

Evaluation of Data Communication Opportunities from Oil Field Locations at Remote Areas

Yang Chen, Jens O. Berg
Schlumberger PCPU
Sugar Land, TX
yangchen, joberg@slb.com

Mostafa Ammar and Ellen Zegura
College of Computing, Georgia Tech
Atlanta, GA
ammar, ewz@cc.gatech.edu

ABSTRACT

Cellular data links are an effective outdoor Internet access solution in urban environments. In this paper, we evaluate cellular data service as a potential data communication solution for oil field crews operating at remote areas in the United States. In our study, we first record the performance of cellular data service at twelve oil field locations. Measurement results show extensive availability of cellular service at those locations making it potentially a data communication solution at field locations. We then design a measurement framework and deploy measurement units to five oil field crews and carry out a side-by-side comparison of two different satellite links and cellular links from two service providers. Analysis of data sets consisting of more than 300 days' measurement shows that the cellular link has comparable or even higher availability than conventional satellite link at many field operations. However, the quality of its coverage is location dependent. This indicates that both cellular and satellite links should be used to provide highly available and cost-effective data communication for such operations.

Categories and Subject Descriptors

C.4 [Performance of Systems]: Measurement techniques;
C.2.1 [Network Architecture and Design]: Wireless Communications

General Terms

Measurement, Experimentation, Performance

Keywords

Cellular, Satellite

1. INTRODUCTION

In today's oilfield operations, data communication at field locations is essential for operation control, equipment monitoring and crew coordination. Data communication is also

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

IMC'11, November 2–4, 2011, Berlin, Germany.

Copyright 2011 ACM 978-1-4503-1013-0/11/11 ...\$10.00.

desirable between oil field locations and applications on the Internet or inside an enterprise network infrastructure. For example, important job parameters during a field job can be sent to a data server where clients are able to view the job progress and provide feedback/comments. After the 2010 Macondo well blowout in the Gulf of Mexico [1], stricter governmental regulations introduced new data communication requirements to field operations. For example, status of oilfield equipment should be reported to a data center regularly. Since most field operations are conducted in remote areas, satellite links have been considered as the only option for a long time. With limited bandwidth provisioning and expensive usage cost, we are evaluating data communication alternatives from field operations to the Internet, in particular, to servers inside our Schlumberger Information Network (SINet).

Recent research evaluated outdoor Internet access from vehicles using 3G cellular service [2–4]. High bandwidth and stable connectivity from 3G data links exist at all those areas. However, all these studies were conducted in urban regions where cellular tower deployment density is high. Performance of cellular data links at remote areas far from cellular towers has not been evaluated in depth.

In this paper, we explore the potential for less costly cellular data usage at remote sites where such usage has not traditionally been contemplated. It is clear that there is significant diversity in availability of cellular and satellite coverage among remote locations. We provide a framework for assessing this availability through systematic measurement and analyze results from this framework at different geographical locations.

We first conducted a series of performance measurements of cellular data links with a field trip covering twelve remote operation locations. Coverage and link characteristics were collected at those locations. Based on analysis of collected information in this trip, we designed the measurement framework and deployed five testing units to field crews to collect performance data of cellular and satellite links over a total time period of more than 300 days. To our knowledge, this is the first side-by-side comparison between cellular data link and satellite links at remote areas.

The rest of the paper is organized as follows. Related work is described in Section 2. We introduce our measurement setup and test scenarios in Section 3 and present analysis of measurement results in Section 4. This paper concludes in Section 5 with summary of findings and plans of future work.

2. RELATED WORK

2.1 Wireless Communication at Oil Fields

Modern oil field operations from exploration to production require a significant amount of interactions between software and hardware as well as collaboration among field personnel. Providing data communication at field locations is not only essential but critical. Wireless technology has been widely adopted since most of oil field operations are conducted at locations where a conventional wired network infrastructure is physically or financially impractical. Wireless LAN (WLAN) and two-way radio are common equipment of field crews with high mobility nowadays. And wireless communication technology continues extending its reach at oil field operations [5, 6]. Correlation between transfer delay and TCP packets' size over GPRS link is studied in [7] for a potential application of real-time oil field monitoring.

2.2 Outdoor Internet Access

Studies of outdoor data access mainly focus on connectivity from moving vehicles. This connectivity can be provided via open WiFi APs [8, 9]. Due to the limited communication range of those APs and their uncoordinated coverage, handoffs with predictive methods and using prefetching to download experience were discussed in [10]. Cellular data service is also considered as a good source for outdoor data access because of its better coverage and less frequent service interruptions. Using tools from information theory, the work in [4] shows a certain level of predictability of bandwidth on 3G links in a metro area. Most recently, performance comparisons between WiFi and 3G connectivity is conducted in urban areas. Although these two types of data links have different characteristics, they can be good supplements to each other based on studies in [2, 3].

2.3 Cell and Satellite Links Characterization

Cellular data service providers normally apply sophisticated control/scheduling method in their networks. The configuration can be highly customized from location to location due to quality of service (QoS) and business concerns [11]. Empirical study of transport layer performance over those managed cellular data links are reported in [12–14].

Satellite networks are under strict regulation and management due to the expensive and limited bandwidth over such links. Research on satellite links focuses on measuring TCP's performance over those links and "tuning" the TCP protocol stack to fit link characteristics such as long end-to-end latency [15–17].

3. MEASUREMENT SETUP

The objective of this work is to evaluate new communication opportunities from oil field locations to the Internet provided by rapidly expanding cellular data service across the continental USA. In the first phase of our study, we conducted a series field tests in northern Texas/Louisiana. During these tests, we evaluated cellular data links on their availability and average TCP throughput. With positive results in this study, we deployed test units to five crews and conducted the second phase of the study. We compare cellular data services and satellite links side-by-side at remote field locations. Our measurement studies target an answer to the question of whether we can cost-effectively improve

the connectivity from remote oilfield operations to the Internet at certain geographical regions.

