isif asia

Project factsheet information

Project title	Improving Internet user experience in Pacific Island countries with network coded TCP
Grant recipient	Pacific Internet Society (PICISOC) Pokinui, Nikao, Rarotonga, Cook Islands PO Box 156, Avarua, Rarotonga, Cook Islands hilyard@oyster.net.ck
Dates covered by this report	1-04-2014 / 30–04–2015
Report submission date	30–04–2015
Country where project was implemented	New Zealand, Cook Islands, Tuvalu, Niue, United States
Project leaders name	Dr. Ulrich Speidel ulrich@cs.auckland.ac.nz 'Etuate Cocker ecoc005@aucklanduni.ac.nz
Team members (list)	Maureen Hilyard hilyard@oyster.net.ck Robert McFadzien robert@telecom.co.ck Atufenua Maui art2dee2@gmail.com Emani Fakaotimanava-Lui emani@niue.nu
Partner organizations	The University of Auckland Telecom Cook Islands Internet Niue Tuvalu Telecommunication Corporation Steinwurf ApS, Denmark
Total budget approved	AUD 30,000
Project summary	The project "Improving Internet user experience in Pacific Island countries with network coded TCP" aims to improve the utilisation of satellite links in and out of Pacific Island countries for whom fibre-optic connectivity is not - or not yet - an option that can be implemented economically. Lack of good Internet connectivity is a significant detriment to economic and social development in the islands. This technology had not been trialled as a tunnel over satellite links before. Our project deployed encoders / decoders for TCP/NC in the Cook Islands, Tuvalu, and Niue on behalf of PICISOC by a team based at the University of Auckland, with logistical assistance from a number of partners in the Pacific. The project also deployed an endpoint in Auckland and has secured a host for another in California. Our experiments show that TCP/NC can successfully mask packet losses to the TCP senders and thus prevent unnecessary slowdown, increasing link utilisation and goodput. The amount of improvement possible depends on the amount of link underutilisation as well as on the amount of residual conventional TCP/IP traffic.

Table of Contents

Project Summary

*Tips: It is recommended to **complete this section once you have finalized the text of the report**. It will be easier to go back through to build the summary based on the highlights of the report the project team just put together. The Project Summary can be up to **one page long**. It should include a brief justification; an outline of the project objectives to be achieved; the project real timeline and the main activities conducted. The abstract of the project written when ISIF Asia initially approved the project and the objectives listed in the Grants Agreement signed by APNIC and your organization should be useful inputs when preparing this section of the report.*

The project "Improving Internet user experience in Pacific Island countries with network coded TCP" aims to improve the utilisation of satellite links in and out of Pacific Island countries for whom fibre-optic connectivity is not - or not yet - an option that can be implemented economically. Lack of good Internet connectivity is a significant detriment to economic and social development in the islands.

Satellite links are generally bandwidth-constrained and have substantial additional latency (signal traveling time between endpoints). Conventional TCP/IP communication does not utilise such links particularly well: The long time that it takes to get data packets acknowledged by the opposite end means that response to congestion is delayed. This causes the packet queue at the satellite gateway to oscillate between "empty" and "overflow", leading to both unnecessary packet loss and significant underutilisation of the link itself - a significant loss in goodput.

In many cases, network coded TCP (now referred to as TCP/NC) can offer relief here. The TCP/NC technology we use in this project was developed at MIT in the United States and at Aalborg University in Denmark. It relies on the well-known concept of sets of linear equations, often taught in senior mathematics in secondary schools. Instead of original IP packets, TCP/NC sends packets with random linear combinations of a number of the original IP packets in the TCP stream. Each such combination represents one linear equation. The transmitter supplies a sufficient number of combinations to enable the receiver to "solve" the set and recover the original data packets, even if some of the combinations are lost enroute, cutting down very substantially on the need for retransmissions.

This technology had not been trialled as a tunnel over satellite links before. Our project deployed encoders / decoders for TCP/NC in the Cook Islands, Tuvalu, and Niue on behalf of PICISOC by a team based at the University of Auckland, with logistical assistance from a number of partners in the Pacific. The project also deployed an endpoint in Auckland and has secured a host for another in California.

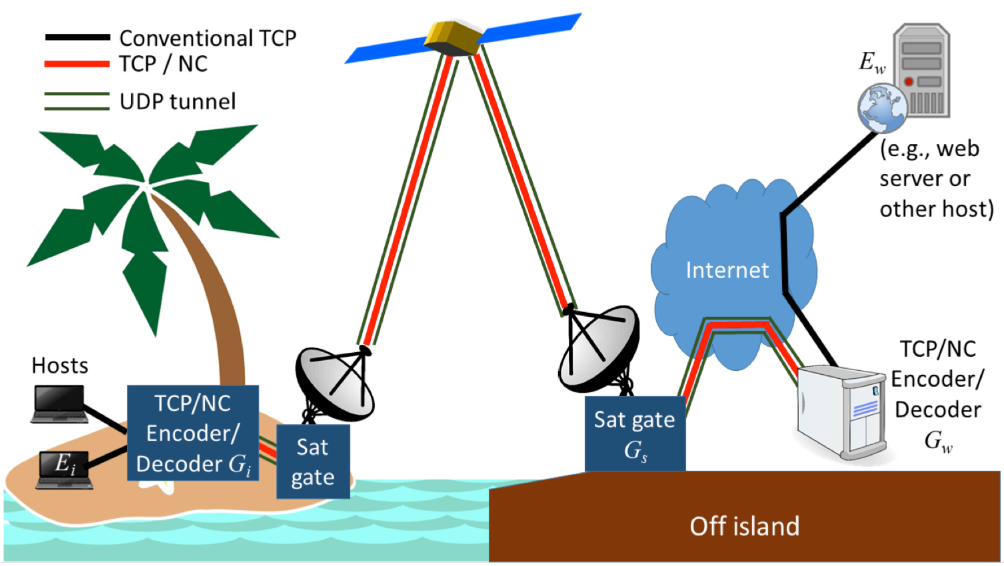

Our experiments show that TCP/NC can successfully mask packet losses to the TCP senders and thus prevent unnecessary slowdown, increasing link utilisation and goodput. The amount of improvement possible depends on the amount of link underutilisation as well as on the amount of residual conventional TCP/IP traffic.

Background and Justification

Tips: *The reader should be reminded of the* ***context*** *your organization is working, and where the project has been developed in. This section provides a window to* ***understand the challenges*** *faced by the community you are working with. Include a detailed description about the situation* ***before the project start****, describing any relevant aspects that make the project relevant in such a particular scenario. The reader should be provided with a clear description about the problem(s) to be addressed through this project and the motivation from your organization and team members to get involved and offer a solution.*

With a few notable exceptions, most Pacific islands cannot afford connecting to the Internet via the international high-speed submarine cable networks. Unless they find themselves in a sufficiently strategic location close to a major international cable corridor or have sufficient population and economic capacity to justify a dedicated connection, satellite connections are the only economically viable alternative.

Satellite connections come in two flavours: geostationary (GEO) and low/medium earth orbit (LEO/MEO). Geostationary satellites are the traditional solution, with links via fixed antenna ground stations to a satellite. The satellite sits 35,786 km above the equator and typically many degrees of longitude away from the ground stations, resulting in a signal propagation time of around 250 milliseconds. In comparison, the latency (delay) on the Southern Cross Cable Network - one of the main providers in the South Pacific - between Australia and the US is just over 62 ms [1]. The large footprint of geostationary satellites means that transmissions can be “heard” across a very large area, which prohibits reuse of the transmission frequency by another satellite in close geographical proximity. As a result, geostationary satellite bandwidth is severely constrained compared to that of fibre optic cables. Traffic backs up at gateways, leading to packet losses and latency variations (jitter).

The only LEO/MEO solution deployed in the Pacific at the start of the project connected Telecom Cook Islands on Rarotonga to O3b Networks’ network of MEO satellites [2]. Due to the lower orbital altitude, LEO/MEO solutions have smaller footprints and shorter signal paths, enabling them to offer significantly higher bandwidths and much lower (but still significant) latencies than GEO satellites. Service levels to run-of-the-mill low volume end users are roughly comparable to those offered by submarine cables. However, LEO/MEO satellites move constantly with respect to the Earth’s surface, and connection requires tracking antennas. As LEO/MEO satellites are not always above the horizon, a continuous ground-to-ground connection requires tracking of multiple satellites and handovers between satellites and ground stations. Ground stations are significantly more complex and set-up and maintenance cost are higher than in the GEO case. This puts the technology beyond the reach of some smaller island communities.

