Networked Systems for Developing Regions

Lakshminarayanan Subramanian, Aditya Dhananjay and Ashlesh Sharma

Courant Institute of Mathematical Sciences
New York University

SIGDev

- Proposed new SIG, in "Computation for development"
- Areas:
  - Networks, Systems, Security
  - HCI
  - AI, NLP, Data mining, Speech, Vision
- Starts this year with DEV 2010
  - http://dev2010.news.cs.nyu.edu
  - July 10, 2010 submission deadline.

Why Computation for Development?

Computation for Development?

- Sustainable Development Theories:
  - Jeffrey Sachs: End of Poverty
  - Bill Easterly: Elusive Quest for Growth
  - C.K. Prahlad: Fortune at the Bottom of the Pyramid
  - Amartya Sen: Development as Freedom
  - Paul Collier: The Bottom Billion
- Commonality: "Rural Empowerment critical to sustainable development"

"Appropriate Technology a potential enabling factor to empower rural markets"

The Bottom of the Pyramid

- 3-4 billion people with per-capita equivalent purchasing power (PPP) less that US$2,000 per year
- Could swell to 6-8 billion over the next 25 years
- Most live in rural villages or urban slums and shanty towns—movement towards urbanization
- Education levels are low or non-existent (especially for women)
- Very hard to reach, disorganized, and very local in nature

The Dharavi case

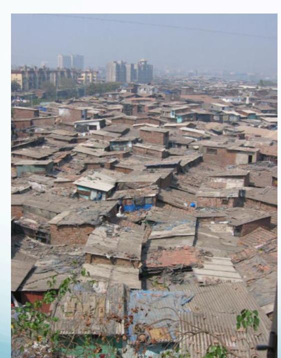

- Largest Slum in India
  - High cost of being Poor!
  - 85% have a TV
  - 50% have a pressure cooker
  - 21% have a telephone
  - … but can't afford a house
- In Bangladesh:
  - Poorest devote 7 percent income to communications (GrameenPhone)
- These are valid markets…

ICT: A Big Missing Piece

- Information and Communication Technologies (ICT) can impact everyone
  - "Bottom of the Pyramid"
- Enable wide range of essential services
  - Not just Internet access:
    - Health, education, financial services, commerce
  - Low-cost
  - Low-power
  - Need for scalable and sustainable solutions
- First World technology is a bad fit
  - New research agenda

How can ICT help?

- Communications
  - Awareness, access to external world, phone calls
- Healthcare
  - "Where there is No Doctor?": Rural healthcare system
    - Telemedicine/consultation
    - Continuing Medical Education for Health-workers
- Finance
  - Microfinance audit, insurance schemes
- Education
  - Educational modules, distance learning
- Others
  - Agriculture, Commerce, Supply chain and E-governance

Key Challenges

We should keep in mind:

- Cost
- Connectivity
- Power
- Reliability
- Accessibility
- Labor
- Socio-cultural issues

Appropriate Computing Solutions

- Minimalistic computing
- Low cost devices
- Ease of deployment and adaptability

CATER: Cost-effective Appropriate Technologies for Emerging Regions

http://cater.cs.nyu.edu

Network connectivity is key!

- Traditional wire-line connectivity solutions are not economically viable!
  - Low user densities, low purchasing power
- Potential options
  - Develop new *low-cost connectivity* solution!
  - Leverage existing *low-bandwidth* wireless solutions
    - Cellular, Satellite, WiMax
  - *Intermittent links are a fact of life*
    - Budget constrained links
    - SMS
    - Power outages
    - Physical transportation links

Connectivity options

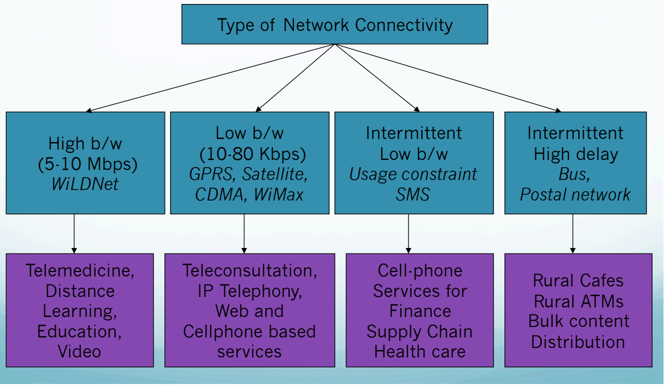

C1: Low-cost Connectivity
C2: Reliability and Power

C1: Low-cost Connectivity

- WiLDNet – WiFi based Long Distance Networks
- ROMA – Multi-radio Mesh Networks
- WiRE Architecture – For Rural Connectivity
- Hermes – Data over Unknown Acoustic Channels

