



DTN: Upgrading Martian Carrier Pigeons


How Delay and Disruption Tolerant Networking works
and how they can be used in various places

Jeroen Massar <jeroen@unfix.org>







Delay/disruption tolerant networking



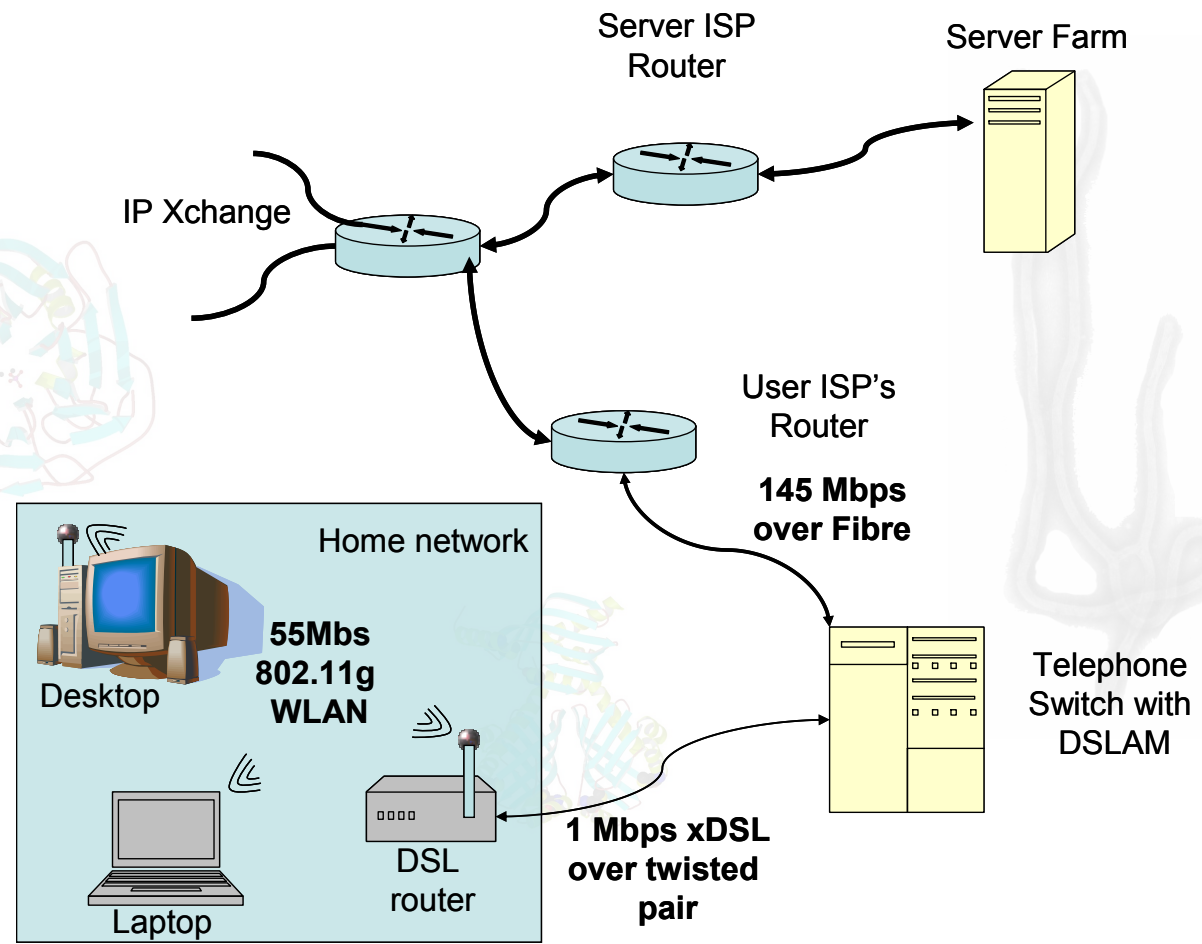
20050305
~350 W-hr/day

20050315
800+ W-hr/day








Typical "Internet" Setup



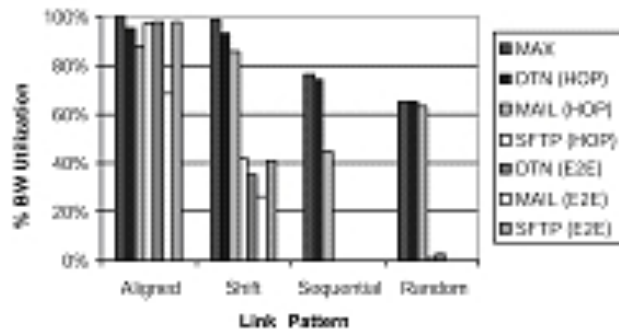


TCP Session






Classic Experiment




(b)

(a)



(c)

Demmer, M., et al., "Implementing Delay Tolerant Networking,"
Intel technical report, IRB-TR-04-020, Dec.28 2004.
<http://www.cs.berkeley.edu/~demmer/papers/dtn-irb-tr-04-020.pdf>





Causes of Delay and Disruption

- Laws of physics
 - Light trip time
- Conserving power
 - Batteries have a crappy Moore's law
- Intermittent availability
 - Lectures 12-1 Tues and 2-4 Thurs only!
- Nothing happening
 - Most of the time, sensors must be really bored!
- Bad things happening
 - DDoS an edge router and what happens?



Laws of physics

- Light trip time
 - 300,000 m/s isn't slow but...
 - Mars varies from 4-20 minutes
- Ephemeris
 - You can't see Mars for half the day
 - You can't see a place on Mars for that long
- Positioning inaccuracy
 - N-body problem
- Visibility
 - Local horizon; rain
 - A Mars orbiter can only see Earth for 2/3 of the time
- Solution: Ephemeris library & (large) data files
 - <http://www.ephemeris.com/>
- Similar effects arise in very, very high speed networks
 - Don't want to wait for signalling traffic
 - <http://www.eecis.udel.edu/~mills/> somewhere in Dave Mills' site



Conserving Power

- For many relevant devices, battery depletion = death