What does this mean for Internet traffic to and from Pacific Islands? No matter whether GEO or MEO satellites provide the connection, there is still substantial latency involved, and the bandwidth that the endpoints can afford is typically significantly below that of the Internet connection of the off-island satellite gateway. The link therefore represents a classic bandwidth bottleneck.

This creates a problem for the mainstay Internet transport protocol, TCP [3,4,5,6,7]. TCP adjusts its packet transmission rate according to feedback from the receiver. If packets are lost due to congestion, it reduces the sending rate. If acknowledgements indicate that there is surplus bandwidth, it increases it. In doing so, TCP effectively attempts to adjust its sending rate to the current conditions on the path between sender and receiver, with the aim of achieving a constant sustainable rate. The transmitter’s view of the conditions becomes progressively out-dated with increasing latency, however. In the case of a GEO satellite, feedback is at least half a second out of date, in the case of a MEO satellite, this delay is still well over 100 milliseconds.

Add multiple TCP senders, and the mix of bottleneck and latency brews up a perfect storm at the input queue of the satellite gateway pointing at the island. It sits in front of the bottleneck, receives the bulk of the traffic and buffers the transition from gigabit/s networks to a link capable of handling only between a few and a few hundred megabits per second. A queue in this spot becomes a prime candidate for a cyclic effect known as queue oscillation.

If the TCP senders see capacity (or rather, that there was capacity when their last packets passed through), they increase their sending rate. As this is (almost) concerted action, the input queue fills quickly and eventually overflows. Packets that arrive at an overflowing queue are dropped. Since the only way for the senders to learn of the problem is from feedback via the island, further packets continue to head for oblivion in large numbers for quite a while, resulting in burst packet losses (packet flows losing a large number of consecutive packets at a time). This is the first part of the cycle.

The other part of the cycle begins once the senders find out that there has been a shocking amount of packet loss in their connections: Sensing congestion, they all slow right down. The effect on the queue is that it now loses most of its data feed, eventually runs empty and leaves the link idle.

Pacific Islands thus have not only some of the longest latency Internet connections in the world with pretty significant jitter, but are also most prone to lose expensive link capacity to queue oscillation. Many Pacific Islands have never experienced a good quality Skype call, let alone with video. Internet activities we take for granted (YouTube, Netflix, Internet Radio) simply don’t work or are prohibitively expensive to run, and even web pages can takes ages to load. Low good put also aggravates the ability to update systems, leading in turn to more malware exposure, which further increases the load on the links.

The result: With a few notable exceptions, Pacific Island communities suffer from some of the worst Internet connectivity in the world, with no affordable alternative. This impacts all aspects of life in the islands: the economy (e.g., tourism and trade), health, education, and social wellbeing. There are also large Pacific Island communities in New Zealand, Australia, the U.S. and Japan, with significant demand for low-cost communication to maintain family relationships.

This was the starting point in 2011 of ‘Etuate Cocker’s PhD project in the Department of Computer Science at the University of Auckland under the supervision of Ulrich Speidel.

Their team at the University of Auckland built and operated a network of measurement stations in numerous Pacific Island Countries (in the Cook Islands, Fiji, Kiribati, Niue, the Solomon Islands, Tonga, Tuvalu and New Zealand) as well as a number of other countries around the globe (Canada, Germany, Japan, Macau, Malaysia, South Africa, Switzerland, and USA) which documented some the connectivity effects from a technical perspective since mid-2011. Another second aim beyond capturing the as-is status was to look at ways in which connectivity could be improved using smart solutions.

Since there is no practical alternative to satellite links, the challenge is to make the best use of the available resource. It was reasonably clear that TCP’s standard behaviour is not well suited to the situation at hand - so what could be done about it?

It became obvious to us that a significant part of the problem was the uneven packet flow caused by TCP’s inaccurate estimation of the current channel conditions, causing the flow to slow down and retransmit large amounts of data unnecessarily while simultaneously underusing the satellite links. If one could prevent packet losses and delays across the satellite link from resulting in TCP timeouts, slowdown and retransmissions, we might be able to achieve a more even flow and higher overall goodput.

In 2010, Ulrich Speidel attended a keynote talk by Professor Muriel Médard from the Massachusetts Institute of Technology (MIT) about network-coded TCP (now, and in this report, referred to as TCP/NC). TCP/NC does not transmit packets in their original formats but rather forms multiple combinations of multiple packets [8]. It then recovers the originals from these combinations at the receiver. As combination packets are interchangeable for the purpose of decoding, a lost combination packet does not require retransmission - the subsequent combination packet may be used in its place. Generally, this packet will already be on its way, so latency ceases to be an issue.

In 2013, the team was able to bring Prof. Médard to Auckland as a Vice Chancellor's Lecturer. Discussions with her in the run-up to the visit convinced us that TCP/NC could assist us in improving throughput between the Pacific and the rest of the world.

The current project had its genesis in discussions between 'Etuate Cocker and PICISOC as well as with a number of potential deployment sites in the Pacific Islands.

Project objectives

***Tips:** Please include here the **original objectives** as listed on the Grant Agreement. If any objectives were modified, added or removed during the reported period this should be explained/justified.*

The recipient's objectives of the project are:

- Commissioning of research services from the University of Auckland in the form of research support to Mr. 'Etuate Cocker and Dr. U. Speidel.
- Completion of a visit to MIT Cambridge by 'E. Cocker to familiarise himself with TCP/NC.
- Implementation of network coding / decoding servers in New Zealand, California, and at least two of Niue, Tuvalu, and the Cook Islands as well as a fourth Pacific Island recipient (candidates included Kiribati and the Solomon Islands as well as a further Cook Islands site).
- Performance testing between these servers
- Publication of the results

Users and uses

***Tips:** Discuss with your project team who would be the future users and how they would use the findings throughout the project lifecycle. The uses identified should relate to the theory of change that you have discussed with your project team. The discussion about theory of change, users and uses, will be a very important input to your communication strategy: depending on who the user is and of what use will be the findings, a communication strategy can be developed. For example, if the users of the findings are policy makers and the use is to influence a change in the regulatory framework, which communication approach will work the best?*

Who will be the user of these findings?

What are the more relevant things the project team wants to learn about or evaluate through the lifecycle of this project?

The main users of our findings will be in the following categories:

- Internet Service Providers in the Pacific and beyond who may wish to protect the capacity of a valuable resource (their satellite link) by tunnelling traffic across it to and from remote gateways via TCP/NC.
- In places where Internet Service Providers will not offer this, organisations such as businesses, schools, government departments, hospitals or universities may wish to make better use of their own connectivity by tunnelling their traffic across the satellite link using TCP/NC (encoders and decoders need not be located at ISPs).
- Technical personnel working and consulting for the above.
- Researchers working on TCP/NC.

The initial phase of the project concerned setup, testing commissioning and configuration, including the fine-tuning of parameters. We wanted to find out how difficult this was and what it took to implement a TCP/NC tunnel. How much performance gain could we expect as a baseline, and how much gain could we get from fine-tuning? Complexity of implementation is part of the cost of any technology, and the project insights in this respect will inform the preparation of local business cases.

In the operative phase, we wanted to learn a number of things. We expect TCP/NC to co-exist compatibly with conventional TCP - could we demonstrate that increased good put with TCP/NC is not outweighed by loss of capacity for conventional TCP, and that there would be benefits in tunnelling all traffic across a satellite connection via TCP/NC? What obstacles would there be to wholesale adoption in a particular location?

How scalable is the technology? The physical limit to the satellite bandwidth still applies. The best we can hope for is that we will be able to use this existing bandwidth much more efficiently with network coded TCP/NC. How much computing capacity do we need for encoders and decoders in order to carry a certain amount of traffic in this context? This is a cost factor, and once again the project insights in this respect will inform the preparation of local business cases.

How robust is the existing technology? Is the existing TCP/NC software sufficiently stable for long-term unattended operation? How stable are our results? How do changes in satellite connectivity (e.g., the addition of extra capacity) and usage impact on our initial findings? Better throughput tends to result in higher usage - can TCP/NC deliver the critical end user benefit without hitting the physical limit to the satellite bandwidth? Could the technology help to create critical mass for increased uptake, and hence result in business cases for upgraded satellite connectivity because it promises users a tangible return on increased investment?