Rural Connectivity in Developing Regions

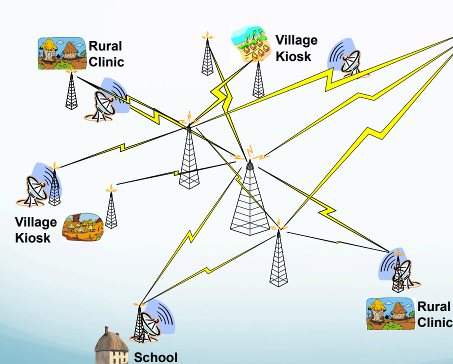

Current solutions:
- Wired
  - Dial-up/DSL
  - Optical fiber
- Cellular
- VSAT
- WiMax

WiFi-based Long Distance Networks

- **WiLD** links use *standard 802.11* radios
- Longer range up to **150km**
  - Directional antennas (24dBi)
  - Line of Sight (LOS)
- Why choose **WiFi**:
  - Low cost of $500/node
    - Volume manufacturing
  - No spectrum costs
  - Customizable using open-source drivers
  - Good datarates
    - 11Mbps (11b), 54Mbps (11g)

AirJaldi Network

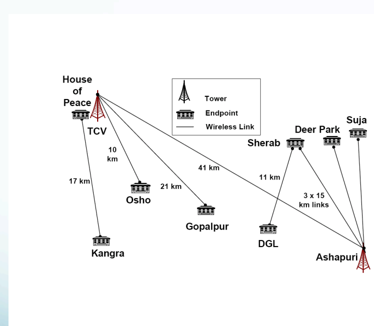

- Tibetan Community
- WiLD links + APs
- Links 10 – 40 Kms
- Achieve 4 – 5 Mbps
- VoIP + Internet
- 10,000 users

Routers used: (a) Linksys WRT54GL, (b) PC Engines Wrap Boards,
Costs: (a) $50, (b) $140

Aravind Eye Hospital Network

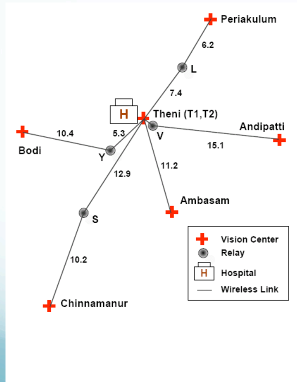

- South India
- Tele-ophthalmology
- All WiLD links
- Links 1 – 15 Kms long
- Achieve 4 – 5 Mbps
- Video-conferencing
- 3000 consultations/month

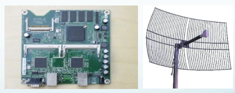

Routers used: PC Engines Wrap boards, 266 Mhz CPU, 512 MB
Cost: $140

Deployment

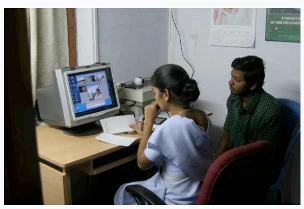

Overall Impact

- Both networks financially sustainable
- 50000 patients/year being scaled to 500000 patients/year
- Over 30000 patients have recovered sight

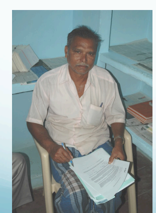

Experience with WiLD Networks

- In the field, point-to-point performance is bad
- On a 60km link in Ghana
  - We get *0.6 Mbps* TCP vs *6 Mbps* UDP
- On a relay (single channel)
  - We get only *2 Mbps* TCP

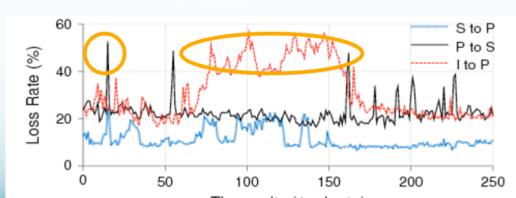

WiLDNet Design Overview

- Fix 802.11 protocol problems
  - Replace CSMA -> TDMA
- Enforce synchronization of multiple links
- Variable channel loss
  - Adaptive loss recovery
  - Combine retransmissions and FEC

Design Constraints

- No hardware changes
- Modify WiLD routers, not endpoints
- Routers are inexpensive machines
  - low processing power
  - low energy budget (solar)
- We want to be *spectrum* efficient

Problem with 802.11: ACKs

- Low utilization
  - Large propagation delays
  - Stop & wait inefficient
  - RTS/CTS makes it worse
- ACK timeouts
  - ACK doesn't arrive in time
  - Retransmissions until retry limit reached

