FORMAL DEVELOPMENT OF A DELAY-TOLERANT MULTICAST PROTOCOL FOR WIRELESS SENSORS

Emil Sekerinski, Tianyu Zhou
Department of Computing and Software

with thanks to

Charles de Lannoy, Erik Frechette
Department of Chemical Engineering

McMaster University

Supported by Global Water Futures 2016-23, 78M C$
https://gwf.usask.ca/

Water quality a concern in Indigenous communities:

Year	# long-term water advisories	# Indigenous communities
2020	160	41
2021	41	30
2024	32	30

Affects ecosystem, mental &reproductive health.

RE:MOTE PROJECT

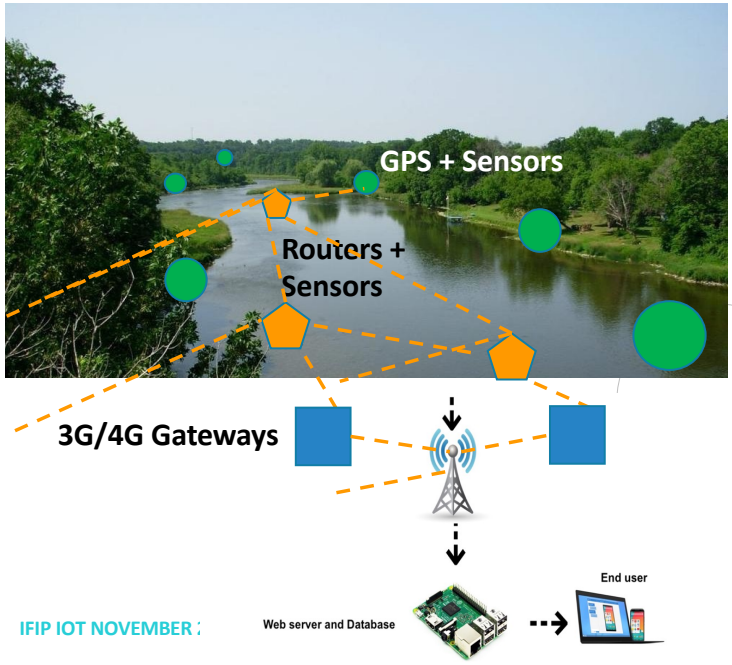

Continuous monitoring of water quality in a wide, remote area:

- Timely water advisories can be issued
- Source of contamination can be traced

Sampling technologies:

- Manual sampling & lab analysis:
  labour intensive, costly, not real-time, impractical
- Motes with satellite connection:
  direct line of sight, power consumption, cost
- Motes with drones:
  weather, power consumption, operation, cost, permission
- Motes on long-range a network with routers and gateways

LOW-COST DESIGN

Sensor Node ≈ C$1400

- C$220 main board, enclosure, power
- C$350 conductivity sensor
  rise indicative of chlorine, chloramine, phosphate, road salt
- C$20 turbidity sensor
  change indicative of increased flow, spill, disturbance to riverbed
- $400 pH sensor
  provides a baseline of the river health; traces imbalances caused by chemical spills, gas leaks, industrial waste dumping, etc.
- $400 dissolved oxygen sensor
  if nutrients are being dumped into surface water (organic matter from wastewater, phosphorous and nitrogen from fertilizers), aerobic bacteria grow, which consume oxygen and deplete the oxygen for other species, including fish
- $10 temperature sensor
  temperature causes changes in DO, helps interpreting DO

Gateway ≈ $500 (optional)

Server ≈ $100 (optional)