Indicators

***Tips:** Indicators help to **measure project's progress**. Indicators help the objectives that were set by the project team to be affordable, tangible, and measurable. They help to verify the success and rewrite the course in case we are not achieving it. An indicator could be quantitative (percentage, amount) or qualitative (perception, opinion). The ISIF Asia secretariat suggests the SMART approach to indicators:*

S** **Specific
M** **Measurable
A** **Achievable (acceptable, applicable, appropriate, attainable or agreed upon)
R** **Relevant (reliable, realistic)
T** **Time-bound

Baseline	Indicators	Progress	Assessment	Course of action
Research agreement: Baseline: Grant awarded, no formal research agreement between PICISOC and University of Auckland place	Research agreement in place? Reports delivered according to agreement?	Agreement is now in place. First report to ISIF was completed. Informal monthly reports to PICISOC have been dispatched.	Due to the unusual triangular nature, it took some time to get the agreement negotiated, which delayed the start of the project at UofA. This delay required the rescheduling of some activities. Core activities are now complete.	Continue to publish
Required equipment: Baseline: required equipment and procurement needs identified (Note: not all equipment here was procured from the ISIF grant)	Island-end encoder/decoder machines acquired? NZ encoder/decoder server in place? Leased server in California in place? Equipment configured with proper accessories and operating system? Ancillary equipment (UPS, time servers, beacons, small parts) acquired & ready for deployment?	The island end machines are in hand and have been configured with Ubuntu. This includes an "engineering machine" that allows us to test on an identical platform. Island end machines have been deployed. NZ encoder/decoder server is in hand and operational. Leased server is now being provided by the Cooperative Association for Internet Data Analysis at the University of California San Diego (USCD) in conjunction with the San Diego Supercomputer Centre. Ancillary equipment has been procured and is in place.	Procurement of the island end hardware took some time - the industrial spec fanless mini PCs are specialised equipment not carried by UofA preferred vendors and vendor setup was complex. This caused more delay than we had anticipated. The island end machines came with US style power cable and had to be replaced with AUS/NZ style cables. UPS equipment was deployed in Niue, Rarotonga and Aitutaki. Deployment in Tuvalu was not possible after all due to Air Fiji's refusal to carry the unit from Auckland. Carriage to Niue was in doubt until departure day due to severe weight restrictions on the flight.	NFA
Software: Baseline: required software identified but not procured or configured	Software license obtained? Software installed, configured, and tested? Configuration fine-tuned after deployment?	We now have access to the software and in fact the source code under a research license. The software has been installed and is operational. The vendor (Steinwurf) provided three customised versions of the software.	Procurement here was similarly cumbersome as for the island end hardware. MIT and their Danish partners at the University of Aalborg have vested the IP in a joint venture startup, Steinwurf ApS. Steinwurf have been very cooperative. They have given us access to their source code, and we are receiving excellent support from them (there is also significant interest in our results from the research groups behind the company). This included two	NFA

			versions of customised software, access to their own in-house testing tools, and last but not least extensive after-hours assistance by their main software developer, including in the running of measurements (they are as interested in the data as we were).	
Training visit to MIT: Baseline: Visit agreed upon in principle	Visit dates set? Visit approved at MIT? Documentation for visit ready? Visit taken place?	Visit took place from November 17 to December 8.	Our original plan at the time of application was to have 'Etuate visit MIT during the (northern) summer break and take one of the "island machines" with him for configuration and testing purposes. However, this had to be abandoned as contract negotiations were taking longer than expected. The next tentative date in September/October also had to be abandoned when we could not be assured that we would have the island end machines in hand at the time. The visit eventually took place during November/December and overlapped with the Rarotonga deployment. MIT kindly waived the hosting fee as the visit was under 5 weeks - this is gratefully acknowledged. The visit was extremely informative and knowledge gained was invaluable during the deployments especially in Niue and in Tuvalu.	NFA

Baseline	Indicators	Progress	Assessment	Course of action
Deployment: Baseline: Deployment sites identified (Cook Islands, Niue, Tuvalu, with Kiribati as backup)	MIT training visit taken place? Equipment ready for deployment? Deployment visit dates set? Visits prepared (tickets booked, liaison with on-site contacts, permissions in place, etc.)? Visits taken place by end of February 2015?	All indicators completed. We were also able to deploy in a fourth location (Aitutaki) in July 2015 with leftover funding.	This stage had to be pushed back as a knock-on effect of the delays in getting the project underway. One problem that this raised was that 'Etuate Cocker's PhD scholarship ran out at the end of 2014. We had anticipated being able to complete the project before then, but given the delays in getting underway we needed to find a way of supporting him to the equivalent of his scholarship for approximately another two months so he could complete deployment and evaluation. ISIF's support in letting us use project savings towards this are gratefully acknowledged. Deployments in Rarotonga and Aitutaki were carried out by U. Speidel.	NFA

Project implementation: understanding the chain that leads to results

Tips: *This is the most important section of the report. Here, the reader will* ***understand the processes and operational issues*** *of your project and how they contribute to the achievement of the objectives and the theory of change behind the project implementation. Is possible that the project team's understanding of the development problems to be addressed with this project will have evolved or* ***changed*** *from those described when the project was originally submitted and approved. If that is the case, please share what motivated the change and what course of action has the project team identified.*

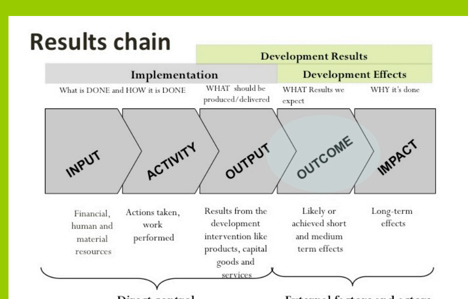

Results chain diagram provided by In Develop

Narrative - project implementation

The project aims to find out whether TCP/NC could allow satellite connected Pacific Island nations to use their satellite links more efficiently. To this end, the partnership between PICISOC and the team at the University of Auckland has built a number of links to other organisations. Some of these, in particular the links to the organisations in the islands that will physically host the equipment, were pre-existing, some have developed as a result of the lead-up to the project and the procurement of the equipment and software. The narrative below follows roughly chronological order.

The partnership between PICISOC and the team at the University of Auckland has strengthened considerably, with much closer and more multilateral contact than beforehand. Getting the agreement into place in consultation with APNIC/ISIF took longer than expected but will hopefully be able to serve all parties as a template for similar projects in future. Most PICISOC members as well as ‘Etuate Cocker are either based in the Pacific Islands or have close connections there, and so there has been close beneficiary involvement since the inception phase. We have been able to widen this circle during the deployment phase with new contacts made.

After much electronic paperwork, the supplier of the industrial grade computers for island deployment, Stealth Computers of Canada, have now become a preferred supplier at the University of Auckland as a result of our project, as have Steinwurf ApS, the supplier of the software which implements the intellectual property contributed by the team at MIT and their joint venture partners at Aalborg University in Denmark. Procurement took longer than expected as both equipment and software were non-standard items and the procurement phase overlapped with a major administrative restructuring process at the university.

We now also have a very strong relationship with the Massachusetts Institute of Technology (MIT), one of the world’s top universities, which ‘Etuate Cocker visited from mid-November 2014 for three weeks. MIT were very generous in waiving the visiting student fees for this visit which they had initially told us would be charged.
We have also built a strong relationship with Internet NZ, who were funding much of this visit and other aspects of the project. Their co-funding helped considerably in enabling us to deploy in four island locations rather than three as originally envisaged.

Our relationship with the ISPs at our deployment sites (Telecom Cook Islands, Internet Niue and Tuvalu Telecommunications Corporation (TTC)) has been strengthened. All bent over backwards to assist us in the project and generously helped with cabling and data as well as a generous amount of staff time.

Among the data we asked for were satellite link utilisation statistics. These showed that Rarotonga had typically about 60 Mbps of available bandwidth in November 2014 that was almost never used. After their total bandwidth was upgraded in 2015, the amount of unused bandwidth rose to over half of what was provisioned. Tuvalu similarly seemed to used only less than a fifth of the bandwidth provisioned. Only Niue seemed to be able to squeeze >95% utilisation out of its 8 Mbps inbound international bandwidth. Aitutaki shares its link with a number of other islands but with a guaranteed minimum rate, such that link utilisation is a more complex concept. At deployment time Aitutaki pulled a relatively steady 12-15 Mbps of data - but see further comments below.

Tuvalu (Funafuti Atoll) proved to be peculiar case that generated a significant amount of discussion. TTC’s satellite system is managed by Pactel (now Speedcast), and includes a SilverPeak WAN optimiser that uses functions such as network memory and parity packets. As a result, some of the data emitted by the optimiser to the island does not originate directly from the link but consists of locally generated / decompressed / reproduced data, which makes direct comparison with TCP/NC difficult. The utilisation figures for Tuvalu above are our own measurement, the correctness of which relies on us having been provided with the right observation point within the TTC network. To the best of our knowledge, the traffic was observed at a point between the satellite receiver and TTC’s network that carried all IP traffic and was no influenced by the SilverPeak device.

