Address and traffic dynamics in a large enterprise network

Richard Mortier, Thomas Karagiannis, Peter Key
Microsoft Research
{mort,thomkar,peterkey}@microsoft.com

July 2008

Technical Report
MSR-TR-2008-98

Microsoft Research
Microsoft Corporation
One Microsoft Way
Redmond, WA 98052
http://www.research.microsoft.com

1. INTRODUCTION

The surge of interest in traffic measurements and characterization over the last decade has lead to a plethora of studies in various aspects of network traffic in the wide-area Internet (e.g., [21]), tier-1 ISPs (e.g., [7]) or university campuses (e.g., [9]). However, despite their significance, enterprise networks have not been the subject of many such analyses ([20] constitutes one of the few exceptions); knowledge of their dynamics is still poor, as limitations such as data sensitivity and access restrictions have inhibited similar progress. Yet, understanding enterprise network traffic patterns is a prerequisite of proper provisioning by network operators, of network dimensioning and network modeling, and can also provide the baseline for anomaly detection. It also opens up the possibility of real-time reaction to on line measurements for network management related actions, such as service migration or load balancing.

This paper is a step towards this direction. Our work provides an extensive description of the address and traffic dynamics of a site that is part of Microsoft's Corporate Network. Through the analysis of a data corpus that spans 3.5 weeks of continuous collection, and contains 13 billion packets (section 2), we address a series of questions pertinent to understanding the aforementioned aspects of today's enterprise networks.

Specifically, we first examine traffic dynamics along three perspectives: a) traffic spread within the enterprise network and its geographical dispersion, b) the relevance of layer-4 port numbers to identify specific applications, and c) the validity of the client-server distinction in terms of traffic volumes. Then, taking into advantage routing configuration files and address allocation information we examine address dynamics along two dimensions: a) the meaning of IP addresses as host identifiers and vice versa, i.e., the interpretation of name to IP mappings, and b) host mobility patterns within the larger enterprise network.

Addressing these questions constitutes the main contribution of this paper. To study traffic dynamics we first divide the observed traffic flows in four classes separating data center, local (intra-site), corporate-wide, and Internet traffic. We then further examine the geographical spread of traffic by mapping data flows to remote enterprise sites. We find that a) temporal patterns depend on the actual traffic classes, b) the majority of the traffic stays within the boundaries of the local site and traffic in the Internet class corresponds to the smallest fraction of all classes, and c) the distribution of the byte contributions per remote site appears heavy-tailed (section 3). We further show how the relevance of layer-4 TCP and UDP port numbers as application discriminators diminishes as we move our observation point away from the local site, to the corporate network, or the public Internet (section 4).

We highlight that defining categories of machines, such as clients, servers or proxies, to account for their traffic contributions seems meaningless within the enterprise. We find that there is a high spatial variability amongst hosts, which naturally suggests the identification of a set of "heavy" users, which contribute most to the overall traffic. This is consistent with characterizing host-behavior as drawn from a sub-exponential distribution (section 5). However, we note that the composition of this set changes with time – a consequence of spatial variability of hosts – and is in general application specific. On the contrary, we show that defining the set of the most "connected" hosts provides a more indicative feature of the functional role of each host in the network.

With respect to address dynamics our findings include the following: a) We observe that approximately one third of IP address to host name, and host name to IP address mappings do not provide a unique identification of hosts or IPs respectively (section 6.1). b) By analyzing DNS responses and distinguishing hosts that appear in various enterprise sites over time, we provide evidence that trip durations (from one remote site to another within the enterprise) follow an exponential distribution. On the contrary, the number of trips between specific site-pairs shows evidence of heavy-tailed distributions (section 6.2).

Practical significance. We believe that the implications of our observations are multifaceted. We highlight here the more direct ones: First, whereas intuitively a distinction of client and server machines may make sense for a single application, we observe that it does not for individual hosts within an enterprise network. Second, the observed extreme variability in the per-host load dictates that any system that attempts to reconstruct network-wide traffic load by sampling must track a very dynamic set of heavy users using a possibly nontrivial set of features. Third, a significant fraction of the traffic stays within the enterprise network, and while it will be opaque to the underlying providers of the network connectivity, it is still distributed far and wide through the network and around the globe. Fourth, address mappings reveal that identification of a host in an enterprise network might be challenging from a network trace alone. Finally, we show that recent findings in opportunistic communication settings [2, 12] seem to also apply when describing mobility within an enterprise network.

2. MEASUREMENT METHODOLOGY AND DATA TRACES

The results presented in this paper are based principally on a single corpus of packet data collected from the network at Microsoft Research Cambridge (*MSRC*). Fig. 1 presents an overall picture of the MSRC network,

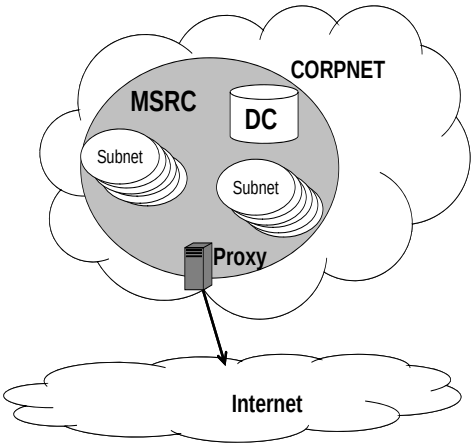

Figure 1: A view of the MSRC network.

and how it fits within the world-wide Microsoft Corporate network (henceforth, *CorpNet*), containing roughly 300,000 hosts connected by approximately 200 routers spread across 100 countries and 6 continents. The MSRC site contains roughly 400 hosts. Hosts run Microsoft operating systems and software suites, and the site contains a mixture of researchers, admin staff, human resources, and developers. The network runs the OSPF and BGP protocols for internal routing, connecting to the Internet through proxies in a small number of places predominantly in the USA.

In the remainder of this section, we will provide a brief description of the measurement methodology and the basic characteristics of the collected data.

2.1 Methodology

The corpus was collected over a period of 3.5 weeks beginning on Saturday August 27, 2005, and stored in 3.4 TB of disk. Packets were captured using custom tools written against the WinPCap [24] packet capture library. After a short testing period we settled on a snaplen of 152 B, as a compromise between total storage required and useful higher-layer header information captured, which was used for the duration of the capture. Our site network is configured with each IP subnet corresponding roughly to a wing of a floor mapped to a single VLAN, and so packets were tapped from the network using VLAN-spanning [3] on our site router. Each VLAN was mapped to one of two ports for load balancing reasons, resulting in two trace sets of roughly 5200 and 3200 0.5 GB files (or *chunks*) respectively.

VLAN spanning applies to all ports on the router and copies each packet that matches against the given VLAN tags to a further designated port. Because of the particular configuration of VLANs in our building we collected duplicate packets both *within* and *between* chunks, since many packets travel between two VLANs that either span to the same collection port, or span to both collection ports. Note that as the "duplicated" packet has been routed in the interim, it is not a byte-for-byte duplicate: for example, its IP time-to-live will have been decremented and its header checksum adjusted. As a result, two further processing stages were applied to remove duplicates within chunks, and then to merge these de-duplicated chunks. The process of merging the two traces also synchronized the timestamps from the two traces, using TCP SYN packets with the same 5-tuple and with equal sequence numbers as synchronization points.

CorpNet uses IPSEC for authentication [11] and so there is a substantial amount of ESP traffic ($\sim$ 85% by packet) in the captures. As a result, the capture tool manipulated captured packets to remove the ESP encapsulation, replacing the IP protocol value of 0x32 (ESP) with the IP protocol value stored in the *next_proto* field of the ESP trailer, before writing the packet to disk. For each packet so manipulated, the IP header checksum was zeroed to signify that the packet was originally ESP.

This process only captures packets observed at the site router, i.e., that enter or exit the site, or that are routed between subnets. To estimate how much traffic we could *not* observe as a result (i.e., intra-subnet traffic), we configured port spanning on one of the Ethernet switches servicing one wing of one floor of our building, producing a third trace. Comparing the data in this trace with the data observed from that subnet in the main corpus, we observe that <0.1% of both packets and bytes from that part of the site were not seen at the main site router. As such, we conjecture that our measurement methodology allows us to observe roughly 99% of all bytes and packets transmitted or received by each end host within MSRC. The bulk of the traffic missed corresponds to local file-sharing traffic and documents being transmitted to local printers, with the remainder made up from remote desktop protocol (RDP, used for remote machine access in Windows), name resolution (particularly NetBIOS name lookup), and ISA Key Management (a necessary component of our IPSEC deployment).