CHARACTERISTICS OF THE RE:MOTE SENSOR NETWORK

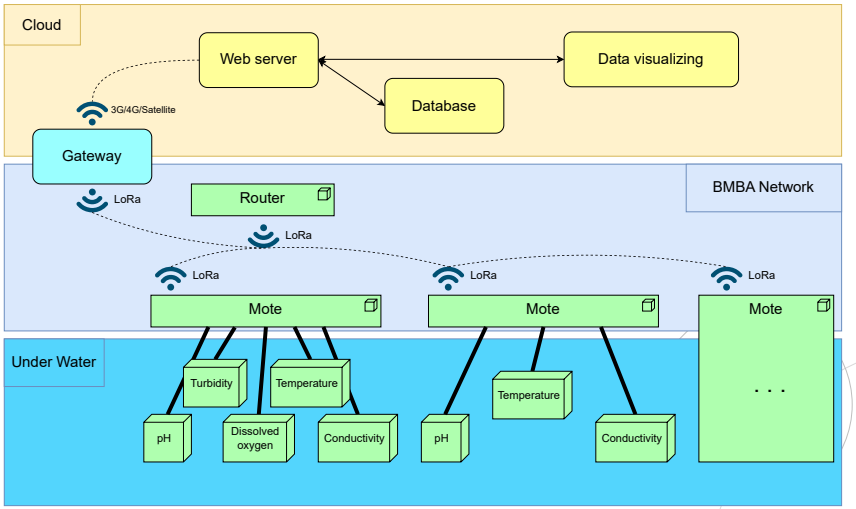

https://macwater.cas.mcmaster.ca

- Motes may become unavailable at any time
- New motes may be introduced
- Network topology may change by motes moving or being obstructed
- Sampling occurs in “large” intervals of about an hour
- Data volume is low
- Delays in transmission can be mostly tolerated
- Power consumption must be minimized

The LoRa (low-power, long-range) protocol:

- Operates on an unlicensed spectrum.
- Longer range than WiFi, Bluetooth, ZigBee.
- Lower power consumption than cellular data.
- No routing of packets, necessitating a star topology
- Range up to 10 km, typically 300m, necessitating routers
- Transmissions are unreliable, necessitating retransmissions.

CONTRIBUTION

3-D printed innards of motes

A novel protocol on top of LoRa:

- Transmission of data and acknowledgements is in large intervals, allowing the motes to sleep in between using an RTC.
- The protocol is based on blind multicast with blind acknowledgements.
- Event-B allows the protocol to be specified abstractly and then refined by correctness-preserving steps.
- The correctness of the protocol does not depend on timing, although the implementation uses timed sleeping for power conservation and random delays to minimize collisions.
- Since Event-B has no notion of fairness, eventual successful transmission cannot be expressed directly. For this, we use finitary fairness as it can be expressed with the existing Event-B proof rules.
- With finitary fairness, an upper bound on the transmission delay can be verified.

RELATED WORK

Test deployment 2024

[Miao et al. 2022] use LoRaWAN for environmental monitoring with battery-powered motes. LoRaWAN uses a star topology, and the gateway must be powered all the time. We use mesh.

[Michalik et al. 2022] propose LoRaLitE for monitoring remote areas: the gateway can enter sleep phases [20]; if a gateway fails, another mote can take over; uses a star topology. We use sleeping and mesh.

[Lee and Ke 2018] use LoRa mesh for environmental monitoring. The gateway maintains the network topology and polls each node periodically; nodes and gateway are wall-powered. We use batteries.

[Cecílio 2021] proposes TDMA based AquaMesh with beacons and routing information. We do not maintain routing information.

[Aranzazu-Suescun and Cardei 2017] propose a routing protocol with reactive routing for detecting composite events. We use periodic sampling.

[Porretta et al 2023] propose a flooding protocol for underwater acoustic sensor networks with dynamic topology. We use limited flooding for information collection: a data point floods the network until it arrives at the gateway; then, an acknowledgement of that data point floods the network.

THE BLIND MULTICAST WITH BLIND ACKNOWLEDGEMENTS PROTOCOL

M.Sc. Erik Frechette, Tianyu Zhou

New mesh protocol where

1) motes can join and leave the network,
2) the topology is dynamic,
3) transmission is highly unreliable due to environmental influences,
4) power is restricted (typically by a battery that has to last for a whole season),
5) data points are sampled in large intervals,
6) the data volume is low,
7) a delay in the reception of data points can be tolerated.

Points (1) and (2) make routing tables unreliable. Points (3) and (7) imply that resending a data point can be delayed. The key observations are that

a) when multicasting, all nodes within the range of a sender receive the sent message anyway
b) all nodes have sufficient memory to buffer all data points and their acknowledgements during the lifetime of the network.

THE SPECIFICATION

Protocol on top of the LoRa:

- The only assumption is that if a message is multicast, it is either correctly or not received.
- LoRa is not explicitly modelled.
- Motes wake up in intervals, typically for an hour, and exchange acknowledgements and data points.
- Timing affects the power consumption. The model does not assume timing.