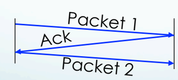

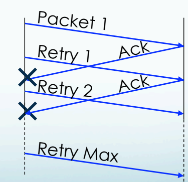

Problem: Propagation Delay

- Large propagation delay → high collision probability

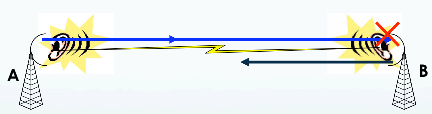

Design Choices for WiLDNet

- Use Sliding Window flow control
  - 802.11 MAC ACKs disabled
  - Packet batches sent every slot
  - Slot allocation determined by demand
- Replace CSMA with TDMA on every link
  - Alternate send and receive slots

Inter-Link Interference

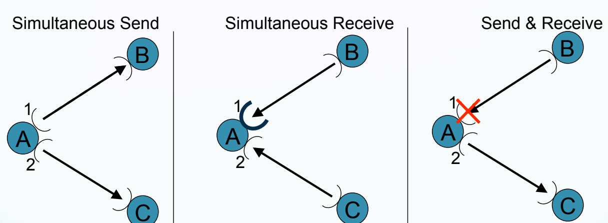

- Disable CCA
- 12dB isolation

Implicit Synchronization for TDMA

- Every packet is time-stamped in TX slot
- Slots are offset because of propagation delay
- We don't use explicit marker packets to signify end of TX slot*

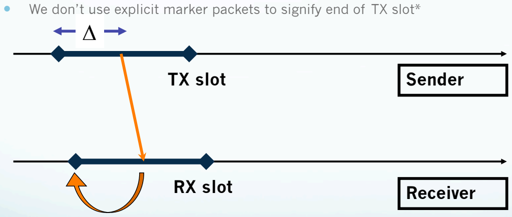

* 2P MAC protocol (Raman et al. Mobicom '05)

Channel Loss: From external traffic

- **Strong correlation between loss and external traffic**
- **Source (A) and interferer (I) do not hear each other**

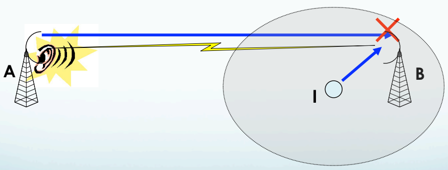

Loss Recovery: Bulk ACKs + FEC

- Bulk ACKs:
  - Aggregate ACKs (bit-vectors) sent with every packet
  - Use retransmissions for loss recovery
  - Retry limit can be per-packet
- Adaptive FEC:
  - Sender performs encoding of packets proactively
  - Packet level FEC
- Tradeoff of BW and Delay
  - Bandwidth efficient: use Bulk ACKs
    - TCP, bulk traffic
  - Delay efficient: use Adaptive FEC
    - Voice, Video

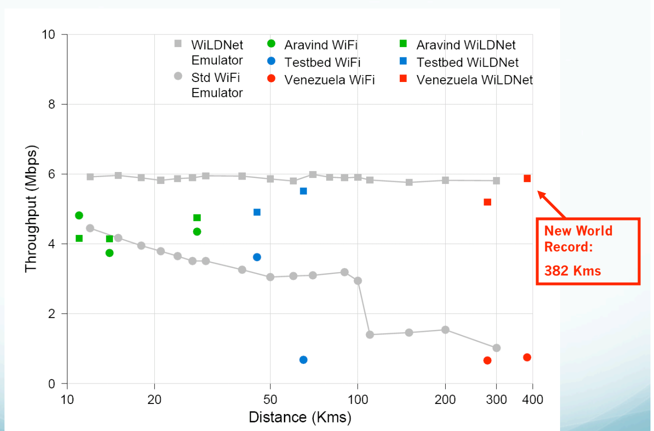

Take-away Message

- Off-the-shelf hardware can be leveraged
- New protocol stack for long-distance environments
- Low-cost and ease of deployment
- Real world deployments with high-impact in many countries across the globe

Wireless Mesh Networks for Local Distribution

Multi-radio mesh promises greater throughput

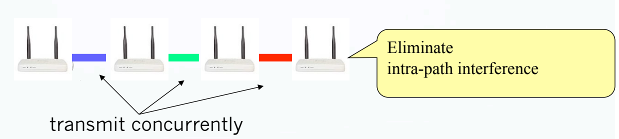

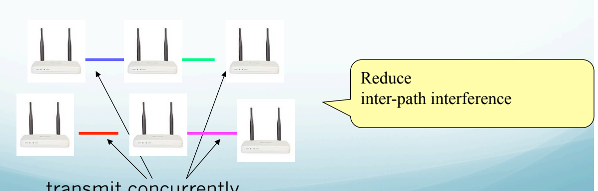

Deployment scenario

- A few radios per node
  - This talk focuses on dual-radio meshes.
- Relatively static channel assignment
  - Assignments last over minutes or hours.
- Mesh access network consists of a few gateways

