. The data was gathered from Cisco's power consumption calculator [22]. Column three represents the maximum capacity (i.e. c_t) of each device, the corresponding maximum power (i.e. p_t) is shown in column four, and the idle power (i.e. p_0), which is typically 90% of the maximum power [23], is denoted in column five. The energy per bit (i.e. slope m) is shown in units of nJ/bit in column six. In the network depicted in Fig. 1, we assume, using the *traceroute* utility, that there are $N_c = 5$ core routers and $N_e = 2$ edge routers on average along the path between the user and the cloud data centre server.

5) Energy per Bit of Access Network: The energy per bit in the case of Ethernet access is approximately 3 nJ/bit; obtained from the data sheet of a Cisco 2960 series switch [24]. The energy per bit for WiFi access is taken to be 128 nJ/bit; obtained from a performance benchmarking study of the Cisco 1250 enterprise WiFi access point [25]. Estimating the energy per bit for a base station is non-trivial since it depends on a variety of different factors such as the number of concurrent users it can support, the deployment area, number of sectors, spectrum allocation, interference, among others. Our energy per bit figures are estimated from [26] by observing that a state-of-the-art 2012-technology 3-sector 2x2 MIMO remote

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/JSAC.2015.2393431, IEEE Journal on Selected Areas in Communications

	Word processing locally (i.e. in Microsoft Word)								
Average power consumed by the Netbook to compose document in Microsoft	**11.3 W**								
	Word Processing in Google Drive (Edit online, Save in the cloud)			**Word Processing in Microsoft Skydrive (Edit online, Save in the cloud)**			**Word Processing in Google Drive (Edit offline, Save in the cloud)**		
Power consumption of data centre server (P_d)	0.25 W			0.25 W			0.25 W		
Power consumption of transport network ($N_c E_c B + N_e E_e B + E_{bng} B + E_{sw} B + E_d B$)	8.4×10^{-3} W			1.7×10^{-3} W			0.2×10^{-3} W		
Access network	**4G**	**WiFi**	**Ethernet**	**4G**	**WiFi**	**Ethernet**	**4G**	**WiFi**	**Ethernet**
Power consumption of access network ($E_a B$)	1.9 W	7×10^{-3} W	0.2×10^{-3} W	0.4 W	1.4×10^{-3} W	35.3×10^{-6} W	0.05 W	0.4×10^{-5} W	4×10^{-6} W
Power consumption of Netbook (P_u)	16.1 W	14 W	13.6 W	16.7 W	14.5 W	14.4 W	15.1 W	13.7W	13.2 W
Average power consumed to use the cloud (i.e. sum of the power consumption of the data centre server, transport network, access network, Netbook)	18.3 W	14.3W	13.9 W	17.4 W	14.8 W	14.7 W	15.4 W	13.9W	13.4 W

TABLE VI: Power consumption per user P_I for using the Word processing application locally and in the cloud.

radio head 4G/LTE base station deployed in an urban environment consumes 528 W under full load, and 333 W when idle. The aggregate achievable throughput of this base station is 72 Mbps with 20 MHz spectrum [27]. Further, [26] also reports that base stations consume different amounts of power in each direction (unlike the equipment listed in Table V); roughly 87% of the energy is consumed in the downlink direction and the remaining 13% in the uplink direction. Considering a typical utilisation of 5% over a 24-hour cycle, the energy per bit of this base station, on average, can be approximated as 76.2 μJ/bit in the downlink and 19.0 μJ/bit in the uplink.

6) Power Consumption Per User P_d of Data Centre Server: Obtaining precise information about Google and Microsoft servers is difficult because this information is not publicly available. We instead resort to the following approach to quantify the server power consumption per user. We note that Google's Word processing, Presentation and Spreadsheet applications are a part of the wider Google Apps service suite [16]. The power consumption of a server per user sharing the compute resources, as reported by Google, for the Google Apps services is about 0.25 W [28]. We therefore use this figure of 0.25 W in our calculations. Further, we assume that the per user power consumption of a server in Microsoft's data centre is also 0.25 W. This is a reasonable assumption because a typical server from Google or Microsoft that supports the types of applications considered in this study consumes about the same amount of power, i.e. $\approx$ 200 W [16], [29].

