

Network Usage and Performance Analysis in Resource-Constrained Systems

Veljko Pejović

Veljko.Pejovic@fri.uni-lj.si

FRI, Ljubljana



University of Ljubljana
Faculty of Computer and
Information Science


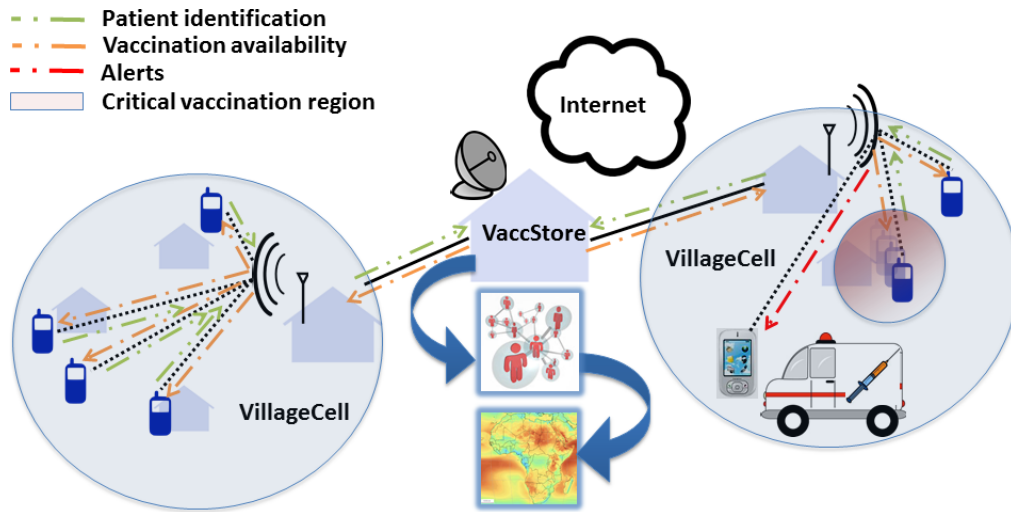
Ljubljana, Jan 2015.

Networked Systems

Systems composed of dynamic units that interact with each other over an information exchange network

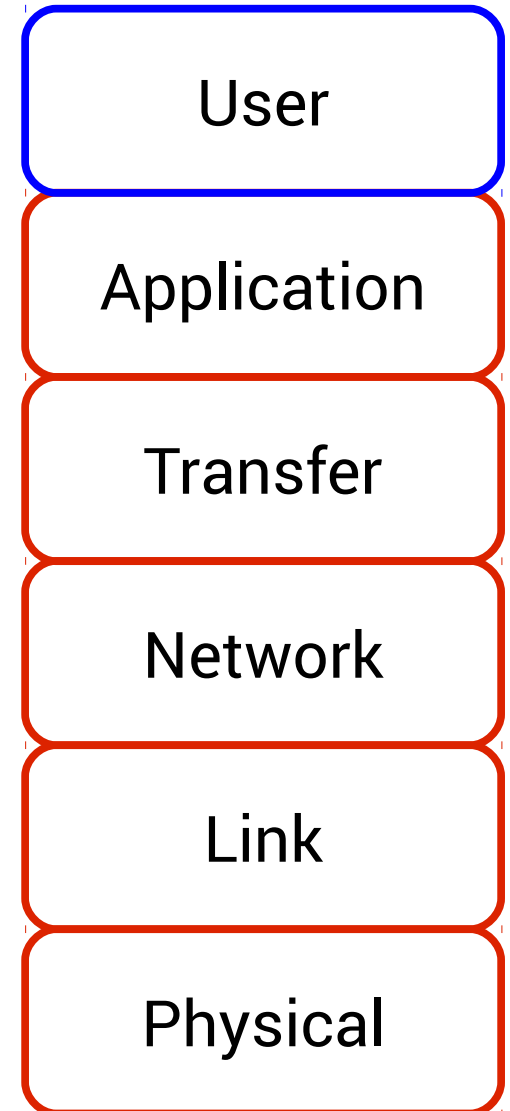


Networked Systems




Networked Systems

- A range of applications
 - Video/voice communication, photo exchange, flight control systems
- Varying in size
 - File sharing via a local server, planet-wide online social network
- Any underlying connection technology
 - Ethernet, WiFi, Bluetooth




Networked System Measurements

- Performance measurements



Networked System Measurements

- Performance measurements
- Usage measurements



Networked System Measurements

- Performance measurements



- Usage measurements



- Surrounding context measurements



Networked System Measurements

- Performance measurements
 - Usage measurements
-
- Surrounding context measurements



Performance Measurements

Measure	Tool
Reachability/routing	ping, traceroute
Wireless signal quality	iwconfig, iwlist
Latency/delay	ping, application level measurements
Throughput	iperf, iptraf
Error rate	netstat, iptraf
Jitter	iperf
...	network programming (sockets, packet injection, ...)



Usage Measurements

Measure	Tool
Traffic content	wireshark, tcpdump
Traffic summary	ipsumdump
Web usage	error logs, proxy logs, cache stats, calamaris
User experience	focus groups, online surveys, interviews
...	Scripting languages, content crawlers, be creative!



Why Measurements?

Troubleshooting

System
management

Innovation

Data-driven
prototype
evaluation



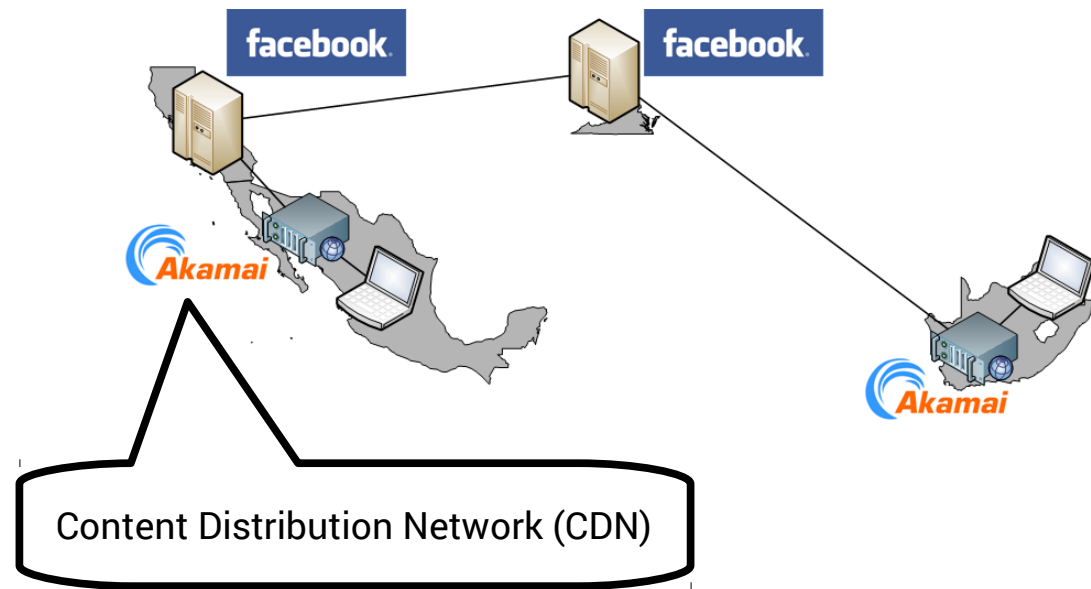
Case Study: Examining Alternative OSN Architectures