3.1 Hardware Platform

During the first phase of the study, we built a mount so that two vertical masts were setup on each back corner of a pick up truck. This allowed us to position antennas at a 4-meter elevation above the ground. This matches the elevation of those antennas in field deployments where they are installed on the roof of trailers.

We use two dual-band flexible omnidirectional antennas. The choice of this type of antenna stems from a specific oilfield operational restriction: The cellular antenna on the roof of trailers must be flexible to prevent damage caused by strong vibrations or impacts from vegetation when field operation are carried out at places with harsh road conditions. This restriction rules out many other choices such as electronically steerable antennas [18]. The gain of this flexible antenna is 5 dB as higher gain antennas are usually longer and associated with a higher probability of breakage during field operations. The AT&T and Verizon 3G modems used in our measurement studies come with one or two (Verizon model) external antenna connectors. Both antenna and cellular data modems installation pass our internal environmental qualification procedure, which involves tests for conditions such as strong vibration and extreme temperature.

A different communication system setup was used in the second phase study. We installed a cellular data modem and a BGAN modem from Inmarsat [19] on trailers of five field crews which used the VSAT modem from iDirect [20] as the data access link to the Internet.

Cellular and satellite data modems are connected to a Dell LatitudeTM laptop via Ethernet port or USB/Ethernet converter. Our measurement application is running on the Windows[®] XP operating system.

3.2 Methodology

In our phase one study, we visited 12 field-operation locations. A commercial Internet speed test server [21] and an internal link speed test server inside SINet were used to measure the upload/download link speed. If data service was available, the average link speed was recorded. Data service coverage and link speed of both service providers were evaluated. Good coverage revealed during this study led us to continue further comparison between cellular and satellite links in field locations.

In our test unit deployed on field crew's trailer, we developed an application to probe for Internet access over each data link every 30 seconds. Furthermore, when the field crew started the field operations, data files with a size of 1 MByte are generated regularly. This is to mimic data generation from a planned application for field service-quality improvement. Our application automatically uploaded those files to a FTP server inside SINet over an available data link. Interrupted uploading was resumed when the connectivity was restored. Each FTP transmission's size and duration was recorded and used to compute average TCP throughput at that time.

Throughout our measurement studies, we used the original TCP stack of the Windows operating system. Our results can serve as the basis for future effort on "tuning" the TCP stack for performance improvement over cellular or satellite links.

End-to-end latency is the link metric which essentially affects TCP stack's performance and determines service quality experienced by many applications. To collect end-to-end latency measurement results, periodical ping's are sent over cellular and satellite links.

3.3 Switching between Network Interfaces

In order to send probe traffic through each network interface, we developed an interface switching mechanism at application layer by manipulating the Windows operating system's local routing table.

Different IP addresses are assigned by DHCP protocols on each of these three interfaces. For example, 192.168.1.100 is the IP address at the interface connected to a satellite data modem which serves as the gateway with the IP address 192.168.1.1. In the meanwhile, 192.168.2.100 is the IP address assigned by a cellular data modem with IP address 192.168.2.1.

Destination	Netmask	Gateway	Interface
10.0.0.1	255.255.255.255	192.168.1.1	192.168.1.100
0.0.0.0	0.0.0.0	192.168.2.1	192.168.2.100

Figure 1: Samples of routing table entry

By setting a routing table entry for all IP addresses, i.e., destination IP is 0.0.0.0 shown in Figure 1, all network traffic is forwarded onto the cellular link. Note that a more specific routing table entry can be created at the same time to specify which interface should be used for the network application with a pre-determined destination IP address. In Figure 1's example, the application with a destination IP address 10.0.0.1 is using the satellite link.

4. MEASUREMENT RESULTS

Results from the Phase 1 study were recorded manually when we visited each field locations. After we deployed measurement units to field crews, measurement results were recorded automatically in a local database on the laptop. We developed scripts to retrieve essential information from database records as log files and upload them to a FTP server from time to time. Logs from both phases are used to present following analysis.

4.1 Phase 1 Study

From March 31, 2010 to April 9, 2010, we visited 12 field locations spread out in an area of more than 35,000 km^2 shown in Figure 2a¹. Selection of those locations is based on operational scheduling and business concerns other than geographical or service coverage consideration.

Coverage (%)	ATT	VZN	Bandwidth (kbps)	ATT	VZN
3G	0	75	Upload	94	323
non-3G	100	0	Download	106	645

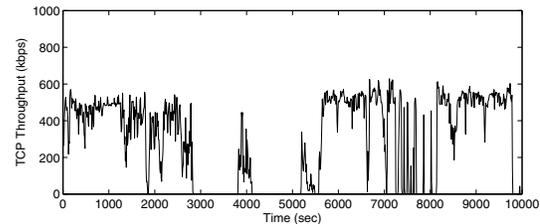
Table 1: Field test summary of cellular links

Field test results are summarized in Table 1. AT&T has EDGE service at all those locations but no high-speed 3G coverage at any of them. On the other hand, Verizon provides 3G EVDO service at nine of the locations and no service in the other three. Therefore, significant differences

¹Due to business concerns, we cannot disclose those fields' exact locations on this map.



(a) Cellular coverage over field tests area



(b) Upload rate on the route

Figure 2: Phase 1 study's locations

in the average upload/download speeds between these two service providers were observed.