 - Even for phones, we cannot deplete battery in hours, absolute requirement that we can carry phone about for days
 - Radio is expensive
 - 1 bit RF \approx 300 instructions on CPU
- Related issue: ground coverage vs. radio range
 - Network density factors
 - How many bluetooth devices/motes to cover the Phoenix park?
 - Major plausibility concern with most ad-hoc networking
 - But not for DTNs:-)



Power in space

- All spacecraft are essentially power limited
 - Solar panels are ok for inner Solar System
 - Outer Solar System: radioisotope (aka nukes)
 - Usually electrical and heat
 - Not nuclear propulsion (yet)
- Mars Exploration Rovers (MER) solar panels degrade
 - In winter
 - They get dusty
- Power constrains bandwidth
- Have to warm-up (or cool down) devices prior to operation



Intermittent availability

- Moving nodes
 - “Drive-thru” Internet, cars drop off and pick up HTTP traffic at motorway service stations as they drive by
 - <http://www.drive-thru-internet.org/>
- Most ad-hoc network proposals/poster-children have to deal with this issue
 - If routes update faster than AODV can handle it then you’re hosed



Not much happening

- It must be very boring to be a node in a sensor network!
 - Remembering to conserve battery means we turn off the radio when there's nothing to say
 - Makes it hard to talk to me!
- Power on/off events vary
 - Beacons, schedules, sun-rise/set
 - Wake on WLAN
 - But potential DoS!



Bad things happening

- Worms, viruses, DDoS etc., when they happen are very disruptive, even to the core Internet
 - Much more deadly for many embedded systems: make a brick today!
- Many standard countermeasures are difficult to apply in DTNs



So...what's a DTN then?

