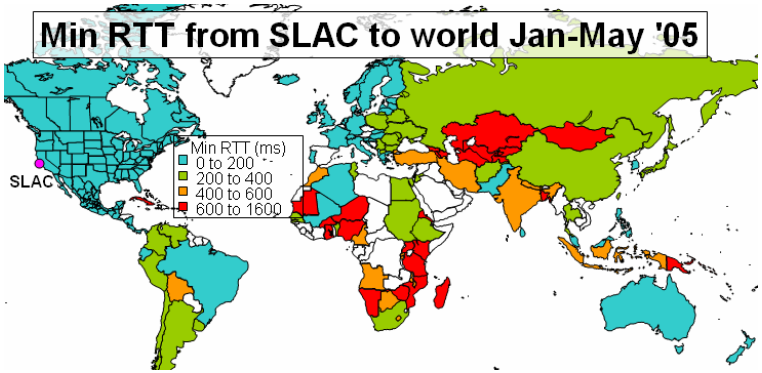

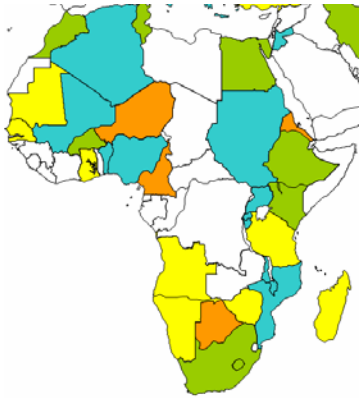

Quantifying the Digital Divide from Within and Without

Les Cottrell, SLAC

International ICFA Workshop on HEP Networking, Grid and Digital Divide Issues for Global e-Science,

Daegu, Korea, May 23-27, 2005

www.slac.stanford.edu/grp/scs/net/talk05/icfa-korea-may05.ppt

Initially funded by DoE Field Work proposal. Currently partially funded by US Department of State/Pakistan Ministry of Science & Technology

Goal

- Measure the network performance for developing regions
  - From developed to developing & vice versa
  - Between developing regions & within developing regions
- Use simple tool (PingER/ping)
  - Ping installed on all modern hosts, low traffic interference,
- Provides very useful measures
- Originated in High Energy Physics, now focused on DD
- Persistent (data goes back to 1995), interesting history

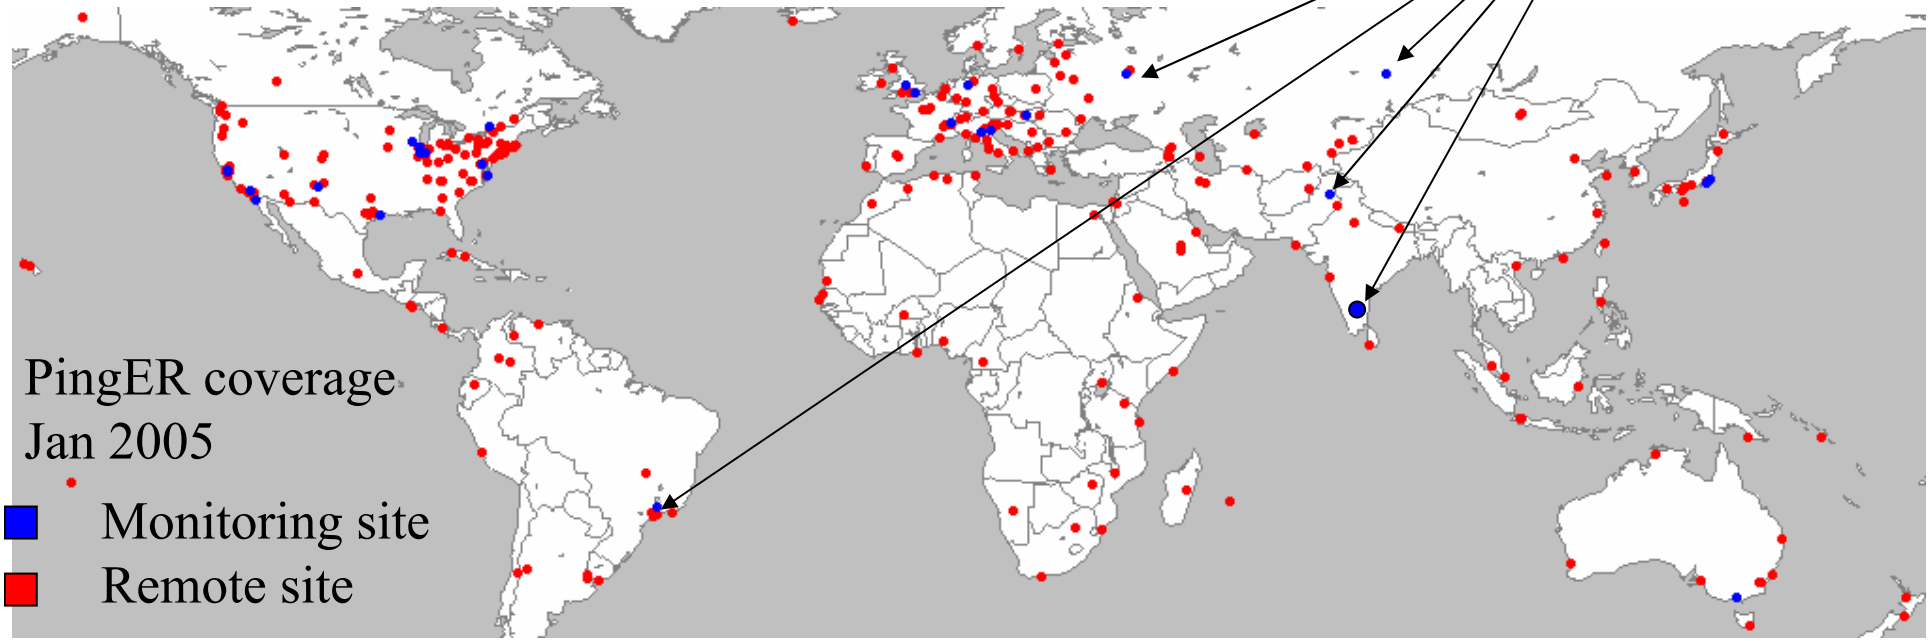

Min-RTT to World

- Measured from SLAC
- 118 countries, 345 sites
- Need contacts in uncolored
- > 600ms ≡ satellite (red)
- <100ms inside N. America

January 2000

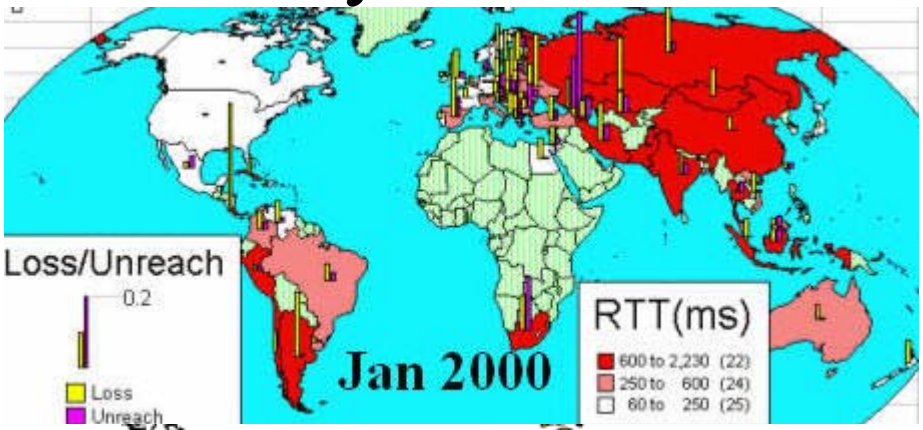

- Japan via NY to SLAC
- Korea via W. Coast

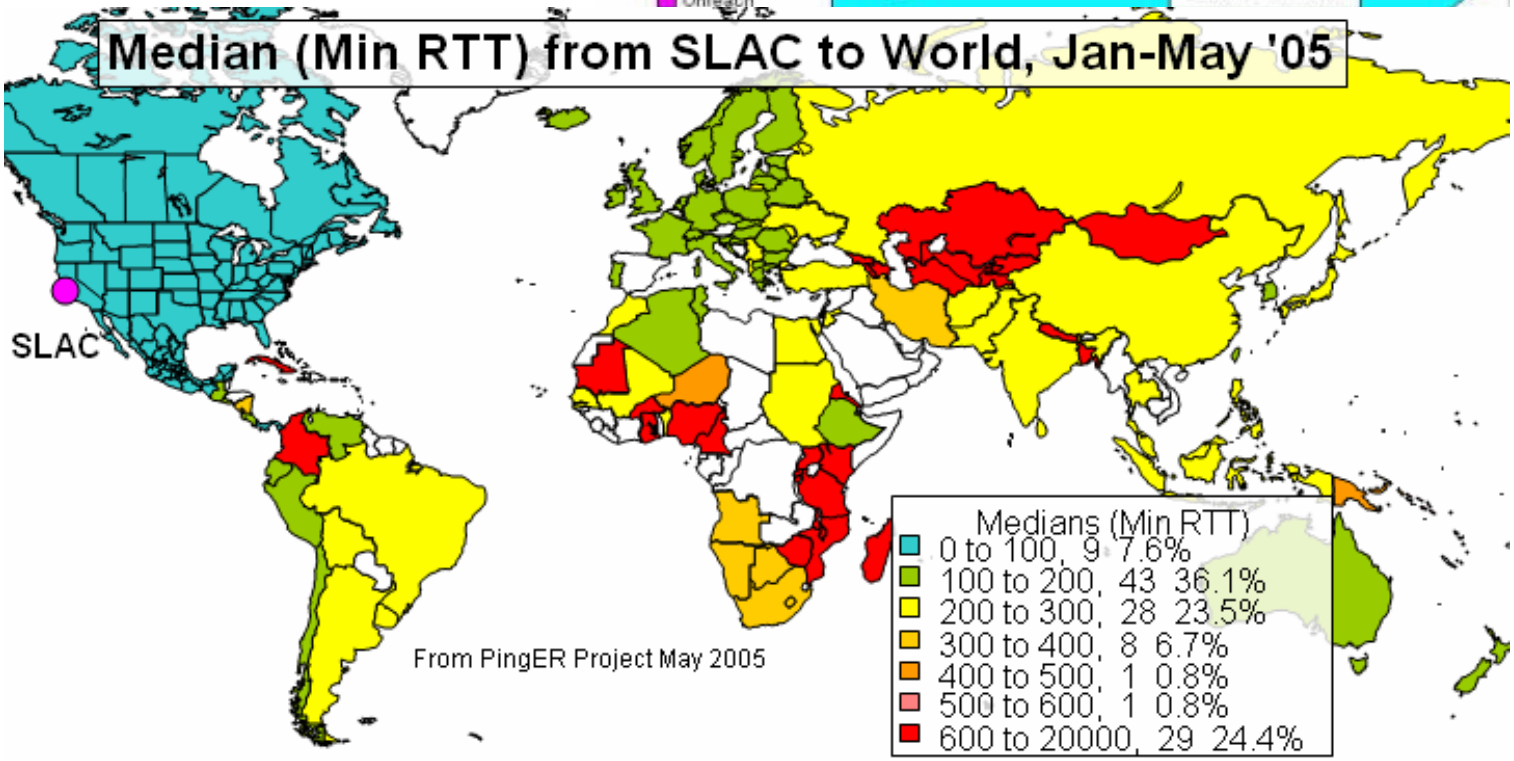

Loss to the World

- Loss is less distance dependent than RTT
- It has a big effect on perceived performance
  - Good < 1%, acceptable < 3%, > 5-12% sessions time out