During our tests, we found out that the coverage map from service providers can be a good reference but disagreements from actual measurement do occur. We overlaid Verizon's coverage map from [22] on the map of our field tests. The shadows in Figure 2a represent areas without coverage. Among the three field locations where we did not get Verizon data service, two are actually inside those shadows. The one location that "should" have service but does not is outside any shadows. However, its topographic situation is not favorable: the surrounding ground is higher than the field location. We also monitored the TCP throughput in a 3-hours drive from point A on the map (not among the field locations) to a nearby city, point D. During this trip, 600 MB-byte of data was uploaded continuously. The TCP throughput over time illustrated in Figure 2b generally agrees with coverage along the route with occasional differences.

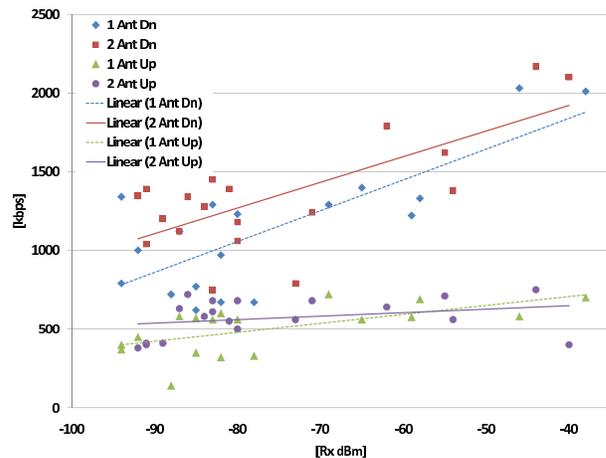


Figure 3: Antenna setup and Verizon's data link speed

We plot download/upload speeds over Verizon's links against

signal strength readings from data modems in Figure 3. Linear fittings indicate that a 1-dB increment on signal strength can lead to an increment of 28.4 kbps in download speed and 7 kbps in upload speed with one antenna. Using two antennas generally increases average signal strength by 1 dB. However, the effect is a 15% download speed increment. We believe this is a result of spatial diversity provided by two antennas. On the other hand, uploading speed is not affected much. Unlike the signal spatial diversity study in [23] where geographic distance between antennas can be up to 500 meters, the distance between antennas in our setting is a few meters. Using two antennas helps on link speed in our field tests but cannot make 3G data service available at locations where a single antenna setup fails.

After we finished this series of field tests, we decided to use a Verizon data modem with the two-antennas setting as the standard 3G data link in our further studies.

4.2 Phase 2 Study

During our phase 2 study, we collected measurement results from four crews stationed at Shreveport, Louisiana. They are named as SLA1 to 4 in this paper. Measurement study was also conducted at one crew stationed at Williston, North Dakota which is named as WIL1 in this paper. Normally, each crew operates within an area of thousands of square kilometer around the city where it is stationed. SLA1 is the pilot crew with which we finalized the hardware and software configuration during the time period from November 2010 to March 2011. Note that the AT&T cellular link was used during that time period. This is because Verizon did not have enterprise-level access to our company’s network until early 2011 after certain business procedures were carried out. Therefore, we name the data set from January 26, 2011 to March 20, 2011 as “SLA1 ATT”. At the end of March, we replaced SLA1 crew’s AT&T link with Verizon service so that it has the same configuration as other four crews. We present analysis of link availability, TCP throughput and end-to-end latency from measurement results collected at those five crews.

4.2.1 Availability

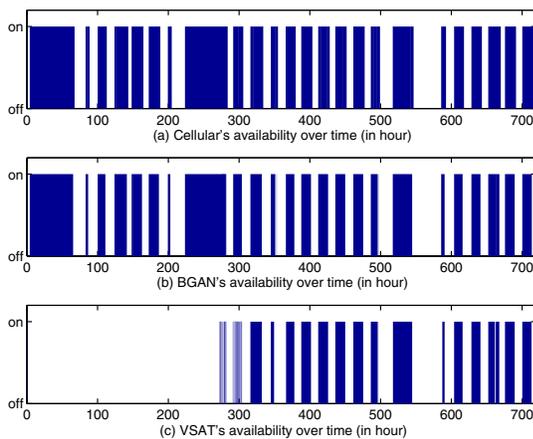


Figure 4: Availability comparison at SLA2 in April 2011

Data link availability via cellular, BGAN and VSAT over time for each crew is the first metric of interest in our measurement study. We find out this metric is affected by the

work shift pattern of each crew. Figure 4 shows the link up and down sequences of SLA2 crew in April 2011. This crew is a day crew which stops operation and shuts down all equipment at night. SLA2 and SLA3 belong to this category and thus their link availability is always interrupted by such operation halts. SLA1 and SLA4 are 24-hour crews which operate continuously at field locations. All crews normally move to next field location after working 5 to 10 days at one place. The job location transition may take 1 to 3 days during which all equipment is powered off.

We retrieved Windows OS’s native logs containing PC power on/off events from different crews. With this type of information, we reprocessed availability logs of cellular and satellite interfaces so that the probability that a link is available is only calculated when the equipment is powered on. Therefore, we eliminated the effect of job transition/equipment down schedule on availability evaluation of those data links.

Difference on costs associated with usage of satellite data links (USD 5/MByte) and cellular one (USD 100 per month with 2-GByte cap) indicates a potential cost reduction if a cellular link instead of a satellite link is used when data communication happens. Therefore, we not only calculate $P(X)$, probability that link X is available, but also the $P(X|Y)$, conditional probability that link X is available given link Y is available. Higher value of $P(X|Y)$ represents more opportunities of using link X instead of Y.

