Analysis of traffic data from a hybrid satellite-terrestrial network

Savio Lau
saviol@cs.sfu.ca

Communication Networks Laboratory
http://www.ensc.sfu.ca/research/cnl
School of Engineering Science
Simon Fraser University

communication networks laboratory

Roadmap

- Introduction
- ChinaSat: network architecture, TCP, and network anomalies
- Mathematical tools for statistical analysis
- Analysis of billing records:
  - aggregated traffic
  - user behavior
- Analysis of tcpdump traces:
  - general characteristics
  - TCP options and OS fingerprinting
  - network anomalies
- Conclusions and future work

Introduction and motivation

- Analysis of traffic data enables:
  - understanding of traffic dynamics
  - characterization and development of new traffic models
  - evaluation of network performance
- Most traffic data are collected at research institutions or from research networks:
  - traffic data from commercial networks are rare
  - commercial network traffic may have different characteristics compared to research networks
- Analysis of traffic data from a commercial network such as the ChinaSat DirecPC network is important

Previous work

- Previous analysis of network traffic focused on:
  - characteristics of TCP connections
  - network traffic patterns
  - statistical and cluster analysis of traffic
  - anomaly detection:
    - statistical methods
    - wavelets
    - principle component analysis

Previous work on the ChinaSat data

- ChinaSat traffic is self-similar and non-stationary
- Hurst parameter depends on traffic load
- Modeling TCP connections:
  - inter-arrival time is best modeled by the Weibull distribution
  - number of downloaded bytes is best modeled by the lognormal distribution
- The distribution of visited websites is best modeled by the discrete Gaussian exponential (DGX) distribution

Q. Shao and Lj. Trajkovic, "Measurement and analysis of traffic in a hybrid satellite-terrestrial network," *Proc. SPECTS 2004*, San Jose, CA, July 2004, pp. 329–336.

Previous work on the ChinaSat data

- Traffic prediction:
  - autoregressive integrative moving average (ARIMA) can be used to predict uploaded traffic but not downloaded traffic
  - wavelet + autoregressive model outperforms the ARIMA model

Contributions: analysis of billing records

- Analysis of patterns and statistical properties of two sets of data from the ChinaSat DirecPC network: billing records and tcpdump traces
- Billing records:
  - daily and weekly traffic patterns
  - user classification:
    - single and multi-variable k-means clustering of traffic volume (packets and bytes)
    - hierarchical clustering of user activity (refined using the three most common traffic patterns)
    - combination of k-means and hierarchical clustering

Contributions: analysis of tcpdump trace

- tcpdump trace:
  - analysis of protocols and applications
  - analysis of TCP options
  - operating system fingerprinting
  - detection of network anomalies
- Developed C program pcapread:
  - processes tcpdump files
  - produces custom output
  - eliminates the need for packet capture library libpcap

Roadmap

- Introduction
- ChinaSat: network architecture, TCP, and network anomalies
- Mathematical tools for statistical analysis
- Analysis of billing records:
  - aggregated traffic
  - cluster analysis
- Analysis of tcpdump traces:
  - general characteristics
  - TCP options and OS fingerprinting
  - network anomalies
- Conclusions and future work

ChinaSat hybrid satellite network

- Employs geosynchrous satellites deployed by Hughes Network Systems Inc.
- Provides data and television services:
  - DirecPC (Classic): unidirectional satellite data service
  - DirecTV: satellite television service
  - DirecWay (Hughnet): new bi-directional satellite data service that replaces DirecPC
- DirecPC transmission rates:
  - 400 kb/s from satellite to user
  - 33.6 kb/s from user to network operations center (NOC) using dial-up
- Improves performance using TCP splitting with spoofing

Characteristics of geosynchronous satellite links

- Large coverage area
- High bandwidth
- Long propagation delay
- Large bandwidth-delay product
- High bit error rates:
  - 10^{-6} without error correction
  - 10^{-3} or 10^{-2} due to extreme weather and interference
- Path asymmetry

DirecPC system diagram

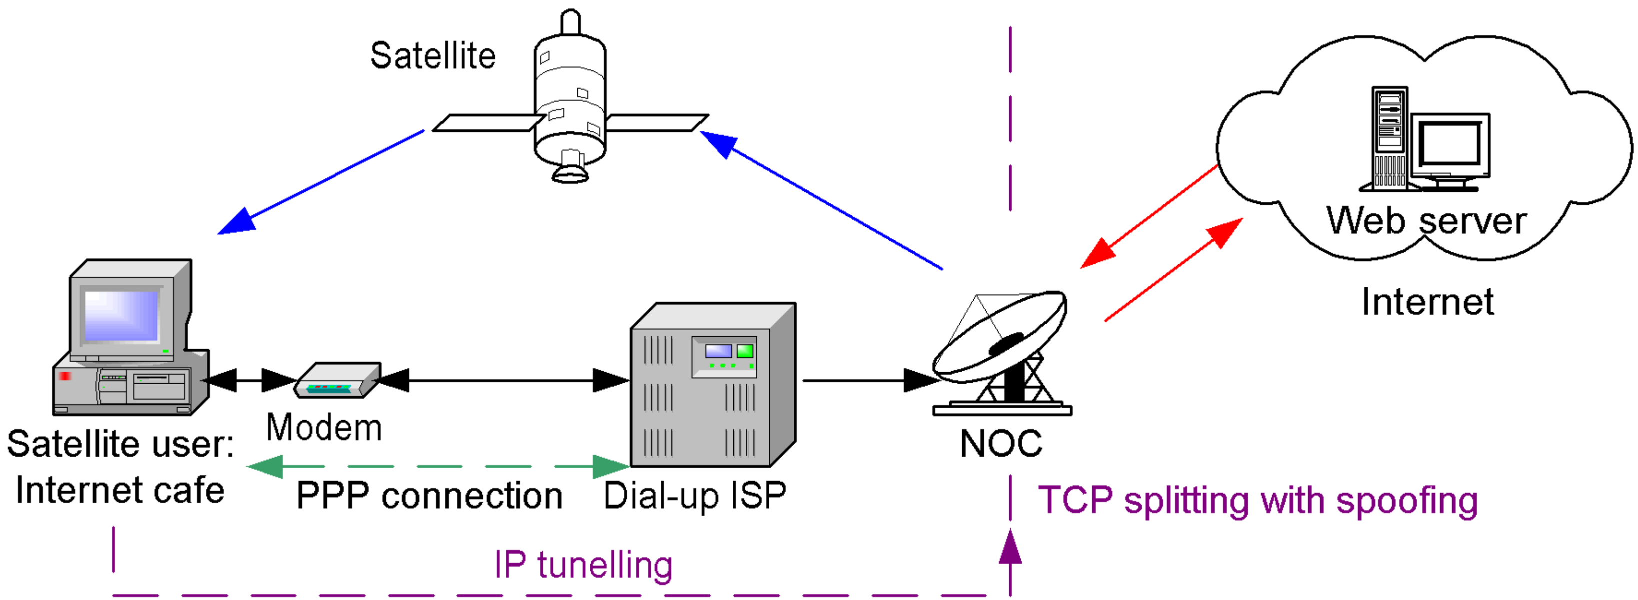

NOC: Network operations center
PPP: Point-to-point protocol

TCP extensions for satellite environments

- Increasing initial TCP congestion window (cwnd)
- Selective acknowledgement option:
  - enables a TCP receiver to acknowledge out-of-order packets
  - allows a TCP sender to identify and retransmit lost segments
  - avoids the performance penalty associated with retransmission timeouts
- Performance enhancing proxies (PEPs):
  - improve TCP performance in specific link environments
  - violate TCP end-to-end semantics
  - example: TCP splitting with spoofing

TCP extensions for satellite environments

- TCP sliding window scale option:
  - expands default TCP window from 16 bits to 32 bits
  - allows greater number of unacknowledged packets
- Path maximum transmission unit (MTU) discovery:
  - determines the maximum allowable size in links between source and destination
  - enables TCP senders to reach maximum throughput earlier

TCP splitting with spoofing

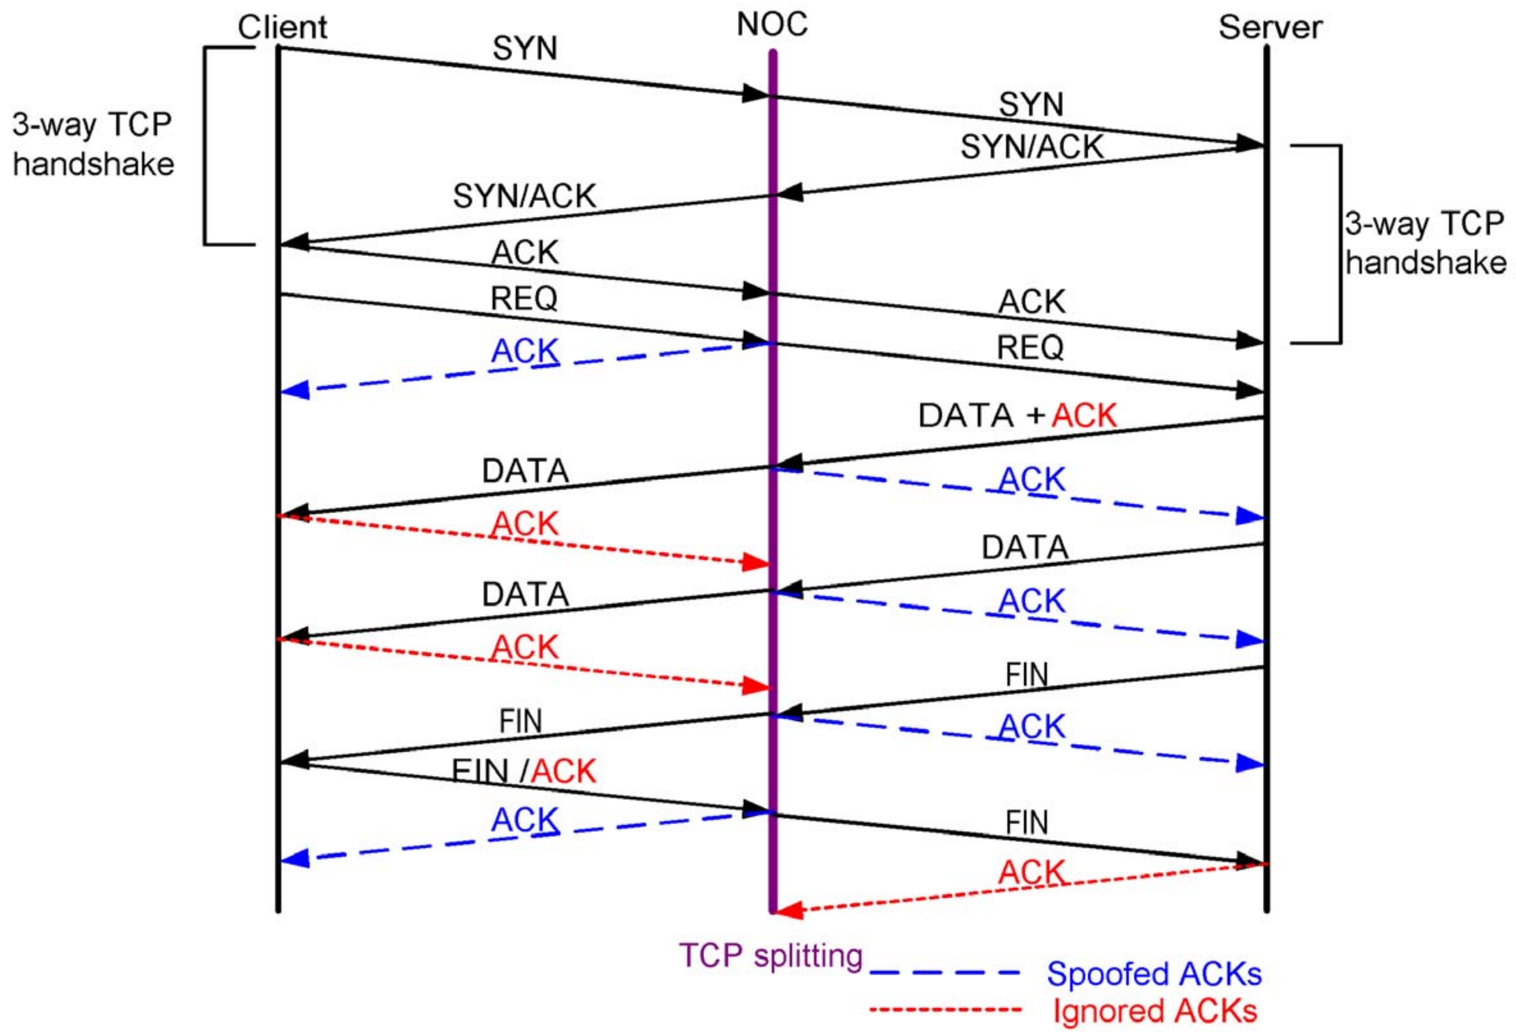

Network anomalies

- Scans and worms:
  - packets are sent to probe network hosts
  - used to discover and exploit resources
- Traffic volume anomalies:
  - significant deviation of traffic volume from usual daily or weekly patterns
  - classified as:
    - outages: caused by unavailable links, crashed servers, or routing problems
    - short term increases in demand: caused by short term events such as holiday traffic
  - involve multiple sources and destinations