V. POWER CONSUMPTION PER USER P_I

We have used the values from the previous section in (1) to estimate the power consumption per user, P_I, incurred in using the cloud applications. The access network power consumption for 4G is calculated as the sum of the power consumption of the 4G base station in the uplink and downlink directions.

A. P_I for Word processing Applications

Table VI summarises our results for the case when the bit-rate B of the online interactive Word processing application from Google and Microsoft is 55 Kbps and 11 Kbps respectively. The bit-rate B of the Word processing application in Google when editing offline and saving in the cloud is 1.3 Kbps. The key points for Word processing from Table VI are:

1) The average power consumption obtained from measurements for composing and saving the document locally on the netbook using Microsoft Word is 11.3 W.
2) When using the cloud, the power consumption of the transport network is small compared to the contributions made by the other parts of the network. This is because the energy per bit of routers and switches is small (in the order of nJ per bit, see Table V), and so is the bit-rate of the applications (a few tens of Kbps, see Table I).
3) The power consumption of the access network is dominated by 4G (i.e. the 4G base stations), which is three to six orders of magnitude more than a WiFi modem or an Ethernet switch.
4) The power consumption of the netbook computer is a significant fraction of the overall power consumption incurred in using the cloud applications.
5) We estimate the average power consumption per user – i.e. sum of the power consumption of the data centre server, access and transport network, as well as the netbook computer – to use Google Drive and Microsoft Skydrive to vary between 13.9 W and 18.3 W for the former, and between 14.7 W and 17.4 W for the latter (depending upon the access technology used). The power consumption is between 13.4 W to 15.4 W for offline file editing and saving in the Google Drive cloud.
6) Most importantly, online editing and saving the document in the cloud consumes more power than offline editing and saving it to the cloud. Both cloud scenarios (online and offline editing) consume more power than processing and storing the document locally.

B. P_I for Presentation Applications

Table VII shows data for the Presentation application when the bit-rate B for online interaction with Google and Microsoft is 38 Kbps and 27 Kbps. The important points for Presentation applications to emerge from Table VII are:

	Processing presentation locally (i.e. in Microsoft PowerPoint)								
Average power consumed by the Netbook to compose presentation in Microsoft	**11.0 W**								
	Processing presentation in Google Drive (Edit online, Save in the cloud)			**Processing presentation in Microsoft Skydrive (Edit online, Save in the cloud)**			**Processing presentation in Google Drive (Edit offline, Save in the cloud)**		
Power consumption of data centre server (P_d)	0.25 W			0.25 W			0.25 W		
Power consumption of transport network ($N_cE_cB + N_eE_eB + E_{bng}B + E_{sw}B + E_dB$)	5.8×10^{-3} W			4.1×10^{-3} W			0.4×10^{-3} W		
Access network	**4G**	**WiFi**	**Ethernet**	**4G**	**WiFi**	**Ethernet**	**4G**	**WiFi**	**Ethernet**
Power consumption of access network (E_aB)	1.4 W	4.9×10^{-3} W	0.1×10^{-3} W	1.4 W	3.5×10^{-3} W	87×10^{-6} W	0.1 W	0.3×10^{-3} W	7×10^{-6} W
Power consumption of Netbook (P_u)	16.1 W	14.2W	14 W	15.8 W	13 W	12.8 W	15.3W	13.9W	13.4 W
Average power consumed to use the cloud (i.e. sum of the power consumption of the data centre server, transport network, access network, Netbook)	17.8 W	14.6W	14.3 W	17.5 W	13.3 W	13.1 W	15.6 W	14.1W	13.6 W

TABLE VII: Power consumption per user P_I for using the Presentation application locally and in the cloud.