	$P(C)$	P(B)	P(V)	$P(C B)$	$P(C V)$
SLA1	98.50%	71.29%	66.42%	98.99%	98.99%
SLA2	97.03%	83.95%	52.86%	99.52%	99.48%
SLA3	93.76%	88.15%	56.99%	98.12%	97.78%
SLA4	89.73%	88.90%	44.79%	91.89%	99.28%
C: Cellular; B: BGAN; V: VSAT					

(a) Measurement results in April

	$P(C)$	P(B)	P(V)	$P(C B)$	$P(C V)$
SLA1	96.86%	26.96%	87.10%	99.67%	98.31%
SLA2	96.07%	80.40%	75.19%	99.05%	99.07%
SLA3	97.28%	91.21%	63.86%	99.27%	99.02%
SLA4	76.11%	75.68%	50.89%	98.03%	97.79%

(b) Measurement results in May

Table 2: Normalized link availability comparison

Table 2 shows the availability analysis of SLA1 to 4 in April and May 2011. Note that cellular data link has high availability comparable to the BGAN data link at those measurement locations. The drop on BGAN’s availability at SLA1 in May is caused by a device mis-configuration which was corrected after a few weeks. Note that the existing data link solution using VSAT is not as reliable as the other two links. One major contributor to the low availability is its bulky and heavy directional antenna (size ranges from 75cm to 1.2m) which fails to track satellite signal from time to time due to reasons such as mechanical parts failure.

High conditional probabilities in Table 2 indicate that a cellular link is always available when satellite links are at locations around Shreveport, an geographical region with good cellular service coverage. The same conclusion might not hold at other regions.

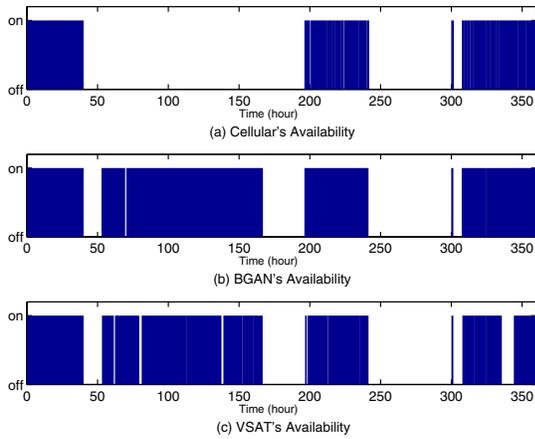


Figure 5: Availability comparison at SLA4 in Aug. 2011

SLA4 crew traveled to southern Texas and operated there for two weeks from Aug. 10 to Aug. 25, 2011. Side-by-side link availability comparison in Figure 5 shows that cellular link was down for a 5 day time period when satellite links were available. Therefore, $P(C|B)$ and $P(C|V)$ are 54.31% and 53.33% in those two weeks. In the meanwhile, $P(B)$ and $P(V)$ are 90.24% and 84.84%.

Measurement results in Table 2 and Figure 5 indicate that no ideal data communication solution can be provided by a single type of data links in our evaluation. A multi-interface system with data link management policy should be developed based on financial consideration and link availability. Implementation on network interface selection can be done via the technique mentioned in Section 3.3.

Both Figure 4 and 5 show regular availability interruptions on those data links. These interruptions may last for hours or even days. There are two reasons for such long interruptions: 1) The crew is working at a field location with no coverage for a few days; 2) The job location transition takes a few days to finish.

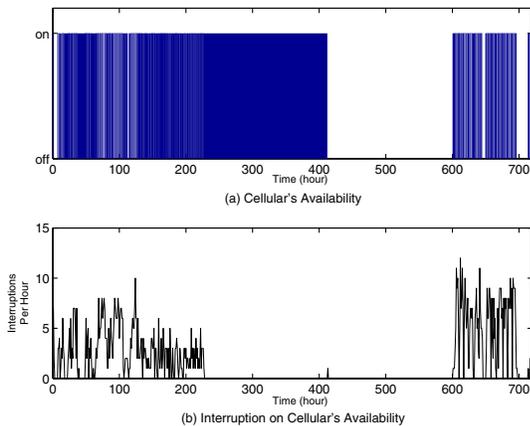


Figure 6: Interruptions on cellular link availability at WIL1

In the meanwhile, interruptions on those wireless links caused by weak signal can be very frequent. We present cellular link's availability from data collected at WIL1 during August in Figure 6. In this month, over 1100 interruptions occurred with a peak rate of more than 10 times per hour.

Overall, data access from field location to the Internet is a delay-tolerant network scenario [24]. Applications running over these data links at field locations should include disruption/delay-tolerance features to handle intermittent link availability as mentioned in [3].

4.2.2 TCP Throughput Profiling

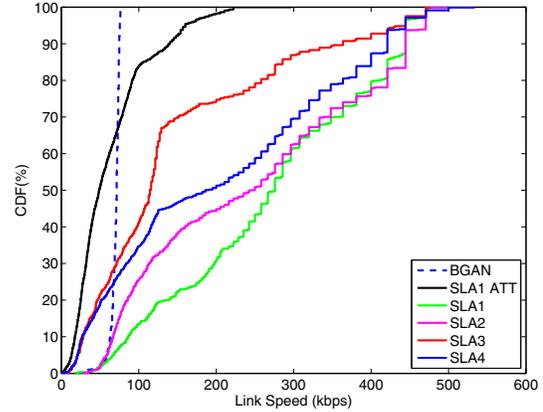


Figure 7: CDF of average TCP throughput

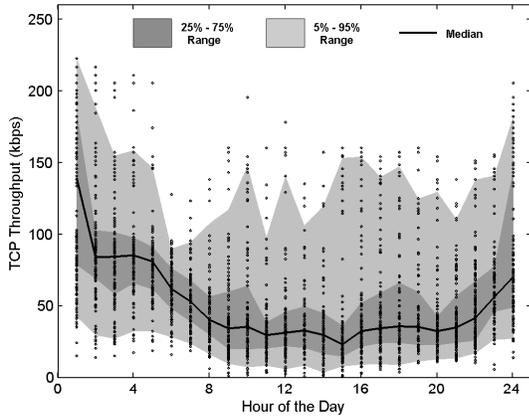
There were 200 to 1000 FTP uploads via Verizon's cellular links in each SLA crew during April 2011. Another 2000 FTP uploads were finished in data set SLA1 ATT. In addition, a BGAN link was used for more than 600 FTP uploads from SLA1 crew from October 2010 to November 2010. No FTP upload was conducted on VSAT link in our measurement study since this link is strictly managed and only business applications with recorded and approved bandwidth consumption profile may access it.