University of Ljubljana
Faculty of Computer and
Information Science

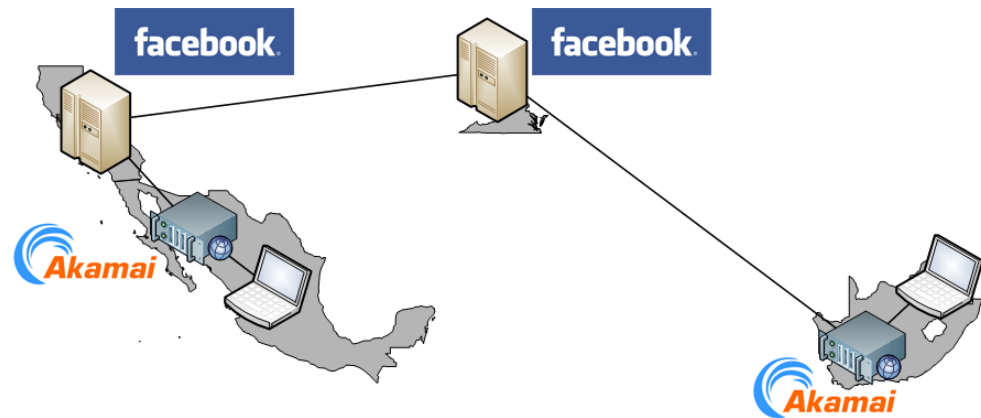
OSN Architecture – Facebook

- Centralised architecture
 - Poor usability outside North America and Europe
 - High cost of infrastructure in new markets
 - Scalability



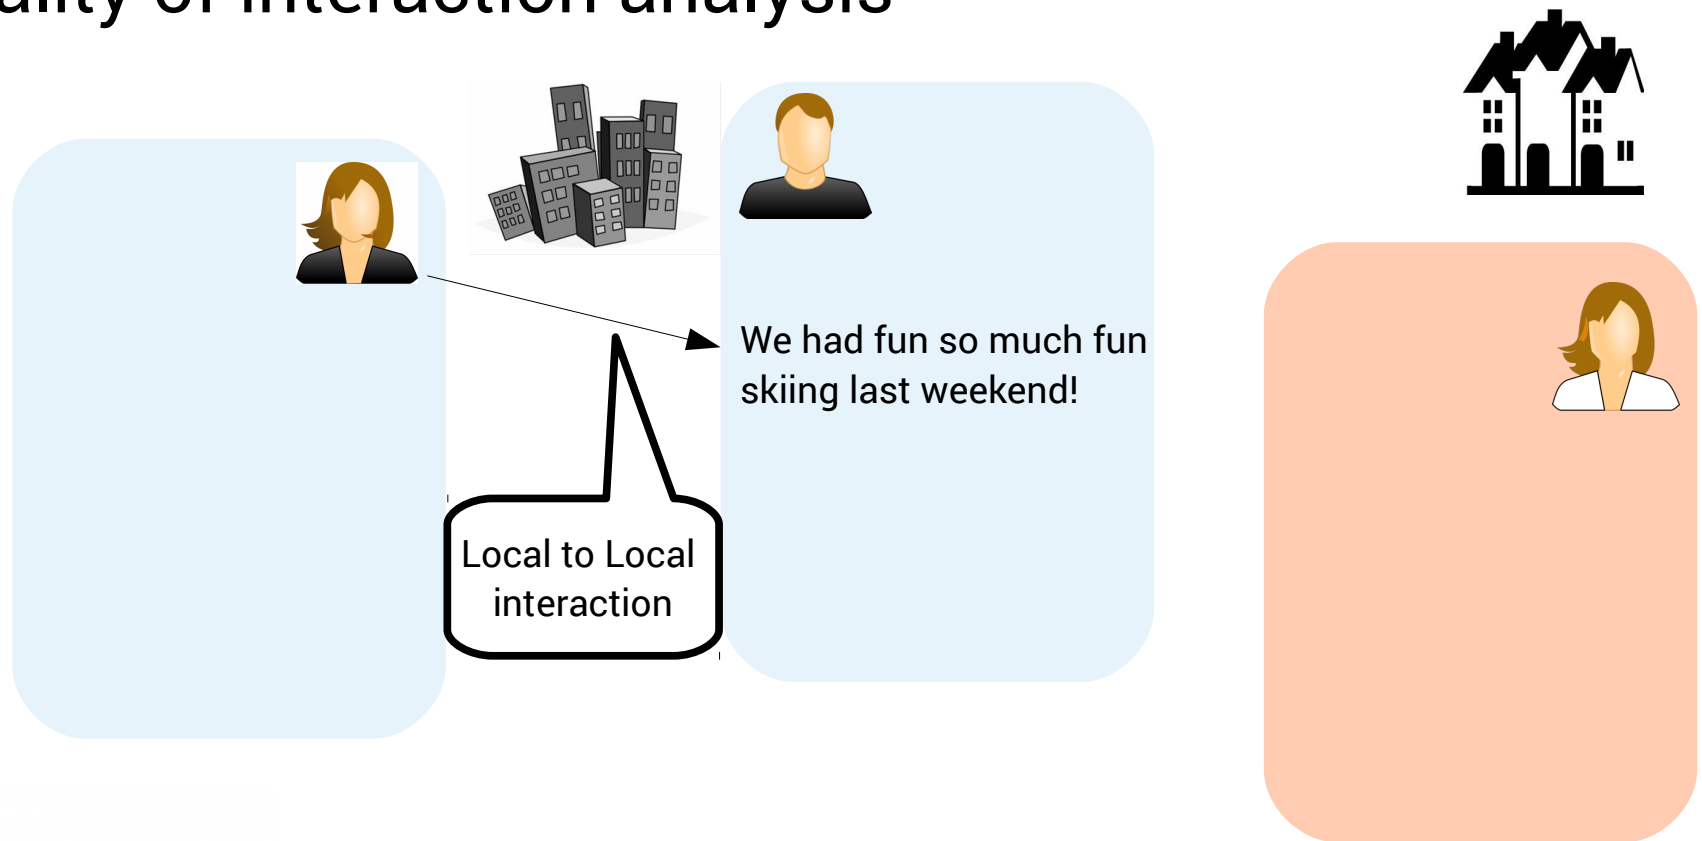
OSN Architecture – Facebook

- Can we improve it?
 - Analyse the usage to identify opportunities
 - Measure the extent of the problem
 - Evaluate solutions with real-world data



