SpaceOps 2006

Jet Propulsion Laboratory
California Institute of Technology

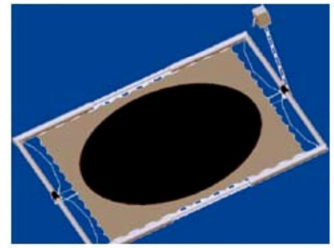
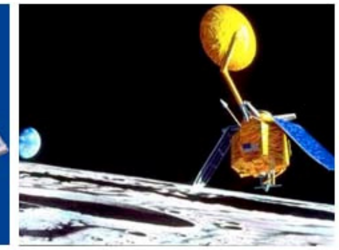

Autonomous Congestion Control in Delay-Tolerant Networks

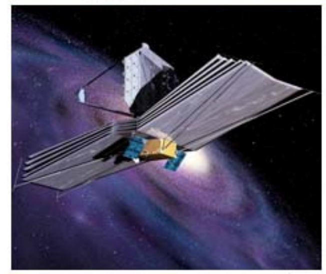


Scott Burleigh
Esther Jennings
Joshua Schoolcraft


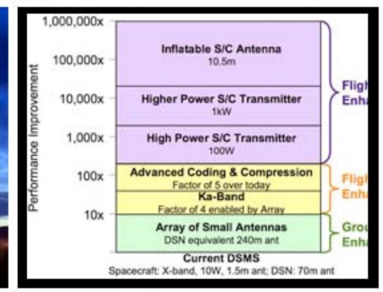

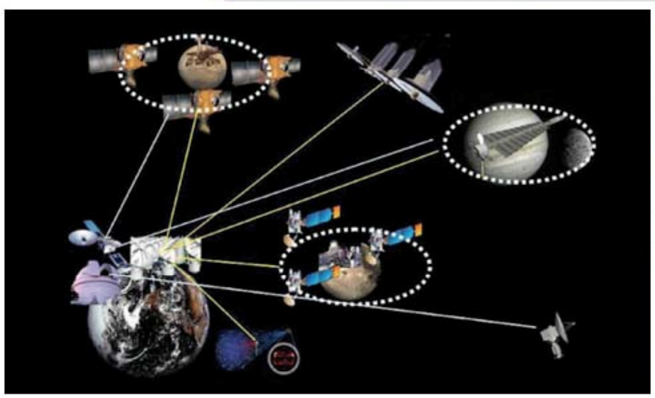

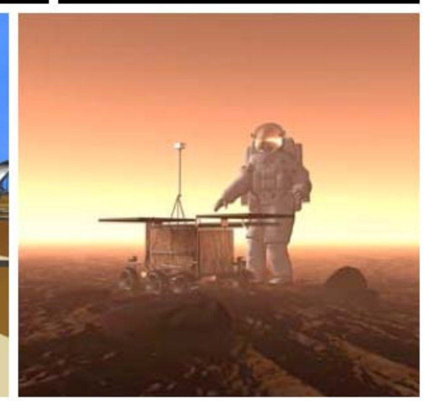

DTN Congestion Control
Motivation

- **Large-scale future space exploration will offer complex communication challenges that may be best addressed by establishing a network infrastructure.**
- **The Internet protocols are not well suited for operation of a network over interplanetary distances; a *Delay-Tolerant Networking* (DTN) architecture has been proposed instead.**
- **Networks derive much of their power from multiplexing over "trunk lines", but this multiplexing can result in *congestion* in the router nodes at the branch points.**
  - **Congestion is an excess of demand for storage resources (forwarding and retransmission buffers) at a router, causing data loss and/or router failure.**
- **Internet techniques for congestion control are, again, not well suited for operation of a network over interplanetary distances.**
- **We present an alternative, delay-tolerant technique for congestion control in a delay-tolerant network.**