Figure 7 illustrates Cumulative density functions (CDFs) of average TCP throughput in those data sets. We notice the significant variation on average TCP throughput over cellular links. This is caused by fluctuation of signal quality as field locations are normally far from cellular towers or by contentions of spectrum usage, which will be discussed later in this section. Verizon's cellular service provides higher maximum achievable uplink speed (500kbps) than AT&T's (250kbps) but is associated with bigger variations. The average TCP throughput is stable over the BGAN link as it is tightly managed and has less background traffic fluctuation due to limited accesses.

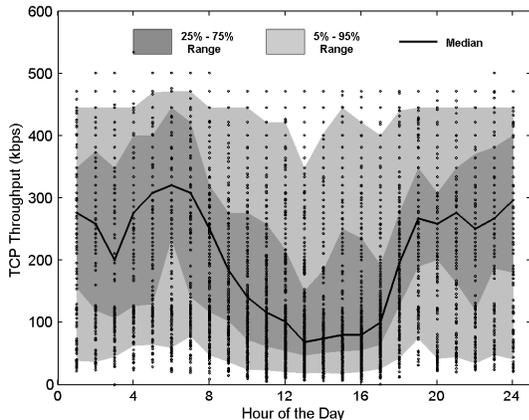
Temporal correlation/predication of TCP throughput over wireless links in small time scales of tens of seconds are investigated in previous work [2,4]. We are interested in average TCP throughput over cellular links and we conduct our study in a much larger time scale.

We show the average TCP throughput of each upload from the SLA1 ATT data set according to the hour when it starts in Figure 8a. Each dot represents one average TCP throughput measured at a certain hour. We also draw the range bounded by 5 percentile and 95 percentile of all values at each hour. Although 5% to 95% ranges are similar across different hours as link quality fluctuates all the time, the 25% to 75% range is narrow which indicates that a significant portion of TCP throughput measurements at each hour are around the median value.

Average TCP throughput is generally low during busy hours of the day and high otherwise. We believe this is re-



(a) SLA1 ATT data set



(b) SLA1-4 Verizon data set

Figure 8: TCP throughput at each hour of the day

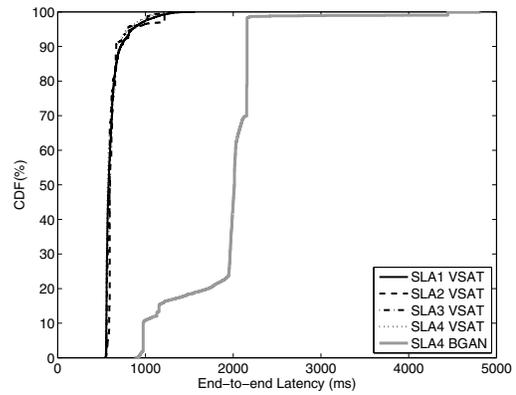
lated to contention of spectrum usage towards the cellular tower from other cellular devices at the same field location or the diurnal congestion on cellular service providers' backbone. As median values change significantly from hour to hour, collecting statistics for each hour of the day would help TCP throughput predication in the same hour of the day.

Similar hourly-shifted pattern can be found in Figure 8b, which illustrates results from the data set consists of over 6000 TCP throughput measurements over Verizon's service from SLA1 to 4 in April and May, 2011.

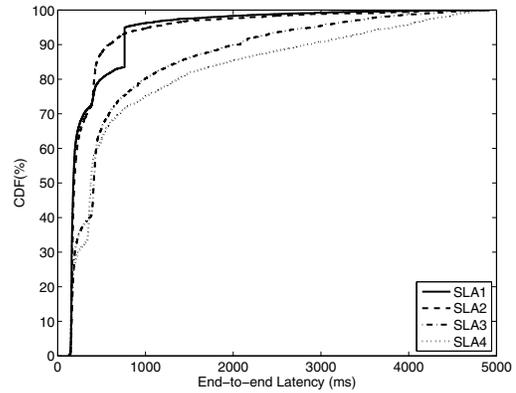
4.2.3 End-to-end Latency

End-to-end latency is the link metric which essentially affects TCP stack's performance and determines service quality experienced by many applications. To collect measurement of this metric, ping packets are sent over SLA1-4's cellular link and VSAT link periodically with an interval of 5 minutes and 15 minutes respectively. Ping packets are also sent over SLA4's BGAN link with an interval of 15 minutes.

Figure 9 shows those measurements in August 2011. Note that VSAT link's end-to-end latency is stable around 500 ms. In the meanwhile, the latency on BGAN link is significantly longer with a lower bound of 1000 ms and 80% measurement results are around 2000 ms. Compared with satellite links, cellular links have a better lower bound around 200



(a) Measurements over satellite links



(b) Measurements over cellular links

Figure 9: CDF's of end-to-end latency

ms. However, measurements of latency on cellular links are heavy-tailed which means long end-to-end latency may occur from time to time.