Link utilisation below 100% indicates that there are times when the off-island satellite gateway has no data to transmit, i.e., the satellite gateway’s packet queue is empty. At the same time, all links considered showed that offering load would result in packet loss, which when load-dependent indicates an overflowing queue. Such packet loss is not normally observed on fibre-based international links. These effects are the main characteristic of full queue oscillation.

In the case of Rarotonga and Tuvalu, we ran multi-day experiments which consisted of a 20 MB (Rarotonga) and 5 MB (Funafuti) download of freshly generated random data, repeated every few minutes for both conventional TCP and TCP/NC.

The Rarotonga time series were taken at the end of January/early February 2015 and end of May/early June 2015 and both show particularly bad conditions for standard TCP over the last days of the month, with goodput during peak time falling to as little as 20% of its off-peak value. In January, TCP/NC was able to recoup the lost goodput in its entirety. In May, the drop was even more dramatic, with some TCP connections between NZ and Rarotonga timing out. The burst losses on the latter time series were sufficiently long to also affect TCP/NC, which still gave around twice the goodput during these times, however, and did not suffer from connection timeouts.

These effects were not as pronounced on the first days of the respective subsequent month. We now know that this is most likely due to users with a monthly data cap using spare credit at the end of the month. A repeat series of experiments in mid-July showed similar effects to the May/June month-end pattern.

In Tuvalu, we found a similar scenario as in Rarotonga at the end of May, except that the burst packet losses here were so bad that they even wiped out the overhead configured for TCP/NC. At time when they did not but standard TCP was affected by queue oscillation, TCP/NC gave around four times the goodput of standard TCP. Note that the standard TCP performance was taken across the SilverPeak device, whereas the TCP/NC setup sits alongside and does not use the SilverPeak at all.

In Niue, the high existing link utilisation prohibited such long time series experiments - they would have seriously impacted on the link. Our first experiments confirmed that the link was indeed carrying almost exclusively goodput and did not see retransmissions. That said - this sounds better than it is in practice: What happens in such scenarios is that the link does not oscillate, but the number of individual TCP connections is so high that any TCP connections that are unable to finish their data transfer within a few packets slow down to a crawl and, in many cases, time out and cannot complete.

To protect the scarce link resource, we first conducted a smaller number of individual experiments in Niue rather than a series. They showed that TCP/NC could deliver around 8 times the goodput of standard TCP for 5 MB downloads on this link, albeit at the expense of other TCP connections sharing the narrow bandwidth. A short series of experiments with four parallel TCP/NC connections showed that this goodput rate did not increase significantly and was in fact shared almost equally between the four parallel connections.

In the Aitutaki deployment, we first tested the available link bandwidth and found this to be higher (up to 70 Mbps) than the shared 40 Mbps advised by Telecom Cook Islands. As mentioned above, only 12-15 Mbps of these are used, with the other islands sharing the bandwidth drawing typically another 3-4 Mbps. Conventional TCP was able to access up to almost 60 Mbps during 20 MB downloads from New Zealand, with no significant packet losses. That is, the queue was clearly not oscillating here - the low utilisation reflects low demand. TCP/NC currently yields no benefit in the context of this link.

This may yet change: Based on the above figures, both Rarotonga and Aitutaki have around 30 kbps of bandwidth per head of population including tourists. However, Rarotonga has a well-established 3G access network, whereas Aitutaki was still on 2G/GPRS at the time of writing and only WiFi hotspots and fixed network connections offer higher speeds. As a result, the number of concurrent users on Aitutaki is only about half that typically observed in Rarotonga.

This may yet change in the coming months as Telecom Cook Islands (which is rebranding as part of being acquired by Bluesky) commissions its 3G network in Aitutaki - at the time of writing, local staff were running a test node and were awaiting the arrival of the equipment. A number of visitors we spoke to on Aitutaki mentioned having bought large data packs on Rarotonga and being unable to use them on Aitutaki: The demand for the last mile bandwidth on 3G is clearly there and once satisfied, will likely have an immediate impact on the upstream satellite link utilisation.

Steinwurf in particular have for all practical purposes become a direct project partner and are making significant amounts of staff time available in tailoring the software for our requirements - something we didn't expect but are immensely grateful for. Special thanks here go to Janus Heide and Péter Vingelmann. Upon seeing that their original proxy software was not an ideal fit for the project, Steinwurf built a Debian/Ubuntu kernel module for the project which they subsequently extended for multi-endpoint use, all without extra charge. Péter (who is based in Hungary and works remotely) in particular made himself available during every island deployment. This involved extremely odd working hours for him due to the time zone difference between the Pacific and Europe.

Steinwurf also made their own measurement software available at no extra charge. Their input and observations have helped us quite considerably, and in the course of the project they have become co-authors of our publications. This included presenting an invited paper at the 7th International Conference on Wireless and Satellite Systems (WiSats) in the UK in July 2015 on our behalf as the conference was too close to our Aitutaki deployment (the workshop was organised by a researcher from MIT, funding from APNIC allowing Péter to attend is gratefully acknowledged). We have since been able to commission additional work from Steinwurf that will make the overhead sent by the software adaptive to the link conditions encountered.

Procuring a server in California took some time, as we wanted to ensure that we had a provider there capable of looking after the server as well as its networking needs (two interfaces, routing to the subnet or autonomous system spanned by the island side encoder/decoder). Commercial server providers we considered were generally unable to provide the networking aspect, and Internet service providers do not normally host servers on behalf of other entities. We are now pleased to have been able to co-opt the Cooperative Association for Internet Data Analysis (CAIDA) at the University of California San Diego (UCSD) who have been able to provide a machine at the San Diego Supercomputer Centre (SDSC) for us. We are currently awaiting the installation of a second network interface there, which will allow us to start TPC/NC operations. At this point, we will also transit the present networking setup from a University of Auckland subnet to a proper autonomous system.