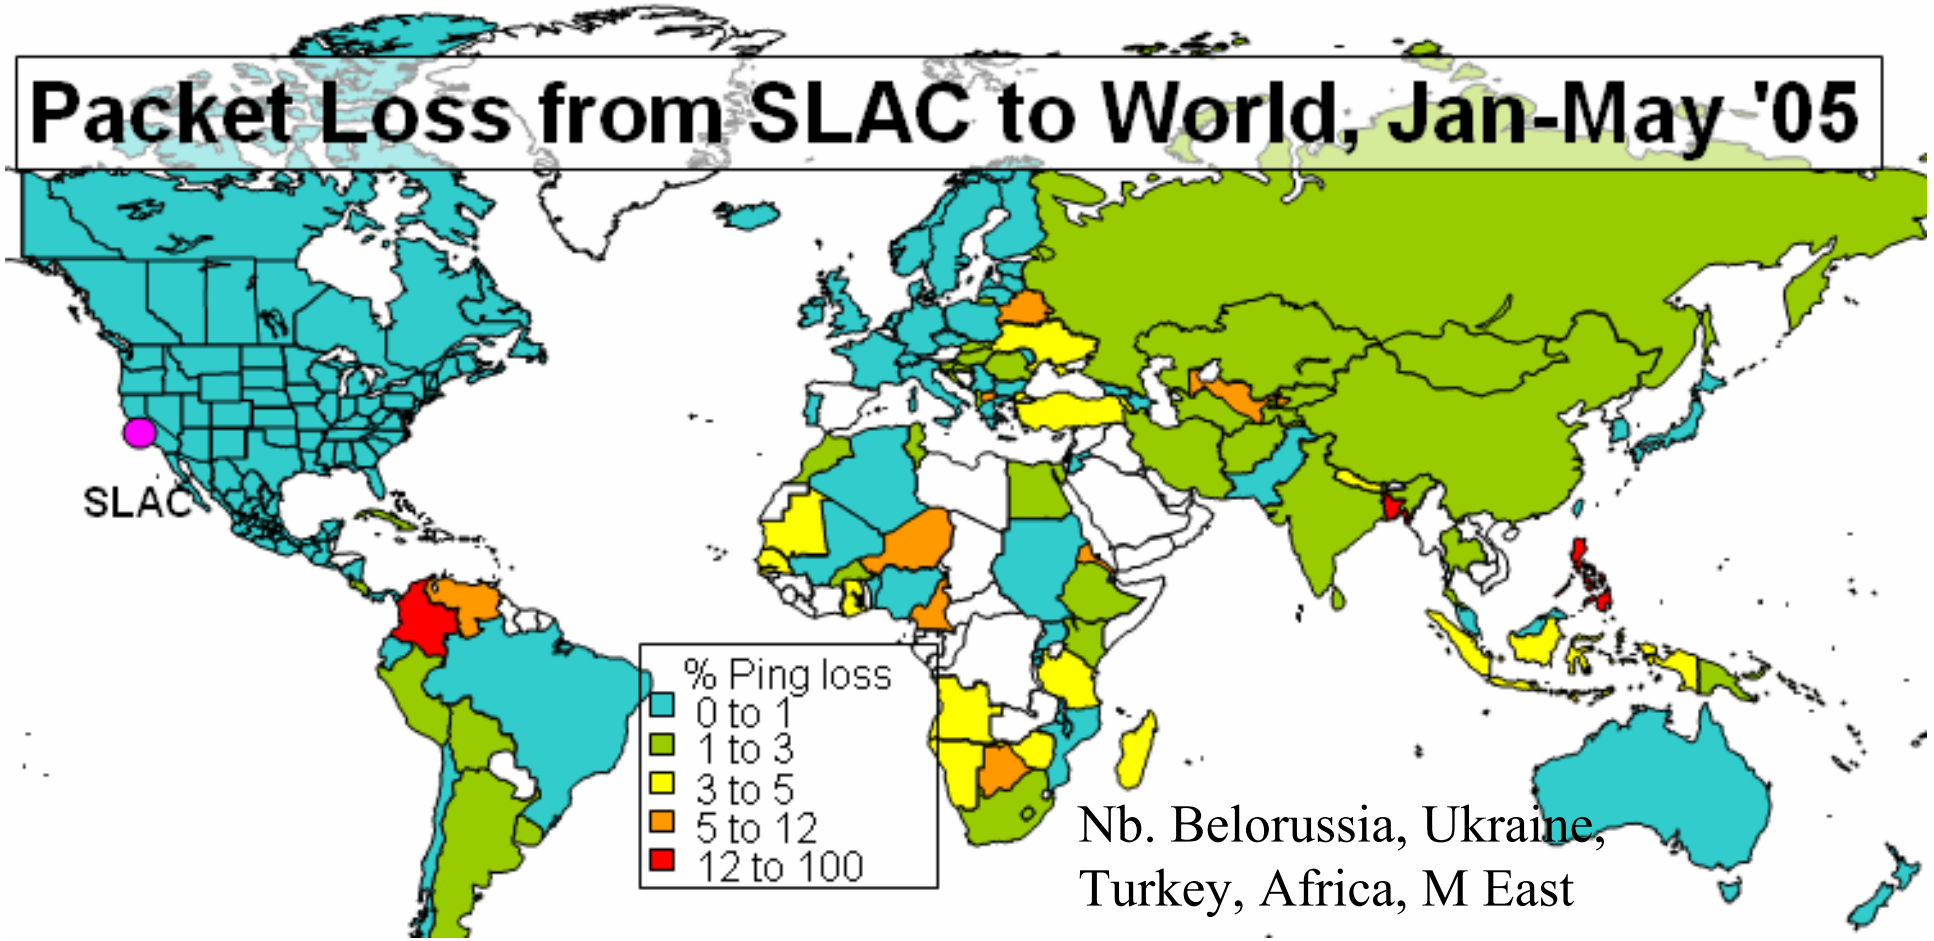

World "Quality"

S.E. Europe, Russia: **catching up**
Latin Am., Mid East, China: **keeping up**
India, Africa: **falling behind**

C. Asia, Russia, S.E. Europe, L. America, M. East, China: 4-5 yrs behind

India, Africa: 7 yrs behind

Important for policy makers

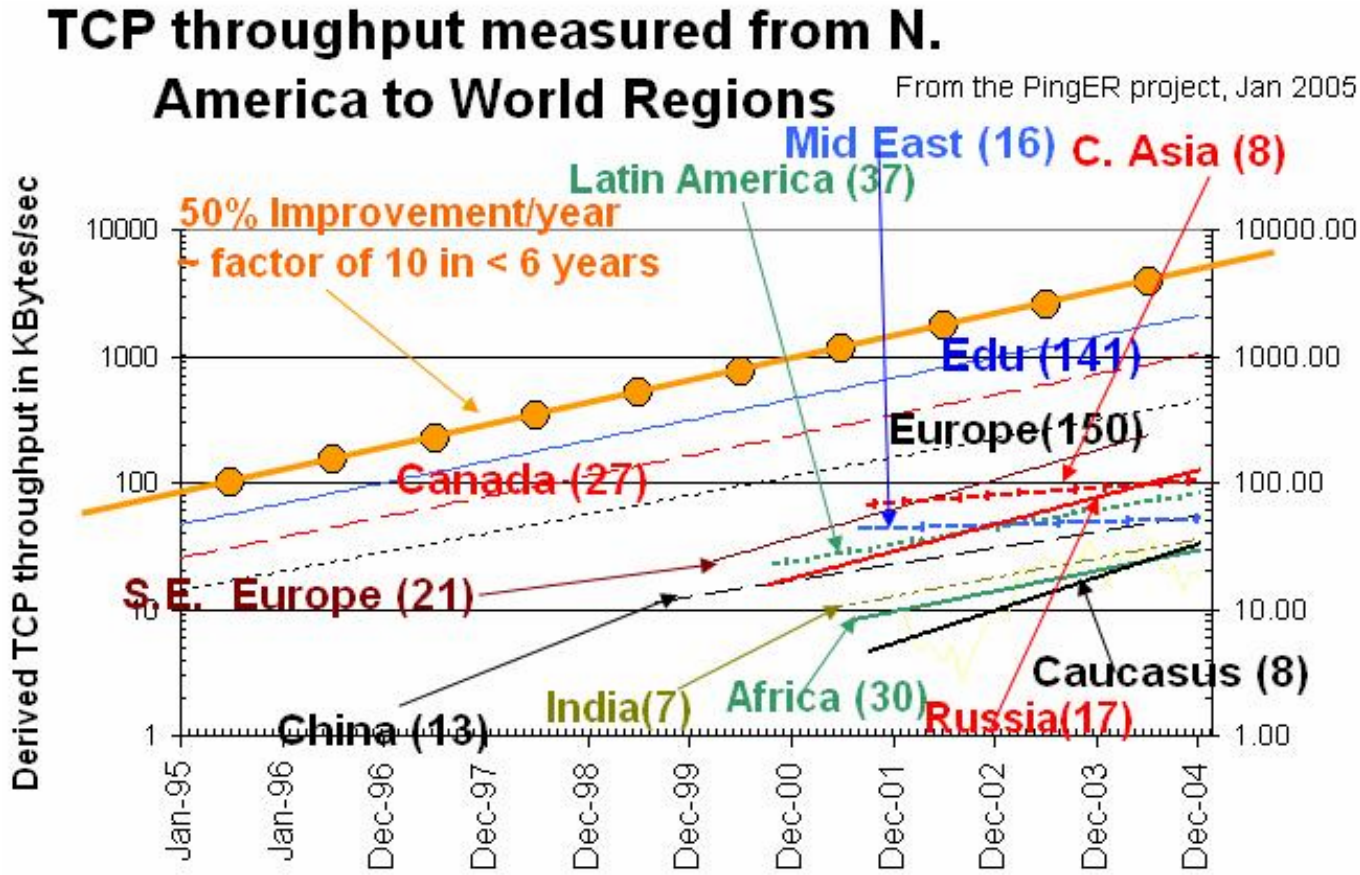

Many institutes in developing world have less performance than a household in N. America or Europe