Challenges in multi-radio meshes

- How to assign channels to radios?
- How to pick high throughput paths?
  - Throughput is no longer limited by sum of ETTs
  - We care more about the bottleneck link in the path

Channel assignment affects routing

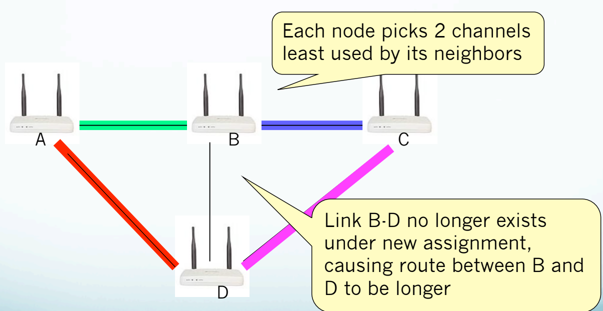

How does ROMA assign channels?

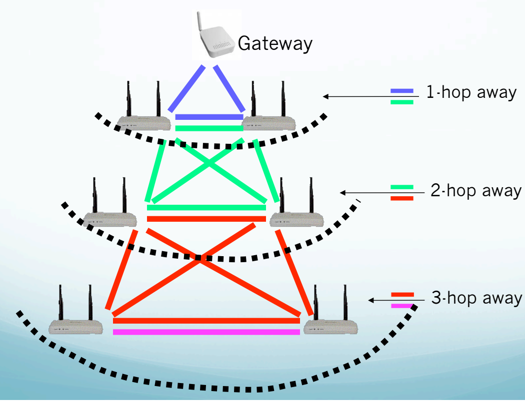

How does ROMA pick routes?

1. Estimate link performance
   - Throughput is dependent on the worst performing link, not route length.
2. Balance the tradeoff of high throughput and low overhead in choosing routes.
3. Discover better routes on not currently assigned channels

#1 Estimate link performance

- ETT/ETX over-estimate link performance.
- Besides average loss, other factors affect performance:
  - Loss variations
  - External load
- ROMA's link metric:

$$ETT = \frac{1}{(p_a - p_v) * (p_a' - p_v')} \qquad L$$

#2 Select routes

- SIM route metric [Das et al. NSDI'08] trades off performance and overhead
- Extend SIM to account for external load and variation

$$\underbrace{0.2 * \sum ETT_i}_{\text{Capture tx overhead}} + \underbrace{0.8 * \max(ETT_i)}_{\text{Capture bottleneck link(s) performance}} *(1+L)$$

- Discover better routes through "investigation"

Testbed evaluations

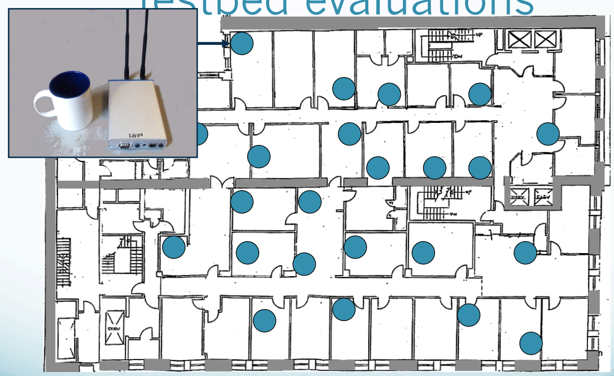

- 24 nodes in downtown manhattan office building
- Software forwarding rate limited by CPU (~6Mbps)

Testbed evaluations

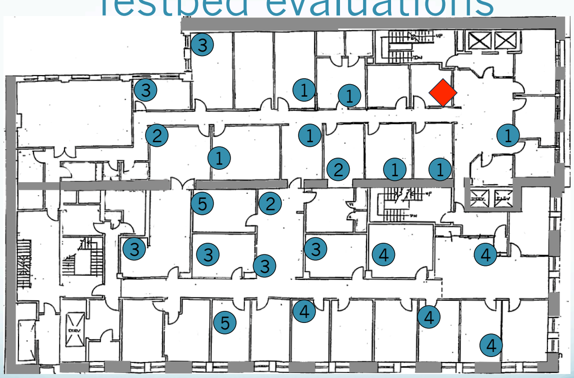

- 24 nodes in downtown manhattan office building
- Software forwarding rate limited by CPU (~6Mbps)

Performance Questions

- Does ROMA have better throughput than alternative channel assignment strategies?
  - Single flow throughput
  - Aggregate throughput
- Does incorporating loss variation (or external load) result in better routes?
- Is channel assignment stable over time?

