

Can network coding bridge the digital divide in the Pacific?

Ulrich Speidel, 'Etuata Cocker, Péter Vingelmann, Janus Heide, Muriel Médard

ecoc005@aucklanduni.ac.nz

Project partners, collaborators and funders

- Aalborg University, Denmark
- Internet Niue
- Internet NZ
- Information Society Innovation Fund (through APNIC)
- Massachusetts Institute of Technology
- PICISOC
- Steinwurf ApS, Denmark
- Telecom Cook Islands
- Tuvalu Communication Corporation
- University of Auckland

Why?

- Almost all Pacific Islands (including NZ) have very long links to the rest of the Internet
- Many rely on geostationary satellites (GEO) with an especially long signal delay
- MEO Systems like O3b have signal travel times comparable to long cables but are a “thin pipe” compared to a cable
- Many satellite ground stations are unable to carry the amount of data that island users request and that arrives at the ground station from the Internet – the data gets “dropped”

Problem

- The TCP protocol, which prescribes how computers talk to each other, is abysmally bad at maintaining good data flow over long links, especially if the link contains a bottleneck where data is dropped (<1% of data loss is enough to make TCP slow to a crawl)
- Net effect: Satellite links are often slow long before they reach capacity – not an efficient use of the resource
- Better alternatives exist, but >99% of computers on the Internet don't use them – TCP dominates
- Bad news if you live in [*insert your island here*]



Tuvalu

Tokelau

Wallis and
Futuna

American
Samoa

Vanuatu

Port Vila

Nadi

Fiji

Niue

Tonga

New
Caledonia

Noumea

Cook Islands

Brisbane

Gold Coast

Newcastle

Sydney

Canberra


ACT

Auckland

Google

Why is TCP so bad?

- *Multiple* TCP senders remotely send traffic to the sat gate
- Sat link is a *bottleneck*. *Queue* at sat gate acts like a *funnel*.
- TCP sender cannot see queue state directly
- Feedback on queue state goes via the satellite to TCP receivers, and from there back to the senders
- Long delays: >500 ms on GEO, >125 ms on MEO



Rarotonga scenario end of 2014

- O3b MEO satellite with 160 Mbps downlink



Graphic courtesy
Telecom Cook Islands

- Burst packet losses and low link utilisation even at the lowest time scale – must have queue oscillation!


Niue scenario

- Connection via geostationary satellite from NZ, 8 Mbps

Graphic courtesy
Internet Niue

- Very high link utilisation – sustained use of around 7.4 to 7.6 Mbps during the day

- When sending data to Niue, we packet loss




- Almost all of the traffic on the link is goodput. Link is **congested**

Tuvalu scenario

January 2015

- Geostationary satellite link, exact downlink BW unknown but probably 16 Mbps
- Tuvalu Telecommunication Corporation was given a Silver Peak WAN Accelerator (SP) by Pactel, which seems to be doing all sorts of smart stuff
 - No detailed documentation of config available, so we don't know what exactly it does
 - Know it uses parity packets and packet caching
- Pactel measurements taken at a router island side of the Silver Peak suggests > 12 Mbps peak hour. Own measurement with tap between sat gate and SP indicate only 2-3 Mbps of this is actual traffic: low link utilisation. We also found packet loss to be high: **queue oscillation!**



Graphic courtesy
Tuvalu Telecommunication
Corporation


Solution: TCP over network coding (TCP/NC)

- Need a way in which we can let the sat gate drop data without causing mayhem in TCP
- Network coding converts IP packets into “combination packets” (“network encoding”)
- Instead of sending N data packets across the satellite link, we send M combination packets – essentially, we generate a system of linear equations whose solution is the set of original packets
- The decoder at the other end of the satellite link can recover the original N data packets from *any* N out of the M combination packets if they are linearly independent

Benefit

- Sat gate can drop up to $M-N$ of the packets without us actually losing any data
- Amount of data going across the sat gate is either the same as unencoded, or takes up what would have been spare capacity anyway
- Receiver almost never has to wait for missing data – TCP can communicate faster
- Technical effort (cost) involved is lower than the equivalent in extra satellite bandwidth or a cable
- No end user needs to upgrade their computers
- Larger end users can use network coded TCP for their networks without involving their ISP

TCP/NC tunnel setup



TCP/NC

- Is actually a bit of a misnomer – we don't code TCP packets as such but IP packets in general
- Also, what we do is tunnel IP packets (including ICMP, UDP, TCP, ...) between encoder and decoder as part of a UDP payload that is network coded – end users are completely agnostic to the coding
- Coded UDP packet payload consists of
 - Coding header identifying coefficients and sizes of the original packets that form part of the coded combination
 - Coding body containing byte-wise linear combination of constituent packets, padded if need be