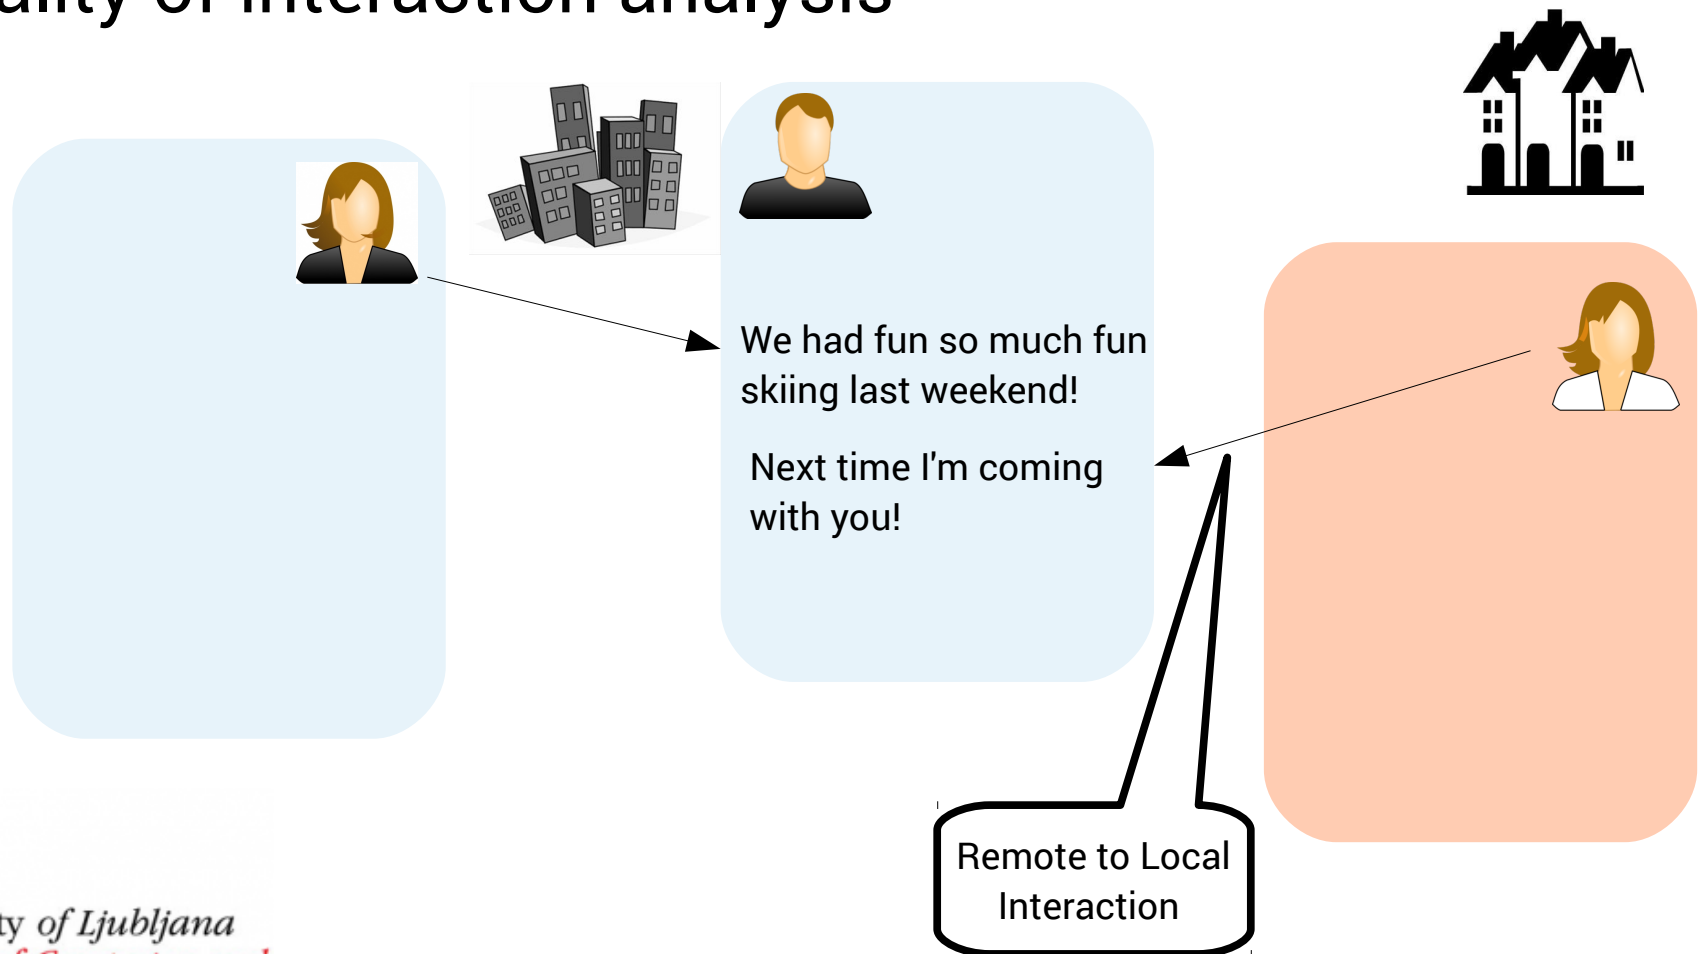
Identify Opportunities – Usage Analysis

- Data: ~10 million profiles from FB [Wilson et al. 2009]
- Locality of interaction analysis




Identify Opportunities – Usage Analysis

- Data: ~10 million profiles from FB [Wilson et al. 2009]
- Locality of interaction analysis