Network anomalies

- Flash crowd:
  - high volume of traffic destined to a single destination
  - caused by breaking news or availability of new software
- Traffic shift:
  - redirection of traffic from one set of paths to another
  - caused by route changes, link unavailability, or network congestion

Network anomalies

- Alpha traffic:
  - unusually high volume of traffic between two endpoints
  - caused by file transfers or bandwidth measurements
- Denial of service:
  - large number of packets directed to a single destination
  - makes a host incapable of handling incoming connections or exhausts available bandwidth along paths to the destination

Roadmap

- Introduction
- ChinaSat: network architecture, TCP, and network anomalies
- Mathematical tools for statistical analysis
- Analysis of billing records:
  - aggregated traffic
  - cluster analysis
- Analysis of tcpdump traces:
  - general characteristics
  - TCP options and OS fingerprinting
  - network anomalies
- Conclusions and future work

Cluster analysis

- Algorithms to group data objects
- Maximization of intracluster similarity and minimization of intercluster similarity
- Goodness of results are measured by cluster quality
- Two methods are employed:
  - partitioning clustering (k-means)
  - hierarchical clustering

Partitioning clustering

- Constructs k partitions of the data from n objects, where $k \leq n$
- Two constraints:
  - each cluster must contain at least one object
  - each object must belong to exactly one group
- Requires exhaustive enumeration of all possible combinations to find the optimal cluster solution
- Heuristic methods such as the k-means algorithm are used in practice

k-means clustering

- Generates k clusters from n objects
- Requires two inputs:
  - k number of desired partitions
  - n objects
- Uses random placement of initial clusters
- Determines clustering results through an iteration technique to relocate objects to the most similar cluster:
  - similarity is defined as the distance between objects
  - objects that are closer to each other are more similar
- Computational complexity of O(nkt), where t is the maximum number of iterations

k-means clustering algorithm

1. Randomly select k objects to be the center of k clusters.
2. Assign each remaining object to the cluster to which it is the most similar.
3. Recalculate the cluster mean after all objects are (re)assigned.
4. Re-evaluate all objects and place them in the cluster to which they are the most similar.
5. Repeat Steps 3 and 4 until no changes have been made (full convergence) or the maximum number of iterations are reached (partial convergence).

Measuring cluster quality

- Silhouette coefficients (SC) may be used to measure cluster quality
- SC of object i (s_i) is defined as:

$$s_i=(b_i-a_i)/max(b_i-a_i)$$

  - a_i is the average distance from object i to all other objects in the same cluster A
  - b_i is the minimum of average distances from object i to all other objects in clusters B, where B ≠ A
- $0.7 < SC \leq 1.0$ indicates high cluster quality
- $0.5 < SC \leq 0.7$ indicates medium cluster quality
- $0.25 < SC \leq 0.5$ indicates low cluster quality
- $SC \leq 0.25$ indicates the absence of cluster structure

Finding natural number of clusters

- The natural number of clusters k is not known a priori
- k-means algorithm is repeated for different k values
- Natural number of clusters is found by comparing average SC value for various values of k:
  - average SC is calculated for all objects
  - the natural number of clusters k is found at the local maxima

SC: silhouette coefficient

Hierarchical clustering

- Objects are grouped into a tree of clusters (dendrogram)
- Two approaches: agglomerative and divisive
- Agglomerative approach (bottom-up):
- Divisive approach (top-down)
- Clusters are merged (or split) based on distance measure
- Four distance measures are commonly employed: minimum, maximum, mean, and average

Distance measures

- Clusters are merged (or split) based on distance measure
- Four distance measures are commonly employed:
  - minimum: distance of two closest objects p_i and p_j, where p_i ε cluster C_i and p_j cluster C_j
  - maximum : distance of two farthest objects p_i and p_j, where p_i ε cluster C_i and p_j cluster C_j
  - mean: distance between the centroid of C_i and C_j
  - average: average distance of objects in C_i to objects in C_j

Distance measures

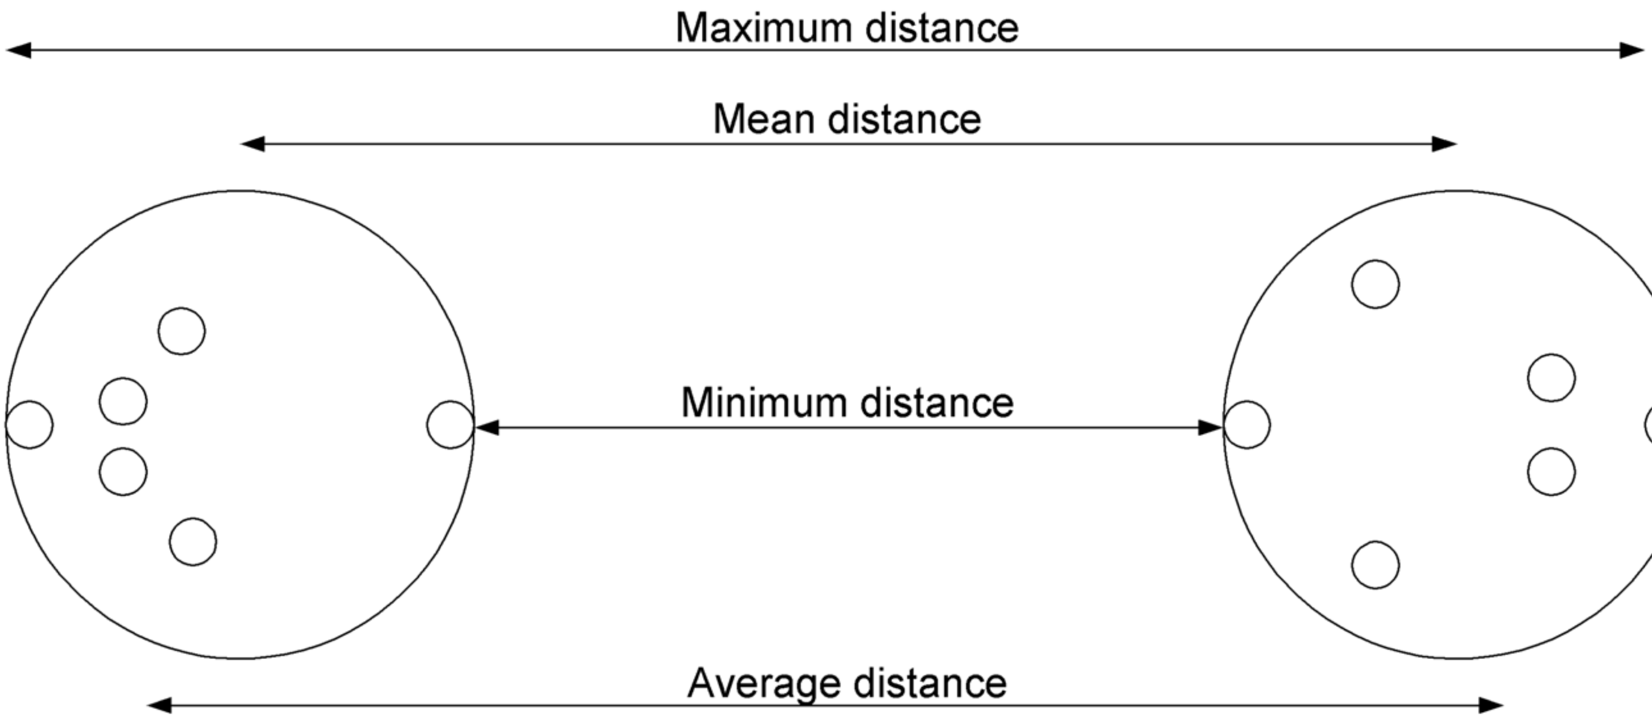

Agglomerative hierarchical clustering algorithm

1. For n objects, a similarity matrix of n x n is generated. Each value records the distance between the two objects or (the number of identical values if a series of values is used)
2. Objects are assigned to clusters from 1 to n.
3. Each iteration merges two clusters that are closest to each other (minimum similarity value)
4. Repeat steps 2 and 3 until all objects are merged into a single cluster or until termination condition is reached.
5. Groups can be found by selecting k or selecting a maximum merge distance.

Measuring cluster quality in hierarchical clustering

- Cophenetic correlation coefficient (CPCC):
  - correlation between the cophenetic distance matrix and similarity matrix
  - used to determine the best distance measure
- Cophenetic distance:
  - defined as the distance between two objects to their common parent
  - measures the mismatch between the distance in the similarity matrix and the distance between clusters
- Higher CPCC values indicate better clustering results

Calculation of CPCC

$$\text{CPCC} = \frac{\sum_{i<j} (Y_{ij} - y)(Z_{ij} - z)}{\sqrt{\sum_{i<j} (Y_{ij} - y)^2 \sum_{i<j} (Z_{ij} - z)^2}}$$

- Y = actual distances between objects
- Z = distances between objects in the hierarchical tree
- Y_{ij} = distances between objects i and j in Y
- Z_{ij} = distances between objects i and j in Z
- y = average distance of all of objects in Y
- z = average distance of all objects in Z

Hierarchical clustering: final clustering results

- Visualized by dendrograms
- Determined by two choices:
  - desired number of clusters k
  - selected cutoff based on inconsistency coefficients:
    - inconsistency coefficient is the difference between the height of a dendrogram link and the average height of links at the same level
    - links connecting two distinct clusters have higher inconsistency coefficient

Calculation of inconsistency coefficients

$$IC = \frac{Z_{ij} - \mu_{z\text{ considered}}}{\sigma_{z\text{ considered}}}$$

- Z_{ij} = link distances between objects i and j in the hierarchical tree Z
- $\mu_{z\text{ considered}}$ = mean of link distances considered in the calculation:
  - links considered are defined as links at the same level as Z_{ij} and links up to depth d below
  - d is chosen as 2
- $\sigma_{z\text{ considered}}$ = standard deviations of link distances considered in the calculation