Seen from Europe

- From CERN similar conclusions

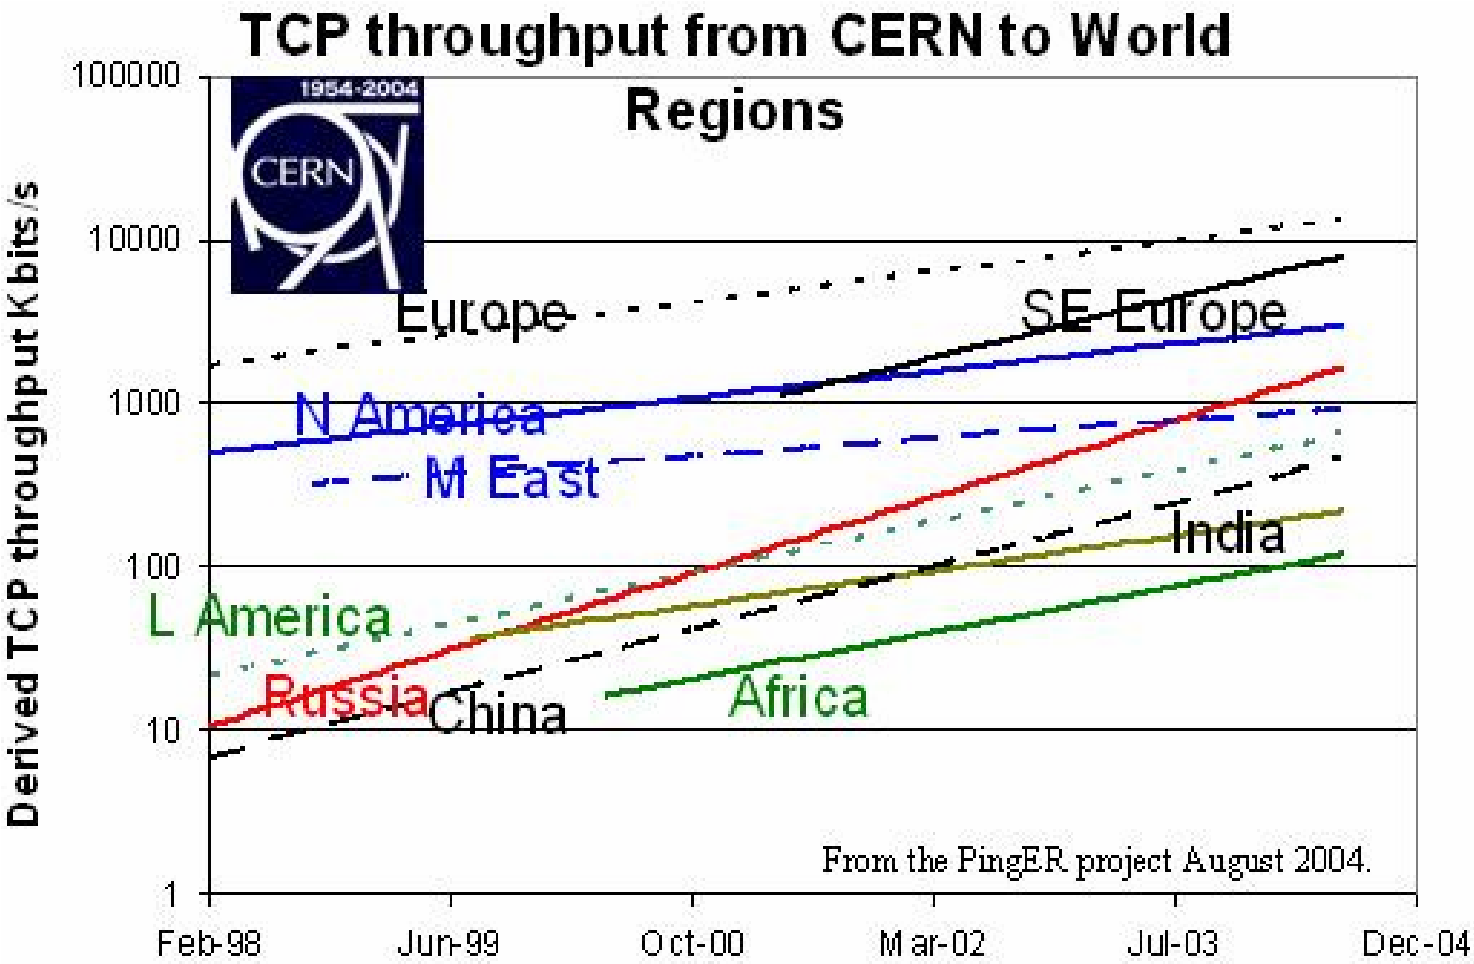

Losses

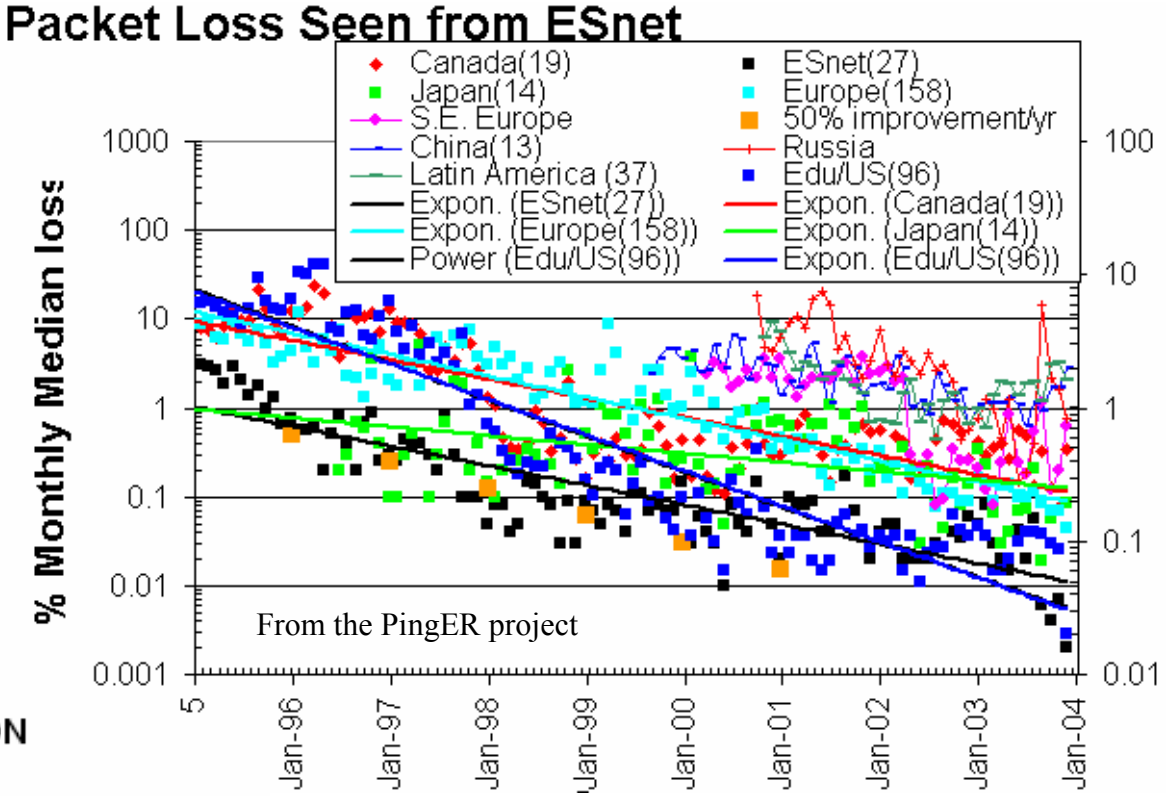

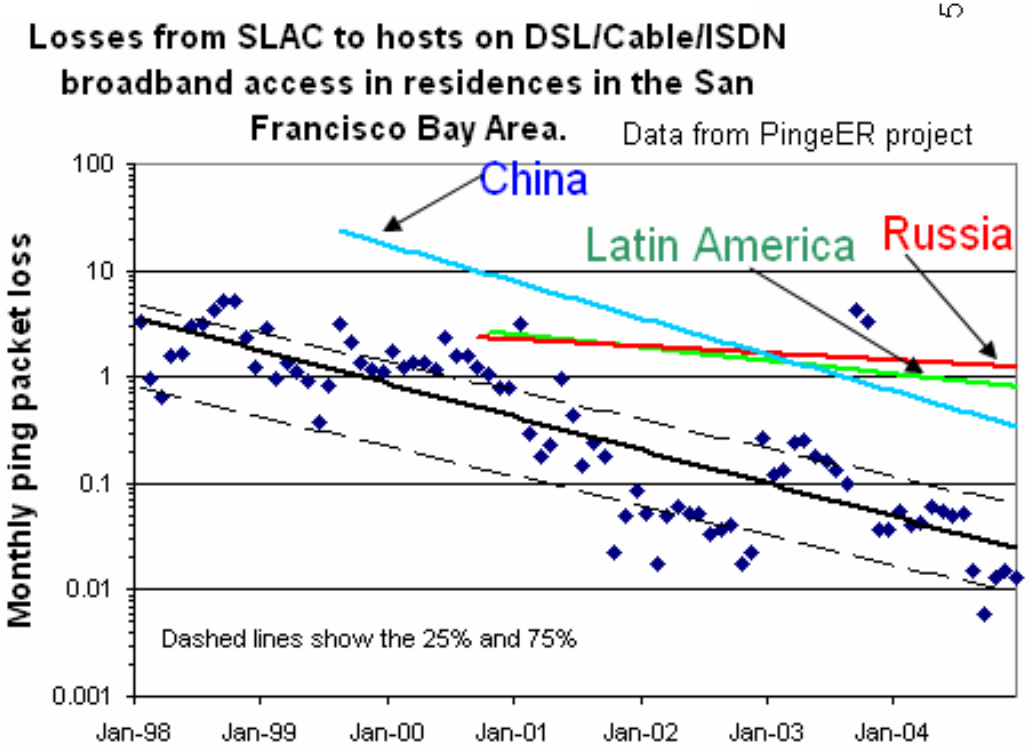

- US residential Broadband users have better access than sites in many regions

Loss to world from US

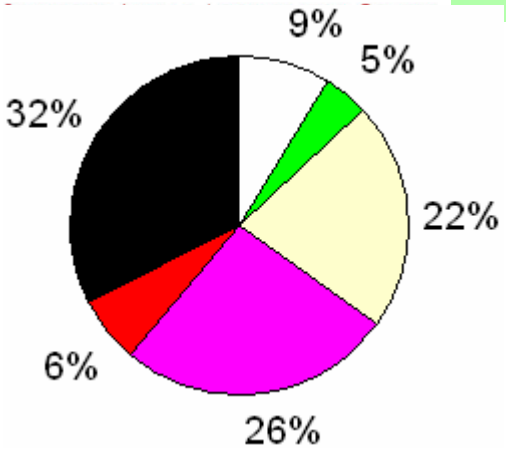

Jan-May 2001

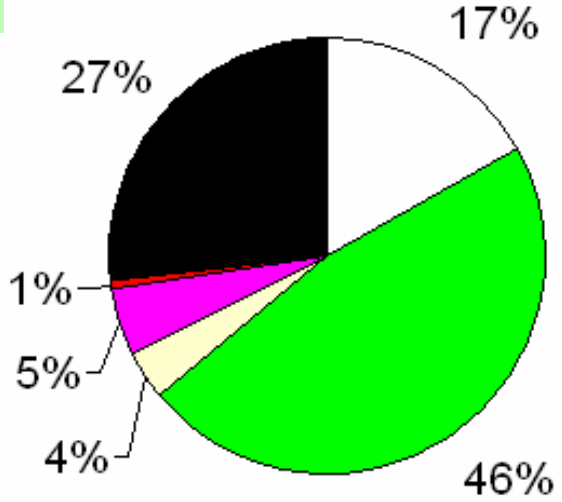

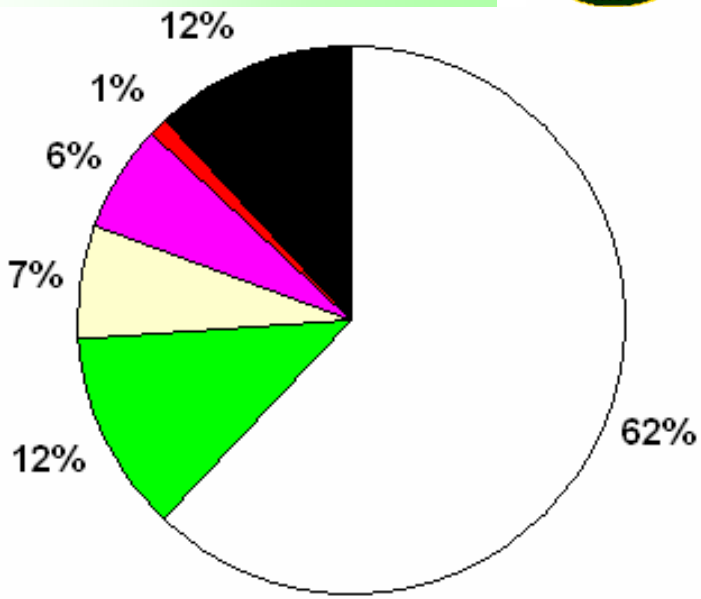

- *In 2001 <14% of the world's population had Good or Acceptable Loss performance*
- *BUT by May 2003 It had improved to 63%*
- *& by May 2005 It had improved to 74%*