	Processing spreadsheet locally (i.e. in Microsoft Excel)								
Average power consumed by the Netbook to compose spreadsheet in Microsoft	**11.0 W**								
	Processing spreadsheet in Google Drive (Edit online, Save in the cloud)			**Processing spreadsheet in Microsoft Skydrive (Edit online, Save in the cloud)**			**Processing spreadsheet in Google Drive (Edit offline, Save in the cloud)**		
Power consumption of data centre server (P_d)	0.25 W			0.25 W			0.25 W		
Power consumption of transport network ($N_cE_cB + N_eE_eB + E_{bng}B + E_{sw}B + E_dB$)	4.1×10^{-3} W			19.8×10^{-3} W			0.07×10^{-3} W		
Access network	**4G**	**WiFi**	**Ethernet**	**4G**	**WiFi**	**Ethernet**	**4G**	**WiFi**	**Ethernet**
Power consumption of access network (E_aB)	1.0 W	3.5×10^{-3} W	87×10^{-6} W	5.8 W	17×10^{-3} W	0.4×10^{-3} W	0.02 W	0.06×10^{-3} W	0.1×10^{-6} W
Power consumption of Netbook (P_u)	17.8 W	16.6W	16.1 W	16.2 W	14.7 W	14.3 W	15.2W	14.3W	13.4 W
Average power consumed to use the cloud (i.e. sum of the power consumption of the data centre server, transport network, access network, Netbook)	19.1 W	16.9W	16.4 W	22.3 W	15.0 W	14.6 W	15.5 W	14.5W	13.7 W

TABLE VIII: Power consumption per user P_I for using the Spreadsheet application locally and in the cloud.

1) The average power consumption for composing 5-slides locally on the netbook computer using Microsoft PowerPoint is 11.0 W.
2) As in the previous example, moving to the cloud consumes very small power in the transport network, 4G dominates the access network power consumption, and the netbook computer's power consumption is a large fraction of the overall power consumption of the service.
3) The power consumption for using the Presentation application on the cloud varies between 14.3 W and 17.8 W (for Google) and 13.1 W and 17.5 W (for Microsoft). The power consumption varies between 13.6 W and 15.6 W for offline file editing and saving on Google Drive.

C. P_I for Spreadsheet Applications

Table VIII summarises the results for the online interactive Spreadsheet application when the bit-rate B is 27 Kbps for Google and 130 Kbps for Microsoft. The bit-rate for the Spreadsheet application in Google Drive when editing offline and saving in the cloud is 0.5 Kbps.

Composing the spreadsheet locally on the netbook computer using Microsoft Excel incurs 11.3 W, while composing the spreadsheet in the cloud could incur an additional 11 W if using Microsoft via a 4G wireless access network. Other observations are similar to ones described above.

D. Key points

These series of measurements using Google Drive and Microsoft Skydrive's Word processing, Presentation and Spreadsheet applications demonstrate that using the cloud could consume more power than local processing, implying that it is not always energy-efficient to adopt the cloud for performing tasks. When making this comparison it is important to note that interactive cloud applications provide many benefits unrelated to energy efficiency. A prime example being collaborative document drafting and editing by geographically spread team

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/JSAC.2015.2393431, IEEE Journal on Selected Areas in Communications

	Word Processing in Google Drive (Edit online, Save in the cloud)			Word Processing in Microsoft Skydrive (Edit online, Save in the cloud)			Word Processing in Google Drive (Edit offline, Save in the cloud)		
Access network	**4G**	**WiFi**	**Ethernet**	**4G**	**WiFi**	**Ethernet**	**4G**	**WiFi**	**Ethernet**
Average power consumed to use the cloud	3.8 W	2.7 W	2.8 W	2.9 W	3.2 W	3.6 W	0.9 W	2.3 W	2.3 W

	Processing presentation in Google Drive (Edit online, Save in the cloud)			Processing presentation in Microsoft Skydrive (Edit online, Save in the cloud)			Processing presentation in Google Drive (Edit offline, Save in the cloud)		
Access network	**4G**	**WiFi**	**Ethernet**	**4G**	**WiFi**	**Ethernet**	**4G**	**WiFi**	**Ethernet**
Average power consumed to use the cloud	3.3 W	3.0 W	3.2 W	3.0 W	1.7 W	2.0 W	1.1 W	2.5 W	2.5 W