- A delay or disruption tolerant network is a network that takes special care to be able to operate in even the most challenging environments
 - But who cares about definitions!
- What applications benefit and how?



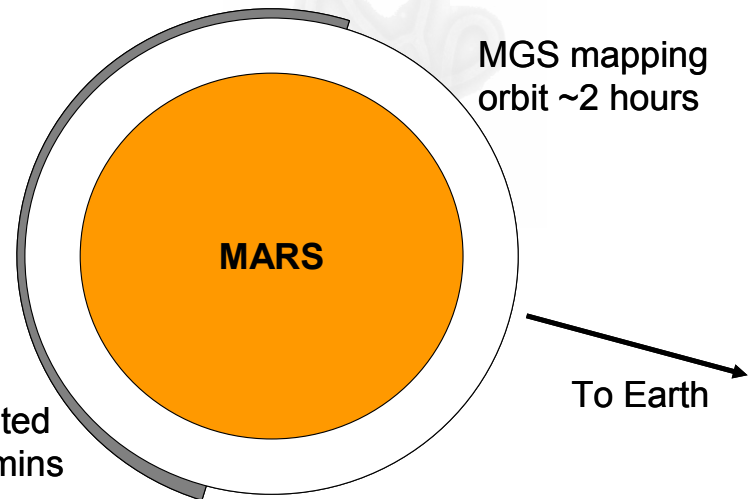
Environmental monitoring

- Imagine a lake somewhere (say in Killeshandra, Co. Cavan)
- Water quality is an issue
- Measuring that is currently expensive
- Would like a sensor network, but:
 - ...scale... $O(100m)$
 - ...its quiet...few visitors
 - ...only GSM coverage...SMS costs add up



Deep Space Networking

- Mars-probes have sort of done networking
 - MEX as relay for MER (\$800M issue!)
- Spacecraft operators hate this idea!
 - They care about mission success only
- Spacecraft in extended mission more likely to be usable
 - I know a JPL guy who can't wait for MER to get a wheel stuck!





More deep-space

- Deep space network (DSN) is critical here
 - 3 main sites (Canberra, Goldstone, Madrid) with 70, 35 & 14m antennae
 - Bottleneck in all deep-space missions
- Interesting question here about range:
 - GEO satellites – TCP with performance enhancing proxies seem to work, mostly
 - Cislunar applications – if there's a BIG pipe to/from Earth we might want to keep that busy
 - Beyond Mars, perhaps DTNs are silly, given there are no plans for multiple spacecraft




SeNDT

- Sensor networking with delay tolerance
 - <https://down.dsg.cs.tcd.ie/sendt/>
 - <http://www.trinityisland.com/>
- Approach is to deploy sensors that communicate via WiFi with a dutycycle of about 10% (i.e. power-off 90%) and use boats as data-mules
 - Sensors in the lake now
 - Instrumenting boats...real soon now
 - Waterproofing;-)





Sensor Node



Intel Xscale, arm-linux
Own PCB with power-off
Solar panel and batteries
Comms. PCMCIA/CF
IEEE 802.11b
Storage: CF disk

Enclosure: various COTS
IP-68 components
Various sensor connections

A/D, RS-232, USB, ...

Nodes are off ~90% of
the time to save power





Data Mules

- Data-mules (aka Message Ferries) are physically moving nodes that provide connectivity for nodes



- Boats are hired for fishing and/or shooting
- Generally taken out daily or weekly depending on the season
- Tend to traverse relatively predictable areas of the lake system



Noise Monitoring

- Also part of SeNDT
 - Was planned as a 2nd pilot but taking over since its much closer to being commercial
- Problem is urban and motorway noise monitoring
- Getting out of car to measure:
 - Expensive and a number of deaths
- Nodes measure sound and data mule (bin lorries, eventually) pick up data
 - Garbage collection!