Input	Project activities	Outputs	Outcomes	Timeline	Status	Assessment
Dr. Ulrich Speidel and 'Etuate Cocker to undertake research project	Subcontract between UofA and PICISOC negotiated and concluded.	Subcontract between UofA and PICISOC in place and executed.	Project implementation able to go ahead.	Jan - Jul 2014	Completed	The subcontract was negotiated between the University of Auckland's research office and the PICISOC board in consultation with and with input from APNIC. The subcontract took a while to get into place due to the unusual nature of the cooperation but is now in place.
Project budget, UoA project team members, UoA IT procurement	Procurement of hardware for island deployment	Computers suitable for deployment, as test bed and spare	Setup and configuration able to go ahead	Jun - Sep 2014	Completed	Procurement took some time as the overseas vendor (Stealth) for the island end machines had not dealt with the University of Auckland before. We acquired five machines within the approved budget: three for island deployment, a fourth as backup / spare and a fifth machine as a test bed here in Auckland. The spare machine was eventually able to be deployed in Aitutaki.
Project budget, UoA project team members, UoA IT procurement	Procurement of NZ endpoint server	Server for NZ endpoint	Setup and configuration able to go ahead	Jul-Nov 2014	Completed	This machine arrived just in time for deployment in Rarotonga and needed last minute servicing by Dell. In hindsight, having two machines in critical locations (single point of failure) would have been desirable.
Project budget, UoA project team members, UoA IT procurement, CAIDA / UCSD / SDSC	Procurement of server in California	Leased server for California endpoint	Server available to act as endpoint in California	Jul 2014 - Jul 2015	Completed	Finding a commercial provider who could ensure both server operation and the necessary routing configurations proved difficult. We were eventually offered a machine by the Cooperative Association for Internet Data Analysis (CAIDA) at the University of California San Diego (UCSD). The server is housed at the San Diego Supercomputer Centre
Project budget, UoA project team members, UoA IT procurement	Procurement of TCP/NC software on research licence	Access to TCP/NC software	Setup and configuration able to go ahead.	Jul 2014 - Jul 2015	Completed	This software is at the core of the project. It consists of a library and several applications built around it, one of which is a pair of proxies that allows tunnelling of standard TCP traffic (and in fact any IP traffic) as TCP/NC over UDP. Procurement took some time as the overseas vendor (Steinwurf ApS of Denmark) of the software had not dealt with the University of Auckland before. Steinwurf are a joint venture between researchers at MIT and Aalborg University, who have vested their IP in network-coded TCP in the company. We have built a good relationship with Steinwurf, who are very supportive of and interested in our work, and have contributed hundreds of hours of staff time at no extra cost. Based on our initial results and suggestions, we have commissioned Steinwurf to add an adaptive overhead feature to the software.
UoA project team members, Stealth computers available, TCP/NC software	Operating system set up on hardware. TCP/NC software compiled, installed and tested. Configuration issues and questions identified.	Computers suitable as a deployment platform for TCP/NC and configured with OS and TCP/NC software.	At least three of the configured computers will be deployed in the Pacific as endpoints for TCP/NC tunnels	Sep 2014	Completed	We now have the hardware and basic software in place that allows us to experiment and familiarise ourselves with the necessary configuration to make the software functional in the range of scenarios we will deploy in (on public IP, private IP, geostationary satellite connectivity, medium earth orbit connectivity, etc.) We know how to drive the software and how to configure the routing on the encoders and decoders to obtain basic functionality as our scenarios are somewhat different from the settings Steinwurf originally tested in. On our suggestion, Steinwurf added the ability to make multiple machines in the islands talk to any machine on the Internet via an encoder / decoder machine each on both sides of the satellite link. This requires a network coded TCP tunnel. In conjunction with Steinwurf, we are also looking at performance data to and from our island sites (which we have a lot of data for) as this will assist us in setting parameters later on.
Project budget, UoA project team members, Network Coding and Reliable Communications Group at Massachusetts Institute of Technology	Deployment visits to the Pacific Island sites.	Encoder / decoder machines deployed at Pacific Island sites	Machines fully configured and operational, initial experiments completed that will give us baseline data.	Nov 2014 to Jan 2015, Aitutaki Jul 2015	Completed	The deployed machines are at the heart of the project as we use them to evaluate how much improvement we can gain from network coded TCP deployment in the islands. Each visit resulted in successful deployment and since each visit took place in a very distinct satellite link scenario, we have been able to learn a lot in each case.
UoA project team members	Dissemination of observations and results	Journal and conference papers, reports, presentations	Users of our findings will be informed and will be in a position to make informed decisions about commercial-scale deployment of network-coded TCP in the islands. Researchers in network coding will gain insights into the behaviour of network coded TCP over very long latency satellite links and satellite links with unstable latencies.	Sep 2014- to date	Ongoing	Initially, we started presenting the project at the University of Auckland and then subsequently to personnel at our deployment locations. We have since also published two conference papers (at Netcod 2015 in Sydney in June 2015 and an invited paper at WiSats in Bradford, UK, in July 2015 - funding from APNIC to attend gratefully acknowledged) and are planning further scientific publications. Ulrich Speidel also presented the project in a lightning session during APrIGF2015 in Macau - funding from Google organised by APNIC gratefully acknowledged.

Project outputs, communication and dissemination activities

***Tips:** Take into account that the reader of your report has not being involved in project implementation, so readers do not have any further knowledge besides the information you are providing here. This section of the report will allow you document the communication and dissemination efforts that the project team has conducted, which might be part of a specific communication strategy design as part of the project, or in place for the organization as a whole. When possible, please provide information about strategies in place and the rationale behind them. Lessons can be learned from many aspects of project implementation, covering a wide variety of aspects such as technical, social, cultural and economical. Taking the rationale behind the project and its objectives can serve as a framework to draw your conclusions. Lessons can be identified by project partners, beneficiaries and general staff from the organization. A project diary and other activity records can serve as a tool to reflect during project team meetings and immediately after project activities are conducted.*

> ***Outputs are immediate, visible, concrete developmental change that is the tangible consequence of project activities, under direct control of the project team**.*

Example of possible outputs to report are:

- *New products and Services (software, online platforms, applications);*
- *Information sharing and dissemination (publications, conferences, multimedia, social media);*
- *Knowledge creation (new knowledge embodied in forms other than publications or reports, such as new technologies, new methodologies, new curricula, new policies);*
- *Training (short-term training, internships or fellowships, training seminars and workshops) and*
- *Research Capacity (research skills; research management capacity and capacity to link research to utilization of research results).*

Project outputs	Status	Assessment	Dissemination efforts
Subcontract between UofA and PICISOC in place and executed.	Completed	This contract allowed the research work to proceed.	At the University of Auckland, we have started a variety of internal and external promotion activities, including regular presentations to the wider network research groups and a cameo appearance at the Undergraduate and Alumni Day to increase awareness of the project within the university and related community. Our media team have been briefed and are also on the prowl for opportunities to promote the project.
Computers suitable for deployment, as test bed and spare	Completed	These machines form the core of our Pacific Island deployment - they will be the encoders / decoders for network-coded TCP at the island endpoints. The machines' arrival in Auckland is a major milestone on the way to deployment.	Dissemination efforts here have been mostly to inform project partners of progress.
Server for NZ endpoint	Completed	This machine forms the core of the NZ endpoint.	Dissemination efforts here have been mostly to inform project partners of progress.
Server for California endpoint	Completed	This machine forms the core of the Californian endpoint	Having the machine at CAIDA at UCSD gives us good visibility in the Internet community even though the environment is not a commercial one.
Access to TCP/NC software	Completed	This software performs the encoding and decoding tasks for TCP/NC. It is the core software for the project.	Dissemination efforts here have been mostly to inform project partners of progress.
Computers suitable as a deployment platform for TCP/NC and configured with OS and TCP/NC software.	Completed	Ready and deployed	Dissemination efforts here have been mostly to inform project partners of progress, especially MIT.
Visit to MIT	Completed	This was the first time a Pasifika graduate student from the University of Auckland's Computer Science department got to visit MIT as a visiting student, albeit for a short time only, to learn about a technology that has not been used in the Pacific before.	This represents a nice opportunity for positive reporting of 'Etuate Cocker as a role model for young Pacific Islanders interested in technology careers and in making a difference for their communities. We intend to share this story both within the university community but also within the Pacific Island community (Pacific Islanders are woefully underrepresented in academia, and especially so at graduate student level). We have reported on this regularly via the Pasifika networks around the University of Auckland and other Pacific Island links we have, e.g., via the NZ Tourism Research Institute.
Encoder / decoder machines deployed at Pacific Island sites	Completed	This will start the main part of the data collection.	We gave relatively immediate presentations / demonstrations to our onsite partners and their communities. Formal reports to some of the project partners are still to follow once we have been able to collect more data. We started to disseminate the project's findings through scientific publications in conferences [9][11] with journals to follow soon. Early findings will also be included in 'Etuate Cocker's PhD thesis. We have also reported on the project at APrIGF in Macau [10].

Project outcomes

***Tips:** This section should be completed **ONLY** for the final report. ISIF Asia expects you to report about the **outcomes** of the project as defined in the table below, based on the project implementation section of this report. Project team is encouraged to discuss the questions provided below to guide the reflection:*

Can you identify and describe the relationships between the activities implemented and the social, economical, cultural and/or political benefits of your project implementation?

Outcomes can be defined as:
• *Medium-term effects* • *Effect of a series of achieved outputs* • *Should capture the changes for the beneficiaries* • *Take place during the life of project/strategy* • *Influence but not direct control*

The primary goal of this project was to find out whether TCP/NC could improve the connectivity for Pacific Island nations dependent on geostationary (GEO) or medium earth orbit (MEO) satellite links. With hardware and software in place in four deployment locations, the first set of outcomes consists of answers to the research questions asked at the outset:

- **How difficult is it to implement a TCP/NC tunnel?**
  In principle, this is easy. Supplying an autonomous system (AS) on an island with traffic through a TCP/NC tunnel requires merely two encoder/decoder gateways on the Internet and an encoder/decoder on the island. If the network supplied is merely a subnet of an AS, a single gateway suffices. Neither gateway(s) nor on-island encoder/decoder require specialised hardware; a sufficiently powerful PC or server running Ubuntu or Debian are all it takes to perform each of the three roles.

  Some degree of network topology planning is required, however: IP traffic to the island must reach the gateway(s) to be encoded, and from there, the UDP carrying the encoded traffic must reach the island encoder/decoder via unencoded IP datagrams. This requires two subnets where previously the entire island could have been treated as a single subnet.

- **How much performance gain could we expect as a baseline?**
  This turns out to be the wrong question to ask. The gain in goodput that can be achieved for individual connections depends on two factors:

  Firstly, it depends on the actual degradation of the unencoded link (i.e., the link utilisation achievable with unencoded traffic). If link utilisation is low due to queue oscillation (e.g., in Rarotonga and Tuvalu), TCP/NC can recapture most of the unused capacity. Improvement factors seen range up to 20, with typical factors of between 2 and 4. If link utilisation is low due to lack of demand (Aitutaki), there can be no gain, and the overhead of TCP/NC in combination with decoding delay leads to a loss of goodput. If link utilisation is high due to strong demand, TCP/NC can be used to prioritise individual flows but cannot provide better goodput overall: TCP/NC cannot create bandwidth - it can only use what already exists and may have been left idle.