	Processing spreadsheet in Google Drive (Edit online, Save in the cloud)			Processing spreadsheet in Microsoft Skydrive (Edit online, Save in the cloud)			Processing spreadsheet in Google Drive (Edit offline, Save in the cloud)		
Access network	**4G**	**WiFi**	**Ethernet**	**4G**	**WiFi**	**Ethernet**	**4G**	**WiFi**	**Ethernet**
Average power consumed to use the cloud	4.6 W	5.3 W	5.3 W	7.8 W	3.4 W	3.5 W	1.0 W	2.9 W	2.6 W

TABLE IX: Power consumption per user for accessing the Word, Presentation and Spreadsheet applications in the cloud assuming the user is already online.

members. Further, the end-user device and the access network, specifically high-speed wireless, can play a major role in determining the overall power consumption involved in using interactive cloud-based applications.

E. Power Consumption when a User is Already Online

When a user is already online (i.e. connected to the Internet) undertaking other tasks, the network interfaces on the end-user device will already be energised irrespective of use of the interactive cloud-based applications. Therefore, one may adopt the viewpoint that when calculating the power consumption for using the cloud applications we should ignore the idle power of the netbook computer as well as the power consumed for enabling the network interfaces. The idle power of the netbook computer is 10.8 W and the power consumed for enabling the Ethernet, WiFi and 4G interfaces are 0.3 W, 0.8 W and 3.7 W respectively. Subtracting these values from the results given in Tables VI, VII and VIII provides an estimate for the average power consumption involved in using the cloud applications when a user is already online. These values are shown in Table IX.

To make the comparison fair, the power consumed for processing the tasks locally should be the results given in Tables VI, VII and VIII for local processing less 10.8 W, the idle power consumption of the netbook computer. Thus, to compose a document, presentation and spreadsheet locally on the netbook would require 0.5 W, 0.2 W and 0.2 W. We note from Table IX that the power consumption for cloud-based processing using any of the three access network technology is still an order of magnitude larger than the power consumed for local processing.

F. Power Consumption using a Tablet as an End-User Device

In addition to using a netbook computer, we carried out measurements using a Samsung Galaxy Tab 3 Lite, 7 inch tablet [17]. We were unable to replicate the scenarios described earlier in the tablet because the tablet-specific offerings of Google Drive and Microsoft Skydrive applications are still under development. For e.g., at the moment, Google does not support inserting pictures or tables in a browser launched from the tablet, and Microsoft does not have the edit online, save in the cloud feature. We therefore composed a text-only document (same number of words as before) in the Word processing application of Google.

The idle power consumption of the tablet with all network interfaces disabled was 2.3 W. Enabling WiFi and the high-speed wireless interface (3G) increased the power consumption to 2.4 W and 2.5 W; these values denote the baseline power consumption of the tablet. This tablet does not have an Ethernet interface. For the edit online, save in the cloud scenario, the increase in the power consumption of the tablet, relative to the baseline, was 1.7 W (with WiFi) and 2.2 W (with 3G). For the edit offline, save in the cloud scenario (performed using the Google Drive app), the increase over the baseline was 1.4 W (with WiFi) and 1.9 W (with 3G). These values give us the P_u in (1). Invoking (1) and noting that the bit-rate B of the application for each of the two scenarios is 28 Kbps and 5 Kbps on average, gives an estimate of the power consumption incurred in using the cloud with the tablet. Assuming the user is already connected to the Internet, the power consumed for editing the document online is 2.0 W with WiFi and 3.3 W with 3G. The power consumption for editing the document offline and then saving it in the cloud is 1.7 W with WiFi and 3.0 W with 3G. The power consumed to compose the document locally in the tablet (using the Polaris Office App) is 1.0 W.

These results show that even when the end-user device is a tablet (an example of a portable mobile device), processing a task in the cloud could be less energy-efficient than processing the same task locally.