Overall, the end result of the all of the aforementioned steps is a trace containing packets forming a data corpus of 13 billion unique packets covering 12.5 TB of data.

Finally, we further extended the analysis of the main corpus by using address allocation information and router configuration files. In particular, we inferred geographic information by extracting the OSPF configuration blocks from the router configuration files. Each such configuration block contains the IP subnets that are originated by the OSPF process at that router. Note that all routers in CorpNet are named according to a convention which encodes their location by city and country and thus traf-

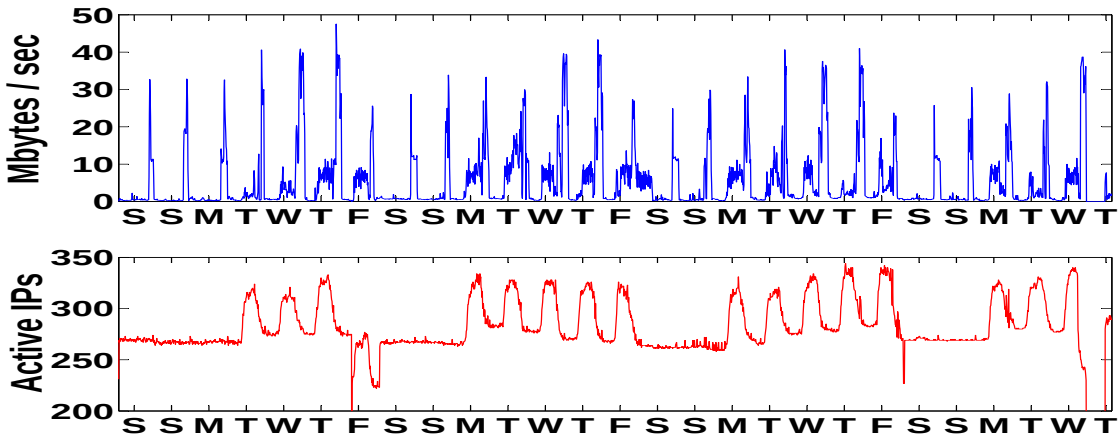

Figure 2: Overall traffic over time and the number of active IPs estimated over 15-minute intervals. The x-axis presents the day of week, with the ticks corresponding to 10am local time (GMT). Diurnal patterns and regular backups are evident in the time-series. The drops observed are due to network maintenance. The first Monday in the trace was a local public holiday.

fic sources and destinations can be straightforwardly mapped to cities.

Although the corpus and the tools cannot be made publicly available at this point, we are happy to consider requests to do so in the future, and we welcome applications from interns and other requests for research collaboration analyzing the corpus. Details of the corpus will be posted in the Internet Measurement Data Catalog (IMDC) `http://www.datcat.org/` in the near future.

2.2 Data corpus

To analyze the collected trace, we constructed flow tables corresponding to 5-minute time intervals, with one record per uni-directional 5-tuple (source IP, destination IP, protocol, source port, destination port) flow observed in each 5-minute period. Each record contains the 5-tuple, the time of the first and last packets in the period, the number of bytes and packets observed, and the application inferred to have generated the traffic. Application was inferred by deep packet inspection, with care taken to track and account for MSRPC invocations appropriately (see section 4 for discussion of MSRPC traffic). Overall, less than 3.5% of the packets in the trace could not be assigned to an application.

We observed 34,397 unique IP addresses in the trace, 591 of which were local to the capture site, MSRC. Of the observed addresses, 23,696 were sources (514 of which were local to MSRC) and 33,885 were destinations (582 of which were local to MSRC). The 77 local addresses that received but never transmitted appear to be the result of automated security tools probing for active addresses; similarly, we observed that 9 addresses only transmitted but never received, and appear to all be single-packet aberrations.

Fig. 2 shows the total traffic volume observed over the 3.5 weeks, as well as the number of active MSRC IPs calculated every 15-minutes. The ticks at the x-axis correspond to 10am GMT (i.e., local time). As expected, the traffic pattern observed at the collectors roughly follows the expected diurnal patterns. The large spikes during the early morning hours of each day correspond to backups in the Data Center (henceforth, *DC*, see Fig. 1), where the vast majority of the MSRC servers reside. The low bandwidth observed on the first Monday occurs due to it being a local public holiday.

The diurnal patterns are more evident in the number of active IPs over time. We observe a "baseline" of roughly 270 IPs that are always active (mostly desktops and servers), with the number of active IPs almost reaching 350 during the working hours. The four sudden short drops in the number of active IPs correspond to network maintenance windows in our local site.

In all, the corpus used in the remainder of this paper is a large, coherent set of data providing a useful window into the behavior of a type of network rarely studied previously.

3. TRAFFIC SPREAD

A distinguishing feature of enterprise networks, when compared to campus networks for example, is that they are both large in size and typically geographically distributed. Furthermore, their configuration, security concerns and the restrictions that these impose, dictate that only a small fraction of enterprise IPs are publicly routable.

Similarly, the configuration of the MSRC network is such that all traffic to the external Internet must be routed via a hierarchy of one or more proxies (Fig. 1). On the other hand, traffic internal to CorpNet, whether

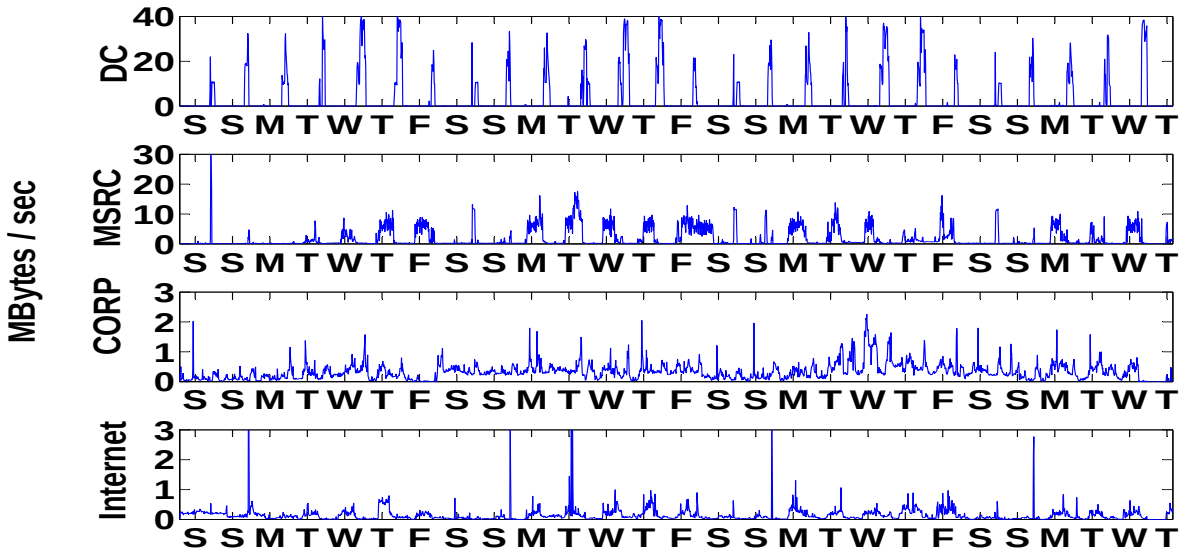

Figure 3: Traffic over time divided in 4 categories: Data center (DC), local traffic (MSRC), traffic to other enterprise sites (CORP) and traffic from and to the public Internet.

it remains within the MSRC site or it goes offsite to other Microsoft installations is routed directly. Thus, in this section we address the question of traffic spread, namely, *what is the network and geographical spread of traffic observed at a site in the enterprise network?*.

To this end, we formulate four different classes of observed traffic:

1. *DC*: Traffic that stays within the data center, and accounts mostly for the large overnight backups.
2. *MSRC*: Traffic that stays local within MSRC, excluding the DC traffic.
3. *CorpNet*: Traffic between MSRC and CorpNet, i.e., intra-enterprise traffic.
4. *Internet*: Traffic destined for or received from the public Internet.

The classes of traffic are separated based on subnet and proxy information. First, isolating DC traffic is straightforward since the data center corresponds to a separate subnet. Second, the use of subnets also denotes CorpNet traffic, while traffic that is received from or is destined to specific proxies reveals Internet traffic.