TCP/NC Protocol stack

TCP, UDP, ICMP,...
IP
Network code
UDP
IP
Data link layer
Physical layer

TCP/NC protocol stack

TCP, UDP, ICMP,...				
IP				
Data link layer	Network code			Data link layer
	UDP			
	IP			
	Data link layer	Data link layer	Data link layer	
Physical layer	Physical layer	Physical layer	Physical layer	Physical layer

Host NC encoder/decoder Sat gate Sat gate NC encoder/decoder Host

Pilot study - questions

- Network coded TCP has run successfully in the lab, across the Atlantic in fibre cables, and end-to-end on satellites
- We're the first to use it in a tunnel arrangement
- We want to know whether it would work for Pacific Island ISPs
- We want to know how much improvement we can expect in practice
- We want to know how easily the technology scales for an island

Pilot study – implementation plan



- Deploy encoder/decoder at the University of Auckland (done)
- Deploy encoder/decoder machines at
 - Telecom Cook Islands, Rarotonga (done)
 - Internet Niue (done)
 - Tuvalu Communication Corporation, Funafuti Atoll (done)
 - Telecom Cook Islands (July 2015)
- Deploy encoder/decoder in California (machine being set up)



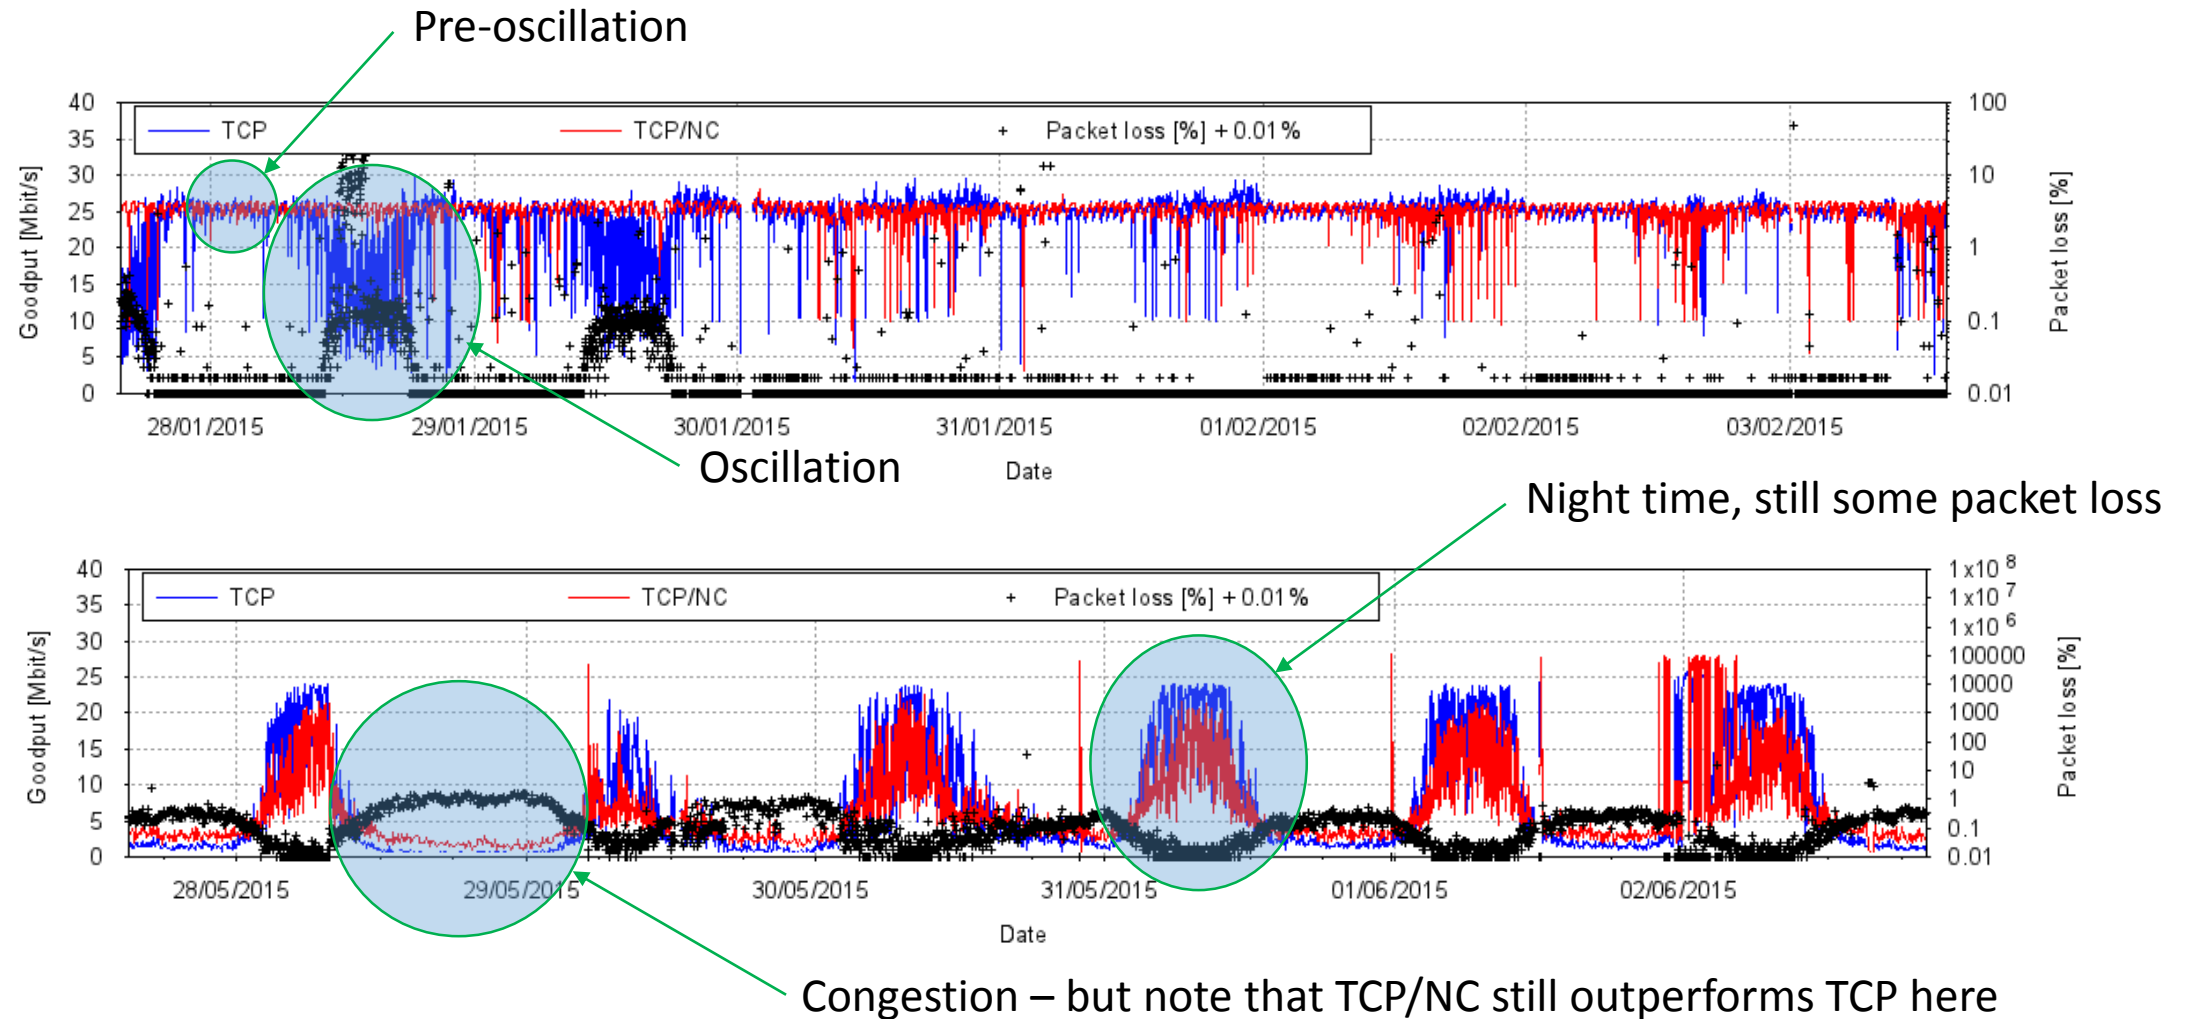
Rarotonga




Rarotonga then and now (N=60, M=90)



Rarotonga then and now (N=60, M=90)




Niue



Niue

- Severe congestion – but there are so many flows that packet loss per flow is manageable with $N=10$, $M=12$
- A single TCP connection running across a TCP/NC tunnel is able to achieve around 2-2.4 Mbps. This does not really improve with multiple parallel connections
- Without tunnel: Typically around 0.3 Mbps
- But: Lack of unused bandwidth means that this eats into the bandwidth of conventional TCP
- Open question: Given that the bulk of bandwidth use are flows that download something, would short fat flows with TCP/NC be better than long thin ones?