Although BGAN and cellular links may have higher availability than the VSAT link according to previous discussions, they might not be used by applications having a tight latency requirement. On the one hand, this indicates multiple network interfaces should be jointly used to provide effective data communication at remote areas; on the other hand, significant differences on end-to-end latency characteristics could make the link management a challenging task as one more criteria should be evaluated besides the cost and availability mentioned in previous discussion. Furthermore, an application's data communication may switch between different interfaces when availability interruptions occur. Whether we should make TCP stack aware of such latency discrepancy in advance so that a faster rate adjustment can be achieved is being evaluated and will be reported in future study.

5. SUMMARY AND FUTURE WORK

We conducted a coverage evaluation in a certain region of the United States in 2010 to confirm the expanding presence of cellular data services in remote areas. Due to the distance from test locations to cellular towers and possibly unfavorable terrestrial conditions, variations in service coverage and data link speed are high. Spatial diversity using two antennas provides improvements over cellular link's speed, but there was no coverage difference in our study. Encouraged

by overall good coverage of cellular service, we conducted comprehensive field tests with five oil field operation crews at the different geographical region in 2011. We evaluated communication opportunities provided by satellite and cellular links in this study of a total over 300 days.

At certain regions, the cellular data link shows comparable availability to satellite links which were once deemed as the only solution for accessing the Internet at remote areas. This indicates a potential operational cost reduction from replacing the satellite link with a more cost-effective cellular data link or developing network policy to manage those data links with financial consideration. We also notice that the operational pattern of the field crew may affect data links' availability over time. Delay-tolerance capability is desirable for potential applications using those links due to intermittent availability at field locations. Average TCP throughput on cellular links shifts according to the hour of the day.

We are selecting crews operating at other districts in the United States to extend our measurement study to other geographical regions. In addition, we will provide more comprehensive characterization on those data links at a finer time granularity. For example, end-to-end latency or jitter performance will be measured so that we are able to evaluate possibilities of introducing more interactive applications such as real-time collaborations between office and field locations. Interaction between TCP/UDP stack and link characteristics will be investigated for potential "tuning" of the transport layer for performance improvements over those links.

6. REFERENCES

- [1] BP Deepwater Horizon Explosion. [Online]. Available: http://en.wikipedia.org/wiki/Deepwater_Horizon_explosion
- [2] P. Deshpande, X. Hou, and S. R. Das, "Performance Comparison of 3G and Metro-Scale WiFi for Vehicular Network Access," in *Proc. of Conference on Internet Measurement*, 2010, pp. 301–307.
- [3] A. Balasubramanian, R. Mahajan, and A. Venkataramani, "Augmenting Mobile 3G Using WiFi," in *Proc. of ACM MobiSys*, 2010, pp. 209–222.
- [4] J. Yao, S. S. Kanhere, and M. Hassan, "An Empirical Study of Bandwidth Predictability in Mobile Computing," in *Proc. of ACM WiNTECH*, 2008, pp. 11–18.
- [5] S. Petersen, B. Myhre, S. Carlsen, J. H. Linden, B. Myhre, M. Sansom, A. Skavhaug, E. Mikkelsen, and D. Sjong, "A Survey of Wireless Technology for the Oil and Gas Industry," *SPE Projects, Facilities and Construction*, vol. 3, no. 4, pp. 1–8, 2008.
- [6] S. Savazzi and U. Spagnolin, "Synchronous Ultra-Wide Band Wireless Sensors Networks for Oil and Gas Exploration," in *Proc. of IEEE Symposium on Computers and Communications*, 2009.
- [7] W. Yin-wen and H. Min-xian, "Remote Monitoring System For Oil Wells Based On GPRS Technology," in *Proc. of International Conference on Computer Engineering and Technology (ICCET)*, vol. 7, 2010.
- [8] A. Balasubramanian, Y. Zhou, W. Croft, B. Levine, and A. Venkataramani, "Web Search From a Bus," in *Proc. ACM workshop on Challenged Networks*, 2007, pp. 59–66.
- [9] V. Bychkovsky, B. Hull, A. Miu, H. Balakrishnan, and S. Madden, "A Measurement Study of Vehicular Internet Access using In Situ Wi-Fi Networks," in *Proc. of ACM Mobicom*, 2006, pp. 50–61.
- [10] P. Deshpande, A. Kashyap, C. Sung, and S. R. Das, "Predictive methods for improved vehicular WiFi access," in *Proc. of Mobisys*, ser. MobiSys '09, 2009, pp. 263–276.
- [11] W. L. Tan, F. Lam, and W. C. Lau, "An Empirical Study on the Capacity and Performance of 3G Networks," *IEEE Transactions on Mobile Computing*, vol. 7, no. 6, pp. 737–750, 2008.
- [12] Y. Lee, "Measured TCP Performance in CDMA 1xEV-DO Network," in *Proc. of PAM conference*, 2006.
- [13] K. Mattar, A. Sridharan, H. Zang, I. Matta, and A. Bestavros, "TCP over CDMA2000 Networks: A Cross-Layer Measurement Study," in *Proc. of PAM conference*, 2007.
- [14] J. Prokkola, P. Perala, M. Hanski, and E. Piri, "3G/HSPA Performance in Live Networks from the End User Perspective," in *Proc. of IEEE ICC*, 2009, pp. 1–6.
- [15] K. Zbynek, N. Marek, and B. Leos, "Optimization of TCP Satellite Communication in Inmarsat Network," in *Proc. of IEEE ICWMC*, 2010, pp. 544–548.
- [16] Y. Chotikapong, H. Cruickshank, and Z. Sun, "Evaluation of TCP and Internet traffic via Low Earth Orbit Satellites," *IEEE Personal Communications*, vol. 8, no. 3, pp. 28–34, 2001.
- [17] I.F.Akyildiz, G. Morabito, and S. Palazzo, "TCP-Peach: a new congestion control scheme for satellite IP networks," *IEEE/ACM Transactions on Networking*, vol. 9, no. 3, pp. 307–321, 2001.
- [18] V. Navda, A. P. Subramanian, K. Dhanasekaran, A. Timm-Giel, and S. Das, "Mobisteer: Using steerable beam directional antenna for vehicular network access," in *Proc. of Mobisys*, 2007, pp. 192–205.
- [19] Broadband Global Area Network (BGAN). [Online]. Available: http://www.inmarsat.com/Services/Land/Services/High_speed_data/default.aspx
- [20] Very Small Aperture Terminal (VSAT). [Online]. Available: <http://www.idirect.net/Products.aspx>
- [21] Speakeasy Speed Test. [Online]. Available: <http://www.speakeasy.net/speedtest/>
- [22] Verizon Coverage Locator. [Online]. Available: <http://www.verizonwireless.com/b2c/CoverageLocatorController>
- [23] F. P. Tso, L. Cui, L. Zhang, W. Jia, D. Yao, J. Teng, and D. Xuan, "DragonNet: A Robust Mobile Internet Service System for Long Distance Trains," in *Proc. of IEEE INFOCOMM*, 2011.
- [24] K. Fall, "A Delay-Tolerant Network Architecture for Challenged Internets," in *Proc. of ACM SIGCOMM*, 2003, pp. 27–34.