Dendrogram example

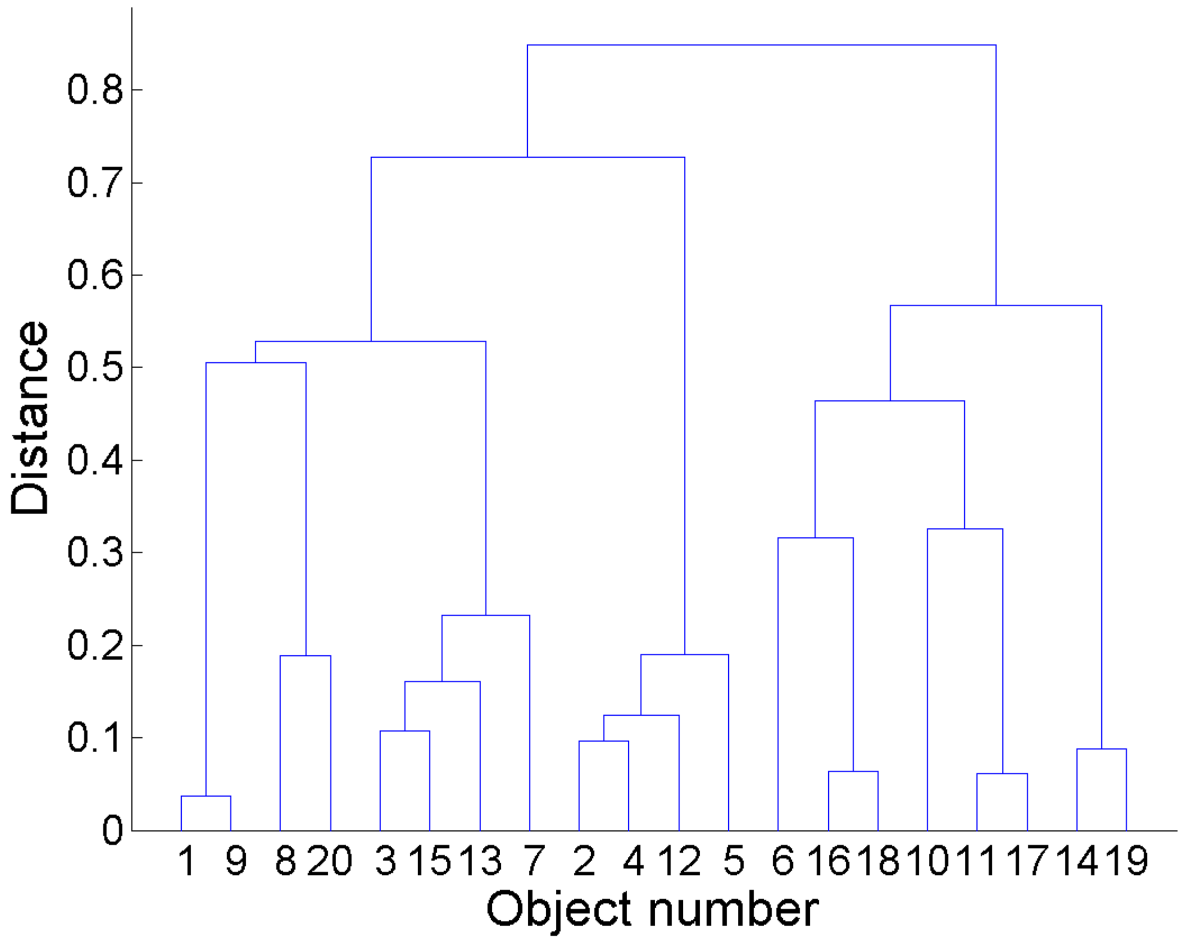

Dendrogram example

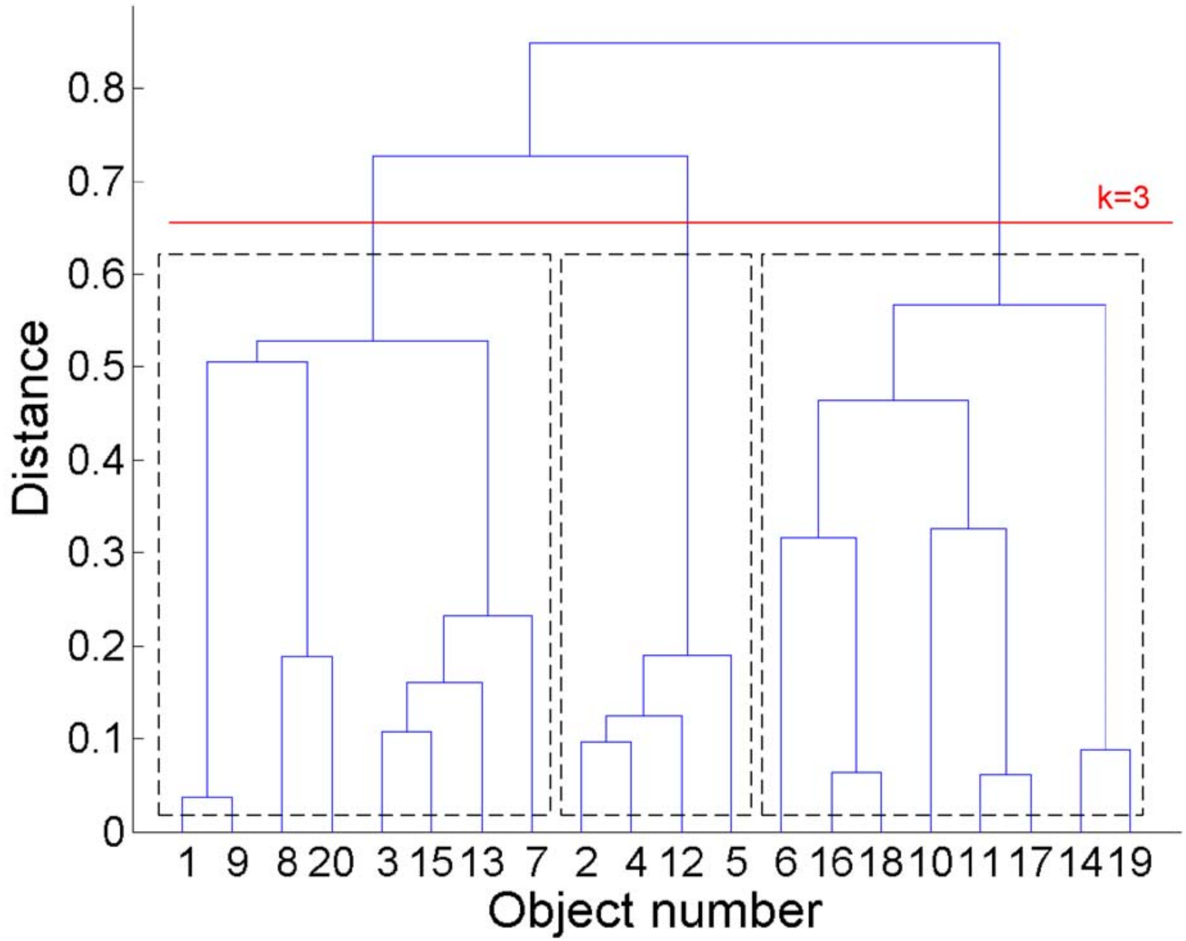

Wavelet transforms

- A time series signal is decomposed into different time scales using wavelet transforms
- Each time scale expresses the original signal at different frequencies
- Coarser time scales contain lower frequency approximations of the signal
- Finer time scales contain higher frequency approximations

Roadmap

- Introduction
- ChinaSat: network architecture, TCP, and network anomalies
- Mathematical tools for statistical analysis
- Analysis of billing records:
  - aggregated traffic
  - user behavior
- Analysis of tcpdump traces:
  - general characteristics
  - TCP options and OS fingerprinting
  - network anomalies
- Conclusions and future work

Billing records

- Records were collected during the continuous period from 23:00 on Oct. 31, 2002 to 11:00 on Jan. 10, 2003
- Each file contains the hourly traffic summary for each user
- Fields of interests:
  - SiteID (user identification)
  - Start (record start time)
  - CTxByt (number of bytes downloaded by a user)
  - CRxByt (number of bytes uploaded by a user)
  - CTxPkt (number of packets downloaded by a user)
  - CRxPkt (number of packets uploaded by a user)

Download: from NOC to user through satellite
Upload: from user to NOC through dial-up

Billing records format

```
RecLen RecTyp SiteID     Start          Stop           Cmin
    Bill CTxByt     CRxByt     CTxPkt     CRxPkt


00100  001    0003809504 20030106130005 20030106140005 060
    2    0000000414 0000017240 0000000007 0000000227
00100  001    0004477001 20030106130005 20030106140005 060
    2    0000000396 0000006084 0000000006 0000000117
00100  001    000456EB01 20030106130005 20030106140005 060
    2    0015844812 0002903556 0000027471 0000034200
00100  001    00045C0002 20030106130005 20030106140005 060
    2    0003061014 0000397334 0000003789 0000004521
00100  001    000455B103 20030106130005 20030106140005 008
    2    0000000120 0000001021 0000000002 0000000009
```

Billing records: characteristics

- 186 unique SiteIDs (users)
- Daily and weekly cycles:
  - lower traffic volume on weekends
  - daily cycle starts at 7 AM, rises to three daily maxima at 11 AM, 3 PM, and 7 PM, then decreases monotonically until 7 AM
- Highest daily traffic recorded on Dec. 24, 2002
- Outage occurred on Jan. 3, 2003

Aggregated hourly traffic

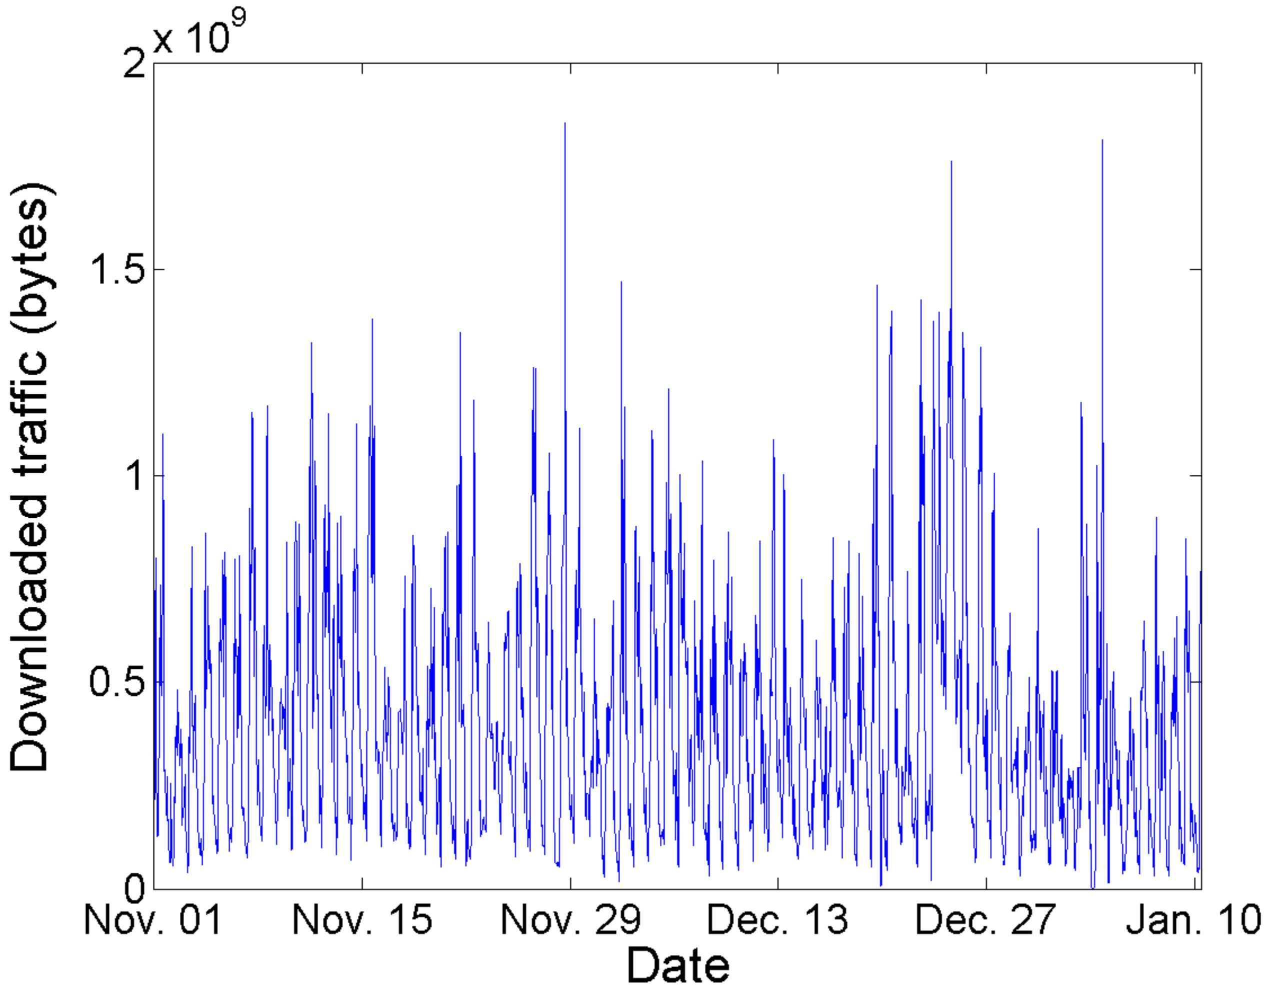

Aggregated daily traffic

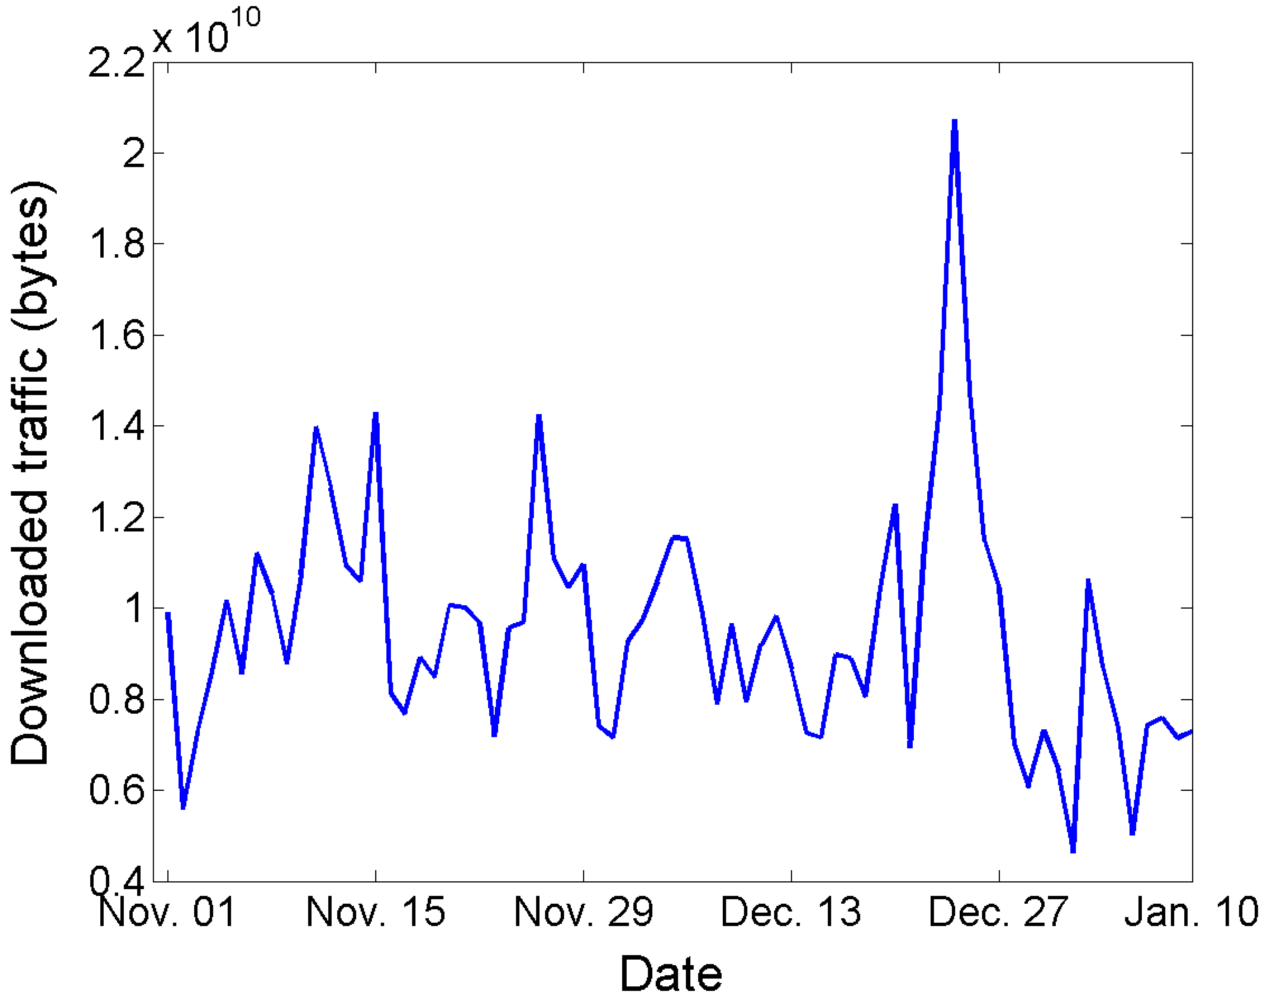

Daily diurnal traffic: average traffic (packets)