Fig. 3 presents what fraction of the traffic shown in Fig. 2 corresponds to each of the four traffic classes. Excluding DC traffic, we observe that on the average 79% of the overall traffic stays within MSRC, while CorpNet and Internet amount to only 14.5% and 6.5% of the total traffic respectively. The fact that traffic stays mostly within the enterprise has been observed before [20]. However, we provide here a further breakdown, by showing that the majority of the traffic is local

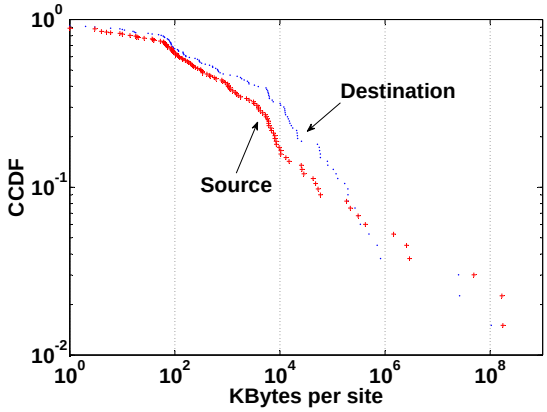

Figure 5: CCDF of traffic sourced at or destined to Microsoft enterprise sites. The distribution shows evidence of a heavy-tailed distribution with a few sites being the largest volume contributors.

within a site of the enterprise, with "intra-enterprise" traffic representing roughly one sixth of the total.

Comparing to Fig. 2, diurnal patterns are observed only for the MSRC and the Internet classes, while they are not as clear in the CorpNet traffic. Absence of such patterns in CorpNet traffic is due to the fact that this traffic class mostly reflects a set of applications that do not require user action (e.g., receiving email). The large occasional spikes in all classes correspond to large file transfers.

Similarly, Fig. 4 presents a breakdown with respect to the number of local IPs active for each class (i.e., for how many local IPs we observe flows from each traffic class). Here, diurnal patterns are evident in all classes except for DC where the set of active IPs is roughly

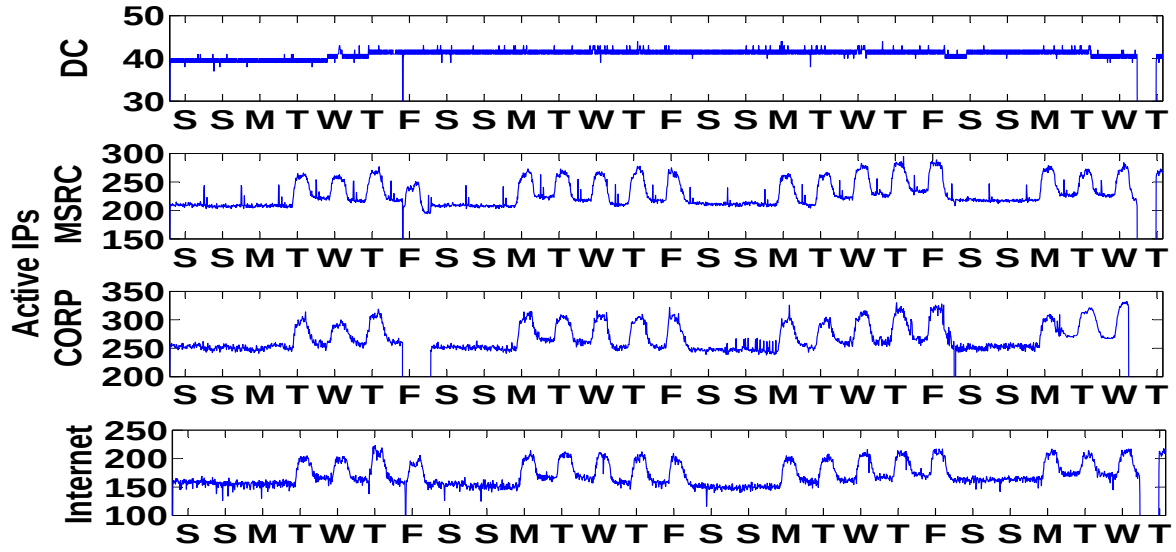

Figure 4: Number of active local IP over time divided in the 4 categories. DC IPs are constant over time representing always-on server machines.

constant over time. Note that the number of IPs in the CorpNet class is higher on the average when compared to the MSRC class. This occurs because of approximately 50 internal IPs that only communicate with other corporate non-local machines and represent networking equipment such as routers.

We further examined the spread of the traffic across the various sites of the enterprise using geographic information derived from router configuration files (see section 2.1). Specifically, we examined the fraction of the traffic sourced at or destined to a particular enterprise site, thus dividing the overall traffic to origin and destination flow pairs between MSRC and remote enterprise sites. Fig. 5 presents the Complementary Cumulative Distribution Function (CCDF) of traffic volumes across all sites observed distinguishing source and destination. Roughly 95% of the traffic is destined to or originating from our local Cambridge site, while the other largest contributors are two US sites, two sites within the UK and one within Europe. The empirical distributions appear heavy-tailed for a range of values (straight line in log-log scale), suggesting a few number of "heavy sites" with respect to their traffic contributions. Fitting however reveals that they are most likely a mixture of distributions since neither Pareto nor Log-normal fitting captures the tail of the distributions.

4. THE RELEVANCE OF PORT NUMBERS

Are port numbers relevant to identify the applications generating traffic within an enterprise network? Lately, there has been extensive evidence that port numbers alone appear insufficient to reveal the various Internet applications [13, 19], especially due to the increasing usage of peer-to-peer applications. However, we expect that due to the centrally-managed nature of enterprise networks, port numbers should allow for a clearer picture of application usage.

Indeed, the majority of bytes and packets in our corpus can be identified using port numbers. Fig. 6 presents the fraction of traffic that cannot be accounted for using port-number analysis. This fraction refers to flows which feature port numbers that do not correspond to any "known" services within the Microsoft enterprise network [1]. With the exception of DC traffic which cannot be characterized by port numbers, we observe a trend in that moving away from the local MSRC network obfuscates port based analysis. The reason for the unaccounted traffic is twofold: First, Windows does not use port numbers consistently in all cases with standard IANA port allocations [10]. Second, Windows makes extensive use of RPC invocation on COM objects for data transfers that results in dynamically allocated ports by the RPC endpoint mapper on the remote machine. This is particularly common for desktop applications such as Outlook.

To have a more complete picture of application usage, we extract the UUID naming of the target service from the MSRPC BIND packet fragments, and maintain a mapping from RPC service UUIDs to flows. In this way we can identify applications such as Microsoft Ex-

[1]Note that using "known" enterprise ports does not guarantee a correct measure of unknown traffic, since users may change the port of a known service. However, in contrast to typical Internet traffic, this would rarely be the case within the enterprise network.

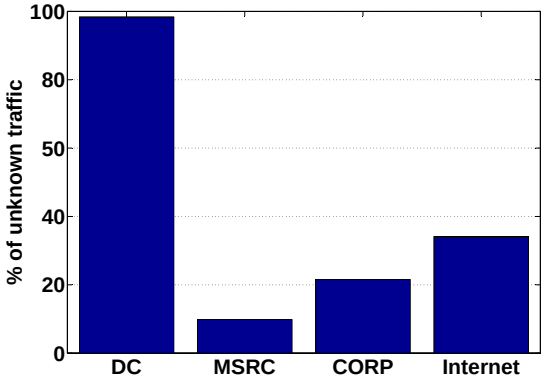

Figure 6: Fraction of unknown traffic according to port-based analysis. Almost all DC traffic cannot be categorized based on port numbers, while the unknown percentage increases as we move away from the local MSRC network.