Summary Review Documentation for

“Evaluation of Data Communication Opportunities for Oil Field Locations at Remote Areas”

Authors: Y. Chen, J. Berg, M. Ammar, E. Zegura

Reviewer #1

Strengths: Novel measurement study that compares satellite and cellular links at remote places. Well written paper and interesting to read.

Weaknesses: Cannot generalize these results as the study is too specific to those particular locations in the southwest US.

Comments to Authors: The paper is an interesting study of cellular and satellite network connectivity at oil field locations. It is of course hard to generalize the results to other locations, but sheds light on the general trends of expanding cellular infrastructure that might prove to be a more compelling alternative to more expensive satellite links. I like this general thesis of the paper.

I have only one major comment regarding the experimentation. It might be good to use the job transition and equipment down schedules from the crews to weed out periods of planned unavailability. Right now, the paper reads a bit speculative. For instance, the reasons for the longer tail in Figure 4 are vaguely attributed to these causes. It might be good to separate them out so that results are more reflective of actual availability of the links. Of course it might be a bit harder since somebody has to collect those logs!

P3: I don't understand how the author's results can serve as the basis for future effort on "tuning" TCP stack. It would be good clarify.

Table 1 needs units for upload and download bandwidths. Are the coverage results based on actual measurements or coverage maps from ATT and Verizon ?

P3: 1dB gain in signal strength translates to 28.4 Kbps in download speed and 7kbps in upload speed. 1dB gain due to multiple antennas is 15% download speed increment. It's not easy to see the comparison between absolute and relative increments. Please standardize.

Reviewer #2

Strengths: Understanding network access in remote/non urban areas sounds like a new emerging area.

Weaknesses: I am not sure any of the findings have much generality or application to other networking problems. The results or techniques are not particularly surprising or interesting.

Comments to Authors: I am all in favor of emerging problem areas and you seem to have the potential to do so with your setup. Unfortunately, there is not much interesting networking in this paper. It would be nice if you can use your measurements to

develop new insights or implications for networking -- e.g., how can you use a hybrid cellular-satellite combination or how can DTN help or how should a TCP stack be optimized etc. Sadly, this is missing here.

You mention "high availability" of cellular, but that seems likely biased by the coverage/workload for the providers and what their economic and business considerations are. I am not sure how we can generalize this to other settings or deployments

Reviewer #3

Strengths: One step in a new environment: how to build communication network at lower cost for remote areas and operational deployment of oil company. The need for information in this system needs to be on the rise. Big deployment effort. Four teams are covered with different schedules, and different providers and satellite are used for completion. Points to need for better accommodation of delay in application, and the role of different factors.

Weaknesses: Apart from the obvious benefit for the operational cost of the oil company, the depth of the problem is not made very precise. In particular, since satellite already have longer delay and more sporadic availability, it seems that switching to cellular is not going much more into the DTN area. The cellular network seems to underperform the satellite in terms of raw bandwidth (through is generally below 70kbps). It is a little frustrating that unavailability due to hardware switched off (or non operation of the team) is not factored out from the experience altogether.

Comments to Authors: I think this paper is a good fit for a short communication at IMC since it deals with measuring availability in a new context (operational team of oil field) and brings several insights on a potential new area of applications for wireless system. This work seems in an early phase, as the result are somewhat limited (old technology from satellite are shown to be more expensive but more predictable). The main surprising message is that availability of cellular networks is higher than satellite.

I think the work can benefit from clarifying a few issues, in particular with regard to its conclusion for future research. This could potentially be included in the paper with a little time.

What are the factors to include in the satellite availability? Is that mostly a hardware and deployment issue (directional antenna and failure) or a fundamental limitation (moving satellites that are not always present). Distinguishing these factors (even roughly) to tell what's counted and what is not would help.

Similarly, as a sanity check it would be important to at least say a word about unavailability due to non-work. Are the trucks left in a place with coverage (in that case, we should see that at night for SLA2). Is that guaranteed that both satellite and cellular are off then?

The time where the power is off, it seems to indicate that the radio is also off. Hence it should be able to find that in the data itself and factor it out.

Can you specify your related improvement w.r.t. [5] and [7]. It seems that [7] already proposed to use cellular network for oil field, so it would be important to contrast your finding or system.