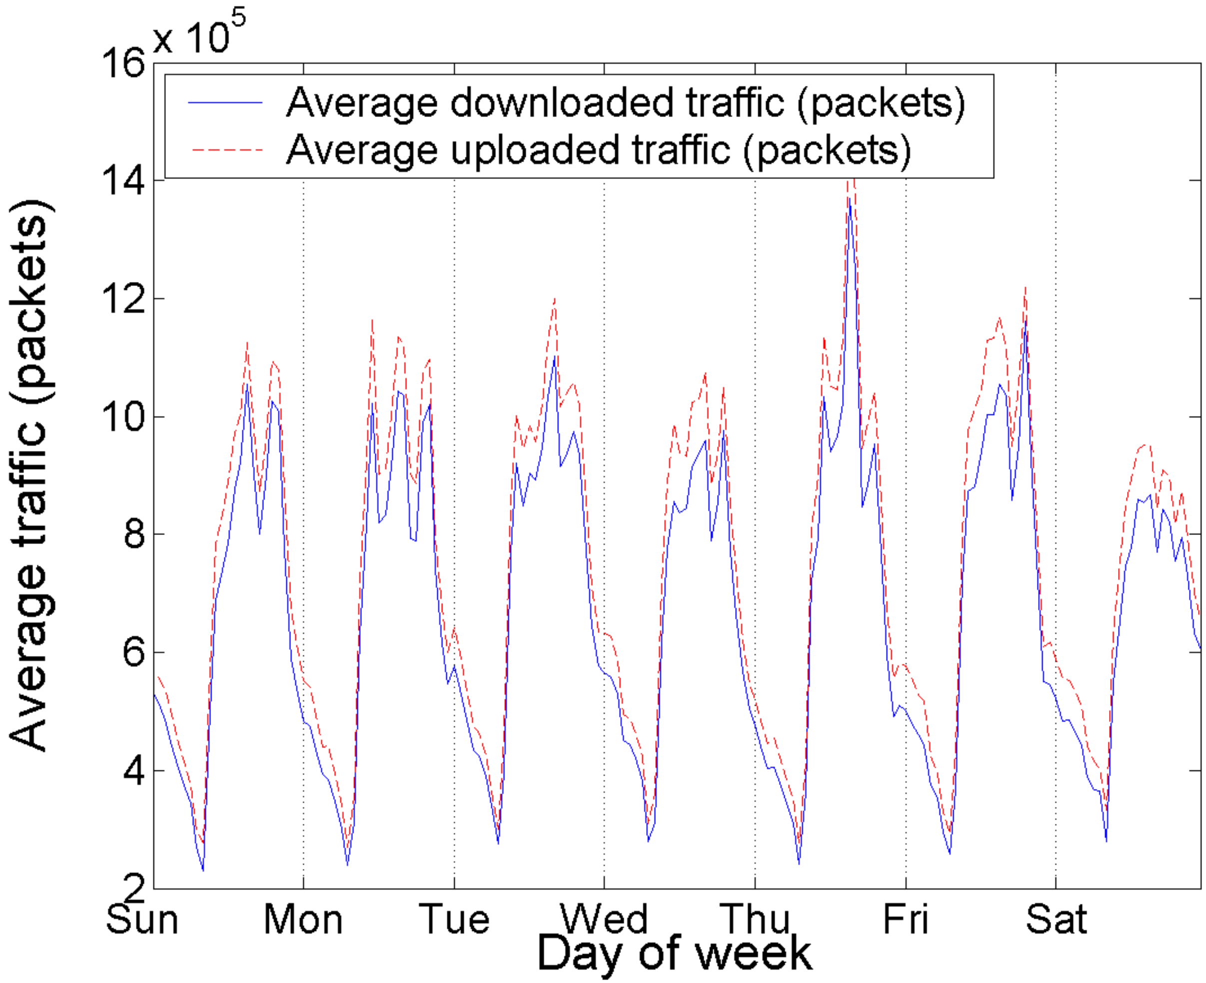

Weekly traffic: average traffic (bytes)

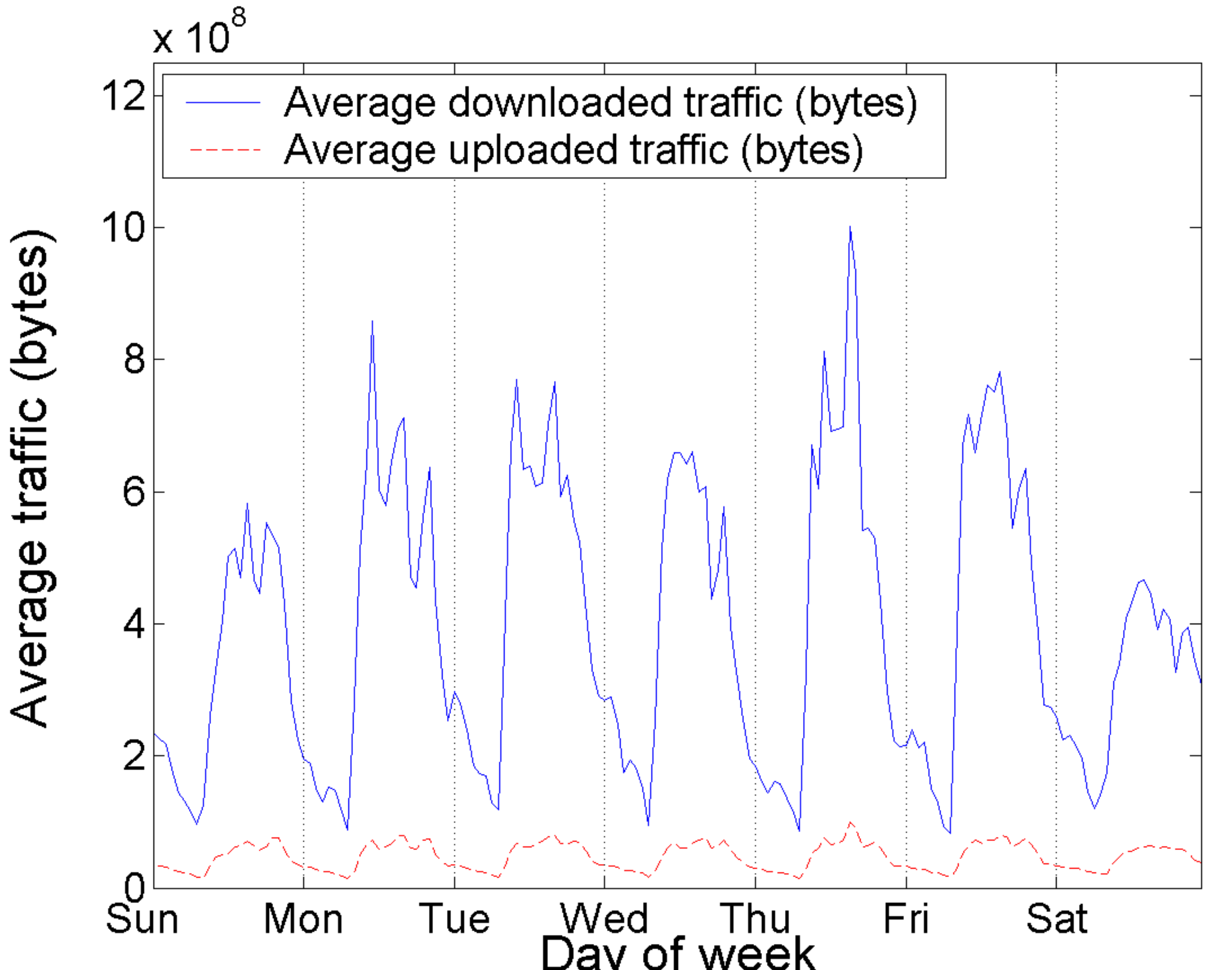

Roadmap

- Introduction
- ChinaSat: network architecture, TCP, and network anomalies
- Mathematical tools for statistical analysis
- Analysis of billing records:
  - aggregated traffic
  - user behavior
- Analysis of tcpdump traces:
  - general characteristics
  - TCP options and OS fingerprinting
  - network anomalies
- Conclusions and future work

Ranking of user traffic

- User traffic are ranked according to the traffic volume
- The top user downloaded 78.8 GB, uploaded 11.9 GB, and downloaded/uploaded ~205 million packets
- Most users download/uploaded little traffic
- Cumulative distribution functions (CDFs) are constructed from the ranks:
  - top user accounts for 11% of downloaded bytes
  - top 25 users contributed 93.3% of total downloaded bytes
  - top 37 users contributed 99% of total traffic (packets and bytes)

Cumulative distribution functions

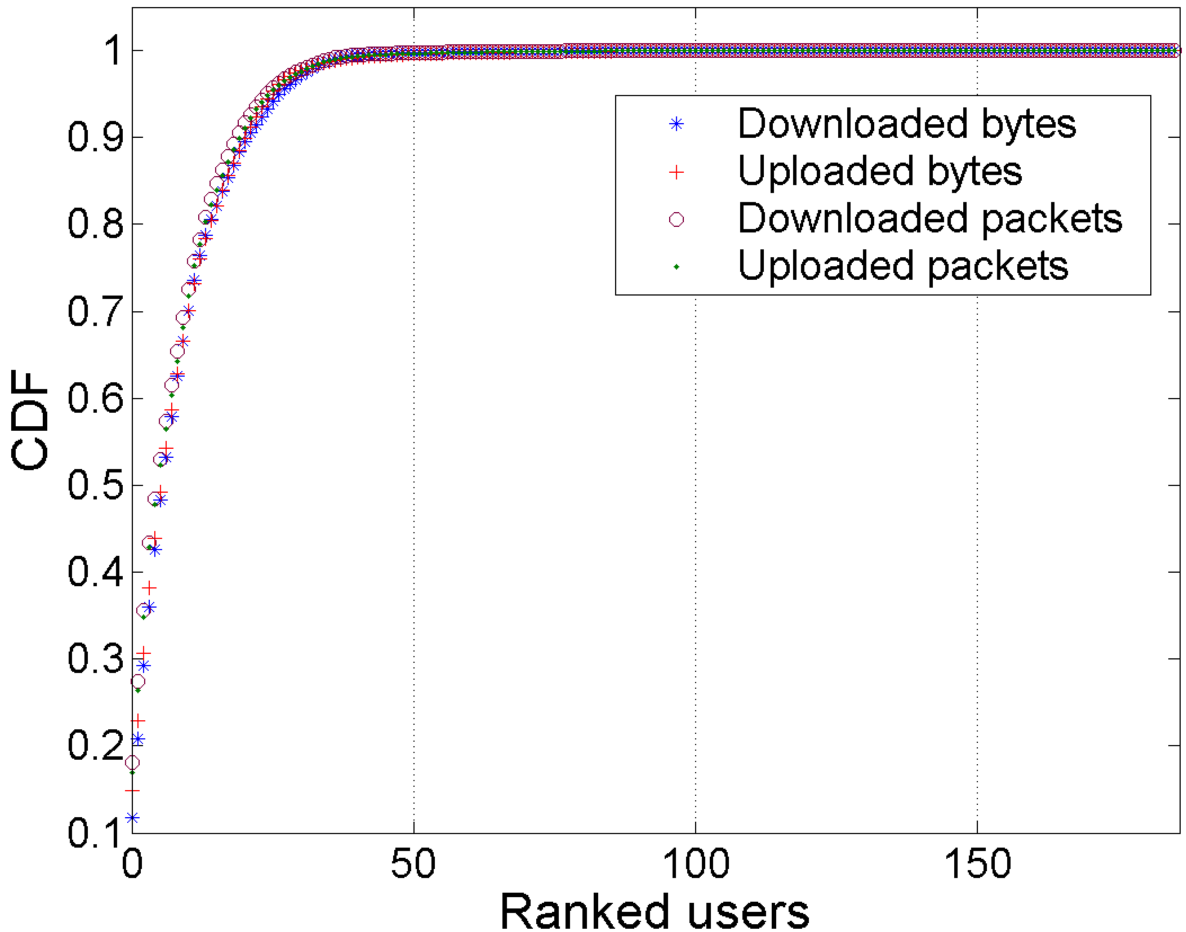

Classification of users with cluster analysis

- k-means clustering:
  - based on the volume of average traffic (downloaded packets, uploaded packets, downloaded bytes, and uploaded bytes)
  - multi-variable
- Hierarchical clustering:
  - clustering of users is based on user activity
  - results are refined by clustering with the three most common traffic patterns

