Next-Generation Wireless Internet: Research Questions and Platform
or
Twelve Steps to Facing the Physical Layer

Brad Karp
UCL Computer Science

NeXtworking 2007, Berlin
19th April, 2007

Scope of "Wireless Internet"

- This talk: wireless networks that allow real-time, end-to-end connectivity (DINs?)
- e.g.,
  - Mesh networks: fixed wireless nodes at edges of Internet, used to carry commodity Internet traffic "last mile"
  - Sensor networks: fixed, resource-constrained nodes, used to collect measured data
- Mobility outside the scope of this talk

Outline

- Lessons from recent history
- Limitations of current wireless research platforms
- Two pressing problems
- Research infrastructure to address these problems

History: Tremendous Focus on Routing

- Roots: ad hoc routing

Original metrics: packet delivery success rate (ignoring link-layer loss), routing protocol overhead, path stretch

Later metrics: state per node, throughput (including link-layer loss)

    - BVR: O(dB) routing state per node; number of beacons B grows very slowly in number of nodes N
  - Non-geo routing
    - VRR, S4: O(sqrt(N)) state per node, very low stretch

Claim: No Need for Another Wireless Routing Protocol!

- Quibbling now is about constants in scaling
- Sensornets hardly ever **use** any-to-any routing (trees suffice for vast majority of apps)
- Flat-routed wireless systems with paths of more than a few hops seem less likely than ever:
  - Gupta and Kumar's capacity result: per-node capacity falls as scale of network increases…
  - …modulo Grossglauser's result: fixed per-node capacity by constraining path lengths under mobility, when load latency-insensitive

History: Simulation-Based Evaluation

- CMU's 802.11 extensions to ns-2 in 1997
- Benefits of ns-2 802.11 extensions:
  - Common environment for apples-to-apples protocol comparisons
  - Lowered barrier to work in area: 1997-2002, explosion in wireless networking research
- Curse of ns-2 802.11 extensions:
  - Simulations extremely poorly predictive of real system behavior (radio propagation, interference)
  - Encouraged "comparison" ungrounded in reality!

History: Testbed-Based Evaluation

- ca. 2002-03 onwards, more and more testbeds
- e.g., MIT Roofnet (now Meraki)
  - COTS 802.11b radios, omni antennas
    - Multi-bit-rate, but rates (coding) fixed
    - MAC nearly entirely rigid (modulo enabling/disabling ACKs, setting CCA threshold)
  - 40-node outdoor urban deployment
- e.g., many mote testbeds
  - Most recent mote generation's radio: 802.15.4
    - Link-layer ACKs in software, can enable/disable CSMA in software
    - Still rigid: fixed, single bit-rate, rest of MAC in hardware
  - 20-100-node indoor and outdoor deployments

Claim: Inflexible COTS Radios Unduly Constrain Wireless Research

- Roofnet: bit-rates only available in powers of 2

COTS hardware has clearly allowed fast building of real systems…

…but how can we build real systems that allow rich experimentation with the MAC and PHY layers?

- WiLDNET: highly variable packet loss rates on long-distance 802.11 links; PHY not intended for use on long-distance links!

Outline

- Lessons from recent history
- Limitations of current wireless research platforms
- Two pressing problems
- Research infrastructure to address these problems

How can we better understand and maximize practical capacity?

- EE understanding:
  - Small cells (low-power transmission) yield spatial reuse, and thus high capacity
- Recent CS result (COPE XOR coding):
  - Overhearing allows enhanced inter-packet coding at layer 3, and thus high capacity
- These two at least partly in opposition:
  - Reducing transmit power reduces overhearing, but increases reuse
  - Increasing transmission bit-rate reduces overhearing, but increases throughput for one (or more) sender-receiver pair(s)

How can we evaluate multiple physical environments on one wireless testbed?

- Physical environment crucial to radio propagation; too nuanced to model accurately in simulation today
  - "Internet in a rack" approach lets you evaluate "this wireless network in this rack in this machine room"
- Promising direction: wireless link emulators
  - Precisely emulate link behavior in hardware, in digital domain, patched between radios
  - Glenn Judd's thesis a compelling start…