- Unmeasure
- Good
- Acceptable
- Poor
- Very Poor
- Bad

Loss Rate

< 0.1 to 1 %
1 to 2.5 %
2.5 to 5 %
5 to 12 %
> 12 %

Loss to Africa (example of variability)

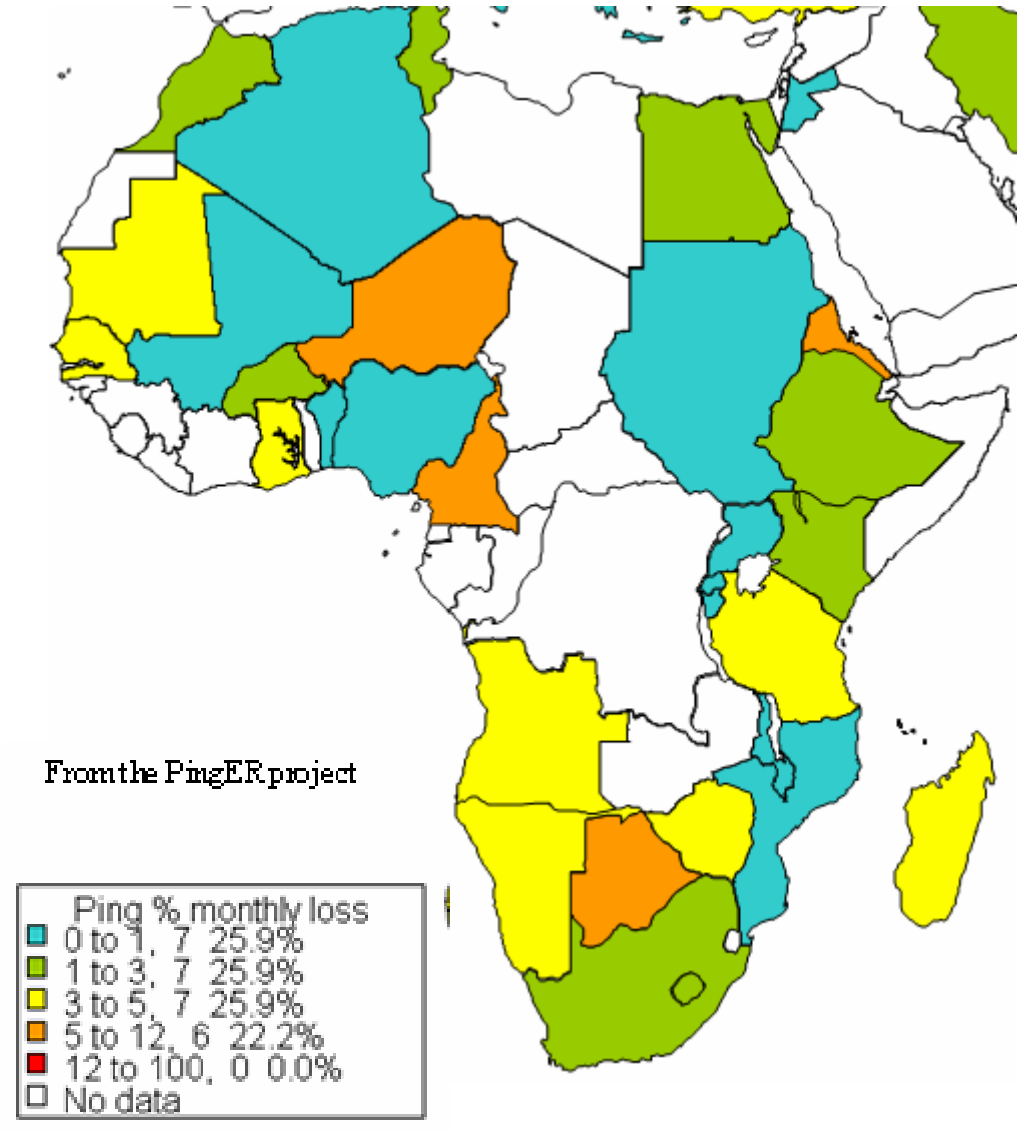

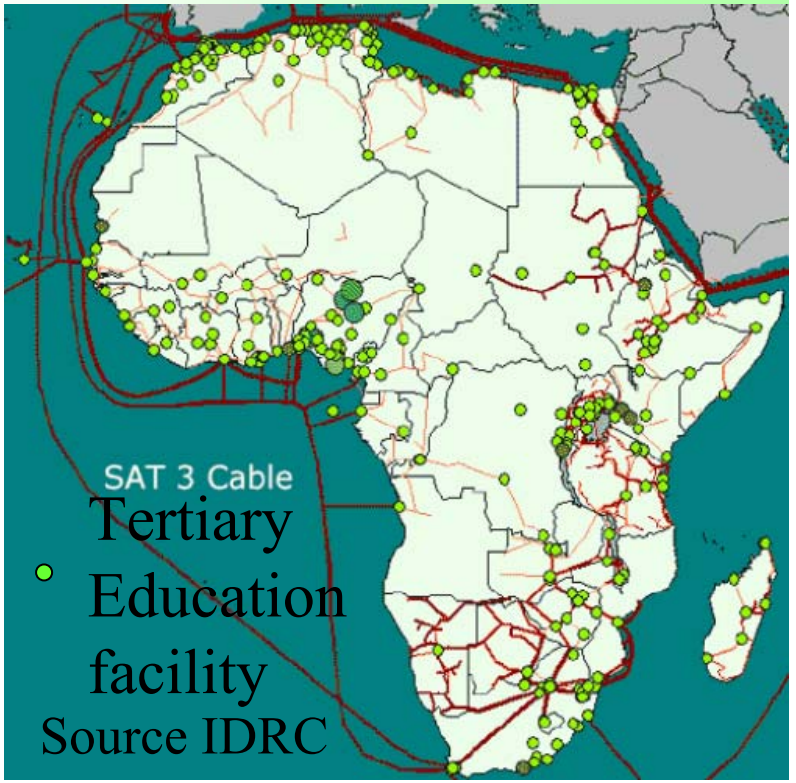

- Note we cover most countries with many tertiary education centers (83% pop)

Digression on problems, esp. for developing regions

- Want > 1 site/country to avoid anomalies
- Hosts block pings or do not respond
  - E.g. of top 25 Korean Universities (by Google search), only 7 respond to ping
  - For Sri Lanka could only find 2 hosts out of 20 that respond
- Web hosts with TLDs in many developing countries have proxies in developed countries
  - Use IP2Location.com,
  - And traceroute to verify location,
  - working on triangulation

From India

- Asia (=India): only to itself 0.04%, i.e. good site
- E.Asia = JP, TW, CN; Balkans=GR,SI,HR;
- L. America=AR,BR,CL; Oceania=AU,NZ