Classification of users with k-means clustering (single variable)

- Single variable k-means clustering is employed for average downloaded and uploaded packets and bytes per hour
- Algorithm is repeated for k=2–10
- Algorithm is repeated 15 times for each k to avoid convergence to local minima
- Maximum number of iterations is set to 500
- Silhouette coefficients (SC plots, average SC) are used to determine the natural number of clusters

Single variable k-means clustering results

- Natural number of clusters occurs at k=3 for downloaded and uploaded bytes
- Most users belong to the group with little traffic
- For k=3:
  - 159 users in group 1 (average 0.0–16.8 MB downloaded per hour)
  - 24 users in group 2 (average 16.8–70.6 MB downloaded per hour)
  - 3 users in group 3 (average 70.6–110.7 MB downloaded per hour)

Classification based on user activity

- Pattern matching of signals with different mean, amplitude, and variance is difficult
- For each hour, user activity is classified as BUSY (1) or IDLE (0):
  - BUSY if a user has either downloaded or uploaded traffic
  - IDLE if a user has neither downloaded nor uploaded traffic

Classification of user activity

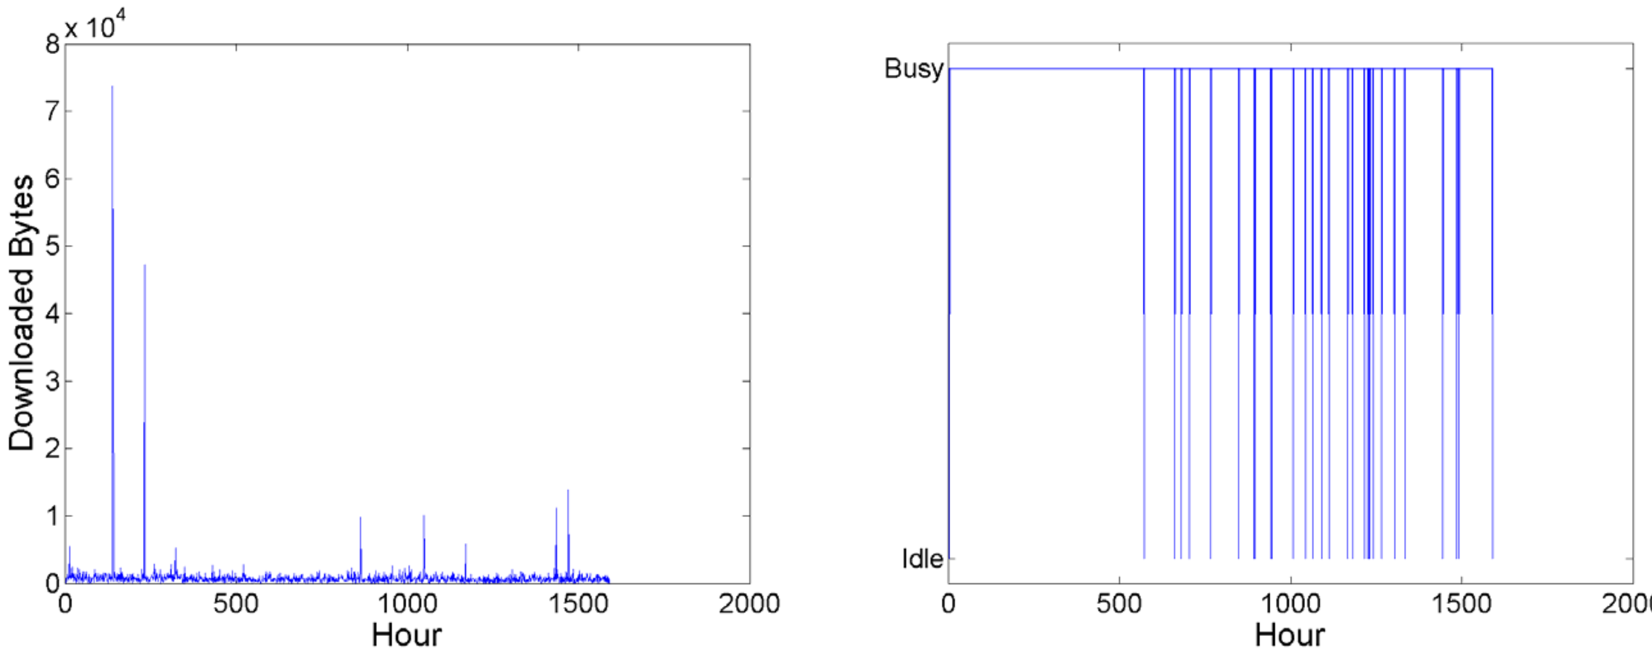

Classification of users with hierarchical clustering

- A similarity matrix is created by comparing the user activity
- Users are compared based on “active period”, which lasted at least 3 weeks (504 hours)
- Four distance measures: minimum, maximum, mean, and average
- Cophenetic correlation coefficients (CPCC) are used to evaluate the quality of distance measures

Comparing user activities

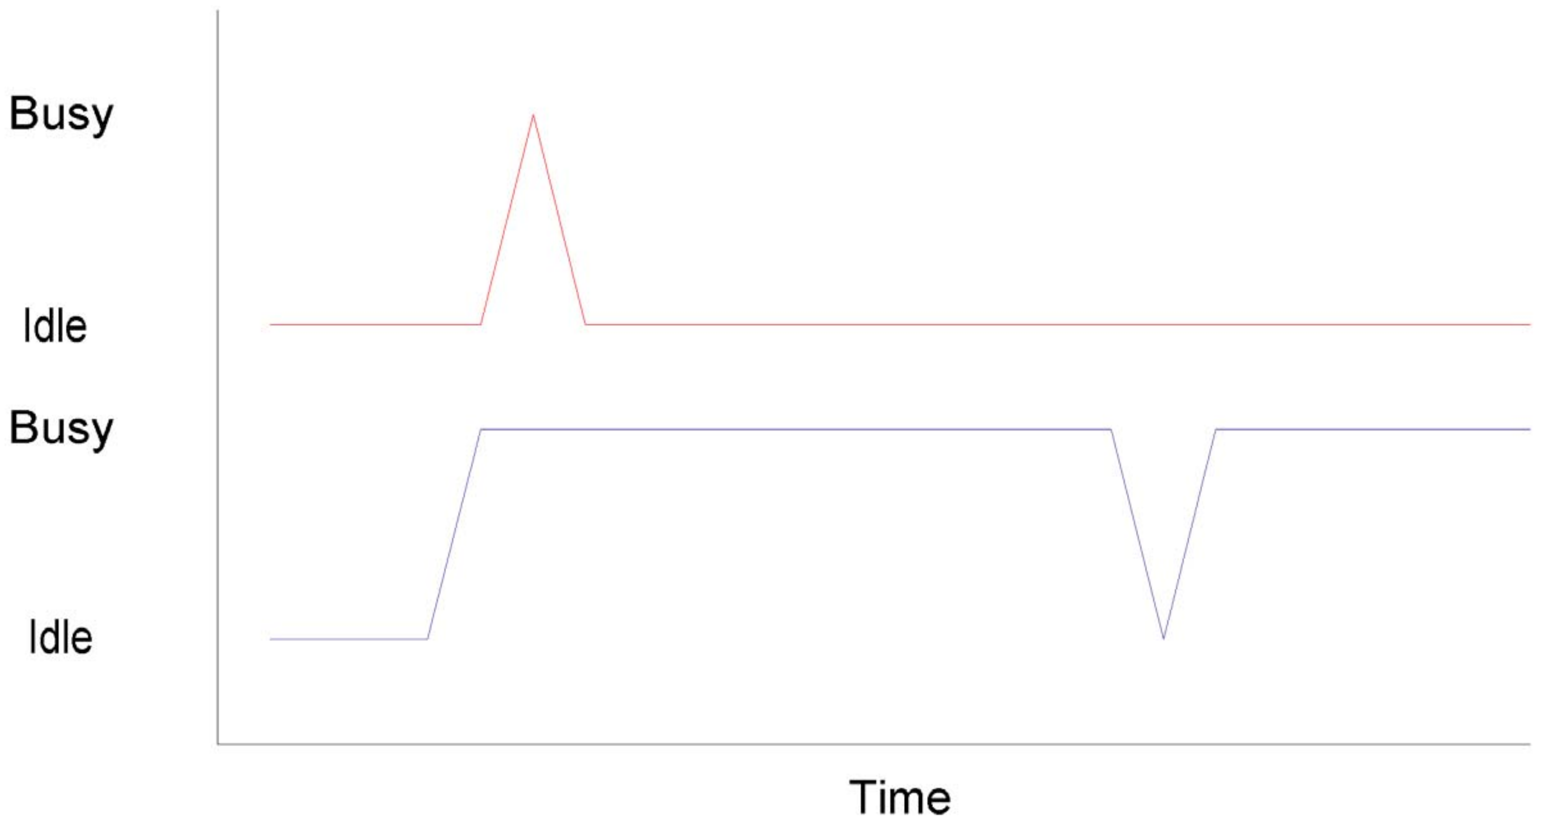

Distance measures used for hierarchical clustering

Distance measure	CPCC
Minimum distance	0.6890
Maximum distance	0.7761
Mean distance	0.9277
Average distance*	0.9363

* Results for the average distance measure is rejected because the result violates the hierarchical property of trees

CPCC: Cophenetic correlation coefficient

Dendrogram (average distance)

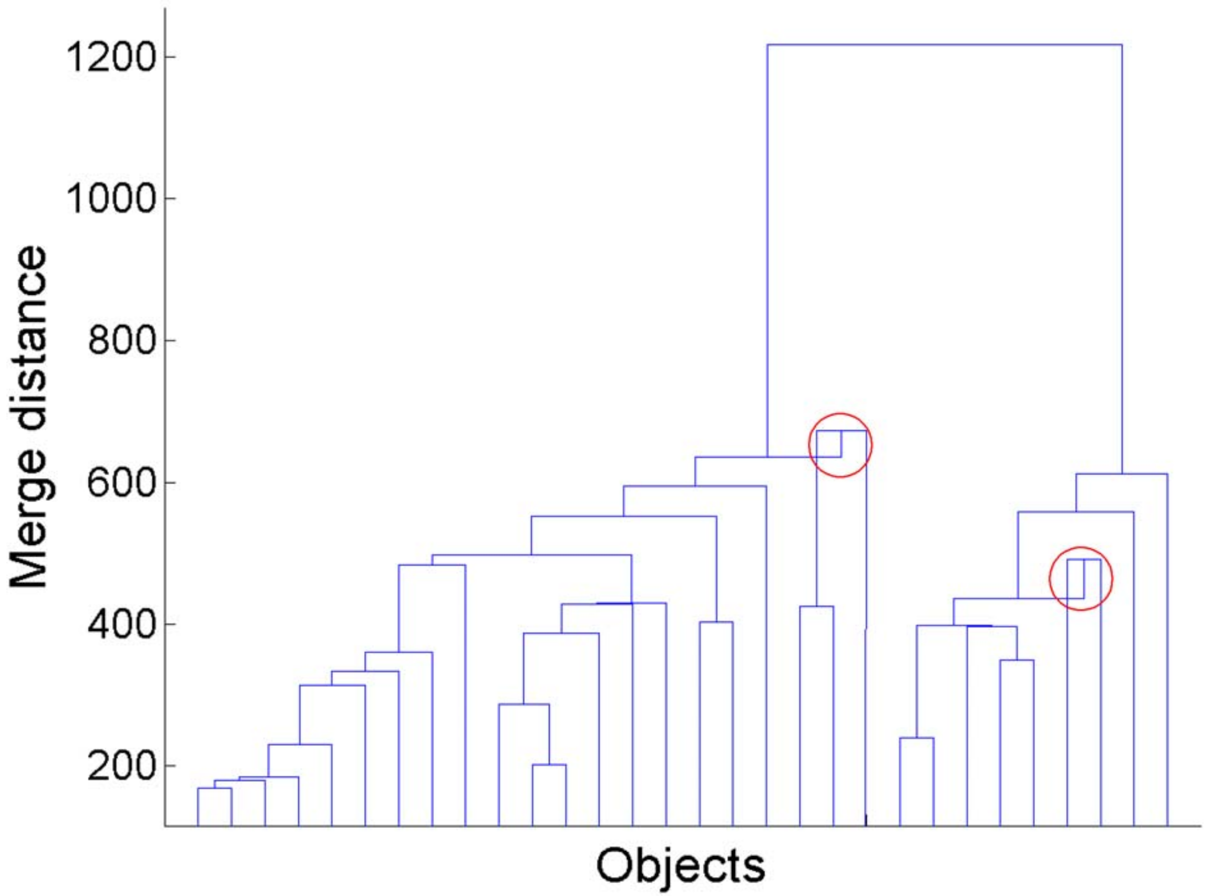

Hierarchical clustering: determining number of groups

- Inconsistency coefficients are used to determine the number of clusters:
  - maximum inconsistency coefficient is 1.1547
  - 90% cutoff value (1.10) generates 68 clusters
  - coefficient cutoff of 0.9 results in 75 clusters
  - large number of clusters is caused by users whose activity do not overlap
- Selecting 3 clusters produces no detectable patterns