Niue (4 parallel connections, $N=10$, $M=12$)



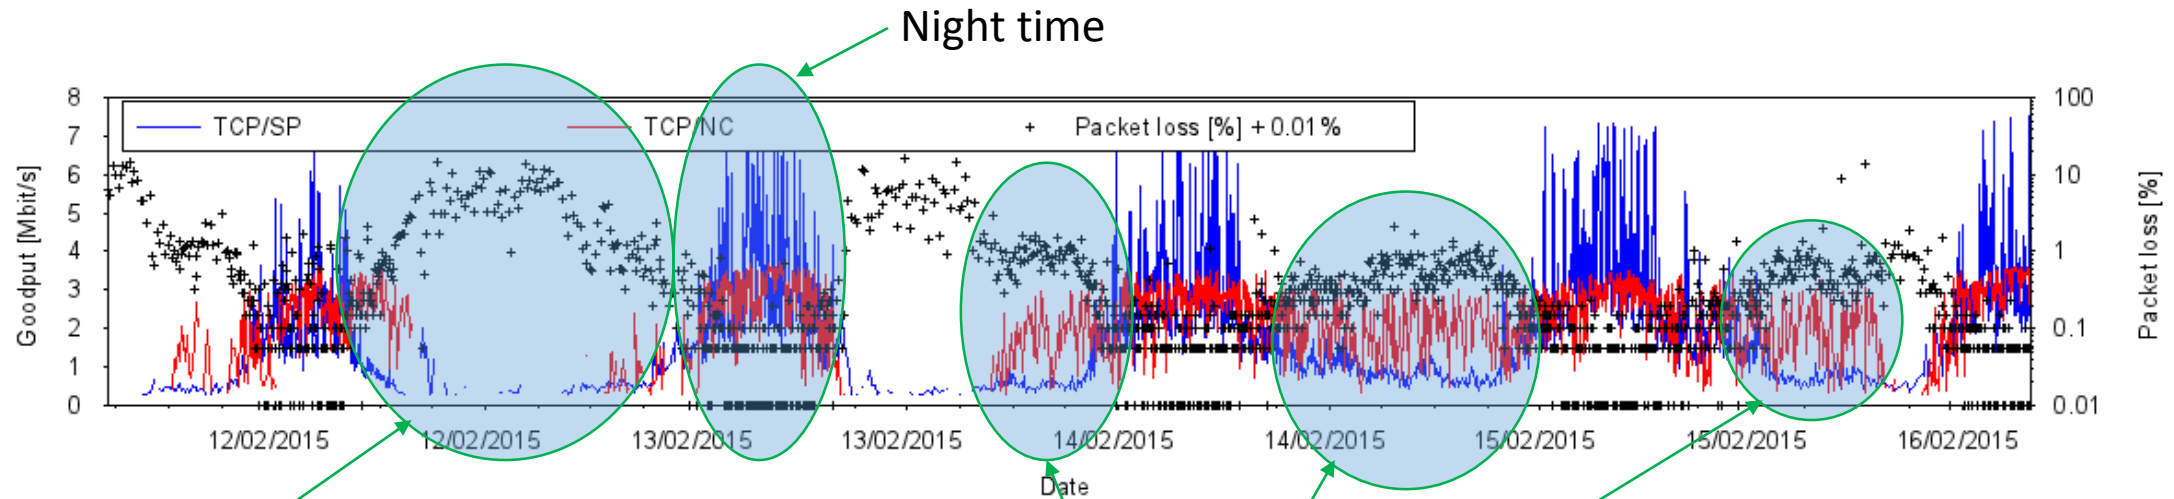
Tuvalu (Funafuti Atoll)



Tuvalu (N=30, M=45)



Tuvalu (N=30, M=45)



Packet loss bursts > M-N
(TCP/NC fails, conventional
TCP gets the odd packet through)

Periods with moderate packet loss
(daytime, TCP/NC has much higher goodput)

Tuvalu observations

- Installed ntop/nprobe with wiretap between Tuvalu sat gate and SP box / TCP/NC encoder-decoder
- We measure peak time traffic from conventional TCP/IP in the order of 2-3 Mbps – which seem to turn into > 12 Mbps after the SP box if Pactel are to be believed
- TCP/NC downloads bypassing the SP box drive traffic up to ~15 Mbps, including (substantial) overhead

Conclusions

- TCP/NC works well under low to moderate queue oscillation – a common daytime scenario in Pacific Islands
- Not so much a matter of “How many times faster?” but more of “How much of my bandwidth can I claw back?”
- No benefit in pre-oscillation conditions, although we may still see some here with lower overhead (M-N) and smaller N
- Can’t fix hopeless congestion such as in Niue fairly, but can get data through at the expense of other flows
- Large burst losses call for larger overhead and larger N, but is this actually worth doing?

Open questions and progress

- What would happen if *all* traffic to an island were encoded?
- Can we simulate satellite connections with and without TCP/NC?

Thank you!