VI. DISCUSSION AND CONCLUSIONS

The results of the modelling have shown that for our set of interactive cloud-based applications, the network transport power is only a small fraction ($< 1\%$) of our estimates of overall power consumption. This finding is consistent with [30] and our conclusion in [11] for low-rate traffic flows between the user and the cloud. As a result, we do not expect our estimates to change significantly if the network topology and/or equipment change. The model also shows that copying and pasting data from the local editor into the browser does not give rise to the large traffic overheads; the overheads arise from real-time interaction with the cloud. Therefore, if one wishes to improve service energy efficiency they would edit locally and only store to the cloud once all the editing is completed. Alternatively there is scope for reducing the traffic overhead multiplier using intelligent client-side caching techniques, and optimising the frequency with which synchronisation of content occurs.

The results in this paper rely on measurements of a netbook computer and a tablet that is representative of low-end user devices for cloud access. Repeating the measurements on other devices could alter the estimates. Similarly, the results show that accessing cloud services via WiFi or Ethernet will generally be less energy consuming than high-speed wireless (3G/4G), however the difference is such that the specific details of the access scenario may change this outcome.

Overall, this work shows that online interactive applications generate high amount of traffic and consume more energy than the same task on a non-interactive environment. Therefore, when online real-time collaboration is not required, it is more energy-efficient to do tasks locally and then save the final version to the cloud.

In conclusion, we have comprehensively examined interactive cloud-based applications and developed a model to estimate the average power consumption per user involved in using these applications. We have shown that the volume of traffic exchanged between the user and the cloud can be considerably larger than that entered by the user, thereby impacting the power consumption of the service. Replacing a 70 W desktop PC (or a 30 W laptop) with a low-power consuming device and adopting the cloud would indeed be energy-efficient. However, our measurements demonstrate that simply migrating to the cloud for processing tasks is not the always energy-wise choice, and it is therefore important to identify the right balance between performing tasks locally and in the cloud for improving energy efficiency.

VII. ACKNOWLEDGMENT

The authors would like to thank Jaime Llorca from Bell Labs for his helpful comments and Yi Liu for her enthusiastic participation in the practical experiments. This work is funded by Alcatel-Lucent, The State Government of Victoria, Australia and The University of Melbourne.

REFERENCES

[1] A. Vishwanath, F. Jalali, R. Ayre, T. Alpcan, K. Hinton, and R. S. Tucker, "Energy Consumption of Interactive Cloud-Based Document Processing Applications," in *Proc. IEEE International Conference on Communications (ICC)*, Hungary, Jun 2013.

[2] Cisco white paper, "Cisco Global Cloud Index: Forecast and Methodology, 2010-2015," 2011.

[3] N. Fernando, S. W. Loke, and W. Rahayu, "Mobile Cloud Computing: A Survey," *Elsevier Future Generation Computer Systems*, vol. 29, no. 1, pp. 84–106, Jan 2013.

[4] R. Buyya, J. Broberg, and A. M. Goscinski, "Introduction to Cloud Computing, Chapter 1," in *Cloud Computing: Principles and Paradigms*, John Wiley and Sons, New York, USA, Feb 2011.

[5] L. Liu et al., "GreenCloud: A New Architecture for Green Data Center," in *Proc. ACM ICAC-INDST*, Spain, Jun 2009.

[6] M. Ali, "Green Cloud on the Horizon," in *Proc. CloudCom*, China, Dec 2009.

[7] Y. Gu, V. March, and B. S. Lee, "GMoCA: Green Mobile Cloud Applications," in *Proc. 1st Int. Wkshp. on Green and Sustainable Software (GREENS)*, Switzerland, Jun 2012.

[8] E. Cuervo, A. Balasubramanian, D. Cho, A. Wolman, S. Saroiu, R. Chandra, and P. Bahl, "MAUI: Making Smartphones Last Longer with Code Offload," in *ACM MobiSys*, USA, Jun 2010.

[9] C. Shi, K. Habak, P. Pandurangan, M. Ammar, M. Naik, and E. Zegura, "COSMOS: Computation Offloading as a Service for Mobile Devices," in *ACM MobiHoc*, USA, Aug 2014.