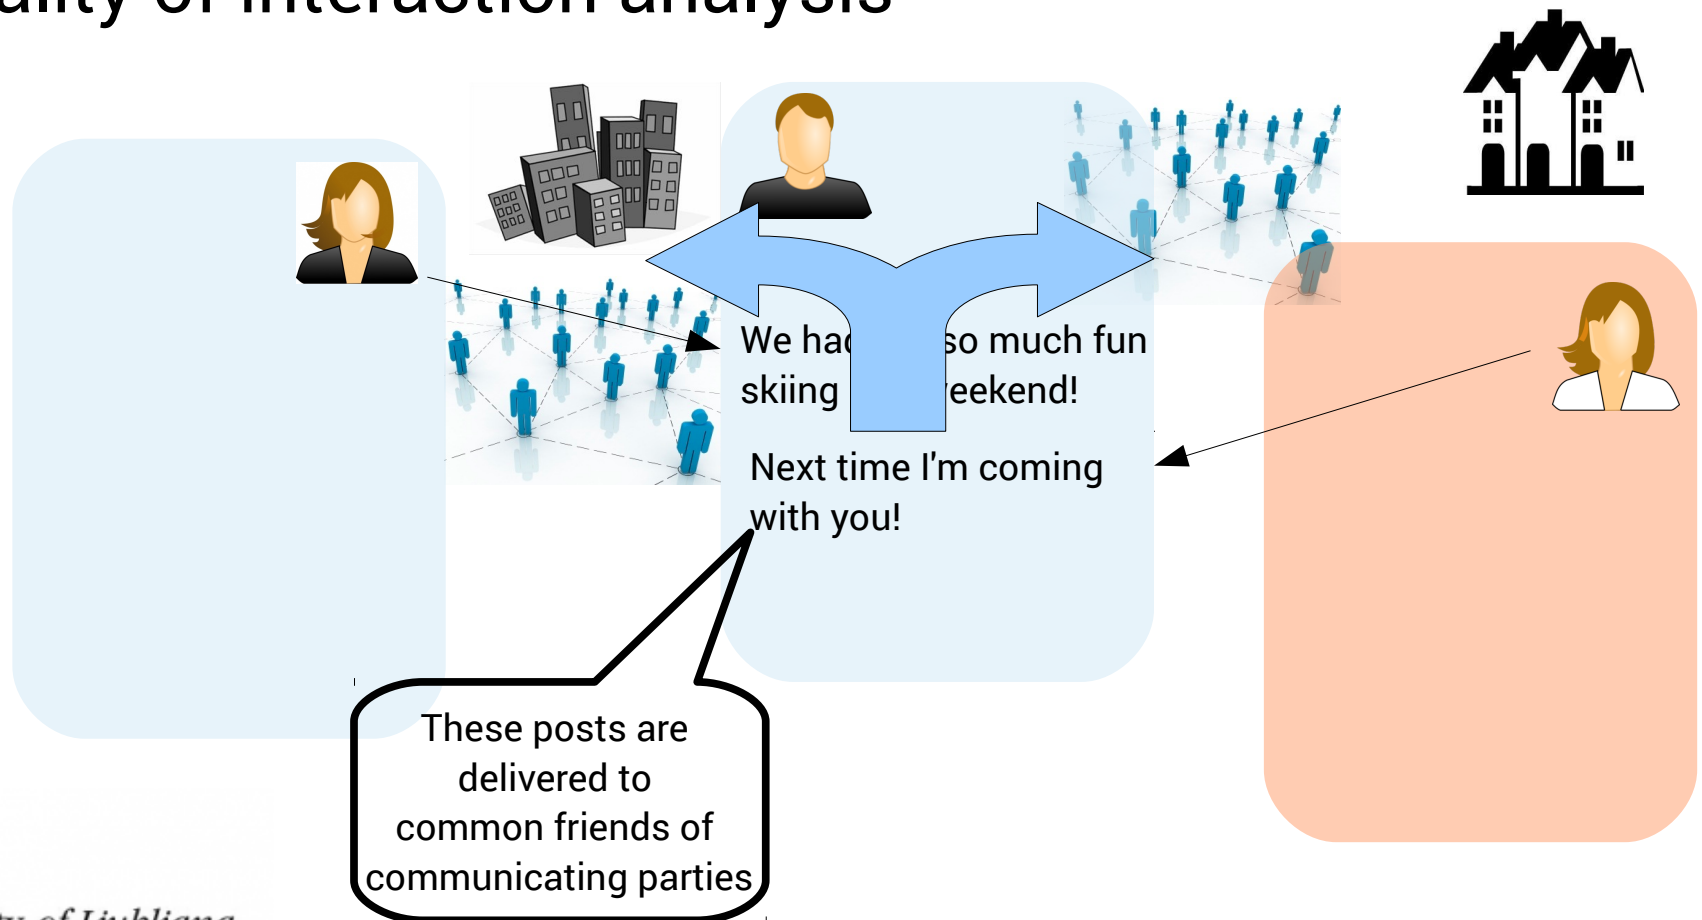
Identify Opportunities – Usage Analysis

- Data: ~10 million profiles from FB [Wilson et al. 2009]
- Locality of interaction analysis
 - Count L-L, L-R, R-L in each region





Identify Opportunities – Usage Analysis

- Data: ~10 million profiles from FB [Wilson et al. 2009]
- Locality of interaction analysis



Identify Opportunities – Usage Analysis

- Data: ~10 million profiles from FB [Wilson et al. 2009]
- Locality of interaction analysis
 - Count L-L, L-R, R-L in each region
 - Count delivery ratio of posts



Usage Analysis Conclusion: High Locality of Interaction



Measure the Problem – Performance Analysis

- Facebook transaction analysis:
 - How much, and what kind of data is transferred on Facebook?
 - Network traces - Wireshark


Photo post

TRANSFER	FROM	TO	BYTES	COMMENT
1	user	CA	1510	#upload script request
2	CA	user	5703	#upload script
3	user	CA	136802*	#photo and post text
4	CA	user	7168	#display markup
5	user	CDN	495	#display image request
6	CDN	user	3819*	#display jpeg