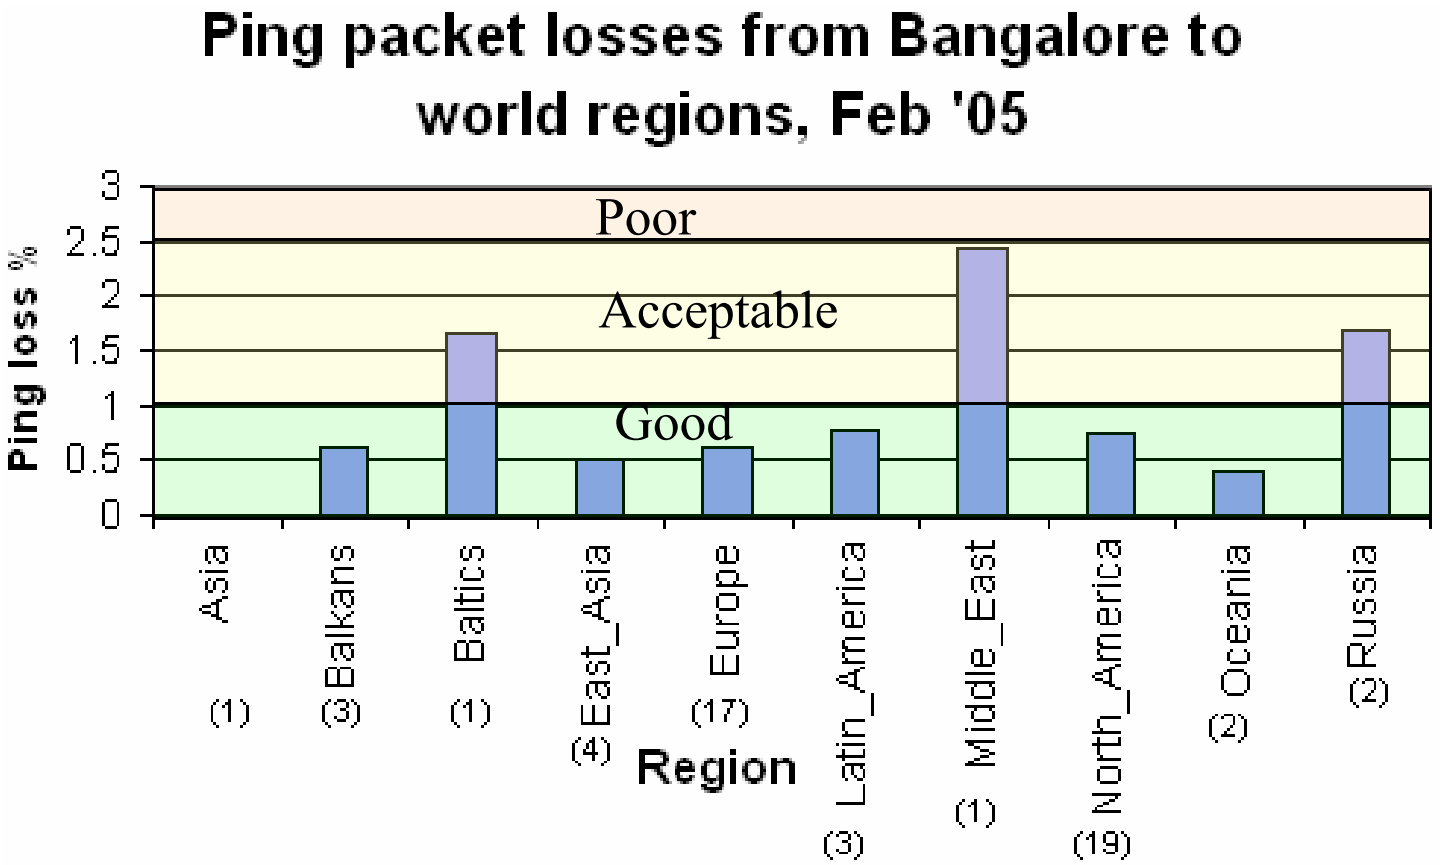

From Pakistan RTT

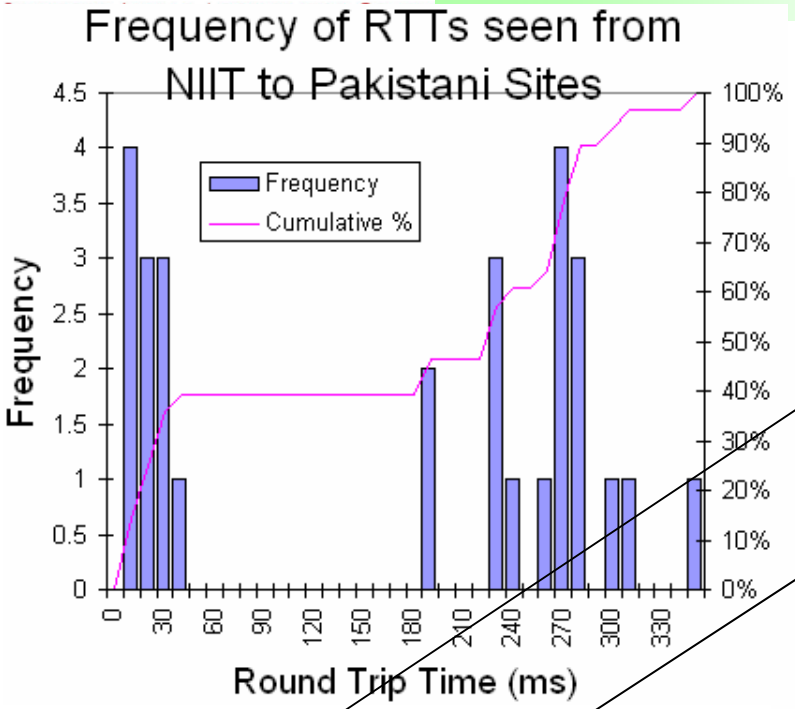

- Some routes direct <40 ms
- Some via outside world > 150ms

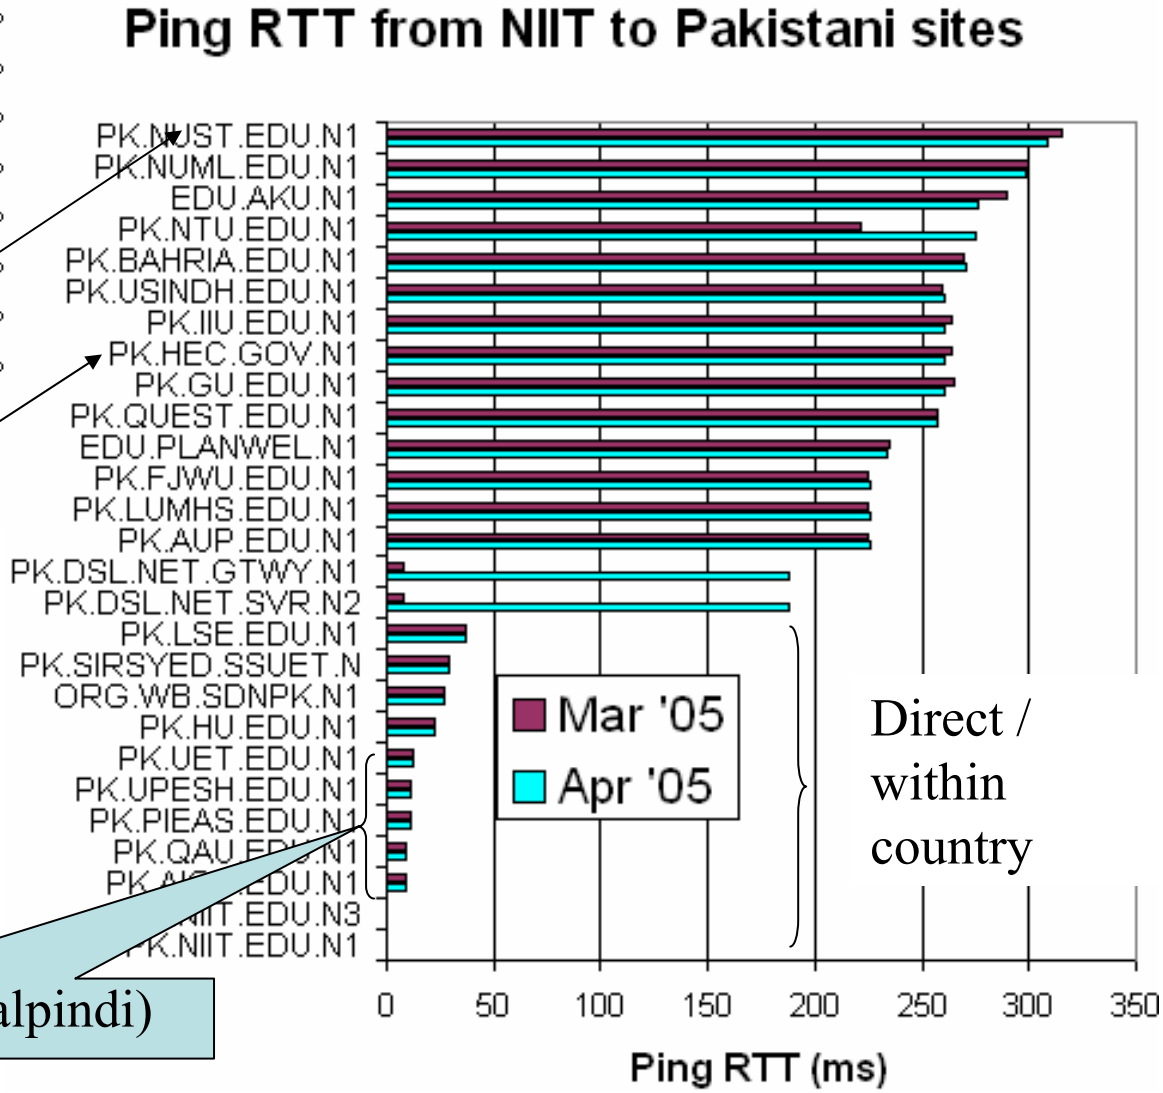

Note NUST (parent organization) but host is in California!

HEC funding agency 10km away in ISL BUT this host is in US

Proxy at NTC (ISP in Rawalpindi)

Pakistan Loss

- NIIT/Rawalpindi since Feb'05 monitoring:

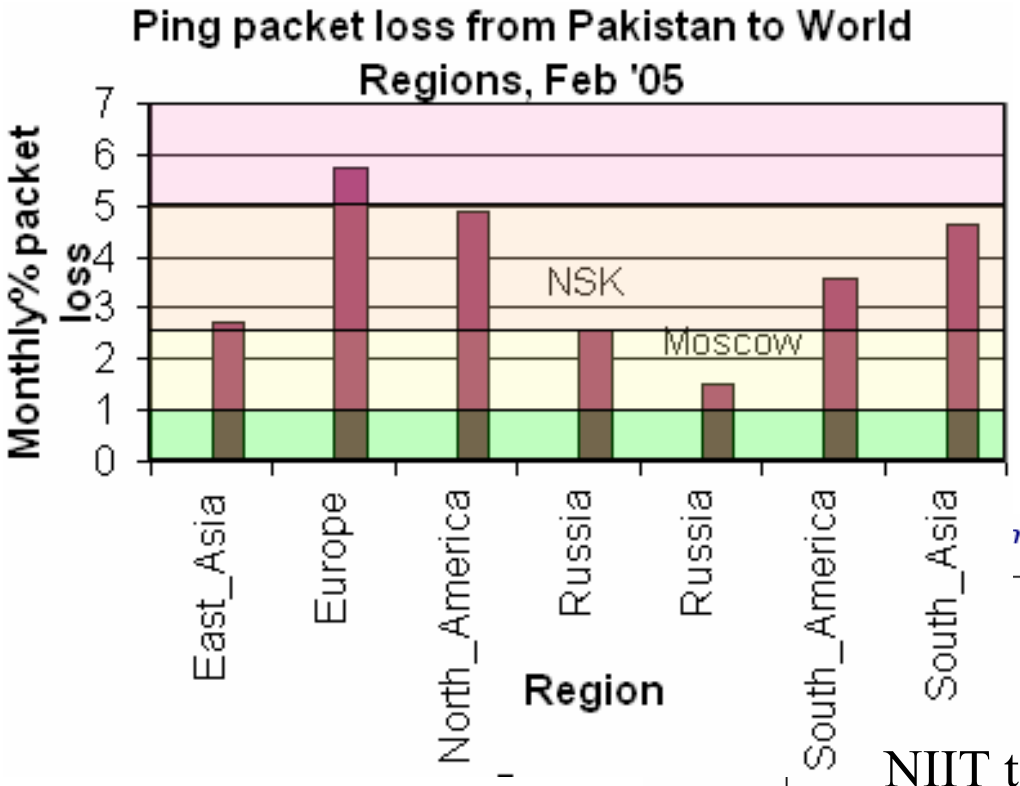

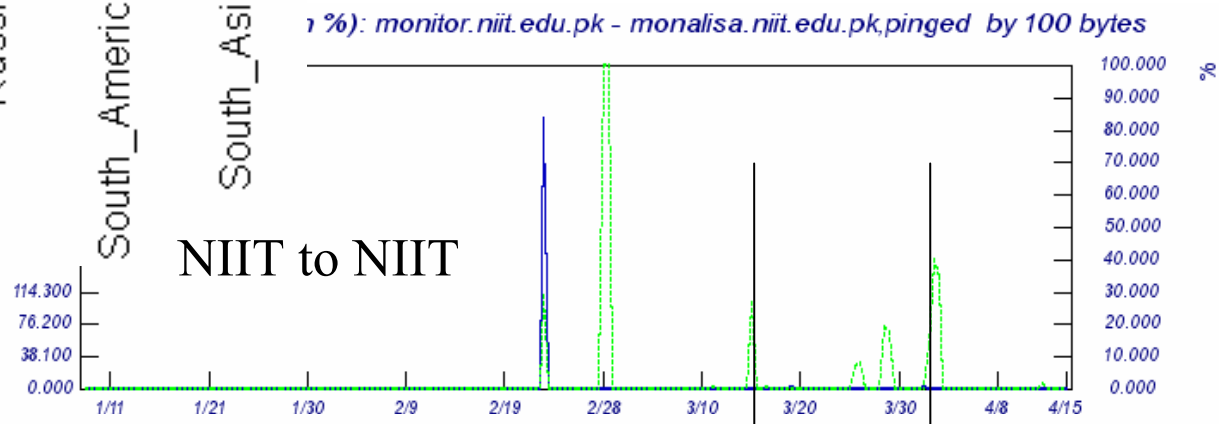

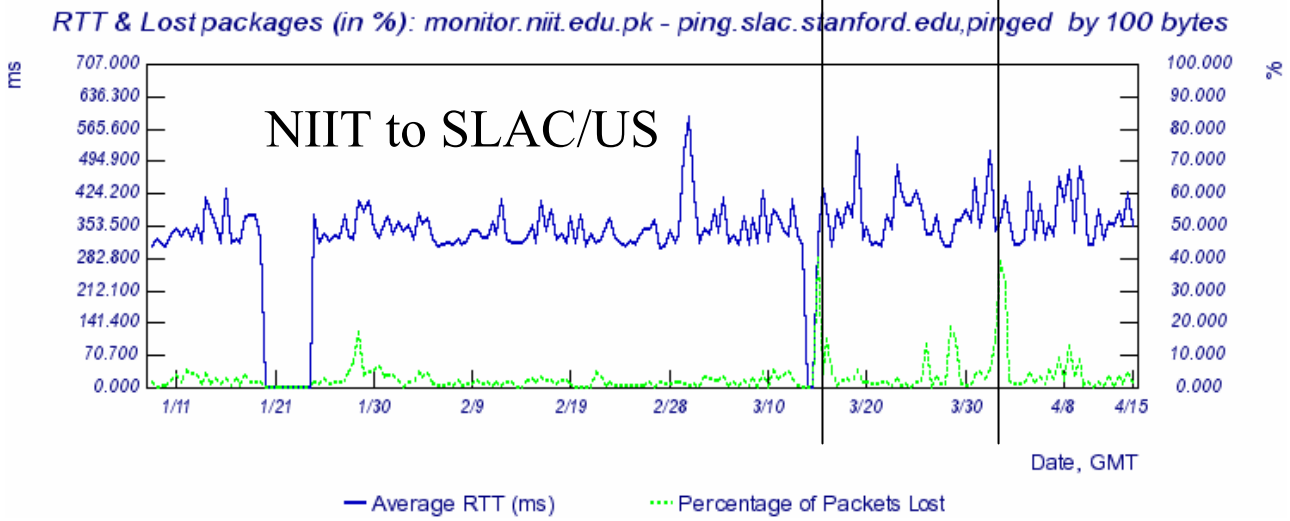

- 36 sites
- 26 in .pk
- **But monitor site problems**

From Russia, Brazil

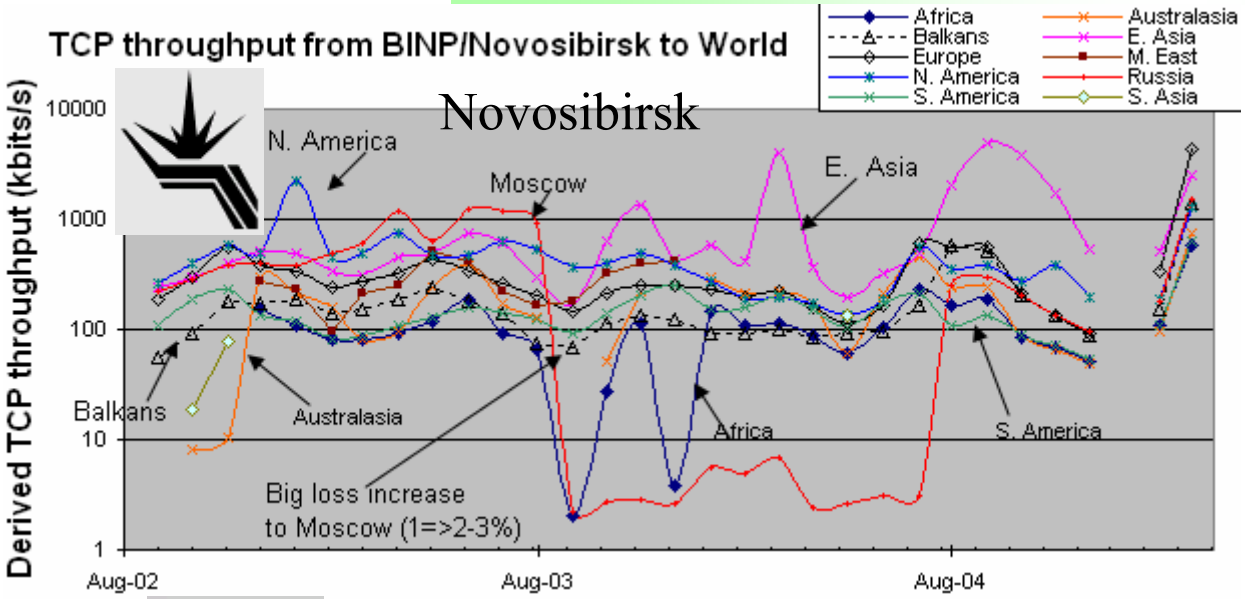

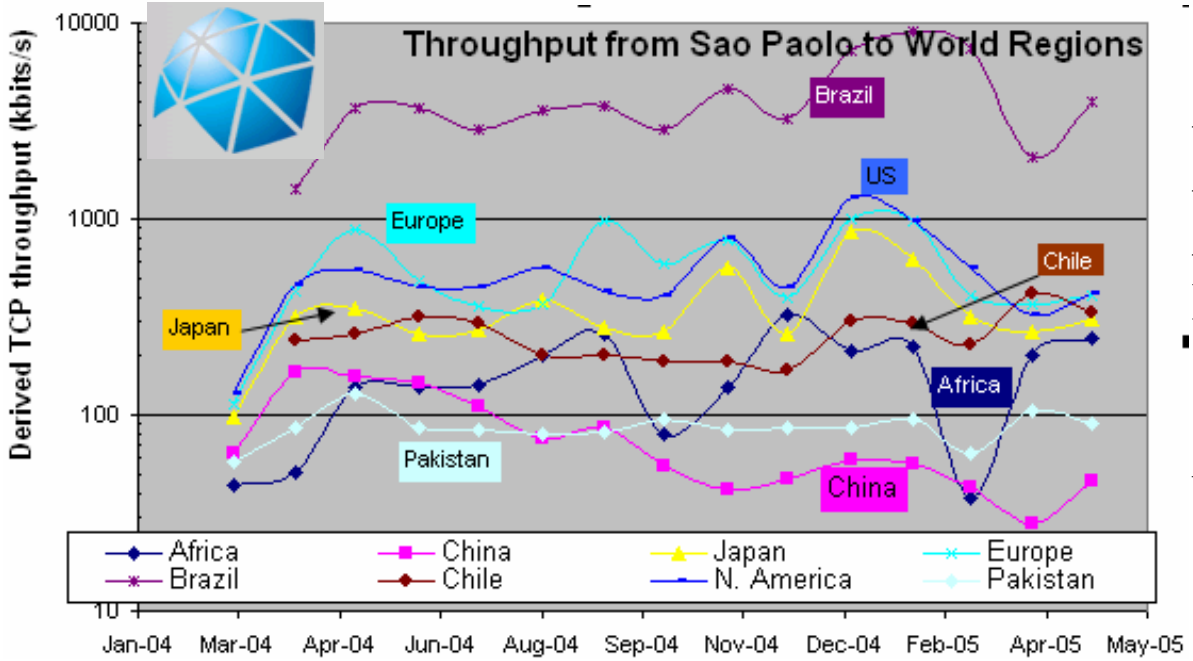

NSK to Moscow used to be OK but loss went up in Sep. 2003
Fixed in Aug 04
GLORIAD kicks in last couple months

As expected Brazil to L. America is good
Actually dominated by Brazil to Brazil
To Chile & Uruguay poor since goes via US (Miami)
US, Europe & Japan similar

Condition in Africa

- Working with Duncan Martin of TENET to get monitoring host in S. Africa
- *Internet connectivity in tertiary education institutions in Africa is in general too expensive, poorly managed and inadequate to meet even basic requirements. As the recent ATICS (Africa Tertiary Institutions Connectivity Survey) survey for the African Virtual University showed, the average African university has bandwidth capacity equivalent to a broadband residential connection available in Europe, pays 50 times more for their bandwidth than their educational counterparts in the rest of the world,* ***and fails to monitor, let alone manage,*** *the existing bandwidth (ATICS 2005).* ***As a result, what little bandwidth that is available becomes even less useful for research and education purposes.***

"Promoting *African Research and Education Networking", IDRC*

Losses to Regions

- Within regions (***bold-face italics***) losses are generally good (<1%)
  - Exceptions L. America, S. Asia
- Africa and S. Asia poor from US & Brazil (& Pakistan for S. Asia)

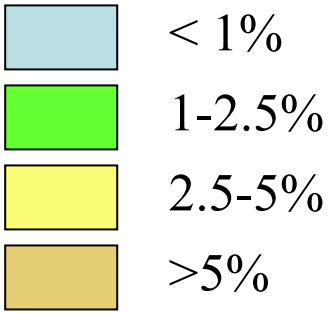

Mar '05	USA	Canada	UK	Denmark	Germany	Italy	Hungary	Russia	Japan	Brazil	India	Pak.	Median
North America	***0.51***	***1.24***	0.76	2.13	0.18	2.17	2.12	1.13	0.39	0.98	0.93	4.87	1.055
Europe	0.23	0.19	***0.24***	***0.16***	***0.18***	***0.25***	***0.15***	0.23	0.24	1.38	3.59	5.75	0.235
SE Europe	1.75	0.95	1.51	1.22	1.35	***0.97***	***0.87***	0.91	1.74		2.03		1.285
Baltics	0.11	0.04	0.17	***0.07***	0.14	0.16	0.04	***0.04***	0.14		1.85		0.125
Russia	0.6	0.52	0.63	0.37	0.51	0.55	0.44	***0.79***	0.59		1.28	2.03	0.59
East Asia	0.18	0.58	0.27	0.06	0.1	0.16	0.05	0.09	***0.59***	9.41	0.79	2.74	0.225
Oceania	0.59	1.22				0.59	0.63	0.37			0.8		0.61
Latin America	1.55	1.18	0.65	0.9	1.04	0.83	0.82	0.89	0.98	***1.63***	2.17		0.98
Middle East	2.51	3.26	1.91	2.03	2.01	2.62	2.24	2.16	2.68	0.62	2.73		2.24
South Asia	3.86									7.14		***4.6***	4.6
Africa	4.69									6.01			5.35
Median	0.6	0.95	0.64	0.635	0.345	0.59	0.63	0.79	0.59	1.63	1.85	4.6	0.6375

Compare with TAI

- UN Technology Achievement Index (TAI)

- TAI captures how well a country is creating and diffusing technology and building a human skills base.
- TAI from UNDP hdr.undp.org/reports/global/2001/en/pdf/techindex.pdf

TAI top 12	
Finland	0.744
US	0.733
Sweden	0.703
Japan	0.698
Korea Rep. of	0.666
Netherlands	0.630
UK	0.606
Canada	0.589
Australia	0.587
Singapore	0.585
Germany	0.583
Norway	0.579

Note how bad Africa is

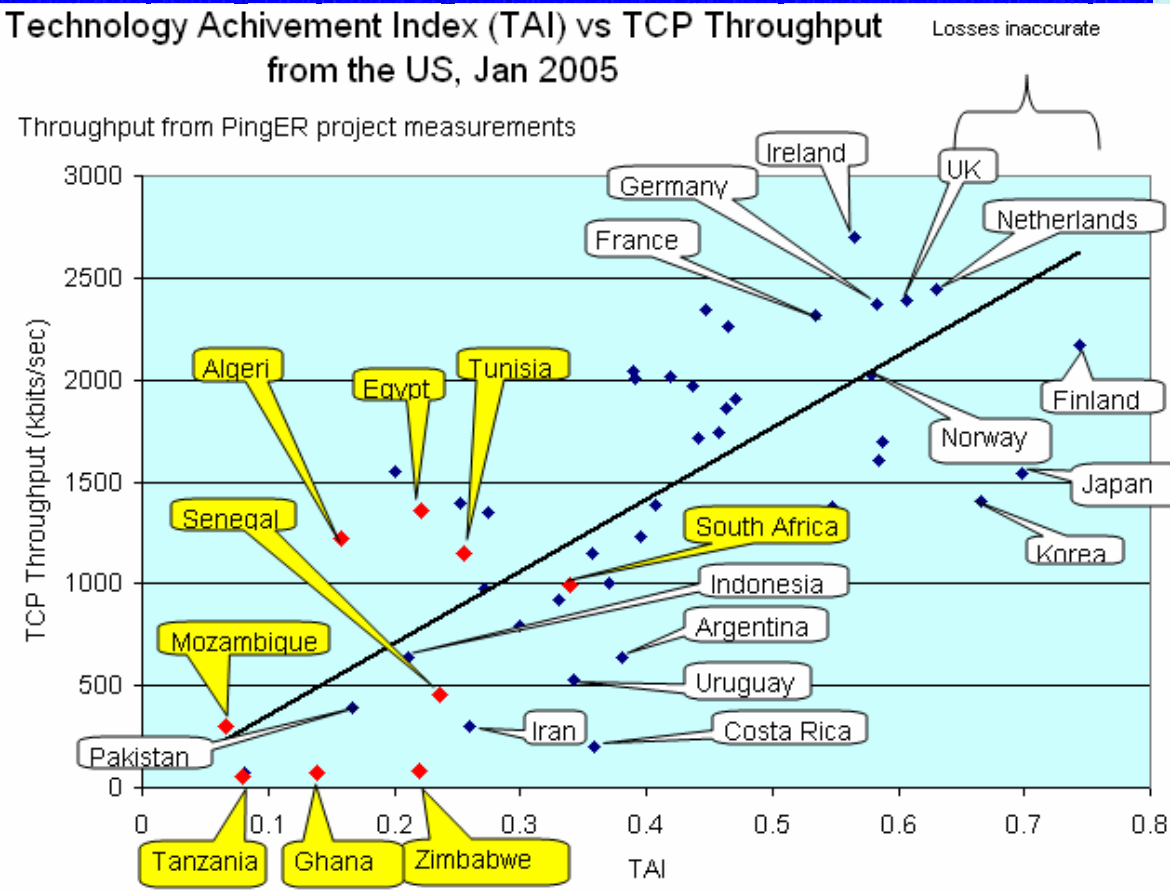

Collaborations/funding

- Good news:
  - Active collaboration with NIIT Pakistan to develop network monitoring including PingER (in particular management)
    - Travel funded by US State department & Pakistan MOST for 1 year
  - FNAL & SLAC continue support for PingER management and coordination
- Bad news (**currently unfunded, could disappear**):
  - DoE funding for PingER terminated
  - Proposal to EC 6th framework with ICTP, ICT Cambridge UK, CONAE Argentina, Usikov Inst Ukraine, STAC Vietnam VUB Belgium rejected, also proposal to IDRC/Canada February '04 rejected
  - Working with ICTP and NIIT on proposals
- Hard to get funding for operational needs (~0.3 FTE)
  - For quality data need constant vigilance (host disappear/move, security blocks pings, need to update remote host lists …), harder as more/remoter hosts