	% Bytes		% Packets		% Flows	
Category	**On**	**Off**	**On**	**Off**	**On**	**Off**
Backup	**24.3**	0.01	**19.9**	0.01	0.50	0.01
Directory	**11.9**	6.16	9.92	4.77	**17.5**	9.14
Email	0.08	4.10	0.08	3.90	0.01	**10.6**
File	**24.4**	**42.4**	**31.9**	**50.5**	**18.6**	**11.6**
Management	0.62	0.95	1.30	2.11	**42.4**	**24.9**
Messenger	0.00	0.27	0.00	0.36	0.01	2.66
RemoteDesktop	0.02	0.24	0.14	0.70	0.06	0.03
RPC	0.01	0.32	0.05	0.43	1.94	2.36
SourceDepot	**36.4**	**13.5**	**32.4**	9.36	0.23	0.58
Web	1.47	**20.8**	2.36	**20.0**	**12.5**	**19.8**

Table 1: Application mix for on- and off-site traffic. Amounts ≥10% are shown in bold.

change email, Active Directory services, and the backup application (as suggested in section 2.1 only 3% of the total traffic cannot be classified to an application). Table 1 presents the volume of bytes, packets and flows of 10 application classes that cover over 98% of the overall byte volume. The on-site traffic refers to the DC and MSRC classes, while the off-site to the CorpNet and the Internet ones. The *File* category refers to applications such as SMB or Netbios, the *Directory* category to services like Active Directory or DNS, and the *Management* one to SMS, IGMP, BGP, NTP, etc. We observe that the heaviest on-site categories consist of *File*, and *SourceDepot* (source code control system), while *File* and *Web* are the off-site heaviest ones. Regarding offsite Internet traffic, it is dominated by HTTP (63%) while one third of the traffic cannot be classified as Fig. 6 also highlights.

5. DISTINGUISHING CLIENTS FROM SERVERS

Since the bulk of the traffic in the network belongs to client-server style applications, the assignation of particular hosts to either "clients" or "servers", i.e., identifying a hosts' functional role in the network, should be straightforward based on observing the byte contributions of the various hosts in the overall traffic. In our network, machines that are physically located inside data center tend to act predominantly as servers for one particular application. Intuitively, such main site servers (e.g., file servers, proxies, etc) should account for the majority of the traffic in the four classes.

To examine this hypothesis we examine the byte contributions per IP over time. To avoid aggregating over the whole trace hiding this way shorter time-scale effects, we limit the analysis in hourly intervals of the third week (which contains no network maintenance intervals). Fig. 7 shows the CDFs of the hourly average of downloaded and uploaded bytes per IP during week three of our trace for the MSRC, CorpNet and the Internet traffic classes (since DC is mostly server-to-server communications and backup traffic we do not examine this class further). Fig. 7(a) and Fig. 7(b) present the download and upload CDFs for the Internet class, along with the CDFs of the 5^{th} and 95^{th} percentiles per IP to provide some evidence of the overall variability across time, which appears nontrivial. Fig. 7(c) and Fig. 7(d) show the uploaded vs. downloaded contributions for the remaining two traffic classes (the percentiles appear similar to the Internet class and were omitted due to space limitations). Indeed, the figures reveal that a small subset of machines contributes most of the traffic overall both downstream and upstream in all traffic classes.

The existence of a small-number of "heavy" hosts in terms of their traffic contributions appears more evident if we plot the corresponding CCDFs in log-log scale, where a straight line would point towards a heavy-tailed distribution. Heavy-tailed as well as sub-exponential distributions decay more slowly than any exponential distribution. A sub-exponential distribution [8] is by definition one where the probability that the sum of independent random variables sampled from this distribution exceeds a threshold is equivalent to the probability that the maximum of these variables exceed the same threshold, when the threshold is large. Examples of such distributions include the Pareto and Lognormal distributions.

The CCDFs for the hourly average of uploaded and downloaded bytes per IP are shown in Fig. 8. Interestingly, while the MSRC and the Internet classes show signs of heavy-tailed distributions (i.e., the tails appear to follow a straight line), the CCDF of the CorpNet class shows less such evidence suggesting that traffic may be distributed more evenly among the local hosts in this class. This observation is especially evident in the upload case, where an exponential distribution appears a better fit for the data.

Intuitively, the set of heavy hosts should consist of the various server machines. Thus, tracking the specific servers over time should allow for a comprehensive view

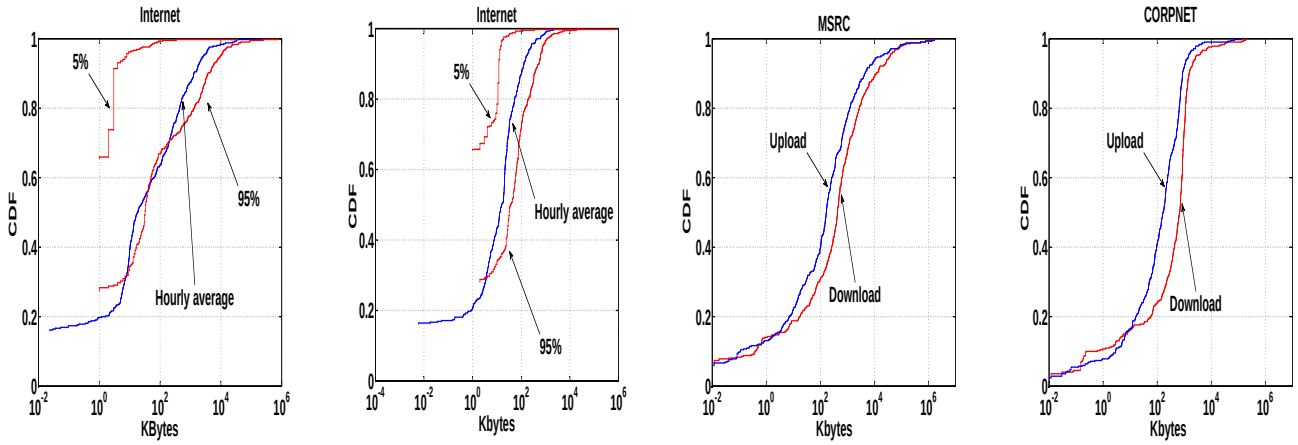

(a) Internet: Downloaded bytes per IP.

(b) Internet: Uploaded bytes per IP.

(c) MSRC: Downloaded & uploaded bytes per IP.

(d) CorpNet: Downloaded & uploaded bytes per IP.

Figure 7: (a) & (b) CDFs of hourly average, 5- and 95-percentile of byte contributions per IP for the Internet class. (c) & (d) CDFs of downloaded and uploaded hourly byte contributions per IP for MSRC and CorpNet.

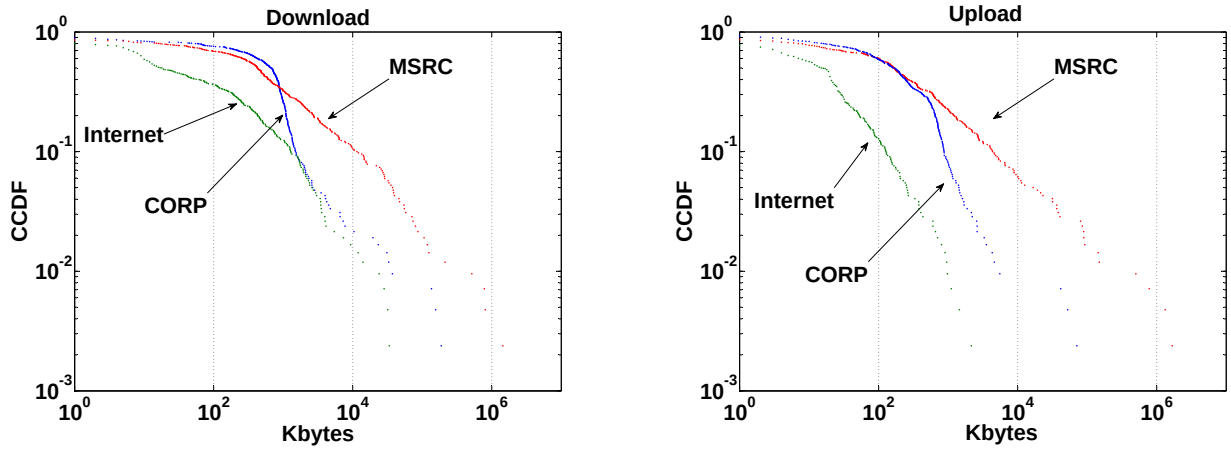

Figure 8: CCDFs of hourly averages for downloaded and uploaded bytes for MSRC, CorpNet and the Internet classes. The distributions appear heavy-tailed, with the exception of the CorpNet traffic which appears closer to the exponential distribution, especially for the upload case.

of the overall traffic volume across the various classes. Surprisingly, this hypothesis does not hold in our data. Examining the set of heavy hosts across time reveals that not only the set comprises both server and client machines, but it is also highly dynamic with its members significantly varying over time. For example, Fig. 9 describes the number of heaviest hosts required to account for 80% and 95% of the total traffic across time for the three classes; that is, tracking the cardinality of the set of heaviest hosts, where this set is defined as the machines required to capture x% of the overall traffic. In all cases, the set varies significantly over time, with diurnal patterns appearing only in the Internet traffic class. Fig. 9 suggests that attempting to predict the overall traffic volumes using a potentially static set of servers will not produce accurate estimates.