MACHINE mac0
SEES c0
VARIABLES gateway_data
INVARIANTS
gateway_data ∈ Nodes ⇸ 1..maxSeq
VARIANTS
card(Nodes) ∗ maxSeq − card(gateway_data)
EVENTS
Initialisation ⟨ ordinary ⟩
gateway_data := ∅
acquire ⟨ convergent ⟩
any mote, s
where
mote ↦ s ∉ gateway_data
s ∈ 1..maxSeq
then
gateway_data := gateway_data ∪ {mote ↦ s}

gateway ∈ Nodes
relation
set of atomic events
decreases variant
event parameters
event guard
pair

FIRST REFINEMENT

- Each node has its own state with the sequence number of its next data point (seqnum)
- mac0.acquire is split into sample and transmit.
- sample stores the next data point into a mote's sending buffer.
- transmit transmits a data point from a mote's sending buffer to the gateway.

```
MACHINE mac1
REFINES mac0
VARIABLES gateway_data, seqnum, sending_buffer
INVARIANTS
    sending_buffer ∈ Nodes ↔ 1..maxSeq
    seqnum ∈ Nodes → 1..maxSeq
    ∀ n · n ∈ Nodes ⇒ seqnum(n) = max(sending_buffer[{n}] ∪ {0}) + 1
VARIANTS
    maxSeq ∗ card(Nodes) − sum({n · n ∈ Nodes | seqnum(n)})
EVENTS
Initialisation ⟨ extended ⟩
    gateway_data : = ∅
    sending_buffer : = ∅
    seqnum := (λ n · n ∈ Nodes | 1)

sample ⟨ convergent ⟩
    any mote where
        seqnum(mote) < maxSeq
    then
        sending_buffer := sending_buffer ∪ {mote ↦ seqnum(mote)}
        seqnum(mote) := seqnum(mote) + 1

transmit ⟨ convergent ⟩
    refines acquire
    any mote, s where
        mote ↦ s ∉ gateway_data
        mote ↦ s ∈ sending_buffer
    then
        gateway_data := gateway_data ∪ {mote ↦ s}
```

relational image

REFINEMENT STEPS

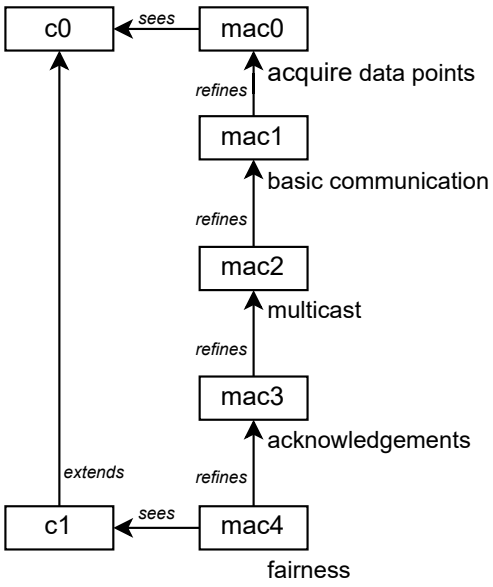

mac0: Gateway acquires data points, no network

mac1: Motes send reliably directly to the gateway

mac2: Mesh network with unreliable connectivity is introduced; receives or multicast are selected nondeterministically

multicast ⟨ anticipated ⟩ — proof of convergence is postponed

any mote, n, s, to **where**

n ↦ s ∈ forwarding_buffer(mote)
mote ≠ gateway
s < maxSeq
to ⊆ Nodes
mote ∉ to

then

forwarding_buffer := forwarding_buffer ∪ {a · a ∈ to | a ↦ {n ↦ s}}

mac3: Acknowledgement messages are introduced. Each mote keeps transmitting sampled and received data points until it receives an acknowledgement for that data point.

mac4: The (un-) reliability of transmission is modelled. After several attempts, one data point or acknowledgement will be transmitted to at least one recipient: weakest possible assumption