Single flow performance

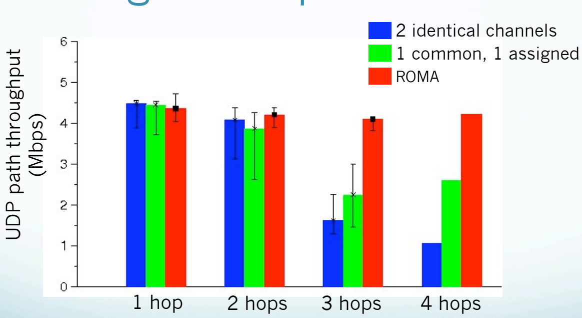

- ROMA's multi-hop path throughput approaches that of single-hop paths

Aggregate performance

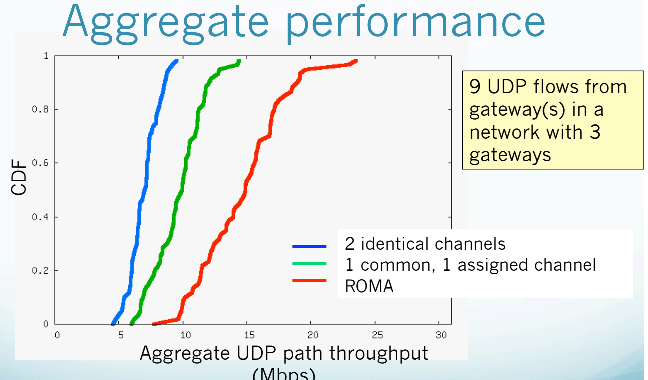

- ROMA can utilize many available channels to improve aggregate throughput

Incorporating link variations result in better routes

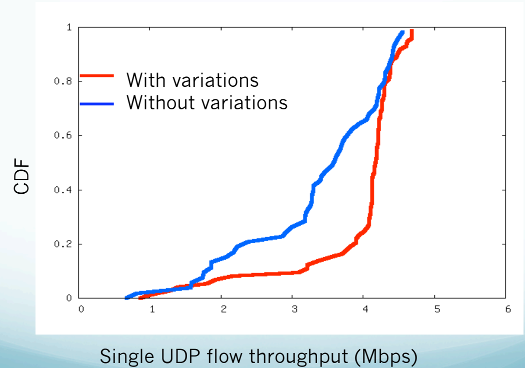

Channel assignments are stable over time

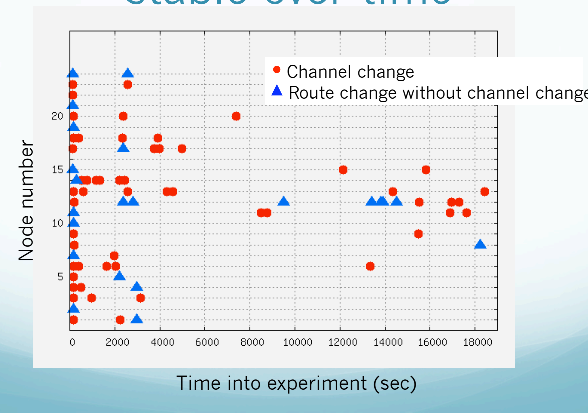

Take-away Message

- Joint channel assignment and routing leads to high performance gateway paths
- ROMA:
  - Addresses practical challenges such as loss rates, loss variations and external load
  - Chooses paths with the best “bottleneck” link
  - Stable routes with predictable performance

WiRE Architecture for Rural Connectivity

WiRE Architecture

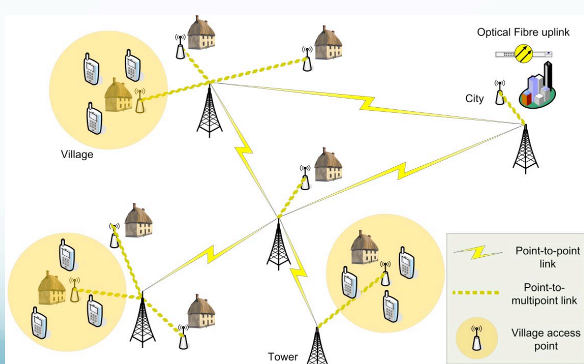

Architectural components

- Point-to-point WiLDNet links
- Point-to-multipoint distribution links
- Multi-radio mesh links
- A large local cache at each node
- Mobile devices as end-points
  - Why? – 40% rural users own a cellphone in Africa!!!

Challenges

- Physical layer
  - Steerable antennas, better radios, 802.11n?
- MAC layer
  - Unified MAC
- Network layer
  - Naming, Addressing, QoS, routing
- Robustness
  - Power, maintenance
- Application layer
  - Security, End-to-end performance

Overall vision

- WiRE architecture – a replacement to the cellular architecture
  - Significantly lower cost
  - Much higher bandwidth
  - Focused coverage
  - Significantly lower power
  - Intermittent operations
  - Economically viable!