Noise Boxes






Developing World Applications

- DTNs might provide a nice way to meet some of the challenges involved in networking in developing regions
 - Wizzy digital courier (South Africa)
 - <http://www.wizzy.org.za/>
 - Postmanet
 - <http://www.cs.princeton.edu/~rywang/distance/>
 - TierStore

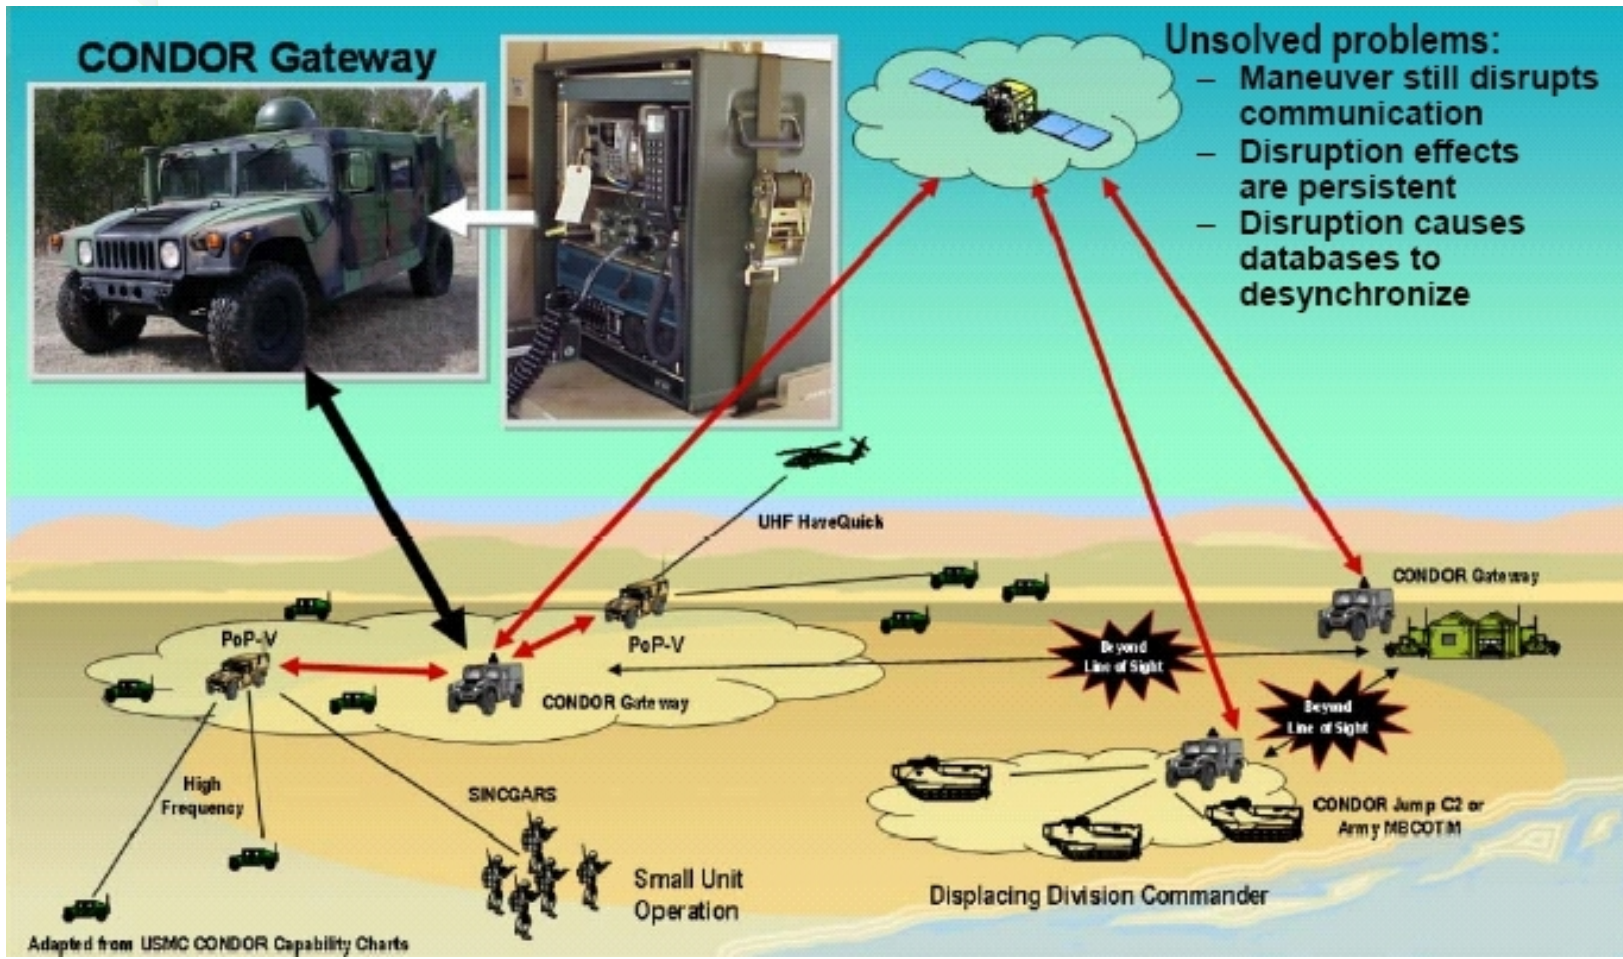


Acoustic networking



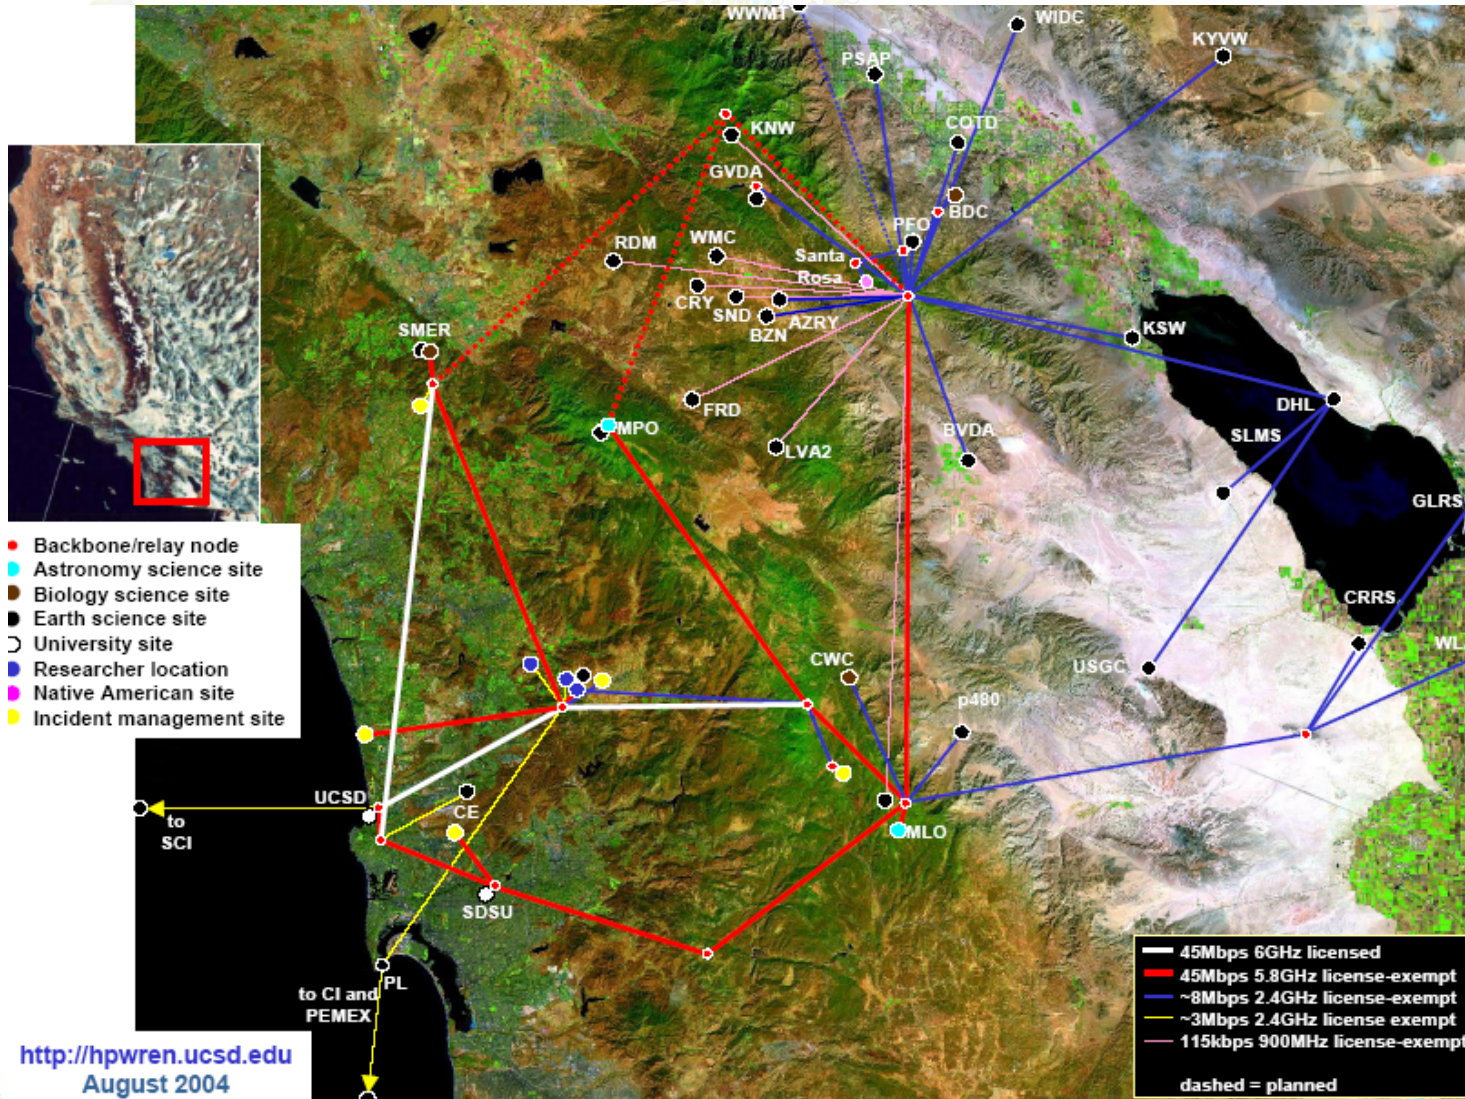


US marine corps.





Forest fire-fighting assistance







DieselNet

- Entirely different, but equally tenable approach to SeNDT
 - Deploying a network with no applications in order to see what's possible
- Instrumented busses in Amherst (UMass)
 - Standard PC106 boxes per bus, each an 802.11 AP & node
 - Busses exchange (random) data as they come into range of one another
- Throwboxes improve throughput in remoter parts of bus routes
 - Probability of contact increases dramatically





DieselNet





Throwboxes






How to DTN?

- “Official” work is done in the IRTF DTN research group (DTNRG)
 - Bundling: an overlay network approach
 - LTP: a point-to-point protocol intended for use on deep-space (long-haul) links
 - Bundling/LTP makes sense
- Most projects are, in fact, using home grown protocols (for now)



Credits

Stephen Farrell
@ Trinity College Dublin




ISBN-10: 1596930632
ISBN-13: 978-1596930636



Descolada

- Name origin: Ender's Game
- The Descolada is first mentioned in "*Speaker for the Dead*", and plays a leading role in the later book "*Xenocide*".
- A quasi-conscious self-modifying organism capable of infecting any form of life.
- "Descolada" is the Portuguese word for "unglued", my private domain is unfix.org





The situation

- A large network of DTN enabled nodes.
- Can't get to all of them, might be hard to reach as the nodes might be in a very dirty place or in a galaxy far, far away.
- But software and/or configuration needs to be upgraded/updated.
- We want to be able to deploy nodes with a minimum amount of effort.