  Secondly, the gain depends on the amount of overhead that TCP/NC needs to use in order to protect against packet loss, and on the degree of protection this gives. Surplus overhead comes at the cost of goodput, but underprovisioning of overhead carries the risk of losing entire systems of equations, which leads to severe burst packet losses in the TCP/NC tunnel and an associated loss in TCP/NC goodput.

- **How much gain could we get from fine-tuning?**
  We have found that links do change over time, in particular with a strong diurnal pattern. In our experiments, we used parameters designed to allow good performance during peak time. During off-peak hours, these represent excessive overhead, however, which may impact on performance. Steinwurf are currently converting our insights into an adaptive version of their software, which we expect to experiment shortly.

- **Can increased goodput with TCP/NC be achieved without loss of capacity for conventional TCP?**
  The answer here is yes as long as the link is affected by queue oscillation. We even believe that the presence of TCP/NC may have a stabilising effect on the queue. Where link utilisation is high as a result of excess demand, TCP/NC pushes conventional TCP aside. However, note in this context that this link condition is already characterised by extremely slow long flows that often time out.

- **Are there benefits in tunnelling all traffic across a satellite connection via TCP/NC?** In the case of queue oscillation, the presence of conventional TCP is the sole cause of queue oscillation and hence of the burst packet losses that slow conventional TCP down. These burst losses also force TCP/NC to operate with significant overhead (extra combination packets).

  If we were to encode all traffic across a connection, the queue oscillation should largely disappear and it should be possible to encode with much smaller amounts of overhead, which should increase the usable capacity. In practice, we were unable to test this hypothesis as all links under consideration were production links and could not have easily been switched between encoding all traffic vs. no traffic.

- **What obstacles would there be to wholesale adoption in a particular location?**
  The main obstacle as it appears to us now is the need for off-island encoder/decoder gateway(s), as these sit topologically outside the network of island-based ISPs, forcing them coordinate adoption with their satellite service provider or at least an off-island data centre able to host and maintain the gateway(s) and the required routing arrangements.

- **How scalable is the technology and how much computing capacity do we need for encoders and decoders in order to carry a certain amount of traffic in this context?**
  The satellite links we investigated ranged from 8 Mbps to over 300 Mbps in provisioned one-way capacity. The Stealth industrial computers we used as island-side encoders/decoders had no problem decoding 50 Mbps of traffic without showing significant load effects - we believe that they would be able to handle links up to half a Gbps with ease.

  Since the decoding is a more computationally intensive task than the decoding, and most islands import vastly more data than they export, the encoders/decoders on either side represent a relatively inexpensive add-on to a satellite link.

  This study is not able to comment on software pricing as it operated with an experimental license. However, conceivable schemes could include fixed-price licences, time-, bandwidth, or volume-based licenses, licenses based on the expected capacity recovery, or the provision of off-island gateways as a service.

- **How robust is the existing technology? Is the existing TCP/NC software sufficiently stable for long-term unattended operation?**

Our experiments have not been able to find any stability issues in Steinwurf's kernel module to date.

- **How stable are our results?**
  All of our results to date show that when low link utilisation coincides with significant packet loss (the two characteristic features of queue oscillation), TCP/NC outperforms conventional TCP. If burst packet losses are well within the overhead of TCP/NC, TCP/NC connections achieve goodput comparable to that of conventional TCP on the link at times of no packet loss and low demand.

- **How do changes in satellite connectivity and usage impact on our initial findings?** We have identified three types of link states:
  - *Pre-oscillation*: In this state, the demand/number of users is too low to trigger queue oscillation. Link utilisation is below capacity and one observes no significant packet losses. Standard TCP is able to claim the whole link capacity on long downloads. At the time of writing, the link into Aitutaki is in this state almost all the time. In this state, TCP/NC offers no advantage, but growth in usage may cause the link to enter a regime where the packet queue at the satellite gateway starts to oscillate over increasing periods of time.
  - *Oscillating queue*: In this state, we observe both low link utilisation and significant packet loss. At the time of writing, the links into Rarotonga and Tuvalu (Funafuti) are in this state, with off-peak returns to pre-oscillation state. In Rarotonga, an upgrade of satellite capacity during the first half of 2015 seems to have resulted in longer burst losses, probably as a result of queue saturation setting in at higher packet rates only, leading to more radical back-off (note that TCP back-off is exponential in nature, whereas slow-start acceleration is generally modelled by a polynomial function). TCP/NC is able to compensate here as long as the burst losses don't exceed the configured overhead.
  - *Saturated/congested*: In this state, the link is dominated by a high rate of new connections, many of which complete before an oscillation cycle. Connections transferring a higher data volume experience oscillations and slow down significantly and/or time out. TCP/NC can help save such connections, albeit at a cost to the shorter flows. Link utilisation is high, as are packet losses once load is offered. The link into Niue was observed in this state for most of the day. Increased connection capacity can return such a link to the two previous states.

- **Can TCP/NC deliver the critical end user benefit without hitting the physical limit to the satellite bandwidth?**
  We can answer this in the affirmative, at least for the link into Rarotonga, where TCP/NC consistently gives access to otherwise dormant capacity. In Tuvalu, we have also seen significant dormant capacity, and believe that if all or at least most of the traffic across the link were to be encoded, TCP/NC would likely be able to prevent the currently seen level of queue oscillation and deliver benefits of perhaps up to 4 times the raw goodput currently seen, well within the currently deployed bandwidth.

  In Niue, the bandwidth currently deployed means that using TCP/NC comes at a cost to existing TCP goodput, some of which is spent on very slow long connections, however, i.e., on connections that may not complete.

- **Could the technology help to create critical mass for increased uptake, and hence result in business cases for upgraded satellite connectivity because it promises users a tangible return on increased investment?**

  We have been able to demonstrate that in locations where queue oscillation prevents standard TCP from accessing the entire capacity provisioned, TCP/NC can claw back a significant part of the idle capacity. In the location where we observed queue oscillation, it occurred predominantly during times of peak demand. On single connection 20 MB downloads into Rarotonga, we observed claw-backs of between 5 and 20 Mbps (25-400% goodput improvement over standard TCP) during peak periods at the end of January 2015 and, with much higher burst losses, of around 1 Mbps at the end of May 2015 (around 100% goodput improvement).

Dedicated satellite bandwidth typically retails to ISPs for hundreds of dollars per Mbps per month. An oscillating link with 100 Mbps nominal bandwidth and an achievable utilisation of 70% with conventional TCP could be operated with, e.g., 10% overhead if fully coded. This would make around 20 Mbps of additional goodput accessible using the same link. At, e.g., US$500/Mbps/month, coding would thus give access to around US$10,000 worth of idle capacity each month.

Hardware cost for the projection of a small AS via TCP/NC (one on-island encoder/decoder unit, two units off-island) could be as little as US$5000. Based on 25 months linear depreciation, this equates to hardware cost of US$200/month, plus off-island hosting costs of perhaps US$200 per offshore unit, i.e. around US$600/month. Installation and commissioning requires around two to three days on the ground on-island. Personnel for deployment would typically include one specialist on the island, a network engineer from the local ISP, and a specialist off-island.

As mentioned above, we are unable to comment on the cost of a software license and support in this context.

As a result of our study, we now understand the causes of impaired connectivity much better. Packet loss can in principle be caused by two phenomena: Bit errors in packets as a result of noise, distortion or interference, or the loss of entire packets as a result of queue overflows in routers and other interface equipment. By simulating the links of our deployment locations, we now know that queue overflow during oscillation can explain both the packet losses and the link underutilisation observed.

Our simulations demonstrate that queue oscillation occurs even when we model the actual satellite link with a zero bit error rate. That is, we now know that bit errors are at best a secondary cause of packet loss, i.e., that the losses are generally not caused by faults on the satellite links. TCP/NC would be able to cope with bit error related losses as well, however. In practice, we observe significant packet loss only on satellite links but not on tests between fibre-connected sites where there is no bandwidth bottleneck shared by a large number of TCP flows. This packet loss also correlates well with times of high demand.