Next Steps

- East African backbone
  - Work in Tanzania this July
- Joint MAC protocol design
- 200 Mbps point-point link
  - MIMO
  - Analog level signal processing

Hermes: Data over Unknown Acoustic Channels

Hermes: Data over Voice Channels

- Scarce / expensive data connectivity
- Ubiquitous cellular connectivity
  - Voice and SMS services.
  - No data connectivity. Why?
  - Cost per bit for SMS is very high.
- Can we modulate data on sounds and send it over a voice call?
  - Functionally like a modem, perhaps?

It is a hard problem

- Why can't we use a traditional modem?
  - No additive noise properties
  - Traditional modulation techniques break down

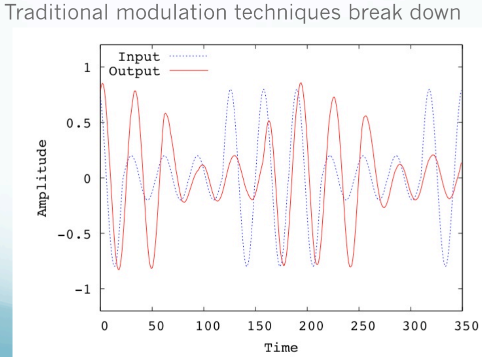

It is a hard problem

- Why can't we use a traditional modem?
  - No additive noise properties
  - Traditional modulation techniques break down
- Challenges
  - Memoryful codec
  - Heterogeneous codecs in use; many are adaptive
  - Narrowband channels
  - Auto gain control
  - Voice activity detection
  - Low quality backhaul links

Key Idea - Transmitter

- Framing: Break input stream into packets and add a CRC to each. Delimiter between packets.
- Transcoding:
  - 0 ➔ 01
  - 1 ➔ 10
- Modulation:
  - Fundamental frequency f_0. Initially, $f = f_0$
  - If input is 0, then $f = f - \delta$. If input is 1, $f = f + \delta$
    - Transmit a sinusoid of frequency f
    - The only frequencies used are: f_0 , $f_0 + \delta$ and $f_0 - \delta$

Key Idea - Receiver

- Demodulation: Differential frequencies
- Reverse transcoding rules: 01 ➔ 0 and 10 ➔ 1
  - Bit flips
  - Bit insertions and deletions
  - Mark 'X' where we are in doubt
- Reverse Framing:
  - Try to substitute the 'X's to recover the original frame
  - Check using the CRC

Parameter Choice: f_0

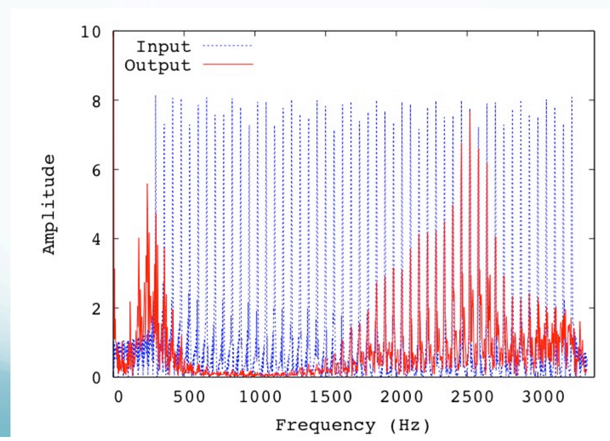

Parameter Choice: δ

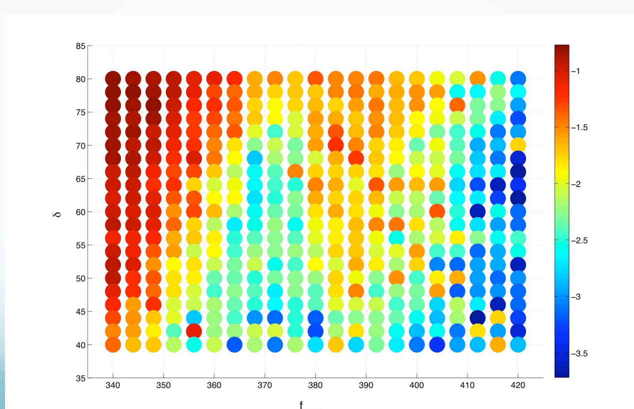

Take-away Message

- Cellular voice codecs are designed and optimized to carry only human voice
- We need to generate "voice-like" sounds to please the codec
- Reduces cost-per-bit by over an order of magnitude over SMS, while providing session-based full duplex connectivity