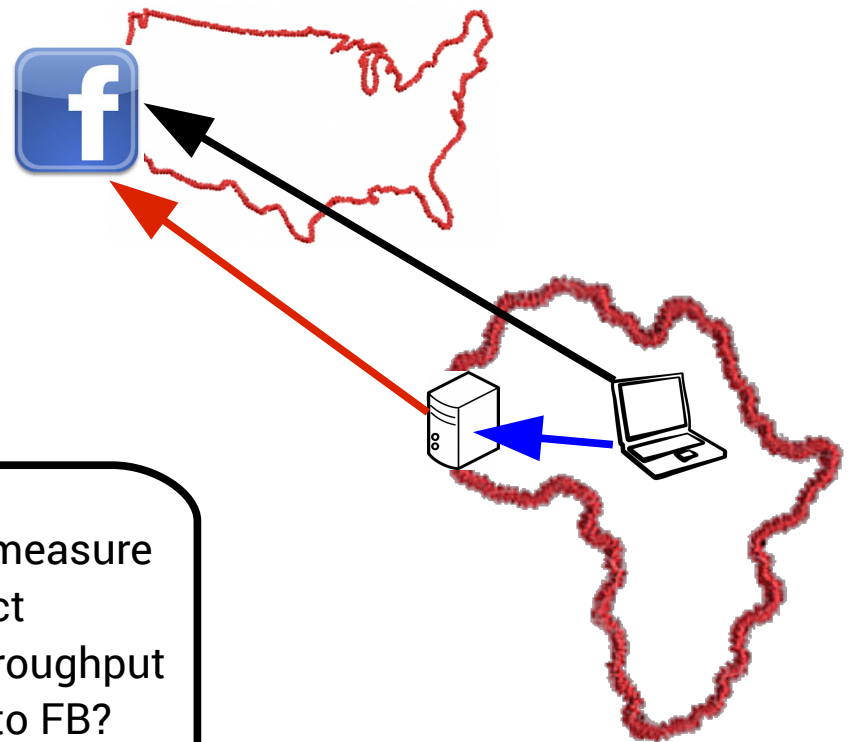
Measure the Problem – Performance Analysis

- Internet path analysis
 - Connection quality between end-users and Facebook infrastructure
 - Latency – `ping`
 - Throughput – `iperf`



Measure the Problem – Performance Analysis

- Internet path analysis
 - Connection quality between end-users and Facebook infrastructure
 - Latency – `ping`
 - Throughput – `iperf`
 - Approximate end points with `planetlab` nodes



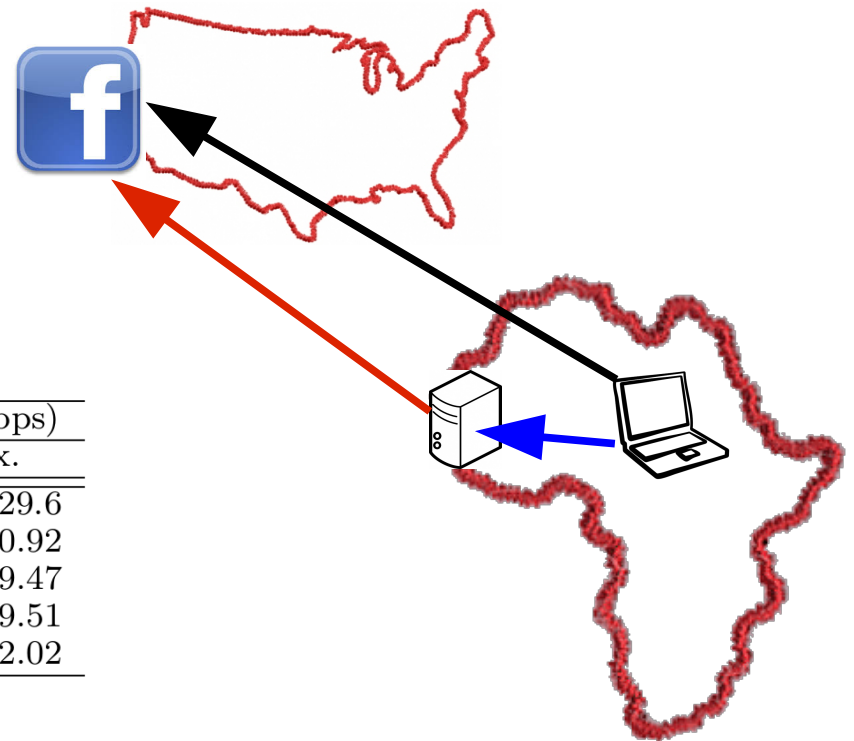
Why don't we measure the exact latency and throughput from users to FB?



Measure the Problem – Performance Analysis

- Internet path analysis
 - Connection quality between end-users and Facebook infrastructure
 - Latency – ping
 - Throughput – iperf
 - Approximate end points with planetlab nodes

Region	Latency (ms)			Loss (%)			Capacity (Mbps)		
	dir.	prox.		dir.	prox.		dir.	prox.	
Russia	148	115	31	6.1	0	1.8	29.6	367	29.6
Egypt	164	176	67	5.8	0	5.8	0.92	736	0.92
Sweden	104	95	14	0.32	0	2.9	9.47	188	9.47
NYC	74	43	33	0.75	0	0.6	9.51	99	9.51
LA	27	9.1	18	0.50	0	0.4	2.02	228	2.02




Performance Analysis Conclusions:

- Parts of a transaction can be sent irrespective of the data
- Server to FB connection is much better than end-user to FB connection



Solution: Alternative Infrastructure

- 1) TCP proxying with regional servers
 - Split connection at a regional server





Region	Latency (ms)		Loss (%)			Capacity (Mbps)			
	dir.	prox.	dir.	prox.	dir.	prox.			
Russia	148	115	31	6.1	0	1.8	29.6	367	29.6
Egypt	164	176	67	5.8	0	5.8	0.92	736	0.92
Sweden	104	95	14	0.32	0	2.9	9.47	188	9.47
NYC	74	43	33	0.75	0	0.6	9.51	99	9.51
LA	27	9.1	18	0.50	0	0.4	2.02	228	2.02