[10] K. Kumar, J. Liu, Y.-H. Lu, and B. Bhargava, "A Survey of Computation Offloading for Mobile Systems," *Springer Mobile Networks and Applications*, vol. 18, no. 1, pp. 129–140, 2013.

[11] J. Baliga, R. Ayre, K. Hinton, and R. S. Tucker, "Green Cloud Computing: Balancing Energy in Processing, Storage and Tranport," *Proceedings of the IEEE*, vol. 99, no. 1, pp. 149–167, 2011.

[12] F. Jalali, R. Ayre, A. Vishwanath, K. Hinton, T. Alpcan, and R. S. Tucker, "Energy Consumption Comparison of Nano and Centralized Data Centers," in *Proc. ACM Greenmetrics*, USA, Jun 2014.

[13] A. Vishwanath, K. Hinton, R. Ayre, and R. S. Tucker, "Modelling Energy Consumption in High-Capacity Routers and Switches," *IEEE Journal on Selected Areas in Communications*, vol. 32, no. 8, pp. 1524–1532, 2014.

[14] A. Odlyzko, "Data Networks are Lightly Utilized, and Will Stay That Way," *Review of Network Economics*, vol. 2, no. 3, pp. 210–237, 2003.

[15] "MSI Wind U100 Netbook," http://techreport.com/articles.x/15291.

[16] "Google Apps: Energy Efficiency in the Cloud," http://static.googleusercontent.com/external_content/untrusted_dlcp/www.google.com/en//green/pdf/google-apps.pdf, last visited Nov 2013.

[17] "Samsung Galaxy Tab 3 Lite 7" Tablet," http://www.samsung.com/au/consumer/mobile-phone/tablet/tablet/SM-T111MDWAXSA, last visited Sep 2014.

[18] "Power-Mate power meter," www.power-mate.com.au.

[19] "Wireshark - Packet Sniffer," www.wireshark.org.

[20] "Robosoft Automatic Mouse and Keyboard Software," www.robot-soft.com/automatic-mouse-keyboard.html.

[21] "Google Docs Internals," http://googledocs.blogspot.com.au/2010/09/whats-different-about-new-google-docs_23.html, last visited Nov 2013.

[22] "Cisco's power consumption calculator," http://tools.cisco.com/cpc.

[23] D. Kharitonov, "Time-Domain Approach to Energy Efficiency in High-Performance Network Element Design," in *Proc. IEEE Globecom Workshop on Green Communications*, USA, Dec 2009.

[24] "Cisco Catalyst 2960 Series Switches," http://www.cisco.com/c/en/us/products/switches/catalyst-2960-series-switches/index.html, last visited Nov 2013.

[25] "Enterprise 802.11n Wireless LAN Access Point Performance Benchmark," http://www.novarum.com/documents/Enterprise802.11nSingleAPBenchmarkTestingv1.3.pdf, last visited Nov 2013.

[26] G. Auer et al., "Efficiency Analysis of the Reference Systems, Areas of Improvements and Target Breakdown, Deliverable 2.3, Energy Aware Radio and neTwork tecHnologies (EARTH)," https://bscw.ict-earth.eu/pub/bscw.cgi/d71252/EARTH_WP2_D2.3_v2.pdf, 2010.

[27] D. Fritz, "The Evolving Wireless World – Alcatel Lucent Presentation."

[28] "Google's Green Computing: Efficiency at Scale," www.google.com/green/pdfs/google-green-computing.pdf, last visited Nov 2013.

[29] "Microsoft Cloud Infrastructure Strategy," https://www.samsung.com/us/business/oem-solutions/events/cio-forum2012/PDFs/Samsung-CIO-Forum-2012_DileepBhandarkar_Microsoft.pdf, last visited Nov 2013.

[30] D. R. Williams and Y. Tang, "Impact of Office Productivity Cloud Computing on Energy Consumption and Greenhouse Gas Emissions," *Environmental Science and Technology*, vol. 47, pp. 4333–4340, 2013.