FINITARY FAIRNESS

gateway

The set to of recipients could be empty each time the event occurs, which means that transmission never succeeds. We use finitary fairness to express that after B events, to will not be empty. General idea:

```
variables
    x
invariants
    x ∈ ℕ

initialisations
    x :∈ ℕ
event L
    when
        x > 0
    then
        skip

fair event R
    when
        x > 0
    then
        x := x − 1
```

```
variables
    x, C
invariants
    x ∈ ℕ
    C ∈ 1..B
initialisations
    x :∈ ℕ
event L
    when
        x > 0
        C > 1
    then
        C := C − 1
event R
    when
        x > 0
    then
        x := x − 1
        C := B
```

fairness counter

unfairness bound

CONCLUSIONS

Event-B Machine	Number of proof obligations	Auto-matically proved	Need interactive proofs	Number of lines
mac0	7	3	4	29
mac1	16	7	9	50
mac2	22	11	11	68
mac3	35	16	19	90
mac4	14	2	12	98
Total	94	39	55	335

ProB was used for debugging,
Rodin for interactive proofs

- The variant of last refinement steps gives a bound on multicast and multicastAck events:

  card(Nodes) * maxSeq * (B – 1) * card(Nodes) * 2

  where B is the number of failed transmission attempts.
- A limited form of flooding based on the observation that all receivers in the range of a sender are listening to incoming messages anyway have the capacity to store them:
  - Each probe requires 4 bytes plus 1 byte for the probe type
  - Each data point additionally 8 bytes for the GPS coordinates
  - Each data point 4 bytes for the time stamp
  - Each data point 1 byte for the mote number.
- For five probes, 38 bytes per data point, or 912 bytes per day if sampled hourly.
- With 1KB per day and 100 motes, a 32 GB SD card can store the data of 320,000 days to be stored.
- With sampling per minute, 14 years of data can be stored