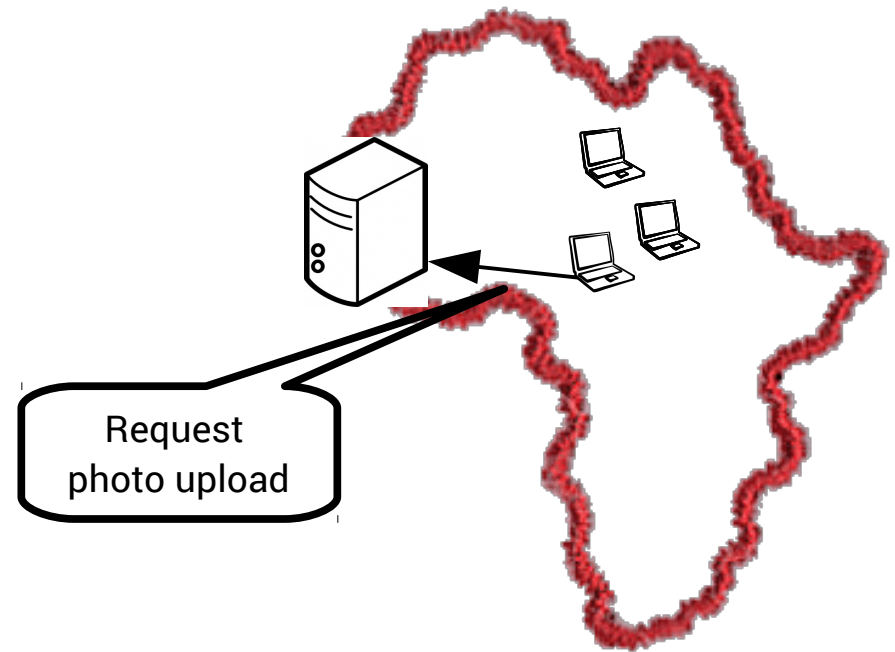
Solution: Alternative Infrastructure

- 1) TCP proxying with regional servers
 - Split connection at a regional server
 - Isolates the effect of poor links
 - Allows a wider TCP window



Solution: Alternative Infrastructure

- 2) Regional OSN cache
 - Regional server:
 - Regional social graph
 - Cache recent local content

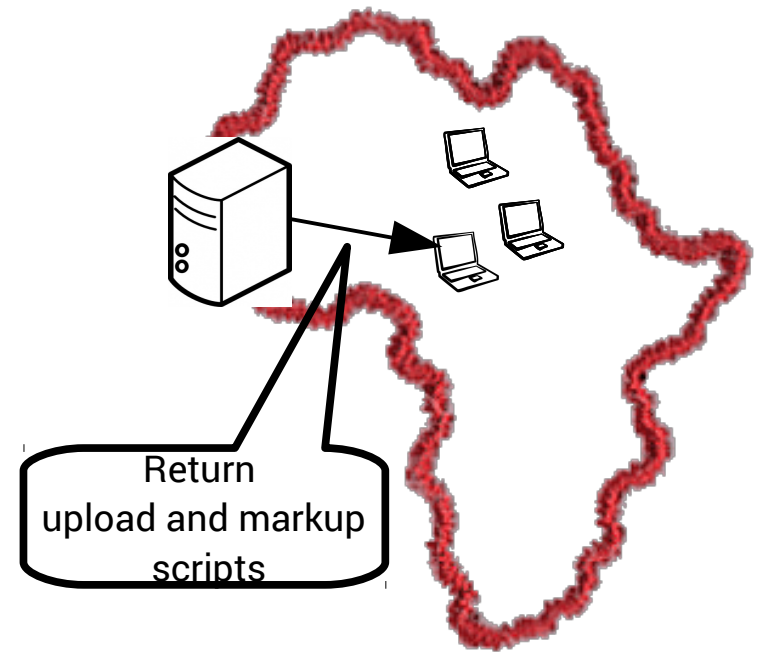


Solution: Alternative Infrastructure

- 2) Regional OSN cache
 - Regional server:
 - Regional social graph
 - Cache recent local content




TRANSFER	FROM	TO	BYTES	COMMENT
1	user	CA	1510	#upload script request
2	CA	user	5703	#upload script
3	user	CA	136802*	#photo and post text
4	CA	user	7168	#display markup
5	user	CDN	495	#display image request
6	CDN	user	3819*	#display jpeg



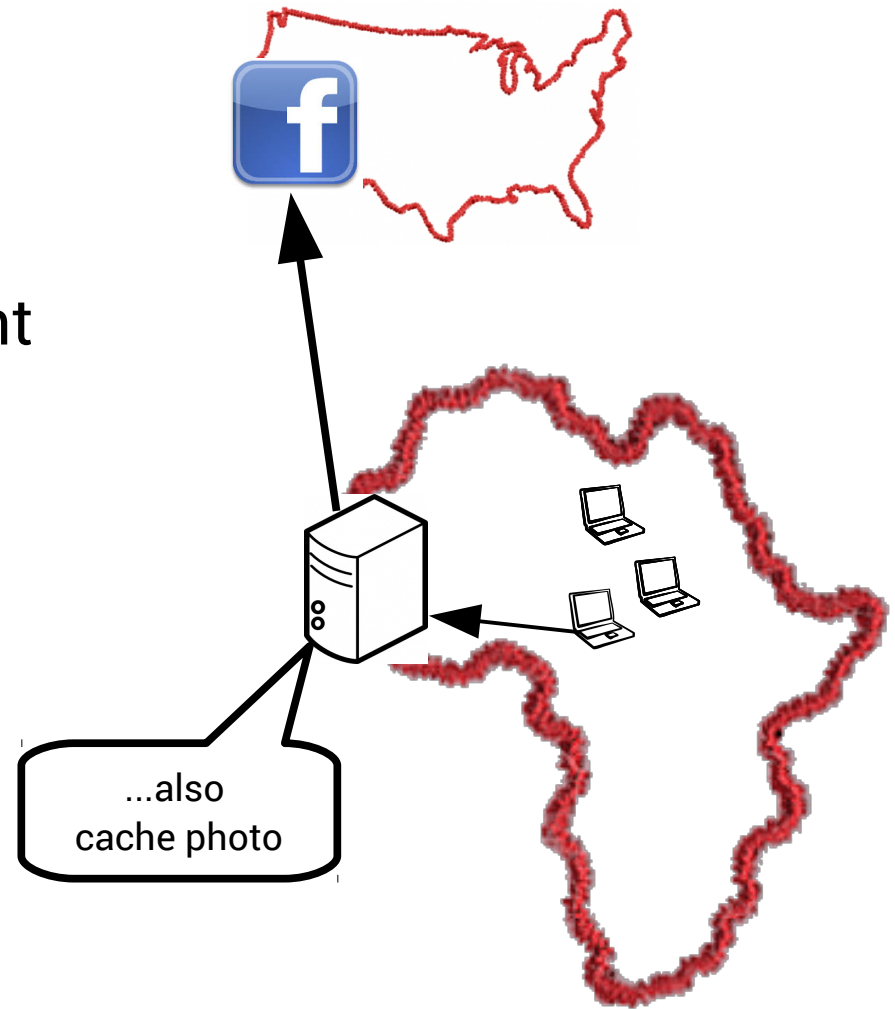
Solution: Alternative Infrastructure

- 2) Regional OSN cache
 - Regional server:
 - Regional social graph
 - Cache recent local content