The above discussion suggests that categorizing hosts as either clients or servers in terms of traffic volumes is not straightforward. While intuitively such a distinction makes sense for a single application, it does not for individual hosts. There are two principle reasons for this: First, hosts invariably behave as both clients and servers in different applications, e.g., a web server will be a client to the directory and management services. Second, other applications may not strongly distinguish between clients and servers, e.g., in an enterprise network many machines may be clients of a central file-server while at the same time themselves acting as file-servers to other hosts. The conjecture that the characteristics of the distributions depend both on the traffic class and the various applications is also supported by Fig. 10, where the distributions for a number of applications for different classes are displayed (the applications were chosen to represent a significant

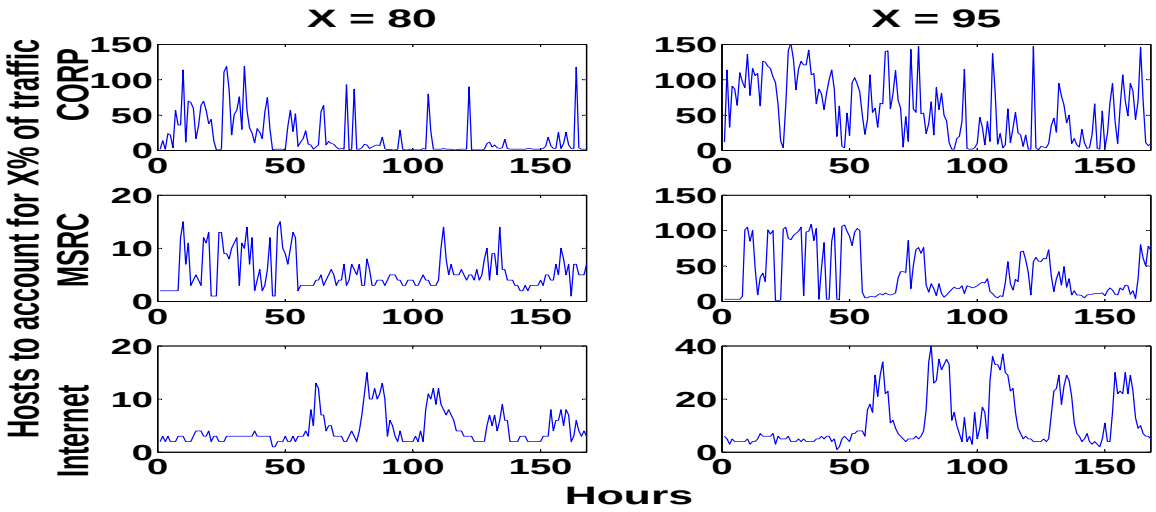

Figure 9: Number of heaviest host to account for 80% and 95% of the total traffic in MSRC, CorpNet and Internet classes. The cardinality of the set of heaviest hosts varies significantly over time.

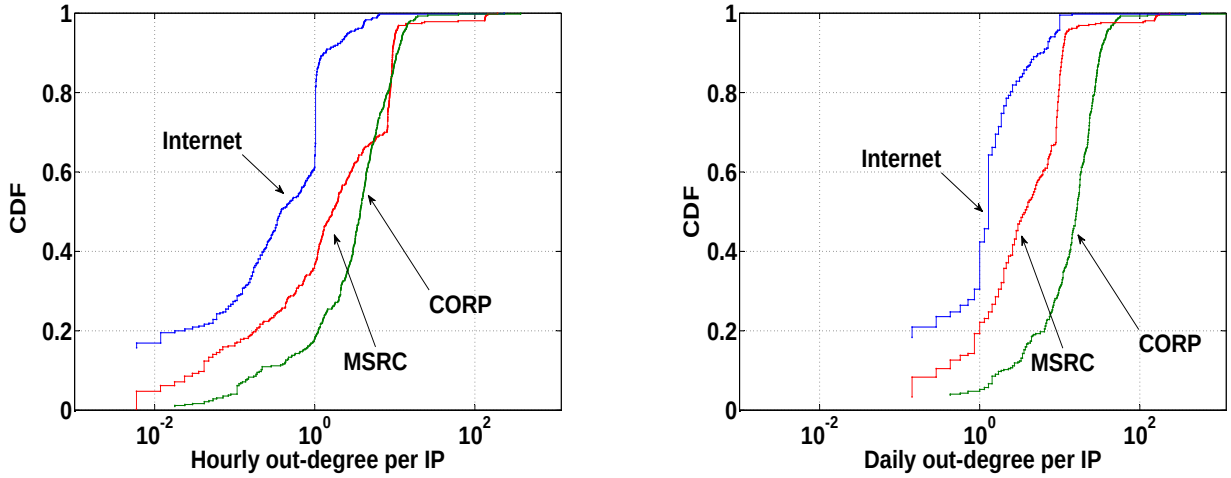

Figure 11: Out-degree of MSRC hosts for two different time-scales (hourly left, daily right). Internal hosts have roughly constant out-degree across time-scales for the MSRC class.

fraction of the traffic for each class). Note that even within the same class, the distribution of per host byte-contributions presents nontrivial variability (e.g., FILE and Directory application categories in MSRC).

While traffic volume is not an efficient distinctive feature to distinguish client from server hosts, activity of hosts appears more stable over time (see for example Fig. 4). In particular, connectivity information (i.e., which hosts communicate with one another) might allow for such a distinction, as intuitively servers should communicate with most of the local active clients. Thus, we define the out-degree of each host to be the number of other hosts it communicates with, and plot the corresponding CDFs in Fig. 11. This is a similar metric as the fan-out used in [20], where the authors observe that most hosts communicate with local hosts rather than non-enterprise ones. Fig. 11 reinforces this observation by using the three classes of traffic (e.g., out-degree of local hosts to other MSRC hosts, to CorpNet hosts and Internet hosts) and introduces a further separation of local to other enterprise offsite hosts, for which the out-degree appears similar as the local MSRC one. Fig. 11 also highlights the effect of the time-scale of observation by presenting the average out-degree measured at hourly and daily intervals. We observe that while the MSRC out-degree remains roughly the same across the two timescales, the CorpNet and Internet ones increase significantly especially for smaller values of the *x-axis*. This observation suggests that most internal hosts will show a roughly constant out-degree across time-scales for the MSRC class since the number of internal MSRC hosts that a host can communicate with is limited. Indeed, analysis across timescales reveals that within a busy hour, hosts will have an out-degree very

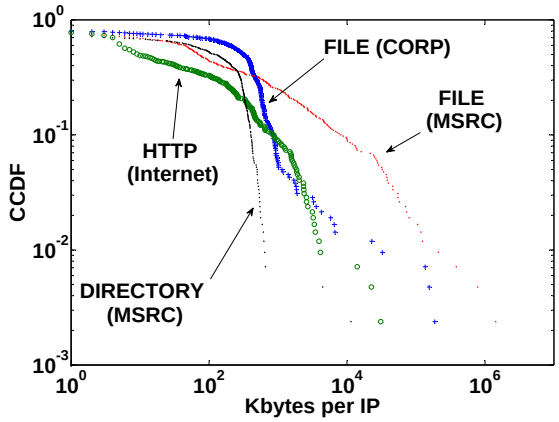

Figure 10: CCDFs of downloaded bytes per application for the three classes. Characteristics depend both on the application and the traffic class.

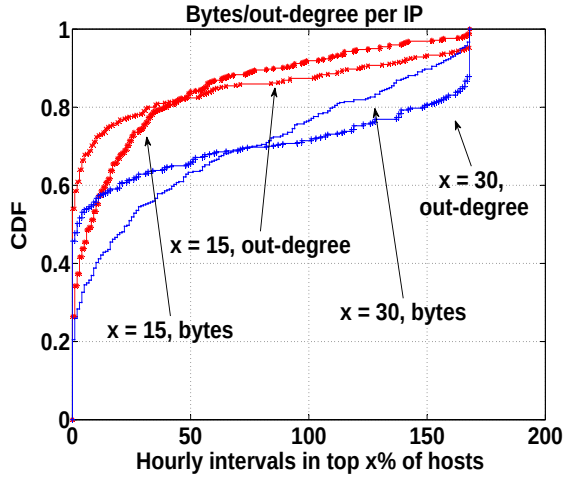

Figure 12: Prevalence of individual hosts in the most connected set and the heaviest hosts set in terms of bytes for hourly intervals during week 3. Connectivity produces a more stable set over time.

close to their maximum out-degree when considering the whole trace.