DTN Congestion Control
Flow Control

- **Congestion can only be controlled (without data loss) by:**
  1. **Increasing the effective available storage at a router, by (for example) adding alternative routes and routers in parallel to the congested router.**
  2. **Reducing the rate at which high-data-volume applications inject new data into the network.**
- **Option 1 is difficult to accomplish dynamically. Most congestion control systems aim for option 2.**
- **Ultimately, option 2 is accomplished by imposing *flow control* on the transmitting application.**
  - **Flow control is the introduction of an incremental delay between the initiation of a data transmission request and the performance of the requested transmission. As the delay is increased, the effective transmission rate of the sending application is reduced.**
  - **Possible triggers for flow control:**
    - **The application's transmission rate exceeds that of an underlying protocol.**
    - **The source application's transmission rate exceeds the sink application's processing (reception) rate.**
    - **The aggregate transmission rate of a set of source applications exceeds the forwarding rate of a router serving all of those data sources – congestion.**

DTN Congestion Control
Congestion Control in an Internet

- In an Internet, flow control is typically exerted by TCP:
  - The interface between the application and TCP is a socket.
  - TCP serving a source socket imposes flow control ("blocking") at the socket in order to limit the source application's transmission rate to TCP's own transmission rate.
- Congestion in an Internet router is handled by directly triggering flow control at source applications in one of two ways:
  - Explicit
    - Router sends an ICMP source quench packet to a packet source.
    - Arrival of the source quench packet causes the source TCP to reduce its transmission rate.
  - Implicit
    - Router simply discards packets.
    - The discarded packets are not TCP-acknowledged.
    - The absence of TCP acknowledgement causes TCP at the source to infer congestion somewhere in the network, so the source TCP reduces its transmission rate.

DTN Congestion Control
Congestion Control in a DTN

- **In a DTN, flow control is typically exerted by the Bundle Protocol:**
  - **BP serving a source application imposes flow control ("blocking") in order to limit source application's transmission rate to BP's own transmission rate.**
- **But congestion in a DTN router can't be handled by directly triggering flow control at source applications, because there's no assurance of continuous or timely end-to-end connectivity on any route.**
- **Instead it must be handled by using the *custody transfer* functions of BP to trigger flow control at source applications <u>indirectly</u>:**
  - **Router discards bundles.**
  - **The discarded bundles are not custody-acknowledged.**
  - **The absence of custody acknowledgement causes congestion at the custodian (an upstream router), eventually causing it to discard bundles too.**
  - **This propagation of bundle distress eventually reaches source nodes, triggering rate control at the source BP and thus flow control.**

DTN Congestion Control
So when do we discard bundles?

- **Unlike in the Internet, rapid instantaneous growth in buffer occupancy in a DTN router doesn't signify congestion; it's a normal effect of disconnection. Only <u>sustained</u> growth in buffer occupancy indicates congestion. But how do we recognize sustained growth?**
- **Answer: a "financial" (not market) model of buffer space management:**
  - **Unoccupied buffer space is taken as analogous to money.**
  - **Routing of network traffic is taken as analogous to investment banking.**
- **A router has limited buffer space, analogous to the fixed amount of capital managed by an investment banker.**
- **Accepting a bundle for transmission is analogous to buying a non-interest-bearing debenture for face value; forwarding the bundle is analogous to selling it for face value.**

DTN Congestion Control
Router's Incentives

- Notionally, a router receives a "commission" for forwarding a bundle, based on the bundle's size and priority.
- Router is at no "financial" risk in accepting a bundle, because each bundle has a TTL (analogous to the due date on a debenture): in the worst case, the "banker" (router) gets his "capital" (buffer space) back when the bundle expires.
- But the router's "compensation" is based on forwarding, and retaining a bundle until TTL expiration ties up "capital", crowding out commission-producing activity. So accepting a bundle has a potential opportunity cost – based on the bundle's size and its residual TTL – which is a risk.
- So the router should accept a bundle whenever possible, but not when it poses a risk that is judged to be too high.