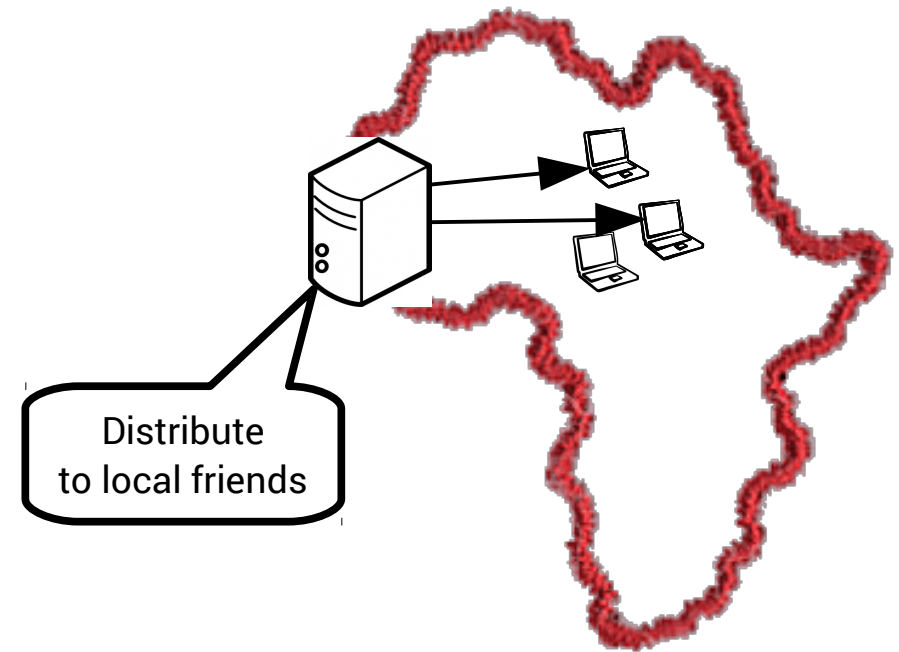
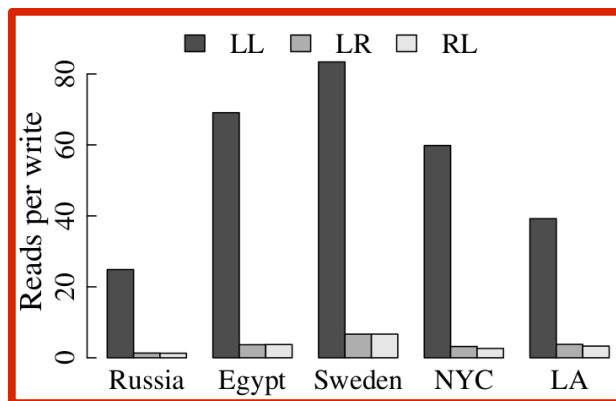
Solution: Alternative Infrastructure

- 2) Regional OSN cache
 - Regional server:
 - Regional social graph
 - Cache recent local content



Solution: Alternative Infrastructure

- 2) Regional OSN cache
 - Regional server:
 - Regional social graph
 - Cache recent local content



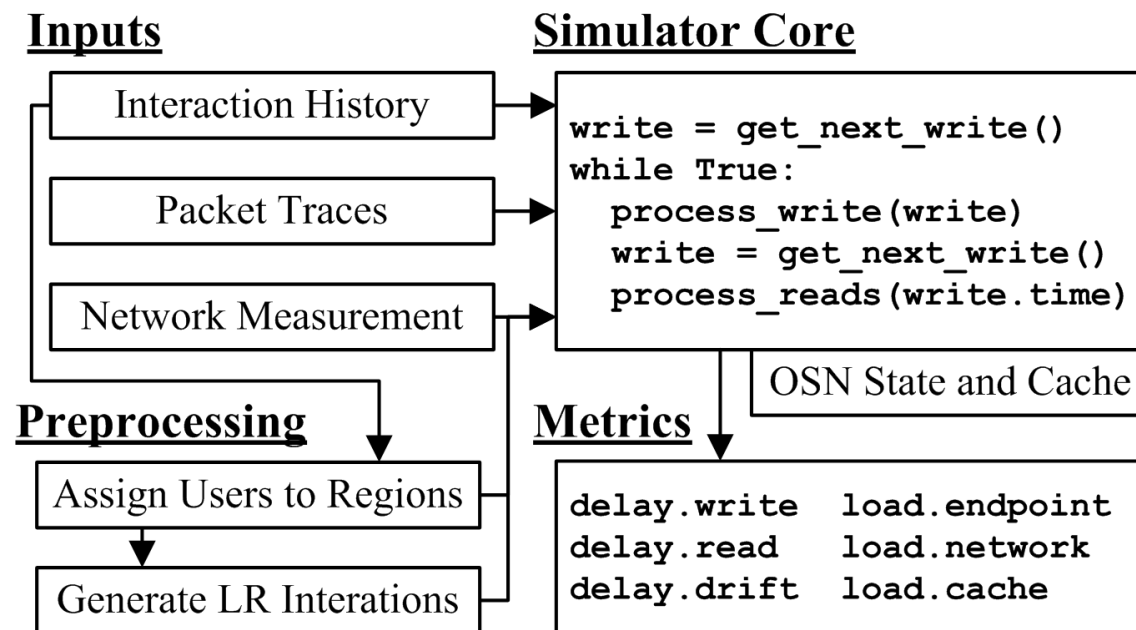
Evaluate Solutions

- Data trace-based evaluation
 - Replay all the interactions identified in the dataset
 - Assign the measured “cost” to each interaction
 - Process interactions according to:
 - Current FB architecture
 - Regional TCP proxy architecture (solution 1)
 - Regional OSN cache architecture (solution 2)