Hierarchical clustering results

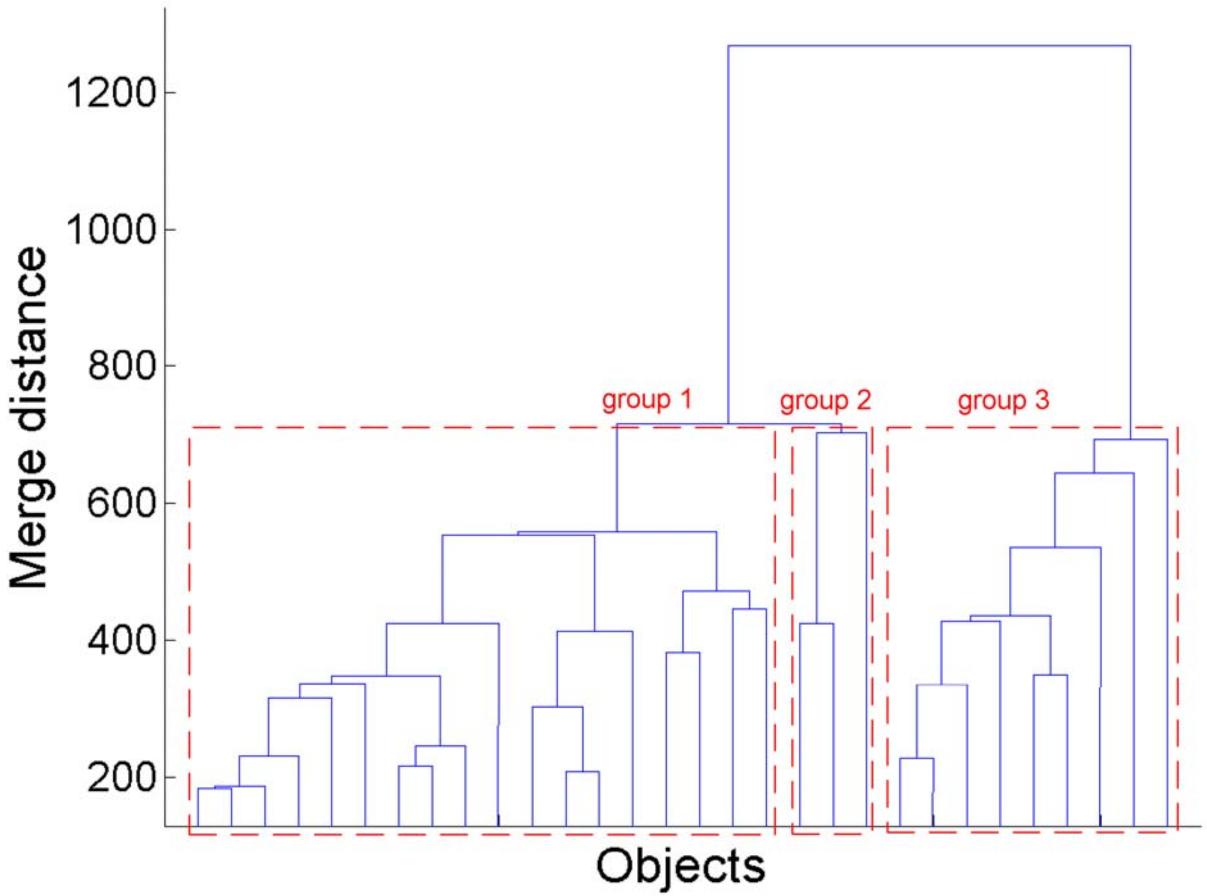

Hierarchical clustering results

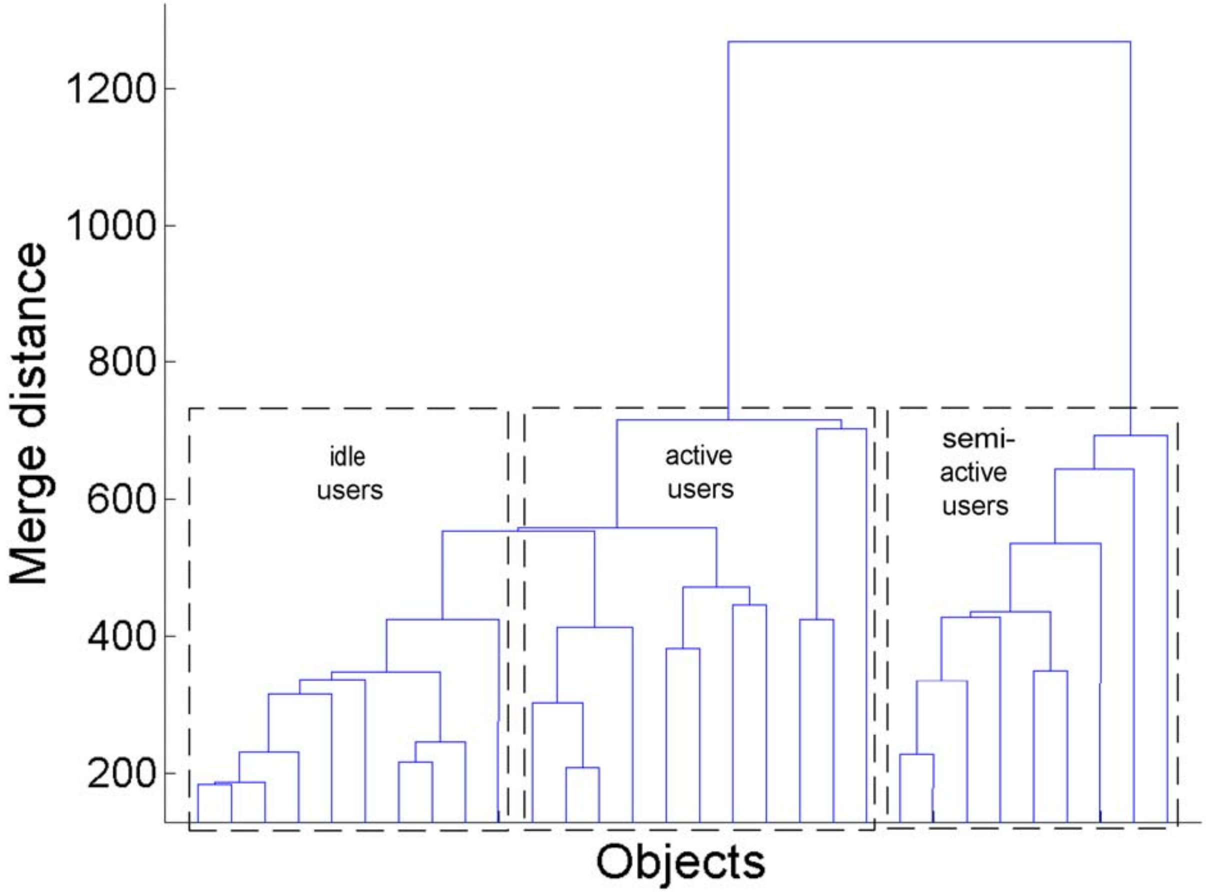

Refinement: three most common traffic patterns

- Inactive users:
  - rarely download/upload traffic
  - represented by zero traffic
- Active users:
  - download/upload traffic for more than 18 hours a day
  - represented by traffic for 24 hours each day
- Semi-active users:
  - download/upload traffic for 8–12 hours a day
  - represented by a cycle of 10 hours BUSY / 14 hours IDLE cycle for each day

Clustering using three most common traffic patterns

- Only the “active period” is compared because some users are not active for the whole duration of the records
- A similarity value of one is added for each hour that the user traffic equals the most common traffic patterns
- The sum of the similarity value is the similarity score
- For the Semi-active traffic pattern, we try to match the cycle phase of the user traffic with the model
- A user is grouped with the model that it has the highest similarity score

Refinement: clustering results

Traffic pattern	Number of users
Inactive	162
Active	16
Semi-active	8
Total number of users	186

k-means and hierarchical clustering combined

- Clustering of users based on average traffic and user activity
- Natural number of clusters using k-means clustering is k=3
- We chose the 3 most common traffic patterns because too many clusters were generated by hierarchical clustering
- The combination of the 3 most common traffic patterns and 3 k-means clusters results in a maximum of 9 groups:
  - one of the groups (high traffic volume and active) in the combined result has no object
  - only 8 groups are present

Clusters: combined results

- Users with low traffic volume:
  - inactive users (150 users)
  - active users (7 users)
  - semi-active users (2 users)
- Users with medium traffic volume:
  - inactive users (11 users)
  - active users (9 users)
  - semi-active users (4 users)
- Users with high traffic volume:
  - inactive users (1 user)
  - semi-active (2 users)

Roadmap

- Introduction
- ChinaSat: network architecture, TCP, and network anomalies
- Mathematical tools for statistical analysis
- Analysis of billing records:
  - aggregated traffic
  - user behavior
- Analysis of tcpdump traces:
  - general characteristics
  - TCP options and OS fingerprinting
  - network anomalies
- Conclusions and future work

tcpdump trace

- Trace were continuously collected from 11:30 on Dec. 14, 2002 to 11:00 on Jan. 10, 2003 at the NOC
- The first 68 bytes of each TCP/IP packet were captured
- ~63 GB of data contained in 127 files
- User IP address is not constant due to the use of the private IP address range and dynamic IP
- Majority of traffic is TCP:
  - 94% of total bytes and 84% of total packets
  - HTTP (port 80) accounts for 90% of TCP connections and 76% of TCP bytes
  - FTP (port 21) accounts for 0.2% of TCP connections and 11% of TCP bytes

pcap file and header format

pcap header section		
pcap data	*pcap* data	
pcap data (cont'd)	*pcap* data	
pcap data (cont'd)	*pcap* data	....

0	16	32
Magic number*		
pcap major version*	*pcap* minor version*	
Local time offset*		
Timer accuracy*		
Snap length*		
Link type*		

tcpdump output example

12/15/2002 04:27:05.328455 192.168.1.83.63260 > 211.167.92.197.6732: . ack 489 win 8192

12/15/2002 04:27:05.331020 211.100.18.48.80 > 192.168.1.164.41842: S 2928120965:2928120965(0) ack 3324468 win 64240 <mss 1460,nop,nop,sackOK> (DF)

12/15/2002 04:27:05.331612 61.135.137.66.9013 > 192.168.1.164.41806: P 3091059901:3091060177(276) ack 11834706 win 5840 (DF)

12/15/2002 04:27:05.343507 192.168.1.164.41806 > 61.135.137.66.9013: . ack 276 win 8192

12/15/2002 04:27:05.343748 192.168.1.242.45045 > 210.51.17.96.9065: P 25309490:25309522(32) ack 1436759200 win 8192 (DF)

12/15/2002 04:27:05.359048 192.168.1.242.44991 > 211.167.92.226.6732: P 17:25(8) ack 16 win 8192 (DF)

12/15/2002 04:27:05.359218 192.168.1.83.64228 > 61.242.153.168.11745: udp 92

12/15/2002 04:27:05.359383 192.168.1.164.9668 > 211.150.186.218.4000: udp 60

12/15/2002 04:27:05.359537 192.168.1.83.64228 > 61.242.153.168.11745: udp 92

12/15/2002 04:27:05.359693 192.168.1.83.64228 > 61.242.153.168.11745: udp 92

12/15/2002 04:27:05.359694 61.152.252.11.55901 > 192.168.1.242.45311: P 48:56(8) ack 1 win 62851 (DF)

12/15/2002 04:27:05.362315 210.51.17.96.9065 > 192.168.1.242.45045: . ack 32 win 32120 (DF)

12/15/2002 04:27:05.366415 61.135.137.26.9013 > 192.168.1.242.45533: P 112:138(26) ack 1 win 6432 (DF)

Roadmap

- Introduction
- ChinaSat: network architecture, TCP, and network anomalies
- Mathematical tools for statistical analysis
- Analysis of billing records:
  - aggregated traffic
  - user behavior
- Analysis of tcpdump traces:
  - general characteristics
  - TCP options and OS fingerprinting
  - network anomalies
- Conclusions and future work

tcpdump trace: TCP options

- Selective acknowledgement (SACK) option: supported by > 60% of connections
- Sliding windows scale option: supported by < 5% of connections
- No instances of path MTU discovery
- Most connections use initial cwnd size: 4 segments or greater
- Observations agree with the TCP implementation in Microsoft Windows

MTU: maximum transmission unit

Operating system (OS) fingerprinting

- Used for intrusion detection, vulnerability discovery, and network auditing
- Based on the principle that TCP/IP implementations are unique
- Identifies an OS using the TCP SYN packet:
  - TCP SYN packet size
  - default TCP options
  - the order of TCP options
  - default TCP window size
  - default IP time-to-live (TTL) value
  - IP “do not fragment” (DF) flag
  - IP type of service (ToS) setting

OS fingerprinting results

- Analyzed 9 hours of tcpdump trace on Dec. 14, 2002 using the open-source tool p0f v2
- Assumed constant IP addresses
- Detected 171 users:
  - 137 users did not initiate any connection and cannot be identified (no SYN packets)
  - 14 users employ Microsoft Windows
  - 2 users employ Linux
  - 1 user employs an unknown OS (identified as an MSS-modifying proxy)