Judd’s Wireless Link Emulator

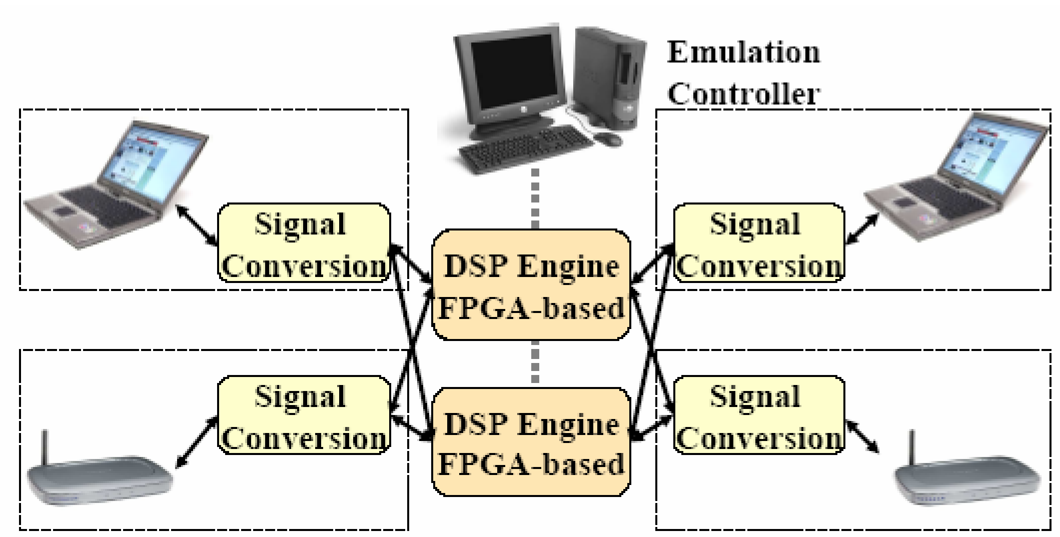

Goals for Flexible Wireless Testbeds

- Software implementation and control of full MAC, channel coding, transmit power…
  - Each radio tailored to full-system design
- Generality to allow hardware sharing across projects
  - One hardware instance, many radios
- Ease of use for “layer 3 and up” collaborators
  - Suggests “library” of link types, MACs
  - “Hackable,” but needn’t write from scratch

One Possible Platform: Ettus Research USRP / GNU Radio

- Modular: different daughter boards for different frequency bands
- Incredibly flexible: PHY and MAC fully in software
- Expensive: ca. $1000 per radio for 2.4 GHz configuration; plus 3 GHz PC saturated to received 1 Mbps 802.11b
- Steep learning curve for CS researchers (benefit! ☺)

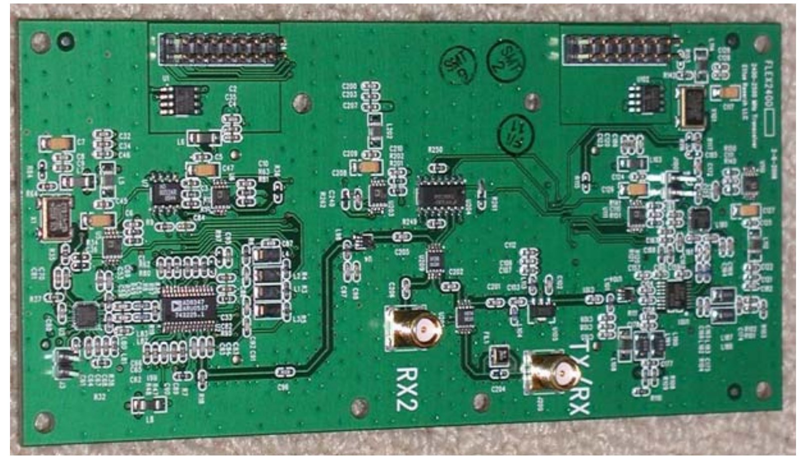

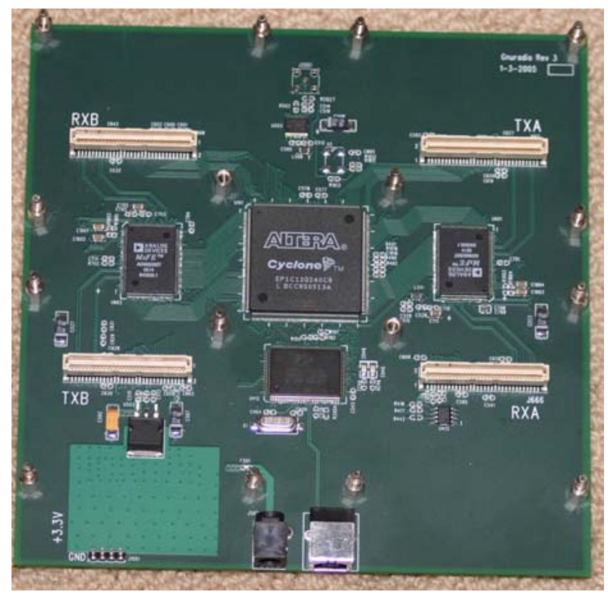

Conclusion

- Lessons from past 10 years of wireless research:
  - Simulation cannot predict system performance
  - COTS hardware constrains system designs
- Pressing questions:
  - How can we build practical multi-hop wireless systems that maximize capacity?
  - How can we evaluate many physical environments with one wireless testbed?
- Promising emerging research platforms:
  - Software-defined radio
  - Hardware wireless link emulators