MOBILE & DESKTOP WEBSITE

- Visualization for community
- Data export for researchers

https://macwater.cas.mcmaster.ca

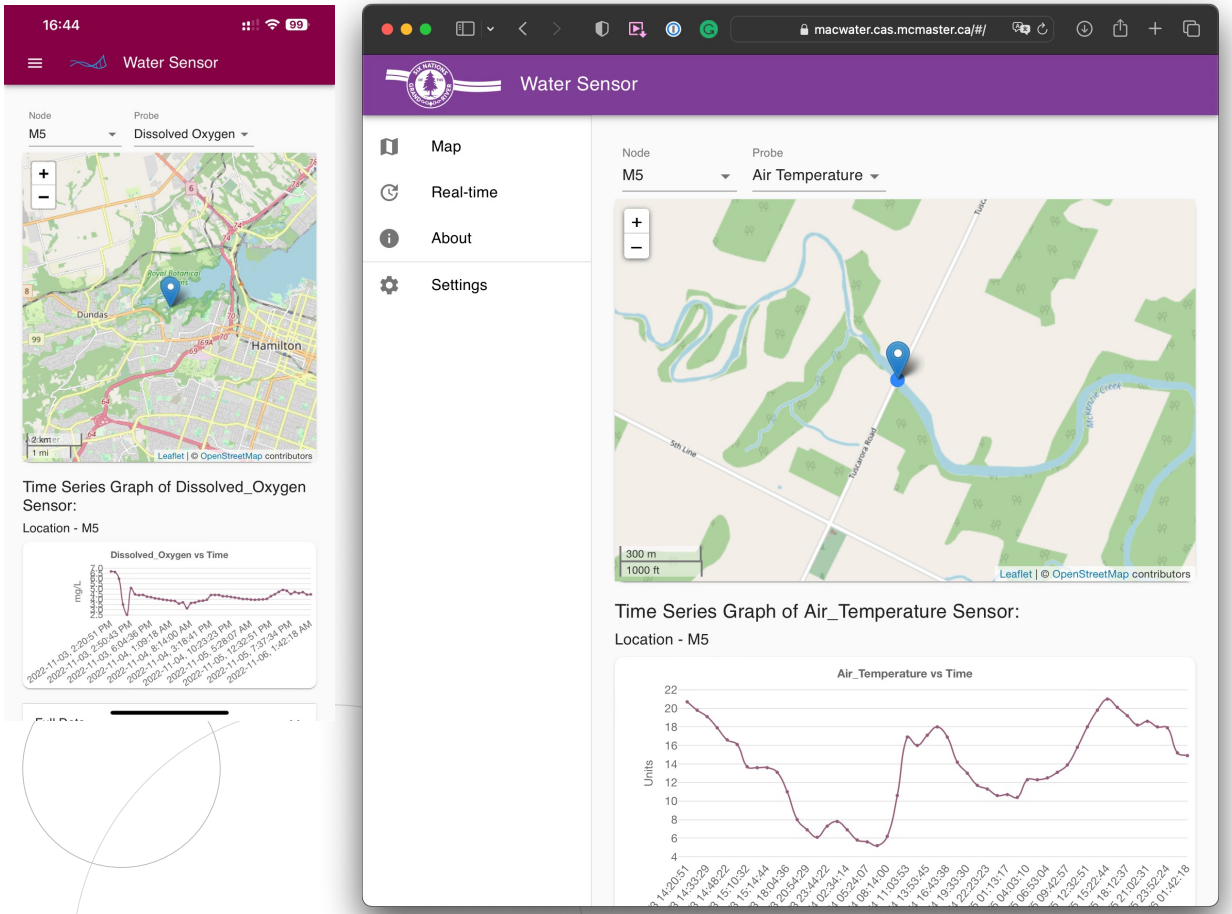

Deployment at Six Nations, M4-M5

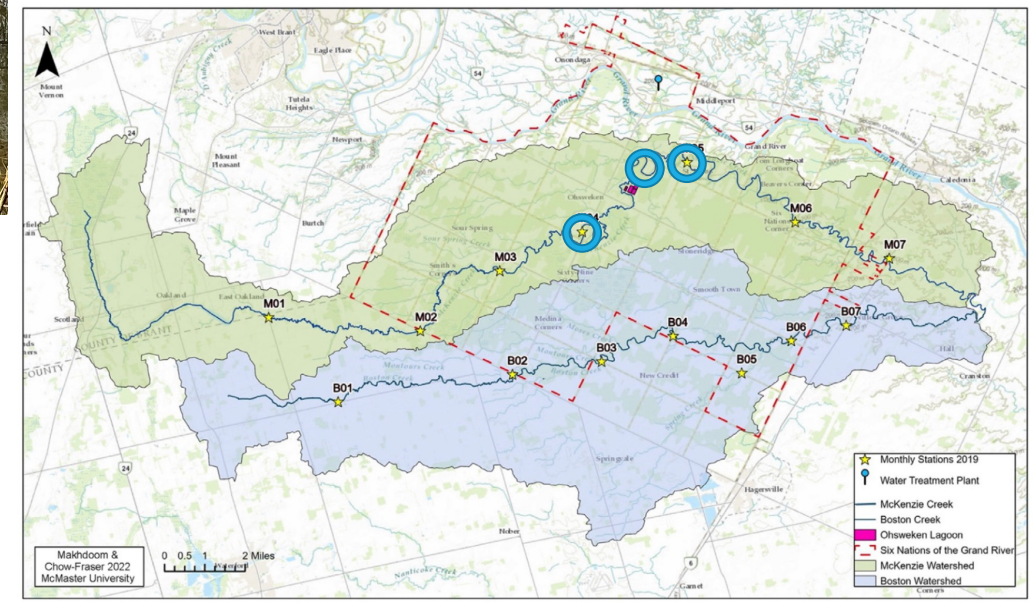

Gateway

NOVEMBER 2024

HIGH SCHOOL OUTREACH AT STEAM ACADEMY, JUNE 16, 2022

Water Festival, Six Nations of the Grand River, August 27, 2023

Six Nations Open House, May 14, 2024

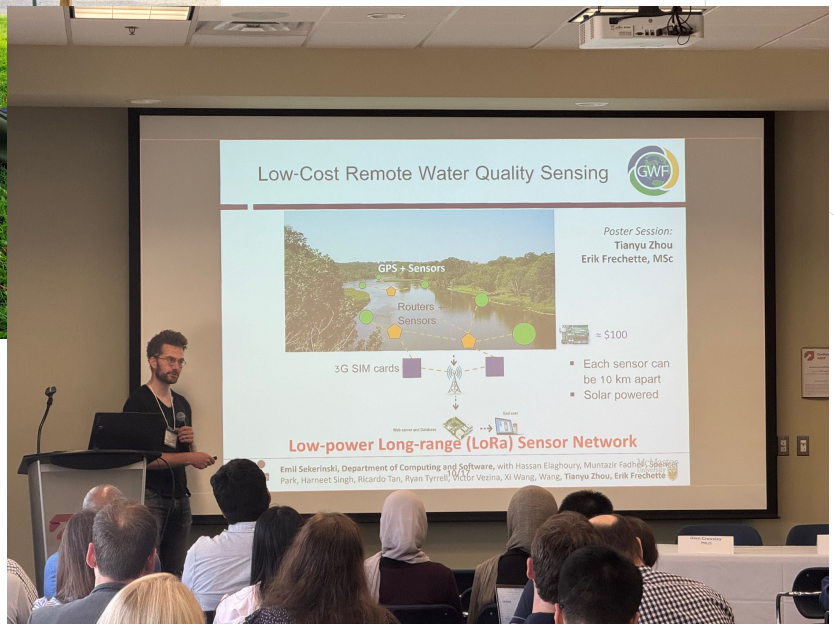