Requirements

- DTN aware, secure, low bandwidth, low latency, update mechanism.
- Be able to update communication (radio) time schedules.
- Use only a small amount of resources, this due to power restraints.
- Fallback for failed upgrades.
- Throw a node into the wild and let it figure out what it is going to do, what software to run and how to update itself.



Existing Methods

- Most systems like rsync, cvs, ftp, http, apt, do not work well with this setup. They do a lot of communications between sender and receiver, which in a DTN environment means a lot of latency and thus slow.
- Thus apply the motto of KISS -> Keep chatter to a minimum.



What we know

- It's DTN, thus low bandwidth, high latency.
- We can rely on versions to be correct, as the installed data was also installed by us.
- As such we can use 'upgrade packages' where applicable, a.k.a. diff's.
- In effect it is a DTN-aware Content Distribution System



Choose your allegiance

- Ask your friendly neighbours to what allegiance they belong.
 - Receive Master Key from neighbours.
- Majority vote must be good: the evil side is always in the minority.
- Join the ‘good’ allegiance.
 - Use the key from the good allegiance, ignoring all other signatures.



Allegiance uses

- In effect a key-distribution way.
- Trust who is trusted the most.
- Of course totally optional: one can bootstrap the machine with a default key which it should trust.
- When the master key expires this system can be used to choose a new allegiance, thus avoiding dead fish^wnodes in the water.



Descolada Roles

- Allegiance Master
 - The package source.
 - Knows all the nodes.
- Package Cache / Node
 - Caches received Descolada Packages.
 - Returns them to speed up retrieval of packages where possible.
 - Only caches packages for their own allegiance.




The Descolada

- Special option/port in LTP.
- Infects every node it can run on.
Okay, only when specified by the configuration etc... it is not viral (yet)
- Caches store an X amount of last recently used packages depending on their storage ability and other factors (power etc).
- DTN packets getting forwarded by a DTN node are checked if they are Descolada packets, if so, the package can be cached.



Typical Communication





Package List

- Bencoded data (len,type,value)
- Header containing labels for the fields:
Node-ID,Master|
Client,Signature,numitems,Package,Version,ExpireDate,Hash,Signature
- Per package record:
record-length,packagename,version,expiredate,hash,signature
- About 100 bytes per package record



Package

- Debian style packages, but with only the minimum amount of contents.
- Scripts can be used for auto-tweaking the update.



Push/Pull Package List

- Push:
 - Allegiance Master pushes Package List with new versions to all the clients that it knows about. (Multicast would help here!)
- Pull:
 - Node notices that a package is about to expire and thus requests a new version of the package.



Optimizations

- When forwarding a package list a node can append it's own package list of the items that are nearing expiry, which are not on the list yet.
- Multicast would be very handy for spreading the updates of the package list.



Simulation Environment


- A tarball that, after extraction, can be easily installed on a variety of platforms.
- Contains a self-contained DTN stack + initial Descolada setup.
- Includes code to 'generate' latency.
- This gets deployed on a variety of hosts
(this explains to some why I was looking around for shell accounts all over the place ☺)
- Code runs and tries to find neighbours, automatically setting up a mini DTN network.



Wildlife Environment

- When the code has matured
- Deploy on the nodes used for the SeNDT project (<http://down.dsg.cs.tcd.ie/sendt/>)
- Test it in that environment and hope that it works.
- This will solve having to recover nodes out of the icky lake for updates 😊





Questions / Comments ?

Remember: The Enemy Gate Is Down!