Roadmap

- Introduction
- ChinaSat: network architecture, TCP, and network anomalies
- Mathematical tools for statistical analysis
- Analysis of billing records:
  - aggregated traffic
  - user behavior
- Analysis of tcpdump traces:
  - general characteristics
  - TCP options and OS fingerprinting
  - network anomalies
- Conclusions and future work

Network anomalies

- Ethereal/Wireshark, tcptrace, and pcapread
- Four types of network anomalies were detected:
  - invalid TCP flag combinations
  - large number of TCP resets
  - UDP and TCP port scans
  - traffic volume anomalies

Invalid TCP flag combinations

- TCP SYN flag: signal to establish connections
- TCP FIN flag: signal to terminate connections regularly
- TCP RST flag: signal to terminate connections when error occurs
- TCP PSH flag: signal to transmit all outstanding packets in the buffer without delay
- Invalid combinations are SYN+FIN, SYN+RST, RST+FIN, RST+PSH, and RST+FIN+PSH
- A single invalid packet may cause a vulnerable TCP/IP implementation to exhibit unexpected behavior

Analysis of TCP flags

TCP flag	Packet count	% of Total
SYN only	19,050,849	48.500
RST only	7,440,418	18.900
FIN only	12,679,619	32.300
*SYN+FIN	408	0.001
*RST+FIN (no PSH)	85,571	0.200
*RST+PSH (no FIN)	18,111	0.050
*RST+FIN+PSH	8,329	0.020
*Total number of packets with invalid TCP flag combinations	112,419	0.300
Total packet count	39,283,305	100.000

Large number of TCP resets

- Connections are terminated by either TCP FIN or TCP RST:
  - 12,679,619 connections were terminated by FIN (63%)
  - 7,440,418 connections were terminated by RST (37%)
- Large number of TCP RST indicates that connections are terminated in error conditions
- TCP RST is employed by Microsoft Internet Explorer to terminate connections instead of TCP FIN

M. Arlitt and C. Williamson, "An analysis of TCP reset behaviour on the Internet," *ACM SIGCOMM Comput. Commun. Rev.*, vol. 35, no. 1, pp. 37–44, Jan. 2005.

UDP and TCP port scans

- UDP port scans are found on UDP port 137 (NETBEUI)
- TCP ports scans are found on these TCP ports:
  - 80 Hypertext transfer protocol (HTTP)
  - 139 NETBIOS extended user interface (NETBEUI)
  - 443 HTTP over secure socket layer (HTTPS)
  - 1433 Microsoft structured query language (MS SQL)
  - 27374 Subseven trojan
- No HTTP(S) servers were active in the ChinaSat network
- MS SQL vulnerability was discovered in Oct. 2002, which may be the cause of scans on TCP port 1433
- The Subseven trojan is a backdoor program used with malicious intents

UDP port scans originating from the ChinaSat network

192.168.2.30:137 - 195.x.x.98:1025
192.168.2.30:137 - 202.x.x.153:1027
192.168.2.30:137 - 210.x.x.23:1035
192.168.2.30:137 - 195.x.x.42:1026
192.168.2.30:137 - 202.y.y.226:1026
192.168.2.30:137 - 218.x.x.238:1025
192.168.2.30:137 - 202.y.y.226:1025
192.168.2.30:137 - 202.y.y.226:1027
192.168.2.30:137 - 202.y.y.226:1028
192.168.2.30:137 - 202.y.y.226:1029
192.168.2.30:137 - 202.y.y.242:1026
192.168.2.30:137 - 61.x.x.5:1028
192.168.2.30:137 - 219.x.x.226:1025
192.168.2.30:137 - 213.x.x.189:1028
192.168.2.30:137 - 61.x.x.193:1025
192.168.2.30:137 - 202.y.y.207:1028
192.168.2.30:137 - 202.y.y.207:1025
192.168.2.30:137 - 202.y.y.207:1026
192.168.2.30:137 - 202.y.y.207:1027
192.168.2.30:137 - 64.x.x.148:1027

- Client (192.168.2.30) source port (137) scans external network addresses at destination ports (1025-1040):
  - > 100 are recorded within a three-hour period
  - targets IP addresses are variable
  - multiple ports are scanned for a single IP
  - may correspond to Bugbear, OpaSoft, or other worms

UDP port scans direct to the ChinaSat network

210.x.x.23:1035 - 192.168.1.121:137
210.x.x.23:1035 - 192.168.1.63:137
210.x.x.23:1035 - 192.168.2.11:137
210.x.x.23:1035 - 192.168.1.250:137
210.x.x.23:1035 - 192.168.1.25:137
210.x.x.23:1035 - 192.168.2.79:137
210.x.x.23:1035 - 192.168.1.52:137
210.x.x.23:1035 - 192.168.6.191:137
210.x.x.23:1035 - 192.168.1.241:137
210.x.x.23:1035 - 192.168.2.91:137
210.x.x.23:1035 - 192.168.1.5:137
210.x.x.23:1035 - 192.168.1.210:137
210.x.x.23:1035 - 192.168.6.127:137
210.x.x.23:1035 - 192.168.1.201:137
210.x.x.23:1035 - 192.168.6.179:137
210.x.x.23:1035 - 192.168.2.82:137
210.x.x.23:1035 - 192.168.1.239:137
210.x.x.23:1035 - 192.168.1.87:137
210.x.x.23:1035 - 192.168.1.90:137
210.x.x.23:1035 - 192.168.1.177:137
210.x.x.23:1035 - 192.168.1.39:137

- External address (210.x.x.23) scans for port (137) (NETBEUI) response within the ChinaSat network from source port (1035):
  - > 200 are recorded within a three-hour period
  - targets IP addresses are not sequential
  - may correspond to Bugbear, OpaSoft, or other worms

Detection of traffic volume anomalies using wavelets

- Traffic is decomposed into different frequencies using the wavelet transform
- Traffic volume anomalies are identified by the large variation in wavelet coefficient values
- The coarsest scale level where the anomalies is found indicates the time scale of an anomaly

Detection of traffic volume anomalies using wavelets

- tcpdump traces are binned in terms of packets or bytes (each second)
- Wavelet transform of 12 levels is employed to decompose the traffic
- The coarsest level approximately represents the hourly traffic
- Anomalies are:
  - detected with a moving window of size 20 and by calculating the mean and standard deviation (σ) of the wavelet coefficients in each window
  - identified when wavelet coefficients lie outside $\pm 3\sigma$ of the mean value