Critical ingredients for queue oscillation are (a) a bandwidth bottleneck, (b) a large number of sufficiently voluminous TCP connections being created at a sufficiently high rate by a sufficiently large number of senders, and (c) a sufficient amount of latency in each connection. If any of these conditions are not met, the queue will not oscillate, and if the connection creation rate on the link is too high, the queue will still overflow but will not drain completely, i.e., the link utilisation will be close to 100%. We can therefore state with some confidence that TCP/NC

Feedback from ISPs in the field indicates that they would like to have access to this technology on a commercial basis, albeit as "part of the routers [they] usually buy". One ISP indicated that they cannot easily request equipment to be added at off-island satellite gateways and that they are apprehensive about the hassles of dealing with a data centre off-island to host the encoders/decoders there or of having to restructure their network topology.

Our work has resulted in two scientific publications to date [9,11].

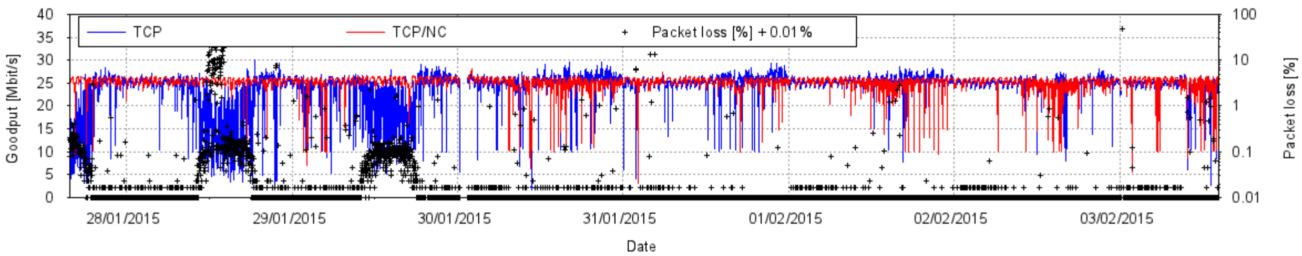

Goodput of standard TCP (blue) and TCP/NC (red) downloads between New Zealand and Rarotonga over the period end of January / early February 2015. Note that TCP/NC maintains high goodput at times of high packet loss.

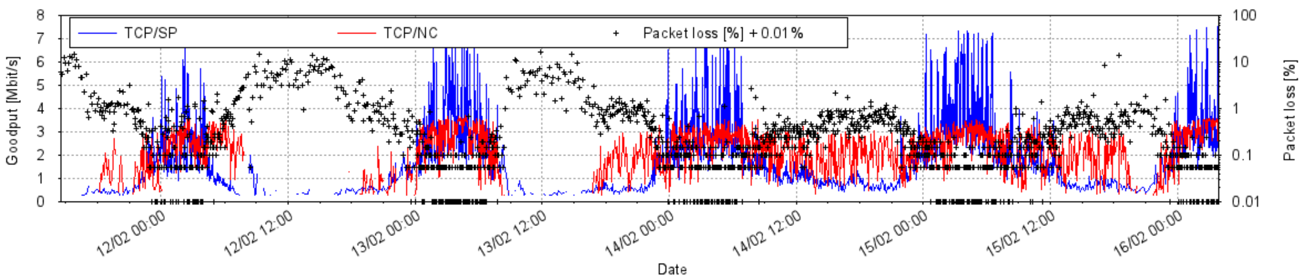

Goodput of standard TCP (blue) and TCP/NC (red) downloads between New Zealand and Funafuti (Tuvalu) in February 2015. At times of high packet loss, neither TCP nor TCP/NC downloads complete. However, unlike standard TCP, TCP/NC maintains high goodput at times of moderate packet loss up to 1%. A scenario such as this would probably benefit if all traffic migrated to TCP/NC, as this would mitigate the queue oscillation and hence the packet loss. Note that the standard TCP goodput measured here includes traffic delivered via SilverPeak network memory, inflating the apparent standard TCP goodput during calmer off-peak periods.

Project management and sustainability

***Tips:** Please comment on the general project **administration, staffing, procurement**, etc. specially those aspects contributing to the fulfilment of the project objectives as well as those that have delay project implementation.*

*Indicate **how the project team has strengthened its capacity** and work towards sustainability with the support provided by ISIF Asia? (new equipment, training, improved administrative skills, lessons learned from the project). Has the organization increased its research or administrative skills of the team involved? Has the project allowed for a particular contribution to capacity building of women or marginalized social groups? Special attention should be paid to the expected or unexpected impact on marginalized social groups.*

*Have you done **anything different** to provide administrative support for this project **besides your "business as usual"** processes and procedures? Has the project inspired change inside your organization?*

> ***Sustainability is to be examined not only in terms of staff retention and financial stability of the organization supporting the project but about the communities' appropriation of benefits perceived from project implementation.***

The ISIF Asia Secretariat is very interest to learn if this project has generated opportunities for future development (new funding from partnerships, sponsorships, investment or other funding mechanisms), please provide details.

Please explain if the ISIF Asia grant has helped to consolidate your organization and how. If any of the project activities will continue after the end of the ISIF Asia grant, please describe how your organization is planning to support future developments.

Administration and procurement: This project is somewhat unusual in that much of its implementation and management was handled by the University of Auckland under subcontract with the grantee, the Pacific Island

Chapter of the Internet Society. The start of the project was later than anticipated at application time, which has had a number of consequences:

- The start of the project coincided with a major administrative restructuring of faculties at the university. This meant that a number of staff members involved in procurement of equipment and software were trying to cope with a changing set of processes while also having to deal with extra load and uncertainty resulting from the changes. This led to additional delay and required a significant amount of follow-up. On the positive side, issues encountered during procurement were noted by various parties at a time when many processes are under scrutiny and we expect this to inform some of these.
- A number of the roles surrounding finance and reporting have changed, so the team had to feel its way forward in the procurement of financial reports for this project as we had to learn how the new processes worked.

The University of Auckland's Research Office allows spending on research budgets based on live contracts even if the granting organisation has not been invoiced yet. This gave the project a head start. We would like to acknowledge APNIC's diligence in this context, which ensured that an outstanding invoice was detected and paid before it was flagged by the Research Office administration.

On the positive side, the delayed start enabled us to take advantage of a fall in computer prices and a strengthening NZ dollar, meaning we were able procure extra machines to serve as a local engineering testbed and a spare, which we were later able to deploy in Aitutaki. On the downside, the NZ dollar's appreciation against the Australian dollar put pressure on the overall budget. This was however more than offset by our ability to procure an overall software license for experimentation covering all equipment rather than the per-machine licenses we had budgeted for, plus an unexpected fees waiver by MIT, which is gratefully acknowledged, savings from not having to lease a server in the U.S., and lower than expected travel cost due to off-season travel and close location of a major conference (Netcod2015).

Staffing: The research team would like to take this opportunity to acknowledge the work done and patience shown by PICISOC board members in a volunteer capacity, especially Ms. Maureen Hilyard, who devoted a significant amount of time on the contractual arrangements. We would also like to acknowledge the staff time and other assistance contributed by our Pacific Island deployment partners.

Mr. 'Etuate Cocker worked on this project in his capacity as a PhD student at the university. The university supported him with a scholarship, which we were able to extend to the end of 2014. We had originally anticipated to start and finish earlier, and therefore did not expect having to find additional support for him. The delayed schedule meant that he needed support for January and February 2015 as well, consisting of approximately NZ$3,800 in living expenses and around NZ$1,200 in fees. We were able to find local funding of around NZ$3,324.34 to cover most of this, and used savings from the project to make up the shortfall (around NZ$1,880).

Sustainability: The project's aim was "proof of concept" and as such the project itself does not aim for sustainability beyond the point where it has delivered the answers and insights we seek. It will however leave a legacy that can be built upon, first of all for further measurement but in future also for potential production use. Our results to date show that we can indeed improve Internet user experience with TCP/NC on at least some links, and we expect there to be a strong incentive for providers in satellite-connected islands to implement this technology in future as it promises significant bandwidth claw-back for relatively little investment. Having our relatively small-scale experimental pilots in place provides a working model for implementers to follow. Moreover, we aim to measure for years to come, and although we expect the bulk of our findings to eventuate early on, the platform deployed can serve both as a small-scale production platform as well as for experiments and measurements and as such builds capability.

Impact

***Tips:** This section of the report does not refer to the project activities, but about the **"bigger picture"**. It will be desirable if the project team can reflect on the **impact that the project has contributed to as part of other actions implemented by your organization and/or your partners**.*

> ***Impact refers to the influence the project may had on the way people does things through the use or adoption of the project outputs; changes in the context the project was implemented; changes in the community the project has been working with; and/or changes inside the organizations that have participated in the implementation or the relationships established through the project's implementation.***

Impact is often impossible to measure in the short term and is rarely attributable to a single activity. Impact can be linked to a vision or long-term development goal that your organization might be working towards.