Close examination of Fig.11 reveals additionally a plateau in the distribution of the out-degree especially for the case of MSRC for larger values of the *x-axis*. This plateau points towards a set of hosts with comparable out-degrees that communicate with most of the internal hosts. Indeed, the IPs comprising this specific component of the distribution correspond to MSRC servers (e.g., proxies, domain controller, etc.) and is stable over time. We can further test this claim by looking at the prevalence of individual hosts in the set of most connected hosts. The prevalence is defined in a similar manner as in [21], and describes the number of intervals a host appears in the most connected set. We define this set as a percentage, x, of the most connected hosts and compare with the same percentage of the heaviest hosts in terms of bytes in Fig.12 (the two sets were calculated in hourly intervals for week three of the trace). We observe that connectivity provides a more stable set of the top hosts across time compared to the set of "heaviest" hosts. For $x = 30$, we observe that roughly 50% of nodes are never in the top-connected set (20% for bytes), while approximately 10% (less than 2% for bytes) of nodes are members of that set for all 168 hourly intervals of week three. Thus, most hosts in the connectivity case are either in or out of the "most-connected" set for all time intervals offering a clear distinction between client and server machines.

Summarizing the discussion throughout the section, we observe considerable spatial and temporal variability, that is, both across time and across hosts with respect to individual hosts' traffic contributions. While it seems intuitive to categorize hosts as either clients or servers based on the largest traffic contributors, examination of the data suggests this is not a fruitful approach to identify the functional role of enterprise network hosts. On the contrary, connectivity information appears as a more efficient alternative since hosts appear to essentially communicate with a stable set of other internal hosts, with server machines being the most connected ones.

6. ADDRESS AND MOBILITY DYNAMICS

In this section we examine address, host naming and host mobility dynamics within the enterprise. We first address the issue of how useful IPs are to uniquely identify hosts, and then we study host mobility patterns within the enterprise.

6.1 Name-address characteristics

It is unsurprising that a large enterprise network will provide wireless connectivity for employees' machines, and will usually allocate addresses via DHCP for all machines, wired and wireless. Thus, a host may be assigned multiple IP addresses over time, and also an IP address may be assigned to multiple hosts. The presence of services that are provided by clusters of machines via a single name further complicates matters. The result is that an IP address does not suffice to uniquely identify a host in general, although for most desktop hosts that connect solely to a wired network it will be a stable identifier. In this section we address the following questions: *What are the characteristics of the name-address mappings in the network? How often should an IP be considered as a unique identifier of an enterprise host?*

To answer these questions, we combine examination of router configuration files and DNS packet information. Specifically, we first parse DNS response packets captured in the corpus and we extract the time-varying mapping of names to addresses. Then, using the subnet

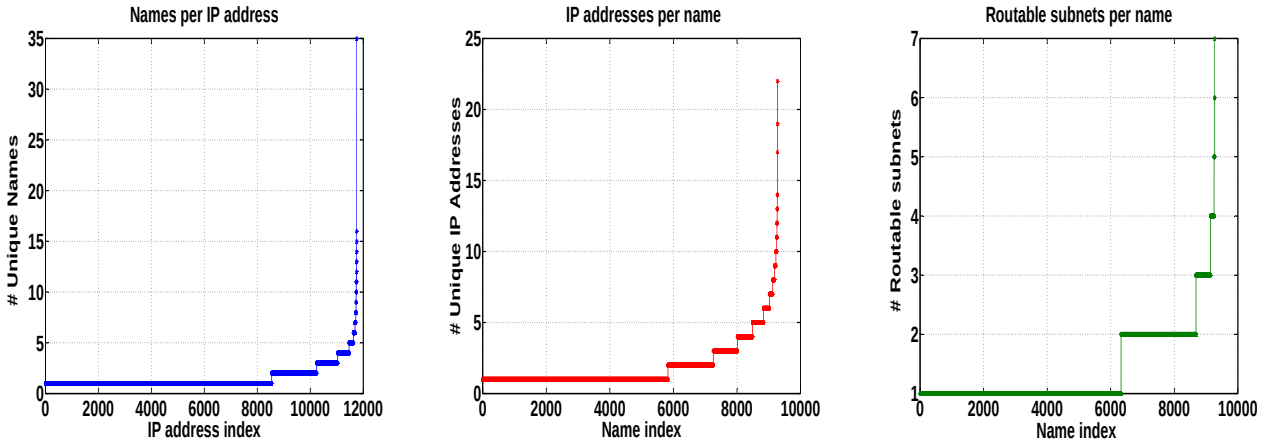

(a) Number of names per IP address. 73% of addresses map to a unique name. 63% map to a single IP address (b) Number of IP addresses per name. (c) Number of subnets per name. 7% of names map to more than two subnets.

Figure 13: Name-Address-Subnet mappings.

allocations obtained from the router configuration files, we can map the addresses in each response to their subnet. We assume that a host's *name* tends to change very infrequently and is thus static for the duration of the trace, allowing us to use a name as a stable host identifier. The nature of our corpus is such that we will see all DNS responses to hosts located on the MSRC network; further most local hosts are configured to utilize the local MSRC proxy, and thus we would still observe DNS packets when these hosts are off the MSRC network in a remote enterprise site. Hosts that temporarily visit the MSRC network will also have their details registered in the local MSRC name server.

Since the name to address mapping is not an one-to-one mapping, we examine three types of mappings:

1. *Name-address* mapping, that reveals the number of unique names per IP address.
2. *Address-name* mapping, that shows how many distinct IPs we observe for a unique name.
3. *Subnet-name* mapping, that describes the number of subnets a unique name has been associated with.

Name-address: Fig. 13(a) displays the characteristics of name-address mappings observed in the corpus. Of the 1,757 unique addresses that were returned as the result of some name resolution, 73% mapped to a unique name, the expected common case. Of the remainder, all but one mapped to 16 names or less, the outlier appearing to be the address of a machine hosting many services in a large datacenter which is thus accessed via a variety of names. The other addresses mapping to roughly 4 names appear to be natural churn due to DHCP. These cases include addresses belonging to subnets used for wireless connectivity or guest access, and so naturally have higher churn and shorter DHCP lease-times.

Address-name: Fig. 13(b) shows the characteristics of address-name mappings observed in the corpus. Of the 9,274 unique names observed, 63% map to a single address, again the expected common case. We cannot directly observe the purpose or intended use of a host, so explaining the reason why a third of hosts appear to have multiple addresses is difficult. From the subnets in question it appears that hosts with multiple addresses are either laptops with both wired and wireless addresses, or are in fact names that correspond to a service provided by a cluster of hosts (e.g., the web-proxy service provided for hosts within Europe).

Subnets-name: Fig. 13(c) shows the number of subnets that each name is mapped into, abstracting away the details of the DHCP assignments. The 63% names that map to a single address obviously map to only a single subnet. Of the other names 30% map to just two subnets, typically one wired and one wireless, which is common behavior for employees using a laptop as their main desktop machine. Finally, the rest of the hosts map to more than two subnets: these are probably laptops moving between sites and they amount to approximately 7% of all names.

The implication of these findings is that proper identification of a host in an enterprise network might be challenging from a network trace alone. Hosts can be accessed via multiple names and single names may map to multiple addresses concurrently (where the name really refers to a service rather than a host). Even when a single name only maps to a single address at any point in time, that address may change either because the name refers to a host which leaves the network for sufficiently long time so that DHCP cannot reallocate the same address, or to a host which moves between subnets requiring a completely different address.

6.2 Host mobility characteristics

# Cities	# Hosts	# Countries	# Hosts
1	8782	1	8939
2	467	2	322
3	18	3	6
4	2	4	2

Table 2: The number of hosts appearing in how many cities and countries.