Wavelet approximation coefficients

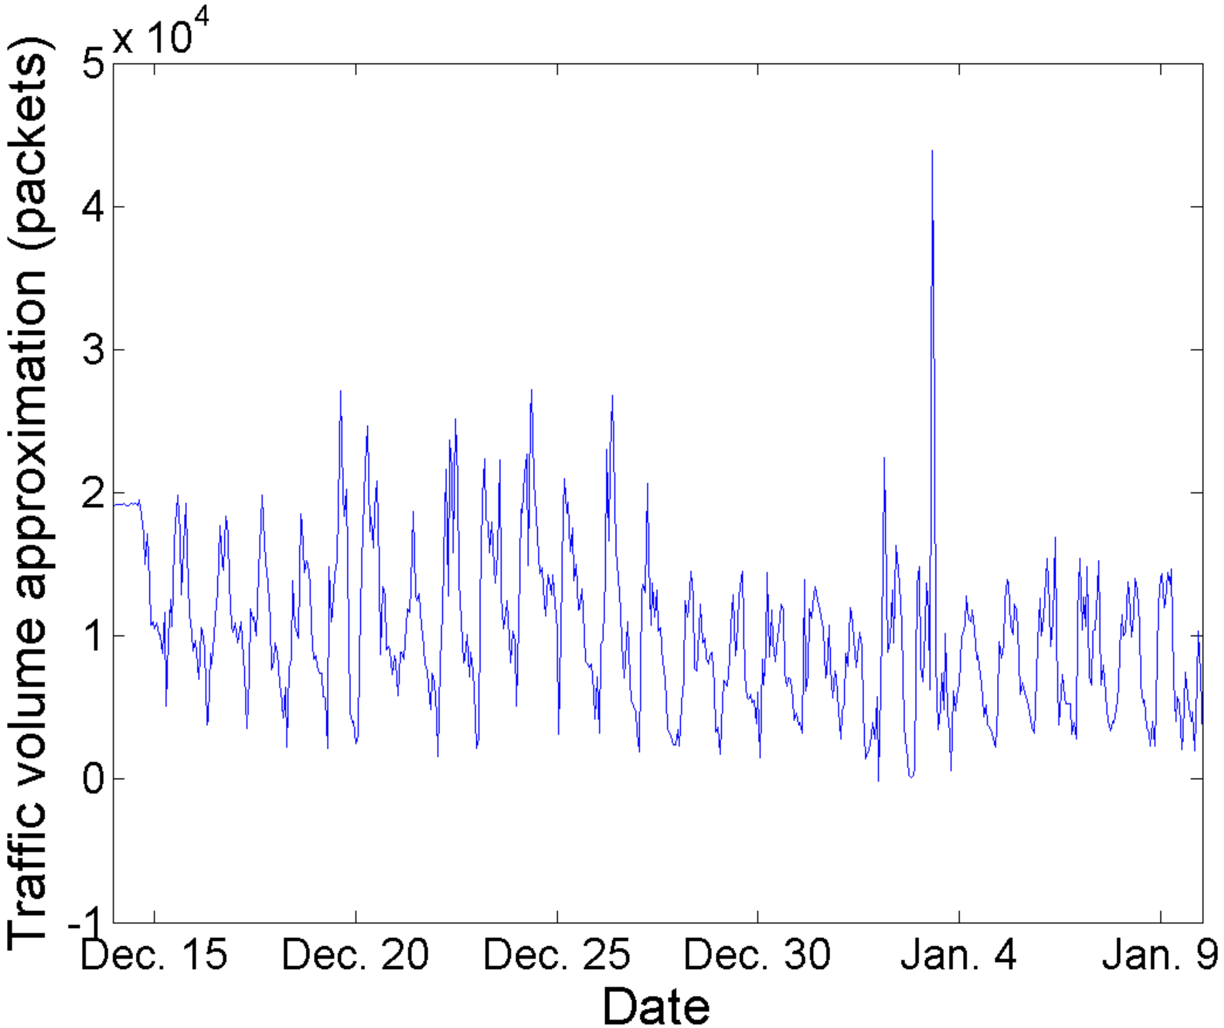

Wavelet detail coefficients: d_9

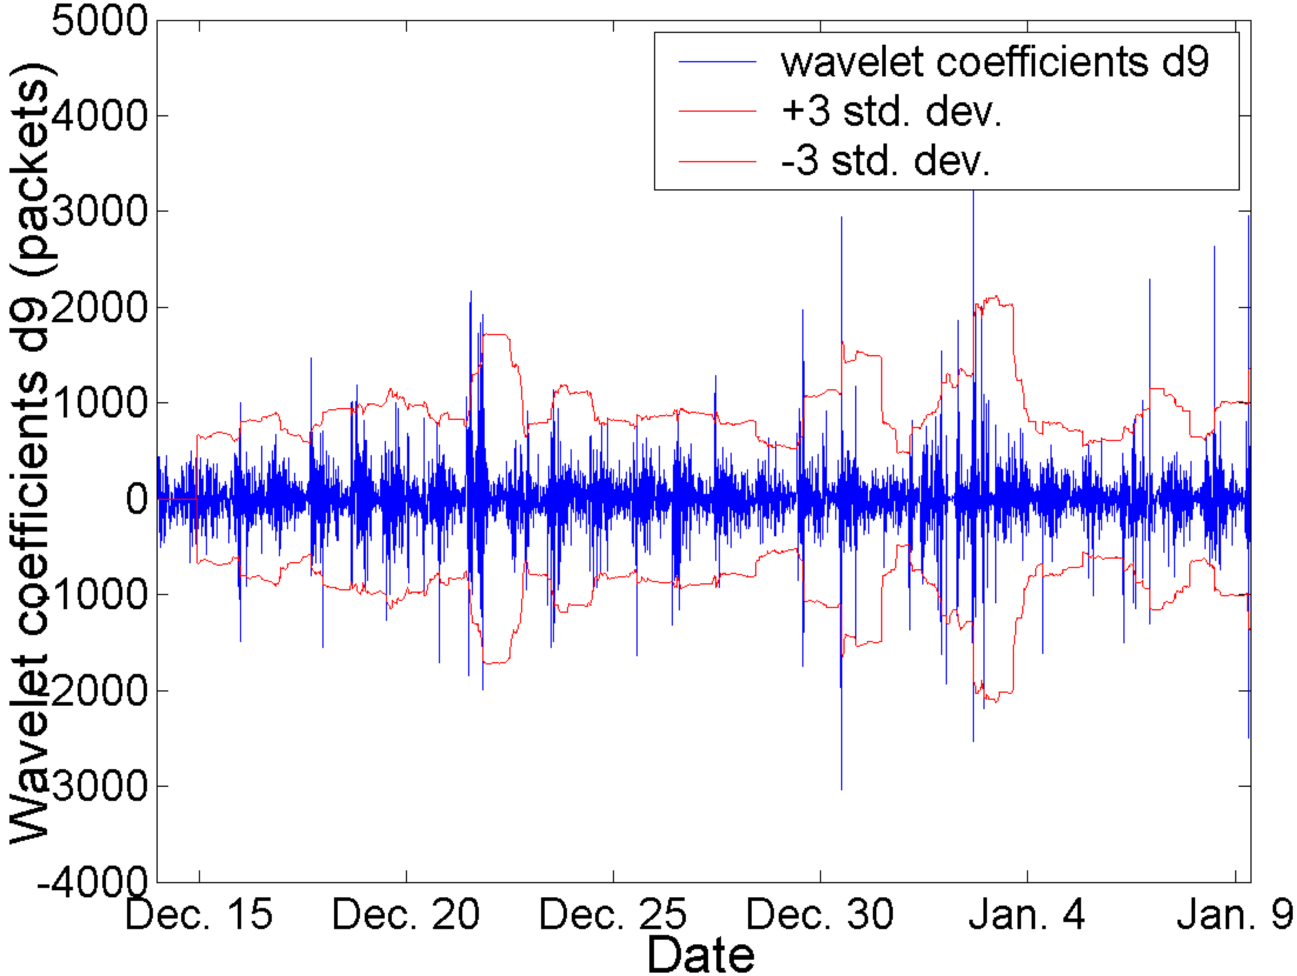

Wavelet detail coefficients: d_8

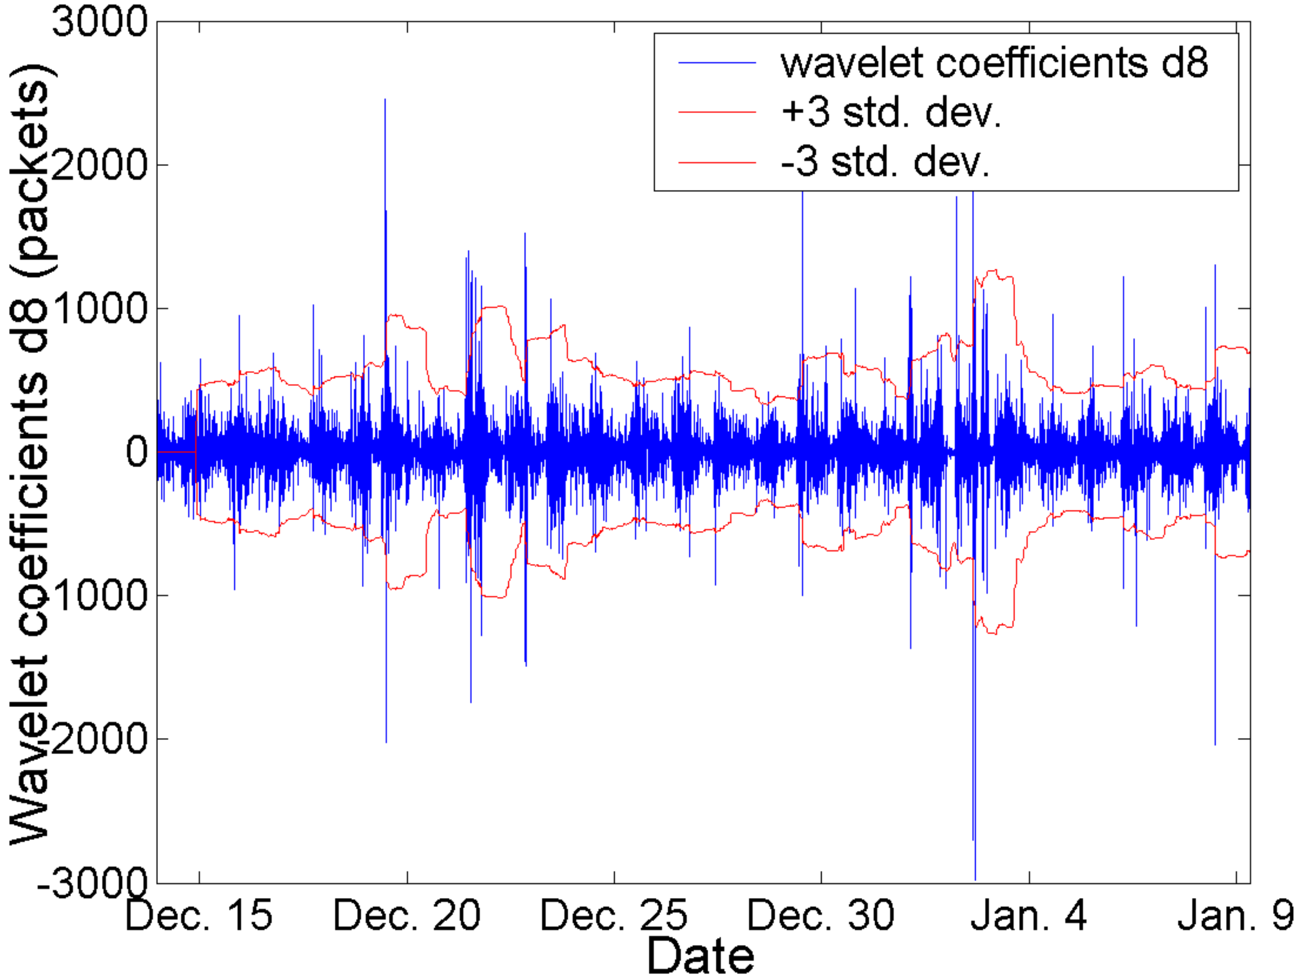

Roadmap

- Introduction
- ChinaSat: network architecture, TCP, and network anomalies
- Mathematical tools for statistical analysis
- Analysis of billing records:
  - aggregated traffic
  - user behavior
- Analysis of tcpdump traces:
  - general characteristics
  - TCP options and OS fingerprinting
  - network anomalies
- Conclusions and future work

Conclusions

- Analyzed billing records and tcpdump traces from a hybrid satellite-terrestrial network operated by ChinaSat
- Billing records:
  - minority of users contributed most of the traffic
  - k-means clustering of average user traffic indicates that there are three natural groups present (k=3)
  - ChinaSat users have three common types of activity:
    - inactive: little traffic throughout the record period
    - active: contribute traffic for > 18 hours a day
    - semi-active: BUSY for 8-12 hours then IDLE for 12-16 hours

Conclusions

- tcpdump trace:
  - TCP accounts for majority of traffic
  - TCP options most widely used to improve performance are SACK and increasing initial windows size
  - ChinaSat DirecPC hosts may be optimized by:
    - ensuring the SACK option is enabled on all hosts
    - enabling the sliding window scale option
  - network anomalies are found using open source tools and wavelet decomposition

Future work

- Use pattern recognition techniques to analyze traffic patterns
- Investigate the effects of illegitimate traffic on the performance of the ChinaSat network
- Analyze traffic data from two-way satellite networks
- Apply analysis techniques to other deployed commercial networks

References

- S. Lau and Lj. Trajkovic, "Analysis of traffic data from a hybrid satellite-terrestrial network," in *Proc. QShine 2007*, Vancouver, BC, Canada, Aug. 2007, to appear.
- Q. Shao and Lj. Trajkovic, "Measurement and analysis of traffic in a hybrid satellite-terrestrial network," in *Proc. SPECTS 2004*, San Jose, CA, July 2004, pp. 329–336.
- J. Han and M. Kamber, *Data Mining: concept and techniques*. San Diego, CA:Academic Press, 2001.
- W. Wu, H. Xiong, and S. Shekhar, *Clustering and Information Retrieval*. Norwell,MA: Kluwer Academic Publishers, 2004.
- Z. Chen, *Data Mining and Uncertainty Reasoning: and integrated approach*. New York, NY: John Wiley & Sons, 2001.
- T. Kanungo, D. M. Mount, N. Netanyahu, C. Piatko, R. Silverman, and A. Y. Wu, "An efficient k-means clustering algorithm: analysis and implementation," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 24, no. 7, pp. 881–892, July. 2002.
- P.-N. Tan, M. Steinbach, and V. Kumar, *Introduction to Data Mining*. Reading,MA: Addison-Wesley, 2006, pp. 487–568.
- L. Kaufman and P. J. Rousseeuw, *Finding Groups in Data: an introduction to cluster analysis*. New York, NY: John Wiley & Sons, 1990.
- M. Last, A. Kandel, and H. Bunke, Eds., *Data Mining in Time Series Databases*. Singapore: World Scientific Publishing Co. Pte. Ltd., 2004.
- W.-K. Ching and M. K.-P. Ng, Eds., *Advances in Data Mining and Modeling*. Singapore: World Scientific Publishing Co. Pte. Ltd., 2003.

References

- J. Postel, Ed., “Transmission Control Protocol,” RFC 793, Sept. 1981.
- J. Postel, “TCP and IP bake off,” RFC 1025, Sept. 1987.
- J. Mogul and S. Deering, “Path MTU discovery,” RFC 1191, Nov. 1990.
- V. Jacobson, R. Braden, and D. Borman, “TCP extensions for high performance,” RFC 1323, May 1992.
- M. Allman, S. Floyd, and C. Partridge, “Increasing TCP’s initial window,” RFC 2414, Sept. 1998.
- M. Mathis, J. Mahdavi, S. Floyd, and A. Romanow, “TCP selective acknowledgment options,” RFC 2018, Oct. 1996.
- M. Allman, D. Glover, and L. Sanchez, “Enhancing TCP over satellite channels using standard mechanisms,” RFC 2488, Jan. 1999.
- M. Allman, S. Dawkins, D. Glover, J. Griner, D. Tran, T. Henderson, J. Heidemann, J. Touch, H. Kruse, S. Ostermann, K. Scott, and J. Semke, “Ongoing TCP research related to satellites,” RFC 2760, Feb. 2000.
- J. Border, M. Kojo, J. Griner, G. Montenegro, and Z. Shelby, “Performance enhancing proxies intended to mitigate link-related degradations,” RFC 3135, June 2001.
- S. Floyd, “Inappropriate TCP resets considered harmful,” RFC 3360, Aug. 2002.

References

- D. E. Comer, *Internetworking with TCP/IP, Vol 1: Principles, Protocols, and Architecture*, 4th ed. Upper Saddle River, NJ: Prentice-Hall, 2000.
- W. R. Stevens, *TCP/IP Illustrated (vol. 1): The Protocols*. Reading, MA: Addison-Wesley, 1994.
- R. Beverly, "A Robust Classifier for Passive TCP/IP Fingerprinting," in *Proc. Passive and Active Meas. Workshop 2004*, Antibes Juan-les-Pins, France, Apr. 2004, pp. 158–167.
- C. Smith and P. Grundl, "Know your enemy: passive fingerprinting," The Honeynet Project, Mar. 2002. [Online]. Available: http://www.honeynet.org/papers/finger/.
- Passive OS fingerprinting tool ver. 2 (p0f v2). [Online]. Available: http://lcamtuf.coredump.cx/p0f.shtml/.
- B. Petersen, "Intrusion detection FAQ: What is p0f and what does it do?" The SysAdmin, Audit, Network, Security (SANS) Institute. [Online]. Available: http://www.sans.org/resources/idfaq/p0f.php.
- T. Miller, "Passive OS fingerprinting: details and techniques," The SysAdmin, Audit, Network, Security (SANS) Institute. [Online]. Available: http://www.sans.org/reading room/special.php/.

References

- P. Barford and D. Plonka, “Characteristics of network traffic flow anomalies,” in *Proc. ACM SIGCOMM Internet Meas. Workshop 2001*, Nov. 2001, pp. 69–73.
- P. Barford, J. Kline, D. Plonka, and A. Ron, “A signal analysis of network traffic anomalies,” in *Proc. ACM SIGCOMM Internet Meas. Workshop 2002*, Marseille, France, Nov. 2002, pp. 71–82.
- Y. Zhang, Z. Ge, A. Greenberg, and M. Roughan, “Network anomography,” in *Proc. ACM SIGCOMM Internet Meas. Conf. 2005*, Berkeley, CA, Oct. 2005, pp. 317–330.
- A. Soule, K. Salamatian, and N. Taft, “Combining filtering and statistical methods for anomaly detection,” in *Proc. ACM SIGCOMM Internet Meas. Conf. 2005*, Berkeley, CA, Oct. 2005, pp. 331–344.
- P. Huang, A. Feldmann, and W. Willinger, “A non-instrusive, wavelet-based approach to detecting network performance problems,” in *Proc. ACM SIGCOMM Internet Meas. Workshop 2001*, San Francisco, CA, Nov. 2001, pp. 213–227.
- A. Lakhina, M. Crovella, and C. Diot, “Characterization of network-wide anomalies in traffic flows,” in *Proc. ACM SIGCOMM Internet Meas. Conf. 2004*, Taormina, Italy, Oct. 2004, pp. 201–206.
- A. Lakhina, M. Crovella, and C. Diot, “Diagnosing network-wide traffic anomalies,” *ACM SIGCOMM Comput. Commun. Rev.*, vol. 34, no. 4, pp. 219–230, Oct. 2004.
- M. Arlitt and C. Williamson, “An analysis of TCP reset behaviour on the Internet,” *ACM SIGCOMM Comput. Commun. Rev.*, vol. 35, no. 1, pp. 37–44, Jan. 2005.