https://macwater.cas.mcmaster.ca

CHALLENGES

- Low-power mesh network for coverage in areas with poor 3G/4G reception
- Tolerant to network faults by buffering data
- Instructional visualization of data
- Auto-configuration of sensor nodes, gateways, routers
- Secure transmission of data
- Authorization with different levels of privilege
- Improved mesh networking
- Extending battery life

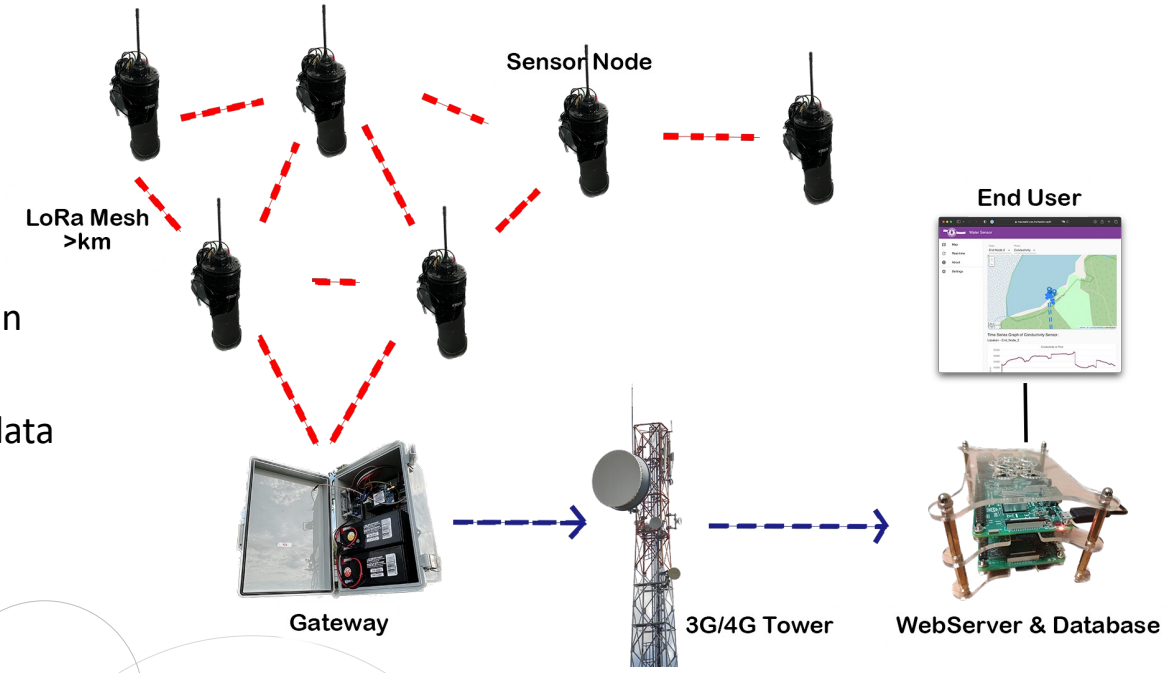

How to build a simple, reliable, teachable system?