C2: Power and Network Reliability

Sustainability Challenges

- Bad quality grid power
  - Higher component failures, more downtimes
- Limited local expertise
  - Local operation, maintenance, and diagnosis difficult
- Lack of alternate connectivity
  - Complicates remote diagnosis and management
- Remote locations
  - Traveling is difficult and infrequent (often once in 6 months)

Poor Quality Power

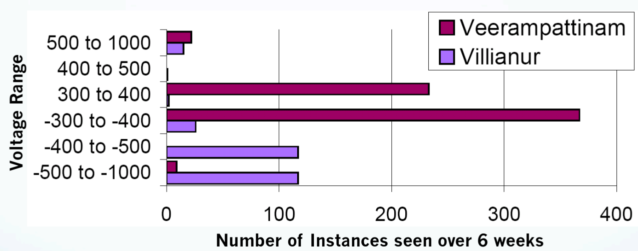

Spikes and Swells:	Low Voltages:	Frequent Fluctuations:
• Lost 50 power adapters	• Incomplete boots	• CF corruptions
• Burned 30 PoE ports	• HW watchdog fails	• Battery Damage

Power: Low Voltage Disconnect

- Low Voltage Disconnect Circuit (LVD)
  - Disconnect load at low voltage
  - Prevent battery over-discharge and hung routers
  - Without LVDs, roughly 50 visits per week for manual reboots at AirJaldi
  - Off-the-shelf LVDs oscillate too much
    - Too many automatic reboots
  - We designed new LVD circuit with better delay
    - No more manual visits or reboots!

Power: Low-cost Solar Power Controller

- Tackle spikes, swells and enable power at remote sites
- *Features*
  - PPT (peak power tracking) => 15% more power draw
  - LVD + trickle charging => Doubles battery life
  - Voltage regulator => No spikes and swells
  - Power-over-Ethernet => Remote Mgmt
  - $70 (compared to $300 commercial units)
- Have not lost any routers yet in 1 year

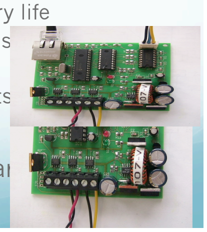

Operational Results

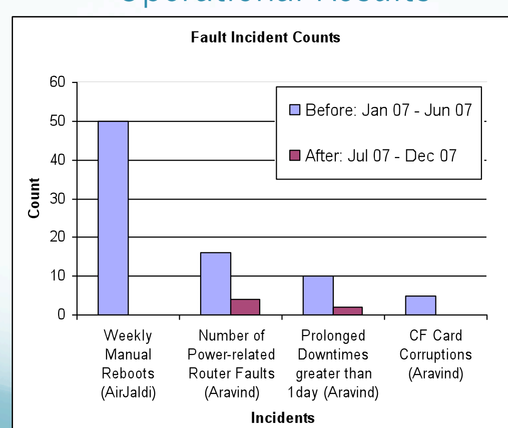

Operational Results

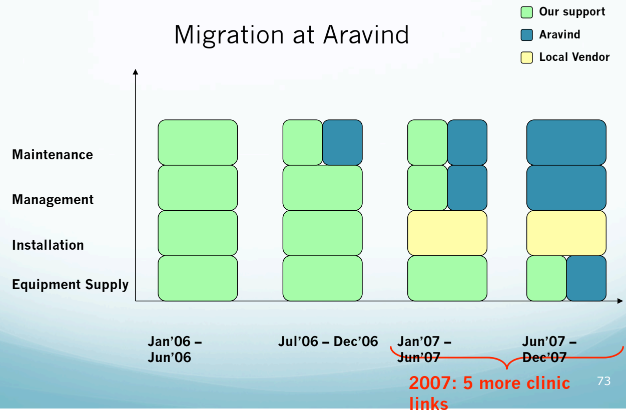

73

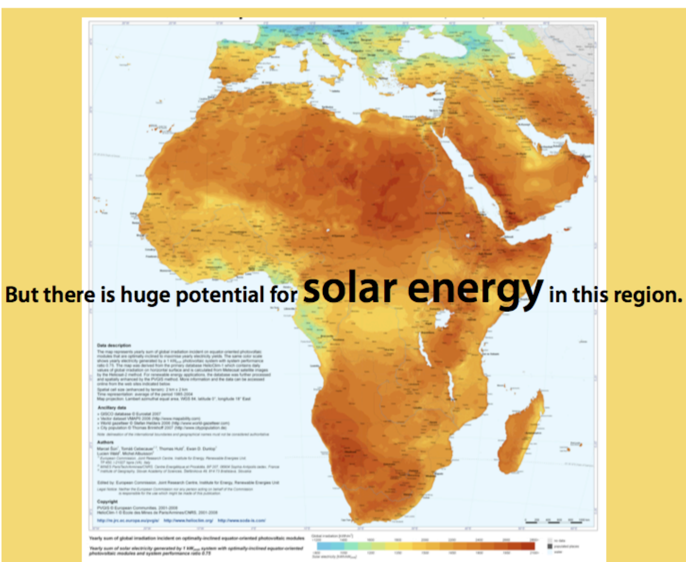

Solar Home Systems

Maintenance is a huge problem!