Summary

- Performance from U.S. & Europe is improving all over, for losses, RTT & throughput
- Performance to developed countries are orders of magnitude better than to developing countries
- Poorer regions 5-10 years behind
- Poorest regions Africa, Central & S. Asia
- Some regions are:
  - catching up (SE Europe, Russia),
  - keeping up (Latin America, Mid East, China),
  - falling further behind (e.g. India, Africa)
- Routing in developing regions may not be optimal
- Within a region can be big differences between sites/countries, due to service providers

Further Information

- PingER project home site
  - www-iepm.slac.stanford.edu/pinger/
- PingER methodology (presented at I2 Apr 22 '04)
  - www.slac.stanford.edu/grp/scs/net/talk03/i2-method-apr04.ppt
- ICFA/SCIC Network Monitoring report
  - www.slac.stanford.edu/xorg/icfa/icfa-net-paper-jan05/20050206-netmon.doc
- ICFA/SCIC home site
  - http://icfa-scic.web.cern.ch/ICFA-SCIC/
- SLAC/NIIT collaboration
  - http://maggie.niit.edu.pk/

Extra slides

Another view of Improvements

- Increase in fraction of good sites

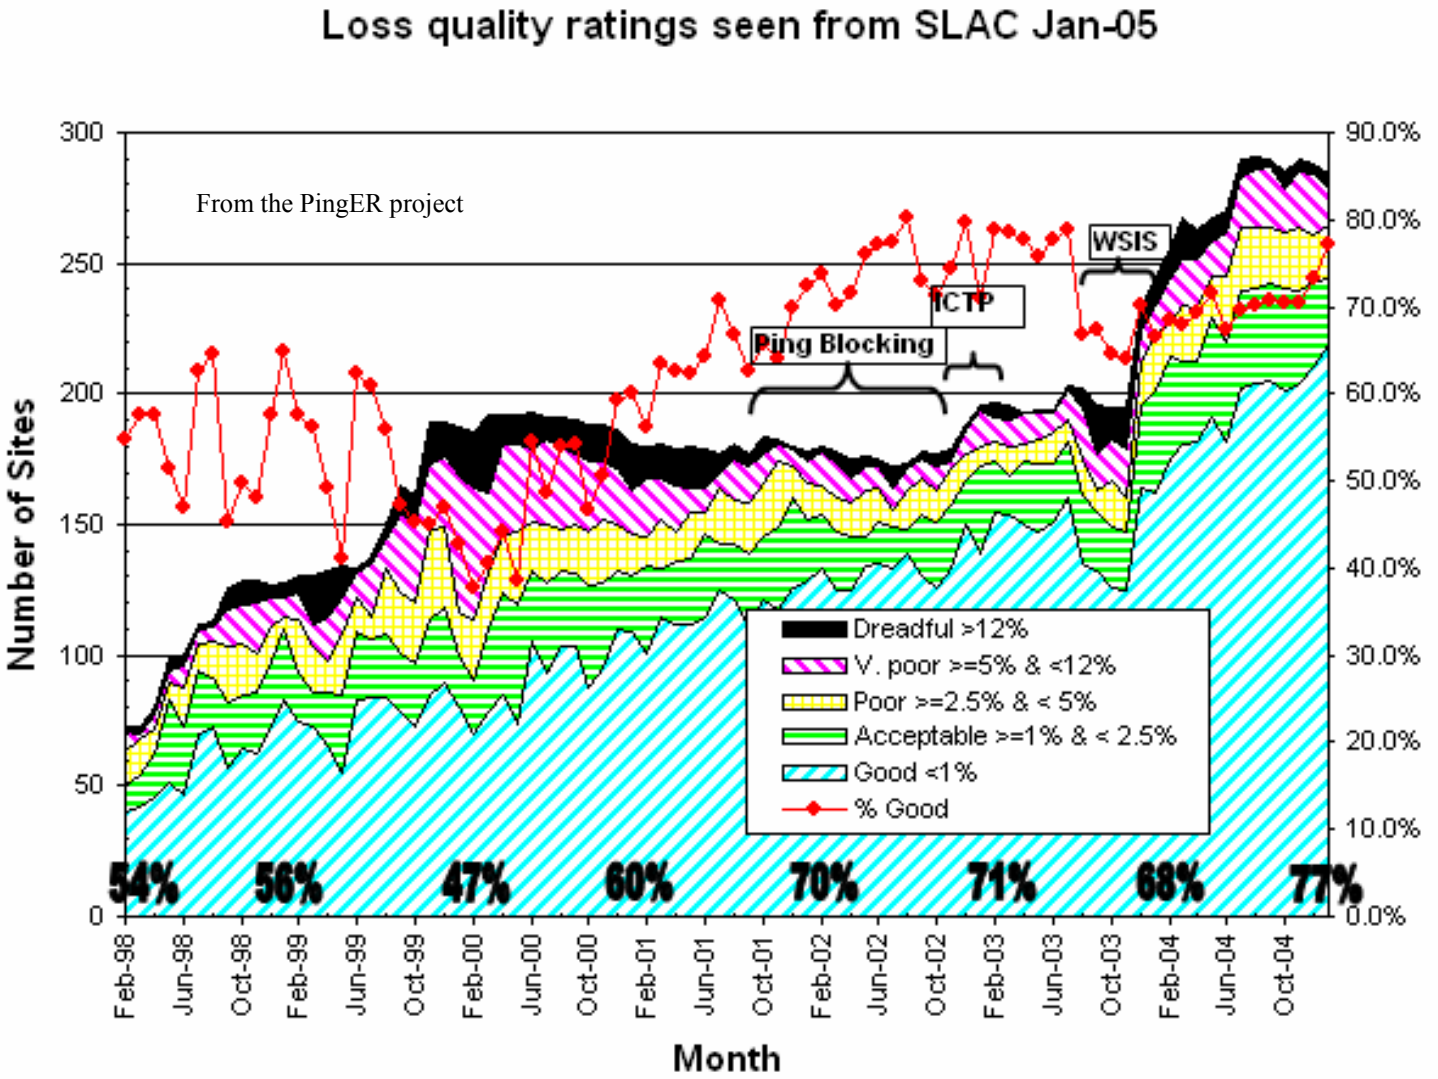

Countries covered

- Sites in 114 countries are monitored
- Goal to have 2 sites/country
  - Reduce anomalies
- Orange countries are in developing regions and have only one site
- Megenta no longer have a monitored site (pings blocked)

Albania	1	Canada	6	Germany	2	Lesotho	1	Papua New Guinea	2	Sweden	1
Algeria	2	Chile	2	Ghana	1	Lithuania	1	Peru	2	Switzerland	4
Angola	1	China	6	Greece	2	Macedonia	2	Philippines	1	Tajikistan	1
Argentina	6	Colombia	1	Guatemala	2	Madagascar	2	Poland	6	Tanzania	2
Armenia	4	Costa Rica	1	Hungary	2	Malawi	1	Portugal	3	Thailand	1
Australia	2	Croatia	3	Iceland	2	Malaysia	3	Reunion	1	Taiwan	1
Austria	1	Cuba	3	India	7	Mauritania	2	Romania	1	Tunisia	3
Azerbaijan	1	Czech Republic	0	Indonesia	5	Mexico	3	Russia	12	Turkey	1
Bangladesh	1	Denmark	1	Iran	4	Moldova	1	Saudi Arabia	4	Turkmenistan	1
Belarus	1	Ecuador	2	Ireland	2	Mongolia	2	Senegal	1	Uganda	2
Belgium	0	Egypt	1	Israel	0	Morocco	2	Seychelles	1	Ukraine	2
Bolivia	2	El Salvador	4	Italy	14	Mozambique	2	Singapore	1	United Kingdom	2
Botswana	2	Eritrea	2	Japan	8	Namibia	1	Slovak Republic	2	United States	80
Brazil	9	Estonia	1	Jordan	1	Nepal	1	Slovenia	1	Uruguay	2
Brunei	2	Finland	1	Kazakhstar	1	Netherlands	1	Solomon Islands	1	Uzbekistan	3
Bulgaria	1	France	7	Kenya	2	New Zealand	1	Somalia	1	Venezuela	3
Burkino Faso	1	French Polynesia	1	Korea, Rep	1	Niger	1	South Africa	4	Vietnam	0
				Kyrgyz Rep	1	Nigeria	1	Spain	1	Yugoslavia	2
						Norway	3	Sri Lanka	1	Zimbabwe	2

Africa RTT (satellite use)