Can you elaborate on the impact of unavailability. Apparently, satellite links are more sporadic (but more predictable when available) than cellular link. That seems to indicate that they would already benefit from DTN. Maybe the cellular network is not making a strong case then. Unless the variability of available throughput is making DTN a must.

For these reasons, the unavailability duration of satellite should be discussed along with Fig.4. while Fig.2 seems not to be critical and left to a tech report.

Clearly the impact of cellular coverage is huge here. If it is modified by 20% then the result of availability will be strongly modified. Can you at least have a quick back of the envelope calculation to prove that other locations considered for oil fields (mentioned in introduction) are not much less exposed to cellular service? That would help making the result representative.

The fact that the entropy reduces as a significant parameter like the time of the day is used seems trivial. This section seems not to add much to the paper at this time (although the general problem of prediction could be important).

Since cellular is better for coverage but less for instantaneous throughput, a total comparison of the volume that can be transported would be useful.

Reviewer #4

Strengths: This is the first paper I am aware of that compares satellite and cellular network performance in remote rural settings. The paper topic is of interest, as it provides insights into a new aspect of Internet use.

Weaknesses: The presentation needs significant copy editing (normally I provide a list of specific issues to correct, but in this case it needs a native English speaker to work directly with the authors (actually, at least one of the authors appears to fit this description, and thus should be able to do this). The analyses are perhaps less rigorous than is typically expected of IMC papers, but given the challenges in obtaining the measurements, are likely acceptable for a short paper.

Comments to Authors: The most glaring issue with the paper is the need for copy editing. There are many, many minor language issues that need to be corrected. Fixing these will make the paper much more enjoyable for a wider audience. There is also quite a bit of repetition in the paper (e.g., about conducting a study with 12 sites, then a follow-up study with four sites).

Please provide more insights on how the initial 12 sites were selected, and then how the final four were selected.

In Section 3.1, please explain in more detail the benefits of an additional 5db from the flexible antenna. Keep in mind that not everyone who reads this paper may understand signal strength measurements.

In Section 3.2, is ftp really used in operations? that is an insecure protocol. if you meant sftp or something similar, please be specific.

It would be helpful to have a table that summarizes the data sets; e.g., how many flows are in each.

The cost information is very interesting; it serves as a good motivation. Thus, you may want to mention this in Section 1.

Figure 6(b) is very interesting. Please make it at least column width. This graph suggests that some changes to business practices/policies may be required; e.g., leaving the IT equipment on at night to take advantage of the increased data rates.

Please clarify what value the entropy analysis provides over the straightforward graph in Figure 6(b).

As future work, I think passive measurement techniques should be explored, since this is clearly a bandwidth constrained environment. Obviously, it would require the network to be in use for business purposes, and then that traffic would be passively monitored.

Reviewer #5

Strengths: The practical nature of this work makes it interesting. It is nice to see work aimed at solving a real problem (data connectivity on oilfields).

Weaknesses: The paper lacks real insight, and it is just too poorly written to really merit attention in IMC.

Comments to Authors: Apart from very frequent problems in English, it is hard from the start to understand exactly what the problem being solved here.

If it is connectivity in remote locations, then how is AT&T or any other service provider relevant. At really remote locations, e.g., and oil rig at sea, such as in the motivating citation to the BP disaster, then clearly a commercial data service is unlikely to be a viable option.

If the question is how well does 3G and related data services work near the edge of their range, then the paper doesn't really answer that because we don't know the positions of cell centers relative to measurements.

Are you proposing that cellular data be installed at remote locations for oil field operators? Even so, the data would have to be backhauled by some technology, e.g., satellite.

These questions aside, a large part of the performance which is diurnal is probably dominated by usage of the cell by other users. That hardly seems relevant to any of the problems above.

Response from the Authors

(1) Generality of contributions (Reviewers 1, 2, 3 and 5): The paper's contribution is in exploring the potential for less costly

cellular data usage at remote sites where such usage has not traditionally been contemplated. It is clear that there is significant diversity in availability of cellular and satellite coverage among remote locations. Our goal in this paper is to provide a framework for assessing this availability through systematic measurement and evaluation and analyzing instances of this framework at different geographical locations.

(2) Filtering data based on job transition and equipment downtimes (Reviewers 1 and 3): We retrieved Windows OS's native logs containing PC power on/off events from different crews. With this type of information, we reprocessed availability logs of cellular and satellite interfaces and eliminated the effect of job transition/equipment down schedule on availability comparison.

(3) About optimizing the TCP stack (Reviewers 1 and 2): Although optimization of TCP performance is outside the scope of this paper, we added end-to-end latency measurements in the final paper, which shows high variation across different interfaces. This indicates a specially designed TCP stack which adapts to high latency variation might have better performance over the default TCP stack of Windows OS in such a multi-interface system.

(4) Discussion on entropy analysis (Reviewers 3 and 4): We removed the discussion on entropy analysis about bandwidth

prediction as the information presented in this analysis can be found in other parts of the paper as reviewers pointed out.

(5) Writing and presentation (Reviewers 4 and 5): Several rounds of editing were applied to the final version of this paper to improve the English and presentation.

(6) Clarification of Table 1 (Reviewer 1): Clarification of coverage results in Table 1 is provided and the unit of measurement is also added.

(7) Hybrid cellular and satellite usage (Reviewer 2): The methodology we developed in this measurement study is actually used to provide a simple link management scheme. Therefore, an application can select among interfaces based on their availability and a pre-determined preference.

(8) Measurement location/crew selection (Reviewer 4): We've added a discussion on the criteria and process of location/crew selection.

(9) Cellular link's availability is questionable at certain locations (Reviewer 5): Our objective is not replacing satellite links with cellular links or vice versa. On the contrary, our measurement framework shows that highly-available and cost-effective data communication may be achieved via a multi-interface system.