DTN Congestion Control
Rules for Bundle Acceptance

1. Insufficient capital: if this bundle's size exceeds the amount of buffer space that is currently available (unallocated), discard the bundle.
2. Risk-free investment: if currently allocated buffer space plus projected growth in allocated buffer space over the residual TTL of this bundle, plus the size of this bundle, is less than total buffer space, then the bundle constitutes no risk. So accept the bundle.
3. Balance of risk:
   - The *risk rate* of a bundle is the risk it constitutes (based on size and residual TTL) divided by the value it represents (based on size and priority).
   - The *mean risk rate* measured at the router over some interval is the total risk of all bundles accepted over that interval, divided by the total value of all bundles accepted over that interval.
   - If the bundle's risk rate exceeds the mean risk rate measured over the bundle's residual TTL, then the bundle is of above-average risk and should be discarded.
   - Otherwise, accept the bundle.

DTN Congestion Control Experiment

- A simple DTN of up to five nodes in series was constructed in JPL's Protocol Test Laboratory.
  - All nodes ran over a Gigabit Ethernet.
- An artificial and variable bundle reception delay of 0 to 50 milliseconds per byte was imposed at the node (E) that was the final destination of the bundles issued by the source node (A). This delay caused congestion at the proximate router, which was propagated ultimately to the source node.
- For each test run:
  - The source node issued 5000 bundles of 61,440 bytes each in custody transfer mode.
  - Total elapsed time to effect delivery of all bundles at the final destination was measured.
  - Bundle delivery throughput rate was calculated.

DTN Congestion Control
Topologies examined

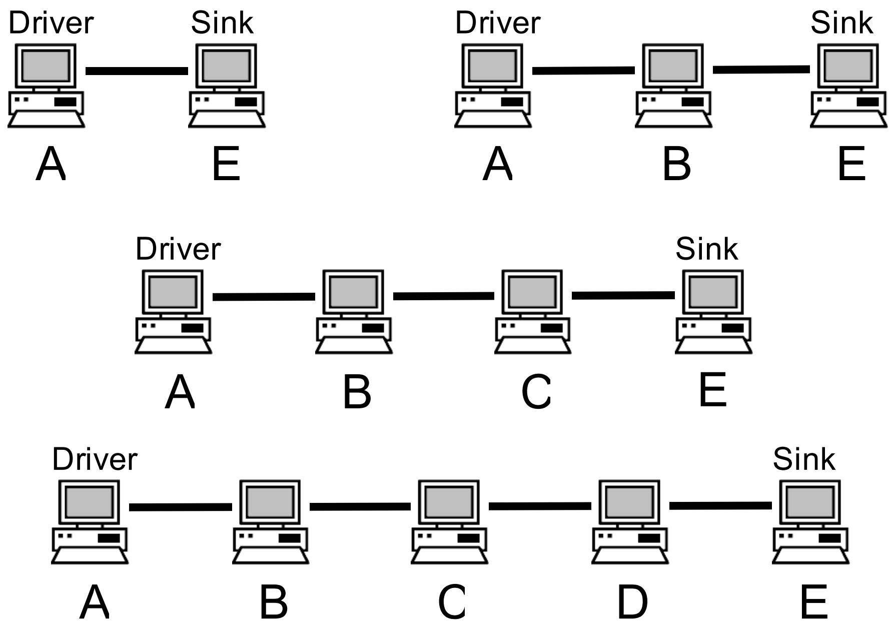

DTN Congestion Control Results

- No data loss and no router failure in any test.
- With zero artificial delay, the throughput rate measured between two nodes with no intervening routers was 300 Mbps.
- Throughput rates for other topologies and imposed delays are as shown:

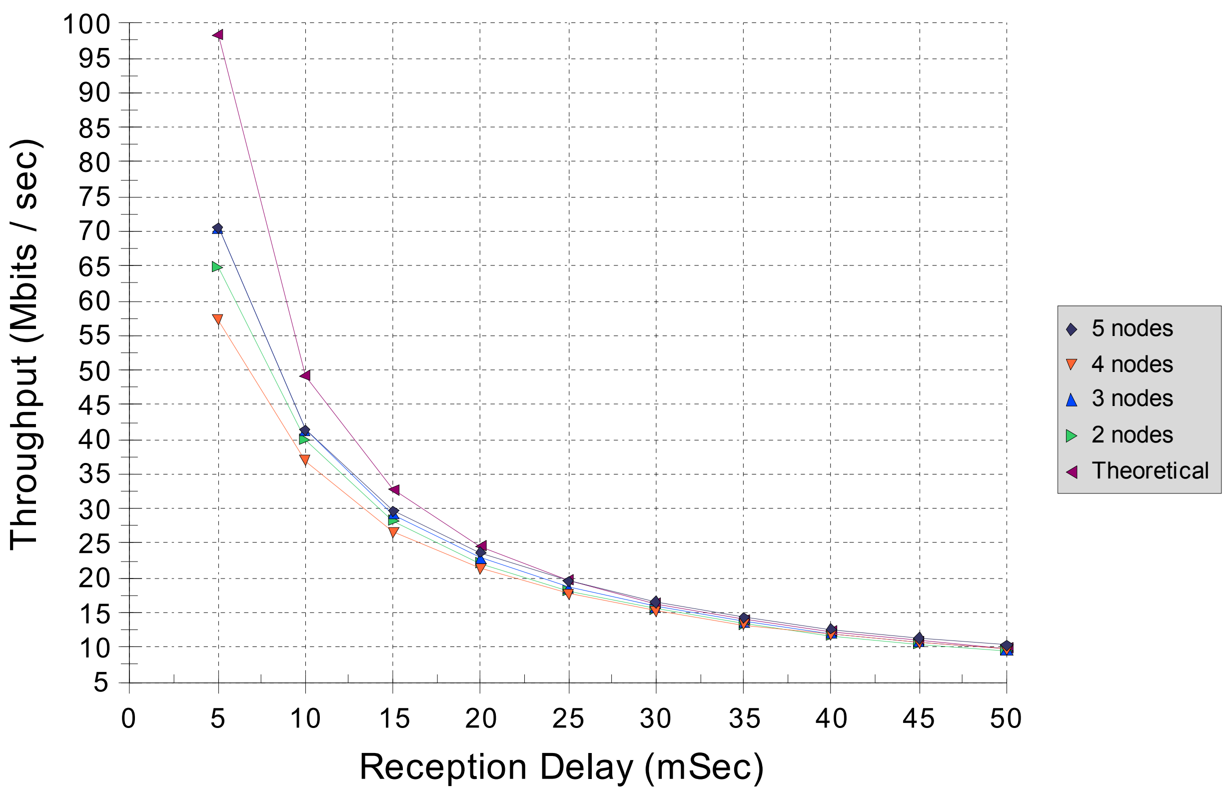

DTN Congestion Control
Remarks

- Reducing the reception rate at the receiver reduces the overall throughput rate, throttling the network in a controlled manner.
- Increasing the number of "hops" in the end-to-end path reduces throughput significantly at low levels of imposed reception delay but less so at higher levels.
  - At low levels of reception delay, the time consumed in route computation and bundle processing at each router is a significant fraction of total forwarding time.
  - As reception delay increases, the time consumed in transmitting the data at the reduced data rate becomes much greater than route computation and bundle processing time, so the number of routers in the end-to-end path recedes in significance.

DTN Congestion Control Conclusions

- **The congestion control mechanism proposed for this study appeared to be effective.**
- **Each router only needed local information in order to make its bundle acceptance decisions autonomously.**
  - **No reliance on continuous or timely communication with any other node.**
  - **No additional protocol traffic.**
- **In future studies we will:**
  - **Examine more complex topologies, including grids and trees, to explore the operation of this congestion control system over multiplexed trunk lines.**
  - **Investigate alternative bundle acceptance rules that might enhance performance.**