- Regular monitoring is critical
  - Batteries
  - Charge controllers
  - Load monitoring
- Companies cannot provide after-sales service
  - Expensive transportation and labor, for rural service
- Solar home power eventually fails
- People begin to mistrust and dislike solar energy

SIMbaLink Usage

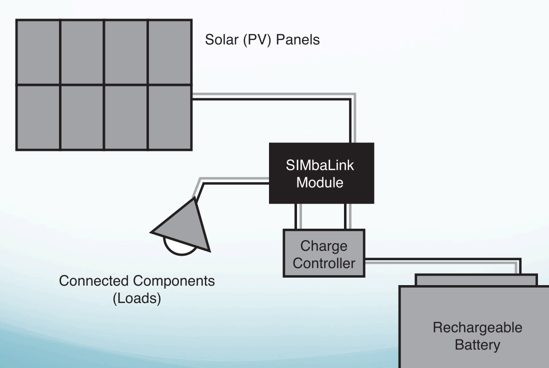

SIMbaLink Prototype

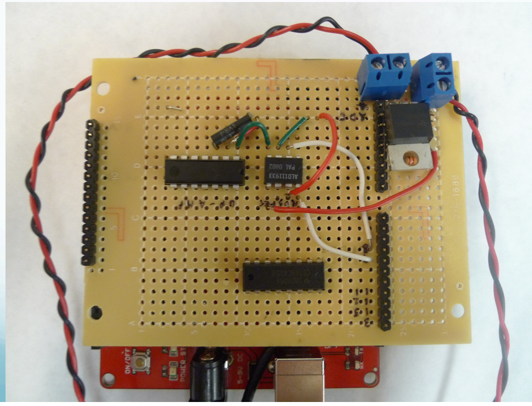

Take-away message

- Solar power is abundant, but they are tricky to manage
- Remote monitoring is the key to lowering costs and keeping the system reliable
- Leverage ubiquitous cellular connectivity to send SMS-based updates to a server
- Improves life-time by a factor of 4
- Reduces cost by a factor of 10-20

Network Reliability

HW Faults

Hardware Faults at Aravind (in 2006)

Instances*	Description	Total Downtime
63	Router board not powered	63 days
7	Router powered but hung	10 days
21	Router powered but not connected to remote LAN (burned ethernet ports)	34 days
3	Router on, but wireless cards not transmitting (low voltage)	2 days
3	Router on, but pigtails not connected	45 days
1	Router on, but antenna Line-of-Sight blocked	8 weeks

***Conservative Estimate**

>90% of faults are power-related

Keeping Rural Wireless Networks Alive

- Different operating environment than we are used to
  - Many pilots relying on standard equipment do not scale
  - *Challenges and faults need to be understood*
- Systems design <u>*can*</u> mitigate these challenges
  - We present our work with 2 networks
- In doing so, we also present 3 broad lessons

SW Faults

Software Faults at Aravind (in 2006)

Instances*	Description	Total Downtime
4	No default gateway specified	4 days
3	Wrong ESSID, channel, mode	3 days
2	Wrong IP address	2 days
2	Misconfigured routing	3 days

***Conservative Estimate**

Data Collection and Monitoring

- Designed Push-based PhoneHome system
  - Measures node, link, network level properties
  - Posts to server every 3 hours
- Also used for remote management
  - With every post, opens reverse ssh tunnel
- Has helped with some diagnosis
  - E.g. 2 wireless cards do not work together during low voltage situations

Alternate Network Entry Points

- Getting into the network is itself a challenge
  - And once in, other points may not be reachable
- Enable alternate points using Satellite (V-SAT) and Cellphone
  - V-SAT up only 60% of time at Aravind
  - GPRS phone (console access) at AirJaldi
- Use in combination with cascaded hop-by-hop logins
  - Advantage: Does not depend on system-wide routing or underlying IP configuration
- Can recover from SW configuration errors
- No recourse for physical faults, except visits
  - But might still gain useful information to help preparation of visit

Final Take-away Message

- Network connectivity is the key
  - Long distance links
  - Mesh Networks
  - Hybrid combinatorial networks (WiRE Architecture)
  - Data over Voice Channels
- Reliability
  - Clean Power
  - Remote power monitoring
- Many exciting CS + Engineering problems to solve
- Large social impact