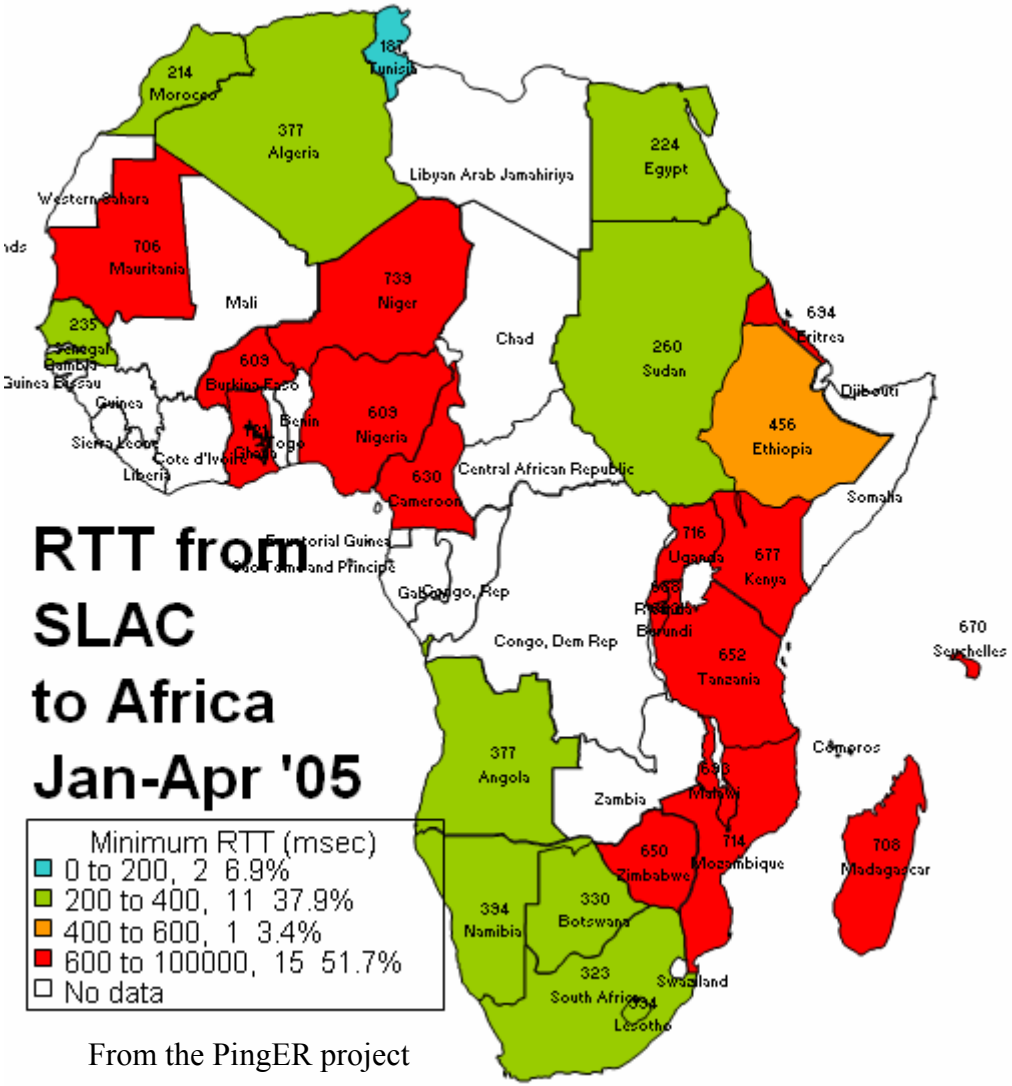

From the PingER project

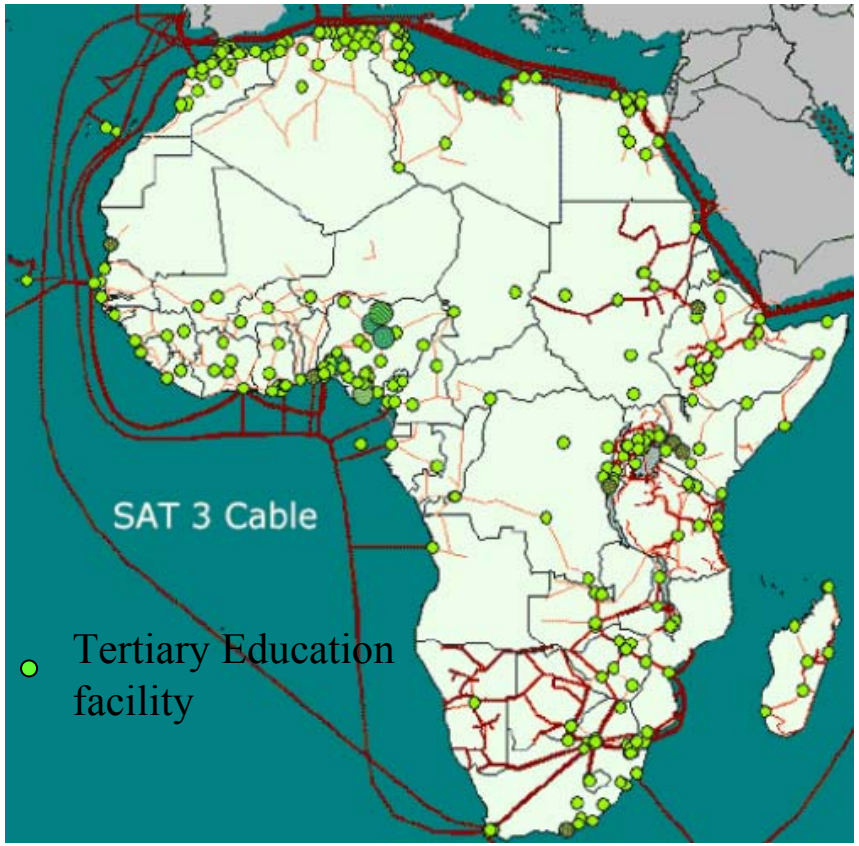

- We are working on ways to determine if a host is really in a country or a proxy host elsewhere

African Region Performance

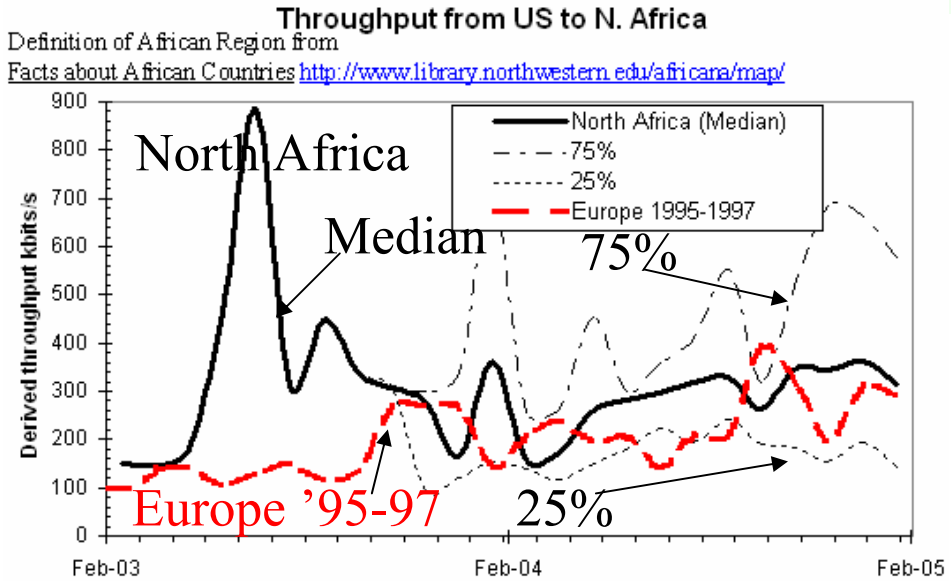

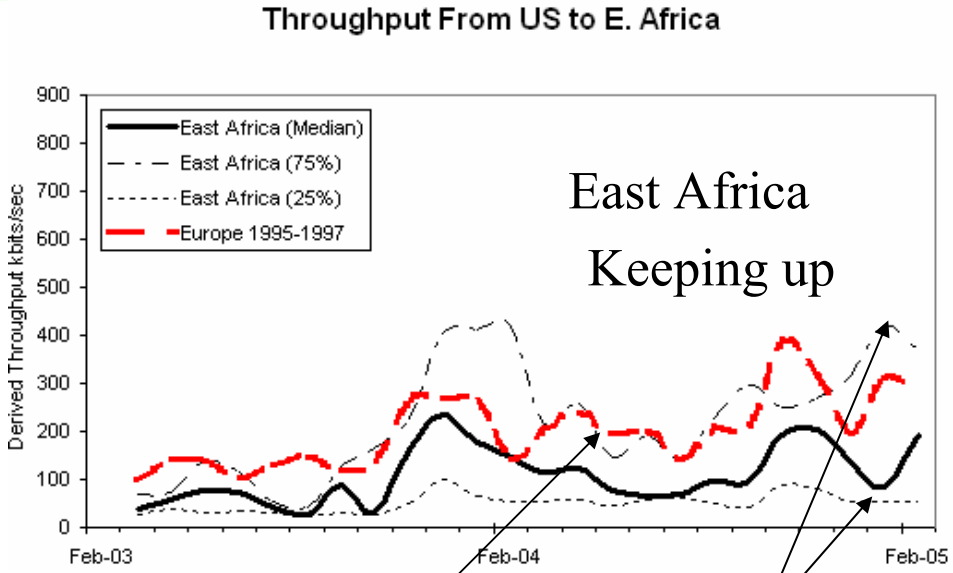

N. Africa has better connectivity; typically 8 years behind Europe, lot of variability

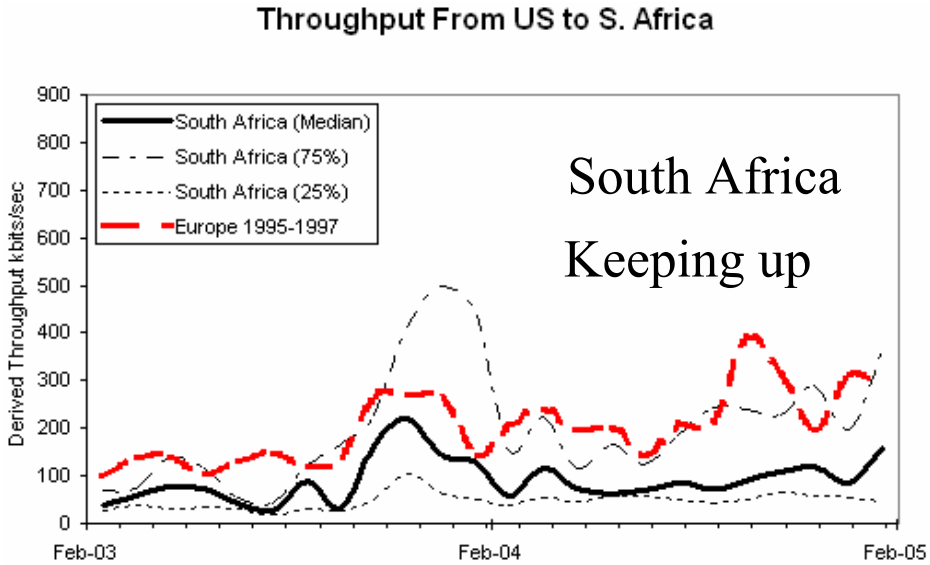

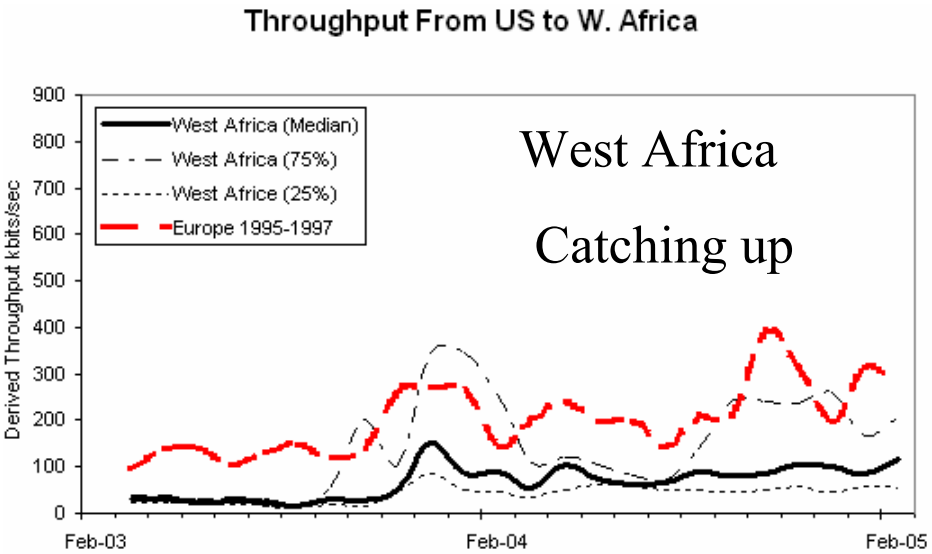

Within Developing Regions

- In ’80s many Eu countries connected via US
- Today often communications within developing regions to go via developed region, e.g.
  - Rio to Sao Paola goes directly within Brazil
  - **But** Rio to Buenos Aires goes via Florida

- Doubles international link traffic, increases delays, increases dependence on others
- Within a region can be big differences between sites/countries, due to service providers