It can be identified as a logical consequence of achieving a combination of outputs and outcomes.

Impact is usually measurable after the project life and is outside the direct control of the project team and the organization.

We observe and anticipate the following impact:

- Our project partners see and understand the benefit of TCP/NC. We would expect routine use of TCP/NC across satellite links in and out of the Pacific and elsewhere in years to come, once the technology can be packaged such a way that it easy to deploy.
- Our project motivated Steinwurf to improve their existing software significantly with a number of new features of substantial practical value, including the ability to terminate tunnels to multiple peers at a single encoder/decoder machine, and the ability to capture the amount of surplus overhead received by a decoder and feed it back to the encoder for the implementation of adaptive overhead schemes. Steinwurf are now contemplating ways of making their product commercially available in the Pacific.
- A better user experience of the Internet on satellite-connected islands with oscillating links, providing better value to end users, with a personal and economic benefit that provides incentive and a business case to seek better connectivity, resulting ultimately in more deployed satellite capacity and perhaps additional cable connections.
- Adoption of the technology in other similar locations, including ship-to-shore satellite links, where our findings inform and contribute to decision making. Presentation of our results to date at WiSats in the UK caught the attention of the European Space Agency (ESA) and Inmarsat, and MIT are trying to broker a collaboration with ESA's Research and Technology Centre (ESTEC).
- The queue oscillation observed in the Pacific has led to a new PhD research project pursued by Mr. Lei Qian at the University of Auckland. It attempts to simulate satellite links, with the goal of being able to predict under which circumstances a queue will oscillate and how much improvement TCP/NC should yield in a given scenario. First results from this research have been published already [11]
- A higher level of awareness of the needs of Pacific Internet users in the communications research community where next generation technologies are being developed, and relationships (such as the ones with MIT and Steinwurf) reflecting that.
- A higher level of awareness of the needs of Pacific Internet users in the Internet governance community (attendance and presentation at APrIGF assisted in that).
- Closer links between PICISOC and the University of Auckland.
- Closer linls between the University of Auckland and Internet Service Providers and other stakeholders in the Pacific.

Overall Assessment

Tips: This section of the report is extremely valuable for the ISIF Asia secretariat as it provides evidence about the role and relevance of ISIF Asia contributions in the Asia Pacific region.

Tips: *Briefly provide* ***your own views*** *on the value and importance of the project relative to the proposed innovation, investment of time, effort and funding involved. Include the strengths and weaknesses of the project and the steps taken to strengthen the credibility and reliability.*

This is your opportunity to conduct a ***team reflection about the value of the project for the organization****. The following questions might help you to prepare a substantive overall assessment.*

- *To what extend the project meet its objectives?*
- *What were the most important findings and outputs of the project? What will be done with them?*
- *What contribution to development did the project make?*
- *Were certain aspects of project design, management and implementation particularly important to the degree of success of the project?*
- *To what extend the project help build up the research capacity of your institution or of the individuals involved?*
- *What lessons can be derived that would be useful in improving future performance?*

PICISOC is an organisation that is distributed across multiple countries in the South Pacific and that relies heavily on volunteer labour. ISIF Asia's willingness to let PICISOC undertake this project in conjunction with the team at the University of Auckland has enabled PICISOC to tap into a technology that might otherwise never have reached the Pacific. It has also enabled PICISOC to handle a project of a significant size that promises significant benefits for end users in the region. In this respect, the project has broken new ground both on the technical and organisational front.

The project team at the University of Auckland has enjoyed an exceptional level of support from various quarters, including significant additional funding from Internet NZ, logistical support from the intended deployment sites, a lot of hands-on advice and support from Steinwurf ApS, the generous offer from Prof. Médard at MIT to host 'Etuate Cocker's visit there in order to help us build capacity, and last but not least from the PICISOC board. The fact that we were able to deploy in four locations rather than the three originally anticipated speaks to everyone's willingness to support us in the field. The project team would like to use this opportunity to extend its gratitude to all parties involved.

This strong support and goodwill has been instrumental in overcoming many obstacles along the way. We had originally hoped to be able to start at the beginning of 2014, but the time it took to get the contractual arrangements in place had knock-on effects on the rest of the schedule that were outside the control of ISIF Asia, PICISOC and the project team. Procurement was also more complicated than anticipated, and this part had to be performed during a period of significant administrative change at the university.

We can now confidently recommend the use of TCP/NC in situations where satellite links show both underutilisation and packet loss, and suggest that careful consideration should be given to its use in scenarios where links are overloaded and a large amount of packet loss prevents TCP connections from completing.

Recommendations

Tips: *Include any recommendations in this section that you and your project team, the organizations supporting the project and the community you worked with, would like to make to other practitioners or researchers on the field facing similar problems or implementing similar solutions. Please take a minute to share recommendations with the ISIF Asia secretariat that might help to improve the support provided.*

At this stage, we have the following recommendations:

- The support ISIF Asia rendered to us in getting the word out was exemplary, and we recommend that these efforts be sustained.

- Project timelines and budgets are generally written at application time. Between application and award of the grant and the eventual conclusion of the award agreement, there can be a substantial time period, during which circumstances and budget may change. Perhaps it would be a good idea to specifically ask for any amendments to timetable and budget planning during conclusion of the agreement, and make the review of both a regular feature with the assumption that most timetables and budgets will need some sort of adaptation over the course of the project.

- We found ISIF to be very flexible and willing to vary the specifics of the contract during the project. However the original contract did not reflect this flexibility, which caused additional work for the project team locally. We recommend that future contracts have a provision for a simple process of variation, such that timelines and partitioning of budget categories may be varied within reason by simple e-mail approval.

- Allow time for procurement that is not straightforward. The University of Auckland changed its procurement procedures quite substantially over the last few years. Reimbursements of expenses other than travel incidentals are now more or less a thing of the past and dealing with non-standard vendors now requires a complex multi-step setup process with several approval levels.

- The report template seems to be written for final reports, not interim reports, and seems unnecessarily complex for small projects, the tables in particular.

Bibliography

***Tips:** Include complete bibliographic references to all sources (printed, on-line, quotes, etc) used to prepare the different sections of this report. The APA style guide offers examples about how to reference a variety of sources. http://www.apastyle.org/learn/quick-guide-on-references.aspx(as accessed on 3/7/2013).*

[1] Southern Cross Cables home page, http://www.southerncrosscables.com/

[2] O3b Networks home page, http://www.o3bnetworks.com/

[3] V. Cerf, Y. Dalal, C. Sunshine, “Specification of Internet Transmission Control Program”, RFC 675, 1974, http://tools.ietf.org/html/rfc675

[4] J. Postel (ed.), “Transmission Control Protocol”, RFC 793, 1981, http://tools.ietf.org/html/rfc793

[5] R. Braden (ed.), “Requirements for Internet Hosts - Communication Layers”, RFC 1122, 1989, http://tools.ietf.org/html/rfc112

[6] M. Allman, V. Paxson, W. Stevens, “TCP Congestion Control”, RFC 2581, 1999, http://tools.ietf.org/html/rfc2581

[7] M. Allman, V. Paxson, E. Blanton, “TCP Congestion Control”, RFC 5681, 2009, http://tools.ietf.org/html/rfc5681

[8] J. Sundararajan, D. Shah, M. Médard, S. Jakubczak, M. Mitzenmacher, J. Barros, “Network Coding Meets TCP: Theory and Implementation”, Proc. IEEE, 99(3), March 2011, pp. 490-512, http://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=5688180

[9] U Speidel, ‘E. Cocker, P. Vingelmann, J. Heide, M. Médard, “Can network coding bridge the digital divide in the Pacific?”, 2015 International Symposium on Network Coding (NetCod 2015), Sydney, 22-24 June, 2015 (to appear in IEEExplore)

[10] U Speidel, ‘E. Cocker, P. Vingelmann, J. Heide, M. Médard, “Can network coding bridge the digital divide in the Pacific?”, lightning presentation, Asia Pacific regional Internet Governance Forum (APrIGF) 2015, Macau, 30 June - 3 July, 2015

[11] U Speidel, L. Qian, ‘E. Cocker, P. Vingelmann, J. Heide, M. Médard, “Can network coding mitigate TCP-induced queue oscillation on narrowband satellite links?”, invited paper, 7th EAI International Conference on Wireless and Satellite Systems (WiSATS 2015), 6-7 July, Bradford, UK