Following from the previous section, it seems that at least a number of hosts move around within the enterprise network. By mapping IP addresses to subnets and then to routers, and thus to cities and countries, we can observe the travel behavior and mobility patterns of particular hosts. We see that, as suggested by the majority of hosts having addresses within a single subnet, most hosts appear to remain tethered in a single location. However, roughly 6% of hosts appear to travel to different cities, and approximately 4% travel to different countries. In this section we ask the question: *how do hosts move around the network geographically?*

Lately, there has been an increased interest in human mobility patterns in the setting of Delay Tolerant Networking (DTN) and opportunistic communications [2, 12]. While the timescales of interest here are not comparable with these studies and the setting is different, our findings in this section provide evidence of similar observations in the context of mobility within an enterprise network. Application scenarios in this context may involve transfers of large volumes of data between two enterprise sites, where utilizing the actual network may simply not be as practical (e.g., transferring terabytes of captured network traces!).

To extract the geographical location of hosts, we take the previously obtained subnet mappings for their addresses, map them to their home routers and decode the city and country codes embedded in each router's name. Note that we remove from consideration hosts with names that are known to refer to clustered service implementations such as our proxies. Overall, we are left with 712,598 name service responses to examine, involving 9,269 names in 110 cities across 63 countries.

Table 2 shows the number of hosts that appear in different numbers of cities and countries. There is no obvious pattern as to the particular countries that are visited, although there is a bias towards the USA and countries in Europe, unsurprising as Microsoft is an American company and MSRC is one of its larger European sites from which many people interact with others around Europe. More detailed examination of the subnets involved reveals that while most hosts which move geographically are doing so on the wireless subnets, this need not be the case: some hosts do appear on the wired network in multiple locations.

Going further, we examine the changes in location (*trips*) visible from these data. Trips are defined as subsequent observations of a host in two different enterprise sites A and B. We can then define the *residence time* [2] at site A as $t_2 - t_1$, to represent the time a host spent in A, where t_1 is the first observation of a host in an origin site A, and t_2, with $t_2 > t_1$, is the first observation at the destination site B. Similarly, if a host follows a travel pattern of $A \rightarrow B \rightarrow A$, we regard as the *return* time to site A (i.e., how much time a host was away from A) to be the residence time at site B. Note that by the definitions above, we will not observe return times for all trips (e.g., when a host appears in more than one site before returning to its origin site A, or if we do not observe the host returning to the origin).

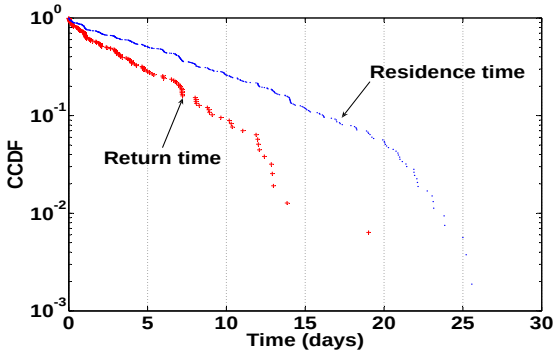

Figure 14: The CCDF of residence time at a site and return time to a site in lin-log scale. The straight lines point towards the exponential distribution with a mean of 5.5 days for residence and 3.8 days for return time.

We assume that trips with residence time less than 5 minutes are spurious and due to convergence among the many name servers in our network. We also deal with the complication of dual-ported hosts, such as laptops with a wireless and a wired interface and many server hosts with two wired interfaces. In such cases, the host may appear in two locations simultaneously, e.g., if a laptop normally based in Cambridge travels to Paris then it may connect to the wireless network in Paris, registering a Parisian address against its name in DNS. However, *at the same time* it may retain registration of its Cambridge address for its wired interface against its name, which will not change unless the laptop is plugged into the Paris wired network, causing it to receive a local address. Thus, the above situation may result in Cambridge and Paris addresses registered against the same laptop.

Since almost all CorpNet wireless addresses are allocated from a single subnet, we resolve this ambiguity by preferring the *wireless* address as a better indicator of the location of a host where there is a choice. In situations where both or neither addresses are wireless addresses, we simply prefer the first one we saw. We

[2]Note that residence time is an approximate metric since it also encompasses travel times, disconnections from the enterprise network, etc.

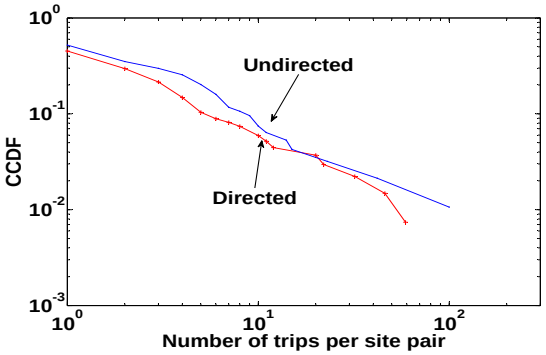

Figure 15: Characterizing host mobility: The CCDF of the number of trips per site-pair appears heavy-tailed both when accounting for the direction of the trip or not.

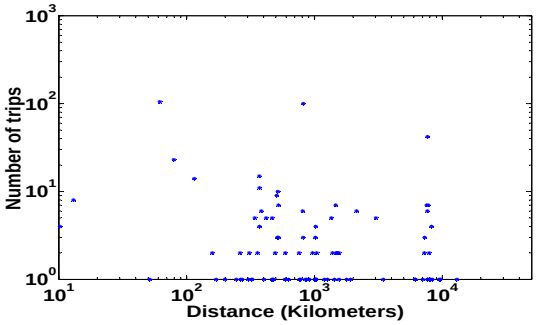

Figure 16: Distance vs. frequency of travel per site-pair. No obvious correlation exists, while the x-axis reveals typical travel distances (Cambridge to other UK sites, Europe and US).

observed 299 ambiguities due to duplication involving 189 unique names. We resolved 205 of these ambiguities by the wireless address heuristic above; the remaining 94 involved 40 unique names of which 25 appear to have multiple wired interfaces and 15 multiple wireless interfaces. The end result is 532 unambiguous observed trips, involving 344 unique names.

The distributions of residence time and return time follow the exponential distribution. Fig. 14 shows the CCDFs of residence and return time in days. The CCDFs are plotted in lin-log scale where a straight line is a sign of an exponential distribution as is the case in our data. We observe plateaus at approximately daily intervals as would be expected if trips are due to people visiting different sites. Slightly longer plateaus, indicating more trips, are observed at one, two, and three day boundaries, and at one and two week boundaries. We hypothesize that these plateaus are the result of common durations for business trips. Overall, approximately 38% of all residence times are less then three days, while the mean residence and return times are approximately 5.5 and 3.8 days respectively.

We further examine how connected the various enterprise sites are in a similar fashion to Fig. 5 in section 3. However, instead of looking at the traffic sourced from or destined to particular sites, we are interested here in connectivity between sites in terms of host trips. Thus, if we consider that the Microsoft enterprise network forms a graph, where the nodes are the various sites and links connect sites between which a trip existed, we can measure the relative importance of the links as the number of times a link was traversed. These links might be directed or not, in which case there is no distinction regarding the direction of travel between two sites.

The distribution of the number of trips per site-pair appears heavy-tailed. Fig. 15 presents the CCDF of the number of trips per site-pair in log-log scale for both the directed and the undirected case of the graph. The distribution follows a straight line in log-log scale suggesting that while most links are traversed only a few times, some routes do appear very frequent. While examination of the most popular routes reveals that distance might play some role in the frequency of trips per site-pair, Fig. 16 illustrates that no obvious correlation between distance and trip frequency appears in our data. Fig. 16 presents a scatter plot showing the effect of the actual distance between two sites on the frequency of trips. Distance is estimated through longitude and latitude coordinates of the various sites.

Overall, our observations are consistent with [12], where the authors observe an exponential tail for the distribution of human inter-contact times and that human contacts occur in a small number of locations. To our knowledge our observations offer the first evidence that mobility within the enterprise network appears to follow the same patterns, and provide some clues as to their potential time-scales of operation (e.g., the mean return time implies roughly three days of inter-contact time with the origin site).