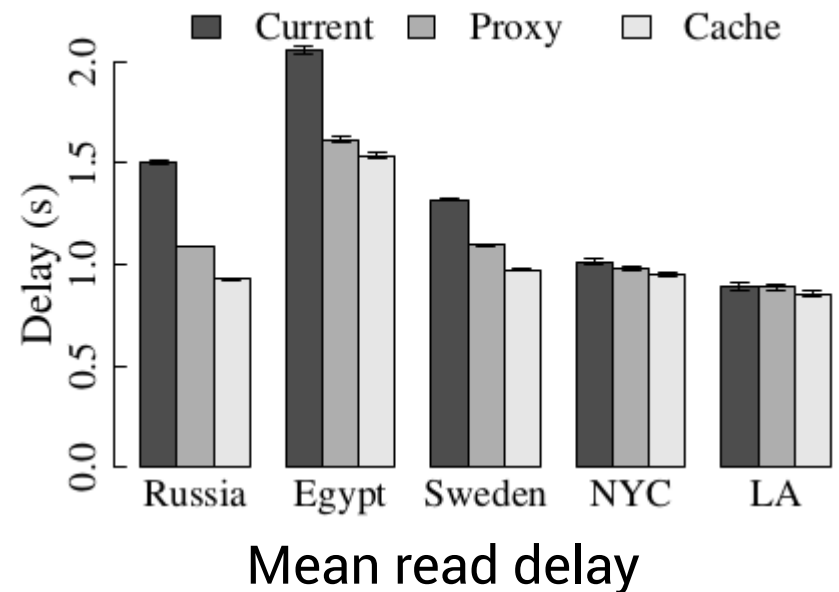
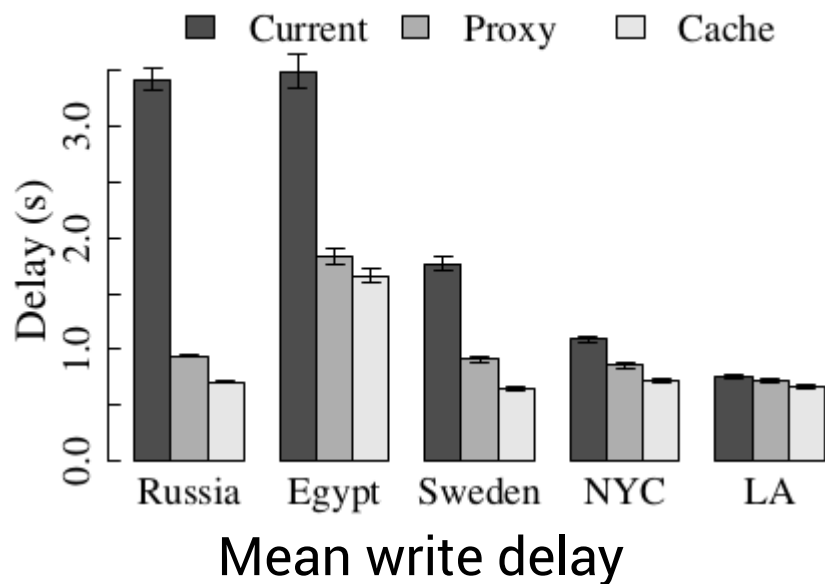
Evaluate Solutions

- Data trace-based evaluation



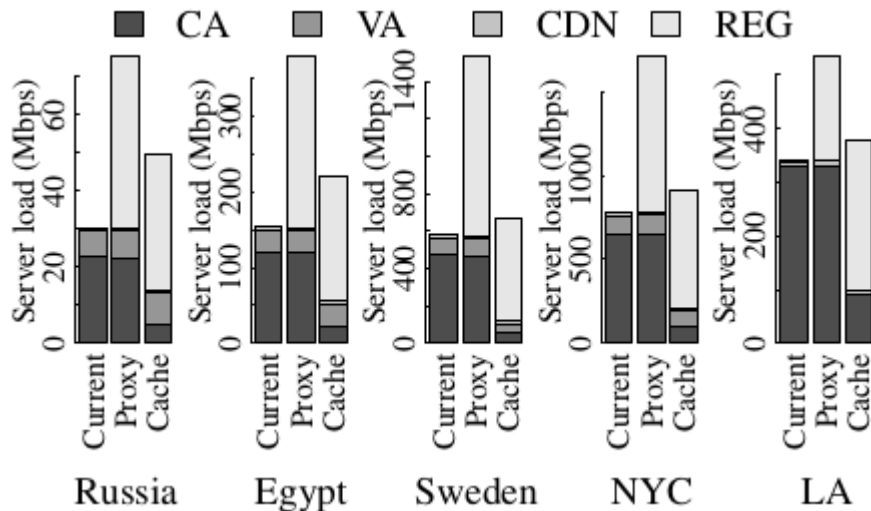
Evaluate Solutions

- Results
 - Delay

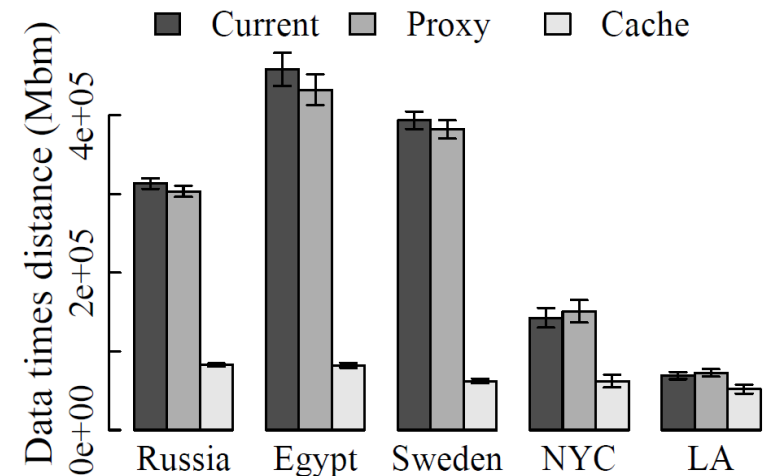


Evaluate Solutions

- Results
 - Server and network load



Server endpoint load



Network load

Think about the economic impact!



Case Study Conclusions

- Network analysis revealed:
 - Bottlenecks in OSN data transfer
 - Locality of interest
- Proposed alternatives:
 - Regional TCP proxying
 - Regional cache servers
- Evaluation:
 - Trace-based simulation
 - Improved user experience and lower system load



References

- Tools and tutorials:
 - www.wireshark.org/docs/wsug_html_chunked/
 - Read man pages for iperf, tcpdump, netstat, ping, traceroute
- Papers:
 - Vern Paxons, “**Strategies for Sound Internet Measurement**”, IMC'04
www.icir.org/vern/papers/meas-strategies-imc04.pdf
 - Christo Wilson, Bryce Boe, Alessandra Sala, Krishna P. N. Puttaswamy, and Ben Y. Zhao “**User Interactions in Social Networks and their Implications**” EUROSYS'09
<http://www.ccs.neu.edu/home/cbw/pdf/interaction-eurosys09.pdf>
 - Mike P. Wittie, Veljko Pejovic, Lara Deek, Kevin C. Almeroth, Ben Y. Zhao “**Exploiting Locality of Interest in Online Social Networks**” ACM CoNEXT'10 www.cs.bham.ac.uk/~pejovicv/docs/Wittie10conext.pdf



Thank you!

Veljko Pejović

Veljko.Pejovic@fri.uni-lj.si

FRI, Ljubljana



University of Ljubljana
Faculty of Computer and
Information Science