Summarizing, our DNS packet analysis and router configuration information highlight a) interesting characteristics of name to address mappings, where approximately one third of all mappings (name to address, address to name) is not unique, and b) that host mobility within the enterprise appears to follow similar patterns to other DTN settings. Throughout this analysis, we also have to bear in mind the potential effect of observing the aforementioned properties for only a local site, thus having limited view of the global behavior and also through the specific methodology. With respect to the latter, examining DNS packets should only present a limited bias since we observe all DNS responses directed to hosts within the MSRC site, and a large number of DNS responses for local hosts when on travel. Estimating the exact fraction of DNS packets observed for local hosts when not in MSRC is not possible with our current dataset, since that would require instrumenting each host machine. Regarding the limited local view of the network, we believe that the properties observed

should still hold in the larger scale, assuming that such patterns in other sites are similar in nature. Confirmation of this hypothesis however is left as future work.

7. RELATED WORK

There have been many detailed studies of large-scale traffic traces at the packet and flow levels lately. Typically, they focus on Internet-wide analyses either using traces gathered from ISP networks [7] or from university campus networks using, for example, traces gathered at a university's border routers [22]. Other research has looked at characterizing coarse features or structures of wide-area traffic [1, 6, 14], inferring a traffic matrix from partial information [18], and predicting traffic [16].

Despite the fact that Leland et al.'s [15] Ethernet analysis in the early 1990s catalyzed a subsequent resurgence of interest in traffic measurement, characterization and analysis, relatively little effort has been expended considering LAN traffic. Similarly, remarkably few studies have considered modern enterprise network traffic although there is a large body of work exploring network behavior of particular application servers, including peer-to-peer [13, 22], Internet chat [5], and of course the Web [1]. Other work [4, 15] has looked at traffic characteristics of specific applications and network types, such as self-similarity or long-range dependence for Web and for Ethernet traffic.

Thus, although enterprise network configuration [17] and routing protocol behavior [23] have been examined, the only recently published work examining enterprise network traffic that we are aware is by Pang et al., using a dataset collected from the LBNL network [20]. Our work differs from [20] in various nontrivial aspects. First, the dataset used in that study differs significantly from the one we analyze here. Their collection was taken at a single central point in the LBNL network, but was restricted to monitoring two router ports at a time, rotating thus collection through different subnets. This resulted in 5 separate datasets spanning about 100 hours and covering 160 million packets in total. Our dataset is substantially larger, generated by continuous monitoring for over 500 hours from a single location towards the edge of our network, giving 12.8 billion packets. Our tracing methodology facilitates studying of almost all packets to and from every host in our local network for roughly 3.5 weeks. Second, our router information allows us to partition the observed traffic in finer classes than just enterprise and WAN, examining also the geographical spread and traffic dynamics to other enterprise sites. Third, we extensively study name to address and subnet mappings, and human mobility patterns within the enterprise network. Last but not least, the focus of our analysis here is notably different by studying address and traffic dynamics, rather than individual application characteristics.

8. CONCLUSIONS

Throughout this paper, we have posed and answered, through the analysis of a substantial collected trace, a series of questions with regards to characteristics that describe the underlying dynamics of modern enterprise networks. The nature of such typically geographically distributed networks that offer numerous diverse services to several clients worldwide renders them remarkably different with specific idiosyncrasies compared to traditional Internet traffic.

Our empirical evidence suggest that a) the majority of the observed traffic stays within the enterprise network with approximately one sixth representing inter-site traffic within the enterprise, b) the meaning of layer-4 port numbers is less helpful to characterize applications as we move away from the local site, and c) significant spatial and temporal variability render valueless the distinction between clients and servers by traffic volume. We further observe that even the identification of an appreciable fraction of individual hosts is not straightforward, complicated by extensive use of DHCP, and that mobility patterns within the enterprise follow properties observed in settings such as opportunistic communications. We believe that our observations provide valuable insights regarding the primary properties of enterprise networking to the research community, for whom such data is rarely accessible.

9. REFERENCES

[1] P. Barford and M. Crovella. Measuring web performance in the wide area. In *Performance Evaluation Review, special issue on network traffic measurement and workload characterization*, 1999.

[2] A. Chaintreau, P. Hui, J. Crowcroft, C. Diot, R. Gass, and J. Scott. Impact of Human Mobility on the Design of Opportunistic Forwarding Algorithms. In *Infocom*, 2006.

[3] Cisco. `http://www.cisco.com/`, May 2007.

[4] M. E. Crovella and A. Bestavros. Self-similarity in World Wide Web traffic: Evidence and possible causes. *IEEE/ACM Transactions on Networking*, 5(6):835–846, Dec. 1997.

[5] C. Dewes, A. Wichmann, and A. Feldmann. An analysis of Internet chat systems. In *Proceedings of ACM/Usenix Internet Measurement Conference (IMC) 2003*, Nov. 2003.

[6] N. Duffield, C. Lund, and M. Thorup. Properties and prediction of flow statistics from sampled packet streams. In *Proceedings of ACM/Usenix Internet Measurement Workshop (IMW) 2002*, pages 159–171.

[7] C. Fraleigh, C. Diot, B. Lyles, S. Moon, P. Owezarski, D. Papagiannaki, and F. Tobagi. Design and deployment of a passive monitoring infrastructure. In *Passive and Active Measurement Workshop*, 2001.

[8] C. Goldie and C. Kluppelberg. *A Practical Guide to Heavy Tails: Statistical Techniques for Analysing Heavy Tails*, chapter 'Subexponential distributions'. Birkhauser, 1997. citeseer.ist.psu.edu/goldie97subexponential.html.

[9] T. Henderson, D. Kotz, and I. Abyzov. The changing

usage of a mature campus-wide wireless network. In *MobiCom '04: Proceedings of the 10th annual international conference on Mobile computing and networking*, pages 187–201, NY, USA, 2004.

[10] IANA port numbers. `http://www.iana.org/assignments/port-numbers`.

[11] Security Architecture for the Internet Protocol. RFC 2401, November 1998.

[12] T. Karagiannis, J.-Y. L. Boudec, and M. Vojnovic. Power law and exponential decay of inter contact times between mobile devices. Technical Report MSR-TR-2007-24, Microsoft Research, March 2007.

[13] T. Karagiannis, A. Broido, M. Faloutsos, and K. Claffy. Transport layer identification of p2p traffic. In *Proceedings of ACM/Usenix Internet Measurement Conference (IMC) 2004*, Nov. 2004.

[14] A. Lakhina, K. Papagiannaki, M. Crovella, C. Diot, E. Kolaczyk, and N. Taft. Structural analysis of network traffic flows. In *Proceedings of ACM SIGMETRICS 2003*, June 2004.

[15] W. Leland, M. Taqqu, W. Willinger, and D. Wilson. On the self-similar nature of Ethernet traffic. In *Proceedings of ACM SIGCOMM 1993*, pages 183–193, New York, NY, USA, 1993.

[16] G. Liang, N. Taft, and B. Yu. A fast lightweight approach to origin-destination IP traffic estimation using partial measurements. In *IEEE Transactions on Information Theory, June 2006*.

[17] D. Maltz, G. Xie, J. Zhan, H. Zhang, G. Hjalmtysson, and A. Greenberg. Routing design in operational networks: a look from the inside. In *Proceedings of ACM SIGCOMM 2004*, Aug. 2004.

[18] A. Medina, N. Taft, K. Salamatian, S. Bhattacharyya, and C. Diot. Traffic matrix estimation: existing techniques and new directions. In *Proceedings of ACM SIGCOMM 2002*, pages 161–174, New York, NY, USA, 2002. ACM Press.

[19] A. W. Moore and K. Papagiannaki. Toward the accurate identification of network applications. In *Proceedings of the 6th Passive and Active Measurement Workshop (PAM'05)*, 2005.

[20] R. Pang, M. Allman, M. Bennett, J. Lee, V. Paxson, and B. Tierney. A first look at modern enterprise traffic. In *Proceedings of ACM/Usenix Internet Measurement Conference (IMC) 2005*.

[21] V. Paxson. End-to-end routing behavior in the Internet. In *ACM SIGCOMM Computer Communication Review*, volume 26,4, pages 25–38, August 1996.

[22] S. Saroiu, K. Gummadi, R. Dunn, S.D.Gribble, and H. Levy. An analysis of Internet content delivery systems. In *Proceedings of the Fifth USENIX Symposium on Operating Systems Design and Implementation (OSDI) 2002*.

[23] A. Shaikh, C. Isett, A. Greenberg, M. Roughan, and J. Gottlieb. A case study of OSPF behaviour in a large enterprise network. In *Proceedings of ACM/Usenix Internet Measurement Workshop (IMW) 2002*, Nov. 2002.

[24] WinPcap: The Windows Packet Capture Library. `http://www.winpcap.org/`.